

# System Identification of the Vestibular Ocular Reflex via Visual and Vestibular Co-Stimulation

by

James Louis Tangorra

B.S. Mechanical Engineering, Cornell University, 1989  
M.Eng. Aerospace Engineering, Cornell University, 1990


Submitted to the Department of Mechanical Engineering  
in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy in Mechanical Engineering

at the

Massachusetts Institute of Technology

June 2003



© Massachusetts Institute of Technology. All rights reserved.

Signature of Author.....

Department of Mechanical Engineering  
23 May 2003

Certified by.....

/ Lynette A. Jones  
Principle Research Scientist  
Thesis Supervisor

Certified by.....

\_\_\_\_\_  
Ian W. Hunter  
Hatsopoulos Professor of Mechanical Engineering  
Professor of Biological Engineering  
Thesis Committee Chairman

Accepted by.....

\_\_\_\_\_  
Ain A. Sonin  
Professor of Mechanical Engineering  
Chairman, Department Committee on Graduate Students

**BARKER**



# **System Identification of the Vestibular Ocular Reflex via Visual and Vestibular Co-Stimulation**

by

James Louis Tangorra

Submitted to the Department of Mechanical Engineering on 23 May 2003

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy in Mechanical Engineering

## **ABSTRACT**

The study of eye motions involved in the vestibular ocular reflex (VOR) is a key tool for understanding the performance of the vestibular system and for the diagnosis of dysfunction. Limitations in experimental equipment and in the analytic methods applied have resulted in VOR testing being conducted under artificial laboratory conditions that do not resemble the conditions under which the VOR naturally functions. The results from these tests are often unreliable, and may, in fact, misrepresent the function and performance of the VOR and the vestibular system.

The purpose of this thesis was to develop the experimental equipment, protocol, and analysis algorithms required to conduct a stochastic system identification of the horizontal, rotational VOR, while it was being used to stabilize gaze during natural, head-free tracking. By providing statistically uncorrelated stimuli to the visual and vestibular systems, estimates of the VOR's impulse response function could be made as subjects tracked a visual target that moved with an unpredictable trajectory. A novel stochastic technique was developed to generate the visual and vestibular input sequences such that they had appropriate amplitude distributions, and auto- and cross-correlation functions.

The results showed that the technique was able to identify the dynamics of the VOR over the frequency range that it naturally functions to stabilize gaze, that is from below 0.5 Hz through 4.0 Hz. Nonlinearities in the head-neck control system limited the analysis at low frequencies, and difficulties in calculating high frequency eye velocities limited the accuracy of the analysis at high frequencies. Unlike the rotational VOR tests that are commonly used today, this technique was able to distinguish between the visually and vestibularly driven eye responses, and was able to show that during head-free gaze tracking, the vestibular system is able to compensate for head disturbances with a near unity gain.

Thesis Supervisor: Lynette A. Jones

Title: Principle Research Scientist



## Acknowledgments

I wish to express my most sincere gratitude to my advisors, Dr. Lynette Jones and Prof. Ian Hunter. They have both been wonderful mentors to me. I am very thankful for their guidance, the patience they showed in allowing me to develop as a researcher, the opportunity that they provided me, and for their never-ending efforts to create an extraordinary research environment in which to work, to explore, and to learn.

I am especially indebted to Lynette. In addition to all the advice, encouragement, and support she has given me these past seven years, she spent countless hours carefully reviewing and editing each page of this thesis to ensure that it had clarity and focus. As a result of her efforts I can be proud of this thesis as a written work.

My committee members, Prof. Dava Newman and Prof. Roger Kamm, made invaluable suggestions that greatly improved the quality of this work. It was a pleasure to learn from, and to interact with them. I greatly appreciate their insight and the time that they devoted to this research.

This research may not have even existed if it were not for Dr. Henrietta Galiana. It was she who first proposed that the VOR should, and could, be tested in this manner. I am very thankful for her advice and expertise which were essential throughout all phases of this research.

The Bioinstrumentation laboratory is a wonderful place to conduct research. The combination of equipment and people that exists there is extraordinary. I have benefited so much from the friendship and knowledge of the people in the lab. They kept my research moving forward and made me look forward to coming to lab each day. In particular, I want to thank Patrick Anquetil, Serge Lafontaine, John Madden, Peter Madden, and Sylvain Martel. Whenever I ran into a hurdle that seemed impassable, I turned to one of them. Each is a remarkable scientist with keen insight and an unlimited supply of ideas. Despite routinely being pulled in tens of directions at once, each always made the time to help me, and others, solve our unsolvable problems.

I am thankful for the support my family and friends have given over these past years. Without their backing, encouragement, and understanding these last years would have been much more difficult to complete.

Lastly, and most importantly, I would like to thank my fiancée Ellie for her unyielding support, her continual encouragement, her unending patience, her help with the manuscript, and her love. The things that I am grateful to her for are too many to list. It may have been possible for me to complete a PhD without her, but I do not imagine it would have felt worthwhile.

This research was funded primarily by the National Institute of Neurological Disorders and Stroke of the National Institute of Health (NS40836). The early stages of this research were funded by the MIT Total Home Automation and Health Care Consortium.

*In memory of Louis D. and Grace d'A. Tangorra.*

# Table of contents

<b>1</b>	<b>Introduction, background, and thesis organization</b>	<b>11</b>
1.1	Introduction.....	11
1.2	VOR and gaze control.....	12
1.3	Traditional VOR testing.....	21
1.4	System identification of the VOR during head-free tracking.....	28
1.5	Research goals and thesis organization.....	36
<b>2</b>	<b>Mathematics of the system identification</b>	<b>39</b>
2.0	Chapter summary.....	39
2.1	Practical head disturbances.....	41
2.2	Linear equations that describe the input-output relations.....	42
2.3	Correlation based convolution equations and the deconvolution process....	47
<b>3</b>	<b>Design and construction of experimental equipment</b>	<b>53</b>
3.1	Introduction.....	53
3.2	Head perturbers.....	54
3.3	Visual target.....	77
3.4	Head tracker.....	78
3.5	Eye tracker.....	81
3.6	Data acquisition and experiment control.....	101

<b>4</b>	<b>Design of visual and vestibular inputs</b>	<b>103</b>
4.0	Chapter summary.....	103
4.1	Input signal requirements.....	104
4.2	Torque perturbation.....	106
4.3	Visual target.....	119
4.4	Tracking unpredictable targets with torque perturbations applied to the head.....	147
<b>5</b>	<b>Generating uncorrelated torque and target inputs</b>	<b>155</b>
5.0	Chapter summary.....	155
5.1	The modified Hunter-Kearney procedure.....	156
5.2	Steps of the modified Hunter-Kearney procedure.....	158
5.3	Defining bounds for zero cross-correlation.....	187
5.4	Swap conditions: number of elements, evaluation criteria, sequence PMF	
5.5	Conclusions .....	202
<b>6</b>	<b>Dynamics of the human head-neck system in the horizontal plane</b>	<b>205</b>
6.0	Chapter summary.....	205
6.1	Dynamics of the human head-neck system in the horizontal plane.....	207
<b>7</b>	<b>VOR experiments</b>	<b>223</b>
7.0	Chapter summary.....	223
7.1	Experimental methods.....	224



7.2	Data processing.....	227
7.3	Results from tracking trials.....	237
7.4	Head response.....	252
7.5	VOR response.....	267
7.6	Predictions of the gaze response.....	300
<b>8</b>	<b>Conclusions and future work</b>	<b>311</b>
8.1	Conclusions of the VOR experiments.....	311
8.2	Other contributions.....	314
8.3	Future work.....	315
8.4	Conclusion.....	316
	<b>References</b>	<b>317</b>
	<b>Appendices</b>	<b>323</b>
A	Additional mathematics.....	323
B	Equipment specification sheets.....	333
C	Code for generating uncorrelated torque and target sequences.....	341
D	VOR experiment conditions.....	425
E	Analysis programs.....	435



# Chapter 1

## Introduction, background, and thesis organization

- 1.1 Introduction
- 1.2 VOR and gaze tracking
- 1.3 Traditional VOR testing methods
- 1.4 System identification of the VOR during head-free tracking
- 1.5 Research goals and thesis organization

### 1.1 Introduction

The goal of this research was to develop the experimental equipment, protocol, and analysis algorithms that were required to conduct a system level evaluation of the vestibular ocular reflex (VOR) while it was being used to stabilize gaze during natural, head-free tracking of a visual target.

The vestibular system, which is composed of the semicircular canals and otolith organs of the inner ear, measures the linear and angular acceleration of the head and transmits these signals to the appropriate cortical and sub-cortical sites. This information is used to maintain balance and posture, enables coordinated body movements, influences how we perceive our spatial environment, and is a crucial input to the reflexes that the body uses to compensate for head motion. Deficiencies and failures of the vestibular system lead to debilitating dizziness, nausea, vertigo, oscillopsia, and spatial disorientation. These symptoms are among the most common medical complaints for people over the age of 65 (Foreman, 2000), and the costs of medical care for these symptoms are estimated to exceed \$1 billion per year in the United States (Vestibular Disorders Association, 2003). Although these symptoms may be associated with other systemic

problems, the majority are associated with deficits in the peripheral vestibular apparatus or in its projections to the cerebellum and brainstem. The study of the eye motions involved in the VOR is a key tool for understanding the performance of the vestibular system and for the diagnosis of vestibular deficiencies. Limitations in experimental equipment, and in the analytic methods applied, have resulted in VOR testing being conducted under artificial laboratory conditions that do not resemble the conditions under which the VOR normally functions. The results from these studies are often unreliable, and may, in fact, misrepresent the function and performance of the VOR and the vestibular system (Leigh, 1996). The focus of this research is on the application of modern system identification techniques to the vestibular system in order to develop the experimental equipment, protocol, and analysis algorithms that will give a valid and reliable evaluation of the VOR by testing the system under its natural operating conditions.

## **1.2 VOR and gaze tracking**

### *The vestibular system*


The function of the vestibular system is to generate information about head movement and to distribute this information to the appropriate sites throughout the central nervous system. The vestibular system gives us our sense of balance, provides the inputs required to control posture, enable us to make coordinated motions, and preserves the stability of the eyes (VOR) (Highstein, 1996). It measures the linear and angular motions of the head using five mechanoreceptors located in each inner ear. Linear accelerations and static position are measured by the two otolithic organs, the utricle and the saccule. Rotational accelerations are measured by three, fluid filled, semicircular canals. Accelerations of the head deflect hair bundles in the epithelial hair cells that line portions of each organ (cristae). The hair cells transduce mechanical stimuli into receptor potentials. A distortion of the hair bundles changes the membrane potential in the hair cells, which affects the discharge patterns of vestibular neurons. These neurons send

head velocity and acceleration signals to the vestibular nuclei in the medulla of the brain stem, which then transmits this information to the higher centers of the brain (Goldberg and Hudspeth, 2000).

The utricle and the saccule are arranged to detect linear accelerations and the static tilt of the head. The utricle primarily detects horizontal head motions (within the transverse plane of the head), and the saccule detects vertical motions (perpendicular to the transverse plane). Each organ consists of an ovoidal sac, about 3 mm long, and within the sac are thousands of hair cells arranged in a region called the macula (Goldberg and Hudspeth, 2000). The bundles of hairs that extend from the hair cells are covered with a gelatinous membrane. Dense, calcium carbonate crystals, called otoconia, are embedded in the membrane and act as masses on top of the flexible hairs. When the head is accelerated or tilts the otoconia shift, deflect the hair bundles, and excite responses from the hair cells which are then transmitted to the vestibular nuclei.

The semicircular canals are closed, circular tubes filled with endolymph fluid. The cross section of the canal, which is not constant over its length, is approximately 1 mm in diameter, and the overall diameter of the semicircular structure is about 8 mm (Figure 1.1). The semicircular canals are shaped like the inner tubes of bicycle tires, but are filled with liquid instead of air (Vestibular Disorders Association, 2003). The three semicircular canals of each ear are arranged in three orthogonal planes which allow the canals to sense the three directions of rotational acceleration. Conceptually, the semicircular canals are aligned with the pitch, twist, and yaw movements of the head, but the canals are not aligned with the anatomical planes of the head. The planes of the anterior and posterior canals are aligned almost vertically, and are set at about 45 deg to the midsagittal plane. The horizontal (lateral) canal is tilted about 30 degrees up from the horizontal plane when the head is in an upright position (Figure 1.1). A simple rotation of the head in one of the anatomical planes therefore elicits a complex response from the canals that must be interpreted by the central nervous system.

When the head is accelerated, the canals, which are within the temporal bone, move with the head. The endolymph fluid within the canals lags the motion of the canal, and, relative to the canal, moves in the opposite direction to the head's motion. The fluid does not rotate freely in the canal because there is a gelatinous diaphragm, the cupula, across the widest section. Hair bundles extend into the cupula in one section of its perimeter. When the head moves, the cupula moves with the head and presses against the endolymph which remains still. This causes the cupula to deflect, the hair bundles to bend, and the hair cells to be stimulated (Goldberg and Hudspeth, 2000).



**Figure 1.1:** Vestibular labyrinth of the inner ear (Vestibular Disorders Association, 2003).

There are two problems associated with the canals' response to head movement. The canals do not respond well to low frequency motions, and they habituate to long duration rotations. During sustained rotation, the mechanical properties of the cupula and the endolymph cause the endolymph to move with the same speed as the head and the cupula to return to its neutral position. As this occurs, the response of the canals decays. The time constant of this decay is approximately 6 s, but processing within the vestibular system is able to extend the time constant of the response to about 15 s (Leigh and Zee, 1999).

### *The vestibular-ocular reflex*

The functional role of the vestibular-ocular reflex is to stabilize vision by creating eye motions that compensate for undesired linear and rotational head motions. In a healthy subject, the VOR interacts with other oculomotor and head control systems so that a person's gaze (eye plus head position) smoothly follows a visual target's trajectory. Ideally, the VOR holds a visual image stable on the retina by moving the eyes with an angular velocity equal and opposite to that of the disturbances to the head's trajectory. The vestibular sensory organs measure the angular and linear acceleration of the head and generate signals that, after processing in the vestibular and abducens nuclei, are used to drive the eyes in a manner that quickly corrects for the head's disturbance. Without this correction, the eyes would move with the head, causing the retina to move relative to the target's position, which would result in the target moving across the retina. The amount of retinal slip that can be tolerated before vision deteriorates depends on what is being viewed, but generally, as the speed and frequency of an image's retinal drift increases, it is seen less clearly. The VOR does not have to be perfect, but to see images that are smaller than 1 degree clearly, retinal slip should be less than 5 deg/s and the image of the object should remain within a degree of the center of the fovea (Leigh and Zee, 1999; Highstein, 1996).

The VOR is indispensable during activities such as walking and running, where unpredictable head perturbations can span a bandwidth of 0.5 to 5.0 Hz and attain rotational speeds of up to 150 deg/s (Demer *et al.*, 1991; Grossman *et al.*, 1988). When the motion of the head cannot be anticipated, such as during locomotion, the brain must rely upon reflexes and feedback mechanisms to correct for disturbances to vision. Of the mechanisms that serve to prevent and correct gaze error, namely the VOR, visual feedback, and the cervicular-ocular reflex (COR), only the VOR has the performance characteristics necessary to compensate for the high frequency perturbations typically experienced by the head. Visual feedback is not effective at correcting gaze error during locomotion because of the time required for the visual system to process the image and then to generate a corrective eye movement (Gauthier and Vercher, 1990).

The VOR can generate eye motions with a latency as brief as 16 ms, while the visually mediated eye motions are initiated at latencies greater than 75 ms (Leigh and Zee, 1999). The cervicular ocular reflex (COR) generates eye motions based on proprioceptive feedback from neck muscles. Although COR eye motions may have as short a latency as those of the VOR, it has been shown that the COR makes a negligible contribution to higher frequency gaze stability in normal subjects (Sawyer *et al.*, 1994). It is possible, however, that in patients with a poorly functioning vestibular system, the gain of the COR could increase over time to help compensate for the loss of the VOR.

### *Gaze tracking*

During natural behavior, humans interact visually with the environment by directing their gaze through coordinated head and eye movements towards visual targets of interest. Images are seen most clearly when the gaze is directed so that the image is positioned at the center of the retina, a region called the fovea. The gaze system controls and coordinates the head and eye motions through the cephalomotor and oculomotor systems. The cephalomotor system controls the movement and stability of the head, and so positions the orbits of the eyes, and the oculomotor system directs the position of the eyes within the orbits. The gaze system also stabilizes the visual image on the fovea when either the object or head moves.

The greatest impediment to clear vision during natural activities, such as locomotion, is disturbances to the head's position (Leigh and Zee, 1999). The cephalomotor system maintains the head as a relatively stable platform for the eyes, but during walking and running the predominant motions of the head in the vertical plane are between 0.5 and 5.0 Hz, and between 0.6 and 3.0 Hz in the horizontal plane (yaw) (Grossman and Leigh, 1989). During walking, the maximum velocities of the disturbances usually remains below 50 deg/s, but when running the maximum velocity may be beyond 150 deg/s. Head stability is controlled using a combination of voluntary mechanisms, neuromuscular reflexes and passive system mechanics. The contribution



of each of these mechanisms to stability is frequency dependent, with each mechanism dominating a different frequency range (Keshner and Peterson, 1995). Voluntary control of the head dominates at frequencies below about 1.0 Hz. This is the range over which there is enough time to recognize that the head has been disturbed, and to either modulate the mechanics of the neck to resist the disturbance, or to make active corrective motions. At higher frequencies, above about 3 Hz, passive mechanics dominate the head's stability. The mass of the head, coupled with the stiffness and damping of the neck resist high frequency disturbances. In the middle range of frequencies the vestibulo-collic reflex, which receives input from the vestibular system, and the cervico-collic reflex, which takes its input from proprioceptors in the neck, are very important to head stability. These reflexes have also been found to be an important mechanism for modulating the dynamic stiffness of the head-neck system (Keshner *et al.*, 1999).

During natural disturbance motions of the head, vision is stabilized mostly by the VOR. The VOR is only one of the six oculomotor systems that control eye movement. There are two primary types of eye motions: those that shift the gaze from target to target, and those that stabilize the gaze on a target that is being tracked. All six systems work synergistically to produce the gaze response, but shifts in gaze are controlled principally through the saccadic control system, and the tracking of a moving visual target is controlled through the combined efforts of the smooth pursuit, VOR, and optokinetic systems.

Saccades shift the fovea very rapidly from one area of visual interest to another. They are used to explore the environment and to bring an image from the periphery of the retina into sharper focus on the fovea. Saccadic eye motions typically follow a standard waveform trajectory, and can attain speeds of up to 900 deg/s (Goldberg, 2000). The maximum velocity that a saccade attains is dependent on the amplitude of the angular position shift. Visual acuity is suspended during large saccades.

The smooth pursuit system allows for clear and continuous vision of a target as it moves in the environment. It keeps the target's image on the fovea by moving the eyes with a velocity

close to that of the target, although small oscillations occur in the velocity of the eyes, especially at the onset of smooth pursuit motions (Leigh and Zee, 1999). Studies have indicated that the smooth pursuit system will respond to both the slip of the image across the retina and to position errors between the gaze and the target. The motion of the target across the retina is the principal input to this system. The predictability of a target's trajectory has a strong influence on smooth pursuit motions. When a target is unpredictable, the smooth pursuit response is driven by visual feedback, and there is approximately a 100 to 150 ms delay between the target's movement and the gaze response. When the target can be predicted, the delay from visual processing no longer occurs and the smooth pursuit eye motion can follow the target with little error or lag. Smooth pursuit eye motions have a maximum velocity of between 100 and 150 deg/s.

The fixation system holds the eyes still when a subject focuses intently on a stationary target. Vision is most acute when the eyes are still, and the fixation system actively prevents the eyes from moving. Fixation is sometimes regarded as smooth pursuit at zero velocity, but studies have indicated that visual fixation is an independent system (Leigh and Zee, 1999).

The optokinetic system is a visually mediated stability system. Because of the mechanical properties of the semicircular canals, the rotational VOR does not stabilize vision well against low frequency or sustained head motions. Together with the smooth pursuit system, the optokinetic system supplements the VOR at low frequencies or during sustained rotation. As the response of the VOR decays, the response of the optokinetic system builds until it reaches a velocity close to that of the rotation. The optokinetic system only functions when there is a visual stimulus.

Most of the eye control systems create conjugate eye motions, that is both eyes move together by the same amount. For objects that are at a distance greater than optical infinity, which is about a meter, the eyes should have the same angular position for the image to fall on the same position on the retina. For objects that are within the near visual field, or moving from the far to near field, the eye must be rotated by different amounts, or even rotated towards each other, so

that the target's image can fall on both foveae. The vergence system is the system that controls these disconjugate eye movements.

Natural head-free tracking is characterized by smooth eye and head movements interspersed with small, corrective saccades. This combination of movements controls both the velocity and position of the gaze response (Carl and Gellman, 1987). During sustained pursuit, the gaze velocity approximates, but often does not precisely match, the target's velocity (Barnes *et al.*, 1987; Pola and Wyatt, 1980). To help correct for the position error that develops between the gaze and the target, small saccades are used to shift the eyes rapidly and bring the target's image onto the fovea. The saccades effectively increase the bandwidth of the tracking system by intermittently resetting the target's image on the fovea when the target is moving faster, or changing more quickly than the velocity controller is able to respond (Bahill *et al.*, 1980).

It is generally believed that visual pursuit is an adaptive process controlled mainly through two mechanisms: one is the feedback of the error between the target and gaze velocities, and the other is a predictive mechanism that creates an estimate of the target's velocity and time course (Wellenius and Cullen, 2000; Barnes and Greal, 1992; Young, 1977). Studies have shown that the pursuit response is also influenced by the target's position, acceleration, and other visual and nonvisual information, but these do not appear to be as significant for visual pursuit as the velocity and predictability of the target (Lindner and Ilg, 2000; Pola and Wyatt, 1980). The motion of a target's image across the retina creates a velocity error signal that initially drives the pursuit system like a conventional servomechanism. The visual feedback is delayed by the 100-150 ms required to process the target visually, so a tracking response driven purely by visual feedback will lag the target's motion. The predictive mechanism builds a model of the target's motion, possibly from an internal copy of the visual feedback signal (Barnes and Greal, 1992), and uses its estimation of future target motion to enhance gradually the visual feedback signal and so overcome the delay of the visual system. The means by which the predictive mechanism generates this signal is not well understood, but the end result is that eye motions are adjusted so

that the gaze velocity gain is improved, and the tracking lag reduced. The visual feedback signal is used to check the validity of the estimate against the visual image, and when there is a conflict, which could be caused by changes in the target's motion, or by a poor model of the target's trajectory, the pursuit system reverts to a purely visual feedback mode (Barnes and Asselman, 1991). Although most studies have been conducted with head-fixed tracking, Barnes and Lawson (1989) demonstrated that the combination of visual feedback and predictive processes also occurs during head-free tracking responses. Unless the target's motion can be modeled with complete accuracy, the predictive mechanism will exert varying degrees of influence on the visual feedback signal, and the response of the tracking system will be time-variant, moving through a continuum from time-delayed, visually driven responses to improved, predictive responses.

#### *VOR cancellation and suppression*

The semicircular canals and the otolith organs of the vestibular system are sensitive to head movement. Head movement will cause a reaction from the mechanoreceptors independently of whether the movement was caused by a disturbance to the head or was an active head movement to follow a target. The VOR cannot respond in the same manner to active motions and disturbances. During combined eye-head tracking, the VOR's response to the active head motions must be nulled for the gaze to follow the target's motion (Leigh and Zee, 1999). If this did not occur, then as the head was actively moved with the target, the eyes would be driven to oppose the head and would be moved away from the target. The VOR cannot be totally suppressed as this would not allow the eyes to be stabilized against disturbances. This implies that the vestibular system can control or limit the vestibular drive that is invoked by voluntary head motions when tracking a moving target (Leigh and Zee, 1999).

The current evidence suggests that there are two mechanisms that affect the VOR signal during head-free gaze tracking. It is likely that the same signal that drives the smooth pursuit response is used to cancel the VOR response to voluntary head motions (VOR cancellation).

Studies have also suggested the existence of a second mechanism called VOR suppression, which makes a direct adjustment of the VOR gain during head-free tracking. This mechanism may also partially disable the VOR during large (> 40 deg) eye-head gaze saccades (Leigh and Zee, 1999).

### **1.3 Traditional VOR testing methods**

Dizziness, vertigo, and balance disorders are among the most common reasons for visits to doctors' offices in the elderly. The Vestibular Disorders Association and the Johns Hopkins School of Medicine Center for Hearing and Balance both claim that 10% of all visits to doctor's offices are for dizziness (Johns Hopkins University Center for Hearing and Balance, 2003; Vestibular Disorders Association, 2003), and the American Academy of Otolaryngology states that more than 2 million people in the US visit their doctor each year for dizziness and balance problems (American Academy of Otolaryngology, 2003). Despite the high incidence of visits for dizziness and related dysfunction, there are few tools available to doctors that allow them to analyze the function of the vestibular system at the bedside or in an office. Of the tools that are available, none allow for a system-level analysis of the vestibular system. At the bedside, doctors and clinicians are limited to making observations about eye motions and a patient's static and dynamic equilibrium. Patients must be referred to specialized vestibular testing facilities to get a system level assessment of their visual and vestibular systems. The Mayo Medical Clinic, the Johns Hopkins School of Medicine, the Massachusetts Eye and Ear Infirmary, and the UCLA Medical Center, are a few of about two hundred clinics listed by the Vestibular Disorders Association as facilities with specialized vestibular and balance laboratories.

At any of these clinics, patients undergo a common series of tests. The most frequent tests that are conducted at the Mayo Clinic, for example, are those that evaluate hearing, measure postural stability with moving platform posturography (MVP), assess oculomotor function using positional nystagmus and electronystagmography (ENG), and assess the vestibular system

response using rotational tests (Mayo Clinic, 2003). Less common tests that may also be performed are MRIs of the brain and structures around the inner ear, CT scans of the temporal bones, blood tests, and cardiac function tests. Of these, the ENG and rotation tests are the main quantitative procedures used to evaluate the vestibular system.

The auditory test is conducted to examine the function of the inner ear, as problems with the cochlea, the hearing apparatus in the inner ear, often accompany disturbances in the function of the vestibular organs. The tests are far more extensive than the hearing tests that are done during a routine physical examination, and include evaluating how sound signals move from the ear to the auditory nerve (electrocochleography), how signals travel from the ear to the brain and within the brain (auditory brainstem response), and how well hair cells within the cochlea are functioning (otoacoustic emissions) (Vestibular Disorders Association, 2003).

Moving platform posturography evaluates a patient's ability to maintain balance. The patient stands on a platform that can move and tilt, and looks at a visual image on a screen that surrounds the patient and that can also move. During the tests, the shifts in the center of gravity that the patient makes to try to maintain balance are recorded. These tests are designed to give the physician information about what parts of the postural control system the patient relies upon most: visual, vestibular or proprioceptive. Although moving platform posturography tests are frequently used to evaluate patients with balance disorders, there is evidence that posturography is insensitive to vestibular disorders, and that it is not an adequate test for vestibular function (DiFabio, 1995).

The positional nystagmus test is a postural maneuver that is done to determine if the patient has benign paroxysmal positioning vertigo (BPPV). The test is also known as the Dix-Hallpike maneuver. BPPV is a disorder caused when calcium carbonate otoconia are dislodged from the utricle or saccule and get trapped in one of the semicircular canals. When the person moves a certain way, the crystals move and this results in the person feeling that they are spinning.

The ENG and the rotational tests are the most frequently performed VOR tests used by vestibular laboratories to assess vestibular function (Furman and Cass, 1996; Wall and Furman, 1989). They evaluate the vestibular system through the eye motions of the VOR, and have been the principal tools from which the system response characteristics of the vestibular system have been understood. The ENG test refers to a series of procedures that measure eye motions that are made by a patient in response to visual targets, changes in body position, and thermal stimuli. Rotational tests directly measure the response of the patient's VOR to rotations. Most often, patients are seated in motorized chairs that rotate them in the dark or as they watch a visual target. Head-only rotational testing (HORT) is an alternative to the rotational chair, and is gaining popularity. In HORT patients actively move their heads rather than having their whole body rotated. HORT is subject to the same analytic shortcomings as rotational-chair testing and has yet to be accepted as a clinically relevant testing method (Hirvonen *et al.*, 1999).

#### *ENG; caloric tests*

ENG tests measure nystagmus movements of the eyes which can indicate the presence of vestibular dysfunction. Although ENG tests are the most popular procedures administered to patients with vertigo, dizziness, and balance disorders, neither the test battery nor the testing methods are standardized (Vestibular Disorders Association, 2003). All ENG testing includes examinations for spontaneous vestibular nystagmus, an "oculomotor screening battery" that includes an assessment of saccades, smooth pursuit, and optokinetic nystagmus, and tests for positional nystagmus, which is caused by the otolith organs and may occur when the patient is moved into a particular position (Furman and Cass, 1996).

Caloric testing is also part of the series of ENG tests and is considered the "mainstay of vestibular laboratories" (Furman and Cass, 1996), and the "gold standard for diagnosis of ear disorders that affect one ear at a time" (Hain, 2002). It uses water, or air, to establish a thermal gradient across the horizontal semicircular canals, which causes convection current to develop in

the canal, pushing the endolymph against the cupula. In a healthy subject, irrigation with warm water will induce a slow component eye movement away from the irrigated ear, with a subsequent quick nystagmus that beats the eye back toward the irrigated ear. Irrigation with cool water will produce similar eye motions, but in the opposite direction. The slow component eye motion is the vestibular portion of the response and its magnitude reflects the amplitude of the vestibular response.

The performance of the right and left semicircular canals can be quantitatively compared by calculating the “percent reduced vestibular response”. The maximum velocity of the slow component motions for the warm and cool irrigations are summed for each ear, and then the velocity sums for each ear are subtracted from each other. This difference is normalized by dividing it by the sum of the maximum slow component velocities for all the cool and warm irrigations. This quantity is then converted to a percentage. A test subject’s response is considered abnormal if the “percent reduced vestibular response” exceeds a threshold value, usually set at around 25 – 30%. Unfortunately, there is great variability in a subject’s response to irrigation and therefore in the subsequent “percent reduced vestibular response”. This makes it very difficult to compare trials from the same subject and to assess improvement or decline. It is also difficult to establish a standard threshold value from which to compare a subject’s “percent reduced vestibular response”.

A limitation to the caloric test is that it gives a signal, and not a system response. The test is only an indicator of how the system responds to the thermal stimulus, which is the equivalent of approximately a 0.03 Hz motion stimulus (Goebel *et al.*, 2000). The response to this stimulus is not necessarily a good predictor of how the vestibular system would respond at any other input frequency. Even though the caloric test provides a poor quantitative measure of vestibular performance, it has, and does serve as a useful diagnostic tool. Its main advantage is that it allows for the left and right vestibular labyrinths to be tested separately. It is also easy to administer, well tolerated, and can be performed at a person’s bedside as well as in a laboratory.



### *Rotational-chair testing*

Rotational-chair testing evaluates the vestibular system by rotating test subjects through a series of rotations, usually while their body and head are rigidly fixed to a large motor-driven chair (Figure 1.2). The eye responses that result from the rotations are measured. Rotational-chair testing has been the main tool used to quantify the system level response of the vestibular system, but the size and expense of the equipment prevent it from being used outside specialized laboratories. The rotations are conducted about the earth-vertical axis and stimulate primarily the horizontal semicircular canal, and the horizontal, rotational VOR. It is most common for patients to be rotated by sinusoidal accelerations and constant velocity ramps (trapezoidal), but some research studies have used other inputs, such as pseudo-random binary sequences, to create unpredictable inputs with wide bandwidths (Wall *et al.*, 1978). Rotations are typically




**Figure 1.2:** Rotational chair (Hain, 2002).

conducted under three visual conditions: with a visual target that moves with the chair and remains directly in front of the test subject; with a visual target that is stationary (earth fixed); and with no visual target. The latter can include testing in the dark, with opaque goggles, or with the test subject's eyes closed. The rotational tests stimulate both canals simultaneously, one is

excited, the other inhibited, and so they do not provide information about each canal as is provided by caloric tests. Instead, rotational tests provide information about how the central nervous system processes vestibular information, and how it is able to integrate signals from the visual and vestibular sensory systems (Furman and Cass, 1996).

A sinusoidal rotation of a patient in the dark is the standard test that is used to measure the default response of the VOR. The intent of conducting the test in the dark is to eliminate the influence of the visual system on the responses of the eyes, and to stimulate only vestibularly driven motions. Typically, the patient is rotated back and forth in a smooth sinusoidal motion at low frequencies (0.01-1.0 Hz). Some laboratories sweep the amplitude of a single frequency so that the patient is subjected to several peak velocities. Eye motions are recorded during the rotations and to prevent the patients from corrupting their eye motions with voluntary movements to imagined targets, they are asked simple questions, or to do simple math, to keep them attentive. In a healthy subject, the rotations will induce a nystagmus response. For example, when the chair moves to the left, a slow phase VOR component will move the eyes to the right, and a quick phase motion will reset the eyes to the left. In order to analyze the VOR eye motions, it is necessary to remove the quick components, and to piece together the slow phase motions as shown in Figure 1.3 (Furman and Cass, 1996). The VOR gain is defined as the ratio of the VOR eye velocity to the patient's rotational velocity, and so the VOR eye velocity must be found by taking the derivative of the constructed slow phase response. A sinusoid is then fitted to the VOR eye velocity so that the variance and noise in the calculated response can be ignored. The VOR gain is estimated by dividing the peak of the best fit sinusoid by the peak velocity of the rotation. The phase response is also estimated by comparing the peaks of the input and output sinusoids. "Perfect" VOR compensation occurs when the eyes move with an angular velocity equal and opposite to that of the head's disturbance velocity, and results in a gain of 1 and a phase of 180 degrees. The standard convention is to bias the phase to 0, so that a gain of 1 and a phase of 0 represent perfect compensation. A crude estimate of a patient's VOR frequency response function

can be made by calculating the gain and phase relations at different frequencies, and then “connecting the dots”.



**Figure 1.3:** Reconstructed slow-phase eye position from a rotational chair test (Leigh and Zee, 2003).

Visual-vestibular interactions are tested by asking the patient to follow a visual target that rotates with the chair during their earth-fixed rotation, and by having the patient focus on a visual target that remains stationary relative to the earth as the patient is rotated. These tests evaluate how the visual system either augments or reduces the VOR response. In a typical series of tests, the patient will be rotated at a single frequency and tested in the dark, with a target fixed relative to the chair, and with a target fixed relative to the earth.

#### *Head-only rotational testing (HORT)*

Head-only rotational testing (HORT) has been proposed as a convenient alternative to rotational-chair testing (Hirvonen *et al.*, 1999; O’Leary and Davis, 1998). Rather than stimulating

the vestibular system with full body rotations created with a rotational chair, during a head-only rotational test, patients will either actively rotate their head, or have their head passively rotated by a technician. The visual conditions are usually the same as with rotational-chair testing. A periodic, audible tone, such as that from a metronome, is sometimes used as a cue to encourage the subject to create head motions with the appropriate frequency content. Studies have found that subjects are able to create active head motions reliably up to 3 Hz, so the bandwidth of the test is broader than with rotational-chairs (Furman and Durrant, 1998). In contrast to the expensive, large, rotational-chair equipment, the head-only rotational test requires little instrumentation and is not limited to being conducted in specialized laboratories. The HORT suffers from problems of reliability, due to the uncontrolled nature of the voluntary head motion, and validity, because the head motions are often artificial and predictable. Since patients are actively creating their head motions, it is easier for them to create corresponding eye motions. HORT is a relatively new method, and its results have not yet been accepted for clinical diagnostic purposes, but it has been shown that in some cases HORT does provide results similar to those from rotational chair testing (Hanson and Goebel, 1998). Two commercially available active head rotation systems are the Vestibular Autorotation Test (VAT) (Vestibular Technologies Inc., Tampa, FL) and Vorteq (Micromedical Technologies, Chatham, IL).

## **1.4 System identification of the VOR during head-free tracking**

### **1.4.1 Criticism of classical testing**

The results from classical rotational and caloric tests suggest that the VOR is not able to stabilize vision adequately when the head moves during natural activities. Most laboratories report that the gain of the VOR is 0.75 or lower in people with apparently normal VOR function (Leigh, 1996; Collewijn, 1989). For clear vision, the gain of the VOR would need to be between 0.9 and 1.1, which is what Grossman and Leigh observed for the VOR that was measured as

subjects walked in place (Grossman *et al.*, 1989). Some researchers have argued that the VOR is indeed inadequate to stabilize vision by itself, and that for an image to remain stable on the retina, the VOR must be supplemented by visual inputs, the cerviculo-ocular response (COR), and the ability to predict the head's disturbance and the visual target's motion (Demer, 1996). Leigh (1996) has proposed that only the VOR is capable of responding to the unpredictable, high frequency, higher speed disturbances that are characteristic of head motions during locomotion. The visual system has too long a latency, and in people with healthy vestibular systems the COR makes a negligible contribution to gaze stability. Furthermore, the disturbance motions of the head are random, and therefore cannot be anticipated. If gaze is stabilized during high frequency head disturbances, then this is accomplished by the VOR. Laboratory tests that do not reflect this ability of the VOR are inadequate, and have underestimated the performance of the system (Leigh, 1996).

The artificial constraints of the classical tests may have prevented them from capturing the functional behavior of the VOR. The experimental conditions imposed during rotation and caloric testing do not resemble the natural conditions under which the VOR serves to stabilize gaze. Under normal conditions, the vestibular and visual system work together to direct and stabilize gaze so that visual targets can be tracked and clearly seen. The two systems can be expected to influence each other under all conditions, as they are not independent and share many neural pathways. Classical rotational protocols are based on the assumption that the visual and the vestibular systems can operate independently, and that by removing the visual target or by delivering a very limited visual stimulus, the interactions between the systems may be mitigated. Unfortunately, the types of visual stimuli used in classical tests allow the subjects to influence their visual responses and so affect the experimental results. In addition to reporting that the VOR gain is lower than would be expected from a healthy VOR response, laboratory tests often report that the gains varied significantly between trials (Wall and Furman, 1989), which implies that the VOR response was influenced by the test subject. In healthy adults the 'default' gain measured in

the dark varies between 0.5 and 0.8, while the gain of the VOR will often rise to 1.0 when tracking an earth-fixed target, or drop to as low as 0.3 when tracking a head-fixed target. Leigh (1996) argues that the frequencies at which rotational tests are conducted, and the conditions under which the tests are conducted, particularly in darkness, are important differences between laboratory and natural conditions. The rotational-chair and the caloric tests stimulate the VOR at frequencies that are much lower than the 0.5 to 5.0 Hz frequency band that is typical of normal head disturbances. The performance of the VOR at low frequencies may not represent what is required of the VOR at higher frequencies. The low VOR gain values may also be directly linked to testing subjects in the dark. Studies have indicated that the VOR performs only at the level required to stabilize vision (Leigh, 1996). If the goal of the VOR is to stabilize the retinal field when there are both visual and vestibular stimuli, then what is its goal when the visual stimulus is removed? Another factor that may affect the function of the VOR and its reported gain is that the sinusoidal rotational stimuli used in rotational-chair tests are predictable. This may result in the test subject generating anticipatory eye motions that mask the actual performance of the VOR, and these cannot be corrected for in the analysis.

Although rotational-chair and caloric testing have provided many of the findings used to characterize the VOR, and have proved invaluable for detecting vestibular dysfunction, limitations in their testing protocols suggest that they do not provide data that reflects the VOR's performance under natural conditions. Leigh (1996) concludes his discussion of the problems with classical testing with his view of how testing of the VOR should be conducted:

“It is unfortunate that, because of technological constraints, it has been possible to study the VOR only during relatively artificial conditions; these studies have, to some extent, misrepresented the role of the VOR during natural activities. As methods for reliably measuring the VOR during locomotion become available, we may improve our ability to account for patients symptoms during natural activities. Until that is possible, it makes sense to use laboratory stimuli that correspond to the head disturbances that occur during natural activities such as locomotion:

rotations with a frequency range of 0.5 to 5.0 Hz, applied in a nonpredictable sequence.” (Leigh, 1996).

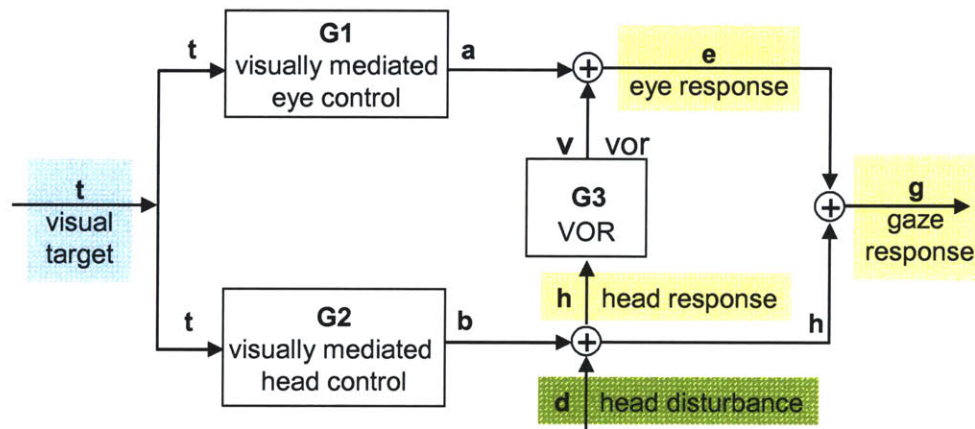
#### 1.4.2 System identification of VOR through visual vestibular co-stimulation

It is apparent from the preceding section that the best way to assess the performance of the VOR is to conduct an evaluation during natural activities such as walking, but obviously there are technical difficulties associated with doing so. It is unfortunate that testing during normal behavior is not usually the most efficient or even possible way of evaluating the dynamic characteristics of a system’s response. For the VOR, the evaluation should involve measuring it as it performs its natural function, which is to stabilize a person’s gaze while he or she tracks a moving target. The visual and the vestibular systems should both be stimulated, and the subject should be able to use natural head and eye motions to follow the target, so that the normal physiological interactions between the tracking and stabilization systems occur. In this type of protocol, external visual and vestibular inputs would have to be delivered in a manner that would allow the vestibular system’s influence on eye motions to be interpreted from the overall eye motions that are used to track the target.

##### *System identification of the VOR during gaze tracking*

The procedure adopted for evaluating the VOR during head-free gaze tracking can best be understood by considering a simplified block diagram representation of the gaze tracking system (Figure 1.4). The system represents the case where the target’s image is presented at optical infinity, and both eyes respond identically to the combined visual and vestibular stimulations. From the block diagram, and from the linear system equations that relate inputs and outputs through the subsystems, it is apparent how an experiment has to be run, and the inputs controlled, so that impulse and frequency response functions can be estimated for the VOR system. The target is the obvious experimental input that needs to be controlled to test the

function of the gaze-tracking system. The other external input to the gaze system is the disturbance motion applied to the head. Almost by definition, disturbances are an uncontrolled input to the system. It was realized that if the disturbance was controlled as an external input then the system can be viewed as a two input (target and disturbance), two output (head and eye) control system, and a response function for the VOR could be identified.



**Figure 1.4:** Block diagram of the gaze tracking system. The visual target ( $t$ ) stimulates responses from the visually mediated eye ( $G1$ ) and head ( $G2$ ) control systems. The VOR ( $G3$ ) creates eye motions ( $v$ ) that compensate for head disturbances such that the target's image is stabilized on the retina.

The purpose of the gaze tracking system is to direct the gaze towards the target. The visual target ( $t$ ) provides a clear goal for the output of the gaze tracking system ( $g$ ). It does not matter, for now, whether this input is considered the target's position, velocity, or an actuation signal based on the error between the gaze and the target. The target acts as a direct input to the visually mediated head and eye control systems, which are under the subject's voluntary control.  $G1$  represents a mapping from the visual goals to the eye trajectory, and  $G2$  is a mapping from the visual goal to the head trajectory. The output of  $G2$  is a head motion ( $b$ ), presumably towards the target. The output of the visually mediated eye control system,  $G1$ , is an eye motion ( $a$ ) measured relative to the head. In this simplified representation, feedback and interactions between the systems that are based on visual signals are assumed to be embedded in the blocks that



represent the visually mediated systems. This includes the 100 ms delay required to process the visual image. Ideally, the output of the eye and head systems would sum to be the gaze response, and the gaze would closely match the target's position. However, if the head position is disturbed from its intended position by disturbances (d), then the total motion of the head (h) is not what was desired by the visually mediated system, but is the sum of the disturbance (d) and the head's intended motion (b). Since eye position is relative to the head, the eye motions that were the output of the visually driven eye-control system will not sum with the disturbed head motion to produce a gaze that will closely match the target. The VOR must stabilize the gaze by creating a VOR eye motion that compensates for the head disturbance. The VOR eye motion sums with the visually driven eye motion to create the total, disturbed eye motion. If the systems are functioning well, the disturbed eye and head motions sum to a gaze response that closely matches the target's movement.

The linear system equations show that consistent, unbiased estimates for both the VOR and visual tracking dynamics are possible, with appropriate control over both available stimuli. The following equations are shown in the Laplace domain, and the systems are assumed to be linear. There are potentially two controlled inputs, the target trajectory (T) and the head disturbance (D). The head trajectory (H) and the eye trajectory (E) are the outputs that can be measured.

The total eye response is due to contributions from the visual system (A) and from the vestibular system (V)

$$E(s) = A(s) + V(s) = T(s) \cdot G1(s) + H(s) \cdot G3(s). \quad \text{Eq. 1.1}$$

This expresses the eye response in terms of the target and the head motion. The two signals can be measured, but the VOR response function,  $G3(s)$ , cannot be solved from this equation. The head response can be easily related to the independent target and disturbance inputs (D):

$$H(s) = B(s) + D(s) = T(s) \cdot G2(s) + D(s), \quad \text{Eq. 1.2}$$

which can then be substituted into Eq. 1.1. Because the relations are linear, Eq.1.2 can be simplified to:

$$E(s) = T(s) \cdot [G1(s) + G2(s) \cdot G3(s)] + D(s) \cdot G3(s). \quad \text{Eq. 1.3}$$

This equation shows how each of the external inputs to the gaze system, the target trajectory which is the input to the visual system, and the disturbance, which is the input to the vestibular system, affects the measured eye motion.

To solve for the VOR response function,  $G3(s)$ , the equations must be expressed as auto and cross-power spectra, the Laplace transform of the auto and cross-correlation functions. If the cross-power spectra are formed between the signals of Eq.1.3, which are  $E(s)$ ,  $T(s)$ , and  $D(s)$ , and the disturbance, then the equation becomes:

$$S_{ED}(s) = S_{TD}(s) \cdot [G1(s) + G2(s) \cdot G3(s)] + S_{DD}(s) \cdot G3(s). \quad \text{Eq. 1.4}$$

If the disturbance is selected optimally so that it is uncorrelated with the target trajectory, then the cross-power spectrum,  $S_{TD}$ , is zero. The transfer function for the VOR can then be estimated as:

$$G3(s) = \frac{S_{ED}(s)}{S_{DD}(s)}. \quad \text{Eq. 1.5}$$

Eq 1.5 shows the VOR as the transfer function between the perturbations and the eye motions that were linearly related to the head's disturbance. It mathematically excludes the eye motions that were driven by the visual target. The visual target still influences the head and eye responses, as the subject tracks the target and so normal physiological interactions occur. When the input sequences are optimally selected, the mathematics allow for the input and output related to the disturbance to be distinguished from those influenced by the visual target.

### *Experiment protocol*

An experiment must be well controlled to create the conditions that enable the analysis to hold. The visual target must travel with a randomized trajectory so that its motion cannot be

predicted, and be designed so that it can be tracked primarily by smooth pursuit. This should result in a clear linear relation between the target motions and gaze tracking and VOR stabilization. The target must be projected at optical infinity to promote conjugate eye motions. While a test subject tracks the target, small stochastic torque perturbations will be applied directly to the head through a head-perturbing device. The torque perturbations must be large enough to create head motions that require a measurable corrective VOR motion, but not so large that the subject is prevented from using natural head motions to track the target. The disturbance motions cannot be explicitly controlled because that would prevent the subject from using the head as desired to track the target. The subject's neck muscles will provide the majority of the torque required to direct the head towards the target, and the perturbations will be just large enough to perturb the head about its natural operating points. In order for G3 to represent the system level dynamics of the VOR, it is necessary for the experiment to stimulate appropriate dynamics from the visual and vestibular systems. Both inputs must span the full range of the visual and vestibular systems, that is, the visual target must have a frequency content from 0 to about 2 Hz, which represents the bandwidth of the visual tracking system and the disturbance motions should have a bandwidth of 0.5 to 3.0 Hz if they are to represent horizontal plane disturbance, or out to 5.0 Hz if they are to represent pitch plane dynamics. Most importantly, the input sequences must be statistically uncorrelated so that the eye motions created by the visual system and those created by the vestibular system can be distinguished.

To keep the analysis as simple as possible, and to minimize the complexity of the equipment that had to be designed to execute the VOR experiments, the protocol was limited to testing the rotational VOR in the horizontal plane. By restricting the protocol in this manner, only horizontal, rotational motions of the eyes and the head had to be measured during the experiments, and the perturbations applied to the head only had to create angular disturbances in the horizontal plane. It is standard practice in clinical evaluations and in research to investigate the rotational VOR exclusively in one plane, either the horizontal or the vertical, as it is much

more difficult to identify the relation between head and eye motions when there are multiple degrees of freedom in the head and eye movements. Because the semicircular canals are not aligned exactly with the horizontal and vertical planes, a horizontal plane rotation of the head will stimulate more than just the lateral (horizontal) semicircular canals. A healthy CNS will resolve the responses of the canals so that only a horizontal eye motion is created in response to a horizontal head disturbance. In the remainder of this thesis “VOR” refers primarily to the horizontal, rotational VOR.

## **1.5 Research goals and thesis organization**

The goal of this research was to develop the experimental equipment, protocol, and analysis algorithms to conduct a system level evaluation of the VOR while it was being used to stabilize gaze during natural, head-free tracking of a visual target. The research can be divided into three main phases: 1. The design and construction of the experimental equipment; 2. The development of the experimental protocol, appropriate input stimuli, and analysis algorithms; and 3. The validation of the system, protocol and analysis through testing with human subjects.

The organization of this thesis reflects these three phases. Chapter 2 explains the mathematics behind the linear system identification techniques. It was found during the research that the gaze tracking and vestibular interaction were modeled well by linear systems. Impulse and frequency response functions were estimated for the VOR and for the visual tracking systems. The calculations were done using the cross-correlation relation between the inputs and outputs that could be measured during the experiments. The technique is explained, as are the requirements that had to be met by the experimental protocol for the analysis to hold. In Chapter 3 the design and construction of the experimental equipment is discussed. Five main systems were required to implement the VOR experiments: a head-perturber, a moving visual target, a sensor to measure head motion, a sensor to measure eye motion, and a computer system for

experimental control, data acquisition, and analysis. The goal set for the experimental apparatus was to create laboratory equipment with the performance and flexibility required to conduct all the experiments necessary to develop the techniques, as well as being used for the final VOR experiments. The head perturber, the visual target, the eye motion sensor, and the software algorithms were developed specifically for this research. The head motion sensors and the data acquisition system were purchased from commercial sources.

The design of the experiments and the development of the appropriate input stimuli are presented in Chapters 4, 5, and 6. Chapter 4 identifies the criteria that had to be met when stimulating the visual and vestibular systems, and presents the results of experiments that were conducted to identify the appropriate characteristics of the torque and target sequences. Chapter 5 presents a novel mathematical technique that was developed for this research to create uncorrelated torque and target sequences. Chapter 6 describes the experiment on the dynamics of the head-neck system that was written and published in the *Annals of Biomedical Engineering*. It describes the head-neck system as a quasi-linear system with dynamics that are dependent on the mean torque applied to the head. Chapter 7 brings together all the elements associated with the development of the experimental system and discusses the results of the experiments that were conducted to evaluate the VOR in healthy subjects. Chapter 8 concludes the thesis with a summary of the results, and recommendations for future work.



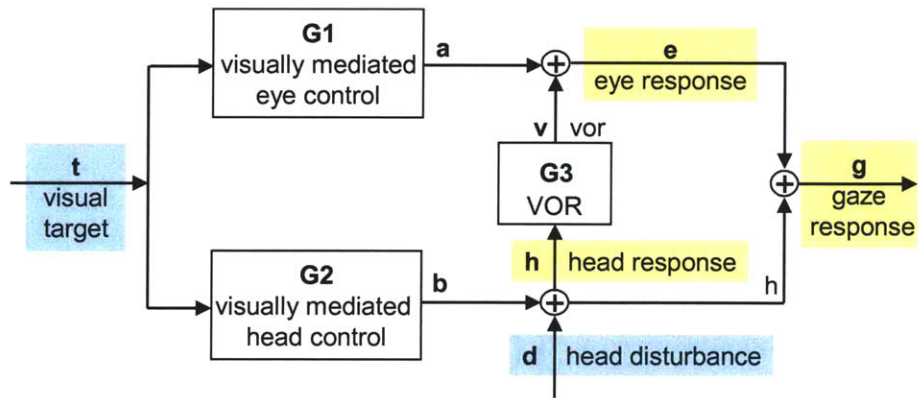
## **Chapter 2**

### **Mathematics of the system identification**

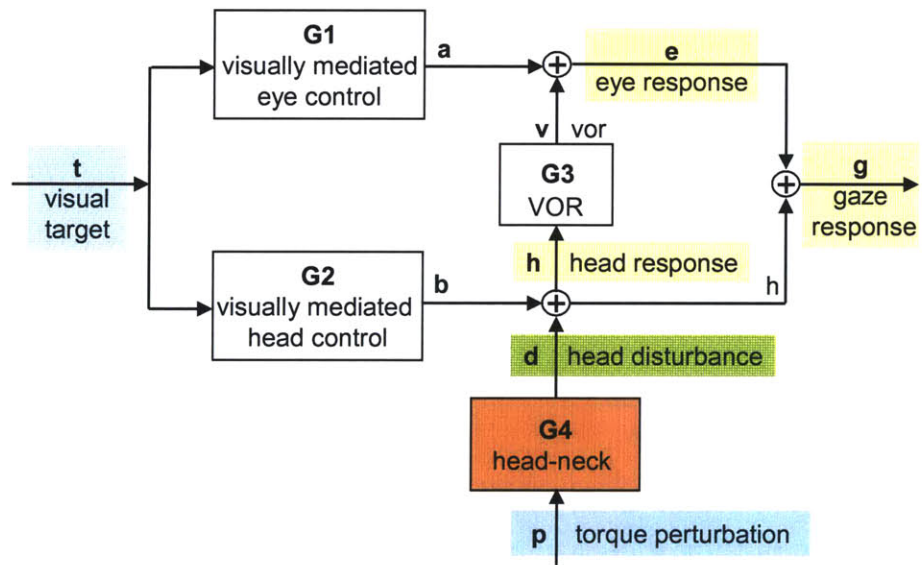
- 2.0 Chapter summary
- 2.1 Practical head disturbances
- 2.2 Linear equations that describe the input-output relations
- 2.3 Correlation based convolution equations and the deconvolution process

#### **2.0 Chapter summary**

As described in Chapter 1, an analysis of the VOR can be made during natural, head-free tracking when the visual target and the disturbance to the head's position are controlled so that they stimulate a wide range of visual and vestibular responses and so that they are not correlated. The disturbance position of the head cannot be explicitly controlled, but must be created by applying a torque perturbation to the head. Because of this, the response of the head-neck system to torque perturbations must be modeled in the block diagram of the gaze tracking system (Figure 2.1 and 2.2). The equations that describe the interactions between the controlled inputs and the subsystems of the gaze tracking system that must be solved to identify the system dynamics of the VOR are thus slightly different from those presented in Chapter 1. In Chapter 2 these equations are developed, and the mathematics of the correlation-based, system identification techniques that were used to calculate the VOR response functions are explained.



**Figure 2.1:** Block diagram of the gaze tracking system.



**Figure 2.2:** Block diagram of the gaze tracking system. A torque perturbation ( $p$ ) is applied to the head to disturb the head's position.



## 2.1 Practical head disturbances

The dynamics of the VOR can be evaluated during head-free gaze tracking only when the disturbance position of the head ( $d$ ) is not correlated with the trajectory of the visual target ( $t$ ). The block diagram that was used to describe the gaze tracking system, and from which the equations that were used to identify the dynamics of the VOR were developed, is shown again in Figure 2.1. When the analysis of the VOR was described in Chapter 1, the disturbance motion of the head was treated as a controlled and measurable quantity. This, however, is not the case if subjects are allowed to control their head movements voluntarily to track the visual target, which is a requirement of the experimental protocol. The disturbance position of the head cannot be explicitly controlled by the experimental protocol, nor can the disturbance be measured independently of the head motion that was produced by the subject. Subjects must be able to move their heads as they desire, and the disturbance ( $d$ ) must be a small deviation to the intended motion of the head ( $b$ ). To accomplish this, the disturbance must be created by applying a small torque perturbation to the subject's head. The amplitude of the torque must be large enough to disturb the head and to create a measurable VOR eye motion, but it cannot be so large that it prevents subjects from using the head as they desire.

When the disturbances to the head's position are created by an applied torque perturbation, the block diagram that describes the gaze tracking system is changed from that in Figure 2.1 to that in Figure 2.2. A block that represents the head-neck system's reaction to the torque perturbation is added ( $g_4$ ), and the applied torque perturbation ( $p$ ), rather than the disturbance motion of the head ( $d$ ), becomes the controlled experimental input. Instead of the requirement that the disturbance and the target's trajectory be uncorrelated, the torque perturbation sequence and the target must be made uncorrelated. The addition of the head-neck system complicates the linear equations that were developed to solve for the VOR response function in Chapter 1, and makes the linear identification of the VOR dependent on the response

of the head-neck system. The head-neck system must be well modeled as a linear system. Nonlinearities in the head-neck system will degrade how well the measured eye motions can be linearly associated with the controlled experimental inputs, and will affect the identification of the VOR's dynamic response.

## **2.2 Linear equations that describe the input-output relations**

The dynamics of the VOR, and of the other systems that comprise the gaze tracking system, were modeled using linear, impulse response functions. Consistent with most biological systems, these systems were not expected to be strictly linear and so it was understood that the linear models would not capture all the dynamics of the systems. The linear system analysis, though, allowed for equations to be established from which the VOR could be identified as it functioned within the multi-input, multi-output gaze tracking system. It is also standard practice to conduct a system identification with linear models first, and then to determine if, and where, the system being assessed would benefit from a nonlinear analysis. How well the linear models had captured the behavior of the systems was evaluated by using the linear models to make predictions of the measured outputs, and by using the coherence-squared function, which quantifies the linearity and noise content of the input-output data (Juang, 1994; Marmarelis, 1988). Although these biological systems were not strictly linear, it was found that the linear models did, in fact, capture the behavior of the systems very well within the bounds of the VOR experiments. This is discussed in detail in Chapters 4, 6, and 7.

When the target's trajectory and the torque perturbations were controlled so that they were not correlated, impulse response functions (IRFs) and frequency response functions (FRFs) for the VOR, as well as for the head's visually mediated control system (g2) and the head-neck system (g4), could be calculated directly from measurements of the controlled inputs and the observable outputs, which are the eye (e) and the head (h) motions. The equations that describe

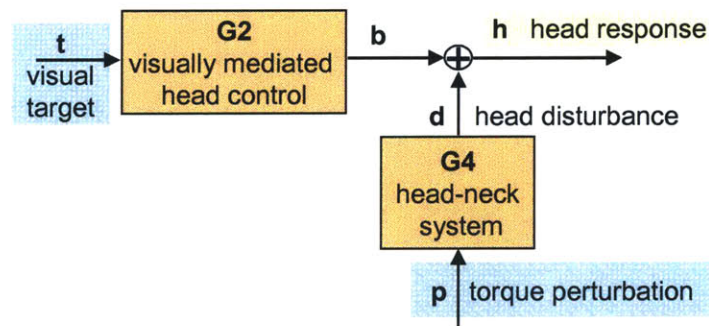
the input-output interactions among the systems will be developed in this section in the time domain, but are viewed as easily in the frequency domain. In the frequency domain, a linear convolution is replaced by a multiplication, the auto-correlation function is replaced with a (auto) power spectral density, and the cross-correlation function is replaced by the cross power spectral density (Oppenheim *et al.*, 1996). To identify the IRF of the VOR, it was necessary to model both the eye and head responses to the torque and target inputs.

### Head control systems

The head response ( $h$ ) can be viewed as the combined output of two head control subsystems. The region of the gaze tracking block diagram that illustrates the head control systems is shown in Figure 2.3. The torque perturbation ( $p$ ) stimulates a response from the head-neck system ( $g_4$ ) and the visual target ( $t$ ) stimulates a response from the visually mediated head control system ( $g_2$ ). The responses of these systems, along with noise, sum to create the measured head response ( $h$ ). The linear convolution equation that relates the head velocity to the inputs is:

$$h = t \otimes g_2 + p \otimes g_4 + n, \quad \text{Eq. 2.1}$$

where “ $\otimes$ ” represents the linear convolution operator.



**Figure 2.3:** The head control systems.

The analysis of the head control systems was made easy because the torque and target inputs were designed to be uncorrelated. The IRF for the head-neck system ( $g_4$ ) can be isolated,

and then solved for, by forming the cross-correlation of the signals in Eq. 2.1 with the torque perturbation. When this is done the linear convolution equation becomes:

$$r_{PH} = r_{PT} \otimes g_2 + r_{PP} \otimes g_4 + r_{pn}, \quad \text{Eq. 2.2}$$

where the  $r_{PH}$  is the cross-correlation function (CCF) between the torque perturbation and the head signal,  $r_{PP}$  the auto-correlation function of the torque perturbation,  $r_{PT}$  the cross-correlation function between the torque and the target, and  $r_{pn}$  the cross-correlation between the torque and the noise. Under the assumptions that the noise is uncorrelated with the input sequences, and the torque and target inputs are, by design, uncorrelated, Eq 2.2 can be reduced to:

$$r_{PH} = r_{PP} \otimes g_4. \quad \text{Eq. 2.3}$$

The impulse response function for the head-neck system can then be estimated by deconvolving  $r_{PP}$  from  $r_{PH}$  via a matrix inversion. The steps involved in the deconvolution process are explained in Section 2.3

A similar development can be used to identify the IRF for the visually mediated head-control system ( $g_2$ ). Instead of forming the cross-correlation of the signals in Eq. 2.1 with the torque input, the cross-correlation functions are formed with the target sequence. Since the torque-target CCF and the noise-target CCF can be approximated as zero, the resultant convolution relation is:

$$r_{TH} = r_{TT} \otimes g_2. \quad \text{Eq. 2.4}$$

The IRF for the visually mediated head response ( $g_2$ ) can be estimated by deconvolving  $r_{TT}$  from  $r_{TH}$ .

### *VOR and eye response*

The IRF for the VOR is identified from the equations that model the eye response. The eye response ( $e$ ) represents the sum of eye motions created by the visually mediated eye control system ( $g_1$ ), eye motions from the VOR ( $g_3$ ), and some noise ( $n$ ) (Figure 2.2):

$$e = a + v = t \otimes g1 + h \otimes g3 + n . \quad \text{Eq. 2.5}$$

Since the head response (h) is stimulated, in part, by the visual target (t), the head signal is not uncorrelated with the target. The response function for the VOR cannot, therefore, be calculated directly from Eq. 2.5 by forming the cross-correlation of the signals in the equation with the head sequence. The eye motion needs to be expressed in terms of the uncorrelated torque and target inputs, and the systems upon which they act. To do so, Eq. 2.1 is used to replace the head motion sequence (h) in Eq. 2.5 with the components of the head motion that were driven by the target and the torque inputs (Eq. 2.1):

$$e = t \otimes g1 + (t \otimes g2 + p \otimes g4) \otimes g3 + n . \quad \text{Eq. 2.6}$$

Since the convolution is a linear operator, Eq. 2.6 can be rearranged to show how the controlled inputs stimulated a cascade of systems that ultimately created the eye motion. This produces:

$$e = t \otimes g1 + t \otimes g2 \otimes g3 + p \otimes g4 \otimes g3 + n . \quad \text{Eq. 2.7}$$

The eye motions that were generated by the VOR to stabilize the gaze against the head motions stimulated by the torque perturbations are represented by  $(p \otimes g4 \otimes g3)$ . The eye motions that were stimulated by the visual target through the visually mediated control systems are represented by  $(t \otimes g1 + t \otimes g2 \otimes g3)$ .

The terms  $(t \otimes g2 \otimes g3)$  represent the response of the VOR ( $g3$ ) to head motions that were voluntarily made to follow the visual target ( $t \otimes g2$ ). The semi-circular canals of the vestibular system are stimulated just the same independently of whether the stimulus comes from a voluntary movement or a disturbance motion of the head, however, as discussed in Chapter 1, the response of the VOR to voluntary head motions is thought to be cancelled, or at least suppressed, by visually mediated mechanisms (Leigh and Zee, 1999). The gaze system functions so that the eyes are not driven equally and opposite to a head movement that is voluntarily generated to follow a target. Although the eye motion from  $(t \otimes g2 \otimes g3)$  is believed to be cancelled, the terms

are included in Eq. 2.7 because the gaze tracking system was not modeled to include components that explicitly cancelled these terms. It can be assumed that the VOR suppression and cancellation systems are among the visually mediated systems embedded in the g1 and g2 system blocks, and also that their direct effect on eye movements is correlated with the visual target. Whether or not these systems are included in Eq. 2.7 does not affect the ability of the analysis to capture the performance of the VOR in response to head disturbances. The torque and target inputs are uncorrelated, and so the analysis can differentiate between eye responses that were driven by the torque perturbation and the eye motions that were created by the visual target.

Eq 2.7 can be changed from representing a relation between eye responses and the torque and target inputs, to a relation that considers only system behaviors that were driven by the perturbation torque by forming the cross-correlation of the input and output sequences with the perturbation torque:

$$r_{PE} = r_{PT} \otimes g1 + r_{PT} \otimes g2 \otimes g3 + r_{PP} \otimes g4 \otimes g3. \quad \text{Eq. 2.8}$$

Since the torque perturbation sequence and the target velocity were made to be uncorrelated, their cross-correlation function can be set to zero ( $r_{PT} = 0$ ), and Eq. 2.8 can be reduced to:

$$r_{PE} = r_{PP} \otimes g4 \otimes g3 \quad . \quad \text{Eq. 2.9}$$

Eq. 2.9 shows the linear relation between eye motions that were correlated with the torque stimulus,  $r_{PE}$ , and torque perturbations that disturbed the head through the head-neck system ( $r_{PP} \otimes g4$ ). Although data were gathered under natural tracking conditions, during which the head and eyes interacted normally to track a moving target, the effects of eye motions that were stimulated by the visual target have been removed and only the interactions between head disturbances and eye motions remain.

Eq. 2.3,  $r_{PH} = r_{PP} \otimes g4$ , can be substituted into Eq. 2.9 to create an equation that relates the VOR's impulse response function ( $g3$ ) to the cross-correlation functions that represent the head and the eye motions that were driven by the torque perturbation:

$$r_{PE} = r_{PH} \otimes g_3. \quad \text{Eq. 2.10}$$

It is from this equation that the VOR's impulse response function ( $g_3$ ) can be identified by deconvolving  $r_{PH}$  from  $r_{PE}$ . Eq. 2.10 can be interpreted as showing the VOR response function as the relation between head motions that were correlated with the torque perturbation, and eye motions that were correlated with the torque perturbation. Functionally, this is the purpose of the VOR. During gaze tracking, the VOR creates eye motions that compensate for disturbances to the head's position, it does not need to correct for voluntary head motions.

### **2.3 Correlation based convolution equations and the deconvolution process**

The correlation-based, convolution equations and the deconvolution process that were used to identify the dynamics of the VOR are developed here to illustrate how the IRFs were estimated. The mathematics of representing system interactions with correlation-based convolutions, and identifying system response functions from the correlation measures, has been well developed in the literature. Texts such as those written by Ljung (Ljung, 1999), Juang (Juang, 1994), Eykhoff (Eykhoff, 1974), and Jenkins and Watts (Jenkins and Watts, 1968), formally develop the linear, correlation based, convolution equations that are solved to identify a system. The texts by Juang and Eykhoff are among the few that develop the discrete time set of equations.

#### *Correlation based, linear, causal convolution equations*

The development of the correlation-based convolution equations, and the deconvolution process from which an IRF can be estimated, can be demonstrated by describing the analysis of the head-neck system. The equations will be developed in a manner similar to that used by Juang (1994).

It is first necessary to define the correlation function estimates. The biased cross-correlation function estimate between sequence X and Y can be defined as:

$$r_{XYj} = \frac{1}{N} \sum_{n=j}^{N-1} X_{n-j} Y_n, \quad \text{Eq. 2.11}$$

and the biased auto-correlation function estimator as:

$$r_{XXj} = \frac{1}{N} \sum_{n=j}^{N-1} X_{n-j} X_n, \quad \text{Eq. 2.12}$$

where the sequences are of length N, and lags are greater than or equal to 0 (Jenkins and Watts, 1968).

The linear, causal, convolution relation between the head response (h) and the torque (p) and target (t) inputs that was expressed in Eq. 2.1 can be written for the head response at time  $n\Delta t$  as:

$$h_n = \Delta t \sum_{i=0}^M g_2 t_{n-i} + \Delta t \sum_{k=0}^M g_4 p_{n-k}, \quad \text{Eq. 2.13}$$

where  $\Delta t$  is the sampling period,  $n$  represents the index of the data record and ranges from 0 to  $(N-1)$ , and the head control systems that are represented by  $g_2$  and  $g_4$  have a maximum duration of  $M$  lags (Oppenheim *et al.*, 1996). For simplicity, noise is ignored, since it is uncorrelated with the inputs and will not be a factor once the cross-correlation functions are formed.

The convolution relation can be expressed directly with input and output signals or with correlation functions. A significant advantage of expressing the input-output behavior of a system with correlation functions is that the influence of noise on the linear system equations is mitigated. This helps the identification process capture the underlying dynamics of the system, rather than model the system behavior that is corrupted by measurement noise. To recast the convolution equation with correlations, Eq. 2.13 is multiplied by a lagged torque input,  $p_{n-j}$ , and averaged over time. When this is done, Eq 2.1 becomes:



$$\frac{1}{N} \sum_{n=j}^{N-1} p_{n-j} h_n = \frac{1}{N} \sum_{n=j}^{N-1} p_{n-j} (\Delta t \sum_{i=0}^M g_{2_i} t_{n-i}) + \frac{1}{N} \sum_{n=j}^{N-1} p_{n-j} (\Delta t \sum_{k=0}^M g_{4_k} p_{n-k}) \quad \text{Eq. 2.14}$$

By definition, the left hand side of Eq. 2.14 is the cross-correlation function between the input torque and the output head response,  $r_{ph}$ . There are two groups of terms on the right hand side of the equation. After a slight manipulation of the order of terms and the indices of the summations, the terms on the right hand side of the equation become the convolution of the  $g_2$  and  $g_4$  response functions with cross- and auto-correlation functions.

Consider, first, the group of terms after the “+”. This group of terms will be shown to be equal to the convolution of the torque auto-correlation with the  $g_4$  response function,  $r_{pp} \otimes g_4$ . Since a summation is a linear operator, the terms can be arranged to look more similar to the convolution summation:

$$\Delta t \sum_{k=0}^M g_{4_k} \left( \frac{1}{N} \sum_{n=j}^{N-1} p_{n-j} p_{n-k} \right) \quad \text{Eq. 2.15}$$

By letting  $v = n-k$ , the index  $n$  can be changed so that only one of the torque sequences ( $p$ ) is expressed with a lag. The new index  $v$  will range from  $(j-k)$  to  $(N-1-k)$ . After this change Eq. 2.15 becomes:

$$\Delta t \sum_{k=0}^M g_{4_k} \left( \frac{1}{N} \sum_{v=j-k}^{N-1-k} p_{v-(j-k)} p_v \right). \quad \text{Eq. 2.16}$$

Eq 2.16 can be seen to be equivalent to the convolution of head-neck IRF ( $g_4$ ) with the auto-correlation function of the torque sequence, as long as the torque sequence is ergodic. When the sequence is ergodic, which can be ensured (influenced) by the proper design of the experiment, the time shift ( $-k$ ) in the range of the index  $v$  does not affect the summation as a statistical estimator of the auto-correlation (Juang, 1994). Eq. 2.16 is therefore equivalent to the convolution of the head-neck system with the auto-correlation of the torque perturbation  $g_4 \otimes r_{pp}$ .

Similar steps can be followed to show that the terms before the “+” on the right hand side of Eq. 2.14 represent the convolution of the response function for the visually mediated head

control system,  $g_2$ , with the cross-correlation between the torque and the target,  $g_2 \otimes r_{pt}$ . Since the torque and target are designed to be uncorrelated, this part of the equation can be approximated as zero.

### Deconvolution

In Eq. 2.3 the impulse response function for the head-neck system,  $g_4$ , can be calculated by deconvolving the auto-correlation function of the torque from the cross-correlation between the torque and the head response. The deconvolution process is conducted by expressing the correlation-based convolution equation as a series of linear, matrix equations, and solving for the vector that represents the terms of the unknown impulse response function.

A set of  $(M+1)$  equations is formed from Eq. 2.3. For convenience the  $\Delta t$  terms that are part of the convolution equation are assumed equal to 1. Formally they should be expressed as a scalar times the identity matrix,  $\Delta t[I]$ . The matrix form of the convolution equations is:

$$\begin{bmatrix} r_{PH0} \\ r_{PH1} \\ \dots \\ r_{PHM} \end{bmatrix} = \begin{bmatrix} r_{PP0} & r_{PP-1} & \dots & r_{PP-M} \\ r_{PP1} & r_{PP0} & \dots & r_{PP-M+1} \\ \dots & \dots & \dots & \dots \\ r_{PPM} & r_{PPM-1} & \dots & r_{PP0} \end{bmatrix} \begin{bmatrix} g_4_0 \\ g_4_1 \\ \dots \\ g_4_M \end{bmatrix}. \quad \text{Eq. 2.17}$$

Since the auto-correlation function is symmetric about the zero lag value, the negative lag values in the auto-correlation matrix can be replaced by positive lag values so that:

$$\begin{bmatrix} r_{PH0} \\ r_{PH1} \\ \dots \\ r_{PHM} \end{bmatrix} = \begin{bmatrix} r_{PP0} & r_{PP1} & \dots & r_{PPM} \\ r_{PP1} & r_{PP0} & \dots & r_{PPM-1} \\ \dots & \dots & \dots & \dots \\ r_{PPM} & r_{PPM-1} & \dots & r_{PP0} \end{bmatrix} \begin{bmatrix} g_4_0 \\ g_4_1 \\ \dots \\ g_4_M \end{bmatrix}, \quad \text{Eq. 2.18}$$

which can be written as:  $\underline{R}_{PH} = [R_{PP}] \cdot \underline{G}_4$  Eq. 2.19

The  $(M+1 \times M+1)$  auto-correlation matrix,  $R_{PP}$ , is of a Toeplitz form (constant values along the negative diagonals). The values of the head-neck impulse response function, which are

in the  $G4$  vector, are estimated by inverting the auto-correlation matrix and multiplying the inversion through the input-output cross-correlation vector,  $R_{PH}$ , so that  $G4 = [R_{PP}]^{-1} R_{PH}$ .

To solve for an IRF that is  $M+1$  elements in length, the correlation function is required only at lag values from 0 to a maximum lag equal to the length of the system's impulse response function. The  $(M+1)$  values of the IRF are calculated from  $(M+1)$  equations formed using the correlation functions. Although  $N$ , input-output, data records were collected during an experiment (with the  $N$  usually much greater than  $M$ ) none of the experimental data is ignored when only  $M+1$  correlation-based equations are used to solve for the IRF. All the information about the system that was in the  $N$  experimental data records is incorporated into the  $M+1$  correlation-based convolution equations, since at each lag value,  $j$ , the correlation function is calculated using  $(N-j)$  elements of the experimental data (Eq. 2.11, 2.12).

#### *Equations for the VOR*

The identification of the VOR IRF was conducted in the same manner as that for the head control systems, but the convolution equation that described the input-output relation of the VOR lead to a different matrix relation. The estimate of the VOR impulse response function was found from Eq. 2.10, which was:

$$r_{PE} = r_{PH} \otimes g3. \quad \text{Eq. 2.10}$$

This equation expresses a relation between two cross-correlations, not between an output-input cross-correlation and an input auto-correlation, as was the case for the head control systems. Thus the matrix equations that had to be solved to identify the VOR IRF had a slightly different form from those that were formed for the head control systems. The matrix form of Eq. 2.10 is:

$$\begin{bmatrix} r_{PE0} \\ r_{PE1} \\ \dots \\ r_{PEM} \end{bmatrix} = \begin{bmatrix} r_{PH0} & r_{PH-1} & \dots & r_{PH-M} \\ r_{PH1} & r_{PH0} & \dots & r_{PH-M+1} \\ \dots & \dots & \dots & \dots \\ r_{PHM} & r_{PHM-1} & \dots & r_{PH0} \end{bmatrix} \begin{bmatrix} g3_0 \\ g3_1 \\ \dots \\ g3_M \end{bmatrix}, \quad \text{Eq. 2.20}$$

which can be written as: 
$$\underline{R_{PE}} = [\underline{R_{PH}}] \cdot \underline{G3}. \quad \text{Eq. 2.21}$$

Since the terms in the  $R_{PH}$  matrix are cross-correlation values, the negative lag values in the upper triangle are not equal to the positive lag values of the lower triangle as in the auto-correlation function. For the cross-correlation function, a negative lag of the CCF when sequence X is lagged relative to sequence Y, is equal to the positive lag of the CCF when Y is lagged relative to X, e.g.  $r_{XY-5} = r_{YX5}$ . Therefore, the matrix equations can be expressed exclusively with positive lag cross-correlation values such that:

$$\begin{bmatrix} r_{PE0} \\ r_{PE1} \\ \dots \\ r_{PEM} \end{bmatrix} = \begin{bmatrix} r_{PH0} & r_{HP1} & \dots & r_{HPM} \\ r_{PH1} & r_{PH0} & \dots & r_{HPM-1} \\ \dots & \dots & \dots & \dots \\ r_{PHM} & r_{PHM-1} & \dots & r_{PH0} \end{bmatrix} \begin{bmatrix} g3_0 \\ g3_1 \\ \dots \\ g3_M \end{bmatrix}. \quad \text{Eq. 2.22}$$

The values of the VOR IRF are estimated by deconvolving the torque-head cross-correlation from the torque-eye cross-correlation function by inverting the  $R_{PH}$  matrix and multiplying it by the  $R_{PE}$  vector.

#### *Alternate development of the equations and the coherence-squared function*

An alternate development of the equations that describe the gaze tracking system can be conducted to show that the solution of the correlation-based convolution equation is similar to finding the least square estimate of the impulse response function directly from the linear convolution equations between the input and output sequences. This is described in Appendix A.

Appendix A also includes the development of the coherence-squared function for a two-input system. The coherence-squared is an indispensable tool for assessing how well a linear analysis can capture the dynamics of a system that may be corrupted by nonlinearities and noise. Coherence squared functions for single-input, single-output systems are explained in many texts, but they are rarely seen for two input systems.

## **Chapter 3**

### **Development of the test equipment**

- 3.1 Introduction
- 3.2 Head perturbers
- 3.3 Visual target
- 3.4 Head tracker
- 3.5 Eye tracker
- 3.6 Data acquisition and experiment control

#### **3.1 Introduction**

Five subsystems of equipment were required to execute the experimental protocol: a head perturber, a visual target, a sensor to measure head motions, a sensor to measure eye motions, and a computer system for experimental control, data acquisition, and analysis. Since this type of testing of the vestibular and gaze tracking systems had not been done before, the specific criteria that had to be met by the equipment to execute a successful VOR experiment were not known. The need, therefore, was to develop equipment with sufficient flexibility to conduct the initial experiments that were required to determine how best to execute the VOR protocol, and with the appropriate characteristics for running the human experiments that would validate the final VOR protocol as an effective test of the VOR.

One of the original desires was to make the test equipment portable. Despite the high frequency of complaints of dizziness and disequilibrium in the elderly (American Academy of Otolaryngology, 2003; Vestibular Disorders Association, 2003), a portable, low cost system for

evaluating the system-level performance of the vestibular-ocular system does not, at present, exist. By making the system portable, vestibular testing could be conducted in the home, at the bedside in the hospital, or even as a routine test administered throughout the day to verify the spatial awareness of people in high risk occupations, such as pilots. But as equipment was developed, it became apparent that it was going to be very difficult to achieve the performance and portability goals simultaneously. To do so successfully, the perturber designs would have to be optimized to meet well-defined performance objectives. Unfortunately, the performance requirements for the equipment and testing protocol were not well defined. The initial criteria were based on estimates of the torque required to perturb the head and assumptions about how to stimulate an appropriate gaze response, but not on experimental data that could be directly related to head-free pursuit of a random target. It was decided that developing a portable system would be deferred until the other performance objectives could be met. It was more important to develop laboratory equipment that was capable of producing a wide range of stimuli that could then be used to develop an experimental protocol and analytic procedures for the VOR experiment, than to have portable equipment that was suitable for use outside the laboratory.

### **3.2 Head perturbers**

The torque perturbation sequence created by the head perturber is one of the two controlled inputs for the experiment, the trajectory of the visual target being the other. The head perturber was thus a crucial element of the testing apparatus. To be acceptable for use in the VOR testing protocol, the perturber had to meet several design criteria. Its torque output had to be produced as a controlled, stochastic output that appeared unpredictable to the subject, but that was defined so that its relationship to the visual target was deterministic. It had to stimulate horizontal plane head rotations throughout the full range of natural head disturbances, from at least 0.5 to 3.0 Hz. The disturbance had to move the head sufficiently to create a measurable VOR response,

but at the same time it could not prevent natural head motions from being used to track the target. Most importantly, the head perturber, which acts directly on the head and neck of a subject, had to meet the safety and comfort requirements of a testing device.

The original performance specifications for the head perturber were that it should be capable of producing a 1 N.m torque perturbation across a 0 to 10.0 Hz bandwidth for 60 to 100 seconds. These criteria were speculative, but derived from sound assumptions. Based on the peak acceleration at which the human head can be driven by the neck, it was thought that an average human neck could produce 10 N.m. It was thought that 10% of this, or 1 N.m, was sufficient to elicit significant head motion, and thereby a measurable compensatory eye movement, but not so strong that a person would be prevented from controlling the head to follow a moving visual target. It was expected that experiments of 60 to 100 s would be sufficient to create the amount of data required to evaluate the full dynamics of the VOR, and that an adequate range of head motions would be created if the applied torque had a 0 to 10 Hz bandwidth.

Five different concepts were explored for the head perturber, and functional prototypes were built for each. The concepts could be divided into two general classes: those that were portable, were fully supported by the subject, and applied torque to the head via a reaction force, and those that were externally supported and applied the perturbation torque directly from a mounted motor. The portable perturbers were much more interesting to develop and to experiment with than the “direct drive” perturbers, but unfortunately, none of the portable prototypes met all the design objectives. It was too difficult to achieve the performance specifications and meet the safety and weight requirements for the perturber if it was to be worn by an elderly patient suffering from a balance disorder. The portability requirement was thus deferred, and a perturber was built that was supported by a frame, and applied a torque to the head directly with a motor. This “direct drive” perturber became the laboratory “mule” with which all the experiments, including the VOR studies, were run.

### 3.2.1 Portable perturbers


Four different devices were investigated as a means for producing the reaction torque for the head perturber: (1) water jets controlled by a Coanda switching valve; (2) an inertial ring accelerated by a servomotor; (3) a liquid metal accelerated by a Lorentz force; and (4) braked, counter-rotating flywheels. Each method had unique advantages, but the inertial ring and the braked flywheel devices seemed to be the most promising perturbation methods. The development that would have been required to implement the water jets and the Lorentz-force liquid-metal perturbers was considered beyond the scope of this research. Perturber prototypes built with the inertial ring and with braked flywheels were successful at producing stochastic torque perturbations, but had difficulty in simultaneously meeting the performance and weight requirements. The limiting factor with both technologies was striking the balance between producing a 0.5 to 1.0 N.m torque across a 10 Hz bandwidth and building a perturber that was light enough to be worn by a patient suffering from a vestibular disorder. An inertial ring perturber that was light was good at producing frequency perturbations above 2.0 Hz, but would have required additional mass if it were to be tuned to generate the torque at lower frequencies. The flywheel prototype was able to generate a full range of torques across a 0.5 to 10 Hz spectrum, but was unable to store the full amount of energy required for a 60 s vestibular test. To do so, flywheels light enough to be worn comfortably by patients would have to operate at nearly the energy density limit of steel flywheels (Genta, 1985).

#### *Coanda water jet*

The first perturbation device that was built used a Coanda-effect valve to direct the flow of a water jet (Figure 3.1). The Coanda-valve was mounted on top of a helmet that could be worn by a subject. Pressurized water was delivered from a faucet through flexible hoses to a Coanda-effect valve, which directed the water jet through one of two outlet nozzles. The emission of the water created a reaction force that could perturb the head. The water jet perturber was unique




among the reaction force perturbors in that it did not accelerate an annulus of material circularly about the head.



**Figure 3.1:** Coanda water jet device. The Coanda perturbation device was constructed from brass tubing and flexible rubber hose. The channel cover was removed in the picture at the right.

A Coanda-effect valve is a bi-stable switching device and allows for very quick switching of a fluid jet between two outlets (Foster and Parker, 1970). It was thus very well suited to creating a binary output that alternated between a torque that twisted the head to the left and a torque that twisted the head to the right. The Coanda effect refers to the phenomenon in which a stream of fluid that emerges from a nozzle tends to follow a nearby curved or flat surface, as long as the curvature of the surface is not too great. When the fluid jet passes close to the surface, a low pressure bubble between the jet and the wall is caused by viscous entrainment of fluid, and the jet is pulled towards the wall (Foster and Parker, 1970). Equilibrium is reached when the radial acceleration of the jet is balanced by the pressure difference across the bubble. The jet can be prevented from curving towards the wall simply by supplying the fluid that the low pressure bubble entrains from the jet, from another source such as an inlet port that is near to where the bubble forms. A bistable, Coanda-effect valve is created by having a fluid jet exhaust into a “Y-gate”. The Y-gate has two walls onto which the flow can attach and that lead to two outlet channels (Figure 3.2). The jet can be directed towards one or the other outlets by providing a

small control flow on the opposite side of the desired outlet (Figure 3.2). Small Coanda-valves can switch a fluid jet between outlets in as little as 100 ms (Foster and Parker, 1970).



**Figure 3.2:** Y-gate of the Coanda water jet valve. The water jet is directed to one outlet or the other by a small control flow.

A functional prototype was built to test the merit of the Coanda device as a head perturber. Coanda air jet actuators have already been used successfully to perturb the human wrist in studies of the dynamics of the forearm (Xu *et al.*, 1991). An air jet was considered to be too noisy to be used as a head perturber, and so a water jet was developed. The Coanda valve was constructed from 11 mm square brass tube, with a 0.5 mm wall thickness. The valve's inlet was shaped into a nozzle by forming the brass tube over a plastic nozzle form. The inlet nozzle was fitted into a Y-gate that had outlet channels that angled from the midline at 20 deg. Slots that acted as inlets for the control flow were cut into the walls of the channels about one nozzle width (3 mm) downstream from the inlet nozzle. The control flow was supplied through nozzles similar to the inlet nozzle. The inlet nozzle, Y gate, and control flow nozzles were joined by a liberal application of silicone caulk. Water was supplied to the inlet nozzle from flexible tubes attached to a faucet. The control flow was supplied from a tap on the faucet line, and was controlled to flow to one side or the other by a manual diverter valve.

The device was successful as a proof of concept, although much would have had to be done to develop the water jet as a routine laboratory device. The Coanda-valve easily switched

the flow between the two outputs, and was capable of producing long periods (low frequency) of continuous torque in either direction. The reaction force was equal to the time rate change of momentum of the water through the nozzle, so as long as the water was accelerated through the nozzle and exhausted, a reaction force was created. The flow switched quickly between the outlets independently of whether the faucet was on full or had low flow rates. The torque amplitude was not quantified, but the device could not be easily stopped when moving if it was held and the diverter valve switch quickly. The torque definitely felt strong enough to shake the head. The factor limiting how quickly the flow could be switched in the prototype was how fast the diverted valve could be moved, but a diverter valve is not how an actual Coanda is controlled. For faster switching, the control water could be pressurized against a controlled port at the inlet, and a quick acting valve used to open the port.

Obvious shortcomings would prevent this type of system from being used as a medical device. A water-driven system would have to be used in a special environment, such as in a shower stall, and therefore has little potential as a bedside or routine laboratory testing device. One technical disadvantage with the system is that the hydraulic lines that carry the water to the Coanda valve on top of the head became stiff when pressurized with water. The stiff lines could inhibit the wearer's head motion, and so could prevent testing the vestibular system with normal head and neck movements.

### *Inertial ring*

The most straightforward, and successfully implemented, portable perturber was an inertial ring perturber (Figure 3.3). A servomotor was used to accelerate a circular ring circumferentially about a helmet that was worn by a test subject. The change in the angular momentum of the ring imparted a reaction torque onto the subject's head, and this torque was controlled by directing the output of the motor.



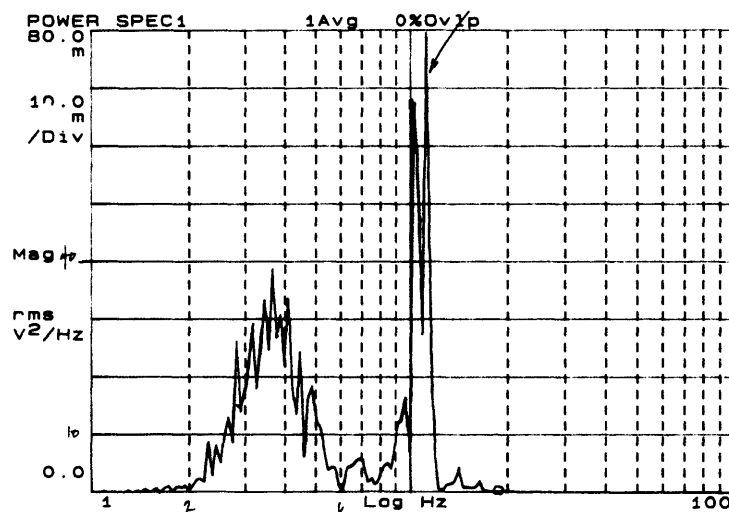
**Figure 3.3:** Inertial ring head-perturber

The initial prototype was built mainly to assist in understanding the issues that had to be addressed when designing head perturbers, such as how well would torque perturbations be tolerated, and how easily could subjects control their head and eye motions when the perturber was activated. A prototype was built using a motor and raw materials that were found in the laboratory, and turned out to be an extremely useful test tool. It was more capable than expected, and without any modification was able to be used for initial investigations into head-neck and vestibular system dynamics. Its shortcomings and successes highlighted issues that have to be considered when designing portable head perturbers.

The inertial ring perturber consisted of a polystyrene bicycle helmet and an aluminum ring that surrounded the base of the helmet. The ring was located so that it was near the same transverse plane as the center of mass of a subject's head, which is in a transverse plane just above the infraorbital notches. The ring had a mass of 0.25 kg and an average radius of 112.5 mm (rotational inertia of  $0.0032 \text{ kg m}^2$ ). The aluminum ring was attached through an aluminum brace to a torque transducer and a servomotor that were mounted on top of the helmet. The servomotor was a Hitec HS-805BB that produces a maximum torque of 1.5 N.m. The servomotor accepted a position command, and moved to the commanded position as quickly as it could. It did not offer torque or velocity control, so little could be done to shape the motor's output other than

influencing its duration. A pseudorandom binary sequence (PRBS) was used to encode changes to the servomotor's position. The duration of a PRB pulse was represented by a proportional change in the servomotor's position with longer pulses resulting in larger amplitude position changes.

The inertial ring perturber was light, comfortable to wear, easy to control, and demonstrated that an unpredictable torque perturbation with an amplitude of 1 N.m was more than sufficient to disturb the position of a subject's head. Subjects easily supported the perturber's 0.83 kg mass, but some subjects commented that it would have been difficult to control their head and support the perturber if it had been much heavier. Subjects were able to control head movements, but they were not able to stabilize their head to resist the torque disturbances. The main limitation of this perturber was that the bandwidth of its torque output was too high. A plot of the power spectral density for a representative torque perturbation is shown in Figure 3.4. There was power between 2 and 6 Hz, and also in a small frequency band just above 10 Hz. The power between 2 and 6 Hz was a direct output of the motor, and the power around 10 Hz was from an oscillation of the aluminum ring. This was not sufficient power at low frequencies to create the head motions required to test the VOR, which requires that head disturbances be towards 0.5 Hz.



**Figure 3.4:** Power spectral density of a representative torque perturbation created by the inertial ring head-perturber

The perturber was unable to produce lower frequency perturbations because of the small rotational mass of the inertial ring and the limited angular displacement of the motor. The perturbation torque was generated for only as long as the ring was accelerated. The greatest angular displacement that the motor could make was 120 deg, and so the longest duration that the torque was produced was the time it took for the motor to move the ring through 120 deg. If the ring had a greater rotational inertia it would have taken longer to traverse the angular displacement, or had the motor had torque control, it could have accelerated the ring with a lower magnitude torque.


The shortcomings of the inertial head perturber illustrated the tradeoffs that had to be made in the perturber's design between the amplitude of the torque, the weight and rotational mass of the inertial ring, and the maximum speed that the motor must attain. When either the rotational inertia of the ring is small or the amplitude of the desired torque is large, the inertial ring tends to be accelerated to a high rotational velocity, or, as in the case of the prototype, it reaches the limits of the motor's displacement. For example, if the ring of the prototype was accelerated with a torque of 1 N.m for 1 s (half the period for a 0.5 Hz perturbation cycle), the angular velocity of the ring would be changed by approximately 313 rad/s, which is 2984 rpm. This might be an acceptable velocity increase if the ring had started from rest, but if the ring was accelerated from an already high speed then it could reach rotational velocities at which the motor was no longer capable of producing the desired torque. Since motors are power-limited, at maximum power as the velocity increases the torque that the motor can produce decreases. There are also safety concerns with having an external ring rotating at such high velocities and with the energy that it has stored. The velocity that the ring attains can be reduced by increasing the ring's rotational inertia, but the ring's weight is limited by how much weight the subject can support. The heavier the perturber, the less likely it is that the subjects can use natural head motions to track the visual target, and the less appropriate the perturber becomes for testing elderly patients with vestibular dysfunction. The maximum speed that the ring attains can also be limited by

reducing the torque with which the ring is accelerated, but there is a minimum torque that is required to create the head disturbance. To design a good inertial ring perturber, the right balance between the rotational inertia of the ring, the perturbation torque, the size and speed of the motor, and the overall weight of the perturber must be found.

### *Liquid metal perturber*

The most interesting device that was considered for the portable perturber was one that accelerated a ring of liquid metal about a subject's head using Lorentz force actuators. The Lorentz force is the force that acts on a charged particle in an electric and magnetic field (Figure 3.5), and is the fundamental force that causes a motor to rotate. The Lorentz force law describes the force that acts on a charged particle moving through a magnetic and electric field:

$$F = q_0 \vec{E} + q_0 \vec{v} \times \vec{B} \text{ (Halliday and Resnick, 1986).}$$



**Figure 3.5:** Liquid metal perturber. The figure on the left is a cutaway of a fluid channel. A magnetic field ( $B$ ) is established across the width of the channel, and a current ( $I$ ) is driven perpendicular to the magnetic field. A Lorentz force is developed that accelerates the liquid metal down the channel.

In the inertial ring perturber, the Lorentz force generated the servomotor's torque, which was applied to the inertial ring through a mechanical connection. Rather than using a mechanical link to connect an inertial ring to a Lorentz force, actuators can be created from magnets and current sources to apply a Lorentz force directly onto a conductive, inertial ring. It was


envisioned that a liquid metal could be housed in a channel that surrounded the head, like a head band, and be accelerated with Lorentz force actuators (Figure 3.5). Since the liquid metal does not have to flow in a circular path, the channel could be made to fit closely to the head, and a portable perturber could be made that did not have the ungainliness of the inertial ring device.

A device to test this concept was created by placing a liquid metal into a circular channel milled into a cast nylon block as shown in Figure 3.6. The liquid metal was an alloy made from 66% gallium, 20.5% indium, and 13.5% tin (MCP Metalspecialties, Inc., Fairfield, CT). It had a eutectic temperature (melting point) of 10.7 deg Celsius, and as claimed by the manufacturer it had low resistivity and was non-magnetic. It looked and felt like mercury, but did not have mercury's toxicity. The reaction torque was created by applying a Lorentz force to the gallium alloy in the direction of the channel. As long as the applied force was greater than the flow's viscous retardation forces, the liquid metal would accelerate around the channel and create a reaction torque that could be used to perturb a head. The Lorentz force was created by driving large currents through the gallium alloy from the top of the channel to the bottom, perpendicular to a magnetic field. A strong magnetic field was established across the channel by neodymium-iron-boron permanent magnets placed on either side of the channel (Figure 3.6). By directing the magnetic field radially across the channel, the resultant Lorentz force was in line with the centerline of the channel and pushed the liquid metal along the channel.

For this system, a 3.175 mm thick, and 45 mm deep channel, with an average radius of 76 mm was milled into the cast nylon block. An octagonal shape was cut from the center of the block to a depth of 45 mm, such that only a 2 mm wall existed at the thinnest point. This was done so that permanent magnets could be placed on either side of the channel. The top of the channel was closed with a 13 mm thick Delrin cover plate into which a 5 mm deep channel was milled. When placed together, the cover and base block had an overall channel height of 50 mm. The Delrin cover was sealed to the cast nylon block with a silicon sealant, and screws were used



to compress the seal and ensure alignment of the channels. The enclosed channel held  $67.0 \times 10^{-6}$  m<sup>3</sup> of the gallium alloy.



**Figure 3.6:** Liquid metal, Lorentz force perturber. A liquid gallium alloy was housed in a channel milled into a Delrin block. The magnetic field was created by neodymium-iron-boron magnets, and was maximized across the channel by using iron “horseshoes” to create a high permeability return path for the magnetic field. Current was delivered through wires immersed in the gallium at the top and bottom of the channel.

Current was delivered to the gallium alloy through 10 gage, wound copper wire, insulated with PVC. The wire’s overall diameter was 4.2 mm, but the conductor’s diameter was only 3.1 mm which just fit into the 3.175 mm channel. Brushes, that made contact with the gallium, were fashioned by stripping the insulation from the end of the wire, and passing the wire through a hole axially machined into a 4 mm diameter nylon screw. The bare wire was sealed to the screw with epoxy. The screw was threaded into the top and bottom of the channel such that a small length of wire extended into the liquid gallium that sat in the channel. Current was supplied by 12 volt, deep cycle, marine batteries, and was switched on and off with automotive ignition switches. The current passed from the cable at the top of the channel, through the gallium, out the cable at the bottom of the channel and through a resistive carbon pile. The resistance of the carbon pile could be adjusted by changing the mechanical force acting on the pile. Changing the resistance of the carbon pile allowed for resistive control of the current flowing through the system.

The strength and proper direction of the magnetic field in the channel was crucial for the perturber to work effectively. Ideally, the magnetic field would be directed radially outward everywhere across the channel, and to complete the magnetic circuit, would follow a return path above and below the channel. A radial magnetic field, perpendicular to the current flow, would ensure that the Lorentz force created by the vertically flowing electrical current would be directed tangential to the centerline of the channel. Any magnetic field lines that flowed through the channel in a direction other than radially outward would change the direction of the Lorentz force away from the tangent to the channel centerline. Ideally, two radially polarized, continuous, cylindrical magnets, one on the inside and one on the outside of the channel, would have been used to direct the field across the channel. Unfortunately, cylindrical permanent magnets of the size and strength needed for this application would have had to have been custom manufactured, at a cost of approximately \$15,000. Instead, strong, 50 mm square, neodymium-iron-boron magnets were used in pairs along the inside and outside of the channel. A finite element, 3d, electromagnetic field software package (Maxwell 3d, Ansoft Corporation, Pittsburgh, PA) was used to analyze different arrangements of magnets and iron around the circular channel. The magnetic field was maximized across the channel by connecting pairs of magnets on either side of the channel with a soft iron horseshoe as seen in Figure 3.6. The iron horseshoe served as a return path for the field after it flowed across the channel. This arrangement resulted in a strong 0.65 Tesla field measured across the channel (air filled, not gallium). The penalty for this large, but necessary, magnetic field was the excessive weight of the magnets and the iron. When implemented in the prototype, each pair of magnets and soft iron horseshoe exceeded 1 kg, which is much too heavy to be used for a head perturber in which the total weight of the system should be less than 2 kg. Unfortunately, there do not seem to be permanent magnets, ceramic, iron based, or rare earth that at present can meet the weight and strength requirements.


The channel dimensions significantly impacted the theoretical torque that could be produced with the liquid metal perturber. It was desirable to make the channel as tall as possible

so that as large a Lorentz force as possible could be applied to the fluid ring. The height of the channel corresponds to the length of the conductor in the  $F = IL \times B$  “force on a conductor” equation, where  $F$  is the force on a conductor in a magnetic field,  $I$  the current,  $B$  the magnetic field, and  $L$  the conductor length. However, the reaction torque that was created by the liquid perturber was not equal to the torque that was applied to the fluid, but to the time rate change in the angular momentum of the fluid. Although larger Lorentz forces could be applied to the fluid ring by making the channel taller, when the channel radius is fixed by the outer dimension of the perturber and the volume of fluid fixed by the allowable weight of the device, the channel must be made thin to accommodate the tall channel. The thinner the channel, the more pronounced the effects of the viscous boundary layer at retarding the flow and restricting the change in the fluid’s momentum. From a simple 2-D fully developed flow model of the channel, it can be shown that the momentum flux of the fluid flow is proportional to the thickness of the channel to the 5<sup>th</sup> power. Therefore as the channel height is increased, more force can be applied to the fluid, but as it does, the fluid channel must get thinner, which increases the viscous effects retarding the flow of the fluid in the channel. A thinner channel, though, makes it easier to support a high magnetic field across the channel. So although the viscous forces are increased as the channel gets thin, a thinner channel will have a stronger magnetic field, and a taller height. Both of which increase the Lorentz force that accelerates the flow.

The thickness of the channel gave rise to another concern, the current carrying capacity of the gallium alloy. To attain the desired 1 N-m torque, under ideal conditions, over 400 A would have to be driven through the metal. With a channel thickness of 3.175 mm, and an average radius of 76 mm, a current of 400 A (if evenly distributed) would yield a current density (flux) of over  $2.6 \times 10^5 \text{ A/m}^2$ . The actual current capacity of the gallium alloy is unknown, but for copper is  $1 \times 10^7 \text{ A/m}^2$ . It was felt that keeping the current flux under  $1 \times 10^6 \text{ A/m}^2$  for the gallium was acceptable.

The perturber was tested by applying brief 1 to 3 s current pulses. Torque was measured by allowing the perturber to push against a force transducer, and the current that flowed through the gallium allow was inductively measured. During testing, the liquid metal perturber produced peak torques near what was estimated for its dimensions and field strength. The device produced nearly 1 N.m when the current passing through the gallium was just over 430 A. In Figure 3.7 it can be seen that the peak torque created by the perturber was nearly proportional to the measured current. Although the perturber created its expected peak torque, the torque could not be sustained. After an initial peak, the torque would decay to zero as shown in Figure 3.7.

Measurements indicated that the current flowing through the system held steady throughout the test, and so the reduction in force was not due to a decrease in current. A possible explanation is that the reaction force decayed because the fluid was actually attaining its steady state velocity and no longer accelerating. The steady state velocity is attained when the viscous forces balance the Lorentz force being applied to the fluid. The viscous forces rise with velocity, and a graph of the fluid speed would resemble the decay curve of the torque. The time that the force took to decay to zero was independent of the amount of current applied to the fluid.



**Figure 3.7:** Torque created by the liquid metal perturber. The peak torque, shown in Plot A, was nearly proportional to the measured current. In plot B it can be seen that the torque could not be sustained, but decayed to zero with time.

The liquid metal perturber was shown to hold some promise; it was able to produce peak forces that matched the requirements of a head perturber, and would be relatively simple to

modulate by controlling current (although the size of the currents that were necessary would require special switches). To maximize the force of the system, there would have to be additional extensive fluid flow analysis and experimentation to determine why the torque decayed quickly. Before any direct Lorentz force based system could be considered for a head-mounted perturber the weight of the permanent magnets would have to be significantly reduced.

### *Braked flywheel*

The final device that was considered for the portable perturber was a braked flywheel. Instead of creating a perturbation torque by accelerating a ring or an annulus of fluid, a reaction force can be generated by decelerating a flywheel from a high initial speed. One of the hardest problems to overcome with the Lorentz perturber, and to a lesser degree with the inertial ring perturber, was the weight of the actuator that accelerated the inertial mass. The braked flywheel device solves this problem by replacing the heavy actuator with a low weight braking device. The energy storage properties of flywheels generally allow a flywheel system to be designed for time-averaged performance requirements rather than for peak performance needs. This often enables a flywheel system to be smaller and lighter than non-energy storage systems which have to be designed to meet peak performance requirements (Genta, 1985). A flywheel could be spun to a velocity for which there was sufficient energy to perturb the subject throughout an experiment by an actuator that was not attached to the perturber. The torque perturbation could then be modulated by controlling the duration and the amplitude of the force used by the braking device to decelerate the flywheel.

It was envisioned that the braked flywheel system would resemble a computer's hard disk drive assembly combined with small, caliper brakes, such as those from a bicycle disk brake (Figure 3.8). Several flywheels would counter-rotate so that reaction torques could be created in both directions. Potential gyroscopic effects from angular movements away from the flywheel's primary rotational axis could be reduced by minimizing the difference in the rotational velocities


of the counter-rotating disks. Small pancake motors could be used to increase the speed on one set of disks if the disparity became too great. The acceleration of disks in one direction would supplement the reaction torque created by the deceleration of disks in the opposite direction. The trick would be to see how closely the speeds of the individual discs could be controlled so that the net rotational momentum of the system was minimized as the kinetic energy of the disks was being removed by the braking force.



**Figure 3.8:** Disks and caliper brake for a braked flywheel perturber

A large prototype was built from high-end bicycle components. The goals for the prototype were to get an empirical feel for modulating the braking of spinning discs, and to see if low braking forces could be created and modulated with frequencies between 0.5 to 10 Hz. It was not thought necessary to build a prototype that met the weight constraints of a head-mounted perturber until it was established that a braking system was a useful way of transferring energy from the flywheel. The intent was to spin two bicycle wheels in opposite directions and measure the force that could be transferred from the wheel to force sensors as the wheels were braked to a rotational speed of zero. Two bicycle wheels were mounted in independently rotating forks, upon which small brake calipers were attached (Hayes Brake L.L.C, Mequin, WI). Mounted on each wheel hub was the disc upon which the brake caliper acted. Pictures of the prototype are shown in Figure 3.9. Disc brakes were chosen in preference to traditional rim mounted, linear push/pull and cantilever brakes because of their superior force modulation. Rim mounted, polymer pad, cable

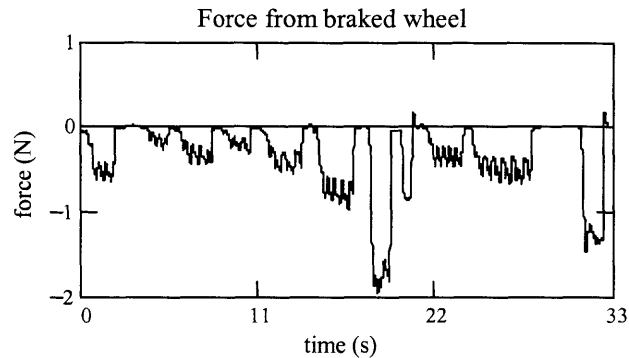
actuated, cantilever bicycle brakes tend to grab the wheel when even slight pressure is exerted upon the brake's hand lever. Hydraulically actuated disc brakes and their hand lever are specifically designed for force modulation, and the manufacturer claimed that the equipment would provide excellent light force application and modulation. The wheels were given an initial velocity by placing a hand drill against the rim and spinning the wheel to a high velocity. When the moving wheel was braked a reaction force pushed the fork on which the caliper was mounted against an Entran force transducer (Entran Sensors and Electronics, Fairfield, NJ) that measured the braking force.



**Figure 3.9:** Prototype built to test the feasibility of producing a perturbation torque by decelerating spinning disks with hydraulic caliper brakes.

Experiments demonstrated that the disk brakes could be modulated to hold a relatively constant force, and that force perturbations could be created with periods that had frequencies between 0.5 and 10 Hz. The braking force as one wheel was slowed in shown in Figure 3.10. To vary the frequency across this bandwidth the brake had to be applied for periods between 0.1 and 2 s. It was simple to modulate the brake's force with the brake handle, but in a head perturber system the hydraulic brake could be actuated by electronically controlling the plunger that creates the pressure in the brake line. As can be seen in Figure 3.10 a high frequency oscillation is apparent in the force output of the brake. The oscillations could be attributed to the slots in the brake disk which are made to allow water and dirt to escape from under the pads and also to

reduce weight. As the holes of the drilled sections passed through the brake caliper and the area between the brake pads changed, the braking force rose and dropped.



**Figure 3.10:** Force created by braking a spinning wheel. The force was modulated using the manual brake handle.

There was little reason to continue with the development of the flywheel perturber until the torque that was required to disturb the head was more precisely defined. The flywheel perturber had great potential as a lightweight perturber, and the actuation device had been shown to be effective, but its feasibility for use in a VOR experiment was directly related to the amount of force it had to produce and the duration of the experiments. The smaller the torque that is required to disturb the head, the lower the energy that has to be stored in the flywheel. Lower energy would mean that a flywheel could have a lower maximum rotational speed, or that the disk could have a lower rotational inertia. Lower rotational speeds can be attained with smaller motors, and lower rotational inertias require smaller (in diameter) disks. The total energy that must be stored in the flywheel is also important because if it increases beyond a certain level, special designs, fabrication methods, and housings are necessary to operate a flywheel safely. For example, assume that a 1 N.m torque is required to perturb a subject's head for 60 s. If two flywheels are used, then each disk must store an angular momentum of 30 N.m.s at the start of each experiment. If the flywheels are each assumed to be thin disks, with a radius of 100 mm and a mass of 0.5 kg, then each must be spun to an initial velocity of 114600 rpm (12,000 rad/s). This




translates to an initial energy density of 360 kJ/kg, which is beyond what can be stored by many steel alloy plates, and so a flywheel made from steel or Kevlar wires would be required (Genta, 1985). However, if the torque required to perturb the head was reduced to 0.5 N.m, then the energy density would decrease by 75%, and the feasibility of creating a flywheel perturber that met the experimental requirements dramatically improves.

### 3.2.2 Direct drive perturber

After designing and building a range of portable perturbers, it became apparent that a more precise set of specifications for the head perturber, based on experimental data, was required. It did not seem feasible to meet the original torque, bandwidth, weight, and safety criteria with a portable design. In addition, the 1 N.m torque perturbation that was created by the inertial ring perturber felt excessive to subjects, and so it seemed reasonable that much lower torques might be sufficient. A permanently mounted, “direct drive” perturber was built to meet these objectives. It was originally intended that this perturber would be used only to identify the performance specifications that were required of a portable perturber, but it was decided that this perturber should be designed so that it could be used in all the studies required to develop the VOR testing protocol.

The permanently mounted, direct-drive head perturber was constructed from an adjustable helmet and a servomotor that were attached to a counterbalanced support frame as shown in Figure 3.11. The motor that drove the helmet was a dc, brushless, resolver-based servomotor (Kollmorgen, MT304A1, Radford, VA) that was configured for torque control. This motor was selected for its ability to switch quickly between clockwise and counterclockwise torque outputs, and for its small size and low stall torque (7.0 N.m). The motor output was driven with an analog command signal produced by a National Instrument’s PCI-MIO-16XE10 data acquisition and output (NI-DAQ) board (National Instruments Corp., Austin, TX). To prevent the head perturber from applying too much torque to the subject’s head, an electronic current limiter

in the motor controller was set to limit the amount of current that could be applied to the motor's coil. A mechanical slip clutch (Hilliard, L2-1-312AC, Elmira, NY) was placed between the motor's output shaft and the helmet and acted as a secondary safety system in case the electronic limiter failed. The slip clutch was set to prevent more than 1.0 N.m from being applied to the helmet by the motor. The torque applied to the helmet was measured with a 6 axis, force/torque transducer (Assurance Technologies Inc., Garner, NC) capable of resolving torques as low as  $2.0 \times 10^{-3}$  N.m and measuring torques up to 4.0 N.m. The manufacturer's specification sheets for the motor, controller, and f/t transducer are included in Appendix B.

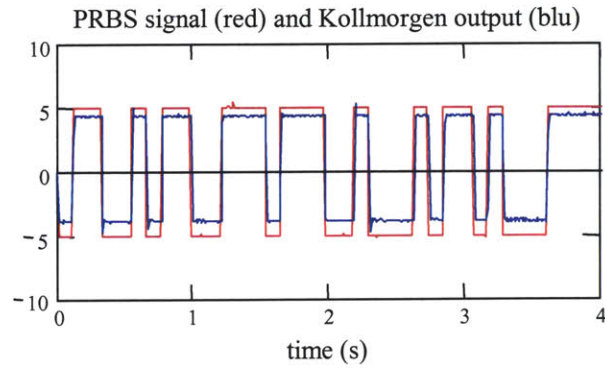


**Figure 3.11:** Direct drive head-perturber

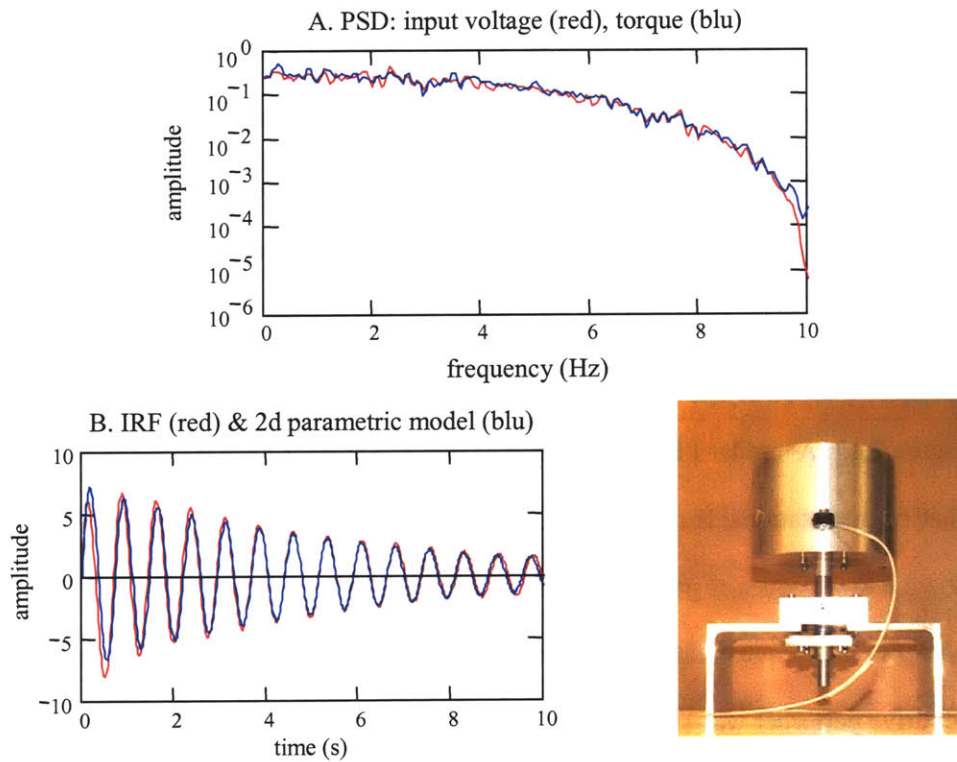
The frame was designed so that the helmet and motor could be lowered onto the head of a seated subject, and adjusted so that the rotational axis of the motor could be aligned with the natural, vertical neck axis of the subject. Weights on the counterbalanced frame were adjusted to balance the weight of the helmet and motor so that the subject did not have to support the weight of the apparatus. This was both for the comfort of the subject and to ensure that the weight of the apparatus did not tire the subject, nor affect how the subject used the head to track the visual targets. When the perturber was properly fitted to a subject, and no torque was applied to the helmet, the subject could freely rotate his or her head in the horizontal plane (yaw), but had little

freedom to rotate the head in the sagittal (pitch) or frontal (roll) planes, or to translate the head in any direction. This simplified the experiment to one of analyzing only the horizontal component of the VOR, rather than having to deconvolve horizontal, vertical, and torsional components of the VOR. With the exception of the shaft bearings and hardware such as screws and nuts, all the pieces of the frame assembly were machined in-house on CNC machining and milling centers from 6061 aluminum and stainless steel stock.

The Kollmorgen motor was capable of producing torque perturbations throughout the 0.5 to 10 Hz bandwidth that was desired to disturb the head. Tests of the motor demonstrated that it could follow square wave pulses with periods as short as 0.05 s. In Figure 3.12 the torque from the perturber is shown following a pseudo-random binary sequence (PRBS) that was executed at 10 changes/s. During the test, the motor applied torque to the torque transducer, which was held at zero velocity. The torque from the motor followed the command voltage very well. The square pulses were executed crisply, although there was some variance in the torque amplitude, the motor followed the command voltage quickly enough that there was not a noticeable lag when data sampled at 100 Hz were inspected. As was done with the inertial ring perturber, the direct drive perturber was tested by using it to evaluate the dynamics of the rotational mass system shown in Figure 3.13. A similar PRBS to the one shown in Figure 3.12 was used to drive the torque applied to the rotational mass. Its PSD, shown in Figure 3.13, shows that the perturber was able to create the low frequency perturbations that were necessary to test the VOR. Whereas there was little power in the torque created by the inertial ring perturber below 2 Hz, the torque output by the direct drive perturber was strong down towards 0 Hz. The PSD of the torque followed the PSD of the voltage command closely. Its decay at high frequencies was due to the command signal, and was not caused by an inability of the motor to follow the command signal.





**Figure 3.12:** A pseudo-random binary input signal (red) and the measured torque output (in volts) of the Kollmorgen motor (blue).



**Figure 3.13:** System identification of a rotational mass. The power spectral density of a PRB command voltage (red) and the torque output of the motor (blue) are shown in plot A. Plot B shows the impulse response function (red) and a second order parametric model (blue) that were calculated for the rotational mass system shown at the bottom right.

### 3.3 Visual target

The visual target was generated by reflecting a red helium laser (wavelength 600-700 nanometer, class IIIb) off a mirror and onto a flat screen that was about 1.5 meters from the test subject. The exact distance between the subject and the screen depended on the seating position adopted by the subject. These distances varied by only a few centimeters, and were accounted for mathematically. The mirror was positioned by a galvanometer (MG350DT, General Scanning Inc, Watertown, MA) that was controlled by a PID controller (AE1000, General Scanning Inc, Watertown, MA). The galvanometer and controller were capable of driving the mirror at frequencies over 10 Hz and at speeds above 200 deg/s with little delay. The frequency response of the galvanometer, mirror, and PID controller is shown in Figure 3.14. Pictures of the laser and galvanometer are also shown in the figure. Position commands were supplied to the galvanometer by the analog output of the National Instruments DAQ board. The laser was equipped with a lens that allowed the beam to be shaped so that the size of its image on the screen could be adjusted. The lens was adjusted and the beam passed through a circular aperture, so that the image subtended about 1 deg of arc on the screen. This was the size recommended by the American National Standard Procedures for Testing Basic Vestibular Function (Acoustical Society of America, 1999). Smaller images can move over the fovea without requiring the gaze to move to follow the target. The screen was 3.0 m long and 1.0 m high, and made from a flat sheet of expanded, high density poly-vinyl chloride (Sintra brand, Alcan Composites, St. Louis, MO). The front of the screen was covered in black vinyl tape to ensure that the screen was opaque. The target's image could travel +/- 35 deg relative to the subject when he or she was seated under the head perturber.



**Figure 3.14:** The visual target. The visual target was generated by reflecting a red helium laser off a mirror that was controlled by a galvanometer (top). The frequency response of the galvanometer and mirror are shown in the bottom two plots. The system's bandwidth exceeded 10 Hz.

### 3.4 Head tracker


It was originally envisioned that subjects would stand in an open testing area and track the moving visual target with whatever head and eye motions they desired. To do this, a head tracking method was needed that could measure the translations and rotations of the head in all three planes. Magnetic sensors that can be attached to the subject's head and body were

considered, such as those from Polhemus™ and Flock of Birds™. At that time, the angular resolution of these tracking systems was not good, and would be further degraded if subjects were tested within a metal frame that would support the various items of test equipment. The measurement of head position was simplified when the perturber design changed from being portable to being permanently mounted. The direct drive perturber restricted the subject's head motions to rotations about the vertical axis, so it was no longer necessary to measure linear translations. The only requirement was that the head tracker could measure one axis of rotation from 0 to greater than 10 Hz, and velocities up to 150 deg/s.

#### *ATA magneto-hydrodynamic angular rate sensor*

The first sensors that were used to track head rotations were magneto-hydrodynamic (MHD), angular rate sensors (ATA, ARS-04E, and Albuquerque, NM). The sensors were small and light, with a diameter of only 20 mm, a height of 16 mm, and a mass of 14 g. The MHD sensor was mounted on a thin aluminum brace that was taped to the bridge of a subject's nose (Figure 3.15). The bridge of the nose has little compliance and the sensor little mass, so it was likely that the sensor's motions reflected those of the head. The sensors worked on the same principle as the liquid metal perturber, or an electric generator. When the sensor was rotated, mercury, which is conductive, moved within a permanent magnetic field. The velocity of the mercury across the magnetic field caused a current flow which was proportional to the mercury's velocity. The response of the sensor was like that of a high pass system; the response was good out to almost 1000 Hz, but decayed at low frequencies. Similar to the semicircular canals of the vestibular system, the MHD sensor was insensitive to low frequency motions and its response decayed during sustained rotation. The sensor had a high pass corner near 0.3 Hz, so was insensitive to low frequency rotations of the head. A simple digital compensation filter was written to lower the effective corners of the sensor, but it could never function as a DC sensor, and the low frequency response of the head was always lost. This was not a significant problem

when only the head's velocity was being considered, but it was problematic when the head's position signal had to be calculated from the velocity. The manufacturer's specification sheet for the MHD sensor and the code for the digital compensation filter are included in Appendix B.





**Figure 3.15:** Magneto-hydrodynamic angular rate sensor

#### *Watson angular rate sensor*

Watson “tuning fork” angular rate sensors were then used as a replacement for the MHD sensors (Watson ARS-C141-1AR, Eau Claire, WI). The Watson angular rate sensors were larger than the MHD sensors, but they were effective over a bandwidth that extended from 50 Hz down to DC, so they could measure the low frequency motions of the head. The Watson sensors had an oscillating piezo-crystal that produced a torque proportional to its angular velocity. The Watson sensors were housed in a protective PVC cylinder that was 20 mm in diameter and 47 mm long, and so was too large to be mounted to the bridge of the nose. Instead, the sensors were mounted in a Delrin plate that was fixed to a mouthpiece. The noise that was measured from the sensor when it was at rest had a standard deviation that was less than the equivalent of 0.05 deg/s, and a maximum amplitude that was usually less than 0.1 deg/s. A representative noise measurement is shown in Figure 3.16. The manufacturer's specification sheet is included in Appendix B.





**Figure 3.16:** Watson angular rate sensor, a representative noise measurement (bottom left), and the power spectral density of the noise (bottom right). In this example the noise measured from the sensor had a standard deviation that was equivalent to 0.034 deg/s.

### 3.5 Eye tracker

#### 3.5.1 Electro-oculograph


Eye position was recorded using electro-oculography (EOG). The EOG quantifies the position of the eyes from voltages that develop on the skin that surrounds the eyes. The cornea and the retina of each eye act as a dipole, with the cornea positive relative to the retina. The corneal-retinal potential is small, about 10 to 30 mV (ADInstruments, 2003). When the eyes move, for example to look left, the cornea of the left eye nears the outer canthus of the left orbit and the retina of the right eye moves closer to the outer canthus of the right orbit. A potential difference of approximately 10 to 500  $\mu$ V develops between the two canthi (Duchowski, 2000).

There is a linear relationship between this voltage and the horizontal eye position up to approximately  $\pm 40$  deg (Leigh and Zee, 1999). For vertical eye motions, the voltage measured above and below the eyes is proportional to eye position up to about  $\pm 30$  deg (Clark, 1995). Silver-silver chloride (Ag-AgCl) surface electrodes are typically used to transduce the ionic potential that develops on the skin into an electric potential that can be amplified, typically through high quality DC instrumentation amplifiers, and output to a strip recorder or sampled by data acquisition devices. Noise is present in this voltage due to the effects of the EEG (electrical activity of the brain), the EMG (muscle activity), and the EOG equipment. This noise is often equivalent to approximately 1 deg of eye movement, thus the spatial resolution of the EOG is about 1 deg (Clark, 1995; Leigh and Zee, 1999).

Several techniques have been developed to record eye motions. The most common methods are EOG, scleral contact with search coils, infrared reflection photo-oculography, and video based oculography. Of these methods the EOG was best suited for the development of the VOR testing protocol as it is widely used in clinical testing, allows for the measurement of a large range of eye movements, and is relatively inexpensive (Leigh and Zee, 1999). The EOG is noninvasive, causes minimal discomfort and can record the position of each eye separately. Most importantly for our VOR protocol, the electrodes that are attached to the test subject to record the EOG are small which means they are able to measure eye position relative to the head even when the head is being shaken. The low mass of the electrodes allows them to move with the head, and they do not interfere with how subjects use their heads to track the moving target. The main disadvantages of the EOG are the noise that lowers its spatial resolution, and a low frequency drift that is often present in the signal. The noise can be especially problematic when the eye's velocity is calculated as the derivative of the eye position signal.

Scleral contacts are often regarded as the most precise means of measuring eye position, but the technique is invasive which makes it one of the more difficult methods to use with human subjects. Subjects wear large contact lenses (larger than the iris to reduce slippage) in which a

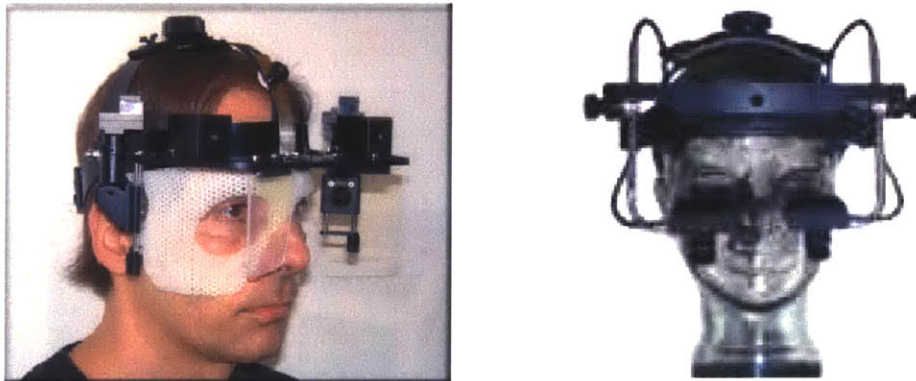
wire coil is embedded and there is also a wire lead that extends from the contact. The wire coil moves with the eye, and its position is tracked as it moves within an electromagnetic field created by a frame that surrounds the subject's head (Figure 3.17). At a cost of \$15,000, the coil and frame systems are expensive, but they are the most precise eye tracking tool available. They are capable of measuring the horizontal, vertical, and torsional position of the eyes with a sensitivity of 5 minutes of arc (Leigh and Zee, 1999). The precision of this method would be ideal for the VOR experiments, where the small eye motions caused by head disturbances are of interest, but the frame could limit the motion of the subject's head, and the use of contact lenses with human subjects outside a medical facility can be problematic. These contact lenses require that a topical anesthetic be applied to the eye, and there is an incidence of scleral abrasions in approximately 1 in 400 to 500 subjects ( Leigh and Zee, 1999; Hain, 2003). The scleral contact tracking method is a very good research tool, but it will probably never become a clinically accepted device for testing human eye motion (Hain, 2003).



**Figure 3.17:** Scleral contact (left) and frame that produces electromagnetic field (Skalar Medical BV, 2003)

Infrared reflection and video based technologies that track the movement of the pupil and iris are rapidly improving and are becoming the method of choice in many research laboratories

for tracking eye motions (Hain, 2003). The systems often have less noise and more precision than the EOG, can track horizontal, vertical, and sometimes torsional motions similar to scleral contact lenses, and are noninvasive. The systems track eye motions using high quality cameras or photoelectric cells that are typically mounted on the subject's head. Pictures of two such systems, the Chronos (Skalar Medical BV, Delft, The Netherlands), and the Eye Link II (SR Research Ltd., Ontario, Canada) are shown in Figure 3.18. Because of the head gear and the mass of the cameras, the systems are inappropriate for use when the head is undergoing fast, high frequency disturbances. The systems are often able to correct for small movements between the camera and the head, but they are not designed to correct for the types of disturbances that were to be used in the VOR testing protocol. The systems can be very expensive. For example, the Chronos eye tracker, which can resolve eye motions as small as 0.1 deg, costs \$40,000.



**Figure 3.18:** Head mounted eye trackers. The eye tracker at the left is the Chronos tracker from Skalar Medical BV (2003). At the right is the Eye Link II from SR Research (SR Research Ltd., 2003).

#### *Construction of the EOG amplifier*

Commercial EOG systems were available, but it was thought that it would be better to build an EOG amplifier so that its characteristics could be specifically selected. The basic layout of an EOG is simple. It requires electrodes that transduce the ionic current from the body to an electronic current, a high quality dc, differential amplifier that can amplify the signal from the


eyes with a gain a 500 to 1000, and a means to measure the output. Additional components, such as isolation amplifiers that protect the subjects from electrostatic charge, and filters and balancing circuits that help condition the signals are nice to have, but are not required to measure a signal that is proportional to eye position. Advice was sought on EOG measurement from the Jenks Vestibular Laboratory at the Massachusetts Eye and Ear Infirmary (Dr. Conrad Wall), and from a vestibular system expert, Dr. H. Galiana (McGill University). Both Dr. Wall and Dr. Galiana agreed that the EOG was the best eye tracking method for our experimental situation. Dr. Wall recommended that the development of the VOR testing protocol should use the EOG, but that different eye tracking methods should be considered as new technology was developed. He cautioned that the noise in the EOG signals would limit the upper frequency at which the eye's velocity could be calculated from the position measurements. Both he and Dr. Galiana recommended measuring the "bitemporal" differential voltage between the outer canthi of the left and right eye, rather than the monocular differential voltage of each eye. This meant that the system would measure conjugate eye motions, rather than the motion of each eye individually. As long as the target was projected at visual infinity, the eyes should move together. The bitemporal measurement would have a higher signal to noise ratio than a monocular measurement, and this would avoid the difficulties that occur when an electrode has to be placed between the subject's eye and nose.

The core of the EOG amplifier that was built was a signal conditioning board designed by Dr. Sylvain Martel (SCM-4X, Unima Project). The board is visible in Figure 3.19 at the back of the EOG amplifier box. The signal conditioning board was designed to amplifying low voltage signals from the heart, and so its components were of very high quality and well suited to amplify the noisy, low voltage signals of the EOG. The EOG amplifier was set up to accept 4 channels of differential input, so potentially the horizontal and vertical movement of each eye could be measured. The first stage of the EOG amplifier was an isolation relay that was designed to protect the subject from any electrostatic charge that may be in the signal conditioning board. After the

isolation relays, the analog signals from the electrodes were input to dc, instrumentation amplifiers (Burr Brown, 1NA11DAG 95394006). These instrumentation amplifiers had a high input impedance so they did not load down the source and a high common mode rejection ratio of 106 dB so that only the difference between the input signals was amplified. The gain of the amplifier could be selected at 200, 600, or 800 times using toggle switches on the EOG amplifier box. This allowed the amplifier's gain to be chosen so that the amplified signal would fill the dynamic range of the data acquisition board that was used to record the EOG. After being amplified, the signals were filtered. Two of the four channels were filtered by a 0.07 Hz high pass filter to remove low frequency drift, and two channels were not high pass filtered so that the dc component of the EOG was preserved. All four channels were filtered with a 60 Hz low pass filter. The last circuit on the signal conditioning board was a unity gain follower which provided low impedance so that the board could be connected to another circuit without it being loaded. The amplified signals were then sent through isolation amplifiers (Analog Devices, AD210JN, Norwood, MA) before being sent to a balancing circuit and then to a data acquisition board in a computer. The isolation amplifiers isolated the subject from the 120 V, power supply of the computer. Although the probability was small that current from the computer would be delivered through the data acquisition board, via the signal conditioning board, its isolation amplifiers, and then to the subject through the electrodes, these isolation amplifiers were added to ensure that the system was safe. A summation amplifier (balancing circuit) was built to generate a balancing voltage that was added to the EOG output so that the signal measured by the data acquisition device began from 0 V, rather than being biased from zero. A common problem with DC electro-oculography is signal drift which is often caused by the half cell potential of the electrodes (Webster, 1995). When EOG gain is set to a moderate 600 times, it only takes a 0.01 V difference between the reference electrode of the forehead and the electrodes attached to the outer canthus of the eye to create a 6 volt offset. The data acquisition board could only accept voltages with magnitudes below 10 V, so a 6 V offset would limit the dynamic range over which the EOG

output could be measured. The balancing circuit was built using precision operational amplifiers (Linear Technology, LT1001, Milpitas, CA) and was designed to add or subtract up to 10 V from the EOG output. The voltage was controlled by adjusting the resistance of a potentiometer. The EOG amplifier was powered by two 12 V lead-acid batteries, and the isolation amplifiers and the balancing circuit were powered by a +/- 15 V Kepko power supply (MPS620M).

Disposable silver-silver chloride electrodes were used to measure the EOG (Hydrotrace, Physiometrix, Billerica, MA). The electrodes had a diameter of 15 mm (10 mm recording surface diameter) and had a cushioned adhesive backing, which made their application simple. They are shown in Figure 3.19.



**Figure 3.19:** Electro-oculgraph amplifier (left) and Hydrotrace Ag-AgCl electrodes (right).

### *EOG noise*

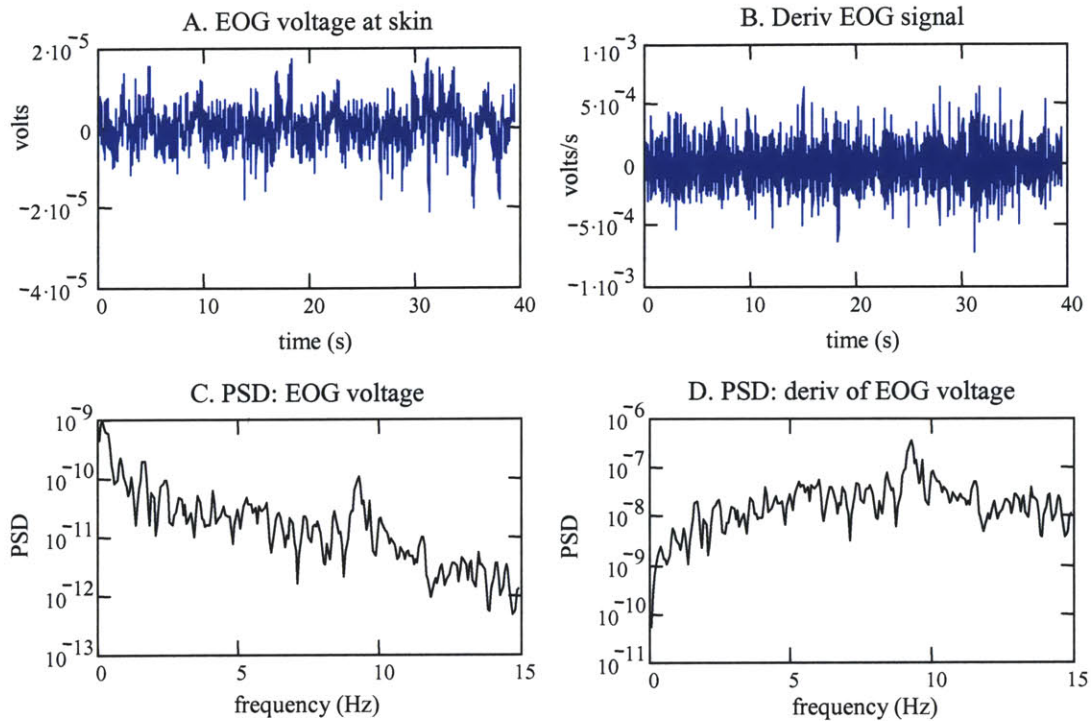
The noise in the EOG signals was found to be slightly smaller than what is usually expected from an EOG system. Typically, the noise of an EOG makes it possible to resolve eye position to within 1 to 2 deg. The noise, which is considered to be any voltage that is not from the corneal-retinal potential, comes from the EOG amplifier, leads, and electrodes, and from the EEG and EMG voltages from the subject. The noise level thus depends not just on the equipment, but on how cleanly the EOG signal can be measured from the subject. The noise in the EOG measurements was quantified from the voltage that was measured when a subject fixated on a

stationary target. During fixation, subjects can suppress many of the smaller eye movements that are associated with pursuit tracking, so fixation represented the condition with the smallest number of eye movements that could be measured by the EOG (Leigh and Zee, 1999). Typically, the noise that was measured in the system was equivalent to a voltage with a standard deviation of less than  $10\mu\text{V}$  at the electrode. In Figure 3.20 the noise measured while one subject fixated a target is shown. The standard deviation of the voltage measurement is  $5.1\ \mu\text{V/s}$ . The PSD of the voltage is largest at 0 Hz, and decays as frequency increases. The noise signal could be converted into a position and velocity measure using a calibration of this subject's EOG measurement to the subject's eye position. The standard deviation of the noise in terms of eye position was to 0.17 degrees. The maximum position amplitude that the noise equated to was about 0.5 deg. This was better than the 1 to 2 deg expected from EOG. When the signal was differentiated, the noise created an eye velocity signal with a standard deviation equivalent to about 5 deg/s. The amplitude of the PSD for this "eye velocity" noise increased with the frequency of the PSD of the "eye velocity" signal until about 10 Hz.

### 3.5.2 EOG gain calibration

The gain relation between the voltage measured by the EOG and the position of the eyes had to be determined. This gain relation was not invariant, and could change over the course of the experiment as the retinal charge varied or as the impedance of the skin under the EOG electrodes changed. By calibrating the EOG immediately before measuring eye positions, the most accurate interpretation of eye position data measured during the tracking trials could be made. Guidance for calibrating the electro-oculograph was found in The American National Standard Procedures for Testing Basic Vestibular Function (Acoustical Society of America, 1999).





**Figure 3.20:** Noise measurement for the EOG. Plot A shows the “noise” measured in the EOG while a subject fixated on a stationary target. The standard deviation of the noise measurement was  $5.1 \times 10^{-6}$  V which was equivalent to 0.18 deg of eye movement. The derivative of the noise is shown in plot B. Its standard deviation was  $1.6 \times 10^{-4}$  V/s which was equivalent to an eye velocity of 5.7 deg/s. The power spectral densities of the signals in plots A and B are shown in plots C and D.

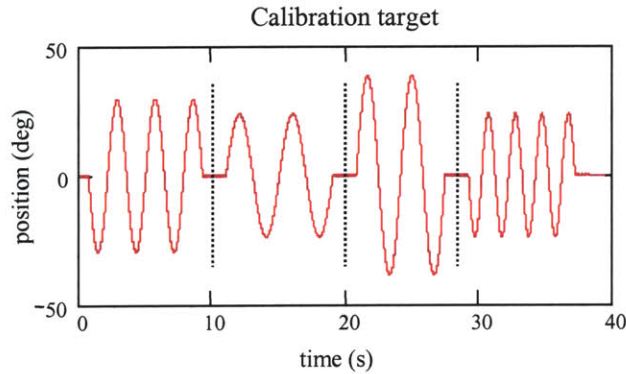
The EOG calibration gain was calculated by directing the subject to known angular positions and measuring the voltage at the EOG electrodes. In accordance with the ANSI standard, it is routine to calibrate an EOG by having a subject scan their eyes between stationary targets placed at -20, 0, and 20 degrees. The EOG voltage-eye position relationship is assumed to be linear, and the gain relation found at these displacements is assumed to hold through approximately +/- 40 degrees of eye motion. This simple procedure was used to calibrate the EOG during the VOR experiments at first, but subjects found it boring and occasionally did not follow the protocol well. The use of only a few static targets also made it difficult to verify that the EOG voltage and eye position had a linear relationship throughout the range of eye motions

that were anticipated to occur during the tracking trials. To make the calibration tests slightly more interesting, and to evaluate the voltage-position relation across a larger range of angular positions, a moving calibration target was used instead of the static targets. The moving target stimulated smooth eye pursuit motions, whereas the static calibration targets stimulated saccades and fixation. Thus the moving target created calibration conditions that more closely resembled the gaze response that would occur during the tracking trials.

The calibration target was made from four, easily tracked, sinusoids (Figure 3.21). The sinusoids had frequencies that were below or equal to 0.5 Hz, and had peak velocities less than 80 deg/s (Table 3.1). Each of the four sinusoids was executed an integer number of cycles, which made the mean position and the mean velocity of the target signal zero. The trajectory started and stopped at zero degrees, which was at the center of the projection screen, and the sinusoidal components were separated from each other by a 1.7 s delay in which the target remained stationary at the zero position. The delays provided an opportunity for subjects to bring their eyes and head to the zero position at intervals during the calibration trials. This was intended to help identify drift and offset in the measured head velocity and EOG voltage signals.

The underlying assumption of the EOG calibration was that the subject would learn the sinusoids quickly and would track the target with little error or lag. The same calibration sequence was used each time the EOG was calibrated which enhanced the subject's ability to predict and track the target with little error. The calibration gain ( $K_{\text{EOG}}$ ) for the EOG was found by fitting the subject's gaze position ( $gp$ ) to the target's position ( $tp$ ) sequence with a least-squares fit. The subjects were instructed to track the target with the head fixed, so ideally their gaze response consisted of only eye motions, and the gaze position could be represented by the scaled EOG voltage ( $gp = K_{\text{EOG}} \times eog$ ). Although subjects generally tried to keep their head fixed, in reality their heads were not stationary, and the gaze position had to be formed by summing the scaled EOG voltage with a measurement of the head position ( $gp = K_{\text{EOG}} \times eog + hp$ ). The head position had to be calculated from the measured head velocity. The EOG calibration gain ( $K_{\text{EOG}}$ )

was determined by fitting the EOG voltage (EOG) to the sequence formed by subtracting the head position (hp) from the target position (tp):  $K_{EOG} \times eog = (tp - hp)$ .



**Figure 3.21:** Trajectory of the EOG calibration target

sine	freq (Hz)	amp (deg)	max vel (deg/s)	cycles
1	0.35	30.0	66.0	3
2	0.25	24.0	37.7	2
3	0.30	39.0	73.5	2
4	0.50	24.0	75.4	4

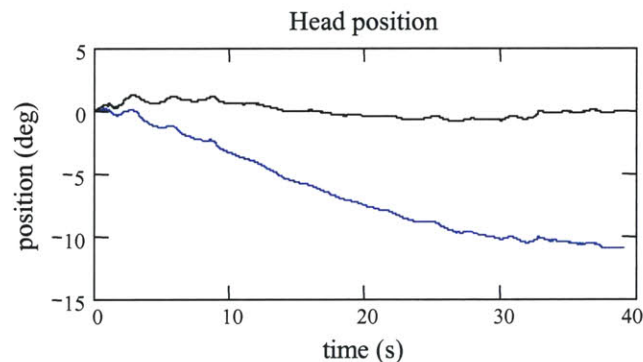
**Table 3.1:** Characteristics of the calibration target’s sinusoids

### *Drift and offset*

The measured head velocity and the EOG voltage signals had offsets and drifts that had to be removed before the subject’s gaze signal could be formed and then fitted to the target. The selection of an appropriate drift model was crucial to creating head velocity and eye position sequences that correctly represented the subject’s gaze during the calibration trial.

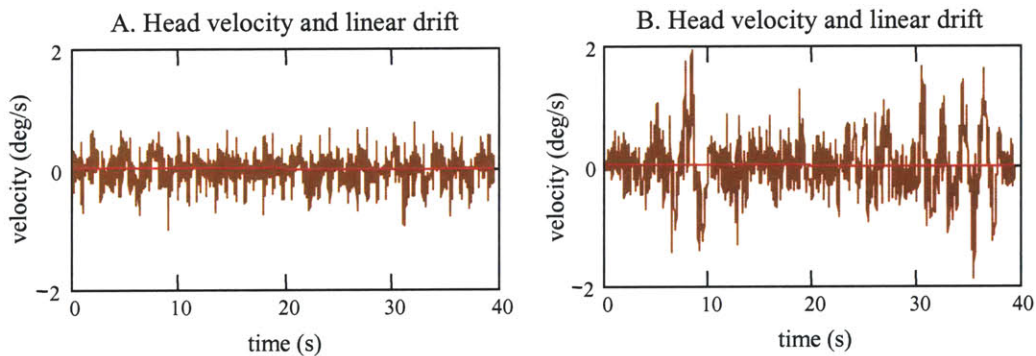
The drift in the head velocity signal was usually small and appeared to be linear. Although the drift was small, the head position signal, which was created by integrating the discrete time head velocity sequence with a trapezoidal integral approximation, was very

sensitive to how the drift and offset were modeled and removed from the measured head velocity sequence. Very different position sequences could be produced when different, but equally reasonable, models of the drift were removed from the velocity sequence. One model might result in a head position that had moved along with the target, while another method might indicate that the head had ended the experiment pointing 10 deg away from the target (Figure 3.22). Two head velocity sequences that are representative of the data collected from subjects are shown in Figure 3.23. The head velocity sequence in Figure 3.23A had one of the smaller drift rates, with a linear drift calculated to be  $0.4 \times 10^{-3}$  deg/s per second, while the plot in Figure in Figure 3.23B shows one of the larger drift rates of  $4.1 \times 10^{-3}$  deg/s/s. These rates are both within the range of the drift claimed by the angular rate sensor's manufacturer which is  $1.0 \times 10^{-3}$  deg/s per second. Ideally, the subjects would have kept the head stationary and at the zero position throughout the calibration trial, but in reality the subjects moved the head along with the target. It can be seen in Figure 3.23B, especially after the 30 s mark that the head moved at the same frequency as the target. If the head had remained still, the measured head velocity data would have been only noise, and the head velocity and head position sequences could have been uniformly set to zero. No processing of the head velocity data would have been necessary, and the gaze could have been determined from the EOG signal alone. But since the head moved, head position was a component of the gaze response.



**Figure 3.22:** Two possible head position sequences after different drift models were removed from the head velocity.

The drift model for the head velocity had to be consistent with assumptions about how the head was used when the subject tracked the calibration target. The calibration trials began and ended with the target at rest, and since the subject had been instructed to keep the head still, it was likely that the head was at zero velocity at the beginning and end of the calibration sequences. The response of the head sensor at these times probably represented 0 deg/s, or at least very close to zero, plus some noise. Any difference in the mean values of the signal during these periods was likely the result of the sensor drifting, and not the result of the head velocity gradually changing over the course of the experiment. Additionally, the mean of head velocity over the entire length of the trial should have been near zero, since the head began and ended the calibration trial in nearly the same position.

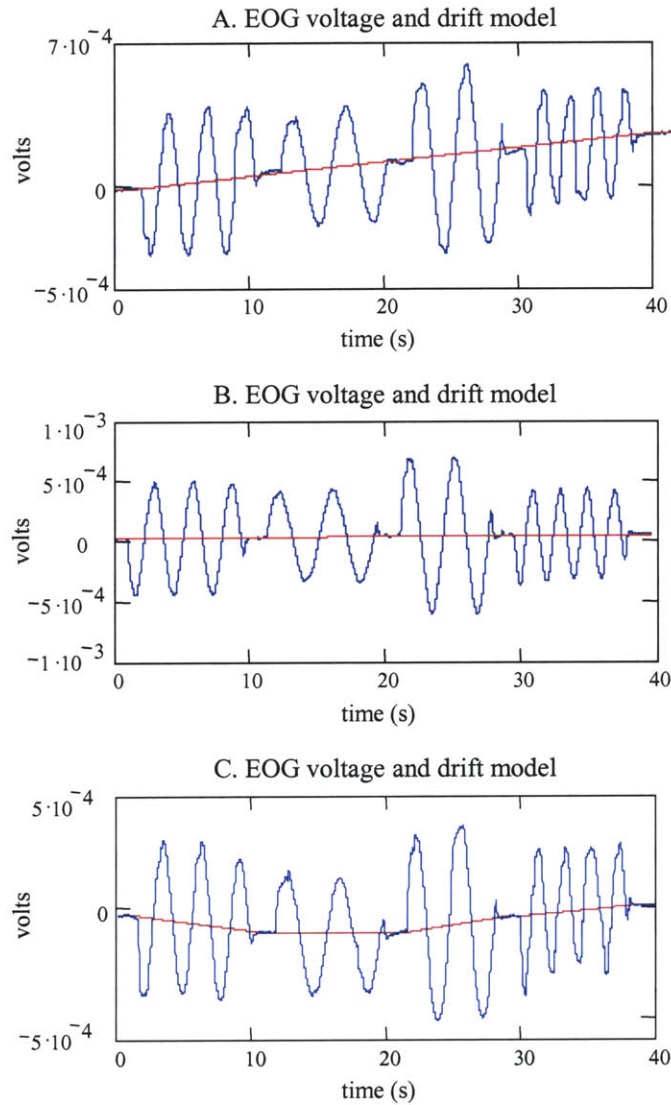


**Figure 3.23:** Head velocity (brown) and linear drift (red). The head velocity sequence in plot A had an average drift of  $-0.36 \times 10^{-3}$  deg/s/s, and the sequence in plot B had a drift of  $-4.1 \times 10^{-3}$  deg/s/s.

The measured EOG signal was also affected by offset and drift. A biasing circuit on the EOG amplifier was used to minimize the offset, so the bias of the EOG was much smaller than for the head velocity signal. Distinguishing signal drift from an actual eye motion was complicated by the fact that the EOG signal was a measure of eye position relative to the head, not a measure of the eye position relative to earth. To determine if the EOG had drifted, the signal had to be considered with regard to the angular position of the head. Although subjects rarely

kept their heads at zero throughout the calibration trial, they did tend to bring the head to rest close to zero during the periods in which the target was stationary at the beginning and end of the calibration trial, and during the rest period between each sinusoid. During these periods the EOG signal more closely represented the position of the eye in space, which assisted in visualizing if the EOG signal had drifted. The measured EOG signals tended to have one of three forms of drift. In some of the eye sequences the drift was significant, but appeared to be linear. Plot A of Figure 3.24 shows a representative EOG signal with a large drift. The EOG voltage began at approximately 0, followed a similar pattern as the target trajectory, but drifted along a line with a slope of  $5.6 \times 10^{-6}$  V/s. By the end of the trial, the drift was of the same amplitude as the sinusoid of the eye voltage, which represented a displacement of approximately 24 deg. The subject's head had remained within 1 deg of the zero position during this calibration trial, so assuming that the subject had tracked the target well, the linear movement of the EOG signal was due to signal drift. It did not represent eye movements that were made by the subject to compensate for the head having been moved away from the zero position. The majority of the calibration trials had EOG signals with a much smaller drift as shown in Figure 3.24B. The shift in the signal from beginning to end still appeared linear, but was much less pronounced than for the signal in plot A. The small shift in the EOG's voltage between the beginning and end of the trial was equivalent to only a few degrees displacement. Although treated as signal drift, this difference in voltage could have potentially represented a small shift in the eyes that compensated for a small head rotation, or even gaze motions that were not directly focused on the target, but were within acceptable gaze error for the subject to see the target clearly. Some of the EOG signals, though, had drifts that were large and did not appear to drift in only one direction between the beginning and end of the trial. Plot C of Figure 3.24 shows a representative signal. The EOG signal appeared to drift downward during the first two sinusoids, and then drifted upward during the last two sinusoids. Especially in these cases, it was necessary to consider carefully the EOG signal with regard to the

position of the head in order to identify what was eye movement and what may have been signal drift.



**Figure 3.24:** Three forms of drift in EOG: A. large linear drift; B. small linear drift; C. nonlinear drift.

### *Drift models*

The drifts and offsets in the measured sequences were estimated by fitting “drift models” to the data. The drift and offset were modeled as a constant, linearly changing function from the beginning to the end of the calibration trial, and also linearly changing between the rest periods of

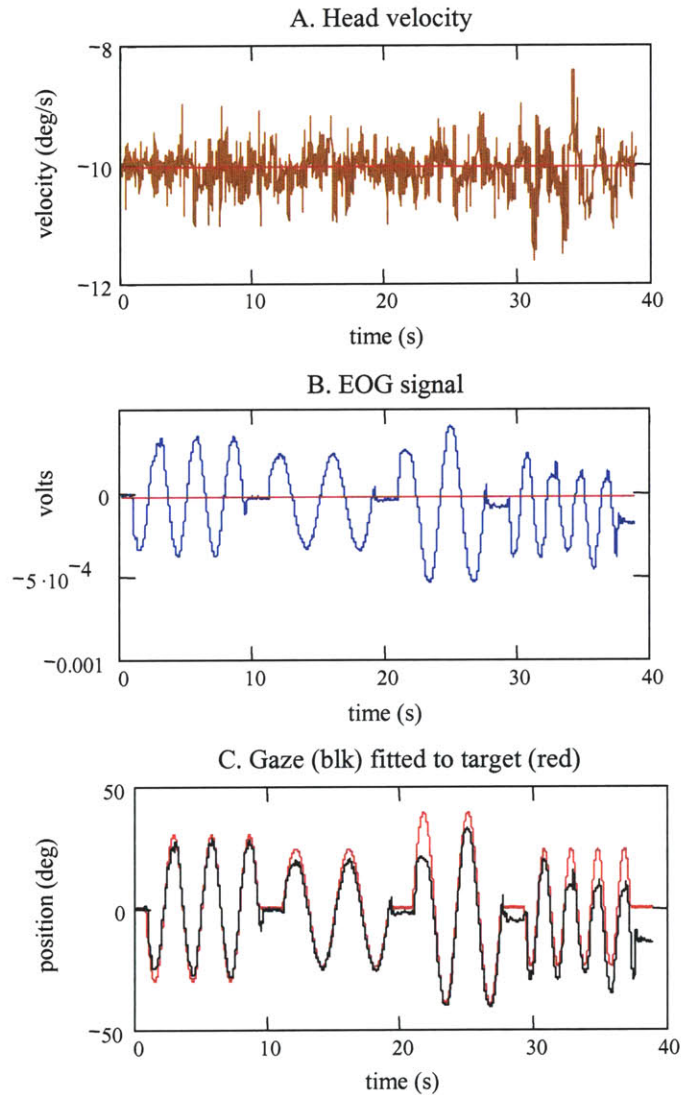
the calibration trial. The estimated drift and offset were subtracted from the measured head velocity and EOG signal to produce what was thought to be a more accurate representation of what the head velocity and the eye position were during the experiment. These sequences were then used to estimate the head position, and to calculate the EOG calibration gain by fitting the gaze position to the target position. A drift model made sense for a data set when certain assumptions about a subject's head and eye performance could be made, and once subtracted from the measured data imposed certain constraints on the resulting data.

The simplest way to model the drift and the offset was to assume, contrary to the evidence, that the head velocity and EOG signals did not drift, but were corrupted by a constant bias voltage. When this model was applied it was assumed that at the beginning of the calibration trial the head was at rest, the eyes were at 0 degrees relative to the head, and the gaze was focused on the target. During the period before the target began to move, the mean value of the head velocity signal represented 0 deg/s, and the mean value of the EOG signal represented an eye displacement of 0 degrees. These mean values were subtracted from their respective sequences. Any movement of the signals away from zero from then on was considered to be from actual head and eye movements. This method produced a good fit of the gaze to the target only when the signals had little drift. It usually produced a gaze that wandered away from the target when there was drift in the signals (Figure 3.25).

A method that was more consistent with the apparent drift in the data was to assume that the drift was linear, and to estimate the drift and offset by fitting a line by least-squares to the head and EOG sequences. By subtracting the fitted lines from the measured data the resulting head velocity and EOG signals have a mean of zero. The zero mean head velocity constrained the head position to end the calibration trial at the same position as it had begun, which may not have been precisely correct. The imposition of a zero mean on the EOG signal did not force the EOG to end and begin the trial at the same value, and so overall the gaze response was not constrained to match the target's zero degree position at the beginning and at the end of the trial. The target



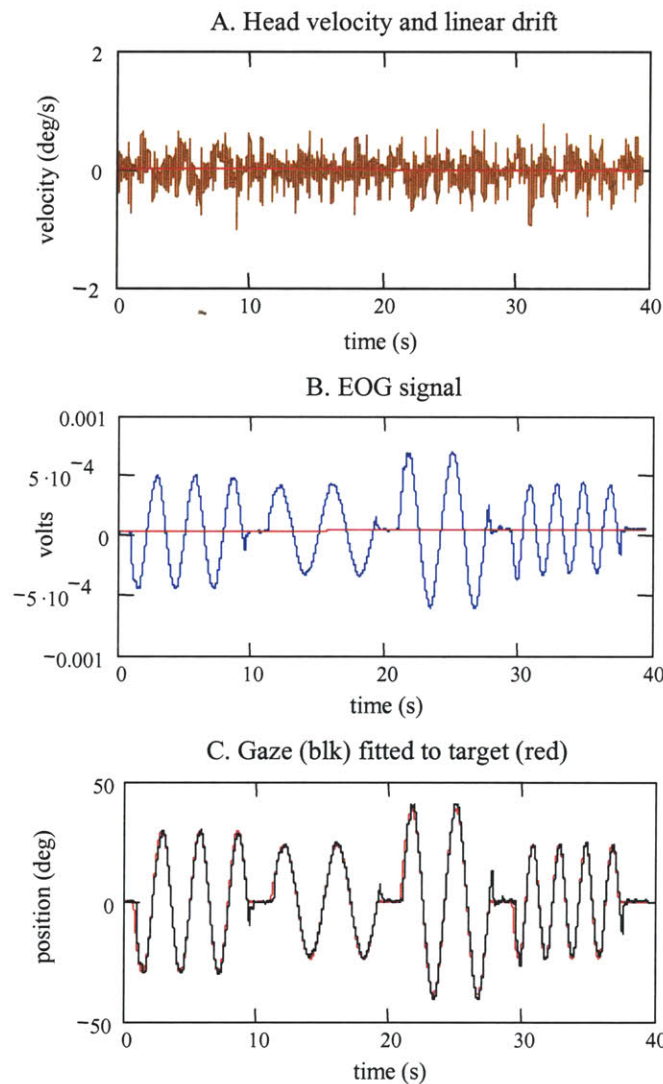
was set to a zero mean so it was not unreasonable to assume that the EOG had a zero mean. An example of sequences that had their drift modeled in this manner is shown in Figure 3.26. Across all subjects, this method tended to produce the best fits of the gaze to the target.



**Figure 3.25:** Gaze fitted to target. Sequences were assumed to have an offset, but not to have drifted. The gaze fit the target with a mean square error of  $41.9 \text{ deg}^2$

A different linear drift model was formed by assuming that, on average, the data collected at the beginning and end of the trial represented the head at 0 deg/s and the eyes at 0 degrees relative to the head. These were periods for which the target was stationary. The difference in the mean value of these regions was due to signal drift. The drift was assumed to be linear and was

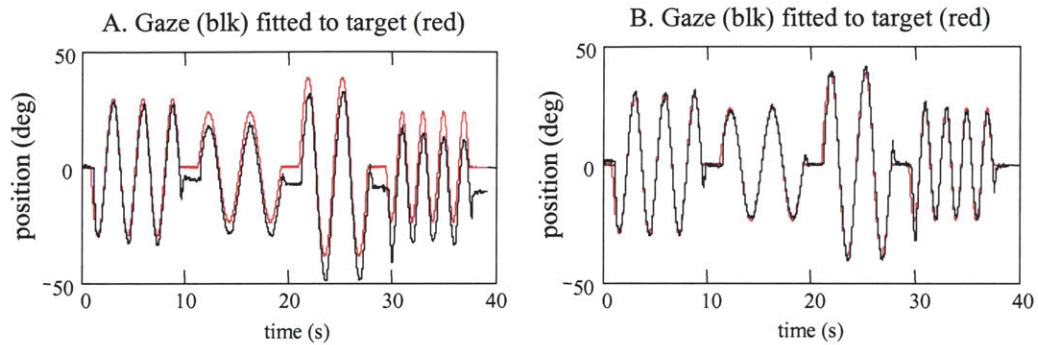
represented as a line fitted between these means. The resulting drift model was not very different from the line found with the least squares fit, but it differed in that it did not impose a zero mean on the head velocity. The final head position was not constrained to be equal to the initial head position, though the initial and final eye position was constrained to be 0 degrees relative to the head. As with the least-squares drift model, this method tended to produce good fits of the gaze to the target.



**Figure 3.26:** Gaze fitted to target. Sequences were assumed to have a linear drift. The gaze fit the target with a mean square error of  $5.0 \text{ deg}^2$

The final drift model estimated the drift with lines fitted throughout each of the five periods during which the calibration target was stationary. The five stationary periods were modeled with constants, equal to the mean of the data, and the drift between each of the stationary periods was modeled by a line joining the end of one constant region to the beginning of the next. The mean value during these stationary periods was calculated from the data after transients had settled, and the head and eyes appeared to be at rest. This model required the most restrictive assumptions about the subject's behavior to be valid, namely that each time the target came to rest the head was held still, and that the head and the eyes were both brought to the zero position. This method was not more effective at modeling the drift and fitting the gaze to the target, for sequences that had linear looking drifts. It was however, very effective when the drift did not appear to be linear but shifted directions during the trial. When that occurred the linear drift models were inappropriate, and this method worked best.

Although the different drift models would sometimes lead to very different fits of the gaze to the target position, the values calculated for the EOG calibration gain were surprisingly insensitive to the drift model. Even in cases where the head and eye signals did not fit the target well, the gain value that minimized the square error between the gaze and the target was often close in value to a gain found from a processing method that created a very good fit (Figure 3.27). It was therefore not crucial to choose the most exact drift model, since only the EOG calibration value was needed, and each method led to the same value being calculated. Understanding how the signal drifted was very helpful though, when identifying an appropriate drift model for the tracking trials. The gain value that was accepted as the "EOG calibration gain" and that was applied to the subsequent tracking trial was found by averaging the EOG gain values from all the drift models that gave reasonable fits and gain results. Methods that were obviously incorrect were excluded.



**Figure 3.27:** Gaze fitted to target. The EOG gain was not very sensitive to the method used to model the drift and fit the gaze to the target. In plot A, the EOG gain was found to be  $6.851 \times 10^4$ , while in plot B the EOG gain was found to be  $6.856 \times 10^4$

### 3.6 Data acquisition and experiment control

Data acquisition was done using a National Instruments data acquisition board (National Instruments PCI-MIO-16XE -10, Austin, TX). The board was capable of  $3 \times 10^6$  samples per second, had 16 bits of resolution and a  $\pm 10$  V dynamic range. The board had a maximum settling time of 10 microseconds between samples. The gain of the board was programmable so that the input signal could be amplified to fill the full dynamic range of the board. Signals were filtered by 6-pole, Bessel, anti-aliasing filters (Krohn Hite Corporation, Avon, MA) with cut off frequencies set to 40 Hz, before they were sampled by the data acquisition board. Bessel type filters were selected because their phase response is more linear than those of other low pass filters. When the phase response is linear, the time shift is the same at all frequencies. The data acquisition was controlled from an application written with Visual Studio 6.0 (Microsoft Corporation, Redmond, WA). The research began with a 180 MHz, Pentium Pro personal computer, and by the time the research project neared completion was run from a 1.6 GHz, Pentium 4.

A second computer was used to output the sequences that controlled the position of the target and the torque output of the head perturber. The output was run from a Visual Studio 6 application that relied heavily upon a custom timer application written by Dr. Serge Lafontaine which enabled the timing of the output signals to be executed precisely.

Data analysis was done offline. Analysis was conducted using applications written primarily in Mathcad 2000 and 2001 (Mathsoft Corp., Cambridge, MA) and by using and developing algorithms in the Plexus analysis library written by Prof. Ian Hunter.



## **Chapter 4**

### **Design of visual and vestibular inputs**

- 4.0 Chapter summary
- 4.1 Input signal requirements
- 4.2 Torque perturbation
- 4.3 Visual target
- 4.4 Tracking unpredictable targets with torque perturbations applied to the head

#### **4.0 Chapter summary**

Chapters 1 and 2 discussed how a system identification of the vestibular ocular reflex could be conducted during natural, head-free, gaze tracking. By applying an unpredictable perturbation to a subject's head, and controlling the trajectory of a visual target being tracked by the subject, the performance of the VOR can be evaluated as it creates eye movements that compensate for the disturbance motion of the head, and so enables clear vision. This chapter discusses the experimentation conducted to identify visual target and torque perturbation sequences that were appropriate inputs for the visual and vestibular systems. A PRBS torque sequence and a random target velocity sequence, with specified amplitude and frequency content, were found to stimulate the appropriate physiological responses from which the VOR could be evaluated.

## 4.1 Input signal requirements

It was shown in Chapter 2 that a linear system identification of the VOR could be conducted by disturbing the position of a subject's head with an applied torque perturbation while the subject tracked a moving visual target. The visual target acts as the experimentally controlled goal for the gaze system and directs the visually mediated head and eye tracking motions. The torque perturbation disturbs the position of the subject's head and stimulates the vestibular system to generate eye motions that compensate for the unwanted head motions, and so maintains clear vision. The performance of the VOR is evaluated by mathematically distinguishing the vestibularly driven eye component from the measured eye tracking response, and evaluating the head velocity – eye velocity relationship.

The appropriate design of the torque and the target input sequences is crucial to the analysis of the VOR. To stimulate adequately the visual and vestibular systems so that an identification of the VOR can be made with linear system identification techniques, the torque and target inputs must be designed to meet four general criteria. First, the inputs must have frequency spectra and amplitude distributions rich enough to stimulate smooth pursuit tracking and VOR responses that span, or exceed, the full range of natural behavior. This ensures that the evaluation reflects the full, functional, system level dynamics of the VOR. To match the disturbance motions experienced by the head during natural activities, the torque perturbation should stimulate horizontal head motions between 0.5 and 3.0 Hz, and create head velocities that are as high as 90 deg/s (Demer *et al.*, 1991; Grossman *et al.*, 1989). The magnitude of the torque must be sufficient to induce a measurable VOR response, but cannot be so high that subjects are prevented from voluntarily controlling head position while tracking the target. To encourage smooth eye and head motions, the visual target should move smoothly across the screen with velocities that approach 30 deg/s and with frequencies from 0.0 Hz to 2 to 3 Hz; the upper limits



are determined by the subject's ability to track the target without excessive error and without relying upon saccades.

Second, the input sequences must be unpredictable. Subjects must not develop a strategy of stabilizing the head against the torque perturbations, nor drive their gaze from a prediction of where the visual target will travel to or what direction the torque will move the head. The eye and head motions must be causal, not anticipatory, responses to the controlled inputs so that the calculated performance of the VOR and the gaze tracking systems reflect the ability of the underlying reflexive systems, not of the systems modified by the subject's ability to model and predict the head and target's motion.

Third, the torque and target inputs must be designed to limit the occurrence of nonlinear head-neck and gaze tracking responses. These include responses, such as making saccadic eye movements and modulating the properties of the head-neck system, that occur naturally during gaze tracking. These nonlinearities can be considered as outputs of systems that are beyond the controlled bounds of the experimental protocol, which is focused on the VOR's stabilization of smooth pursuit eye motions. When present, these nonlinearities change the dynamic relation between the controlled inputs and the measured response, and reduce how well a linear systems analysis can capture the dynamics of the VOR. Last, and perhaps most important for the mathematics of the VOR analysis, the torque and target sequences must be statistically uncorrelated so that the visually mediated and reflexive vestibular components of the eye motions can be quantified, and the dynamics of the VOR identified.

To conduct the proposed VOR testing protocol, a large set of input signals, composed of pairs of torque and target sequences that met the four design criteria, was needed. The set had to be extensive so that several subjects could be tested, and so that the limits of each subject's tracking and VOR performance could be explored. A protocol similar to this had not been used before to test the VOR or gaze tracking. It was not known, therefore, what types of signals would best stimulate the natural responses of the visual and vestibular systems and still permit the

experiment to be controlled in such a way that the analyses were accurate. Thus, the first step in designing the input signals was to determine experimentally the fundamental signal characteristics (e.g. signal type, frequency spectrum, amplitude distributions, duration) that the torque and target inputs needed to have in order to stimulate, individually, the head-neck system and the gaze tracking system. Once it was understood how to stimulate the head and the tracking systems separately, the torque and target inputs were delivered to the subject simultaneously to validate that the systems still responded within the bounds required by the protocol.

It should be stressed that the goal was to generate a set of signals, and a systematic method for creating the signals, that would satisfactorily meet the needs of the VOR testing protocol. The goal was not to identify the optimal signals to test the systems, though of course, the desire is to achieve that eventually.

## **4.2 Torque perturbation**

Torque sequences that were appropriate for use in the VOR testing protocol were identified primarily from the results of two sets of experiments. In the first set, unpredictable torque perturbations were applied to a subject's head while the subject tracked a target that moved with a sinusoidal trajectory. These experiments were conducted to identify the amplitude of the torque perturbation necessary to disturb the head as subjects tracked a moving target. The second set of experiments, in which subjects actively corrected for disturbances to the head's position that were caused by a pseudo-random binary torque perturbation, was conducted to validate that the pseudo-random binary sequence (PRBS) was an acceptable input for the VOR protocol, and to characterize the dynamic response of the head-neck system to the horizontal plane perturbations. The experiment and its results are fully described in Chapter 6.

#### 4.2.1 Tracking sinusoids while head was disturbed by torque perturbations

The first experiments that were conducted with the head-perturber were designed to quantify the magnitude of the torque perturbation that was needed to disturb the head as the head was voluntarily moved to follow a sinusoidal target. During the development of the head-perturber, it had been observed that broadband inputs, with amplitudes as low as 0.1 N.m, delivered enough power to a subject's relaxed head and neck for the dynamics of the head-neck system to be identified. This magnitude was only ten percent of what had been estimated to be needed, and indicated that the relaxed head and neck were very compliant and could be sufficiently disturbed by a small amplitude torque. It was not known if the torque required to perturb the head would increase when the neck was not relaxed. The stiffness and damping properties of muscle are known to change with the load and displacement of the muscle (Kearney and Hunter, 1990). The required perturbation magnitude could change dramatically when the neck muscles were being used to move the head to follow a target. The level could also vary as the control and stabilization requirements of the head changed. Experiments were conducted to identify the minimum perturbation amplitude needed to elicit a measurable VOR response while the subject used head and eye motions to track a sinusoidal target.

#### *Experimental methods*

Subjects were seated in a rigid chair under the helmet of the head perturber. The helmet was lowered onto the subject's head, and tightened with head and chin straps so that it did not slip when the torque was applied or when the subject moved the head. The head perturber was adjusted so that the head and the helmet could rotate without there being any significant eccentricity between the neck's and the helmet's rotational axes. This allowed the subjects to rotate the head without interference from the perturber. Only two subjects were used in this experiment, as the goal was to develop an appreciation for the level of torque necessary to disturb the head. A laser dot was projected onto a flat screen 1400 mm in front of the subject, and moved

in the horizontal plane with a sinusoidal trajectory. The sinusoidal trajectory was chosen because it is one of the easiest trajectories to track. Single frequency sinusoids are very predictable, and can be learned by a subject within one half cycle. Subjects can often track sinusoids with little gaze error and slip, and at speeds much higher than that for unpredictable targets (Baloh *et al.*, 1980; Young and Stark, 1963; Rashbass, 1961). The sinusoidal target trajectories were moved through 83 degrees at frequencies of 0.15, 0.30, 0.45, and 0.60 Hz. These frequencies resulted in peak target velocities of 39.4, 78.8, 118, and 157.6 deg/s respectively, which exceeded what is considered to be the upper velocity limits for smooth pursuit tracking of predictable targets (Leigh and Zee, 1999). The subjects were required to track the target using natural head and eye motions. The large angle through which the target traveled discouraged the subjects from tracking the target using only eye movements. While the subjects tracked the target, a pseudo-random binary torque perturbation was applied to the subject's head with the head perturber. Pseudo-random binary sequences (PRBS) are commonly used for disturbance signals because they approximate an unpredictable, bandlimited, white sequence even though they are deterministically generated, and have a binary amplitude distribution (Godfrey, 1993; Davies, 1970). At each frequency tested, perturbation amplitudes from 0.05 to 0.25 N.m were applied in increments of 0.05 N.m. The maximum amplitude was limited to 0.25 N.m, because at that amplitude it was obvious that the head was being disturbed and the VOR was creating eye motions to maintain clear vision. The applied torque perturbation, the target's trajectory, the subject's head velocity, and the binocular eye position were measured and sampled at 100 Hz.

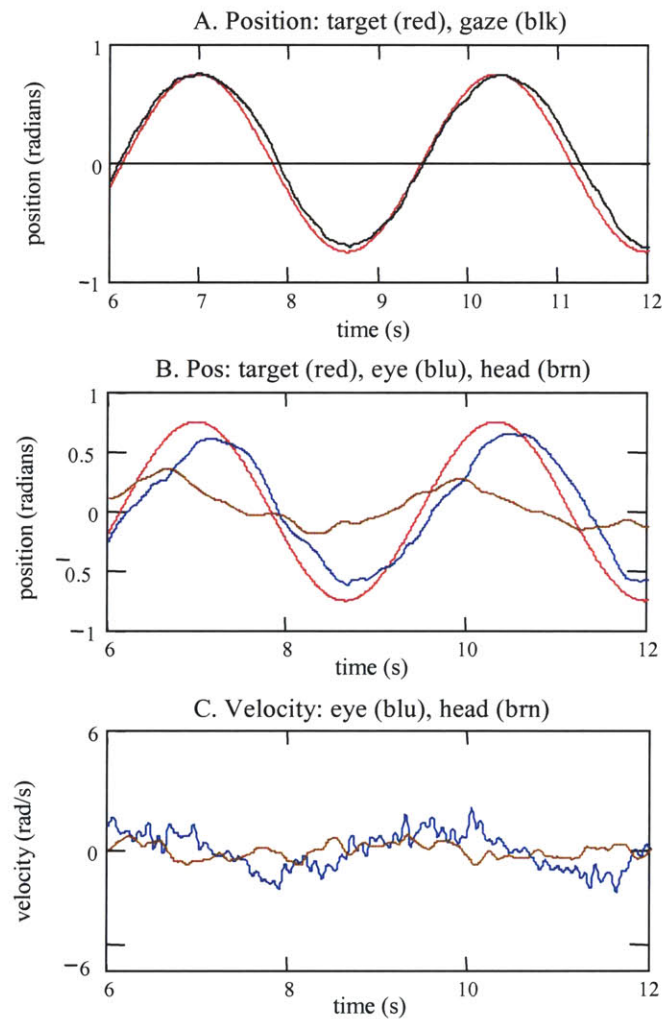
### *Results*

As expected, as the level of torque was increased within trials conducted at the same frequency, the VOR driven eye motions became stronger, and an equal and opposite relationship between the head and eye motions became more apparent. The plots of Figures 4.1 and 4.2 show data from a subject tracking a 0.3 Hz sinusoid. At low amplitudes, the torque perturbations did

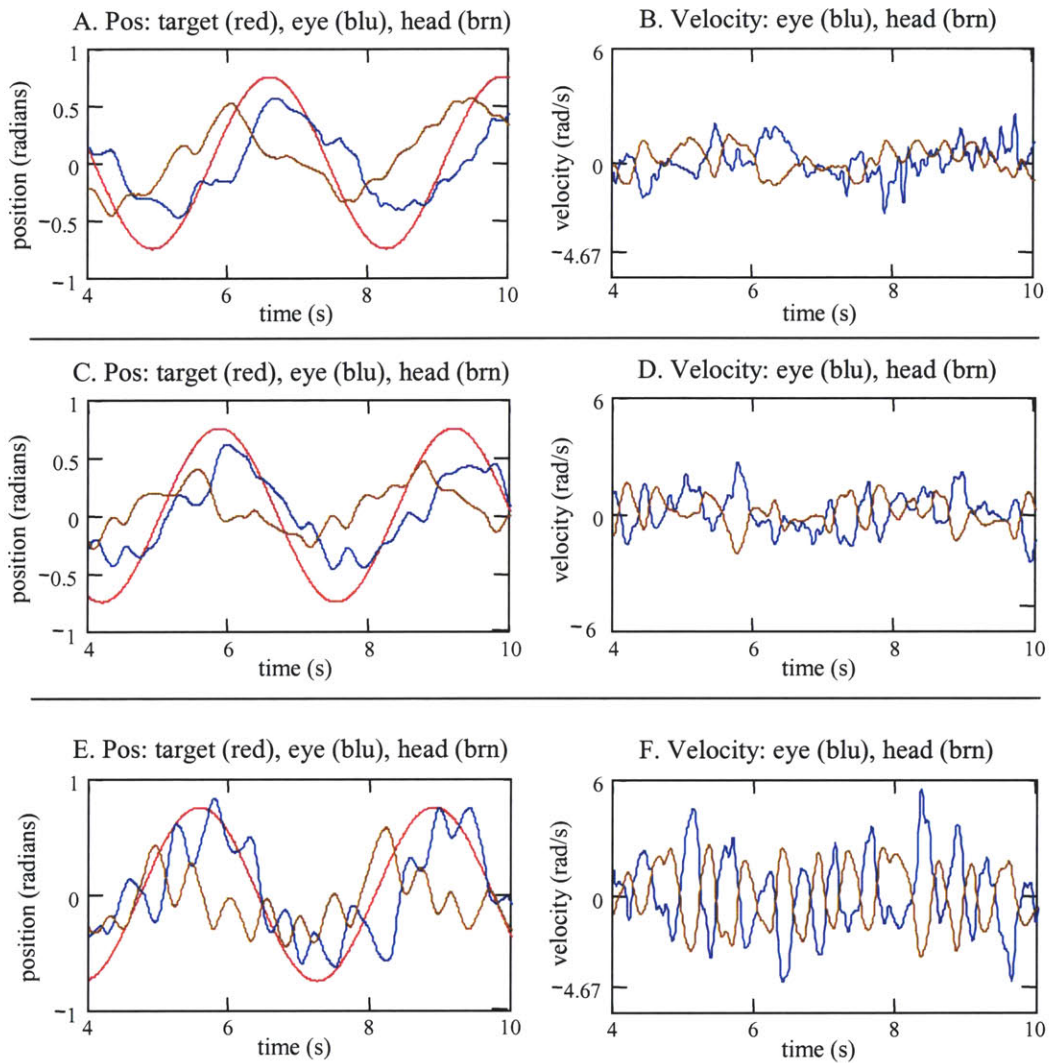
not disturb the head's position enough for the compensatory VOR to be observable, as shown in Figure 4.1 for a torque of 0.05 N.m. The subject's gaze did not match the target's motion exactly, but it did follow the target well. Although at times it led, and at other times lagged the target, the gaze response moved approximately in synchrony with the target. This is a result of the subject being able to predict the sinusoidal trajectory of the target. The head led the target, and its peak amplitude was about half that of the target. Eye position, which was measured relative to the head, lagged both the head and the target, and its amplitude was much closer to that of the target. Both the head and the eye moved in a smooth periodic path, similar to the sinusoidal movement of the target. Together, the eye and the head summed to a gaze that moved with the target, and equaled its amplitude. The trend of the head leading the target and the eye lagging the target was seen at all frequencies, although as the target frequency was increased the head's lead and the eyes' lag were less evident. VOR induced eye motions began to develop as the torque amplitude was increased. When the torque was 0.10 N.m, the perturbation affected the position of the head and eyes, and it became evident that the eyes were compensating for the head's motion as shown in Figure 4.2. Once the torque had reached 0.2 N.m, the head perturbation produced an obvious change in head and eye position, with the eye velocity and the head velocity moving in opposite directions. The head and eye position plots no longer looked like the sinusoidal movement of the target. The eye velocity was almost equal and opposite to that of the head. If they were exactly equal and opposite, the velocities would sum to zero, and would not create a gaze response that followed the target's sinusoidal trajectory.

As the frequency of the target was increased, and therefore also its speed, higher amplitude torque perturbations were needed to produce an obvious VOR response. In the plots shown in Figures 4.1 and 4.2, in which the target was traveling at a frequency of 0.30 Hz, the 0.15 N.m torque was sufficient to produce a VOR response. A perturbation of 0.15 N.m was still adequate when the target traveled at 0.45 Hz, but was not sufficient when the frequency increased to 0.60 Hz. Figure 4.3 compares the effect that the 0.15 N.m torque had on the subject when

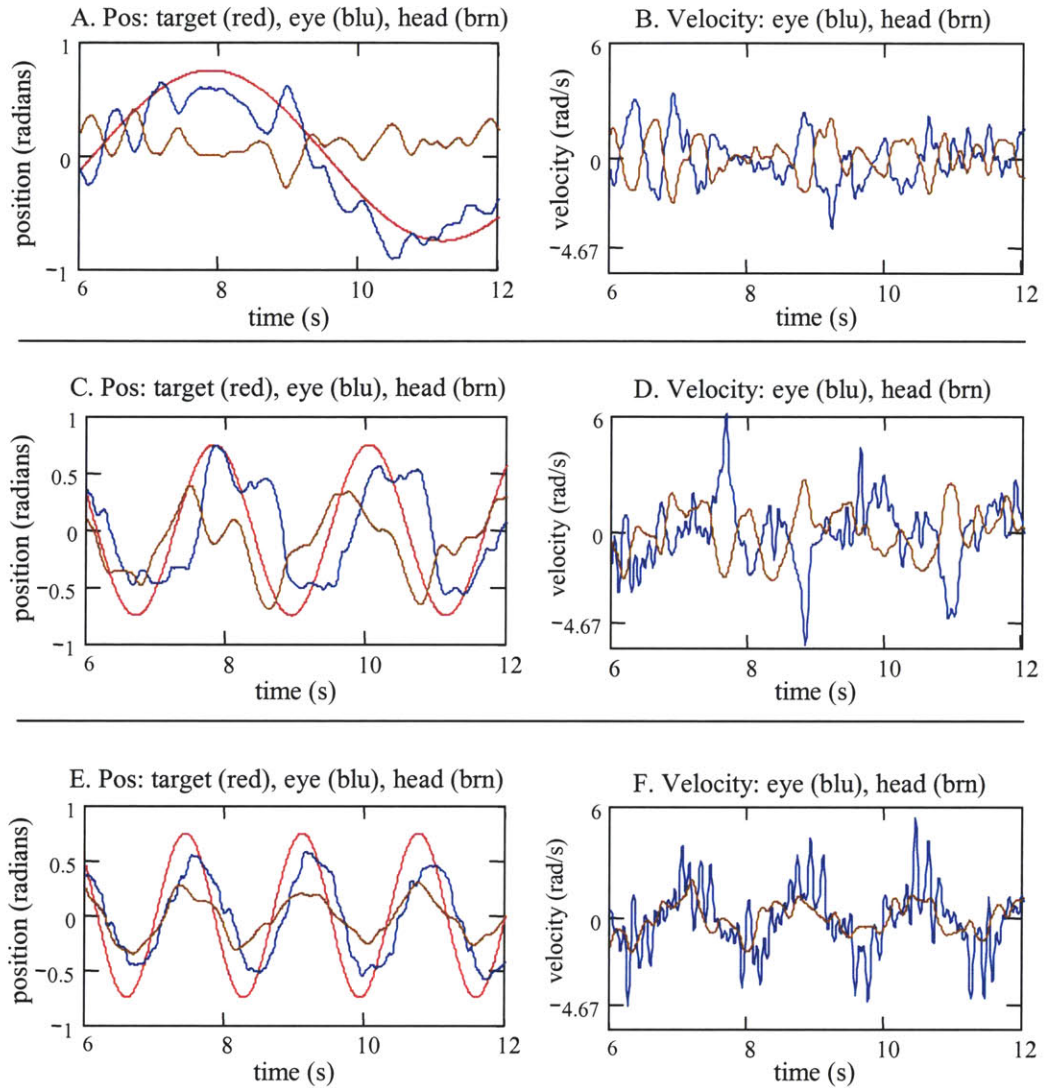
tracking the 0.15, 0.45, and 0.60 Hz targets. At 0.15 Hz, the 0.15 N.m torque created a significant disturbance to the subject's head. In both the position and the velocity plot it can be seen that the head has been displaced and is not moving on a smooth trajectory. The VOR eye motions are obvious in the velocity plot, and are even clearer in the position plot. When the target was increased to 0.45 Hz, the head did not appear to be as strongly perturbed as when the target was traveling at 0.15 Hz (Figure 4.3). The eyes still moved with a component opposite to the head, but the velocity was not a smooth mirror image of the head. For targets traveling at 0.6 Hz, the head was not disturbed much by the torque.



**Figure 4.1:** Tracking response to a 0.3 Hz sinusoid when a 0.05 N.m PRB torque perturbation was applied to the head.



**Figure 4.2:** Progression of eye and head motions as the amplitude of the PRB torque perturbation was increased. The target frequency was held constant at 0.3 Hz. The torque amplitude was 0.10 N.m in plots A and B, 0.15 N.m in plots C and D, and 0.20 N.m in plots E and F.



**Figure 4.3:** Progression of eye and head motions as the target's frequency was increased. The amplitude of the torque perturbation was held constant at 0.15 N.m. The target frequency was 0.15 Hz in plots A and B, 0.45 Hz in plots C and D, and 0.60 Hz in plots E and F.

Overall, torque perturbations with amplitudes of 0.20 to 0.25 N.m created measurable VOR responses across all frequencies tested. These were about the minimum amplitudes needed at the highest frequency of 0.6 Hz, and although higher than required at the lower frequencies,



subjects were still able to control their head and to track the target well at these low frequencies. The 0.20 to 0.25 N.m amplitude was significantly lower than the original estimate of 1.0 N.m. This estimate was based on the assumption that a torque between 5 and 10 % of the maximum torque that could be generated by the neck in rotation would be required to create an adequate disturbance. Little was found in the literature, however, about the maximum rotational torque that can be generated by the neck, although there were data on the maximum neck forces that could be generated in flexion, extension, and lateral flexion (Vasavada *et al.*, 2002). The maximum rotational torque was therefore calculated from estimates of the head's rotational inertia and peak acceleration. It was assumed that an average male skull has a radius of 100 mm, and a density equal to that of water. This results in a rotational inertia of between  $1.57 \times 10^{-3} \text{ kg.m}^2$  (if approximated as a sphere) and  $1.68 \times 10^{-3} \text{ kg.m}^2$  (if approximated as a cylinder). These assumptions were confirmed by a study of 21 cadaver heads that found that the mean rotational inertia of the head about the vertical axis is  $1.5 \times 10^{-3}$  (Beier *et al.*, 1980). Since the head can attain peak accelerations between 1000 and 1500  $\text{rad/s}^2$ , the estimated peak torque was between 16 and 24 N.m. The selection of 1.0 N.m fell on the low side of 5 to 10% of this maximum torque range. However, it was later observed that a healthy 32 year old male found it difficult to rotate his head in the horizontal plane against more than 4.5 N.m of torque. If 4.5 N.m is assumed to be near the maximum torque that can be generated by the neck in pure rotation, then the neck is not nearly as strong as had been estimated, and it seemed more reasonable that a torque of 0.2 N.m would be adequate to disturb the head's position during tracking.

#### 4.2.2 Dynamics of the head-neck in response to horizontal plane torque perturbations

Torque perturbations that were driven from a PRBS worked well during the tests in which the subject tracked the sinusoidal targets, and so appeared suitable for use in the VOR protocol. They were easy to generate with the characteristics needed to disturb the head's position, were well tolerated by the subjects, perceived as unpredictable, and were reproduced

consistently by the controller and motor. The one potential drawback to the PRBS was its limited binary amplitude spectrum, but since 0.2 N.m had been found to be a good amplitude over a range of voluntary head motions, a single amplitude PRB torque sequence seemed adequate for testing the VOR under natural, head-free, gaze tracking conditions. It still needed to be determined whether the PRBS stimulated the desired frequency response from the head, and if there was a linear relationship between the torque input and the head's response. Since the same torque amplitude could create different responses at different frequencies, it was possible that the amplitude that stimulated a linear response under some conditions could drive the head-neck system beyond a linear range under other conditions. A linear relation between the head-neck response and the torque input was necessary for the linear analysis of the VOR to be valid.

A study was conducted to examine the dynamics of the head-neck system in response to horizontal-plane torque perturbations, and to understand how the dynamics of the system changed as a function of the magnitude of the static torque applied to the head. A quasi-linear testing and analysis approach was used to quantify the dynamics of the head-neck system at different levels of mean (static) torque. An operating point for the head-neck system was established by applying a static torque to the head with the head perturber. The system dynamics were then probed by disturbing the head-neck system with a rich spectrum, PRB torque disturbance added on top of the static torque. What follows is a summary of the experiments, results, and their implications to the VOR testing protocol. A complete description of these experiments and the experimental results can be found in Chapter 6.


The physical set-up of the head-neck study was similar to that used when the subjects tracked the sinusoidal target, but the testing procedure and the task given to the subjects were different. Each subject was seated under the head perturber in front of the screen, and the head-perturber was attached to the head. Once the helmet was adjusted, subjects were asked to rotate their head to a natural, "zero rotation" position. A laser pointer mounted on top of the helmet was illuminated, and its projection was marked on the screen as a visual reference. During

experimentation the projection of the laser relative to this mark acted as visual feedback so subjects would know where their head was pointing. Each experimental trial began by applying a static (constant) torque to the helmet, which tended to displace the subject's head away from the zero-rotation position. The static torque ranged from 0.0 to 0.65 N.m, and was applied in increments of 0.05 in both the clockwise and counterclockwise directions. Subjects were instructed to compensate for the torque, and to rotate their head so that the laser projection was re-aligned with the zero-rotation mark. A small amplitude, zero mean, PRBS perturbation torque was added to the static torque. This perturbation shook the subject's head about the zero-rotation point, effectively acting as a disturbance to a system defined at an operating point established by the static torque. Subjects were instructed to correct for any position disturbance, and to try and keep the laser pointer aligned with the zero-rotation marker. Trials were run for 60 to 90 s. Six subjects, one woman and five men, participated in this study. Their ages ranged from 22 to 33. All were healthy graduate students with no history of or apparent neuromuscular or vestibular problems, and had normal to corrected-to-normal vision.

The torque applied to the helmet, the helmet's rotational velocity, and the head's rotational velocity were sampled at 100 Hz. The data were analyzed as representing measurements of a quasi-linear system that had been disturbed about an equilibrium point. The system's equilibrium point and its associated dynamics were set by the system's reaction to the applied static torque, and the disturbance about this operating point was created by the torque perturbation. The admittance impulse response function (IRF) was calculated for each experiment, and a second-order parametric model was fit to the IRF. Admittance is defined as the relation between a torque input and a velocity output, which is the mechanical analog of the electrical admittance.


The PRB torque sequence was executed at 10, 15.4, and 20 changes per second, and scaled for a magnitude of +/-0.1 or 0.2 N.m. The execution rates were selected to disturb the head throughout the 0.6 to 2.7 Hz range of frequencies that Grossman et al. (1988) had found to

be predominant for the head in the horizontal plane during locomotion, and for which the VOR is normally required to stabilize vision. Representative power spectral densities (PSD) of the torque input and the head velocity output are shown in Figure 4.4. The torque input was successful at stimulating the head throughout, and beyond, the 0.6 to 2.7 Hz bandwidth. The PSD of the torque was highest near 0.0 Hz, gradually decayed as the frequency increased, and then dropped sharply near the frequency at which the PRBS was executed. The PSD of the head velocity was strong from approximately 0.5 Hz to beyond 5.0 Hz, and its peak value usually occurred between 1 and 3 Hz. Below 0.5 Hz the power in the head velocity decayed rapidly. In addition to disturbing the head across a wide frequency band, the randomness of the PRBS prevented the subjects from predicting the sequence of torque perturbations. Subjects reported that, although the perturbations were mild and did not create large head motions, the randomness of the sequence made it very difficult to bring the laser back to the zero position after being disturbed. By the time they were attempting to correct for the disturbance, and had generated a motion towards the zero-rotation position, another perturbation would displace the head. A critical analysis of whether the torque had stimulated the head velocity sufficiently comes from the impulse response analysis that relates the torque input to the velocity output, and from calculating the coherence function which can identify if nonlinearities or a poor signal to noise ratio corrupted the data.



**Figure 4.4:** Power spectral densities of a PRB torque perturbation input (plot A) and a head velocity output (plot B).


The IRF analysis indicated that between 0.5 and 10.0 Hz, the head-neck system behaved like a quasi-linear, passive system in response to small stochastic inputs applied to the head in the horizontal plane. When the mean torque was held constant, and the system perturbed with the PRB torque sequence, the head-neck system responded much like a linear, second-order, underdamped system. The system was not, however, strictly linear. The properties of the head-neck system were sensitive to the static component of the input torque, and the characteristic impulse response function changed markedly as the mean torque level was changed (Figure 4.5). As the static torque was increased, the effective stiffness and damping of the system, which represented the combined effects of neuromuscular reflexes, voluntary control, and passive viscoelastic properties of muscle, progressively increased, but did so such that the system's second-order damping ratio remained essentially constant across a mean torque of 0.0 to 0.65 N.m.



**Figure 4.5:** Non-parametric IRFs (plot A) and second-order parametric IRFs (plot B) for the head-neck system. A mean-torque of 0.10 N.m (solid), and a mean torque of 0.55 N.m (dotted) were applied to the head.

At frequencies below 0.5 Hz and above 10.0 Hz, the characteristics of the system were not captured by the linear input-output analysis, and the system's behavior could not be described by constant parameter, second-order models. Low coherence-squared values indicated that at these frequencies the analysis was corrupted by system nonlinearities and/or measurement noise.

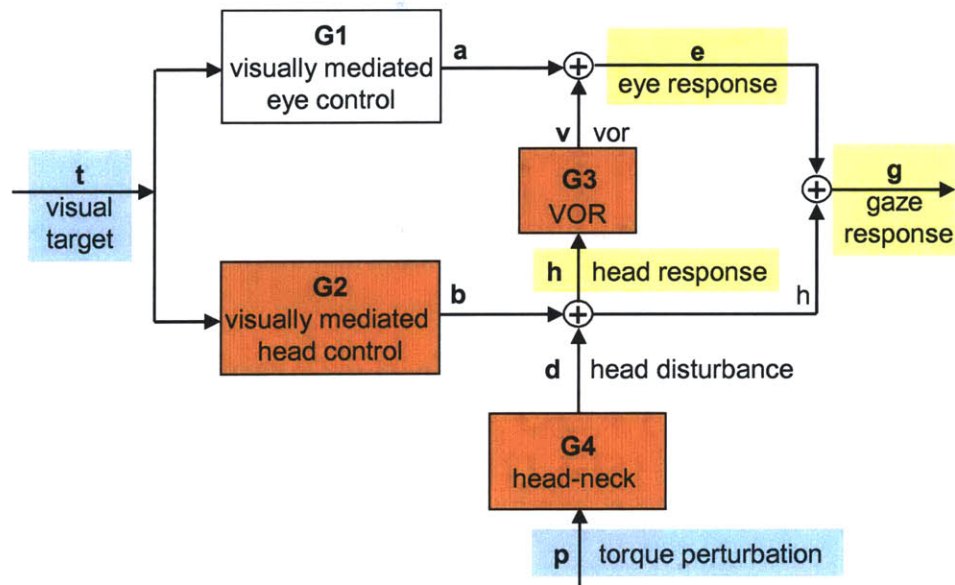
Above 10 Hz it is evident that the drop in coherence was due to low input power, which reduced the signal-to-noise ratio. As can be seen in Figure 4.6, around 10.0 Hz the drop in coherence mirrors the drop in input power. The input power remained high below 0.5 Hz, so the low coherence-squared measures at these frequencies were not caused by a low signal to noise ratio, but by the nonlinearity of the active control system that dominates the head's response at low frequencies.



**Figure 4.6:** Coherence-squared function for a torque input and the head velocity output (plot A), and the power spectral density for the torque input (plot B).

The results of the tracking and the head-neck experiments supported two important conclusions with regard to the VOR testing protocol. A PRBS input can stimulate the head through the desired frequency band for which the VOR naturally functions to stabilize vision, and does so in a manner that is unpredictable, but still allows a subject to control the head and to track a target. When stimulated by a low amplitude PRB torque sequence, a linear relation was preserved for the head-neck system, and the VOR response could be analyzed with the linear method's described in Chapter 2.

### 4.3 Visual target



**Figure 4.7:** Block diagram for the gaze tracking system.

#### 4.3.1 Requirements for the visual target

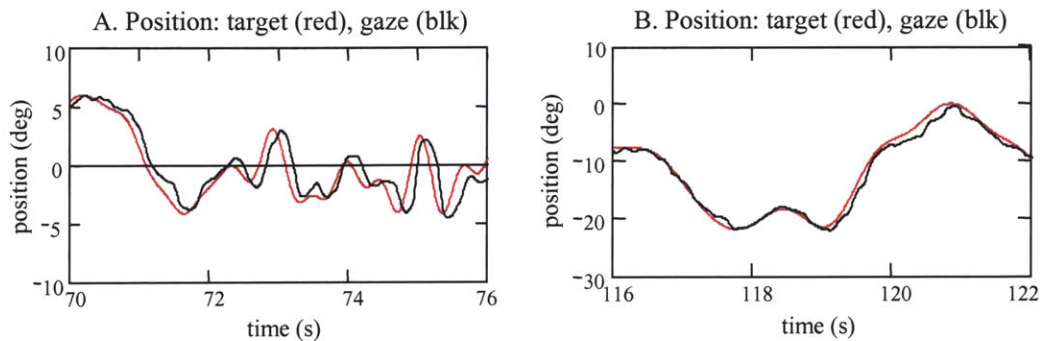
A visual target provides a clear, controlled goal for the gaze tracking system, and ensures that the VOR is evaluated under conditions that encourage natural system behaviors. The target drives the visually mediated eye and head sub-systems, which are shown as  $g1$  and  $g2$  in Figure 4.7. During the experiment, the target must be tracked accurately so that a strong target-gaze relation exists, and it should also stimulate a range of gaze movements so that the resulting evaluation reflects the VOR's performance throughout its full range for which it stabilizes gaze. A linear identification of the VOR requires that the target's trajectory is unpredictable, and that the target is tracked primarily with smooth eye and head motions. This minimizes the use of predictive tracking mechanisms and saccades which are natural oculomotor responses that enhance the capabilities of head-free tracking, but which can introduce time-variance and non-

linearities into the gaze response. These dynamics would not be captured or modeled well by the proposed linear systems analysis.

It is difficult to create a target that is both unpredictable and that stimulates smooth pursuit tracking, as these two requirements are somewhat contradictory. The characteristics that make a target unpredictable also tend to lead to a tracking response that is populated with saccades. Some researchers use the breakdown in the smooth-pursuit response, as marked by the occurrence of saccades, to classify a target as unpredictable (Barnes *et al.*, 1987). The solution comes from understanding what is achieved by making the target unpredictable. As discussed in Chapter 1, natural gaze tracking is driven by a mixture of visual feedback and predictive mechanisms, and is characterized by smooth eye motions interspersed with saccades (Carl and Gellman, 1987). When a target first appears, it is tracked using visual feedback. Since 70 to 150 ms are required to process the target's image and to generate head and eye tracking movements, the visually driven gaze response lags the target (Figure 4.8a) (Leigh and Zee, 1999). As the target is followed, its motions are learned and a model of its trajectory is developed by predictive mechanisms in the tracking system (Wellenius and Cullen, 2000; Barnes and Grealy, 1992; Young, 1977). The signal driving the gaze system changes from one derived from the visual feedback signal to a signal based on the internal model of the target's trajectory. This often reduces tracking error by advancing the phase and improving the gain of the gaze so that it better coincides with the movement of the target (Figure 4.8b). Visual feedback is not ignored, but is used to monitor the error between the gaze and the target. If the target changes course, or if the model of the trajectory is inaccurate, the system reverts back to being driven by visual feedback, and the modeling process starts over (Barnes and Asselman, 1991). When the target's motion is predicted well, little error exists between the gaze and the target, and the gaze tracking response is smooth with few saccades. When the target is not predicted well, the gaze's position and velocity do not match those of the target, and saccades are used more frequently to keep the gaze close to the target. When the target cannot be predicted some saccades will occur. Since they cannot be



eliminated, the trajectory must be generated so that their use is minimized. The target must be tracked with visual feedback, but the error cannot be so great that saccades are frequently triggered to interrupt the smooth tracking.



**Figure 4.8:** Gaze tracking. In plot A the target was unpredictable and the subject's gaze (black) lagged the target (red). In plot B the subject was able to predict the target and there was little lag between the target and the gaze.

Rather than stating that the goal is to make the visual target unpredictable, it is more correct to state that the goal of the target is to stimulate a tracking response driven predominantly by visual feedback, and not by predictive mechanisms. Predictability is a difficult characteristic to quantify, and it is often not an “either-or” characteristic of a sequence (Michael and Melvill Jones, 1966). Unless the sequence is a strictly random, white sequence, it is likely that it will have some level of predictability. By making the target “unpredictable”, the anticipatory motions of a predictive response, and the time variance that occurs as the gaze tracking system switches between prediction and visual feedback, are avoided. A consistent response from the tracking system is desired throughout the experiment, not a response that varies as components of the tracking response are intermittently enhanced by predictive mechanisms. An invariant drive to the gaze tracking system helps the linear analysis, and also makes it easier to distinguish between eye motions that were driven by visual feedback and eye motions caused by the VOR, which have a much shorter latency of about 15 ms. If a tracking response is driven by visual feedback, and not

by predictive mechanisms, a consistent 50 to 150 ms lag should exist between the target and the gaze response at all frequencies for which the target was controlled.

#### 4.3.2 Generating experimental target trajectories using the Hunter-Kearney stochastic swap

It was anticipated that targets that would satisfy the visual system's input requirements could be created by driving the target from a random sequence with the correct balance of velocities and frequencies. Gaze tracking studies have demonstrated that the dominant factors that influence whether tracking is maintained by smooth pursuit or by saccadic motions, and whether the gaze is driven by visual feedback or by predictive mechanisms are the target's velocity, frequency spectrum, and whether its trajectory follows repeated patterns. The smooth pursuit system responds to both target position and target velocity, but studies indicate that it is the velocity of the image on the retina that is most important, and therefore the more appropriate characteristic to control to drive smooth pursuit tracking (Leigh and Zee, 1999). Smooth pursuit tracking has been found to be effective only with target velocities up to 10 to 40 deg/s; above that, saccades are needed to help keep the gaze on the target (Baloh *et al.*, 1980; Young and Stark, 1963; Rashbass, 1961). The predictability of the target, or, more precisely, whether gaze tracking was enhanced by predictive mechanisms or driven by visual feedback, is determined by the signal's frequency content, not its complexity (Barnes and Lawson, 1989; Barnes and Ruddock, 1989). For pseudorandom targets made from mixed sinusoids, the tracking of all the sinusoidal components is enhanced by predictive mechanisms when the maximum frequency is below 0.4 to 0.8 Hz. However, when the highest frequency component is above this critical range, only the tracking of the highest frequency component is enhanced by predictive mechanisms; the tracking of the lower frequency components is driven by visual feedback, and smooth pursuit motions are observed to break down (Barnes and Asselman, 1991). When the highest frequency component of a target made from as few as two sinusoids exceeds the 0.4 to 0.8 Hz, the lowest frequency component can be tracked as if it were unpredictable. By balancing the speed and the frequency

content of the target, and by using a random sequence, it was expected that the target could be made to stimulate smooth pursuit tracking responses driven by the visual feedback.

It seemed reasonable to use the sums of sinusoids as the basis for the set of target velocity sequences. “Sum-of-sines” sequences are often used in oculomotor studies because they are smooth, generally well tolerated by subjects, easy to define and to manipulate, and, as demonstrated by Barnes et al. (Barnes and Ruddock, 1989), simple to arrange so that the tracking of all frequency components, except the highest, is driven by visual feedback. Along with “constant velocity ramps”, mixed sinusoids have been used in many of the studies that established the body of knowledge about the gaze tracking system. It was thought that the findings from those studies would provide the basis for the development of targets that would satisfy the needs of the VOR testing protocol.

In practice, the sum-of-sines sequences had shortcomings that made them awkward to design and inadequate for use as the target’s velocity sequence. The obvious drawback to sum-of-sines sequences is that they stimulate a linear system only at distinct frequencies, rather than continuously across a frequency range. In the context of the VOR testing protocol, the role of the visual target is to control, or at least to direct, the gaze throughout the full bandwidth of smooth pursuit tracking. With a sum-of-sines sequence, the gaze system is only controlled at the discrete frequencies of the individual sinusoids, leaving many frequency bands for which the VOR and gaze tracking performance would not be explicitly analyzed. A large number of sinusoids had to be combined for the target to have a frequency spectrum that could approximate a continuous spectrum across the bandwidth that was to be tested. This created a practical challenge in generating a sequence that would remain within the physical bounds of the experiment. As more and more sinusoids were added together to create a more continuous frequency spectrum, it became increasingly difficult to combine their amplitudes, frequencies, and phases so that the target’s image remained within the limits of the screen onto which it was projected. The target could always be kept on the screen by abruptly stopping the sequence when it reached the ends of

the screen, but this affected the smoothness of the trajectory, and altered the target's frequency spectrum. Another simple, but also inadequate, solution was to keep the velocity amplitude of each frequency component small. When this was implemented, however, the target tended not to move over a sufficiently wide area of the screen, and did not stimulate much of a range of head and eye motions. These trajectories were often not interesting for the subjects to track, and as a consequence subjects were not engaged in the tracking task. It was not difficult to create a good sum-of-sines sequence, but each sequence's design became a laborious, manual process of deciding which frequencies and amplitudes to choose so that when projected onto the screen the target appeared appropriate for the subject and stimulated the right type of head and eye motions. It became apparent that it was impractical to define the large set of target trajectory sequences that were needed to conduct the VOR experiments with this method.

The use of sum-of-sines sequences to create the targets also failed to address the cross-correlation requirements for the target and the torque perturbation sequences. This was not important for creating a target that was only to be tracked, but was vital for creating the targets for use in the VOR experiments. The torque and target velocity sequences needed to span similar frequency ranges, and so it was unlikely, though possible, that when they were generated independently, they would be sufficiently uncorrelated for the VOR analysis to be conducted. When two sequences have similar frequency content, they have a non-zero cross spectral density (Jenkins and Watts, 1968). The cross-power spectral density (cross PSD) is the Fourier transform of the cross-correlation function, and so a non-zero cross-power spectral density results in a non-zero cross-correlation.

When two different types of signals are required to be uncorrelated, as is the case when using a PRBS for the torque and a sum-of-sines for the target, it is not uncommon to create many sequences that meet the individual stimulus criteria, and then to mix and match the sequences until two are found serendipitously that are uncorrelated. Although possible, and ultimately effective, this would not be an efficient approach for defining a large set of torque and target

sequences that are needed to investigate the tracking and VOR systems. A more direct way is to design specifically the sequences so that they meet the testing criteria.

It is relatively easy to generate two uncorrelated sum-of-sine sequences by ensuring that the sequences do not share the same, or harmonically related, frequencies (called interweaving). It is also possible to “interweave” the elements of two PRBS inputs so that the sequences are uncorrelated, but still approximate bandlimited white signals (Davies, 1970). The interweaving of sinusoids in the frequency domain, and the interweaving of PRB sequences in the time domain are very different processes, but are given the same name, and can result in uncorrelated sequences. The interweaving techniques do not work, however, when both signals are not PRB sequences or when both are not sinusoids.

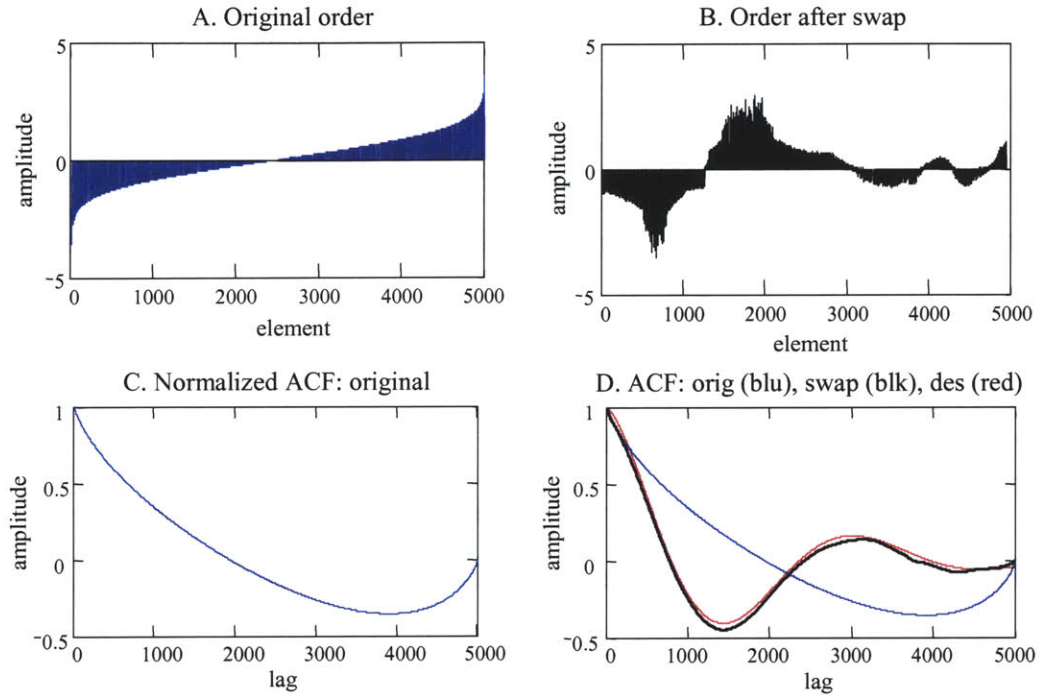
The solution to generating appropriate target sequences was found by using a method to generate random sequences that was developed by Hunter and Kearney (Hunter and Kearney, 1983). Their method uses an efficient, recursive, double stochastic interchange to generate random, discrete time, sequences with jointly specified auto-correlation and probability density functions. This allows for direct control of the target’s velocity distribution and frequency spectrum, the two characteristics that had been observed to affect most directly the target’s predictability and whether saccades were triggered during gaze tracking. Target sequences were generated in which the target’s velocity amplitudes were defined by a specific probability mass function (PMF, or probability density function for continuous time signals), such as being uniformly distributed at distinct velocities between -20 deg/s and 20 deg/s. The trajectory’s frequency content was defined by a desired auto-correlation function (ACF). Since the time domain ACF and the frequency domain power spectral density (PSD) are a Fourier transform pair, the target’s frequency spectrum could be defined to be from 0 to 2 Hz, for example, and the desired ACF calculated by taking the inverse Fourier transform of the PSD.

The Hunter-Kearney algorithm first creates a discrete time sequence with a desired amplitude probability density function (PDF) by transforming a uniformly distributed sequence

with a static filter made from the reversion of the desired cumulative distribution function (CDF). The algorithm then arranges the elements of the sequence with a recursive, double stochastic, interchange of elements so that the sequence's ACF converges to the desired ACF as shown in Figure 4.9. This is done by randomly choosing two elements of the sequence, swapping their positions, and then calculating the ACF of the reordered sequence. If the resultant ACF is closer to the desired ACF, the swap is kept. If not, the elements are returned to the positions they had before the swap was initiated. The Hunter-Kearney algorithm significantly reduces the computational time needed to calculate the ACF by updating the value of the ACF, rather than by explicitly recalculating the ACF, after two elements are swapped. The update is efficiently accomplished by considering only the products in the ACF summation that are affected by the exchange of elements. The effort to compute the ACF is independent of the number of elements in the sequence since a maximum of eight terms in the ACF summation (at any lag) are affected when the two elements are swapped. The process of interchanging elements and evaluating the ACF continues until the sequence's ACF is sufficiently close to the desired value, or until there is no more improvement. Reordering the elements affects the sequence's ACF and spectral content, but it does not alter the PMF of the elements. The PMF is determined only by the values of the elements in the sequence and the frequency with which they occur, not on how the elements are ordered. Although reordering the sequence to have a specific ACF does not affect the PMF, the PMF does affect the ACF that can be achieved by reordering the elements. The potential of the sequence to attain a desired ACF is dependent on the elements that are available in the sequence.

The Hunter-Kearney algorithm made it easy to generate target sequences that satisfied the criteria specified for stimulating the visual pathways. Because the velocities and frequencies were directly controlled, it was simple to manipulate target characteristics and investigate what properties lead to the target being tracked with smooth pursuit, but with the gaze response consistently lagging the target. Targets could be generated that had continuous spectra between 0.0 and a maximum frequency, which permitted the full range of gaze tracking movements, and

were therefore appropriate for system identification. A clear practical advantage of this approach over the sums-of-sines sequences was the direct control over the velocity which made it simpler to define the sequence so that the target remained on the screen when projected.



**Figure 4.9:** Stochastic swap. The sequence in plot A was reordered to that in plot B using the Hunter-Kearney stochastic swap. The sequence’s auto-correlation was changed from that in plot C, to the black plot in plot D.

The Hunter-Kearney technique did not, in its original form, address the correlation between the target and the torque inputs any better than when the target sequences were created from sums-of sines. As with sum-of-sines sequences, or with any sequence designed to control the visual path, a target sequence generated by the Hunter-Kearney technique would have to share a similar frequency band with the torque PRBS, and would likely be too highly correlated with the torque input to be used in the VOR analysis. By making a modification to the Hunter-Kearney algorithm, a target velocity sequence could be generated so that its cross-correlation with a torque perturbation input could be manipulated. This modification made it possible to create

torque and target sequences that met all four of the criteria required to conduct a system identification of the VOR. This modified version of the Hunter-Kearney technique is a very powerful method for developing input signals for system identification studies, in that it allows two very different sequences to be created with arbitrarily defined ACF, PDF, and CCF. It does not rely upon chance for the sequences to meet correlation requirements, nor is it restricted to de-correlating sequences of the same type, such as when interweaving two sum-of-sines or two PRBS. The technique is not limited to making the sequences uncorrelated, but can give the sequences any desired cross-correlation over the defined lag period. The modified Hunter-Kearney algorithm is more thoroughly discussed in Chapter 5.

#### 4.3.3 Gaze tracking experiment and results

Pseudorandom target sequences, with various frequency spectra and velocity distributions, were generated with the Hunter and Kearney technique and were judged as to their appropriateness by having subjects track their motion and evaluating the gaze-target, output-input, relation. The primary goal was to identify what velocities and frequencies were good target trajectories. The secondary goal was to learn how best to create the sequences so that all the testing requirements for the stimuli were met. A good target was defined as one that was well-tracked and that induced a gaze response that was driven predominantly by visual feedback and not by predictive mechanisms. “Well-tracked” meant that the gaze followed the target’s motions closely, that the gain relation between the gaze and the target was near unity over the frequency range for which the target was controlled, and that few saccades were needed to maintain the target’s image on the fovea. The position error between the target and the gaze is an indicator of how well a target is tracked, but since the error is caused both by the gain and phase response of the gaze, the magnitude of the error does not indicate how well the gaze matched the target’s motion. A gaze response that followed the target well could have a larger error than one that followed the target poorly, if there was a consistent time lag between the target’s motion and the




gaze reaction. The lag separates the gaze from the target, and may mask a gaze that closely followed the target, but that was delayed by the time required for the visual system to process the target's image. The phase lag, or alternately the time lag between the target and the gaze, was used as the primary indicator of whether the tracking was driven by the visual response or by predictive mechanisms. A phase lag that increased linearly with frequency would indicate that there was a constant time lag between the target's motions and the gaze response at all frequencies. When this time lag was near that expected for the visual delay, which is between 50 and 150 ms, it was a good indication that the gaze response was consistently driven by visual feedback. If predictive mechanisms affected tracking, it would be expected that the phase would be advanced at the frequencies for which the predictive mechanism was able to lock on to the target motion, and anticipate its movements.

### *Methods*

The target's velocity was created as a random sequence with the Hunter-Kearney algorithm, and then integrated into a position sequence using a discrete-time, trapezoidal, integration approximation (Press *et al.*, 1997). The position sequence was input to the galvanometer which positioned the target's image on the screen. This sequence was given a mean of zero, so that it was balanced on the screen, and was sometimes scaled so that it spanned the full width of the screen. Giving the target a zero mean position did not affect the velocity sequence, but the velocities were affected when the amplitude of the position sequence was scaled. The sequence was scaled so that the maximum velocity remained within the desired bounds. The first second and last second of the position sequence were replaced by a constant velocity ramp so that the sequence started and ended at the center of the screen (Figure 4.10). The sequences were generated so that the position command was updated at least 200 times per second. Targets with much slower update rates did not appear to move smoothly across the screen. There was little noticeable improvement in the target's smoothness when targets were

updated at rates faster than 200 times per second, but there was a definite advantage to keeping the update rates as low as possible. Faster update rates had more elements per second in the discrete time sequence, and therefore required more processing time and power when the sequence was being ordered to calculate its ACF and CCF.

The target's image was projected onto the screen, 1600 mm in front of seated subjects. The maximum velocity of the sequences ranged from 10 to 40 deg/s, and the frequency spectrum of the sequences began at 0 Hz and extended to between 1.0 and 4.0 Hz, although some tests were run with targets that extended to 10 Hz. Subjects were asked to follow the target's motion as best they could with natural head and eye motions. During these trials the subjects' heads were placed into the helmet of the head perturber, but the perturber was not activated. This prevented the subjects from translating their head and changing the distance between their eyes and the screen, but allowed for free rotation. The subjects' gaze was calculated from measurements of the head's velocity and the eyes' position. Individual trials lasted between 30 to 90 s, and subjects completed up to 15 trials. Seven different subjects were tested over the course of the experiment; some during the initial development, and others as the testing protocol and development of the input sequences evolved. A few of the subjects were tested more than once during the course of the experimental development. All the subjects reported that they were healthy. Three of the seven wore corrective lenses or contact lenses during the experiment.



**Figure 4.10:** Target position (plot A) and target velocity (plot B) sequences.

## *Results*

Targets that were tracked well, with a good gaze-target gain relation, small error, without many saccades, and with a gaze response that had a consistent lag and appeared to be driven by delayed visual feedback, had frequency spectra that were continuous from 0 Hz to between 1 and 2 Hz, and velocity distributions in which the majority of the sampled amplitudes were under 10 to 20 deg/s. The particular trajectory characteristics that led to good tracking were dependent on the subject, but the values specified above generally agree with what has been found in other studies (Barnes and Lawson, 1989; Baloh *et al.*, 1980; Young and Stark, 1963; Rashbass, 1961). When the target's spectrum was close to or extended beyond 2 Hz, subjects often reported that the target was difficult to follow, and the tracking data showed that the gaze response contained more saccades. When the frequency spectra were reduced to below 1 Hz, subjects found the trajectories easier to follow, and saccades were used less frequently than for the higher frequency targets, but it was evident that predictive mechanisms enhanced the tracking responses. The phase delay between the gaze and the target was small, and did not consistently increase with frequency. The target's velocity also affected how well it was tracked. Gaze slip generally increased steadily with frequency and with the target's velocity, but increased more rapidly as the target's velocity rose over 15 to 20 deg/s.

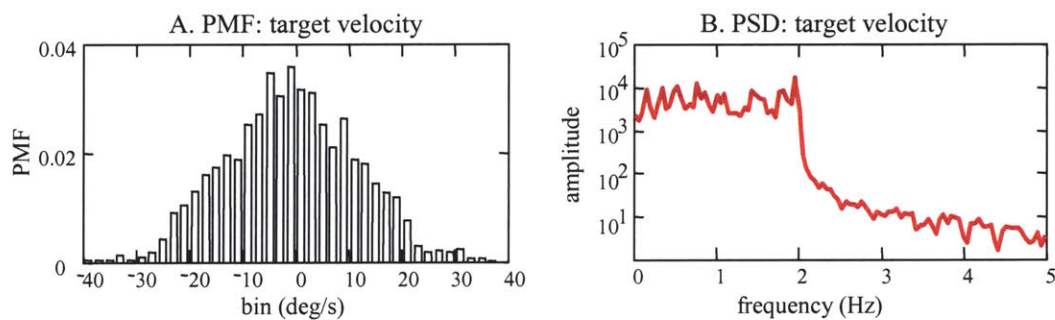
### *Well- tracked, unpredictable targets*

Figures 4.11 through 4.18 show representative plots of targets that were considered to have been well tracked. Data were sampled at 100 Hz, filtered with a small, three or five-element neighborhood, median filter, and then low-pass filtered at 15 Hz. The 15 Hz cutoff was high enough to retain all the information about smooth pursuit motions, and although it removed some of the detail in quick phase saccades, it was low enough to remove some of the inherent noise that corrupts EOG measurements. Filtering greatly improved the time domain plots of the data, but did not have a great impact on the linear, input-output analysis, as long as the cutoff remained

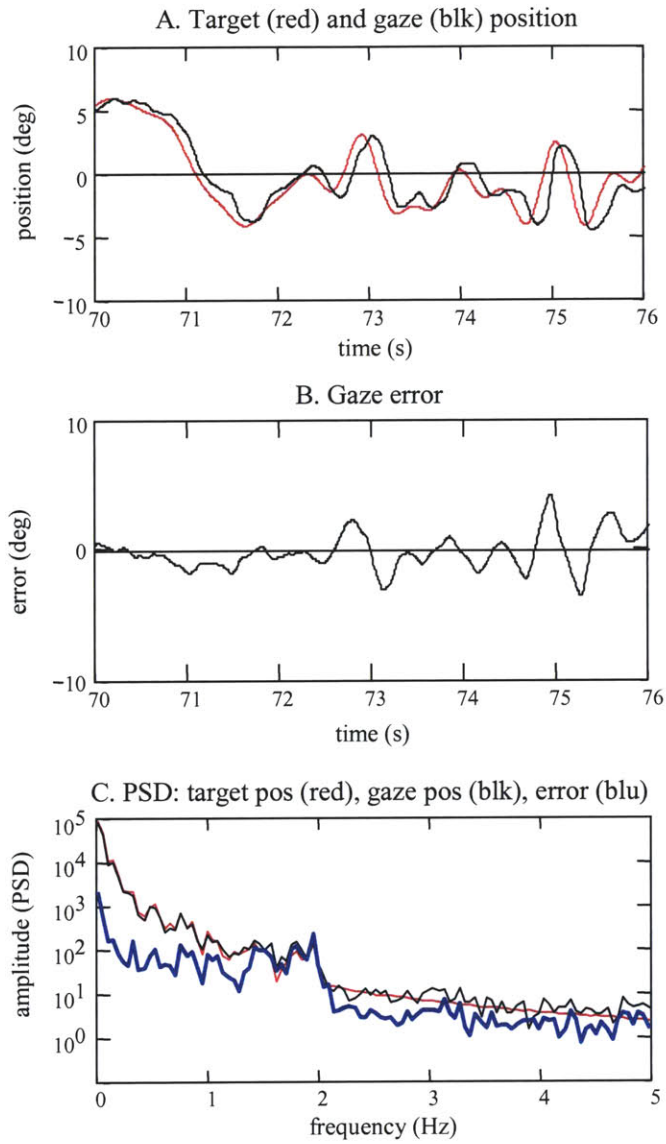
higher than the bandwidth of the target. The interpretation of the frequency and correlation relations for the gaze tracking response were unchanged independently of whether the data were unfiltered, filtered at 15 Hz, or even at 2 Hz, as long as the filter's corner was above the target's frequency spectrum.

The data shown in the plots in Figures 4.11 to 4.18 are from a target sequence with characteristics that were closer to the upper frequency and velocity limits that generally produced a good tracking response. The frequency spectrum of the target's velocity extended to 2.0 Hz, and the probability mass function (PMF) showed that the majority of the target's velocity was below 15 deg/s (Figure 4.11). It can be seen in Figure 4.12 that the subject's gaze position followed the target's position well. The tracking error was small, with a standard deviation of only 1.7 deg. The subject's gaze was never as smooth as the target's motion, as the gaze tracking system was reacting to the unpredictable movements of the target, and it is obvious that the gaze lagged the target. The lag was most consistent at times when the target's position changed most rapidly, between 74 to 76 s for example, and there was little opportunity for the subject to learn the trajectory. The gaze and target traces are similar, but are offset. A portion of the gaze error (gaze error = target position – gaze position) can be attributed to the delay between the gaze and target, although much of the error was due to the gaze system simply not being able to, or needing to, match precisely the target's motion. When the gaze position was advanced by 100 ms, the mean square error between the gaze and the target was minimized. The standard deviation of the error was reduced from 1.7 to 1.2 degrees, not to zero, which would be the case if the gaze had followed the target position exactly (Figure 4.13). Although the subjects reported that they felt they had tracked these targets well, the gaze position did not mirror the target's position. However, it is not necessary to maintain zero gaze error when tracking a target. Studies have concluded that visual acuity can be maintained when the gaze is within 3 deg of a target, and the gaze slip less than 5 deg/s (gaze slip = target velocity – gaze velocity (Leigh and Zee, 1999)). Since the target was a blurred circle, the gaze tracking system did not need to focus the image

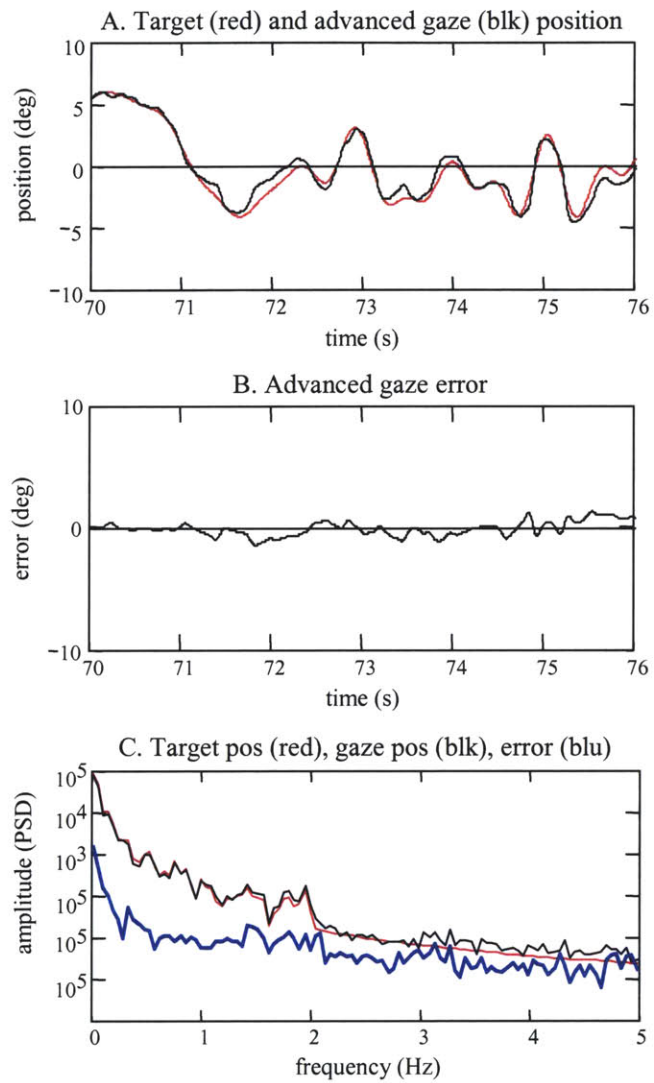
precisely at the center of the fovea. The 100 ms delay between the target and the gaze fell within the 50 to 150 ms normally expected to process the target visually, and implied that, on average, the tracking was driven by visual feedback, not by predictive mechanisms. If the target's motions had been consistently predicted by the subject, the gaze would have been advanced so that it was more in phase with the target, which did occur when the same subject tracked a sinusoid (Figure 4.14). The subject anticipated the sinusoidal trajectory very well and there was nearly zero lag between the target and the gaze. Gaze and target velocity plots for the random target are shown in Figure 4.15. The gaze velocity did not match the target velocity nearly as closely as the gaze position matched the target position which supports the historical claim that the gaze system is a "sloppy" velocity controller (Williams and Fender, 1979). Some of the high frequency eye velocity and gaze slip can be attributed to the "calculation noise" that was generated when the eye velocity sequence was calculated as the derivative of the sampled eye position sequence. The velocity was maintained usually to within 10 deg/s of the target velocity. This compared well to the slip that occurred when the same subject tracked the sinusoid (Figure 4.16). The sinusoid was tracked with less position error, but the gaze velocity was still seen to follow the target with slips often in excess of 10 deg/s. Even though the target's movements were anticipated, and the gaze appeared locked on to the target, high frequency gaze motions were used to keep the gaze in the vicinity of the target.



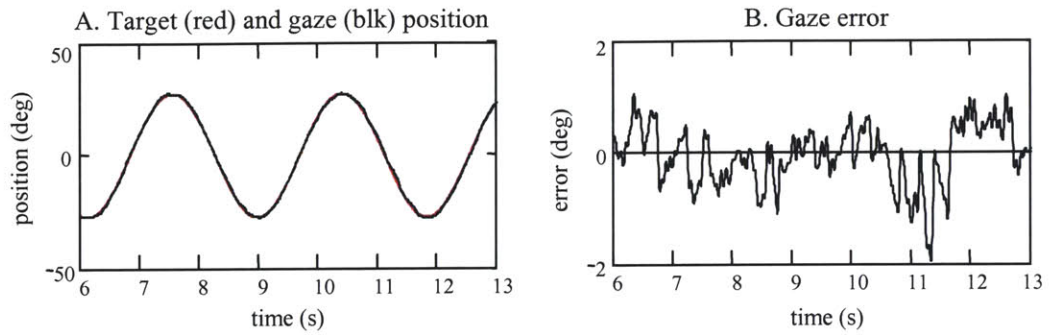
**Figure 4.11:** Probability mass function and power spectral density for a target velocity sequence that was tracked "well".



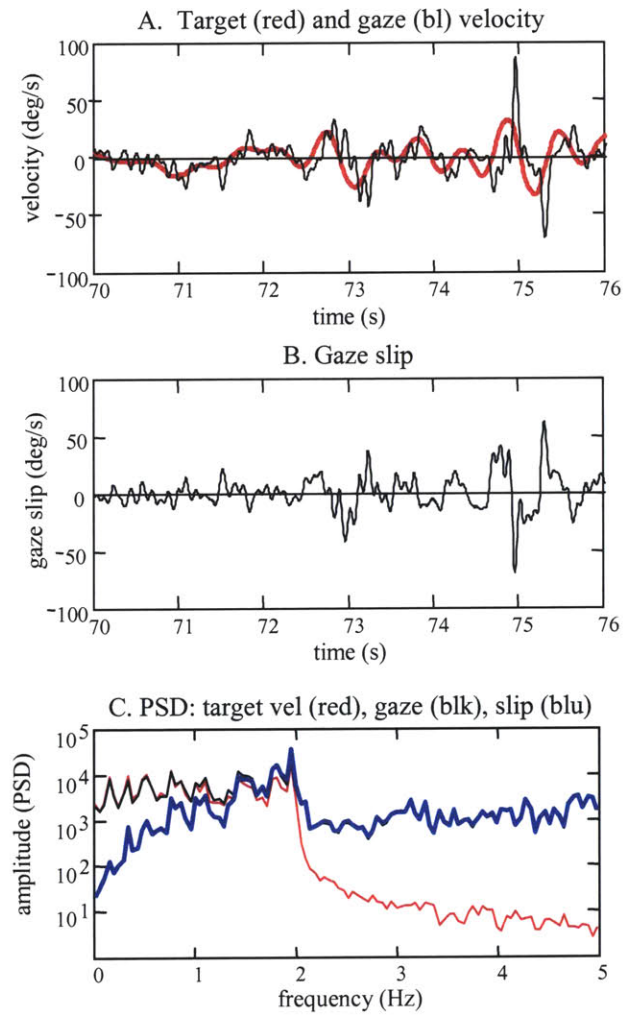
**Figure 4.12:** Gaze tracking of a target with a 0 to 2.0 Hz PSD. The standard deviation of the gaze error was 1.7 deg.



**Figure 4.13:** Gaze tracking of the same target as in Figure 4.12. The gaze was advanced by 100 ms to minimize the mean square error. The standard deviation of the gaze error, after the gaze was advanced, was 1.2 deg.




**Figure 4.14:** Gaze tracking of a sinusoid. The subject tracked the target with little lag, and an error with a standard deviation of 1.8 deg.



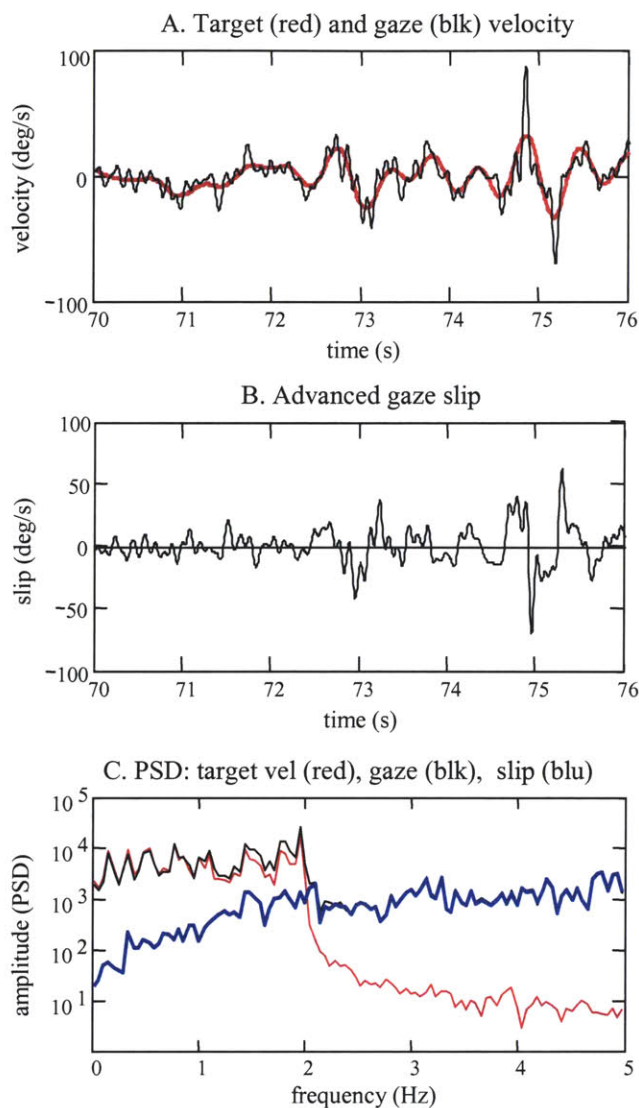
**Figure 4.15:** Gaze and target velocity for a target with a 0.0 to 2.0 Hz PSD.





**Figure 4.16:** Gaze and target velocity for a sinusoidal target.

Although the random target was tracked as if it were unpredictable, it did not induce excessive gaze slip or require frequent, quick phase, saccades. When the gaze velocity was advanced by 100 ms, as was done for the position response, the velocities matched up well (Figure 4.17). The gaze velocity had the same low frequency motion as the target, but had additional high frequency motions. These oscillations reflect an underdamped servo-system that does not have the time to settle to a steady state response. It is likely that these high frequency oscillations were due to the subject's tracking response being in a continual state of "smooth pursuit onset". Oscillations are often observed during the onset of smooth pursuit during the first few hundred milliseconds after the gaze reacts to a target's motion (Leigh and Zee, 1999). The random changes in the target's trajectory prevented the subject's gaze from ever moving away from an onset response. Although the target was not predicted well, only a few saccades are evident in the gaze response. Saccades are visible in the data just before and just after the 75 second mark. These high velocity spikes occurred just as the target's velocity peaked near 40 deg/s, and quickly brought the gaze position back toward the moving target. These movements are identified as saccades because of their high velocity, which is about 50 deg/s greater than the velocity of the moving target. Their shape however does not have the characteristic, sharp edge of a saccadic motion, but this can be attributed to the data having been low-pass filtered.



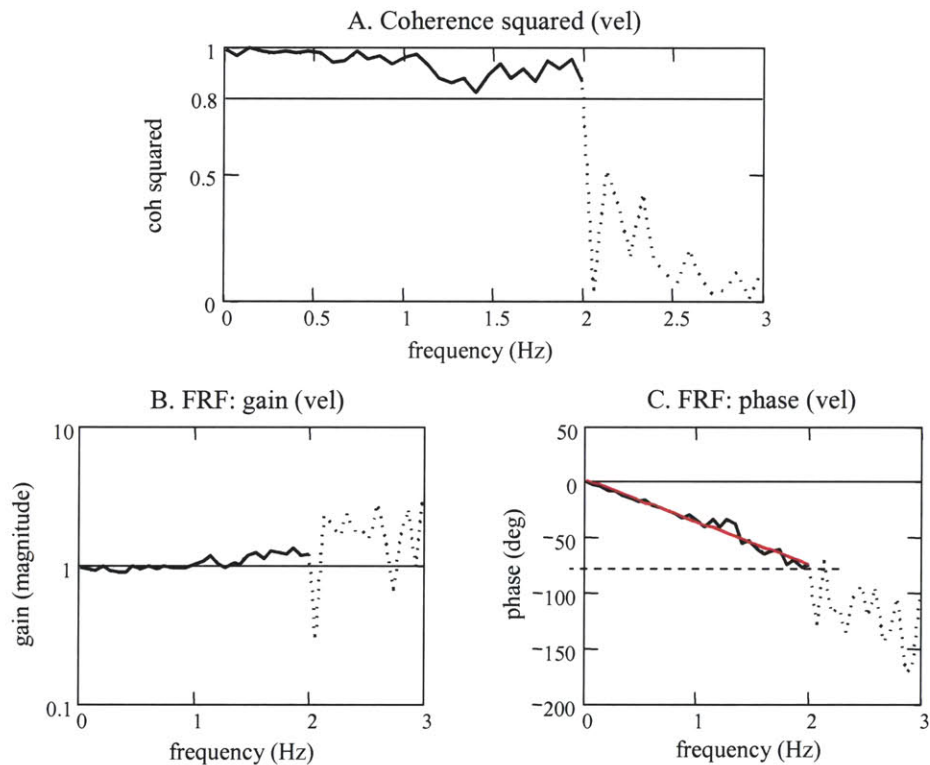
**Figure 4.17:** Advanced gaze velocity. The gaze velocity was advanced by 100 ms for the same gaze and target that was shown in Figure 4.15. Advancing the gaze velocity reduced the standard deviation of the gaze slip from 19 deg/s to 13.8 deg/s.

Although saccades were present in these data, they were infrequent and did not significantly affect how well the linear analysis captured the dynamics of the gaze response. The frequency response and coherence squared functions indicated that there was a strong, linear relationship between the gaze velocity output and the target velocity input over the bandwidth of the target (Figure 4.18). The coherence squared remained above 0.9 from 0.0 to 2.0 Hz, with the exception of a small region between 1.15 and 1.45 Hz, where it was a bit lower. This provided

strong evidence that saccadic motions did not significantly degrade the linear gaze-target relationship over the frequency range for which the visual target was controlled. Beyond 2.0 Hz, the target velocity had very little power, and the coherence dropped below 0.25, which indicated that a low signal to noise ratio, and/or nonlinearities, had degraded the linear input-output relationship. The gain plot shows that the gain was just under unity until 0.8 Hz, and then gradually rose to approximately 1.2 at 2.0 Hz. Gains near unity were to be expected, since the subject tracked the target well. The crossover to above unity near 0.8 Hz was an interesting characteristic observed in the response of many of the subjects. Above 2.0 Hz the FRF gain plot rose quickly. This was not the result of the gaze tracking system responding more strongly to the target, but occurred because the gaze system followed the target with higher frequency motions than existed in the target trajectory. There was little input power at these frequencies, but, as is evident from the gaze velocity plot, there were measurable high frequency gaze motions. The high frequency velocities are greater than can be attributed to noise from the eye velocity calculation, but some of the high frequency responses were from the EOG noise. The dramatic drop in the coherence squared at these frequencies is consistent with the idea that the gain relation is degraded by nonlinearities and a low signal to noise ratio.

The phase plot in Figure 4.18 illustrates that the target was tracked as if it were unpredictable, but also shows what occurs when predictive mechanisms enhance the tracking response. The phase plot shows that the phase lag between the gaze and the target increased steadily with frequency. A phase lag that increases linearly with frequency translates into a constant time lag between the input and the output, where the slope of the line equals the time delay. Between 1.15 and 1.45 Hz the phase was not fitted well by a straight line. At these frequencies, the phase lag was slightly reduced, indicating that the gaze responded more quickly to the target motion. The slope of the line fitted to minimize the sum of squared error is 100 ms. This matches the delay previously found between the gaze and the target when the gaze velocity was advanced so as to be best aligned with the target velocity. The average lag is determined by

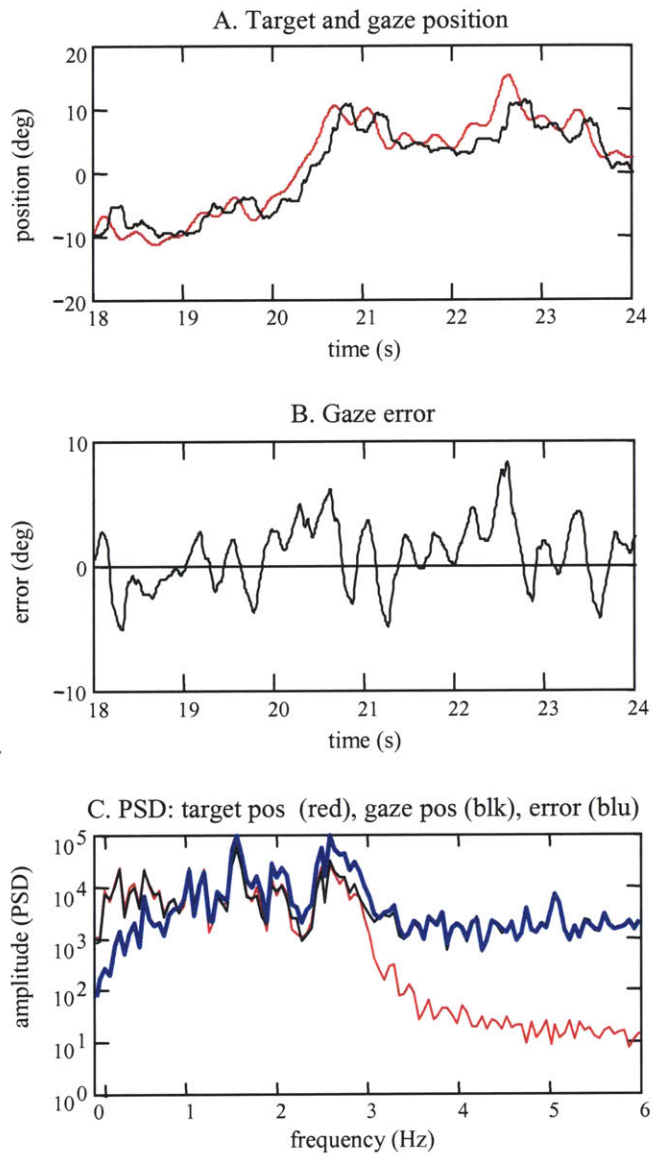
advancing the gaze to best match the target, and the phase response examines this lag as a function of frequency. If the lag is the same at all frequencies then the gaze is being driven by the same mechanism. A delay of 100 ms is presumed to result from visual feedback, and not predictive mechanisms. Changes in the gain plot and the coherence squared plot were also seen at these frequencies. The gain plot shows a gain very close to unity at these frequencies, whereas before and after, it is above unity. The coherence squared also drops to its lowest level in this frequency band. These are all indicators that suggest that in this small frequency band predictive mechanisms did enhance the tracking response.



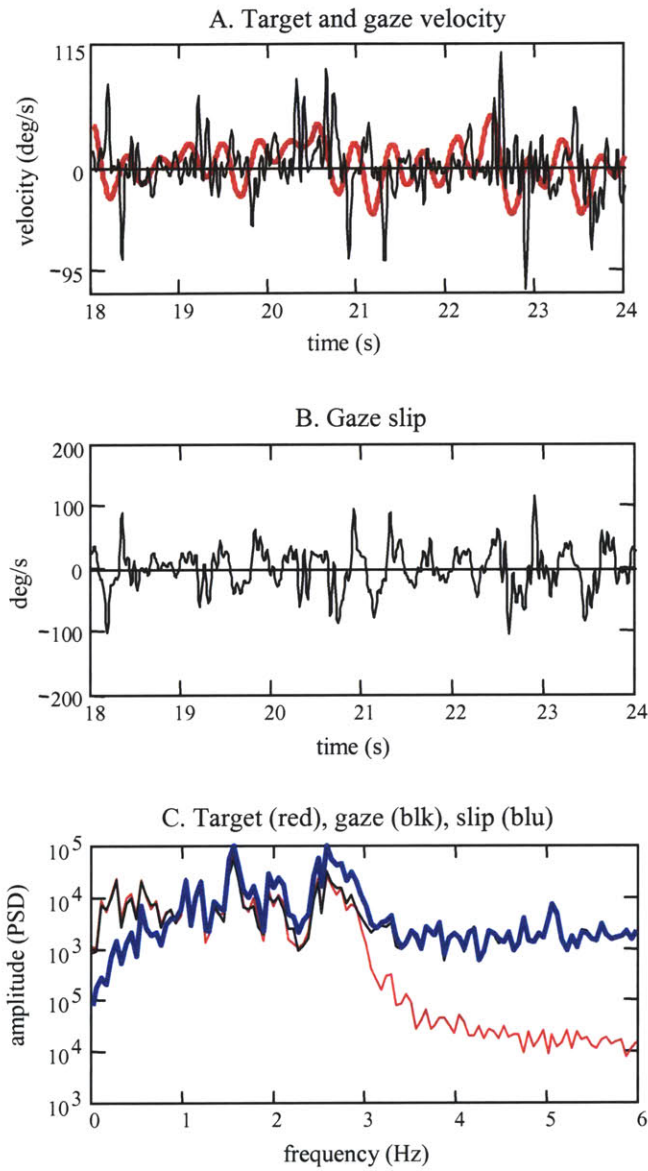
**Figure 4.18:** Frequency response function (plots B and C) and coherence-squared function (plot A) for the gaze tracking response to a target with a 0.0 to 2.0 Hz PSD

### *Target velocity and frequency too high*


Subjects reported that it felt harder to track targets with frequencies that extended beyond 2.0 Hz. Targets with frequencies as high as 10 Hz were presented, but these trajectories changed much too rapidly for subjects to react to the movements. A maximum frequency of 3.0 to 3.5 Hz was about the limit at which subjects thought that they could follow some of the higher frequency motions. Subjects reported that they would often ignore the higher frequency movements, which, unless they had high velocity, did not move far. The trajectories were generally too fast and too difficult to focus upon. However, the data do not indicate that the higher frequency targets were tracked particularly poorly, but it was evident that these targets were tracked with greater gaze slip and more frequent saccades. Figures 4.19 and 4.20 show the response for one subject to a target with a frequency spectrum from 0.0 to 3.0 Hz. The phase lag shows that the target was tracked with a consistent time lag, in this case, 130 ms. Visual feedback did not advance the tracking response at any frequency. The gaze position and velocity were aligned best with the target motion when they were advanced by 140 ms. The gaze position plot is not as smooth as the target's, but it generally follows the target's motion. To accomplish this, the subject used higher velocities and generally more saccades than when the lower frequency trajectories were tracked. The effect of the saccades, which are a nonlinear response, can be seen in the coherence-squared and FRF functions in Figure 4.21. The coherence-squared, which relates the target velocity input to the gaze velocity output, was high, but did not remain as high as when the target's frequency spectrum was lower. For the lower frequency targets, it was common for the coherence-squared function to be above 0.9 across the full spectrum over which the target was controlled. For the higher frequency trajectories, the coherence squared was high at the lower frequencies, but had greater variance and decreased as the frequency increased. The high velocity of the gaze slip, and the increased number of saccades associated with the higher frequency targets, degraded the linear analysis. Although the saccades affected the linear input-output relationship, they functionally allowed the gaze to remain near the target, and so the gain stayed close to unity.



**Figure 4.19:** Gaze tracking of a target with 0.0 to 3.0 Hz PSD. The standard deviation of the gaze error was 3.3 deg.



**Figure 4.20:** Gaze and target velocity during tracking of a target with 0.0 to 3.0 Hz PSD



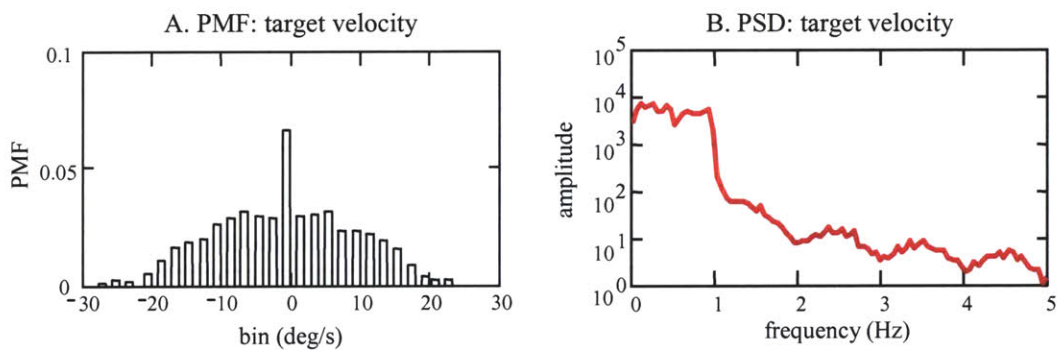
**Figure 4.21:** Frequency response function (gain in plot B and phase in plot C) and coherence-squared function (plot A) for the gaze tracking response to a target with a 0.0 to 3.0 Hz PSD.

#### *Target frequency too low*


When the target spectra were limited to below 1 Hz, and the velocities kept below 20 deg/s, the targets were tracked with little position error and little gaze slip, and the number of saccades was lower than when the targets had higher frequencies. It was also evident that predictive mechanisms enhanced the gaze response. As can be seen in Figure 4.23, which is an example of one subject tracking a target with a 0 to 1 Hz spectrum, the gaze did not lag the target by much. If the gaze had been driven by visual feedback, the time lag would have been between 75 and 150 ms, as was seen when the higher frequency targets were tracked. To minimize the mean square difference the gaze only had to be advanced by 30 ms. Although the gaze lagged the target less than for the higher frequency targets, it still did not move as smoothly as the target, and




the gaze velocity oscillated around the target's velocity, as had been seen during tracking of the higher frequency targets. The coherence-squared function was nearly 1 throughout the frequency range of the target, which indicated that neither nonlinearities nor noise had degraded the linear analysis. This would seem ideal for the VOR analysis if it were not for the phase plot which had an inconsistent lag across the frequencies as shown in Figure 4.25. Between 0 and 0.15, the phase was nearly zero. At these frequencies, where both the gain and coherence were nearly 1, the predictive mechanisms had anticipated the target's movements very well. The phase increased steadily above 0.15 Hz, but the lag at 1 Hz is only 27 ms, which is shorter than the time required for visual system processing. When the target did not exceed 1 Hz, the gaze was enhanced by predictive mechanisms and these targets were therefore inappropriate for the VOR protocol.



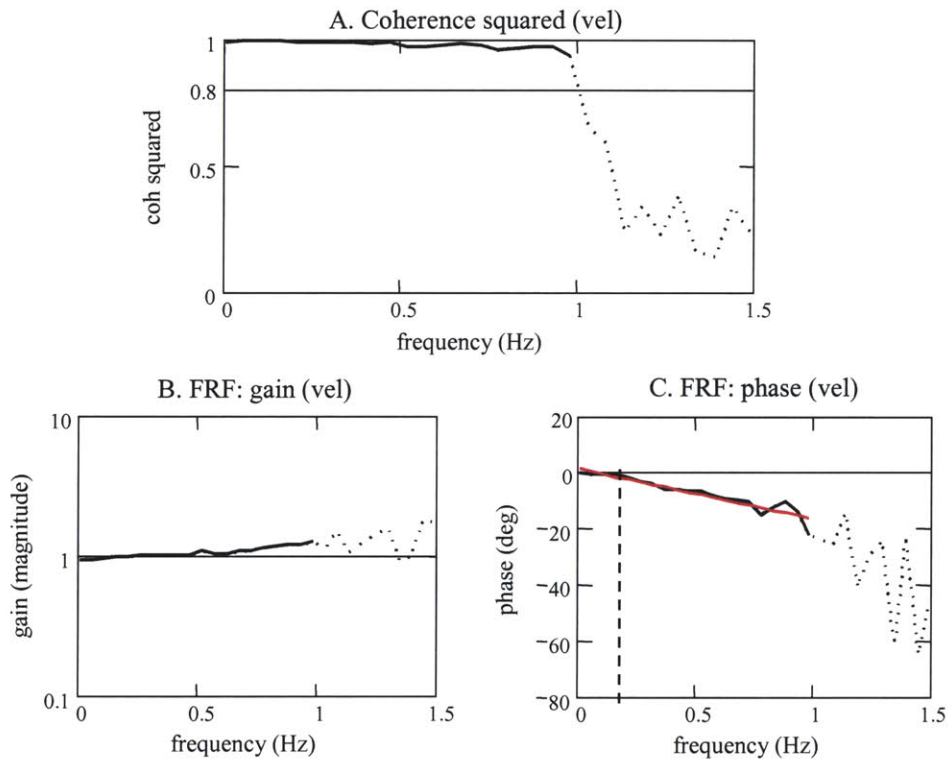
**Figure 4.22:** Representative probability mass function (plot A) and power spectral density (plot B) of a target velocity sequence that was tracked as if it were predictable.



**Figure 4.23:** Gaze tracking a target with a 0.0 to 1.0 Hz PSD. The standard deviation of the gaze error was 2.1 deg.



**Figure 4.24:** Gaze and target velocity for a target with a 0.0 to 1.0 Hz PSD.



**Figure 4.25:** Frequency response function (gain in plot B and phase in plot C) and coherence-squared function (plot A) for the gaze tracking (velocity) response to a target with a 0.0 to 1.0 Hz PSD.

## **4.4 Tracking unpredictable targets with torque perturbations applied to the head**

The torque and target input sequences had to be validated for use in the VOR experiments by applying them simultaneously to the subject. The experiments that were conducted to evaluate the head-neck system's response to torque perturbations, and to identify the characteristics of a visual target that resulted in it being tracked as if it were unpredictable, had successfully identified input sequences that were appropriate stimuli for the visual and vestibular systems, but none of the experiments had been conducted with the same protocol that would be used to test the VOR. It seemed possible that the ability of a subject to track an unpredictable target might be altered when torque perturbations were applied to the head. The torque perturbation sequences generated considerable disturbances and covered a range of potential disturbances that a subject might experience during natural activities into a short time span (60 to 90 s). Although a subject with a healthy vestibular system should be able to generate eye motions that compensated for the head disturbances, the ability of the subject to track the target would be degraded if the perturbations saturated the vestibular system or if the perturbations prevented the subjects from using the head as they needed to track the target. A related concern was whether the PRB torque inputs would still stimulate the desired range of head disturbances when the head was being actively driven to follow the unpredictable target. The torque inputs had been shown to be adequate when the subjects tracked a predictable sinusoid, but the mechanisms involved in controlling the stability of the head could change with an unpredictable target. If the performance of the tracking and the head-neck systems changed when the torque and target inputs were applied to the subject simultaneously during the VOR experiments, then the input sequences would have to be refined to accommodate the changes.

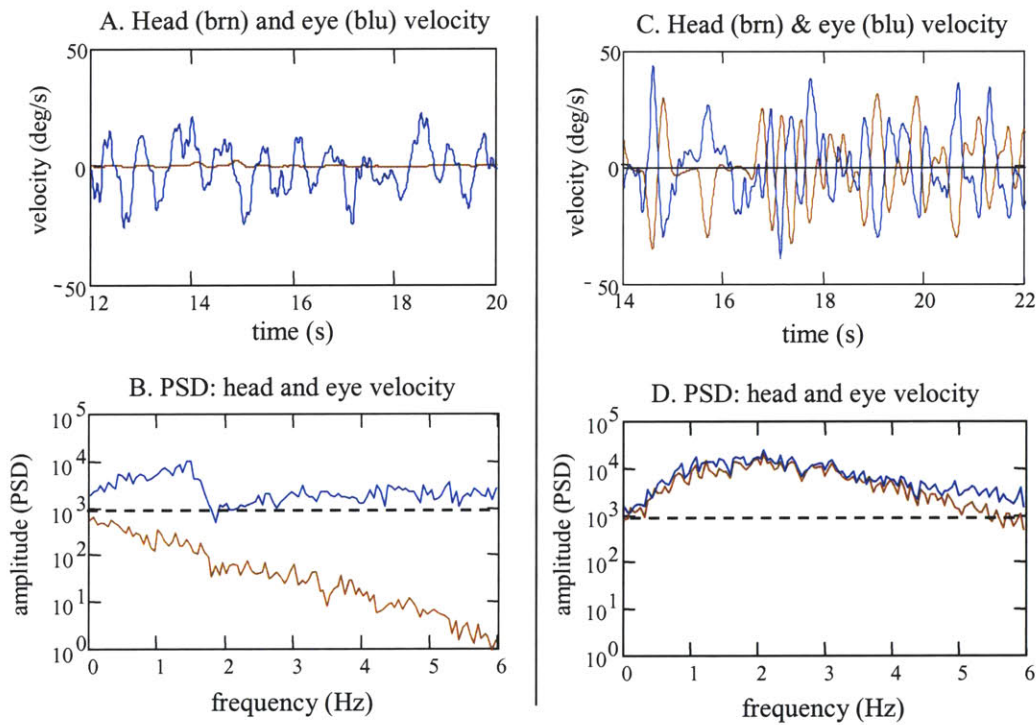
The response of the gaze tracking system, and of the associated head and eye motions, were therefore evaluated as subjects tracked unpredictable targets while the head was perturbed

by PRB torque perturbations. The experimental procedure was similar to that used when the subjects had tracked the unpredictable targets without the applied head disturbances. The subject was seated 1600 mm in front of the screen onto which the target was projected, and the head perturber was lowered and attached to the subject's head. The subjects were required to track a randomly moving target, as best they could, with natural head and eye motions. The target trajectories were selected from the trajectories that had previously been used to create visual-feedback driven, gaze responses. PRB torque perturbations, with amplitudes of approximately 0.2 N.m, were applied as the subject tracked the target. Head velocity, eye position, target position, and the torque applied to the head-perturber's helmet were sampled at 100 Hz.

It was found that the torque and target sequences that were good stimuli when they were applied separately to the visual and vestibular systems were still appropriate when applied simultaneously. The torque perturbations disturbed the head sufficiently, but the VOR was able to compensate for these disturbances, and the tracking response was not noticeably affected. The subjects tracked the targets as well, and in the same manner, as when no perturbations were applied.

The PRB torque perturbation sequences created head disturbances with an appropriate magnitude, and desired frequency range, as the subject controlled the head's motions to track the target. The disturbance to the head position and velocity was obvious, and it induced easily measured VOR eye motions. Plots of the head and eye velocities from two representative trials, one with perturbations and one without, are shown in Figure 4.26. When no perturbation was applied to the head, the subject tracked the target without moving the head rapidly. Throughout the 1.8 Hz spectrum for which the target was controlled, the PSD of the head velocity was much lower than that for the eyes, and decreased as the frequency of the target increased. The head appeared to move slowly to direct the eyes towards the target, and the eyes did most of the work to follow the target's motions. The velocity and PSD plots for the head and eyes were very different when the perturbations were applied to the head. It is obvious that the torque

perturbations created strong head disturbances. The head velocity rapidly changed direction and was obviously driven by the PRB torque input. It is also evident that the VOR created eye motions to compensate for the unwanted head disturbances. The plot of the eye velocity appears close to being equal and opposite to the head velocity, and from 0.0 through 4.0 Hz the PSD for the eye and the head velocities are very similar. The PSD of the head velocity when the head was perturbed is much higher than when the head was not perturbed. Beyond about 4 Hz, the PSD for the eyes and head begin to deviate.

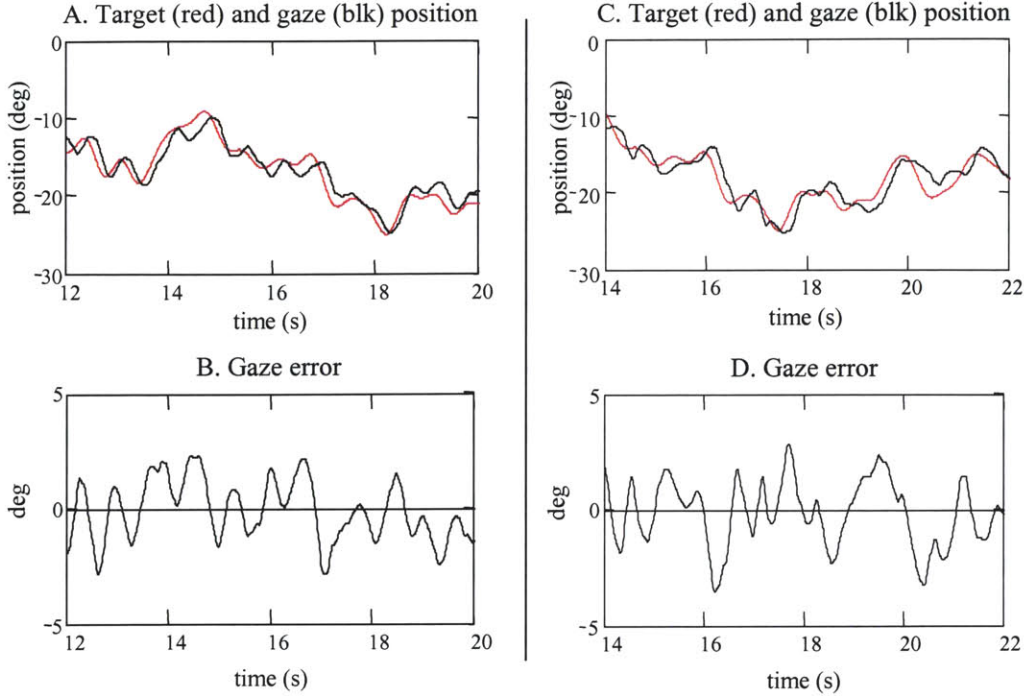


**Figure 4.26:** Head (brown) and eye (blue) responses during gaze tracking. No torque perturbations were applied to the head in plots A and B. Perturbations were applied to the head in plots C and D.

There were no noticeable differences in the subjects' overall tracking performance when the torque perturbations were applied to the head. The torque perturbations significantly disturbed the head position and created the desired span of disturbance motions, but did not interfere with the subjects' ability to track the target. In fact, some subjects reported that they found it easier to

track the targets while their head was perturbed as compared to when no perturbations were applied. One subject claimed that it felt as though a different system had taken over her tracking response when the torque sequences were shaking the head. When no perturbations were applied, she felt that she had to make a conscious effort to focus on and track the target. When the perturbations were applied, she no longer felt as if she had to concentrate to follow the target, and could just let her gaze system react to where the target was moving. It has been reported that smooth pursuit tracking is sometimes improved when subjects are distracted by mental arithmetic (Wells and Barnes, 1999; Van Gelder and Lebedev, 1995). This may offer some explanation as to why target tracking felt easier, rather than more difficult, when the perturbations were applied to the head. Humans are not good at voluntarily controlling smooth pursuit eye motions. Very few people can create smooth pursuit eye motions unless there is a smoothly moving visual target to track (Leigh and Zee, 1999). It is possible that the perturbation torque acted as a distraction to the subjects, and stopped the subjects from trying to actively enhance their smooth pursuit response.

The head and eye responses were very different when the torque perturbations were applied to the subjects' heads as compared to when there were no perturbations, but the net result of how the eye and head motions summed to produce a gaze response, was indistinguishable. The vestibular systems of these healthy subjects were obviously able to stabilize the subjects' gaze response against the disturbances caused by the torque perturbations. There were certainly differences between the responses when results from the two sets of experiments were compared, but the differences were not such that it was possible to determine reliably whether the tracking trial had occurred with, or without, applied torque perturbations. The gaze tracking responses that were associated with the head and eye velocity plots of Figure 4.26 are shown in Figure 4.27. Both gaze responses lag the target's motion, and the error plots are of about the same magnitude. The standard deviation of the error was 2.1 deg when the perturbations were applied to the head, and 2.0 deg when no perturbations were applied.




**Figure 4.27:** Gaze (black) and target (red) position during gaze tracking. No torque perturbations were applied to the head in plots A and B. Perturbations were applied to the head in plots C and D.

The results for the gaze velocities were also very similar. The gaze velocity, gaze slip, and PSD plots in Figure 4.28 are from a trial when torque perturbations were applied, and are similar to the plots shown in Figure 4.15, which shows the gaze velocity response to a 0.0 to 2.0 Hz target when no perturbations were applied. The gaze slip did not change significantly when the head was perturbed, with a standard deviation of 15.2 deg/s with perturbations, and 17.2 deg/s without the perturbations. The PSD of the gaze slip in both cases increased steadily with the target's frequency, and the magnitude of the PSD at each frequency was nearly identical. The gain and phase plots of the frequency response function, and the coherence-squared measures, are also remarkably similar. The FRF and coherence-squared plots for a trial with torque perturbations are shown in Figure 4.29. Similar to the situation when there were no perturbations applied to the head (Figure 4.18), the gain is near unity until around 0.8 Hz, and then increases to

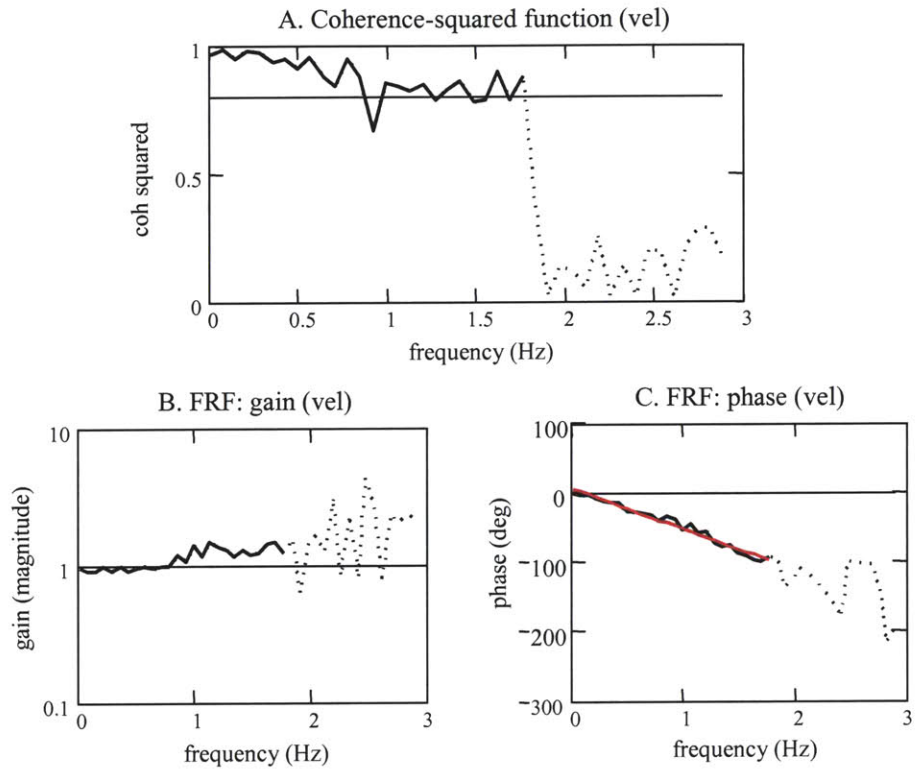
just above unity across the 1.8 Hz bandwidth of the target. As would be expected, the addition of perturbations did not improve the subject's ability to predict the target motions. Targets that were tracked as if they were unpredictable when the head was not perturbed remained unpredictable when the perturbations were applied. The phase lag shown in Figure 4.29 increased nearly linearly with frequency, just as in Figure 4.18. For these examples, the time lag between the target and the gaze response was 140 ms when the perturbations were applied, and 130 ms when they were not. The coherence-squared measure was also not affected much, if at all, by the addition of head perturbations. The coherence-squared measure in Figure 4.29 is not quite as high as that shown in Figure 4.18, but it still is above 0.8 across most of the 1.8 Hz spectrum for which the target was controlled, and actually above 0.9 for frequencies below 0.8 Hz. The disturbances applied to the head did not change the linearity of the response between the target and the gaze, and did not add a considerable amount of noise to the measurements. In some cases the torque perturbations may actually have improved the experimental measurements since more pronounced head velocity, eye position, and torque measurements occurred when the head was being perturbed.

The torque perturbations significantly disturbed the head and eye motions, but did not change how well the subjects tracked the targets, nor did they change the velocity content of the gaze response. Saccades did not appear to be more frequent, nor did the perturbations seem to alter the signal that drove the gaze response. It was evident that when the VOR was functioning properly, it was able to compensate for the level of disturbance created by the torque perturbations and maintain the subject's gaze. At least among healthy subjects, the PRBS torque perturbations and random target sequences that had been found to stimulate individually the visual and vestibular pathways worked very well when applied simultaneously under the VOR testing protocol.





**Figure 4.28:** Gaze tracking response (black) when perturbations were applied to the head. The target had a 0.0 to 1.8 Hz PSD



**Figure 4.29:** Frequency response function (gain in plot B and phase in plot C) and coherence-squared function (plot A) for the gaze tracking (velocity) response. A PRB torque perturbation sequence was applied to the head. The target had a 0.0 to 1.8 Hz PSD.

## Chapter 5

### Generating uncorrelated torque and target inputs

- 5.0 Chapter summary
- 5.1 The modified Hunter-Kearney procedure
- 5.2 Steps of the modified Hunter-Kearney procedure
- 5.3 Defining bounds for zero cross-correlation
- 5.4 Swap conditions: number of elements, evaluation criteria, sequence PMF
- 5.5 Conclusions

#### 5.0 Chapter summary

It was demonstrated in Chapter 4 that a PRBS torque perturbation and a random target velocity trajectory would stimulate the physiological responses from which the performance of the VOR could be evaluated. The analysis of the VOR required that the torque and target inputs be uncorrelated. This chapter discusses the techniques that were developed to generate torque and target sequences that were uncorrelated and still retained the properties required to stimulate the visual and vestibular systems. The core procedure that was used to generate the signals was a novel stochastic interchange method that allowed two sequences to be generated with jointly specified probability density, auto-correlation, and cross-correlation functions.

## 5.1 The modified Hunter-Kearney procedure

Uncorrelated torque and target sequences were generated using a novel process developed for this research. The technique allows two discrete time sequences to be created with individual spectral and temporal characteristics, and with a specified cross-correlation. One of the sequences can be chosen to be of any signal type, and there are no restrictions as to how it is created. The second sequence is generated with a stochastic swap technique and thus, by the nature of the stochastic swap process, is a random sequence. To fulfill the requirements of the VOR testing protocol, the technique was used first to create a PRBS torque sequence and then to manipulate a target velocity sequence so that it had the characteristics required to stimulate the gaze system but was uncorrelated with the torque sequence.

The core of the technique is based on the Hunter-Kearney stochastic swap algorithm, which was developed to create a single random sequence with jointly specified auto-correlation and probability density functions (Hunter and Kearney, 1983). The Hunter-Kearney algorithm creates a discrete-time sequence with an amplitude distribution that approximates a desired probability density function (PDF), and then orders the elements with a recursive, double stochastic, interchange of elements, so that the sequence's auto-correlation function (ACF) converges to a desired ACF. Details of the algorithm are described more fully in Chapter 4.

It was recognized that the Hunter-Kearney algorithm could be modified so that instead of arranging the elements of a sequence to have a desired auto-correlation function, the elements could be arranged to have a desired cross-correlation with a different, previously defined sequence. After all, a sequence's ACF is simply the CCF of the sequence with itself. In addition, the elements of the sequence could be arranged so that the sequence's ACF, and its CCF with another sequence, simultaneously converged toward desired values. With this modified procedure, a target velocity sequence could be generated with the frequency and amplitude characteristics that were appropriate for stimulating the gaze tracking system, and also be made

so that it was uncorrelated with a PRBS torque sequence that had the appropriate amplitude and temporal characteristics for stimulating the vestibular system.

The development of the modified Hunter-Kearney technique was an important step towards being able to execute the VOR testing protocol. It provided an efficient and systematic means of generating the experimental torque and target sequences so that all stimulus criteria were met. It was, of course, crucial to identify the types and characteristics of signals that stimulated appropriate responses from the head-neck and gaze tracking systems, but it was not sufficient. Although an experimental trial could generate data that represented a full spectrum of natural gaze tracking and corrective VOR responses, and was thus informative enough for a system identification (Ljung, 1999), when the inputs are correlated it is not possible to distinguish the eye response that is due to the visual input from that resulting from the vestibular input, and so the performance of the VOR cannot be extracted from the data. Without the modified Hunter-Kearney technique, torque and target sequences would have had to have been created independently, and then mixed and matched until pairs were found with a sufficiently low cross-correlation; a very time consuming and tedious process.

The modified Hunter-Kearney procedure is a powerful technique for creating sequences for system identification studies. The method has few restrictions, and is very general in that it can manipulate many different types of signals so that they have arbitrarily defined auto-correlation and cross-correlation functions. The technique is not restricted to ordering a sequence to be white, nor to making the two sequences uncorrelated. It can be used to generate sequences with arbitrary spectral content, and with shared power at defined frequencies (cross-power spectral density). There are several other methods that can be used to create signals with a specific ACF, or PDF, or that are uncorrelated, but these methods tend to be more limited in the types of signals they can use and in the ACF and CCF that they can generate. For example, two PRBS can be made uncorrelated through interweaving (Davies, 1970), but the technique cannot be extended to generate anything other than binary sequences with a band-limited white spectrum

(impulse ACF) and a CCF that is uncorrelated. Sums-of-sines can be made to be unpredictable and to have arbitrarily defined spectral content (ACF) and cross spectral densities by carefully choosing the amplitude and frequency of each sinusoid, but the method is less general than the modified Hunter-Kearney procedure because it allows only sinusoids to be manipulated. The modified Hunter-Kearney technique is an important development for generating inputs for multi-input systems. Rather than having to use the same type of signal along each input path, the technique allows the most appropriate type signal for each input path to be generated with specified ACF and CCF properties. This is particularly important for physiological studies, where the sensory or motor systems may be able to tolerate only certain types of signals. In the past, the most common means used to create two different signal types with cross-spectral and cross-correlation properties was to find two signals, by chance, that met the requirements.

The main limitations of the modified Hunter-Kearney technique are that it requires a significant amount of computing time and power to arrange the elements of the sequences, and how well the final ACF and CCF meet their desired values depends on the characteristics of the initial sequence of elements. The modified Hunter-Kearney technique is not a closed form solution that creates sequences with exact ACF and CCF values. Owing to its stochastic nature, its success at creating sequences with acceptable characteristics is somewhat dependent on how good the user is at defining favorable initial conditions.

## **5.2 Steps of the modified Hunter-Kearney procedure**

The torque and target sequences used in the VOR experiments were created with an application coded using Visual Studio 6.0. The code is shown in Appendix C. The process is first described briefly, and then each step is discussed in more detail.

There were five fundamental steps to generating the torque and target input pairs with the modified Hunter-Kearney process. First the PRBS torque sequence was created. It was generated

to have the amplitude and temporal characteristics that were appropriate for perturbing the head-neck system, but without regard for its relation to the target sequence. Second, a sequence was generated that was the same length as the torque sequence, but with an amplitude distribution that represented the velocities the visual target would travel during the experiment. Third, this velocity sequence was arranged using the modified Hunter-Kearney stochastic swap algorithm until it has the desired auto-correlation and cross-correlation function. The ACF was defined by taking the inverse Fourier transform of a power spectral density function, which was easier and more intuitive to define directly than the ACF. The desired CCF was set to zero. Fourth, the target sequence was low-pass filtered to remove any unwanted high frequency content. Finally the velocity sequence was integrated, and became the target position sequence used to move the target on the screen.


#### 5.2.1 Generate the PRBS torque sequence

The first step of the process was to create the PRBS torque perturbation sequence. In terms of a generalized “modified Hunter-Kearney process”, in which two sequences are designed to have particular ACF and CCF properties, it is not necessary that this first sequence be a PRBS. It can be any signal type that is desired, such as PRBS, sinusoids, random, or deterministic signals. The only requirement is that the sequence must be designed to meet the stimulus requirements for the system on which it acts. It is created without regard to its cross-correlation with the random sequence that will be manipulated with the stochastic swap algorithm. The characteristics that are given to this first sequence will, however, affect how well the modified Hunter-Kearney algorithm can order the second sequence to achieve the desired spectral and cross-correlation properties.

The example shown in Figure 5.1 is representative of the torque command sequences used during the VOR experimental trials. The torque PRBS had an amplitude of 0.2 V, and was executed at a rate of 10 elements per second. The sequence represented the voltage reference that

was applied to the motor of the head-perturber. The amplitude was adjusted so that the motor's torque output was near 0.2 N.m and comfortable for the subject. Most of the subjects found torque amplitudes between 0.2 and 0.3 N.m strong enough to disturb the head, but low enough so that they were still able to position the head where they wanted. A PRBS execution rate of between 10 and 15 elements per second was determined to be appropriate for the protocol. Slower execution rates did not produce power across a sufficient frequency band to stimulate all the desired motions of the head, and faster execution rates were not always followed faithfully by the motor. When the perturber did not follow the torque command faithfully, the torque that was delivered to the head did not have the same spectrum as the command signal, and so the cross-correlation relation that had been carefully designed into the sequences was degraded. The sequence in Figure 5.1 is 30 s long, which is shorter than most trial lengths which lasted from 60 to 90s. It was selected as the example sequence because its short length made it more convenient to show the intermediate steps associated with generating the target velocity sequence. The PRBS was created from a nine-element seed. The number of seed elements was chosen so that the PRBS did not complete a full cycle during the experimental trial. The PRBS repeats itself after  $2^n - 1$  elements and becomes periodic. Unless the protocol is designed to conduct a "periodic" system identification, where periodic based ACF and PSD estimators are used to account for the periodicity of the signal (Godfrey, 1993), the ACF will show that the sequence became correlated with itself at a lag value equal to  $2^n - 1$ , and the PSD will no longer look as if the signal approximates a band-limited white sequence. The torque sequence was "filled in" to have 200 elements per second (Figure 5.1b), so that it could be output by the experimental control program at the same rate as the target sequence. This update rate was necessary for the target to appear to move smoothly across the screen.





**Figure 5.1:** Pseudo-random binary torque sequence. The sequence was created so that potential changes in the PRBS were executed at 10 elements per second (c/s) (plot A). The sequence was then “filled-in” so that the sequence had 200 elements per second (plot B).

### 5.2.2 Generate target velocity sequence with a desired amplitude distribution

The target velocity sequence was created using the modified Hunter-Kearney algorithm and was designed to have the spectral and amplitude characteristics required to stimulate the gaze tracking system, and to be uncorrelated with the PRBS torque sequence. The modified Hunter-Kearney procedure requires that this second sequence be a random sequence, whereas the first sequence can be of any signal type.

The first step in generating the target velocity sequence was to form a discrete-time sequence that was the same length as the torque sequence, with a desired amplitude probability distribution. The amplitude distribution of the sequence represented the velocities that the target would travel during the experiment. In the method developed by Hunter and Kearney, a sequence is made to have a probability mass function that approximates a desired PDF by filtering a sequence of numbers, uniformly distributed between 0 and 1, with a static nonlinearity. The static nonlinearity is created from the reversion of the desired cumulative distribution function (CDF), which is the integral of the probability density function. The cumulative distribution function, evaluated at  $x_0$ , represents the probability that an outcome of an experiment (random variable) will be less than the value  $x_0$ . The CDF always falls between 0 and 1, where the probability that

an outcome is less than  $-\infty$  is 0 and the probability that the outcome is less than  $\infty$  is 1. The reversion of the function  $y = f(x)$ , solves the independent variable in terms of the dependent variable,  $x = g(y)$ . For the CDF, the independent variable is the value that the random variable will be equal to or less than, and the dependent variable is the probability that the experimental outcome will be lower than the value of the independent variable. The reversion of the CDF therefore uses the probability value as the argument, and the value of the random variable as the dependent variable.


A uniformly distributed, discrete time, sequence was first created. It was required to have the same number of elements as the torque sequence, and to be equally spaced with values between, but not equal to, 0 and 1 (Figure 5.2a). The range (0 to 1) corresponds to the probability limits of the cumulative distribution function. The uniform sequence was then given the desired amplitude distribution, which was approximately Gaussian in this example, by transforming the sequence with the reverted cumulative distribution function. Each element of the uniform sequence, a value between 0 and 1, was used as an argument into a function created from the reverted CDF. The outputs of the function formed a sequence with a probability mass function that approximated a Gaussian PDF.

The use of the reverted CDF to transform the sequence avoided having to rely upon the central limit theorem and a random number generator to create a sequence that approximated a Gaussian distribution. The Gaussian CDF can be expressed explicitly as the integral of the closed form Gaussian PDF (Eq. 5.1), but is often expressed in terms of the complementary error function (Eq. 5.2).


$$CDF_{Gauss}(x, \mu, \sigma) = \int_{-\infty}^x \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right) dx \quad \text{Eq. 5.1}$$

$$CDF_{erfc}(x, \mu, \sigma) = 1 - \frac{1}{2} \operatorname{erfc}\left(\frac{x-\mu}{\sqrt{2\sigma^2}}\right) \quad \text{Eq. 5.2}$$

The CDF for a Gaussian random variable with zero mean and a standard deviation of 2.0 is shown in Figure 5.3.




**Figure 5.2:** A uniformly distributed sequence (plot A), its probability mass function (plot B), and its cumulative distribution function (plot C).



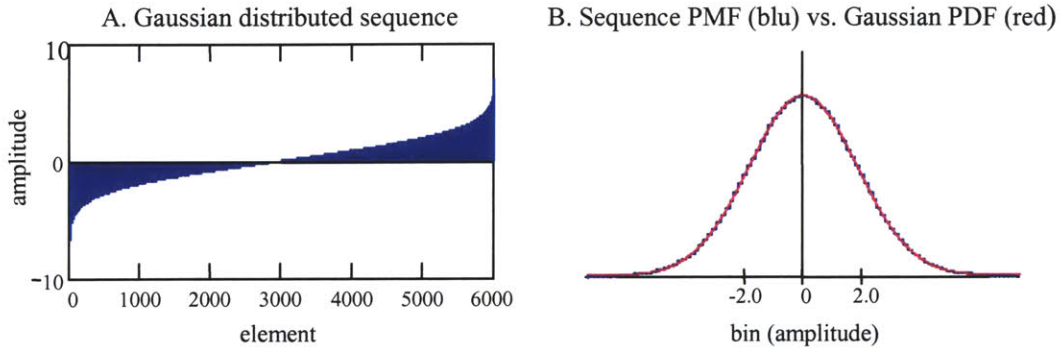
**Figure 5.3:** Cumulative distribution function (CDF) for a Gaussian sequence with a mean of 0 and a standard deviation of 2.

There is no elementary solution for the reversion of the Gaussian CDF, which is of the form  $f(x) = \int_a^b \exp(-x^2) dx$ , and so the definite integral must be approximated (Thomas and Finney, 1984). Fortunately, the reverted Gaussian CDF is commonly tabulated and explicitly defined in mathematical software packages. In the Mathcad mathematical package (Mathsoft Inc., Cambridge, MA) the reverted Gaussian CDF is expressed by the function QNORM( $p, \mu, \sigma$ ) (Figure 5.4).



**Figure 5.4:** The reversion of the Gaussian cumulative distribution function.

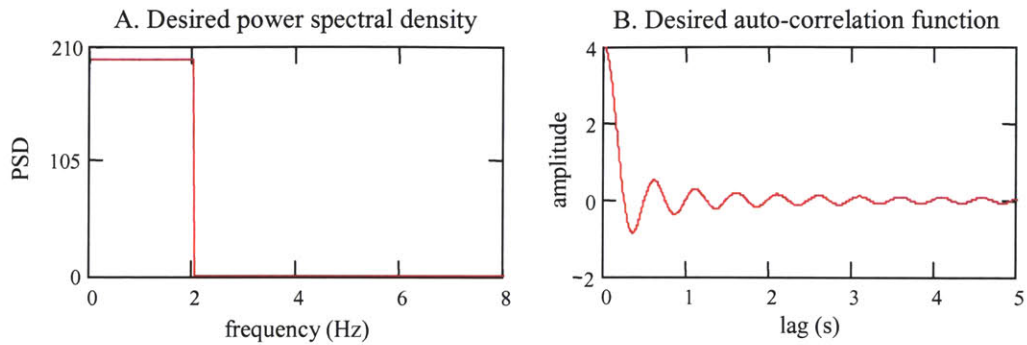
The Gaussian distributed sequence was formed from the output of the QNORM function, where the input arguments were the elements of the original, uniformly distributed, sequence. The end result was a sequentially ordered sequence, with a PMF that approximates a Gaussian PDF (Figure 5.5).



**Figure 5.5:** The outcome of transforming the uniform sequence with a static filter made from the reversion of a Gaussian cumulative distribution function (plot A). The sequence's probability mass function and the Gaussian PDF is shown in plot B.

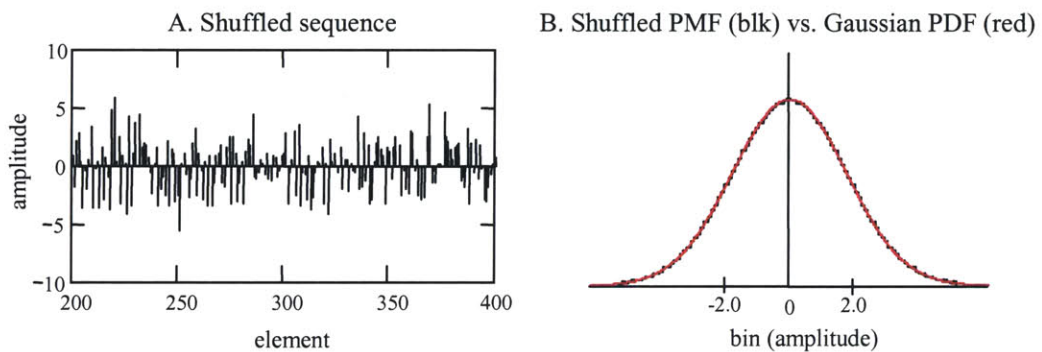
### 5.2.3 Arranging the velocity sequence for ACF and CCF with modified Hunter-Kearney swap

After the target velocity sequence had been transformed to have the desired amplitude distribution, the elements were ordered to meet the ACF and CCF criteria using the modified Hunter-Kearney stochastic swap algorithm. The desired ACF was not directly specified, but was calculated by taking the discrete, inverse Fourier transform of a power spectral density function. It is more natural to define the desired spectral content for the target sequence through the power spectral density function than through the ACF. A 0.0 to 2.0 Hz PSD is shown in Figure 5.6, along with its corresponding ACF. The PSD represented the same total power as was in the target velocity sequence. This ensured that the ACF, calculated from the PSD, could potentially be achieved by a configuration of the sequence's elements.

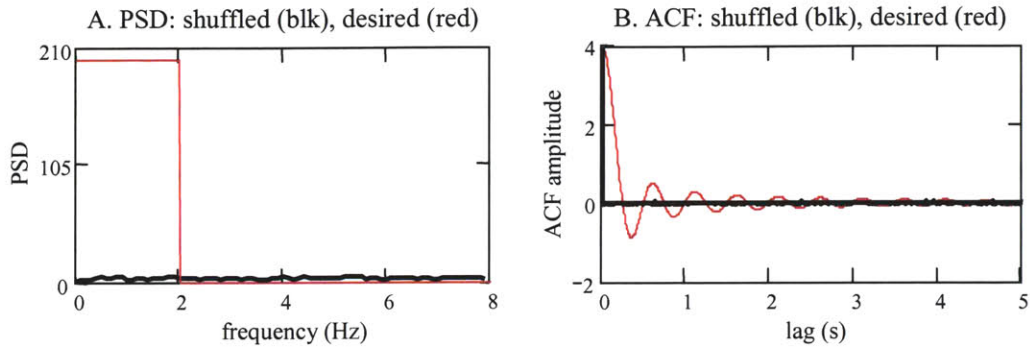


**Figure 5.6:** The desired power spectral density (plot A) and the corresponding auto-correlation function (plot B).

The sequentially ordered sequence was randomly shuffled and “whitened” before being manipulated to attain the desired ACF and CCF (Figure 5.7). Hunter and Kearney (1983) found that this is required to ensure that the final sequence is stationary. As with the original Hunter-Kearney technique, rearranging the elements did not affect the velocity distribution because only the order, not the value, of the elements was changed. A comparison of the desired PSD and ACF, and the PSD and ACF for the shuffled sequence is shown in Figure 5.8. At this point in its development the shuffled velocity sequence is a good approximation of a band-limited white sequence. Its PSD is uniform, and its ACF is close to an impulse.



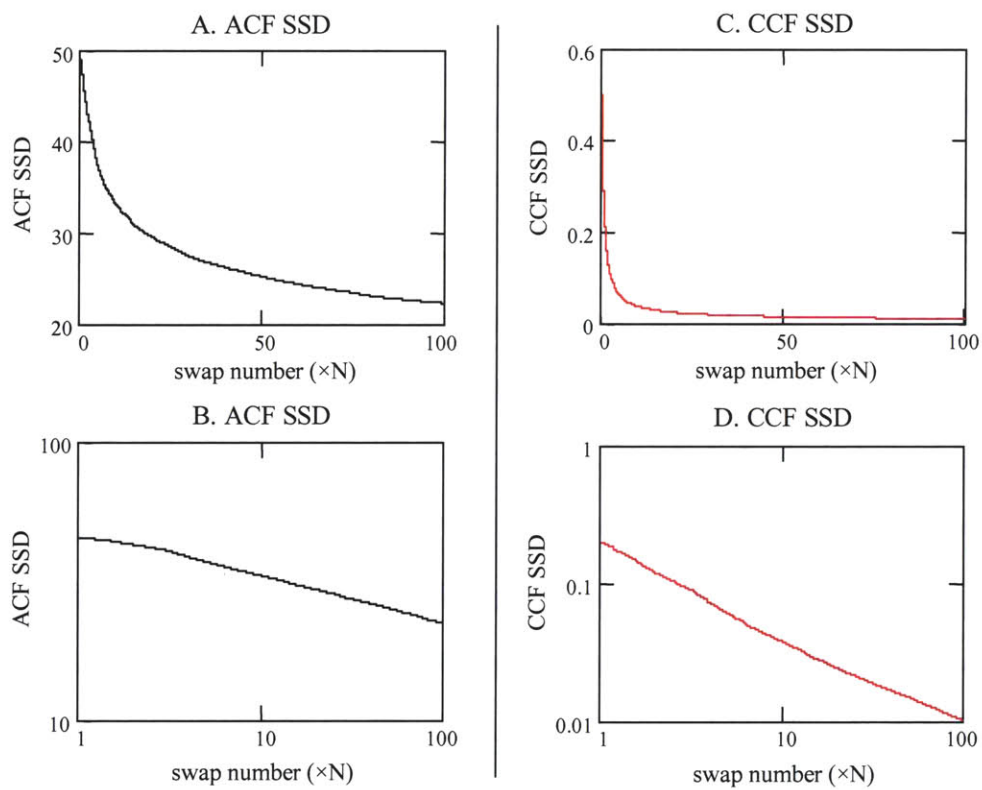
**Figure 5.7:** The sequence after being shuffled. Reordering the elements does not affect the PMF of the sequence.



**Figure 5.8:** The PSD and the ACF of the shuffled sequence (black) and the desired PSD and ACF (red).

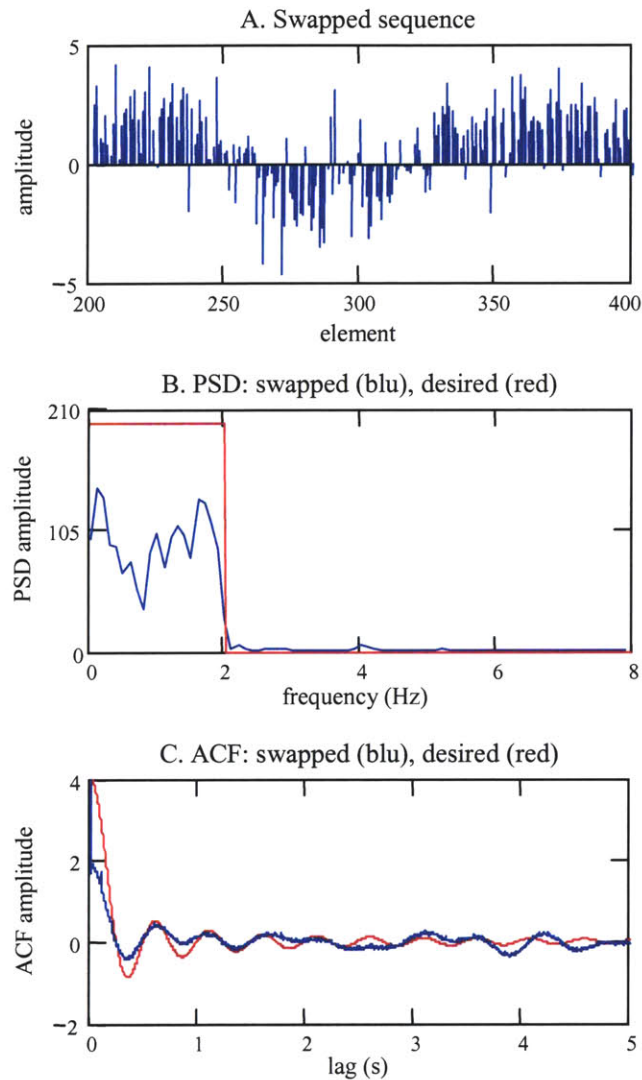
Two elements of the sequence were then randomly chosen, and the effects of swapping their position on the target velocity sequence's auto-correlation and the torque-target cross-correlation were evaluated. The efficient algorithm that Hunter and Kearney had developed to update the ACF was extended so that the CCF could also be evaluated without having to calculate explicitly the entire CCF. The sum of the squared differences (SSD) was used as the criterion to compare the actual and the desired ACF and CCF. In the example, the elements were swapped when interchanging their positions decreased the CCF SSD and did not increase the ACF SSD. For other sequences, the decision to accept a swap and to reposition two elements was made on the basis of different combinations of the correlation SSDs decreasing or not increasing. This was done to help identify which criteria resulted in the ACF and CCF converging most rapidly to their desired values. Surprisingly, none of the combinations consistently produced better results than the others (this is discussed in Section 5.4). The stochastic interchange was conducted a prescribed number of times, usually from 250 to 1000 times the length of the sequences. The rate at which the ACF and the CCF approached their desired values generally decayed exponentially (Figure 5.9). Most of the improvement was realized by the time a number of swaps equivalent to 20 to 30 times the length of the sequence had been evaluated. The remaining swaps were still necessary to get the ACF and CCF close to the desired functions, but the rate of improvement was

lower than at the beginning. The interchanges could also have been conducted until the ACF and CCF were within some range of the desired values, or until improvements no longer occurred. After evaluating 600,000 swaps (100 times the length of the sequence), and accepting those that decreased the CCF SSD and did not increase the ACF SSD, the example sequence was rearranged to that plotted in Figure 5.10. Its ACF is similar to the desired ACF, and its PSD shows that its power is highest between 0 and 2 Hz.



**Figure 5.9:** The convergence of the ACF and the CCF as a function of the number of swaps. Plots A and C show the sum of the squared difference (SSD) between the actual and the desired correlation functions as a function of the number of swaps. The functions are plotted on a log-log scale in plots B and D. The nearly linear behavior of the SSDs in the log-log plots indicates that the ACF and CCF converged to their desired values approximately exponentially.





**Figure 5.10:** The sequence, its ACF, and its PSD after being reordered by the Hunter-Kearney stochastic swap (blue). The sequence is shown with the desired PSD and ACF (red).

### *Defining the desired ACF*

It was more intuitive to specify the ordering of the target sequence through the power spectral density function than directly from the ACF. The gaze tracking studies had identified the frequency content of the target that lead to good tracking, and so the PSD was the characteristic

that the sequence was being ordered to meet. The desired ACF is simply the inverse Fourier transform of the desired PSD.

The sequence's auto-correlation was estimated with the biased autocorrelation estimator:

$$r_{xx}[j] = \frac{1}{N} \sum_{i=0}^{N-1-j} x[i]x[i+j], \quad \text{Eq. 5.3}$$

and the PSD from an N to N Fourier transform of the ACF:

$$C_{xx}[n] = \sum_{k=0}^{N-1} r_{xx}[k] * \exp(-j2\pi n \frac{k}{N}). \quad \text{Eq. 5.4}$$

The ACF was calculated from the PSD with the inverse transform:

$$r_{xx}[k] = \frac{1}{N} \sum_{n=0}^{N-1} C_{xx}[n] * \exp(j2\pi \frac{n}{N} k) \quad \text{Eq. 5.5}$$

(Press *et al.*, 1997).

The biased auto-correlation estimator was used because, although biased, it has a smaller mean square error than the unbiased estimator (Jenkins and Watts, 1968). In order to preserve amplitude information in the calculation, the ACF estimator was not normalized by the sequence's variance.

The N to N transformation of the ACF into the PSD (Eq. 5.4) produces an estimate of the power at each frequency with a large variance (100%). The variance can be reduced through windowing, and by dividing the time series into smaller records and averaging calculations of the PSD. It is not improved by increasing the number of elements in the sequence since the transformation creates N discrete frequency values from N data points. More elements mean more frequency bins, and not a better estimate at each frequency. This method was used to keep the estimate of the sequence's total power consistent when transforming between the time and frequency domains, as specified by the Parseval and Plancherel relations, and not to get the best estimate of the power at a particular frequency (Press *et al.*, 1997). These definitions allow consistency when taking either the inverse transform of the PSD to get the ACF or when taking

the transform of the ACF for the PSD. This did not, however, prevent the use of a windowed, averaged power spectrum estimate when considering the final spectrum of a sequence.

*Efficient evaluation of ACF and CCF after two element swap*

It is inefficient to calculate explicitly the ACF and CCF after each swap. Each time the ACF of an N-element sequence is calculated to a maximum lag of M,  $\sum_{n=N-M}^N n$  products must be evaluated. Twice as many products have to be calculated for the CCF, because unlike the ACF, the negative lag values of the CCF do not equal their positive lag counterparts. For the short, 6000 element sequence used as an example, over  $1.3 \times 10^7$  products had to be calculated for the ACF after every swap when the maximum lag was set to half the sequence length. This brute force method is very time consuming considering that the number of stochastic swaps that must be conducted in order for the sequence to have a satisfactory ACF and CCF may exceed 100 times the length of the sequence.

Hunter and Kearney (1983) recognized that a maximum of eight terms in the biased auto-correlation estimator were affected when the position of two elements in the sequence were swapped. It was therefore not necessary to recalculate the entire auto-correlation function after each swap. The ACF, at a particular lag value, could be updated by adding the correction, D/N, to the correlation value. N is the length of the sequence, and D is found with the following algorithm. The indices for the elements being considered for the swap are j1 and j2.

1.  $D = 0$
2. If  $(j1 + k \leq N - 1)$  and  $(j1 + k \neq j2)$  then  $D = D - x_{j1}x_{j1+k} + x_{j2}x_{j1+k}$  else  $D=D$
3. If  $(j1 - k \geq 0)$  and  $(j1 - k \neq j2)$  then  $D = D - x_{j1}x_{j1-k} + x_{j2}x_{j1-k}$  else  $D=D$
4. If  $(j2 + k \leq N - 1)$  and  $(j2 + k \neq j1)$  then  $D = D - x_{j2}x_{j2+k} + x_{j1}x_{j2+k}$  else  $D=D$
5. If  $(j2 - k \geq 0)$  and  $(j2 - k \neq j1)$  then  $D = D - x_{j2}x_{j2-k} + x_{j1}x_{j2-k}$  else  $D=D$

The efficient update for the ACF is not unique to the ACF, and was modified for this research to update the cross-correlation between two sequences. The algorithm assumes that elements of sequence Y are being swapped, the element indices range from 0 to N-1, and that the

biased cross-correlation estimator is defined as:  $r_{yx}[k] = \frac{1}{N} \sum_i x[i-k]y[i]$  Eq. 5.6

where  $i = k$  to  $(N-1)$  for positive lags, and  $i = 0$  to  $(N-1-k)$  for negative lags. The CCF, at lag  $k$ , is updated by adding the correction,  $D/N$ , where  $D$  is calculated by:

1.  $D = 0$
2. If  $(j1 - k \geq 0)$  and  $(j1 - k \leq N - 1)$  then  $D = D - x_{j1-k}y_{j1} + x_{j1-k}y_{j2}$  else  $D=D$
3. If  $(j2 - k \geq 0)$  and  $(j2 - k \leq N - 1)$  then  $D = D - x_{j2-k}y_{j2} + x_{j2-k}y_{j1}$  else  $D=D$

The algorithm for the cross-correlation has fewer update comparisons than for the auto-correlation algorithm. Two sequences are involved in the CCF calculation, but only the elements of the Y sequence are being interchanged. Only one term in each product in the cross-correlation sum is being changed, rather than potentially both as for the auto-correlation sum. The swap has an effect only when the potential index of the X element,  $j1-k$  or  $j2-k$ , falls within the range of the sequence elements, 0 to N-1.

#### 5.2.4 Low pass filtering the target velocity sequence

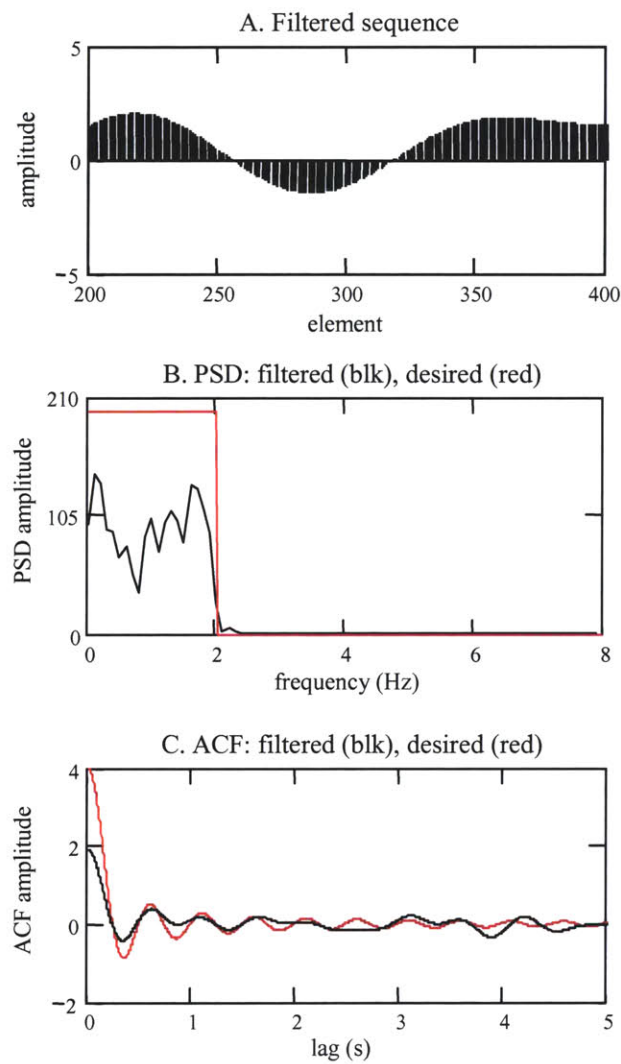
The target velocity sequence was low-pass filtered to remove unwanted high frequency content that remained in the sequence after the stochastic swap was complete. Frequencies outside of the desired frequency range made the target trajectory jittery, and more difficult to track, and were removed so the visual target was not tracked using saccades or with too much error. It was difficult to arrange all the elements of the sequence so that they fell within the constraints imposed by the desired ACF and CCF. It was more difficult than when trying to meet a desired ACF alone. The stochastic swap does the best it can, and can be helped by a judicious

choice of the PDF, but the position of some elements will cause the ACF and CCF to differ from that desired. It can be seen in Plot b of Figure 5.10 that after the swap the target velocity sequence had only about half the power of the desired PSD from 0 to 2 Hz. The total power of the velocity sequence was the same as that represented by the PSD, so about half the sequence's power fell outside the desired frequency band. It was not uncommon when developing signals for VOR experiments for only 30 to 40% of the total power of the sequence to be in the desired frequency range.


The low-pass filter was designed using the Kaiser-window filter design method (Oppenheim *et al.*, 1999). It was designed as an acausal, finite impulse response (FIR) filter, symmetric about the time-0 element. By making it symmetric about time 0, the filter did not introduce a phase change into the velocity sequence. If the filter had introduced a phase change into the target velocity sequence, the temporal relationship between the torque and the target sequences would have been shifted and their cross-correlation function would have been altered significantly.

The filtered velocity sequence shown in Figure 5.11 is much smoother than after it had been arranged by the stochastic swap algorithm (Figure 5.10). The same underlying pattern can be seen in both plots, but the filtered sequence makes no abrupt jumps and the velocities transitioned nicely into one another. The filtered sequence had the same PSD as the unfiltered sequence between 0 and 2.0 Hz, but it had little power outside this frequency band. The low-pass filter did not change the sequence's ACF dramatically. The most obvious change was in the zero lag value, which is proportional to the total power of the sequence. After the filter, the zero lag ACF value was half the value before the filter, just as the total power was halved. The rest of the ACF remained virtually unchanged. The shuffling and swapping procedures had only affected the order of the elements in the velocity sequence and so the sequence had retained its original amplitude distribution, which is Gaussian in this example. The low-pass filter affected the magnitude of the elements in the sequence, and therefore altered the sequence's amplitude

distribution (PMF). The PMF of the filtered sequence did not approximate a Gaussian PDF as did the original distribution (Figure 5.12). How great of an effect the low-pass filter had on the sequence's PMF depended on how much of the sequence had remained in higher frequency bands after the stochastic swap was complete. The requirement of filtering the sequence was a necessary limitation associated with wanting the frequency criteria to be strictly met.



**Figure 5.11:** The sequence, its PSD, and its ACF after being filtered by a 0.0 to 2.0 Hz low-pass filter




**Figure 5.12:** The PMF of the arranged sequence (black) is shown with a Gaussian PDF (red).

#### 5.2.5 Integrate the velocity sequence into a position command signal

The target velocity sequence was integrated into a position sequence using an extended trapezoidal rule (Press *et al.*, 1997). This position sequence was the command signal that directed the galvanometer and moved the target's image across the screen. Additional processing was done to make the position sequence more useable for the experiments. It was given a zero mean so that it was well balanced relative to the center of the screen. This effectively shifted the target's initial position, but did not affect the target's velocity. It therefore did not alter the torque-target velocity relationship that had been carefully designed into the target. In some cases, after the position was shifted, its amplitude was increased so that the target spanned the full width of the screen. This also amplified the velocity of the target, so was used carefully, and checked to ensure that the target's velocity was not increased to magnitudes that would encourage the use of saccadic eye motions. The velocity sequence's range had generally been decreased by the low pass filter, so there was room to amplify the velocity and still remain within the 10 to 40 deg/s limit of smooth pursuit movements. Finally, the first and last second of the position sequence was replaced by a constant velocity ramp so that the target began smoothly from, and ended smoothly

at the zero position. This helped the subjects acquire the target at the beginning of the experiment, and brought their gaze back to the zero position at the end.

The sequence shown in Figure 5.13 was a 90 s target position command used in one experiment to move the target throughout a +/- 30 degree displacement on the screen. There was approximately a 10 to 1 ratio between the target's position on the screen, in degrees, and the voltage signal applied to the galvanometer that moved the target's image.



**Figure 5.13:** A target position sequence

### 5.3 Defining bounds for zero cross-correlation

It has been stated repeatedly, and was shown mathematically in Chapter Two, that in order for the system level performance of the VOR to be analyzed, the cross-correlation between the torque and target input sequences should be zero. Although this was desired, in practice it could not be accomplished. The input sequences could not meet the stimulation requirements for the visual and vestibular systems, and also have a cross-correlation function (or cross-covariance function) that was identically zero. Instead, the sequences had to be designed so that their cross-correlation function, which was calculated to a limited lag, was lower than a bound that represented an acceptable level of correlation. When the magnitude of the CCF was smaller than this bound, it could be approximated as zero, and the VOR analysis could be conducted as if the



sequences were ideally uncorrelated. Since the PRB torque sequence approximated a bandlimited white sequence, its ACF was a suitable bound to decide if the sequences were uncorrelated.

### 5.3.1 The torque-target CCF cannot be made identically zero

The fact that the torque and target sequences could not be made so that their CCF was identically zero can be understood by considering the practical difficulty of making the cross-correlation function equal to zero at all lags (Eq. 5.9), or more rigorously by considering that the torque and target sequences had to have power in similar frequency bands.

#### *Cross-correlation calculation*

The cross-covariance (CCVF) of two random processes,  $x(t)$  and  $y(t)$ , is defined as:

$$C_{XY}(t_1, t_2) = E\{(x(t_1) - \mu_X)(y(t_2) - \mu_Y)\}, \quad \text{Eq. 5.7}$$

where the processes are assumed to be wide sense stationary,  $E\{\}$  is the expected value, and  $\mu$  is the process mean. The cross-correlation (CCF) function is defined as:

$$R_{XY}(t_1, t_2) = E\{x(t_1)y(t_2)\}. \quad \text{Eq. 5.8}$$

Two random processes are defined as “uncorrelated” when the cross-covariance function is zero for all times,  $t_1$  and  $t_2$ . The two processes are defined as “mutually orthogonal” when the cross-correlation function is zero for all times  $t_1$  and  $t_2$ . When the processes are zero-mean, the cross-correlation and cross-covariance functions are the same. In this case, two sequences that are uncorrelated are also mutually orthogonal, and vice versa. Since zero-mean sequences were used for the torque and target velocity sequences, “uncorrelated” has been used to describe the sequences when referring to either the cross-correlation or cross-covariance function being zero.

The torque and target sequences that were being generated for the VOR protocol were finite length, discrete-time sequences, not probabilistically defined random processes, and therefore could not have their CCF calculated via the function definitions of Eqs.5.7 and 5.8.

Their CCF, which equaled their CCVF, was calculated using the biased cross-correlation estimator, which, for positive lag values ( $j$ ), is defined as:

$$c_{XY}(j) = \frac{1}{N} \sum_{i=0}^{N-1-j} x(i)y(i+j), \quad \text{Eq. 5.9}$$

where  $N$  is the length of the sequences,  $X(i)$  can be thought of as representing an element of the torque sequence, and  $Y(i+j)$  an element of the target sequence.

The elements of the target sequence, which represented the velocities of the visual target, were arranged by the modified Hunter-Kearney algorithm so that the sequence's order had a particular spectral content and so that its CCF with the torque sequence, calculated with Eq. 5.9, equaled zero. Unless all the elements of the target velocity sequence were zero, it is unlikely that the elements could be arranged so that the cross-correlation sum was identically zero at all lag values,  $j$ . It is not unreasonable to expect that the sequence could be arranged so that the CCF was smaller than a defined bound, but it cannot be expected that the CCF would be identically zero and the sequences strictly uncorrelated.

### *Common frequency spectra*

The fundamental reason why the torque and target sequences could not be made to have zero cross-correlation was that the sequences needed to span similar frequency spectra. The PRB torque sequences were created to have power between 0 and 5 Hz, and the target sequences were generated to have a bandwidth from 0 through 1 to 3 Hz. Two sequences that have power in the same frequency bands have a non-zero cross-power spectral density (CPSD) (Jenkins and Watts, 1968). The CPSD is the Fourier transform of the cross-correlation function. Since the total power of a sequence must be the same no matter in what domain the sequence is represented, a non-zero CPSD implies that the sequence must have non-zero CCF. The torque and target sequences cannot, therefore, be strictly uncorrelated.

### 5.3.2 ACF of the PRB torque as a correlation bound

The auto-correlation function of the PRB torque sequence can be used as a bound to decide if two sequences are sufficiently uncorrelated. A full cycle of a PRBS is accepted as being approximately band-limited white, and is considered uncorrelated with itself at lags greater than zero. Theoretically, the auto-correlation of an infinite length, continuous time, white sequence is the delta function. Its only non-zero value is its zero lag value. For a finite, discrete time PRBS, the off-impulse values are not zero, but have a theoretical value of  $-a/n$ , where  $a$  is the amplitude of the binary sequence and  $n$  is the length of the seed sequence from which the PRBS originated (Davies, 1970). Although the off-impulse ACF values are not zero, the finite length, discrete time, PRBS is considered uncorrelated with itself at these lags. The off-impulse values of the auto-correlation function can therefore stand as a measure of two sequences being uncorrelated. If the target sequence can be designed so that the magnitude of the target-torque CCF is smaller than the magnitude of the off-impulse values of the PRB torque ACF, then the torque and target sequences can be considered uncorrelated, and any dynamic effects caused by the cross-correlation of the inputs ignored.

### 5.3.3 Evaluate CCF over a limited lag

In order for the target sequence to be arranged so that the ACF and CCF design goals were met satisfactorily, the maximum lag over which the CCF between the torque and target sequences was evaluated had to be limited. Since the two sequences had power in common frequencies, their CCF could not be zero. The requirement that the sequences spanned a common spectrum did not, however, dictate that the sequences had to be strongly correlated at all lags. The total power of the cross-spectral density function will also be the same as that in the cross-correlation function, but where it arises in the CCF can be controlled by how the sequences are designed.

A key to making the sequences appear uncorrelated is to separate temporally the common frequencies so that strong correlations occur at long lag values. By making the lag values greater than the settling times of the subsystems in the gaze tracking system, the high correlation levels should not have a measurable influence on the dynamics of the systems. Consider the discrete-time, linear, convolution equation that relates the inputs ( $u, x$ ), the output ( $y$ ), and the subsystems ( $g, h$ ) of a causal, linear, two-input, single-output system:

$$y[n] = \Delta t \sum_{k=0}^{M-1} g[k]u[n-k] + \Delta t \sum_{k=0}^{M-1} h[k]x[n-k]. \quad \text{Eq. 5.10}$$

The subsystems  $g$  and  $h$  are approximated by finite length impulse response functions (IRF) with settling times of length  $M$ . This implies that the output at time  $n$  is only affected by inputs to the subsystem that happened later than time  $(n-M)$  (e.g.  $u(n-M+1), u(n-M+2) \dots u(n)$ ). Inputs that occurred at times earlier than time  $(n-M)$  are assumed not to have a measurable effect on the output at time  $n$ . The two inputs can therefore be considered uncorrelated when their cross-correlation is near zero at all lags up to a maximum lag equal to the length of the system's IRF. At lags that are greater than the duration of the IRF, the interaction between the inputs does not matter because the inputs no longer contribute to the systems' output. The two sequences may be highly correlated at some lag values, but as long as the correlation occurs at lag times that are greater than the duration of the system's IRF, the CCF can be mathematically approximated as zero. For the VOR experiment, the cross-correlation between the torque and target sequences had to be near zero for lags only as long as the IRFs of the subsystems in the gaze system.

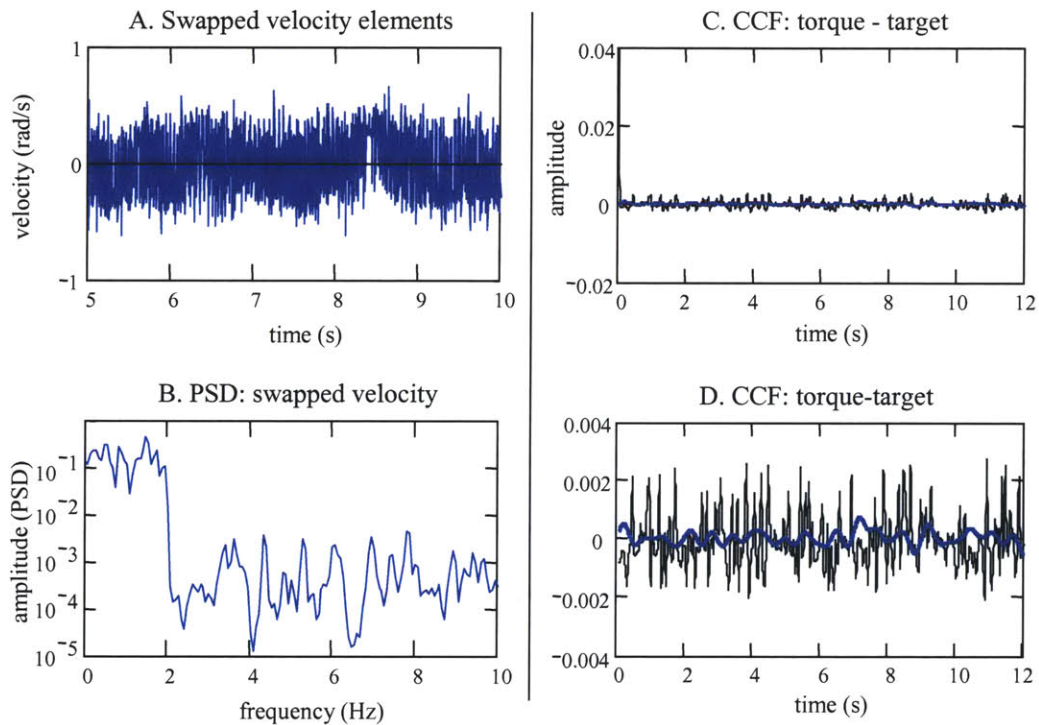
It can also be seen from the correlation-based convolution equations that were developed in Chapter 2 (Eqs. 2.17, 2.18, 2.20) that the cross-correlation function at lags greater than the duration of the systems' IRF do not have to be considered in order to identify the dynamics of the IRF. The maximum lag of the correlation functions that are used in the equations from which the IRF is calculated is equal to the length of the IRF. The input sequences can therefore be treated as uncorrelated as long as their cross-correlation function is near zero out to a max lag of  $M$ .

*Swapping for ACF and CCF; evaluate CCF over all lags*


When the modified Hunter-Kearny method was used to try to make the torque and target sequences uncorrelated at all lag values, the resultant CCF was low enough to consider the sequences uncorrelated, but the target sequence was inadequate for stimulating the gaze system. The algorithm was not able to arrange the target velocity sequence so that the sequence had sufficient power in the desired frequency range. The target velocity sequence shown in Figure 5.14 was stochastically swapped to have a 0 to 2.0 Hz frequency spectrum and to be uncorrelated with a torque PRBS. The decision to accept a swap of two elements was based on an evaluation of the target's ACF and torque-target CCF out to a maximum lag of 90 s, a lag equal to the length of the sequences. The final torque-target CCF is plotted together with the ACF of the PRB torque sequence in Figure 5.14. The magnitude of the CCF was smaller than that of the torque's ACF, which indicates that the torque and target sequences were less correlated than the PRB torque sequence was with itself. In fact, the mean squared amplitude of the off-impulse ACF values was 15 times the mean squared amplitude of the CCF. Although the target velocity and torque sequences can be considered to be uncorrelated, too little of the target velocity's total power was in the desired 0 to 2.0 Hz frequency band. The power spectral density (PSD) of the target velocity sequence is also shown in Figure 5.14. The PSD of the sequence was greatest from 0 to 2.0 Hz, but most of the sequence's total power was in frequencies above 2.0 Hz. This is obvious in the plot of the target velocity which is not smooth and so is unlikely to be followed by a smooth gaze response.

The velocity elements had to be filtered by a LPF to have the smoothness required to produce a position sequence that would stimulate the appropriate gaze tracking responses. However, too much of the sequence's total power was lost when the sequence was filtered. The outcome of the low pass filter is shown in Fig 5.15. The filtered sequence retained only 12% of the total power of the original sequence. The original random sequence had a zero mean and a


distribution that looked like a combination of two Gaussian sequences; one with a mode at 0.28 rad/s (15 deg/s) and the other with a mode of -0.28 rad/s (Figure 5.16a). The filter removed so much of the velocity content, which had remained above the desired 2.0 Hz limit, that the sequence's amplitude distribution was destroyed. Even the range of velocities that were required for testing the VOR was lost. The sequence originally had velocities that reached 0.5 rad/s, but after being filtered the sequence's maximum velocity was only 0.28 rad/s. Consequently, the position sequence that was created from this velocity sequence was inadequate for testing.



**Figure 5.14:** The target velocity sequence in plot A was arranged to have a 0.0 to 2.0 Hz PSD (plot B), and to be uncorrelated with a PRB torque sequence over all possible lag values. The CCF between the target and the torque is shown in plots C and D (blue) along with the ACF of the torque (black).



**Figure 5.15:** The sequence (plot A) and its PSD (plot B) after being filtered by a 0.0 to 2.0 Hz low-pass filter.




**Figure 5.16:** Sequence probability mass functions. The plot on the left shows the sequence’s PMF before the sequence was low-pass filtered, and the plot on the right shows the PMF after the sequence was filtered.

*Swapping for ACF and CCF; evaluated over limited lags*

The solution to meeting both the ACF and CCF criteria was to limit the lag over which the torque and target sequences were made uncorrelated. This eased the requirements that had to be met when two elements in the target sequence were swapped, and allowed for more swaps to be accepted that improved the spectral content of the target signal.

Since the torque and target sequences had power in common frequency bands, there was no way to avoid some degree of correlation. The stochastic swap algorithm was thus limited to

arranging the sequence so that the power in the cross-correlation function was distributed over all lag values. However, when the maximum lag over which the sequences had to be uncorrelated was reduced, the power in the cross-correlation function could be pushed to greater lag values. Arrangements of the target sequence that had a good ACF and a good CCF over the limited lag range, but that were highly correlated with the torque sequence at large lags, could now occur.



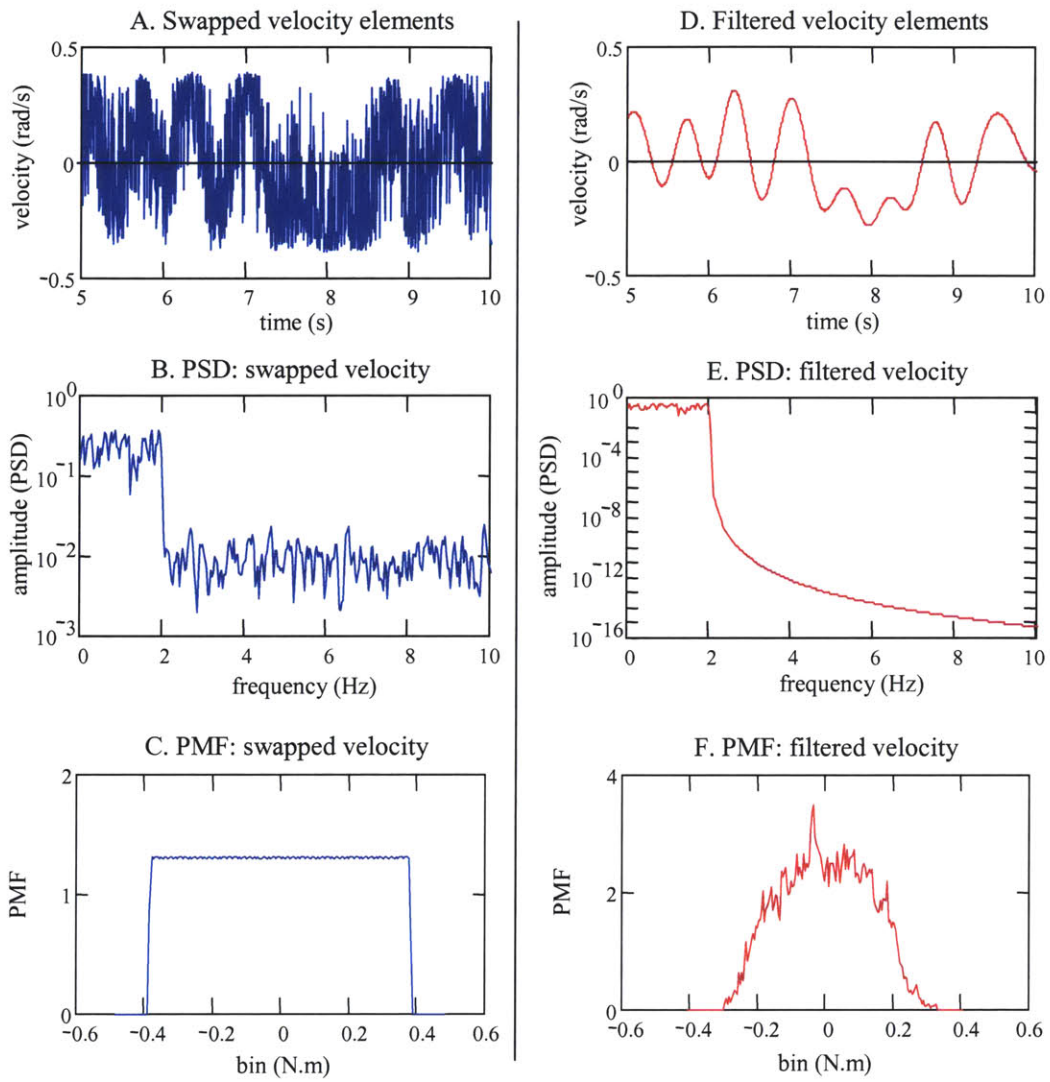
**Figure 5.17:** CCF between the target and the torque (red), and the ACF of the torque (black). The target sequence was arranged to meet ACF and CCF criteria over a limited range of lags. After the swap, the ACF of the PRB torque sequence was 1015 times greater than the CCF between the torque and the target. Plot B shows a magnified view of plot A.

Figure 5.17 shows the cross-correlation function that resulted when the torque and target sequences were made to be uncorrelated only out to a limited lag value. The same sequence of velocity elements was used in this example as in the example shown in Figure 5.14. As before, the goal for the target's frequency spectrum was to have power from 0 to 2.0 Hz. The CCF was evaluated over 10 seconds of lag, a lag time that was longer than the settling time of the subsystems in the gaze response system. Between 0 and 10 seconds of lag, the magnitude of the cross-correlation function was very small. The mean squared amplitude of the torque's ACF was 1015 times greater than the mean squared amplitude of the torque-target CCF. The mean squared amplitude of the ACF was only 15 times that of the CCF when the CCF had been evaluated to a lag of 90 s. The mean squared values were evaluated after the peak ACF value had decayed, and



so the CCF was only compared against ACF values that represented “uncorrelated” values. At lags greater than 10 s, the torque and target sequences were more correlated, although the mean squared amplitude of the CCF was never much greater than that of the ACF. The power in the correlation, or cross-spectrum, of the torque and target sequence had been pushed to these greater lag values, which is what allowed the magnitude of the torque-target CCF to be so small at lags less than 10 s.

It can be seen in Figure 5.18 that after being reordered by the stochastic swap algorithm, the target sequence had its highest power between 0 and 2.0 Hz, but it still had some power outside the desired range. The velocity sequence was smoother than when it had been created with the CCF evaluated to a lag of 90 s (Figure 5.14), but it still needed to be low-pass filtered before it could be integrated into an acceptable target position sequence. However, the low-pass filter did not remove nearly as much power from the sequence, nor did it affect the sequence’s amplitude distribution as much as when the CCF was evaluated over all lag values. The sequence retained 38% of its original total power when the CCF was evaluated out to 10 s, as compared to only 12% when the CCF was evaluated to 90 s. In some of the sequences created for the VOR experiment, the sequence retained as much as 60% of its power. The sequence’s amplitude distribution was altered by the filter, but the sequence retained velocities close to the desired range. In both Figure 5.18 and Figure 5.16 the filter appears to have made the sequence’s amplitude distribution more Gaussian. This was common to nearly all the sequences created for the VOR experiments.



**Figure 5.18:** The sequence, its PSD, and its PMF after being ordered by the modified Hunter-Kearney swap routine (plots A, B, C) and then after a low-pass filter (plots D, E, F)

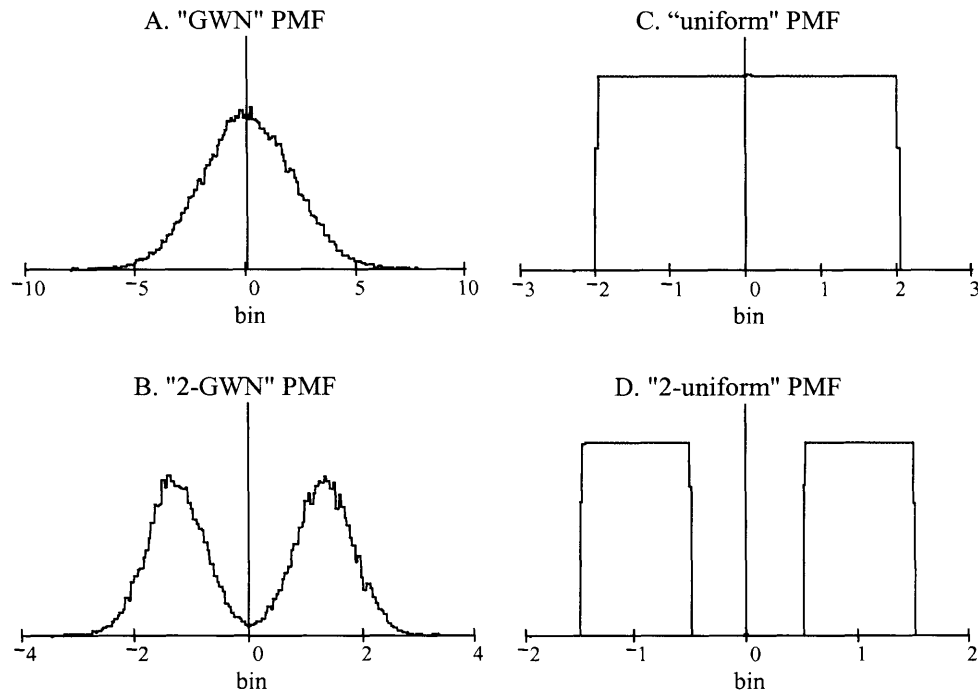
## 5.4 Swap conditions: number of elements, evaluation criteria, sequence PMF

### 5.4.1 Optimizing the swap routine

The modified Hunter-Kearney method manipulates the order of the elements in a sequence so that the sequence's spectral and correlation measures gradually progress toward desired values. It is not a closed form technique that produces a specific solution, but it does not fail if an exact solution cannot be found. How well the modified Hunter-Kearney technique can arrange a sequence to achieve desired properties is dependent on the goals set for the sequence (e.g. the desired ACF, CCF, and lag over which the correlations are evaluated), the raw elements available to be swapped (e.g. the PMF of the sequence), and how the stochastic swap is conducted (e.g. the number of elements in the swap, how the success of the swap is evaluated, and the number of swaps executed). To help optimize the modified Hunter-Kearney method for creating target velocity sequences appropriate for testing the VOR, we investigated how the convergence of the ACF and CCF to their desired values was affected by the number of elements in the swap, how the decision to accept the swap was made, and the PMF of the sequence being swapped.

Fifty-four target velocity sequences, each 6000 elements long, were manipulated with stochastic-swap routines with conditions specifically designed to investigate each of the above concerns. All the sequences were arranged to have power between 0 and 2.0 Hz, and to be uncorrelated with the same PRB sequence. Each sequence had an amplitude distribution that approximated one of four PDFs: a Gaussian (GWN), bimodal GWN (2-GWN), uniform, or "bimodal" uniform (2-uniform) distribution (Figure 5.19). The rationale in selecting the bimodal distributions was that they would have some qualities of a binary sequence, which can be ordered to have a broad range of frequency spectra, but would have a richer velocity content. The characteristics of the amplitude distributions were chosen so that the sequences had appropriate

velocities for testing the VOR. Although the sequences were not as long as those used for the VOR experiments, their other characteristics were similar. The stochastic swap routines were executed for 360,000 swaps, which was 60 times the sequence length. This number of swaps was chosen because it produced an amount of data that was within the memory limits of the computer system and software used to analyze the data.



**Figure 5.19:** Four amplitude distributions. A. Gaussian (GWN); B. bimodal Gaussian (2-GWN); C. uniform; and D. “bimodal” uniform.

The findings of this investigation apply directly only to the modified Hunter-Kearney process as it was used to create the torque and target sequences for the VOR experiment. The results should not be interpreted broadly to apply in general to the modified Hunter-Kearney process, because the ability of the method to produce an outcome, and the influence of each of the three characteristics, will vary greatly depending on the objective of the swap.

#### 5.4.2 Two versus three-element swaps

One hypothesis was that the potential of a sequence to be arranged using the modified Hunter-Kearney technique to meet ACF and CCF goals was limited when only two elements were considered to be swapped. The procedure might be able to arrange the sequence better if the positions of three, or four, elements were exchanged at once. Simultaneously interchanging more than two elements might result in a “non-linear” dynamic where, had the same of elements been exchanged sequentially, two at a time, there would have been no improvement in the ACF and CCF after the first two element swap, and therefore no reason to retain the new element positions as the initial stage for the next two element swap.


A three-element stochastic-swap routine was created. Three elements in the sequence,  $j_1$ ,  $j_2$ , and  $j_3$ , were randomly selected to be swapped. There are two permutations where the position of each of the three elements is changed;  $j_3, j_1, j_2$  and  $j_2, j_3, j_1$ . The elements were reordered into the permutation that improved the ACF and CCF the most. If neither permutation improved the ACF and CCF, then the elements remained in their original positions,  $j_1, j_2, j_3$ . The algorithm designed by Hunter and Kearney to update the sequence’s ACF after two elements had swapped positions was used as a building block to update the ACF and CCF when three elements changed positions. The ACF and CCF for each three-element permutation were calculated by using three, two-element updates. First, an intermediate reordering was done. Elements  $j_1$  and  $j_2$  were swapped to create the order  $j_2, j_1, j_3$ , and the ACF and CCF were calculated using the efficient, two-element, update. The two, three-element permutations were then made by switching elements  $j_1$  and  $j_3$  to get the order  $j_2, j_3, j_1$ , and by switching elements  $j_2$  and  $j_3$ , to get  $j_3, j_1, j_2$ . The ACF and CCF for each permutation were then calculated using another two-element update. This three-element update required about three times as many calculations as that needed to evaluate a two-element swap, but it was still more efficient than explicitly recalculating the ACF and CCF after each potential swap.

It was expected that a sequence’s ACF and the CCF would converge to their desired values more rapidly with the three-element swap, since more was manipulated and evaluated in a

“single”, three-element swap than in a single, two-element swap. A “single swap” in the two-element swap routine compared the ACF and CCF after the positions of two elements in the sequence were changed with the ACF and CCF of the original order. A “single swap” in the three-element swap routine involved choosing the better of two, three-element permutations, and comparing its ACF and CCF with the ACF and CCF of the original order. Surprisingly, the three-element swap routine did not lead to a better convergence of the ACF and CCF towards their desired values. Representative results from the two- and three-elements swap routines are shown in Table 5.1 and in Figures 5.20 and 5.21. The plots show the progression of the sum of squared differences between the actual and desired correlation measures. The final outcomes for the two-element and three-element swaps were not very different for any of the sequences. The two-element swap consistently arranged the sequence to have slightly more power in the desired frequency bands, and in the majority of cases, to have a lower CCF with the torque sequence.

PDF	GWN	GWN	2-GWN	2-GWN	uniform	uniform	2-unif	2-unif
swap elements	2	3	2	3	2	3	2	3
% power	33	27	34	26	34	32	25	24
CCF swap	10.7	14.2	3.9	5.3	4.0	4.5	5.1	5.1
CCF lpf	4.4	3.5	1.5	2.0	1.3	1.9	0.8	1.0

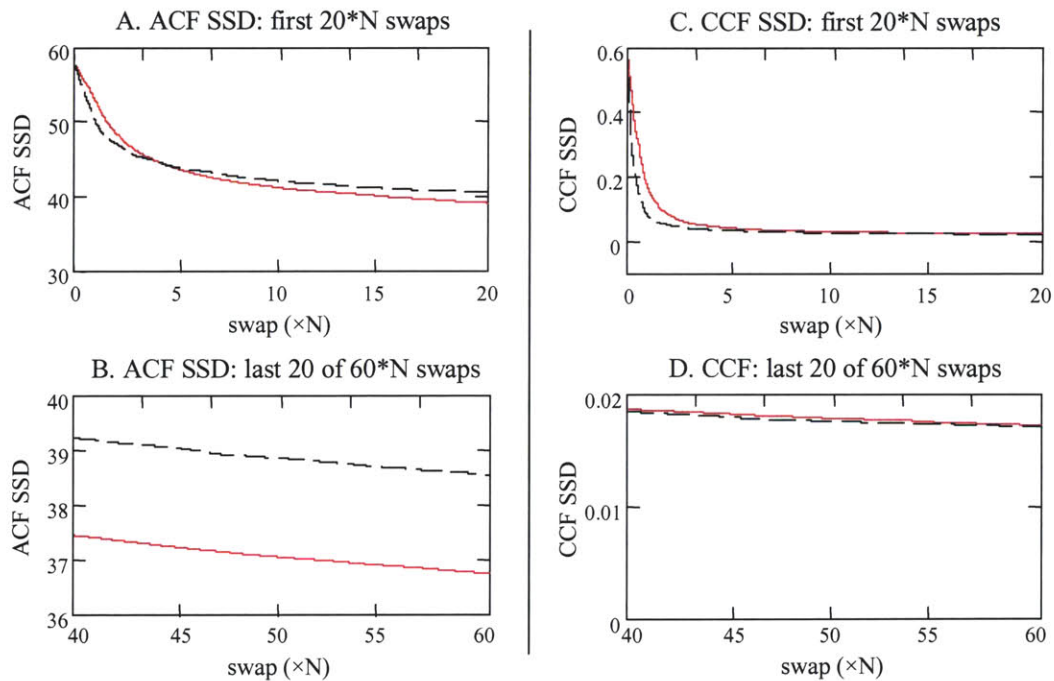
**Table 5.1:** Comparison of the results from two- and three-element swaps. “% power” indicates the percent of the sequence’s total power that was in the desired 0 to 2.0 Hz frequency range after the swap routine finished. “CCF swap” is the mean squared amplitude of the torque-target cross-correlation function immediately after the swap-routine that was evaluated to a maximum lag of 10 s. “CCF lpf” is the mean squared amplitude of the CCF after the sequence was low-pass filtered to remove power in frequencies beyond the desired 0 to 2.0 Hz frequency band.



**Figure 5.20:** Two- vs. three-element swap routines for a Gaussian sequence. The plots show the convergence of the ACF (plots A and B) and the CCF (plots C and D) to their desired values. The sum of the squared difference (SSD) between the actual and desired values for the functions are plotted as a function of the swap number. The solid red plots are for the two-element swap, and the dashed black are for a three-element swap. Plots A and C show the SSD for the first  $20 \times N$  swaps, where  $N$  is the length of the sequence, and plots B and D are for the last  $20 \times N$  swaps.

Although the three-element swap routine did not produce a better final outcome, the ACF and CCF were driven towards their desired values faster in the beginning than for the two-element swap routine. In the example shown for the 2-uniform sequence (Figure 5.21), the final SSD for the two- and three-element routines were close for the ACF (38.6 vs. 36.8) and nearly identical for the CCF. The two-element routine generally caught up to and surpassed the three-element routine once the number of swaps exceeded five to ten times the sequence length. It makes sense that the three-element swap would begin faster than the two-element routine. Each potential swap selected the best of three, three-element, permutations (including the original order), rather than the best of two, two-element permutations. It is possible that once the sequence started to take shape and began to have the desired spectral and cross-correlation characteristics, that it was harder to find a three-element permutation, than a two-element permutation, that

further improved the sequence order. Three-element permutations disturbed too much of the sequence's order. At this point, fine tuning was needed, and this was better accomplished by swapping two elements at a time. Mathematically, this can be seen from the perspective of there being a limited set of two- and three-element permutations that improve the characteristics of a sequence towards the desired values. Once the sequence begins to have the desirable ACF and CCF characteristics, there are more two-element permutations, than three-element permutations, among the set of all sequences that further improve the ACF and CCF. A randomly selected two-element swap has a higher likelihood of being in the set that improves the sequence's characteristics than a randomly selected three-element swap, and therefore the sequence improves faster by selecting two elements at a time.



**Figure 5.21:** Two- vs. three-element swap routines for a 2-uniform sequence. The plots show the convergence of the ACF (plots A and B) and the CCF (plots C and D) to their desired values. The sum of the squared difference (SSD) between the actual and desired values for the functions are plotted as a function of the swap number. The red plots are for the two-element swap, and the black are for the three-element swap. Plots A and C show the SSD for the first  $20 \times N$  swaps, where  $N$  is the length of the sequence, and plots B and D are for the last  $20 \times N$  swaps.



Although the three-element swap routine improved the sequences faster initially, there would be little advantage to incorporating the three-element swap into a routine for producing sequences for the VOR experiments. Each three-element swap required about three times the calculations for each two-element swap. Considering the extra time needed to carry out these calculations, it is far better to conduct three times the number of two-element swaps than to use the three-element routine. However the three-element routine cannot be dismissed entirely. There may be some cases when swapping three, or more, elements at a time allows for arrangements to develop that would not otherwise occur if the sequence was ordered two elements at a time. These arrangements might be the only way to achieve some ACF and CCF requirements.

#### 5.4.3 When should a swap be accepted and the sequence reordered?

When the two- and three-element swap routines were directly compared in section 5.4.2, the elements of the sequence were repositioned when the sum of the squared differences between the actual and the desired cross-correlation functions (CCF SSD) decreased ( $<$ ), and the ACF SSD did not increase ( $\leq$ ). This was not the only way that could have been selected to accept the swap and reposition the elements. Any combination of the ACF SSD and the CCF SSD that resulted in a decrease following the swap might have sufficed. It would not have been logical to accept a swap when either the ACF or CCF was made worse, and so the least restrictive criterion that could be used to accept a swap was that the SSD did not increase. To help optimize the Hunter-Kearney procedure for creating the torque and target sequences for the VOR experiment, it had to be determined if there was one way to decide if a swap should be accepted that consistently lead to the best ACF and CCF values.

Velocity sequences were arranged using stochastic-swap routines that repositioned the elements using one of four decision criteria: a) ACF SSD decreased, CCF SSD decreased, b) ACF SSD did not increase, CCF SSD decreased, c) ACF SSD did not increase, CCF SSD did not

increase, and d) ACF SSD decreased, CCF SSD did not increase. At least one velocity sequence from each of the four amplitude distributions described in Section 5.4.1 was arranged using every decision criteria. Each of the sequences was swapped to have a 0 to 2.0 Hz frequency spectrum, and to be uncorrelated with the same PRB torque sequence. The sequences were swapped 360,000 times ( $60 \times N$ ).

It was expected that the fastest way for the ACF to converge would be to accept swaps that decreased the ACF SSD but did not increase the CCF SSD. If the CCF SSD had to decrease as well, swaps that decreased the ACF SSD, but that did not change the CCF SSD, would be rejected. By putting the least restrictive requirement on the CCF, fewer swaps that improved the ACF would be rejected because they failed to meet the CCF criterion. Although the practical difference between the CCF SSD decreasing and it not increasing was small, it was thought that it could make a noticeable difference in how the ACF and CCF converged over the course of hundreds of thousands of swaps. This would also apply to the CCF. To give it its best chance of converging, it was decided that elements in the sequence should be swapped when the CCF SSD decreased and the ACF SSD did not increase.

However, no particular set of the decision criteria used to accept a swap resulted in the best outcome for either the ACF or the CCF in all the sequences tested. Contrary to what was hypothesized, the least restrictive bound on the CCF did not always yield the best ACF, nor was the best CCF always created when swaps were accepted that decreased the CCF SSD and did not increase the ACF SSD. Table 5.2 and Figures 5.22 and Figure 5.23 show representative results from the four sequences. Direct comparisons of the results were made only among the sequences with the same amplitude distribution, since only then were the initial conditions set for the swap routines and the sequences' total power the same.

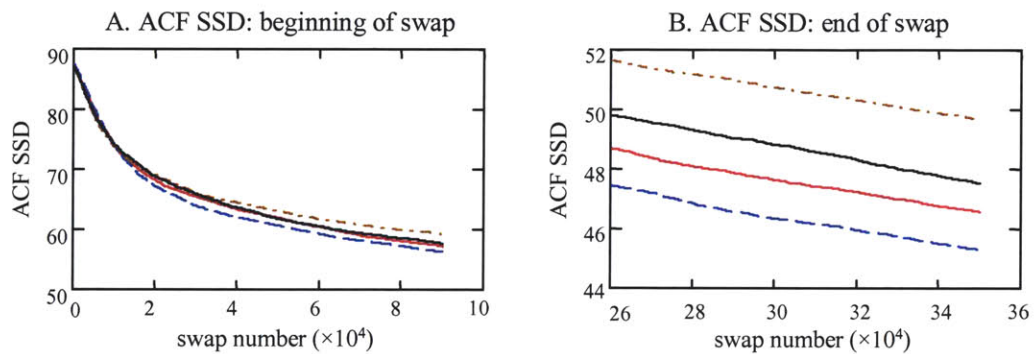
PDF	GWN	GWN	GWN	GWN		2-GWN	2-GWN	2-GWN	2-GWN
ACF cond	<	<=	<=	<		<	<=	<=	<
CCF cond	<	<	<=	<=		<	<	<=	<=
% power	30	33	31	<b>40</b>		31	34	<b>40</b>	33
CCF swap	10.8	<b>10.7</b>	10.9	11.2		5.6	3.9	3.9	<b>3.4</b>
CCF lpf	3.7	4.4	<b>3.3</b>	5.0		2.7	1.5	<b>0.6</b>	1.5
PDF	unif	unif	unif	unif		2-uni	2-unif	2-unif	2-unif
ACF cond	<	<=	<=	<		<	<=	<=	<
CCF cond	<	<	<=	<=		<	<	<=	<=
% power	33	<b>34</b>	30	32		20	25	24	<b>31</b>
CCF swap	5.0	4.0	3.7	<b>3.3</b>		5.2	<b>5.1</b>	<b>5.1</b>	6.1
CCF lpf	2.1	<b>1.3</b>	1.5	1.9		1.2	<b>0.8</b>	1.8	0.9

**Table 5.2:** Results of swap routines based on different decision criteria. GWN, 2-GWN, unif, and 2-uni sequences were arranged to meet ACF and CCF criteria using four different swap criterion. “% power” indicates the percent of the sequence’s total power that was in the desired 0.0 to 2.0 Hz frequency range after the swap routine finished. “CCF swap” is the mean squared amplitude of the torque-target CCF immediately after the swap routine. “CCF lpf” is the mean squared amplitude of the CCF after the sequences was low pass filtered to remove power in frequencies beyond the desired 0.0 to 2.0 Hz band.

The results for the GWN and the 2-unif sequences were consistent with the hypothesis regarding creating the best ACF or CCF. Both sequences had their lowest “CCF swap” value when the decision criterion accepted swaps that decreased the CCF SSD (<) and did not increase the ACF SSD (<=). Both sequences had their highest “% power” when the ACF SSD decreased (<) and the CCF SSD did not increase (<=). This decision criterion also created the worst (highest) “CCF swap” value for both amplitude distributions, so the tradeoff for giving the sequences a good spectral content was a poor cross-correlation.


The results of the 2-GWN and the uniform sequences were not consistent with the hypothesis. The 2-GWN sequence had its highest “% power” when both the ACF SSD and the

CCF SSD did not increase ( $\leq$ ), and had its best “CCF swap” value when the CCF SSD did not increase ( $\leq$ ) and the ACF SSD decreased ( $<$ ). The criteria that gave the best ACF and CCF results for the uniform sequence were contrary to what the hypothesis predicted. The uniform sequence had its best spectral power when the swap did not increase the ACF SSD ( $\leq$ ) and decreased the CCF SSD ( $<$ ), and its lowest “CCF swap” value occurred when the CCF SSD did not increase ( $\leq$ ) and the ACF SSD decreased ( $<$ ).



**Figure 5.22:** Convergence of the ACF for a uniform sequence. The SSDs between the desired and actual ACFs are plotted as a function of the number of swaps that were evaluated. The decision criteria that were used to accept a swap were: ACF SSD decreased and CCF SSD decreased (blue dash), ACF SSD decreased and CCF SSD did not increase (red solid), ACF SSD did not increase and CCF SSD did not increase (black solid), and ACF SSD did not increase and CCF SSD decreased (brown dash-dot).

Figure 5.22 shows the progress of the ACF SSD for the uniformly distributed sequence, and Figure 5.23 shows the progress of the CCF SSD for the GWN sequence. The plots are not identical to those for the other distributions, but do represent their trends. There was little difference between the four decision criteria in how the ACF or the CCF converged, particularly for the CCF. In fact, the differences between the ACF SSD and the CCF SSD curves created with different decision criteria were no greater than the variance between the ACF SSD and the CCF SSD curves when the same sequence was arranged several times using the same decision criteria.



**Figure 5.23:** Convergence of the CCF for a Gaussian sequence. The SSDs between the desired and actual CCFs are plotted as a function of the number of swaps that were evaluated. The decision criteria that were used to accept a swap were: ACF SSD decreased and CCF SSD decreased (blue dash), ACF SSD decreased and CCF SSD did not increase (red solid), ACF SSD did not increase and CCF SSD did not increase (black solid), and ACF SSD did not increase and CCF SSD decreased (brown dash-dot).

The reason that no decision criteria stood out as being the best method to arrange the sequence to meet the ACF or CCF requirements was probably because there was little effective difference between the “decreased” and the “did not increase” criteria. The only difference was that the “did not increase” criterion allowed for a swap to be accepted when the correlation measure remained the same. It is unlikely that this ever occurred when the positions of two elements in the sequence were swapped. The elements of the target velocity sequence that were selected to be swapped would have had to have been either of the same value, or their product with the elements in the torque sequence would have had to have remained the same. Because the velocity sequences were not truly random sequences, but were created by transforming uniform sequences with static filters, few, if any, of the elements in the sequence were of the same value, so when swapped, the correlation values would always have changed. The “remained the same” portion of the “did not increase” criterion probably never occurred.

Overall, the ACF and CCF converged to their desired values at about the same rate no matter what decision criteria were used to evaluate the swap. The results of the GWN and the 2-unif sequence showed that sometimes the cost of using the decision criterion that arranged the

sequence so that it had the best spectral content was that the sequence had the worst CCF measure. However, this was not reflected in the results for the other sequences. No method gave significantly, or consistently, better or worse results than another method. Each of the decision criteria was equally appropriate for creating target and torque sequences for the VOR experiment.

#### 5.4.4 The effect of the amplitude distribution on the ACF and CCF

In their 1983 paper, Hunter and Kearney described how a signal's auto-correlation function and its probability distribution function (PDF) were independent descriptions of the signal, and therefore could both be defined arbitrarily. The order of the elements in a discrete time sequence can be manipulated to give the sequence an arbitrary ACF without affecting the sequence's (amplitude) probability mass function (PMF), which is the discrete time analog of the PDF. Although this is true, the ACF and the PMF are related. The amplitude distribution of the sequence affects how well the sequence can be arranged to have a desired ACF. The amplitude distribution defines what elements are available in the sequence to be arranged. The potential of the sequence to be arranged to have a particular ACF will depend upon the availability of appropriate elements so that the desired spectral content can be created.

Although the amplitude distribution and the ACF are separate characterizations of a sequence, the potential of the sequence to be ordered to have a desired spectral content can be helped by a judicious choice of the sequence's amplitude spectrum. The corollary is that an inappropriate amplitude distribution can make it difficult to achieve the ACF and CCF requirements, no matter how the swap routine is executed. The relation between a sequence's amplitude distribution and the spectral content that the sequence can be arranged to have is not, however, an obvious one. A rich amplitude distribution does not guarantee that the sequence can be ordered to meet a range of ACFs, nor does a limited amplitude distribution prevent a sequence from achieving a range of ACFs. In fact, a binary sequence, which, other than a constant sequence, can be considered to have the most limited amplitude distribution, can be arranged to

have the widest range of frequencies. The highest discrete time frequencies, 0.5 c/s, can be created by alternating the elements of opposite sines, and the lowest frequency that it can have is limited only by the length of the sequence.

In order to identify if one of the four amplitude distributions tended to meet the ACF and CCF criteria better than the others, or if one of the distributions consistently lead to poor results, comparisons were made between the results of the swap routines conducted on the 54 6000-element sequences. Each sequence was arranged using one of the several iterations of the modified Hunter-Kearney method to have a frequency spectrum between 0 and 2.0 Hz, and to be uncorrelated with the same PRB torque sequence. Since sequences with different amplitude distributions had different initial swap conditions and different total powers, normalized correlation measures had to be used to compare the results of the swap routines conducted on different sequences. The results were compared using the percent of the sequence's total power that was in the desired 0 to 2.0 Hz frequency range (percent power), and the mean squared amplitude of the normalized cross-correlation function right after the swap routine (mCCF-swap), and after the sequence was low pass filtered (nCCF-lpf). The CCF was normalized by dividing the CCF by the product of the torque and target standard deviations.

Tables 5.3 through 5.7 show representative results from sequences arranged with five iterations of the modified Hunter-Kearney stochastic swap routine. Each table summarizes the results for sequences arranged with the same swap routine. The columns list the amplitude distributions from best (at top) to worst (at bottom) for nCCF-swap, nCCF-lpf, and percent power.

For all five iterations of the stochastic swap routine, the GWN sequences had the normalized cross-correlation function (nCCF-swap) with the lowest magnitude, and the 2-uniform sequences had the worst nCCF-swap value immediately after the swap. This changed once the sequences were low-pass filtered, and now the results were not as consistent. The GWN sequences had the lowest nCCF-lpf value in only two out of the five swap routines, and the 2-

uniform sequence had the worst nCCF-lpf only once. The target velocity sequences had to be low-pass filtered before they could be integrated into a position sequence and used in the VOR experiments, so the nCCF-lpf value is more relevant to the VOR experiment than the nCCF-swap value. None of the amplitude distributions consistently resulted in the highest percent power. The uniform sequence had the highest percent power thrice, the 2 GWN once, and the GWN also once. The 2-uniform distribution, however, did have in the lowest percent power all five times.

None of the amplitude distributions was clearly the best for creating the best ACF or CCF results. There was no typical order for which the GWN, 2-GWN, and uniform sequences were ranked from best to worst, but the 2-unif sequences consistently had the worst percent power. Any one of the three distributions that gave good results (the GWN, the 2-GWN, and the uniform distributions) could be used to create acceptable target velocity sequences for the VOR experiments. The decision as to which one to use for the VOR experiments could be based on which sequence the subject tracks best.

3-element swap: ACF SSD not increased, CCF SSD decreased							
PDF	nCCF swap		PDF	nCCF lpf		PDF	% power
GWN	3.4		GWN	3.2		unif	32
unif	5.7		2-unif	4.1		GWN	27
2-GWN	7.0		unif	4.5		2-GWN	26
2-unif	12.4		2-GWN	5.6		2-unif	24

**Table 5.3:** 3-element swap; swap accepted when ACF SSD was not increased and CCF SSD decreased. “nCCF swap” represents the mean squared amplitude of the normalized cross-correlation function after the sequence was reordered, and “nCCF lpf” after the sequence was low-pass filtered. “% power” is the percent of the total sequence power that was in the desired frequency range.



2-element swap: ACF SSD decreased, CCF SSD decreased							
PDF	nCCF swap		PDF	nCCF lpf		PDF	% power
GWN	2.5		GWN	3.1		unif	33
unif	6.6		unif	4.7		2-GWN	31
2-GWN	7.7		2-unif	5.7		GWN	30
2-unif	12.6		2-GWN	6.4		2-unif	20

**Table 5.4:** 2-element swap; swap accepted when ACF SSD decreased and CCF SSD decreased.

2-element swap: ACF SSD not increased, CCF SSD decreased							
PDF	nCCF swap		PDF	nCCF lpf		PDF	% power
GWN	2.4		unif	2.8		unif	34
2-GWN	4.5		2-unif	3.1		2-GWN	33
unif	4.7		2-GWN	3.3		GWN	33
2-unif	11.2		GWN	3.4		2-unif	25

**Table 5.5:** 2-element swap; swap accepted when ACF SSD was not increased and CCF SSD decreased.

2-element swap: ACF SSD not increased, CCF not increased							
PDF	nCCF swap		PDF	nCCF lpf		PDF	% power
GWN	2.5		2-GWN	1.1		2-GWN	40
2-GWN	4.1		GWN	2.7		GWN	31
unif	4.5		unif	3.7		unif	30
2-unif	11.8		2-unif	6.9		2-unif	24

**Table 5.6:** 2-element swap; swap accepted when ACF SSD was not increased and CCF SSD was not increased.

2-element swap: ACF SSD decreased, CCF SSD decreased							
PDF	nCCF swap		PDF	nCCF lpf		PDF	% power
GWN	2.2		2-unif	2.6		GWN	40
2-GWN	4.0		GWN	2.7		2-GWN	31
unif	4.2		2-GWN	3.7		unif	30
2-unif	11.3		unif	6.9		2-unif	24

**Table 5.7:** 2-element swap; swap accepted when ACF SSD decreased and CCF SSD decreased.

## 5.5 Conclusions

The modified Hunter-Kearney technique can be used to create sequences with arbitrary spectral and cross-correlation characteristics. It is a powerful tool for creating input signals for multi-input and physiological systems, where each input may have to be of a different signal type in order for it to be tolerated well by the system under study. As a stochastic technique, the modified Hunter-Kearney process does not provide a closed form solution, but arranges the sequence to meet design goals. The ability of the algorithm to arrange the elements of the sequence to meet the desired spectral and correlation goals depends on the raw signal that is being manipulated, and how the stochastic swap is executed and evaluated.

Fifty-four velocity sequences, each 6000 elements long, and with four amplitude distributions, were arranged using several iterations of the modified Hunter-Kearney stochastic swap method. The goal was to investigate how effective the modified-Hunter Kearney technique was in creating sequences for the VOR experiment as a function of the number of elements involved in each swap, the decision criteria used to evaluate each swap, and the amplitude distribution of the sequence.

The two-element swap routine produced better results than the three-element swap routine. When the number of swaps executed was less than ten times the length of the sequence, the three-element swap routine often reduced the ACF and CCF SSD slightly more than the two-element swap routine, but the two-element routine surpassed the three-element routine when a greater number of swaps were executed. After a large number of swaps, however, the differences between the results from the two- and three-element routines were not dramatic. The most significant advantage of the two-element swap over the three-element swap was that it was more efficient. A single two-element swap required about a third of the processing of a single three-element swap.

How well the sequences were arranged to have desired characteristics was not significantly affected by the decision criterion used to evaluate and decide if a swap should be accepted and the elements repositioned. The differences observed between the four decision criteria was no greater than the variance observed when the same sequence was arranged several times using the same decision criterion.

Of the four amplitude distributions of the sequences, none stood out as the best for creating the characteristics for the VOR experiments. The results of the swap routines that manipulated the 2-uniform sequences consistently had the lowest percent power in the desired frequency range, and so the 2-uniform distribution was viewed as the least appropriate distribution. Any of the other three distributions were appropriate for generating sequences for the VOR experiment. This allowed for target velocity sequences to be created with a broader range of velocity distributions than if only one amplitude distribution had been found to be optimal.



## Chapter 6

# Dynamics of the human head-neck system in the horizontal plane

### 6.0 Chapter summary

Chapter 6 is a reprint of a paper that was published in *Annals of Biomedical Engineering*, Vol. 31, pp. 606-620, 2003. The study was conducted to determine the dynamics of the head-neck system in response to stochastic, horizontal plane, torque perturbations. The linearity of the head-neck system was important to being able to associate the VOR eye motions to the controlled torque input. If the system was strongly nonlinear, then nonlinear system identification techniques would have been required to accurately assess the performance of the VOR. Fortunately, it was found that the response of the head-neck system to stochastic perturbations could be well represented by a quasi-linear, second-order system. When the mean torque was held constant, and the disturbances to the head's position were small, the viscoelastic properties of the system were not dependent on the stimulus frequency between 0.5 and 10 Hz. The viscoelastic properties were, however, sensitive to the mean torque. As the mean torque increased, both the effective stiffness and damping of the system increased, and did so in a manner that maintained the system's damping ratio nearly constant.



## Dynamics of the Human Head-Neck System in the Horizontal Plane: Joint Properties with Respect to a Static Torque

JAMES L. TANGORRA, LYNETTE A. JONES, and IAN W. HUNTER

Bioinstrumentation Laboratory, Department of Mechanical Engineering, Massachusetts Institute of Technology,  
77 Massachusetts Avenue, Cambridge, MA

(Received 25 February 2002; accepted 14 February 2003)

**Abstract**—The vestibular system has often been studied by perturbing the position of the head. This study was conducted to identify the dynamic properties of the head-neck system in response to horizontal plane perturbations. A quasilinear approach was used to quantify the dynamics of the head-neck system at different levels of static torque. An operating point was established by applying a static torque to the head with a helmet-based perturber. The head-neck dynamics were then probed with a rich spectrum, stochastic, torque perturbation. Impulse response functions (IRFs) were estimated from correlation measures, and parametric models were fit to the IRFs. The results indicated that when the mean torque was held constant, the head-neck system behaved like a second-order, underdamped, passive system between 0.5 and 10.0 Hz. The system was not strictly linear, however. The properties of the system were sensitive to the static component of the torque. As the mean torque increased, the effective stiffness and damping progressively increased, and did so such that the system's damping ratio remained essentially constant. The findings of the study will assist in designing stimuli that are well tolerated by subjects and can induce head motions that span the performance capabilities of the vestibular system. © 2003 Biomedical Engineering Society. [DOI: 10.1114/1.1566772]

**Keywords**—Helmet, Impulse response, Stochastic, System identification, Vestibular, VOR.

### INTRODUCTION

The vestibular ocular reflex (VOR) stabilizes vision by creating eye motions that compensate for undesired linear and rotational head motions. The head serves as a relatively stable reference platform for the eyes, but is often disturbed during normal activities so that without a properly functioning VOR, people find it difficult to read, focus, or track moving objects. We are developing a testing method that evaluates the system-level performance of the horizontal, angular VOR while it stabilizes gaze during normal target tracking. Small amplitude, stochastic, torque perturbations are applied to a subject's head while the subject tracks an unpredictable, moving

target with natural head and eye motions. The torque perturbations are designed to disturb the motion of the subject's head and to make it necessary for the VOR to generate corrective eye motions so that the subject's gaze can remain focused on the moving target. The amplitude of the perturbation applied to the head must be large enough to induce a measurable VOR eye response, but not so large that the test subjects are prevented from moving their heads as they desire while tracking the target. In order to determine the appropriate amplitude for this perturbation, the dynamic relationship between the torque applied to the head and the resultant motion of the head and neck must be quantified.

The properties and the response of the head-neck system can be expected to have a nonlinear, time variant relation with the torque stimulus applied to the head, and will be sensitive to the conditions associated with target tracking, such as head velocity and gaze error. The stability of the head is controlled using a combination of voluntary control, neuromuscular reflexes, and passive system mechanics. The contribution of each of these mechanisms to stability is frequency dependent, with each mechanism dominating in a different frequency range.<sup>16</sup> It seems reasonable to expect that the properties of the head-neck system will change as the mechanisms that contribute to head stability increase or decrease their influence. Studies of the head-neck system have shown that properties such as stiffness are sensitive to external loading, as is the case for many joints in the human body.<sup>1,18,30</sup> Keshner *et al.*<sup>15</sup> concluded that the stiffness of the head-neck system was neurally mediated, primarily through a modulation in the gain of the vestibulo-collic reflex (VCR), which helps to maintain similar head dynamics under different inertial conditions. They found that as the inertial load on the head was nearly doubled, the effective stiffness of the head-neck system increased considerably, but the dynamics of the system, as measured by the system's natural frequency and settling time, remained essentially the same. Viviani and Berthoz<sup>28</sup> had also shown that the stiffness of the head-

Address correspondence to: James Tangorra, 77 Massachusetts Ave,  
3-147, Cambridge, MA, 02139. Electronic mail: jtangor@mit.edu

neck system was sensitive to the level of continuous force applied in the sagittal plane, but contrary to Keshner *et al.*,<sup>15</sup> found that significant changes in system dynamics occurred as the viscoelastic properties of the head-neck system changed.


This paper focuses on the response of the head-neck system in the horizontal plane to torque perturbations applied directly to the head. A quasilinear testing and analysis approach was used to quantify the dynamics of the head-neck system at different levels of mean (static) torque. An operating point for the system was established by applying a static torque to the head. The system dynamics was then probed by disturbing the head-neck system with a rich spectrum, stochastic, torque disturbance added on top of the static torque. The impulse response function (IRF) for the head-neck system was calculated at each level of static torque, and a parametric model was fitted to the nonparametric IRF. The focus of these experiments was on the overall response of the head under different loading conditions, not on the individual contributions made by different mechanisms towards the head-neck system's stability. The contributions of all mechanisms were lumped together, thus the parameters calculated for the system represent the combined contributions of all the mechanisms participating in the head and neck's response.

Active and passive stimulation of the head has been gaining popularity as a means of investigating vestibular responses,<sup>4,6,10</sup> but little has been published that characterizes the dynamic response of the head in response to forces and torques applied in the horizontal plane. The aviation and automotive industries have extensive data on the head's response during crashes and in response to vertical accelerations, but these motions are generally considered for the sagittal plane, and for much greater loads and larger motions than are needed to probe the VOR. Studies that have looked at the head's response in the horizontal plane often use full-body motion as the input, rather than a force or torque.<sup>16,26</sup> Full-body motions are a more natural input to the head than are directly applied torques, but may be better suited for stability studies than for identification of the system, as it is difficult to quantify the magnitude of the input that is transferred to the head. The present study will assist in the design of unpredictable torque sequences that can be used to stimulate the head for studies of the vestibular system.

#### METHODS: APPARATUS, EXPERIMENTATION, AND ANALYSIS

##### *Apparatus*

A head perturbing mechanism was developed to apply torque disturbances to a subject's head that induced small head rotations about the vertical neck axis. The head



**FIGURE 1.** Head perturber designed to apply torque perturbations in the horizontal plane to a subject's head. The perturber is composed of an adjustable helmet driven by a computer controlled torque motor.

perturber was constructed from an adjustable helmet and servomotor that were attached to a counterbalanced support frame as shown in Fig. 1. The frame was designed so that the helmet and motor could be lowered onto the head of a seated subject, and adjusted so that the rotational axis of the motor could be aligned with the natural, vertical neck axis of the subject. When the perturber was properly fitted to a subject, and no torque was applied to the helmet, the subject could freely rotate his or her head in the horizontal plane (yaw), but had little freedom to rotate the head in the sagittal (pitch) or frontal (roll) planes, or to translate the head in any direction. Weights on the counterbalanced frame were adjusted to balance the weight of the helmet and motor so that the subject did not have to support the weight of the apparatus. This was both for the comfort of the subject and to ensure that the weight of the apparatus did not tire the subject nor affect the dynamics of the head and neck.

The motor that drove the helmet was a static, brushless, resolver-based servomotor (Kollmorgen MT304A1, Radford, VA) that was configured for torque control. The motor output was driven with an analog command signal produced by a National Instrument's PCI-MIO-16XE10 data acquisition and output (NI-DAQ) board (National Instruments Corp., Austin, TX). To prevent the head perturber from applying too much torque to the subject's head, an electronic current limiter in the motor controller was set to limit the amount of current that could be applied to the motor's coil. A mechanical slip clutch



(Hilliard L2-1-312AC, Elmira, NY), placed between the motor's output shaft and the helmet, acted as a secondary safety system in case the electronic limiter failed. The slip clutch was set to prevent more than 1.0 N m from being applied to the helmet by the motor. The torque applied to the helmet was measured with a six axis, force/torque transducer (Assurance Technologies Inc., Garner, NC) capable of resolving torques as low as  $2.0 \times 10^{-3}$  N m and measuring torque up to 4.0 N m. A laser pointer was mounted on top of the helmet to give visual feedback of head position. The laser projected onto a screen placed 1.6 m in front of the subject. Head and helmet motions were measured with either small (6 g, 12 mm diameter), magneto-hydrodynamic (MHD) angular rate sensors (ATA ARS-04E, Albuquerque, NM), or slightly larger (14 g, 20 mm), piezoelectric "tuning fork" angular rate sensors (Watson ARS-C141-1AR, Eau Claire, WI). The ATA MHD sensor was mounted on the bridge of the subject's nose with a thin aluminum brace, and the Watson sensor was mounted on a bite plate. Since the helmet perturber restricted movement of the head to rotations about the neck's primary axis, only the rotational velocity of the head was measured. Data were gathered with the same NI-DAQ board used to create the motor's reference signal. Data were analyzed offline with Mathcad 2001 and with software developed in the Visual Studio 6.0 environment.

### *Experimentation*

Six subjects, one woman and five men, participated in this study. The subjects' ages ranged from 22 to 33. All were healthy graduate students with no history of or apparent neuromuscular or vestibular problems, and had normal or corrected-to-normal vision. No quantitative test was conducted to assess the subjects' vestibular performance before assessing their head-neck response, but subjects were questioned by the experimenter and excluded if they reported any present or past neck, back, or vestibular problems. All subjects had a range of motion with their neck that exceeded the rotational displacements that would be experienced during the experiments. Three potential subjects were excluded: two had previous back-related injuries, and one had suffered a head injury within the past year. This study was approved by the Massachusetts Institute of Technology Committee on the Use of Humans as Experimental Subjects. All subjects gave written consent before participating in this research study.

Subjects were seated in a rigid chair just under the helmet of the head perturber. The helmet was lowered onto the subject's head and tightened with head and chin straps so that it did not slip when torque was applied or when the head was moved. The primary axis of the subject's neck was not explicitly measured, but the test-

ing apparatus was adjusted for each subject so that the head and the helmet rotated without there being any significant eccentricity between the neck's and the helmet's rotational axes. There were independent adjustments for the lateral and horizontal positioning of the helmet, the angle of the motor axis in the vertical plane, the helmet's attachment and alignment to the motor, and a variety of adjustments within the helmet allowed the helmet to be securely attached to the subject's head. Subjects were instructed to adjust the head perturber so that they felt that their head rotated comfortably with the helmet, and so that they were able to rotate their heads with the helmet through  $\pm 30$  deg without feeling that the primary axis of their neck was eccentric to the axis of the head perturber. This was verified visually by the experimenter as the subjects rotated their heads. During the experiment, the torque perturbations rarely created disturbances to the head's position that exceeded 5 deg, so any translation that would have been caused by a deviation between the rotational axis of the helmet and the subject's head was minimal. The compliance of a subject's skin did permit the helmet to move slightly relative to the skull. This made it necessary to measure both the helmet and head velocities during the experiments, and to calculate the resultant torque that acted on the head, which differed slightly from that measured at the helmet. Once the helmet was adjusted, subjects were asked to rotate their head to a natural "zero-rotation" position. The laser pointer on top of the helmet was illuminated, and its projection marked on the screen as a visual reference. During experimentation the projection of the laser relative to this mark acted as a visual reference so subjects would know where their head was pointing.

Each experimental trial began by applying a static (constant) torque to the helmet, which tended to displace the subject's head away from the zero-rotation position. Subjects were instructed to compensate for the torque, and to rotate their head so that the laser projection was realigned with the zero-rotation mark on the screen. Once the subject was accustomed to compensating for the static torque, a small amplitude, zero-mean, stochastic torque perturbation sequence was added to the static torque. This perturbation shook the test subject's head about the zero-rotation point, effectively acting as a disturbance to a system defined at an operating point established by the static torque. Subjects were instructed to correct for any position disturbance and to try and keep the laser pointer aligned with the zero-rotation marker. It was stressed to each subject that they should not resist the torque perturbations by cocontracting their neck muscles, which has been shown to increase the stiffness of a joint,<sup>14</sup> but that they should actively correct for any position error. Training runs were conducted to try to illustrate to subjects the difference between actively cor-

recting the head's position and resisting the disturbance with cocontraction. All subjects kept their eyes open during the experiment.

The static torque applied to the helmet ranged from 0.0 to 0.65 N m and was applied in increments of 0.05 N m in both the clockwise and counter-clockwise directions. Although the neck is capable of producing rotational torques well above 0.65 N m, 0.65 N m was used as an upper limit in the interest of safety, and to ensure that all subjects would be able to resist the torque during the experiment without tiring. The stochastic torque perturbation sequence was generated with a 12-seed pseudorandom binary sequence (PRBS), executed at 15 or 20 elements per second, and scaled for a magnitude of  $\pm 0.1$  or 0.2 N m. This execution rate produced sufficient input power through 10.0 Hz. As the predominant frequencies of head motions in the horizontal plane during locomotion are in the range of 0.5–3.0 Hz,<sup>9</sup> this input spectrum seemed sufficient to stimulate the head motions necessary for system identification. The torque perturbation was gradually applied to the helmet by increasing the level from 0.0 to 0.2 N m, and in a few experiments was limited to 0.1 N m. The stochastic PRBS was chosen so that a subject would not be able to predict the sequence of torque perturbations being delivered and potentially develop a strategy to compensate for the disturbance. Subjects reported that although the perturbations were light and did not create large head motions, the randomness of the sequence made it very difficult to bring the laser to the zero position after being disturbed from the neutral position. By the time they were attempting to correct for a disturbance and had generated a motion towards the zero-rotation position, another perturbation would displace their head before the correction could be made. Experimental trials were run for 60–90 s. To help prevent fatigue, rest periods of equal duration were given between each trial, 10 min breaks were taken after every six experimental trials, and the direction in which the static torque was applied alternated between clockwise and counter-clockwise directions. Additionally, the order in which the static torques were applied to the head was randomized, so that the subject would not get accustomed to a continuously increasing level of torque, and to balance low and high levels of torque throughout the experiment. A total of 150 experimental trials were conducted on the six subjects.

### Analysis

The torque applied to the helmet (tor), the helmet's rotational velocity (helm), and the head's rotational velocity (vel) were each sampled at 100 Hz and analyzed offline. Data were median filtered to remove aberrant data points and the input and output power spectral den-

sities were calculated to verify that appropriate power was available throughout all frequencies of interest.

The data were analyzed as representing measurements of a quasilinear system that had been disturbed about an equilibrium operating point. The system's equilibrium point and its associated dynamics were set by the system's reaction to the applied static torque, and the disturbance about this operating point was created by the torque perturbations. The admittance IRF was calculated for each experiment, and a second-order parametric model was fit to the IRF. Admittance is being defined as the relationship between a torque input and a velocity output, the mechanical analog to the electrical admittance, the relationship between a voltage input and a current output. The admittance IRF is simply the first derivative of the more common compliance IRF, which relates the position output to the torque input. The head and helmet were instrumented with angular rate sensors, so the admittance IRF was the natural output-input relationship to calculate.

The admittance IRF estimate,  $h_{est}$ , was calculated by deconvolving the input (torque) auto-covariance estimator from the output-input (head velocity–torque) cross-covariance estimator. Sequence means are written as  $\mu$ .

Biased auto-covariance estimator

$$c_{tt}(j) = \frac{1}{N} \times \sum_{i=0}^{N-1-j} [\text{tor}(i+j) - \mu_{\text{tor}}] \times [\text{tor}(i) - \mu_{\text{tor}}]. \quad (1)$$

Biased cross-covariance estimator

$$c_{vt}(j) = \frac{1}{N} \times \sum_{i=0}^{N-1-j} [\text{vel}(i+j) - \mu_{\text{vel}}] \times [\text{tor}(i) - \mu_{\text{tor}}]. \quad (2)$$

Identification of the IRF from the covariance estimators, rather than directly from input and output measurements, gives a more accurate estimate of the dynamics of the underlying system by reducing the effect of measurement noise.<sup>12,20</sup> The linear relation between the covariance functions and the system IRF is

$$c_{vt}(n) = c_{Nt}(n) + \sum_{k=0}^n h_{est}(k) \times c_{tt}(n-k), \quad (3)$$

where  $c_{Nt}$  represents the cross-covariance between the measurement noise and the input torque. When the measurement noise and input torque are statistically uncorrelated, the noise-torque cross-covariance may be set to zero. The underlying system response, uncorrupted by noise, therefore represents the relation between the input covariance and the output-input cross-covariance.

The input auto-covariance was “deconvolved” from the cross-covariance by first expressing  $m$  lags of the input auto-covariance as an  $(m \times m)$  Toeplitz matrix,  $\Phi$ , where the  $(i, j)$ th term was set equal to the covariance at lag  $|i - j|$ . The Toeplitz matrix was inverted and multiplied by the  $(m \times 1)$  output-input cross-covariance estimator. The resulting IRF estimate was then scaled by the sampling interval,  $\Delta t$ , to make its amplitude consistent with the originally sampled continuous time signal. The maximum lag value,  $m$ , was iteratively chosen to be approximately six times the apparent settling time of the IRF, so that all relevant system dynamics would be included in the calculations.

Admittance impulse response function estimate

$$h_{\text{est}} = \frac{1}{\Delta t} \times \Phi^{-1} \times c_{vt}. \quad (4)$$

The coherence-squared function was calculated from the auto- and cross-spectral densities to assess the significance of linearity and noise in the analysis, and to identify the frequency range over which the linear IRF represented well the dynamics of the underlying head and neck system about the operating point.

The admittance IRF had the characteristic shape of a linear, underdamped, second-order system (Fig. 2), and was well approximated by the second-order parametric model

$$h_{\text{mod}} = \text{gain} \times \omega_n^2 \times e^{-\zeta \times \omega_n \times t} \times \left[ \cos(\sqrt{1 - \zeta^2} \times \omega_n \times t) - \frac{\zeta}{\sqrt{1 - \zeta^2}} \times \sin(\sqrt{1 - \zeta^2} \times \omega_n \times t) \right], \quad (5)$$


where  $\omega_n$  is the undamped natural frequency, and  $\zeta$  is the damping ratio.

This velocity/torque IRF is the first derivative of the more familiar position/torque IRF

$$h_{\text{mod}} = \text{gain} \times \omega_n \times e^{-\zeta \times \omega_n \times t} \times \frac{1}{\sqrt{1 - \zeta^2}} \times \sin(\sqrt{1 - \zeta^2} \times \omega_n \times t). \quad (6)$$

The rotational inertia ( $J$ ), damping ( $B$ ), and stiffness ( $K$ ) of the head and neck are related to the model's gain, natural frequency, and damping ratio by

$$J = \frac{1}{\text{gain} \times \omega_n}, \quad B = \frac{2\zeta}{\text{gain} \times \omega_n}, \quad K = \frac{1}{\text{gain}}. \quad (7)$$



**FIGURE 2.** Admittance impulse response function estimates (solid) and their second-order parametric models (dashed) obtained from two subjects. The data represent the impulse relation between a torque input and a velocity output. The characteristics of the IRF shown in plot (a) are very similar to that of the ideal second-order model. The IRF shown in plot (b) begins at a low value, uncharacteristic of the ideal second-order response, but then quickly rises to a value close to the time=0 value of the IRF in plot (a). After this initial spike, the IRF has a similar oscillatory, exponential decay as the ideal second-order model.

The second-order parametric models were fit to the IRF by minimizing the mean sum of square difference between the IRF and the parametric model with a nonlinear, Levenberg–Marquardt, minimization technique.<sup>24</sup>


Estimates of the head velocity measured during each experiment were made by convolving the measured input torque with the nonparametric and parametric IRFs. The “percent variance accounted for” (VAF) and the mean square error of these estimates were calculated to assess how well the IRF captured the dynamics of the underlying head-neck system, and to compare how much of the dynamic response was lost by modeling the nonparametric IRF with the three-parameter second-order model.

To determine the characteristics of the head-neck system, the influence of the head perturber must be removed. Ideally the analysis of the admittance IRF of the head and neck would be made using a direct measurement of the torque applied to the head. Unfortunately, we are limited to measuring the torque applied to the helmet

and estimating the torque that is transferred to the head. These torques, although close in value, differ because of the dynamic response of the helmet, and the compliance of the skin, which permits the helmet to have a slightly different motion from the head. The analysis began by modeling the head perturber components located after the torque transducer, the helmet and its connecting shaft, as a first-order rotational mass and damper, the compliance of the skin between the helmet and the head as a spring and damper, and the head and neck system as an unknown dynamic system. Because the skin between the head and helmet is being modeled without inertia, the IRF of the head and neck can be calculated from a measurement of the torque applied to the helmet minus the “dynamic torque” used by the helmet. The helmet’s dynamic torque is estimated by convolving a measurement of the helmet’s velocity with an *a priori* IRF that represents the velocity-torque relation of the head perturber.

This calculation can be simplified when the difference between the velocities measured for the head and for the helmet are small enough to be neglected, and if the head-neck system is treated as a linear, single mass, lumped parameter system. When the measured velocities are set equal, a single IRF representing the response of the combined helmet and head-neck system can be found using the torque measured by the transducer at the helmet. Parameters representing the properties of the head-neck system can be found by subtracting the known mass and damping of the head perturber from the parameters that describe the combined system.

Significant differences were not found between trends in the parameters estimated when the head-neck IRFs were calculated by removing the dynamic torque or when the parameters were calculated by subtracting the perturber’s parameters from the parameters describing the combined system. There were differences in the specific values of the parameters found with the two calculation methods, but the differences were small and consistent, and the general trends exhibited by the parameters for a series of experiments were the same. Because of intersubject variability, and variability in an individual subject’s response, particular values for a property at a specific mean torque are less significant than the trend describing how the property changed as the mean torque applied to the head was varied. This information can be gathered with either the analysis that removes the helmet’s dynamic torque, or the analysis that simply subtracts the head perturber’s mass and damping. For consistency, the results reported throughout this paper are those found without removing the dynamic helmet torque. This allowed the inclusion of data for cases where a sensor was not available to measure independently the helmet’s velocity.



**FIGURE 3.** Coherence-squared (a) and input power spectrum (b) from a representative subject. The drop in the coherence squared function below 0.5 and above 10 Hz is an indication that nonlinearities and/or a low signal to noise ratio are adversely affecting the linear input-output analysis at these frequencies.

## RESULTS

### *Nonparametric IRFs*


The data displayed in Fig. 2 are typical of the nonparametric admittance calculated from the experimental measurements. Both curves show that in response to an impulse input, the head-neck system exhibits exponentially decaying oscillations. In some cases the IRF decayed from a high initial value, but in others the IRF began with a quick rise from a low initial value that was followed by an oscillatory decay. Both curve shapes were found during the analysis of data from each subject, and neither was more prevalent over any particular range of static torques. Although the head-neck system possesses nonlinear, time variant, adaptive qualities, coherence measures and the VAF in velocity predictions showed that, under the constraints of the experiment, the quasilinear assumption for the system and analysis was sound, and the resulting nonparametric IRFs accurately represented the dynamics of the head-neck system at the different levels of mean torque.

Figure 3 shows a representative coherence-squared function from the linear analysis of the system's input and output measurements. In general, coherence values were high (above 0.9) at frequencies from as low as 0.5 to 10 Hz. This indicates that, at least in this frequency range, the linearity assumption was valid, and the noise content of the experimental data was low.<sup>20,21</sup> The high coherence-squared is evidence that a single IRF represents the system well from 0.5 to 10 Hz. Below 0.5, and above 10 Hz, the coherence drops rapidly, indicating that at these frequencies, either the system would have been better modeled with higher order, nonlinear functionals,<sup>17</sup> or that the signal-to-noise ratio of the measurements was too low to assess system performance accurately. The drop in coherence above 10 Hz can be attributed to a drop in the input power at higher frequencies. As can be seen in Fig. 3, the spectral density of the input torque dropped significantly by 10 Hz. Coherence is sensitive to the signal to noise ratio, and was adversely affected by the nearly two orders of magnitude drop in power by 10 Hz.

The nonparametric IRFs were convolved with the measured input torque sequences to create estimates of the head velocity. The resulting calculations matched the measured head velocity well, indicating that the nonparametric IRFs captured the dynamics of the head-neck system. The mean VAF for the velocity estimates for all 150 experimental trials was 92.2% [standard deviation (s.d.) 2.4%]. In the majority of instances, the same velocity sequence that was used to calculate the cross covariance function, and subsequently the IRF estimate, was used in the calculation of the VAF. In only a few instances were data records long enough to be split so that a velocity prediction could be made with data that were not also used to estimate the system's IRF.

*Modeling the Head-Neck System as an Underdamped, Second-Order, Linear System*

The IRF and the frequency response functions' (FRF) gain and phase plots showed that, at all levels of static torque tested, the head-neck system exhibited characteristics of a second-order underdamped system. The fundamental features of the nonparametric IRFs were consistent with the characteristics of ideal second-order models, with visually similar oscillatory frequencies, decay rates, and settling times. As would be expected, the nonparametric IRFs were not as smooth as the parametric IRFs, nor were the decays of the nonparametric IRFs ideally exponential, but the same fundamental dynamics could be seen in the nonparametric and in the second-order parametric IRFs. The spike seen in the initial response of some of the nonparametric IRFs deviated from the ideal, but even these IRFs quickly returned to a more idealized oscillatory, exponential decay after the spike.



**FIGURE 4.** Gain (a) and phase (b) plots of the frequency response function calculated from a single subject. These show the characteristic traits of a second-order response from the perspective of a torque input and a velocity output.

In the frequency domain, gain and phase plots calculated from experimental data also supported modeling the nonparametric IRFs with second-order, parametric models. The data displayed in Fig. 4 are typical gain and phase results, and exhibit the traits of an ideal second-order system when considered from the perspective of a velocity output. The experimental gain response can be seen to increase with frequency with a slope near 1 until reaching a resonance peak between 2 and 3 Hz. After the resonance peak, the gain drops with a slope near  $-1$ . The phase begins near 90 deg, decreases and crosses 0 deg near the frequency of the resonance peak, and then approaches  $-90$  in the decades after the resonance peak. Although not exactly the same as that for an ideal

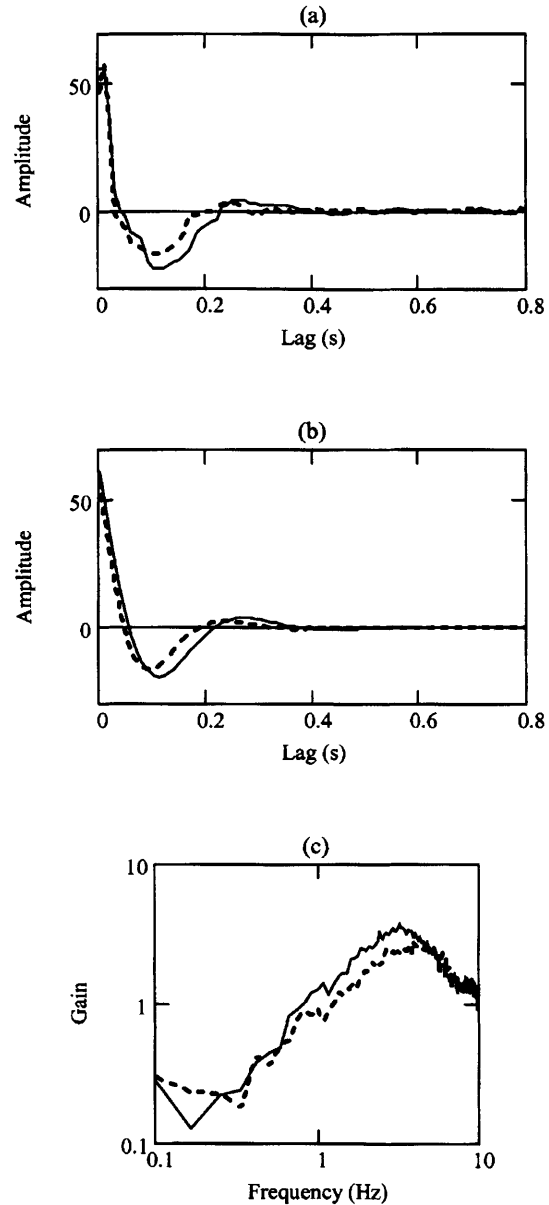
second-order system, the gain and phase plots strongly suggest that second-order modeling of the head-neck is appropriate.

Second-order parametric models were fitted to the nonparametric IRFs by minimizing the mean square error with a nonlinear Levenberg–Marquardt minimization scheme (Mathcad 2001, Mathsoft Inc., Cambridge, MA).<sup>16</sup> As with the nonparametric IRFs, the second-order parametric IRFs were used to calculate estimates of experimental velocities by convolving the model IRF with measured input torque sequences. The mean VAF for the all 150 experiments was 86.4% (s.d. 3.9%) whereas when the velocity was calculated using the nonparametric IRF the mean VAF was 92.2% (s.d. 2.4%). The parametric outcomes are exceptional considering that the parametric model is described by only three terms, whereas the nonparametric IRF is defined by a value at each sampling interval over its entire length, which is over 200 terms in this analysis. The high VAF indicates that the second-order models captured the dynamics of the system, and that the model parameters were accurate descriptors of head and neck properties.

#### *Second-Order Parameters versus Mean Torque*

Significant changes in the response of the head-neck system occurred as the static torque applied to the head was varied from 0.0 to 0.65 N m. The amount the subject's head was disturbed in response to the perturbations was noticeably smaller at the higher levels of static torque than at the lower levels, implying that the effective stiffness of the head-neck joint had increased as the static torque increased. Changes in the system's dynamics were apparent in the IRFs and FRFs, but were most readily seen, and best quantified, by the changes in the parameters of the second-order IRFs.

*Representative Data from One Test Subject.* The results from one test subject, shown for static torque levels of 0.10 and 0.55 N m, are displayed in Fig. 5. The changes illustrated by these IRF and FRF curves are representative of the changes found in other subjects. In each graph the curves have a similar shape, and if separated from a time and amplitude scale, would appear nearly identical. However, the differences in amplitude, and crossover values, indicate that although the system remained fundamentally the same, its speed and magnitude of response were a function of the static torque applied to the head. It is apparent from the IRF that the system tested at 0.55 N m had a higher natural oscillation frequency and settled more quickly than when the input torque was 0.10 N m. From the FRF gain plot, the system is clearly less compliant at the higher level of static torque. The parameters of the second-order models that were fitted to the nonparametric IRFs support these observations.



**FIGURE 5.** Impulse response functions (a), second-order parametric models (b), and the frequency response function gain response (c) of trials conducted at 0.10 N m (solid) and 0.55 N m (dashed).

When the mean torque was increased from 0.10 to 0.55 N m, the calculated natural frequency increased from 3.6 to 4.3 Hz, and the settling time decreased from 0.38 to 0.33 s. Stiffness increased from 8.06 to 12.46 N m/rad and damping increased from 0.33 to 0.43 N m s/rad. As expected the calculated rotational inertia remained virtually unchanged [0.012–0.013 kg m<sup>2</sup> (N m s<sup>2</sup>/rad)] across the static torque levels tested.




FIGURE 6. Rotational inertia data (crosses and pluses) and mean rotational inertia (lines) from two subjects.

Although there was significant intersubject variability in the values of the parameters calculated at each torque level, the same general trends held for each subject. In general, stiffness ( $K$ ), damping ( $B$ ), and natural frequency ( $\omega_n$ ) all increased with increasing mean torque, settling time ( $t_s$ ) decreased, and the rotational inertia ( $J$ ) and damping ratio ( $\zeta$ ) remained essentially constant over the range of mean torques used.

*Parameter Relations with Mean Torque. Rotational Inertia ( $J$ ,  $\text{kg m}^2$ ).* For each subject, the rotational inertia remained nearly constant over the range of static torques tested. This was expected since the torque applied to the head should have had little or no effect on the rotational mass of the head and neck. Small deviations in head and neck inertia could occur if the test subject's head was positioned differently between tests, or could result from changes in blood volume and muscle mass that might be required to compensate for different levels of torque, but a subject's head and neck inertia should remain nearly the same from test to test. For each subject, the calculated rotational inertia did vary from test to test, but the variance was small, and there was no trend in the data to indicate that the rotational inertia had changed in a predictable manner with an increase in mean torque (Fig. 6). The rotational inertia data were best modeled by their mean value. For each subject, the mean rotational inertia fell between 0.011 and 0.017  $\text{kg m}^2$ , with variances ranging from  $0.75 \times 10^{-6}$  to  $2.3 \times 10^{-6}$ .

*Damping ( $B$ ,  $\text{N m s rad}^{-1}$ ).* For all subjects, the damping parameter,  $B$  ( $\text{N m s rad}^{-1}$ ), of the system increased as the static torque applied to the head increased. Across the range of 0.0–0.65 N m, the smallest mean change in damping for a subject was  $0.2 \text{ N m s rad}^{-1} (\text{N m})^{-1}$ , from a minimum of  $0.25$  to a maximum of  $0.38 \text{ N m s rad}^{-1}$ , and the greatest mean increase was  $0.6 \text{ N m s rad}^{-1} (\text{N m})^{-1}$ , from  $0.31$  to  $0.70 \text{ N m s rad}^{-1}$ . These values did not necessarily occur at the lowest and




FIGURE 7. Damping data from two subjects shown with first-order polynomial models fitted independently to data at positive and negative mean torque values.

highest static torque values. These data are shown in Fig. 7. The data for each subject were fitted by first and second-order polynomial models to determine which model best represented the data with respect to mean square error.

Physiologically, the head and neck system is symmetric, and should exhibit a similar response to a torque applied in either the clockwise or counter-clockwise direction. Therefore, models were applied that allowed for, and assessed, symmetry about zero mean torque. First-order polynomials were fitted to the data from the positive mean torque and negative mean torque,  $B = c \times \text{tor} + d$ , separately. The two sides were fitted independently to identify if, in fact, the system dynamics changed symmetrically about zero static torque. The second-order polynomial,  $B = a \times \text{tor}^2 + b \times \text{tor} + c$ , was fitted to all data points. Its first-order coefficient allowed it to bias how quickly the model increased with positive or negative torques.

The pairs of first-order models and the second-order models both fit the data reasonably well. On average, the first-order models fit the data with a slightly lower mean square error than did the second-order polynomials. The pairs of first-order models were not perfectly symmetric. Typically there was a slight offset between the zero torque values for the negative and positive sides of the model, and one side had a slightly greater slope than the other as shown in Fig. 7. There was no particular bias, however, when the models for all test subjects were considered, as to whether the higher slope or higher zero offset occurred in the direction of clockwise or counter-clockwise neck rotation. The slope, which is the proportionality constant between the system damping and the

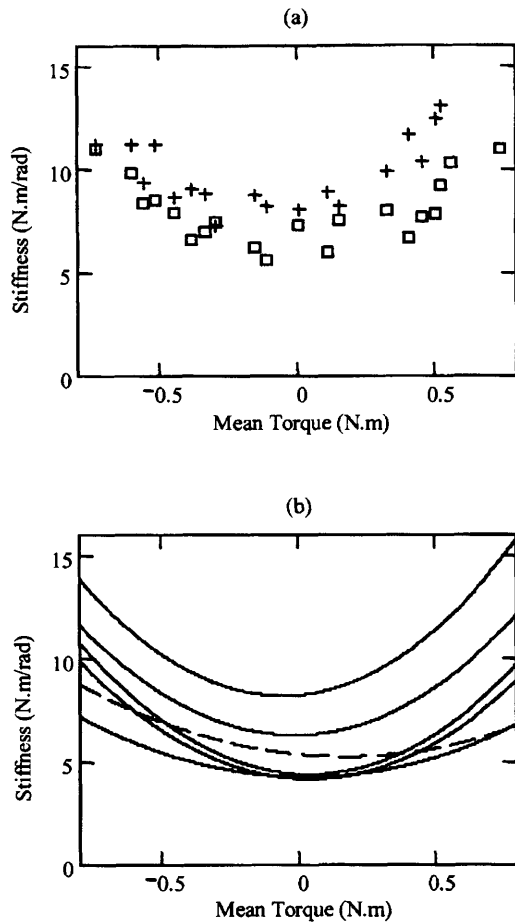


FIGURE 8. Stiffness data for two subjects (a) shown with second-order polynomial models (b) for all six data sets.

mean torque, ranged from a low of 0.07 to a high of  $0.27 \text{ N m s rad}^{-1}$  per N m. The results of the second-order models were consistent with those from the first-order models. The models were nearly symmetric, but each had a slight offset towards one side.

**Stiffness ( $K$ ,  $\text{N m rad}^{-1}$ ).** The stiffness of the head-neck system increased as the static torque was increased. The relation between stiffness and torque is shown in Fig. 8. Intersubject variability was considerable in terms of the absolute values for stiffness, but the same trend held for the relation between stiffness and mean torque for all subjects. Typically, the stiffness seemed flat, almost constant, around static torque values near zero, but as the static torque increased beyond 0.25 N m, the stiffness increased rapidly. The change did not appear to be a linear, but looked very much like a second-order polynomial. The subject with the smallest change in stiffness over the 0.0–0.65 N m torque range, showed an increase




FIGURE 9. Undamped natural frequency data for two subjects.

from  $3.7$  to  $6.7 \text{ N m rad}^{-1}$ , with a mean of  $4.7 \text{ N m rad}^{-1}$  per N m, whereas the subject with the greatest change in stiffness showed an increase from  $7.2$  to  $13.0 \text{ N m rad}^{-1}$ , with a mean of  $8.9 \text{ N m rad}^{-1}$  per N m. Both increases represented changes of over 80% of the stiffness at zero static torque. It can be seen in Fig. 8 that the change in stiffness was nearly symmetric about zero. For all but one subject, second-order polynomial models fitted the data with lower mean square error than did pairs of first-order models.

**Natural Frequency ( $\omega_n$ , Hz).** For most subjects (5 of 6) the head-neck system's natural frequency tended to increase with an increase in the static torque applied to the head. The natural frequency values did not rise smoothly with incremental increases in the static torque, but varied about a local mean that increased with the applied torque. This trend can be seen in Fig. 9. The one test subject who deviated from this pattern had a natural frequency that seemed to vary between high and low values across the entire range of negative and positive static torques. Although the increase in natural frequency was not smooth, there was a significant change in the values over the relatively small range of torques applied to the head. The natural frequency for five of the six test subjects was between 2.0 and 2.5 Hz (12.6 and 15.7 rad/s) at 0 N m, and by 0.65 N m it had reached between 2.9 and 4.0 Hz (18.2 and 25.1 rad/s).

**Damping Ratio ( $\zeta$ ).** There was little change in the subjects' damping ratios across the range of torques tested. The data did not visually exhibit any consistent change with mean torque, and first- and second-order polynomials fit to the data modeled the data as essentially unchanging with mean torque. The mean values for the subjects ranged from 0.42 to 0.64. With the exception of



TABLE 1. Second-order damping ratio.

Subject	Mean	Variance
1	0.45	$0.7 \times 10^{-3}$
2	0.59	$4.9 \times 10^{-3}$
3	0.64	$15.0 \times 10^{-3}$
4	0.42	$0.8 \times 10^{-3}$
5	0.53	$4.7 \times 10^{-3}$
6	0.54	$3.5 \times 10^{-3}$

a few outliers in the 150 experimental trials, the damping ratio was always found to be less than 1, indicating that the head-neck system tends to respond to torque disturbances in an underdamped fashion. The mean and variance for each subject's damping ratio data are shown in Table 1.

## DISCUSSION

This study examined how the dynamics of the head-neck system changed as a function of the magnitude of static torque applied to the head. The results indicated that between 0.5 and 10.0 Hz, the head-neck system behaved like a quasilinear, passive system in response to small, stochastic torque inputs applied to the head in the horizontal plane. When the mean torque was held constant, and the system perturbed, the head-neck system responded much like a linear, second-order, underdamped system. The system was not, however, strictly linear. The properties of the head-neck system were sensitive to the static component of the input torque, and the characteristic impulse response function changed markedly as the mean torque level was changed. As the static torque was increased, the effective stiffness and damping of the system, which represented the combined effects of neuromuscular reflexes, voluntary control, and passive viscoelastic properties of muscle, progressively increased, and did so in a manner such that the system's second-order damping ratio remained essentially constant across a mean torque range of 0.0–0.65 N·m.

### *Single IRF Holds from 0.5 to 10.0 Hz for a Constant Level of Torque*

When the mean torque applied to the head was held constant, the behavior of the head-neck system could be represented by a single, linear IRF for frequencies between 0.5 and 10.0 Hz. High (above 0.9), relatively constant values of the coherence-squared function indicated that system nonlinearities and measurement noise were not significant, and that the linear covariance analysis and the resulting linear IRFs applied well throughout the 0.5–10.0 Hz range. The IRFs had the characteristics of, and were well modeled, by underdamped, second-order, parametric models. On average, the second-order models

captured 94% of the response described by the nonparametric IRFs (as measured by the ratio of the average variance accounted for calculated when the IRFs and models were used to predict the system's output), clearly indicating that the parameters of the second-order models were good measures of the head and neck's inertial and viscoelastic properties.

The physiology of the head and neck supports modeling the system's reaction to perturbations in the horizontal plane with a lumped parameter, damped, mass-spring model, which has a characteristic second-order response. When the head rotates in the horizontal plane, the skull and the first cervical vertebra, the atlas (*c1*), rotate about a pivot created by the dens of the second cervical vertebra, the axis (*c2*). Although a large displacement rotation of the head may require other cervical and thoracic vertebrae to move along their articular facets, and for intervertebral disks to twist, rotation of the head in the horizontal plane is predominantly manifested by a rotation of the skull about the axis.<sup>8,27</sup> This can be simplified as a rotational inertia moving on a single, vertical axis. Damping and stiffness can be attributed to the muscle, connective, and soft tissue surrounding the neck, the properties of which will change as more or fewer muscle fibers are involved in creating motion or resisting disturbance forces.

The head-neck system's behavior from 0.5 to 10.0 Hz was well described by a single linear, time-invariant model which indicated that when the mean level of torque was held constant, the viscoelastic properties of the head and neck did not change as a function of the stimulus frequency, but remained nearly constant. This is somewhat surprising considering that between 0.5 and 10.0 Hz head stability is not controlled by a single mechanism, but through a combination of voluntary control, neuromuscular reflexes, and passive system mechanics.<sup>23</sup> The contribution to stability from each of these mechanisms is frequency dependent, with each mechanism dominating stability in a different portion of the frequency range. Stability at low frequencies, around 1.0 Hz and below, is dominated by voluntary control, whereas stability at higher frequencies, 2.5–3.0 Hz and above, seems to be dominated by passive system mechanics. The reflexive mechanisms, namely the vestibulocollic reflex VCR and the cervicocollic reflex, have been found to be most significant between the low and high frequency ranges.<sup>16</sup> The expectation that the properties of the head-neck system would change with the frequency of the applied stimulus is based on the finding that the mechanical stiffness of muscle fibers increases with the frequency of a sinusoidal force applied to the muscle.<sup>25,28</sup> Despite the frequency-dependent nature of the mechanisms that control head stability, the present results indicated that when the mean load on the head remained constant, the control and stability mechanisms

functioned to create an overall system with physical properties that were relatively invariant between 0.5 and 10.0 Hz. This finding of invariance has been shown in other joints that were tested in a similar fashion. As this study found for head-neck responses in the horizontal plane, Viviani and Berthoz<sup>28</sup> found that the properties of the head-neck system in the sagittal plane remained nearly constant between 0.5 and 5.0 Hz, and could be modeled by a single, linear, second-order system. In their review of joint dynamics, Kearney and Hunter<sup>14</sup> concluded that "linear models provide excellent descriptions of joint dynamics provided that conditions (e.g., activation level, mean position) are maintained approximately constant throughout the experiment" (Ref. 14, p. 73).

The finding that the viscoelastic properties were invariant probably reflects the fact that the subjects' heads were not displaced through large rotations by the torque perturbation. Over the range of experimental conditions employed the head was rarely, if ever, displaced more than 5 deg. The stiffness of isolated muscle, as well as the overall stiffness of human joints, is known to vary with displacement.<sup>7,14</sup> If the head had been moved through a much larger range of motion, it is likely that the viscoelastic properties of the head-neck system would have varied, and the system would not have been well modeled by a constant parameter, linear model.

At frequencies below 0.5 and above 10.0 Hz, the characteristics of the system were not captured by the linear input-output analysis, and the system's behavior could not be described by constant parameter, second-order models. Low coherence-squared values indicated that at these frequencies, the analysis was corrupted by system non-linearities and/or measurement noise.<sup>12,21</sup> Above 10.0 Hz, it is evident that the drop in coherence was likely due to low input power, which reduced the signal to noise ratio. As can be seen in Fig. 3, near and above 10.0 Hz the drop in coherence mirrors the drop in input power. This drop in power was intentional, a result of the design of the input sequence which attempted to simulate the frequency of motions encountered in the natural environment. During locomotion the predominant frequencies of head perturbations in the horizontal plane are between 0.5 and 3.0 Hz.<sup>9</sup> The PRBS torque sequence applied to the head was therefore executed at only 15 or 20 elements per second to concentrate input power at lower frequencies. Input power remained high below 0.5 Hz, so the low coherence-squared measures at these frequencies were not caused by a low signal to noise ratio, but by the nonlinearity of the active control system that dominates the head's response at low frequencies. Below 0.5 Hz, a single torque stimulus will act on the head for over 2 s, and so subjects have the time to compensate actively for the applied torque and move the head back to the zero position. At higher frequencies, the subjects can see and feel that their head has been displaced, but

do not have the time to correct for the disturbance before the next perturbation acts on the head. Although the initial response to long duration perturbations may be the same as that for the short duration perturbations, the linear IRFs and the passive, second-order models do not model the active correction that occurs after a time delay.

#### *Stiffness Changes with Mean Torque*

Although the head-neck system behaved like a second-order linear system when the mean torque was held constant, the overall dynamics of the system were not strictly linear, and cannot be sufficiently represented by a single linear IRF or a second-order, constant parameter model. The physical properties of the head-neck system were sensitive to the static torque, and underwent significant changes as the mean torque was varied. On average for the six subjects tested, as the mean torque was increased from 0.0 to 0.65 N m, stiffness increased by 80% and damping increased by 50%.

The data supported modeling head-neck stiffness as increasing with the square of the mean torque. As can be seen in Fig. 8, at mean torque levels of less than 0.25 N m, stiffness did not increase with mean torque, but fluctuated around a local mean value. In this range, the relation between stiffness and mean torque appeared flat. Above about 0.25 N m a change was seen in most subjects' data, and the stiffness increased significantly with increases in torque. Second-order polynomials visually fit the data well, and fit the data with slightly lower mean square error values than did pairs of first-order polynomials fitted separately to the data for the positive and negative torque levels.

The squared relationship found between stiffness and mean torque for the head and neck differs from that commonly reported for other joints. In most cases a joint's stiffness has been reported to increase linearly with the load applied to the joint.<sup>14,18</sup> This linear relationship is often attributed to motor unit recruitment: as the load applied to the muscle is increased, motor units are recruited to compensate for the increased force. The force, and also the stiffness, of each muscle fiber of the motor unit add in parallel, so doubling the load on the muscles doubles the stiffness of the joint about which the muscle acts.<sup>11</sup> This theory cannot, however, be assumed to account for the behavior of the muscles at all levels of contraction. The number of fibers activated in a muscle does not increase linearly across the full range of forces that can be generated by a muscle. It is generally thought that up to a certain percentage of its maximal load, a muscle will react to increases in load with an orderly sequence of recruitment of motor units. Beyond this load, rate coding is used to increase the force of contraction. The load at which the transition from recruitment to rate coding occurs, as well as the recruitment order of

the motor unit types, is dependent on the size, location, and complexity of the muscle.<sup>2</sup>

The difference between the torque-stiffness relation found for the head and neck and that for other joints may be attributable to a combination of the relatively narrow range of mean torques over which the head-neck system was tested, and the complicated musculoskeletal structure of the neck. Typically, the joints for which stiffness has been shown to increase linearly with torque, such as the ankle, knee, and elbow, have been tested across a much larger range of mean torques than was done for the head and neck, in some instances with loads exceeding 50% of the joint's maximum output.<sup>29</sup> For safety, comfort, and test repeatability, the maximum mean torque applied to the head-neck system in any of our tests was limited to 0.7 N m, which is much lower than the maximum rotational torque that the neck is capable of generating. In preliminary experiments, we found that a healthy male could rotate the head in the horizontal plane against a torque of 4.5 N m. Above 4.5 N m it was difficult for the subject to produce the appropriate movement, and the motion became a mixture of rotation and lateral flexion of the neck. Assuming that 4.5 N m is the maximum torque that can be generated by the neck in pure rotation, the 0.7 N m load applied in our experiments is less than 16% of the neck's maximum torque output. A more linear relationship between head-neck stiffness and mean torque may be evident if the system is tested with higher levels of mean torque. Since the maximum load applied to the head was kept low, the "flat" region seen between 0.0 and 0.25 N m had a significant bearing on the type of model that best fit the data. It is possible that the nonlinearity exists only at low torques, as if the system has a "bias" torque that has to be overcome before the linear relation between stiffness and mean torque can be seen. This flat, nonlinear region might also be characteristic of other joints, but is not apparent in the data because the joints are not typically tested at as many low amplitude torques. The few data points in the nonlinear region at low loads then become insignificant compared to the number of data points in the more linear region that exists at higher loads.

The neck's complicated muscle architecture may also account for this relation between stiffness and mean torque. Most of the joints that have been studied, such as the ankle, elbow, and knee, are hinged joints for which movement is primarily controlled by the action of an agonist/antagonist pair of muscles. The properties of these joints are therefore dominated by the characteristics of the agonist/antagonist pair. These muscles may compensate for added load with a linear increase in the number of motor units activated, which results in a linear relation between stiffness and mean load. Rotation of the head about the atlantoaxial joint, a monaxial pivot joint, is accomplished with a much more complicated arrange-

ment of muscles. There is not a single pair of agonist/antagonist muscles arranged to directly pull the head around the atlantoaxial joint. Rotation of the head towards one side is accomplished mainly by the coordinated action of the sternocleidomastoid, semispinalis capitis, splenius capitus, and the longissimus capitis muscles. These muscles are arranged somewhat vertically in the neck, originating either near the sternum or one of the lower cervical vertebrae, and inserting in the head at the mastoid process of the temporal bone or the occipital bone. This arrangement allows these muscles to control also head tilt and extension.<sup>8,27</sup> In addition to these muscles, there are many other muscles that pass through the neck to control the oral-facial musculature. There are also major arteries, veins, the trachea, esophagus, and connective tissue. As the head is rotated, the motion of all these structures contributes to the overall viscoelastic properties of the head-neck system. This complicated arrangement of muscles and soft tissue may well account for the difference in the stiffness-mean torque relation observed for the head and neck as compared to the arm and the leg.

#### *Damping Ratio, $\zeta$*

The head-neck system's stiffness and damping properties changed considerably as the mean torque was varied, but did so in a manner that maintained the second-order damping ratio virtually unchanged for each subject. In all but a few of the 150 experimental trials, the head-neck system was calculated to be well underdamped, with the average damping ratio for each subject falling between 0.42 and 0.65. These values are higher than that reported for the forearm (0.10–0.17),<sup>18</sup> the ankle (read from plot as 0.2–0.45),<sup>29</sup> and the knee (0.2).<sup>30</sup> The damping ratio calculated for the head and neck seemed, at first consideration, low for a system whose primary task is to stabilize the head. Purely passive systems with damping ratios between 0.4 and 0.7 respond to step inputs with multiple oscillations before settling to a steady state. It seemed reasonable to expect that the head-neck system would have evolved to be overdamped ( $\zeta > 1$ ), or at least critically damped ( $\zeta = 1$ ), so that the energy of disturbances would be dissipated quickly and without oscillations. Although low as a measure of a system's relative stability, the damping ratios found for the head-neck system may allow the neck to maintain a good balance between stabilizing against disturbances and responding quickly for active positioning. The head serves as the frame of reference for the visual, vestibular, and auditory senses, and must be both well stabilized and well controlled for us to interact effectively with the environment. The advantage of maintaining a relatively low damping ratio is that the head can be quickly positioned without having to overcome high damping loads.

Underdamped systems with damping ratios between 0.5 and 0.8 are able to get close to their final value more rapidly than either critically or overdamped systems.<sup>22</sup>

The head and neck's damping ratio remained nearly constant across a range of applied loads, an observation also found for the ankle<sup>13,29</sup> and knee.<sup>30</sup> This invariance may simplify the planning necessary for movement.<sup>14</sup> The motion of a limb is usually carried out by sequences of agonist/antagonist muscle activations. These phasic activations control the limb's acceleration and deceleration, and are combined as needed to produce motion with the appropriate temporal characteristics.<sup>5</sup> It can be assumed that the neck is positioned using a similar sequence of phasic muscle activations, although they are not generated by an agonist/antagonist muscle pair. By maintaining a constant damping ratio under different loads, the head-neck system retains similar fundamental response characteristics for different muscle activation levels. Once an appropriate command sequence has been developed for one activity, or for one level of activation, the sequence of phasic activations can be generalized for other activities by scaling the amplitude and time course.<sup>14</sup>

The head does not react to disturbances as a purely passive system, and so does not rely upon static physical properties to stabilize against, or to mitigate the effects of, these disturbances. They are initially resisted by passive properties and reflex action,<sup>3</sup> but the overall response of the head and neck is the result of a strategy that involves active intervention. After the initial reaction to a disturbance, an active modulation of the neck's viscoelastic properties can be used, if needed, to reduce the magnitude and time course of the disturbance to the head. Cocontraction of the agonist/antagonist muscles acting on hinged joints has been shown to increase the joint's effective stiffness and damping.<sup>14</sup> Lacquaniti *et al.*<sup>18</sup> showed that when subjects actively resisted perturbations applied to the forearm both the stiffness and the damping of the elbow increased without significantly changing the system's damping ratio.<sup>18</sup> As a result, the magnitude of the forearm's response decreased, and the settling time of the system, which describes the time the disturbance causes the system to oscillate, was reduced. By contracting the muscles responsible for the head's rotation, the viscoelastic properties of the head-neck system could be actively modulated to change the dynamic response of the head and neck. The present study showed that the head and neck responded like a second-order system with invariant parameters under the conditions that a subject was to correct for position error, but was not to resist the torque disturbances by co-contracting the neck muscles. This restriction may have prevented the subjects from using a natural strategy that involved co-contraction to minimize the head's position error during the experimental trials, but was necessary so that the

activation level of the neck muscles would be dictated by the torque applied to the head and not arbitrarily by the subject.

## CONCLUSION

The motivation for this study was to develop an understanding of the relation between small torque perturbations applied to the head and the resulting disturbance to the head's position, so that appropriate torque sequences could be designed for stimulating the human angular VOR. Directly applied, horizontal plane, torque perturbations are not frequently experienced by humans, but unpredictable disturbances to the head in the horizontal plane are experienced actively during locomotion and experienced passively when traveling in a car, plane, or any moving vehicle. The VOR (horizontal) is often considered to be capable of responding to unpredictable head motions that range from 0.5 to beyond 10.0 Hz and attain rotational velocities of up to 170 deg/s. Laboratory testing of the VOR is rarely, if ever, conducted so that the full performance range of the VOR is explored. More often than not, testing protocols use predictable stimuli, relatively low frequencies, and slow velocities.<sup>19</sup> The findings from this study will assist in designing torque stimuli that are well tolerated by subjects and can induce head motions that span the performance capabilities of the horizontal, angular vestibular system. This study demonstrated that although the head-neck system is stabilized by frequency dependent mechanisms, and is in general a nonlinear system, the response to stochastic perturbations can be well represented by a quasilinear, second-order system. When the mean torque applied to the head was held constant, and the resulting position disturbances small, the viscoelastic properties of the system were not dependent on the stimulus frequency between 0.5 and 10 Hz. The viscoelastic properties were, however, sensitive to the mean torque. As the mean torque was increased, both the effective stiffness and damping of the system increased, and did so in a manner that maintained the system's damping ratio nearly constant. The finding that the head neck system has a quasilinear, second-order response from 0.5 to 10 Hz should, however, not be interpreted to imply that under natural conditions the head and neck will behave like an oscillatory, second-order system. The models and trends found in this study apply directly to the specific conditions of this study. The natural behavior of the head and neck involves both active and passive stabilization, as well as active positioning. The response to torque perturbations observed in this study are an indication of the expected short duration response of the head, before active intervention and modulation of the system properties occurs.

## ACKNOWLEDGMENTS

The authors would like to thank Dr. Henrietta Galiana for her expert advice and Rachel M. Peters for her contributions to this project. This research was supported in part by the Total Home Automation and Health Care System Consortium and by a grant from NINDS of NIH.

## REFERENCES

- <sup>1</sup>Bennet, D. J., J. M. Hollerbach, Y. Xu, and I. W. Hunter. Time-varying stiffness of human elbow joint during cyclic voluntary movement. *Exp. Brain Res.* 88:433–442, 1992.
- <sup>2</sup>Binder, M. D., C. J. Heckman, and R. K. Powers. "The physiological control of motoneuron activity." In: *Handbook of Physiology: Section 12: Exercise: Regulation and Integration of Multiple Systems*, edited by L. B. Rowell and J. T. Shepherd. New York: Oxford University Press, 1996, pp. 3–53.
- <sup>3</sup>Bizzi, E., P. Dev, P. Morasso, and A. Polit. Effect of load disturbance during centrally initiated movements. *J. Neurophysiol.* 41:542–556, 1978.
- <sup>4</sup>Collewijn, H., and J. B. J. Smeets. Early components of the human vestibulo-ocular response to head rotation: Latency and gain. *J. Neurophysiol.* 84:376–89, 2000.
- <sup>5</sup>Cooke, J. D., and S. H. Brown. Movement-related phasic muscle activation. II. Generation and functional role of the triphasic pattern. *J. Neurophysiol.* 63:465–472, 1990.
- <sup>6</sup>Furman, J. M., and J. D. Durrant. Head-only rotational testing in the elderly. *J. Vestib. Res.* 8:355–361, 1998.
- <sup>7</sup>Gottlieb, G. L., and G. C. Agarwal. Compliance of single joints: Elastic and plastic characteristics. *J. Neurophysiol.* 59:937–951, 1988.
- <sup>8</sup>Gray, H. *Gray's Anatomy*. London: Churchill Livingstone, 1986.
- <sup>9</sup>Grossman, G. E., R. J. Leigh, L. A. Abel, D. J. Lanska, and S. E. Thurston. Frequency and velocity of rotational head perturbations during locomotion. *Exp. Brain Res.* 70:470–476, 1988.
- <sup>10</sup>Hanson, J. M., and J. A. Goebel. Comparison of manual whole-body and passive and active head-on-body rotational testing with conventional rotary chair testing. *J. Vestib. Res.* 8:273–282, 1998.
- <sup>11</sup>Huxley, A. *Reflections on Muscle*. Princeton: Princeton University Press, 1980, pp. 68–79.
- <sup>12</sup>Juang, J. N. *Applied System Identification*. Englewood Cliffs, NJ: Prentice Hall, 1994.
- <sup>13</sup>Kearney, R. E., and I. W. Hunter. Dynamics of human ankle stiffness: Variation with displacement amplitude. *J. Biomech.* 15:753–756, 1982.
- <sup>14</sup>Kearney, R. E., and I. W. Hunter. System identification of human joint dynamics. *Crit. Rev. Biomed. Eng.* 18:55–87, 1990.
- <sup>15</sup>Keshner, E. A., T. C. Hain, and K. J. Chen. Predicting control mechanisms for human head stabilization by altering the passive mechanics. *J. Vestib. Res.* 9:423–434, 1999.
- <sup>16</sup>Keshner, E. A., and B. W. Peterson. Mechanisms controlling human head stabilization. I. Head-neck dynamics during random rotations in the horizontal plane. *J. Neurophysiol.* 23:2293–2301, 1995.
- <sup>17</sup>Korenberg, M. J., and I. W. Hunter. The Identification of Nonlinear Biological Systems: Volterra Kernel Approaches. *Ann. Biomed. Eng.* 24:250–268, 1996.
- <sup>18</sup>Lacquaniti, F., F. Licata, and J. F. Soechting. The mechanical behavior of the human forearm in response to transient perturbations. *Biol. Cybern.* 44:35–46, 1982.
- <sup>19</sup>Leigh, R. J. "What is the vestibular-ocular reflex and why do we need it?." In: *Disorders of the Vestibular System*, edited by R. W. Baloh and G. M. Halmagyi. New York: Oxford University Press, 1996, pp. 12–19.
- <sup>20</sup>Ljung, L. *System Identification: Theory for the User*. Upper Saddle River, NJ: Prentice Hall, 1999.
- <sup>21</sup>Marmarelis, V. Z. Coherence and apparent transfer function measurements for nonlinear physiological systems. *Ann. Biomed. Eng.* 17:143–157, 1988.
- <sup>22</sup>Ogata, K. *Modern Control Engineering*. Upper Saddle River, NJ: Prentice Hall, 1997.
- <sup>23</sup>Peng, G. C. Y., and T. C. Hain. A dynamical model for reflex activated head movements in the Horizontal Plane. *Biol. Cybern.* 75:309–319, 1996.
- <sup>24</sup>Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C*. Cambridge: Cambridge University Press, 1999.
- <sup>25</sup>Roberts, W. J., N. P. Rosenthal, and C. A. Terzuolo. A control model of the stretch reflex. *J. Neurophysiol.* 34:620–634, 1971.
- <sup>26</sup>Shupert, C. L., and F. B. Horak. Effects of vestibular loss on head stabilization in response to head and body perturbations. *J. Vestib. Res.* 6:423–427, 1996.
- <sup>27</sup>Tortora, G. J., and S. R. Grabowski. *Principles of Anatomy and Physiology*, 7th ed. New York: Harper Collins, 1993.
- <sup>28</sup>Viviani, P., and A. Berthoz. Dynamics of the head-neck system in response to small perturbations: Analysis and modeling in the frequency domain. *Biol. Cybern.* 19:19–37, 1975.
- <sup>29</sup>Weiss, P. L., I. W. Hunter, and R. E. Kearney. Human ankle joint stiffness over the full range of muscle activation levels. *J. Biomech.* 21:539–544, 1988.
- <sup>30</sup>Zhang, L. Q., G. Nuber, J. Butler, M. Bowen, and W. Z. Rymer. *In vivo* human knee joint dynamic properties as functions of muscle contraction and joint position. *J. Biomech.* 31:71–76, 1998.



## **Chapter 7**

### **VOR experiment, analysis and results**

- 7.0 Chapter summary
- 7.1 Experimental methods
- 7.2 Data processing
- 7.3 Results from the tracking trials
- 7.4 Head response
- 7.5 VOR response
- 7.6 Prediction of the gaze response

#### **7.0 Chapter summary**

Chapter 7 discusses the results and the analyses of the experimental trials that were conducted to test the VOR. The goal of this research was to develop the methods required to evaluate the VOR under natural tracking conditions. The results show that a consistent identification of the VOR system can be made over a broad frequency range during natural head-free pursuit of a visual target. The testing protocol stimulated a sufficiently wide range of tracking and VOR responses from subjects a frequency range of 0 Hz to beyond 8 Hz. The linear analysis, though, was limited to evaluating the VOR over a smaller frequency range, from approximately 0.5 Hz to between 4 and 5 Hz. Nonlinearities in the head-neck system affected the low frequency linear analysis, and a poor signal-to-noise ratio degraded the analysis at the higher frequencies. Very consistent behaviors were seen in the VOR in both the frequency and time domain analyses. Both indicated that the VOR had a high-pass system response, and that the VOR gain during

head-free tracking was near one. In addition, estimates of the impulse response functions could be made for the head control systems as well as for the VOR. These models were used to reconstruct the gaze tracking system during the VOR experiment, and to predict the eye and head responses that occurred during the experiments. The results showed that the linear analysis was able to capture the underlying dynamics of the gaze tracking system during head-free tracking.

## **7.1 Experimental methods**

### **7.1.1 Subjects**


Eleven VOR experiments were conducted on seven subjects: five men and two women aged 21 to 34 years. All were healthy graduate students with no history of or apparent neuromuscular or vestibular problems, and had normal or corrected-to-normal vision. No quantitative test was conducted to evaluate the subject's vestibular performance before conducting the VOR experiments, but the subjects were questioned by the experimenter and excluded if they reported any past neck, back, or vestibular problems. All subjects had a range of motion with their head and neck that exceeded the rotational displacements that would be experienced during the experiments. The study was approved by the Massachusetts Institute of Technology's Committee on the Use of Humans as Experimental Subjects. All subjects gave their written consent before participating in this research.

### **7.1.2 Procedure**

Subjects were seated in a rigid chair, just under the helmet of the head perturber, and 1600 mm in front of the screen onto which the visual target was projected (Figure 7.1). The helmet was lowered onto the subject's head and tightened with head and chin straps so that it did not slip when torque was applied or the head was moved. Electrodes (15 mm diameter) were attached beside the outer canthus of the eyes and on the forehead to record the EOG. The head



perturber was aligned with the subject's head and neck as described in Chapter 6. The head perturber was adjusted so that the subjects were able to rotate their head comfortably through  $\pm 30$  deg without feeling that the rotational axis of the perturber was eccentric to the primary, vertical axis of their neck.



**Figure 7.1:** Top down view of the experimental testing area.

Each VOR experiment consisted of between 8 and 20 tracking trials. During each trial a visual target was projected onto the screen, and the subject was asked to follow the target with natural head and eye motions. In about half the trials torque perturbations were applied by the head perturber to disrupt the position of the subject's head, and in the other half no perturbations were applied. The amplitude of the torque perturbations was between 0.2 and 0.3 N.m, depending on what the subject found comfortable and what was required to induce a strong VOR eye motion. The tracking trials lasted from 30 to 120 seconds. The shorter tracking trials were used early in the development of the experimental protocol. Trials that lasted from 60 to 120 seconds produced sufficient data for analyzing the system characteristics of the VOR and the gaze system. The torque and target velocity sequences that were used in the experiments were selected from a large set of sequence pairs that were created in advance to have the required ACF, CCF, and amplitude characteristics. These torque and target signals were generated using the techniques

described in Chapters 4 and 5. The particulars of each experiment, such as the target PSD, torque amplitude and PRBS execution rate, experiment length, and sampling rate, are listed in Appendix D. Immediately before each tracking trial, the gain relation between the EOG's voltage measurement and the subject's eye position was calibrated by having the subject track a series of slowly moving sinusoids (Chapter 3). To help prevent fatigue, rest periods of at least 1 minute were given between each trial, and 5 to 10 minute breaks were taken after every four tracking trials. Data were collected on the torque applied to the helmet of the head perturber, the position of the visual target, the head velocity, and the position of the eyes relative to the head. The sensors were energized and allowed to warm-up for at least 30 minutes prior to collecting data. The equipment and sensors used in the VOR experiments are fully described in Chapter 3.

### 7.1.3 EOG measurements

Prior to recording the EOG, the subject's skin had to be prepared for the EOG electrodes, and the subject's vision had to be acclimated to the testing environment. These preparations helped improve the amplitude and stationarity of the voltage signal measured at the eye. Subjects first washed their forehead and the outer canthus of each eye. The skin of these areas was then scrubbed with a lightly abrasive electrode gel (Nuprep, D.O Weaver, Aurora, CO) that removed oils, dirt, and some of the stratum corneum – the top layer of the epidermis. This lowered the impedance of the skin and increased the strength of the voltage signal acquired by the EOG. Since only horizontal, binocular eye movements were evaluated during the VOR experiments, an electrode was placed at the outer canthus of each eye socket in line with the horizontal movement of the eye. A third electrode was placed on the forehead as the reference electrode for the EOG (Figure 7.2). Once the electrodes were attached and connected to the EOG amplifier, the subject was seated in the testing area under the head perturber, and the testing area was made dark. The only source of light seen by the subject during this period and during the experiment was the visual target. The subject sat for 20 to 30 minutes so that their retinae acclimated to the lighting

conditions in the testing area. Since the EOG does not directly measure the position of the eyes, but measures the voltage that develops on the bones surrounding the eye, it is necessary to mitigate the effects of, if not eliminate, any transients that could affect the voltage measured by the electrodes. The voltage level measured by the electrodes was often less than 1 mV, so even very small voltage transients could have a significant impact on the measurement of eye position. Twenty to thirty minutes was sufficient time for the retinae to adjust to the low light level, for the half-cell potential of the electrodes to equilibrate, and for any transients from the placement of the electrodes on the skin to settle (Acoustical Society of America, 1999). After the acclimation period was over, the EOG signal was checked to verify that it had settled, and the gain of the EOG amplifier adjusted so that the EOG signal filled the dynamic range of the data acquisition board that sampled the data.



**Figure 7.2:** Subject with electrodes placed at the outer canthus of each eye socket and a reference electrode placed on the forehead.

## 7.2 Data processing

The data that were collected during the tracking trials had to be processed before they could be used to analyze gaze tracking and VOR performance. The measured voltages had to be converted from voltages into physical units, filtered for noise, and corrected for transients, and drift. Also, data that were needed for the analysis, but that had not been measured directly by

sensors, such as the head’s position and the target’s velocity, had to be calculated. The data were processed using a dedicated application, and the processed data were saved to files. An example of the program used for this initial processing called “Make Files” is in Appendix E. The raw data were particularly affected by how the data were filtered, and how drift was modeled and removed. These steps are discussed in more detail in the following section.


### *Scaling data*

The raw data collected included the target’s position, the torque applied by the head-perturber, the position of the eyes relative to the head, and the head’s velocity. These data were inspected to ensure that they had not saturated the data acquisition board, and were converted by the appropriate sensor and acquisition gains into physical units. For convenience, angular positions and velocities were converted into deg and deg/s, but were later converted to radians and rad/s for analysis. Corrections also had to be made to the head velocity and to the target position sequences. The head velocity sequence had to be corrected to compensate for a non-linear calibration error in the angular rate sensor that measured the head’s velocity (Chapter 3). The target position sequence had to be converted from the position of the galvanometer into the angular position of the visual target as seen by the subjects. It was the target’s position on the screen relative to the subject, not the galvanometer’s position that was the input that drove the subject’s gaze response. These differed because the subject’s head, and the mirror that positioned the visual target, were not at the same distance from the screen. The angular positions were related by:  $tp = \arctan\left(\frac{md}{sd} \tan(mp)\right)$ , where  $tp$  is the target’s position relative to the subject,  $mp$  the galvanometer’s (mirror’s) angular position,  $md$  the slant distance between the mirror and the position of the target on the screen, and  $hd$  the slant distance between the subject’s eyes and the target.

### *Median and low-pass filtering*

The data were median filtered to remove aberrant data points and then low-pass filtered to remove high frequency noise. Unlike mean filters, which are frequently used to smooth data, the median filter can remove outliers without affecting the sharp, higher frequency characteristics of a sequence. Median filters are nonlinear so they can affect signals within the same data set differently. They must be used carefully so that they do not alter the linear relation between data sequences. The neighborhood of the median filter that was applied to the tracking data was kept small so that its effect on anything but outliers was small. When the data had been sampled at 100 samples/s, the neighborhood of the median filter was set to 3, 5, or 7 elements. The size of the neighborhood was chosen so that it did not have a noticeable effect on the PSD of any of the data sequences below 5 Hz, which was the upper limit for the analysis of the linear relationships. A representative eye position sequence and PSD are shown in Figure 7.3. The PSD of the sequence that was filtered with the 5-point median filter deviated from the unfiltered PSD between 6 and 8 Hz. The time domain sequence was changed, but retained many of the sharp characteristics of the unfiltered sequence. The median filters seemed to have had the greatest effect on the eye position sequences. This may have been because the eye position's power spectral density dropped rapidly with increasing frequency. When the PSD of the eye position indicated that the median filter had affected the eye position at frequencies as low as 5 Hz, the PSD of the other sequences typically were not affected much below 10 Hz.

When the median filter neighborhood was set too high, the characteristics of the sequences were noticeably changed, and the sequence's PSD was affected at low frequencies. The blue plots in Figure 7.3 show the eye position sequence after it was filtered by a median filter with a 15-point neighborhood. The neighborhood is too large and has had a great impact on the sequence. The two small peaks between time 16.4 and 16.6 have been replaced by a constant. The PSD of the sequence deviates from the unfiltered PSD at frequencies as low as 3 Hz.



**Figure 7.3:** Eye position (plot A) and its power spectral density (plot B). The eye position sequences are shown when they were not filtered (black), and when they were filtered by a median filter with a 5 point (0.05 s) neighborhood (red), and a 15 point (0.15s) neighborhood (blue).


After the data were median filtered, they were low-pass filtered (LPF) to remove high frequency noise. The LPF was created with the Kaiser window filter design method (Oppenheim *et al.*, 1999). The great advantage of this method as compared to designing an ideal low pass filter and applying a more common window, such as Hamming or Hanning, is the ease with which the impulse response function (IRF) of the filter can be generated. The Kaiser method allows the user to specify the filter's desired cutoff frequency, transition band, and peak approximation error and then produces the elements of the discrete-time impulse response function that satisfy the specifications. Any of the traditional filter windows can be created with the Kaiser method by simply specifying the characteristics associated with that window. To prevent the filter from introducing a phase change to the data, the IRF was made to be acausal, and to be symmetric about time zero. Depending on the conditions of the particular tracking trial, the filter was given a cutoff frequency between 10 and 15 Hz. The filter's cutoff was usually set to be just above the execution rate of the PRB torque sequence. A cutoff frequency of 10 to 15 Hz was also high enough to retain the important characteristics of the smooth pursuit eye motions (Dr. H.L. Galiana, McGill University, personal communication, April 2002). It was important that the cutoff of the LPF was set this high so that on visual inspection the time and frequency

domain plots of the sampled experimental sequences looked similar to the designed command signals, as this allowed the signals to be easily compared. During the analysis it was found that noise in the experimental signals degraded the linear analysis at frequencies above 5.0 Hz. The low-pass filter cutoff frequency could be reduced towards 5.0 Hz to smooth the time domain picture of the experimental sequences without having an obvious impact on the results of the linear frequency analysis. After the data sequences were filtered, transients were removed from the beginning and end of the sequences. The transients were caused by there not being sufficient data to convolve the full length filter IRF, and were also due to starting and stopping the tracking trial.

#### *Drift and offset*

Drift and offset had to be removed from the data sequences as was done for the calibration trials (Chapter 3). This was particularly important so that the calculated gaze responses reflected what occurred during the tracking trials. The head and eye position sequences were most sensitive to drift in the sensors. A low frequency drift in the head velocity data might not have affected the velocity measure, but since head position was calculated as the integral of the velocity, drift could result in large displacements in head position. The also occurred for eye position and eye velocity sequences, as eye velocity was calculated from the derivative of the measured eye position. A slow drift in the eye position signal might indicate that the eyes had moved through a large displacement, but the drift would not be reflected in a significant change in the eye velocity. It was more difficult to identify drift and offset in the tracking trials than it was in the calibration trials. During the calibration trials, the head movement was usually small and the gaze response was composed mainly of eye movements. During the tracking trials, subjects were encouraged to use their heads naturally to track the target, and the head's position was disturbed by the torque perturbation.

The measured torque and target sequences were biased, but did not appear to drift over the course of the tracking trial. The target sequence was measured from the galvanometer's position feedback signal, and the torque was measured using the force/torque sensor. Both sensors had small offsets, but neither had appreciable drift. The offsets in the torque and the target signals were estimated by taking the mean value of the sequences at the beginning and end of the trial, when the target, the head perturber, and the head were at rest (Figure 7.4). There was usually very little difference between the means calculated at the beginning and end of the trial and so the two were averaged. The mean offset value was then subtracted from all elements in each sequence.



**Figure 7.4:** Regions where target was at rest.

As was discussed in the EOG calibration section in Chapter 3, the head velocity and the eye position signals were corrupted by both drift and offset. Since the target drove the subject's gaze, and not the head or eye motions directly, the head and eye movements were coupled and could not be processed independently. Several drift models were fitted to, and removed from, the head velocity and eye position sequences. Head position sequences were calculated from the head velocity sequences, and were added to the eye position signals to form potential gaze responses. These gaze responses were then compared to the target's position. The error was calculated between the gaze and the target, and summarized by the mean squared error and standard

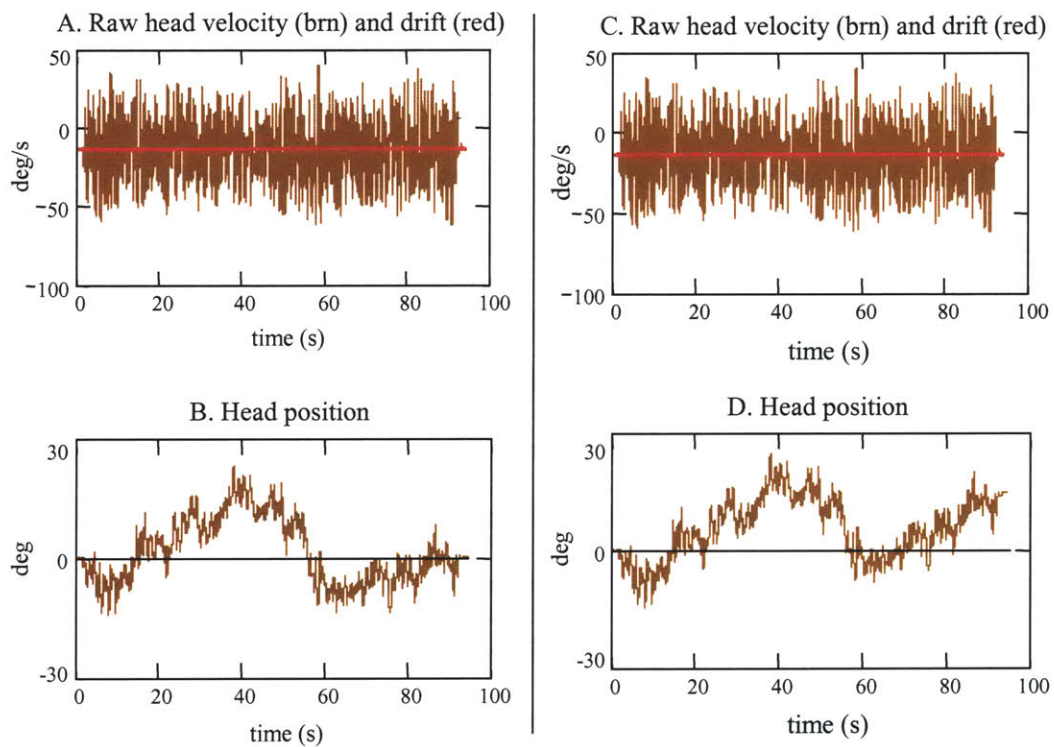


deviation, as these gave different estimates of the error which did not have a zero mean. The head velocity and eye position sequences that created the gaze that seemed most sensible in terms of the how the gaze had followed the target, not necessarily with the lowest error, were selected as representative of what had occurred during the experiment.

The head and eye sequences were analyzed with the assumption that the subject had tracked the target, and when the drifts and offsets were removed, the head and eye sequences should produce a gaze response that looked similar to the target's trajectory. A subject might not have tracked all the motions of the target during the trial, especially the high frequency and high velocity motions, and so some error was expected to exist between the gaze and the target. Additionally, the target began and ended the trial at the zero position and at rest. Subjects were instructed to start the trial with their head and eyes at rest at the zero position, and to end the trial the same way.

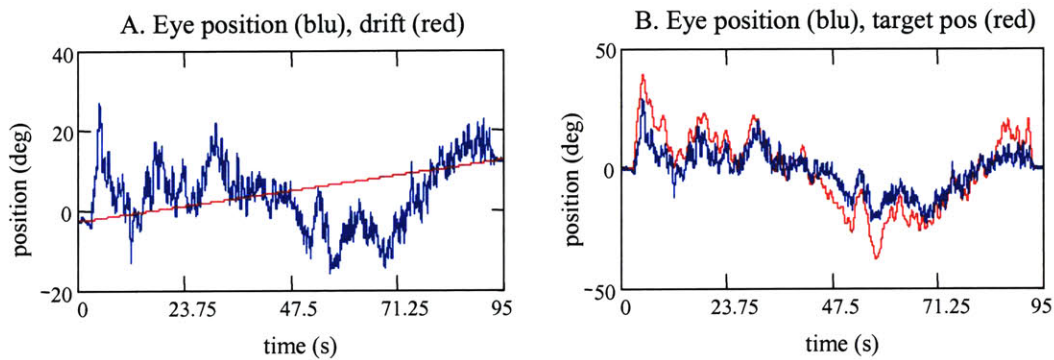
Three different methods were used to approximate and remove the drift and offset from the head velocity sequence. The drift in the head velocity sequence was usually small, but the position sequence was definitely affected by how the drift was removed. The methods, and the drift models that they assumed, were similar to those used for the EOG calibration. The first and easiest method was to give the head velocity a zero mean, which is based on the assumption that the velocity signal had an average offset, but did not drift. The zero mean velocity constrained the head's position to end the trial at exactly the same position as it began, which were the instructions given to the subjects, but probably did not reflect the reality of the head's movement. The second method was to remove a least-squares line from the data that was fitted to the velocity sequence. This model assumed that the drift was linear, and also gave the head velocity sequence a zero mean (Figure 7.5.A). The third method assumed that the head was at rest at the beginning and end of the trial. The mean of the data during these periods represented the sequence's offset during that time. Between these regions the offset was assumed to drift linearly. It was modeled by connecting a line to the means of the two regions. This model gave the head velocity a zero

mean during the regions at the beginning and end, but did not give the overall sequence a zero mean. The head position was therefore not constrained to end the trial at zero degrees. Since it is unlikely that the subject's head would have ended the trial at exactly the same position as it began, the assumptions of this drift model probably best reflected what occurred during the trials. However, this method most often produced a head velocity sequence that was inconsistent with what was observed. Often the calculated head position indicated that the head had gradually twisted during the trial, and had ended the trial 10 to 30 degrees away from the zero position (Figure 7.5.B). The head moved throughout the tracking trials, but subjects generally brought their head back to within a few degrees of zero at the end of the trial.



**Figure 7.5:** Head velocity, drift, and the calculated head position. In plot A the drift (red) was modeled by a line fitted to the velocity sequence (brown). The resulting head position is shown in plot B. In plot C the drift was modeled by a line that connected regions at the beginning and end of the experiment when the head was assumed to be at rest. The resulting head position is shown in plot D.

Two drift models were fitted to the eye position sequences. For the first, there was assumed to be no drift in the eye position data, and the offset was calculated as the mean of the regions where the target was at zero. For the second, the drift was modeled by fitting a line between the offsets calculated at the beginning and end of the data record (Figure 7.6), as was done for the head velocity (Figure 7.5.B).



**Figure 7.6:** The drift (red) in the eye position sequence was modeled as a line that connected the regions at the beginning and end of experiment when the eyes were assumed to be at zero deg (plot A). The resulting eye position sequence (blue) and the target position (red) are shown in plot B.

The six possible gaze responses from the three head position sequences and the two eye position sequences were calculated and each was compared to the target position sequence. It was usually obvious which sequences, and which combinations, did not make sense within the context of the experiment. Sequences that indicated that the gaze, or the head, had deviated far from the target were discarded. Of the sequences that seemed to make sense, the MSD and standard deviation of the error were used to decide the best combination of head velocity and eye position. These were the sequences that were saved to files and used in the data analyses.

### *Velocity calculations*

The eye velocity and target velocity sequences were calculated from their position data using a five-point, equal abscissa, numerical derivative algorithm (Abramowitz and Stegun,

1972). The algorithm calculated the derivative of the discrete-time sequences using the coefficients of a polynomial expansion that was fitted through 5 data points. The data points were usually centered around the data point for which the derivative is desired. The expansion is shown in Appendix E.

The five-point, equal abscissa algorithm was selected from several numerical methods that were evaluated to calculate the derivatives of the target and eye position data. The target and eye position sequences were not smooth functions, but were the sampled output of random processes and were, as well, noisy. Although the underlying process might not change quickly, noise in discrete-time sequences can lead to high values being calculated for the numerical derivative by causing there to be greater changes in the values of successive data points (Press *et al.*, 1997). Thus it was necessary to select a numerical derivative algorithm that was somewhat insensitive to the noise in the position data. The numerical derivative methods that were evaluated were: 1. fitting a 2<sup>nd</sup> order polynomial to groups of three data points, and finding the derivative of this polynomial, 2. fitting a 2<sup>nd</sup> order polynomial with least squares to groups of 5 data points, and calculating the derivative of the polynomial (Savitzky-Golay filtering), and 3. the five-point, equal abscissa expansion. For comparison, a two-point difference method, which is very sensitive to the noise in the data, was also tested. In general, the results of the three derivative methods were quite similar. The Savitsky-Golay method tended to produce the smoothest derivative results, but the velocities were not necessarily more accurate than from other methods, and the method required the most processing time. The five-point equal abscissa algorithm was selected as the method to be used to calculate the experimental eye and target velocities because it was well balanced between being insensitive to noise and having a low calculation time, and was robust in that it was able to find a derivative for all the data sequences (Press *et al.*, 1997).

## 7.3 Results from the tracking trials

The accuracy of the VOR system identification was dependent on having controlled the experimental inputs successfully, and having stimulated the appropriate head and eye motions from the subjects. Before the data from a trial were analyzed, the input and output sequences were inspected. The input sequences should have maintained the spectral characteristics and cross-correlation relations that they had been designed to have. In addition, the subject should have tracked the target well so that there was confidence that the gaze response had been driven by the visual target. Finally, the torque perturbations should have disturbed the position of the head so that the VOR was needed to stabilize the gaze.

The outcome of the tracking trials was consistent with goals of the protocol. The performance of the subjects was consistent with the behavior discussed in Section 4.4, where it was shown that the perturbations did not degrade the subject's ability to track the target. It was apparent during the tracking trials that the torque perturbations had disturbed the motion of the head, but that the subjects were able to track the targets well. It was obvious that the subjects' VOR systems had stabilized their gaze. If the experimental protocol and the analysis were to be valid, then the system identified by the analysis must reflect the performance of the VOR across the full frequency spectrum of the head disturbances.

### 7.3.1 Input sequences: experimental vs. designed

The torque and target inputs that were measured during the experiments were not identical to the "designed" torque and target input signals. The differences were small, but resulted in the experimental signals having slightly different spectral characteristics and being more correlated than the designed signals. The differences occurred because the ideal, high precision torque and target sequences that were designed were transformed into physical realizations by actuators that were not 100 % accurate, and were sampled by noisy sensors. The

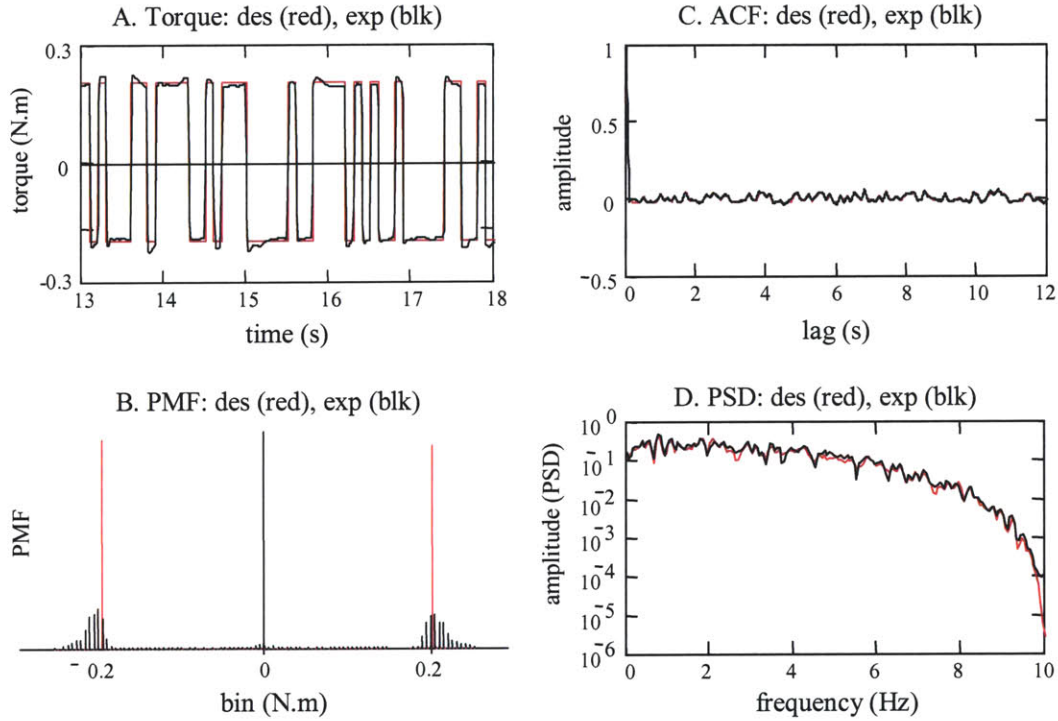
characteristics of the experimental signals were still good and satisfied the requirements for testing the VOR.

### *Input torque*

The torque signal was designed as an ideal PRBS. It had two amplitude values, and the switches made between amplitudes in the sequence were immediate. Each torque sequence was converted to a continuous-time voltage by a digital to analog converter (National Instruments), and applied as a command to the motor of the head-perturber. The motor's controller controlled the torque at the motor's rotor. The torque was delivered to the subject through a slip clutch, across the force/torque ( $f/t$ ) sensor, and then to the helmet that was fitted to the subject's head. The experimental torque signal was sampled from the  $f/t$  sensor's output. The  $f/t$  sensor measured the torque as close as possible to the head, which was required for the analysis. Since the  $f/t$  sensor was not collocated with where the motor controlled its torque output, the torque that the  $f/t$  sensor measured was different from the torque applied by the motor. The torque that was measured by the  $f/t$  sensor and recorded in the experimental torque sequence was a combination of the torque applied by the motor and the torque created by movements of the subject's head that accelerated the masses along the head perturber's shaft.

The experimental torque sequences matched the shape of the designed PRB torque sequences well and had similar power-spectral density and auto-correlation functions. The most obvious difference between the experimental and the designed torque sequences was that the experimental sequences did not have a binary amplitude distribution. The amplitudes of the torque pulses were not constant, but varied around the commanded torque level. A representative pair of experimental and designed torque and target sequence is shown in Figure 7.7, along with plots of their amplitude distributions, PSD, and ACF. The designed signal had two amplitudes, + 0.20 and - 0.20 N.m. The amplitudes of the experimental signal were not binary, but the majority of them were spread between magnitudes of 0.18 to a 0.23 N.m. The plot of the amplitude

distribution shows that the experimental sequence had some values with magnitudes less than 0.18 N.m. This was due to the  $f/t$  sensor being sampled as the motor switched from applying torque in one direction to the other.



**Figure 7.7:** Comparison of the experimental (black) and the designed (red) torque perturbation sequence (plot A). The PMF (plot B), ACF (plot C), and PSD (plot D) for the experimental and designed sequences are also shown.

Although the experimental torque sequences did not match the designed signals exactly, the degradation of the torque signal because of the non-binary amplitudes was unimportant. The important design goals for the torque signals were its spectral and auto-correlation characteristics, and very little difference could be seen between the PSD and ACF for the experimental and designed sequences (Figure 7.7). The experimental sequences approximated band-limited white sequences, and had ACFs that could be used as a bound to decide if the torque and target sequences were uncorrelated. The experimental sequences had the same temporal characteristics

as the designed sequences which helped preserve their cross-correlation with the target sequences.

The only times that the characteristics of the experimental torque sequence did not match those of the designed torque inputs were during the few trials where the PRB torque sequence was executed at 20 changes/s. The servomotor could not always respond quickly enough to follow the highest frequency amplitude changes. The high frequency amplitude switches were not accomplished, and the torque was held at the same amplitude for a longer duration than was intended by the command signal. This effectively changed the high frequency switches into lower frequency patterns. This altered the ACF of the torque sequence, and affected the torque-target cross-correlation function since the torque sequence did not have the same temporal arrangement of frequencies that had been designed.

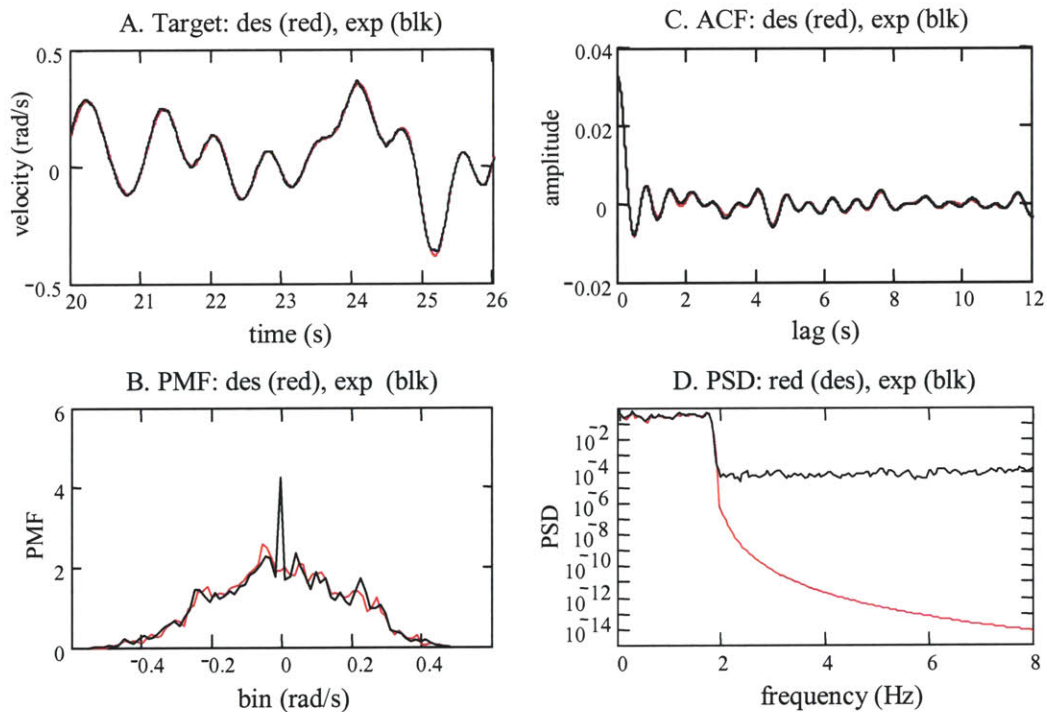
#### *Visual target*

The trajectories of the visual targets were designed as discrete time velocity sequences. Each velocity sequence was integrated into a position sequence, which was then converted into a continuous time command signal. The command voltages were applied to the galvanometer which moved the visual target's image by rotating a mirror off which the laser was projected. During the experiments, the galvanometer's position was measured. The angular position of the galvanometer was converted into the angular position at which the subject saw the target on the screen. The experimental target velocity sequences were calculated by taking the discrete time derivative of these target position signals.

The experimental target velocity trajectories indicated that the target had followed the trajectory of the designed velocity sequence. When the experimental sequences were overlaid on the designed sequences the sequences matched closely, but they were not exact. The experimental sequences were usually not as smooth as the designed velocity sequences. This was partly because the designed sequences had 200 elements/s whereas the experimental data sequences



usually had 100 element/s (due to sampling rate), and partly because the experimental velocity sequences were calculated by differentiating the target position sequence. As discussed in Section 7.3, differentiating a discrete time signal tends to amplify high frequency noise. The high frequency noise was evident in the power spectral density plots of the sequence, but did not have much effect on the target velocity's auto-correlation function. Figure 7.8 shows an example of a velocity sequence that was designed to have its power between 0 and 1.8 Hz. The PSD of the experimental velocity matches the PSD of the designed sequence in the desired frequency range, but beyond 1.8 Hz it contains more noise. The most apparent difference between the amplitude distributions occurred at the 0 deg/s bin. Most of the experimental sequences had a peak in their amplitude sequence at 0, which was simply due to the sequences being sampled when they crossed between positive and negative values.



**Figure 7.8:** Comparison of the experimental (black) and the designed (red) target velocity sequence (plot A). The PMF (plot B), ACF (plot C), and PSD (plot D) for the experimental and designed sequences are also shown.

### *Torque-target cross-correlation*

The experimental torque and target sequences were more correlated than the designed torque and target sequences. Over the lag values for which the torque and target sequences had been designed to be uncorrelated, the mean squared amplitude of the experimental torque-target CCFs were often 10 to 20 times greater than the mean squared amplitude of the designed torque-target CCFs. Although the CCFs were larger, their mean squared amplitude was still much smaller than that of the torque ACF. The torque and target sequences could therefore be considered as uncorrelated. The mean squared amplitude values for the CCFs and ACFs from two experiments are summarized in Table 7.1. The differences between the mean squared amplitudes of the ACFs for the experimental and for the designed torque sequences were not as great as those for the CCFs. The mean squared amplitudes of the experimental torque ACFs were rarely more than 20% larger than those for the designed sequences. The magnitude of the experimental ACF was still a good measure of the correlation between sequences that were uncorrelated. The ratio of the mean squared amplitude of the ACF to the mean squared amplitude of the CCF was used to verify that the torque and target sequences had a low enough correlation to be accepted as uncorrelated. The functions were evaluated over the lags for which the sequences were designed to be uncorrelated, but to be conservative, were not compared at lags that included the impulse in the torque ACF. For the designed sequences, it was common for the mean squared amplitude of the ACF to be 50 to 500 times that of the CCF. For the experimental sequences, the magnitude of the ACF was often only 3 to 40 times the magnitude of the CCF. It can be seen in Figure 7.9 that for either the designed or the experimental sequences, the CCF was small relative to the ACF.

The experimental torque and target sequences were more correlated than the designed sequences because of power that was in the experimental target sequences at frequencies higher than the desired frequency spectrum of the target's trajectory. This power was caused by noise, sampling, and from calculating the target velocity as the derivative of the position sequence. This added power lead to the cross-power spectral density (cross-PSD) between the target and the


torque sequence being higher at these frequencies. Because the same total power is in the cross-correlation function as is in the cross-PSD, the experimental sequences had to have higher cross-correlation values than the designed sequences.

Exp: 14Jun02	4	6	10	12	14	18	20	24
ACF des ( $10^{-6}$ )	0.54	0.76	0.20	0.20	0.16	0.65	0.78	0.67
ACF exp ( $10^{-6}$ )	0.64	0.84	0.23	0.22	0.20	0.71	0.86	0.78
CCF des ( $10^{-9}$ )	6.5	1.5	17.6	0.72	7.8	2.1	1.5	2.2
CCF exp ( $10^{-9}$ )	46.6	26.6	82.1	149	42.1	28.5	22.4	42.0
ACF/CCF des	82.7	513	11.2	273	20.9	305	514	312
ACF/CCF exp	13.8	31.6	2.8	1.5	4.8	25.1	38.3	18.5
Exp: 04May03	2	4	6	8	10	12	14	16
ACF des ( $10^{-6}$ )	1.32	1.00	0.23	0.61	0.36	0.98	0.67	0.67
ACF exp ( $10^{-6}$ )	1.36	1.04	0.24	0.69	0.41	1.06	0.74	0.75
CCF des ( $10^{-9}$ )	2.35	1.46	5.25	2.68	0.58	4.80	2.16	1.0
CCF exp ( $10^{-9}$ )	60.1	23.9	30.2	63.3	111	39.3	28.0	25.2
ACF/CCF des	564	686	43.0	226	617	203	310	646
ACF/CCF exp	22.5	43.5	8.1	10.9	3.7	27.0	26.6	29.8


**Table 7.1:** Mean squared amplitude of the PRB torque auto-correlation functions and the torque-target cross-correlation functions. The values are listed for the designed (des) and the experimental (exp) sequences for experimental trials run on two subjects. The trial's number is listed in the top row of each table.

In the example shown in Figure 7.8, the target sequence was designed to have power from 0 to 1.8 Hz. Between 0 and 1.8 Hz the PSD of the experimental target velocity was nearly identical to the PSD of the designed sequence. But at higher frequencies, the magnitude of the experimental sequence's PSD was greater than that of the designed target sequence. This power occurred at frequencies for which the PRB torque sequence also had power (Figure 7.7) which resulted in the experimental torque and target sequences having a higher cross-power spectral density at frequencies above 1.8 Hz than did the designed sequences (Figure 7.10). Since the

power above 1.8 Hz was created after the design process was complete, the target sequence had not been arranged so that the power occurred in the cross-correlation function at large lag values. It therefore corrupted the cross-correlation function at lags where the sequence had been designed to be uncorrelated, and made the experimental sequences more correlated than the designed sequences (Figure 7.10).



**Figure 7.9:** The torque ACF (red) and the torque-target CCF (black) for the designed (plot A) and the experimental (plot B) torque and target sequences. For the designed sequences, the mean squared amplitude of the torque sequence’s ACF was 310 times the mean square amplitude of the torque-target CCF. For the experimental sequences, the mean square amplitude of the ACF was 27 times that for the CCF.



**Figure 7.10:** Cross-power spectral density (plot A) and cross-correlation function (plot B) for the experimental (black) and the designed (red) torque and target sequences.

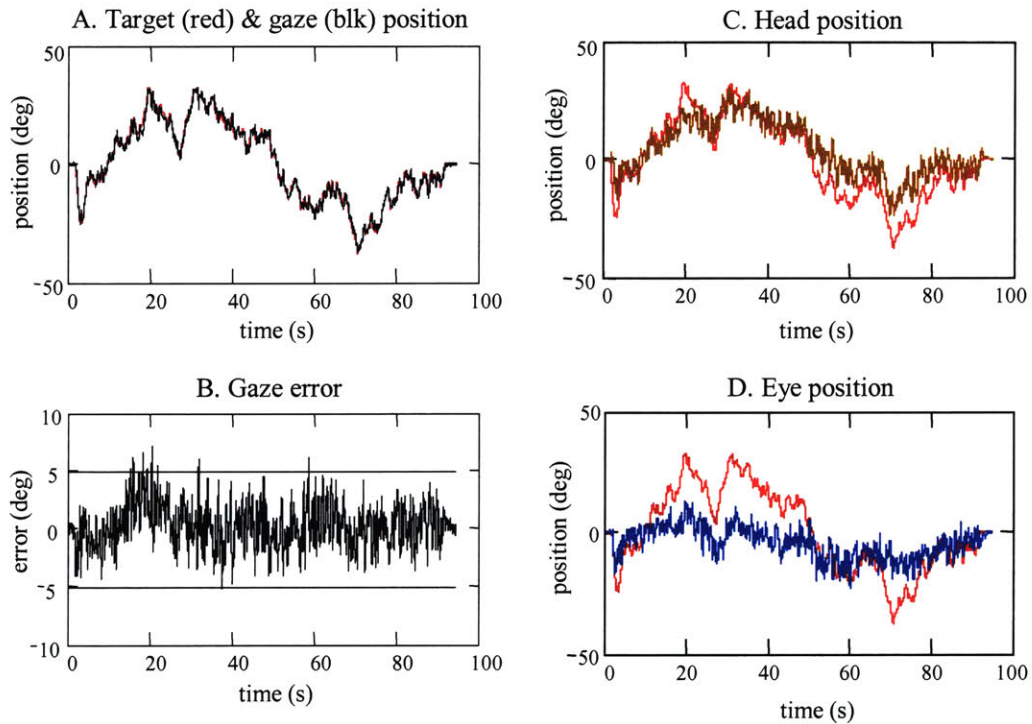
### 7.3.2 Output sequences: tracking, head and eye motions

The targets were well tracked by the subjects, and stimulated predominantly smooth pursuit tracking responses. The subjects demonstrated a range of tracking behaviors, but they were generally able to compensate for the disturbance to the head and tracked the targets without too much error or need for saccades. The head and eye movements were considerably disturbed, but combined to produce a good gaze response. There was not an observable difference in how well the subjects tracked the target, or if the tracking was done smoothly or with saccades, as a function of whether perturbations were applied or not during the tracking. The tracking behavior was consistent with what was observed during the investigations that were conducted to identify the characteristics of the visual target that were to be used in the VOR experiments (Chapter 4).

In the majority of trials the subject's head moved with the low frequency motions of the target. The eyes followed the higher frequency target movements and compensated for the disturbances to the head's position. The plots of head and eye position shown in Figure 7.11 are typical of trials in which the head was moved through nearly the same range as the target. The plot of the head position has the same gross shape as that of the target, but its motions are not smooth and are obviously affected by the torque perturbation. The eyes also have movements that are characteristic of the target, but the magnitude of the eye movements is smaller than that of the head. The eye motions are disturbed by the same amount as the head. The head and eye motions combined to form a gaze position that replicated the target's position. In this example, the mean squared difference between the gaze and the target was less than  $3.5 \text{ deg}^2$ , and the standard deviation of the gaze error was only 1.7 deg.

Not all subjects moved their head with the target, and a subject would not necessarily use the same head strategy in all tracking trials. In some cases, the subject tracked the target mostly with eye movements. The head's position remained close to zero, although it was affected by the torque perturbation. The eyes compensated for the disturbances to the head's position and moved with the target (Figure 7.12). How well the subjects were able to track the target did not seem to

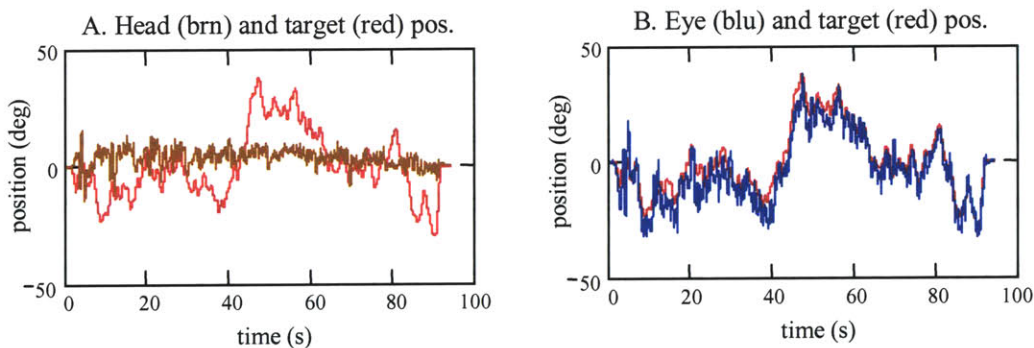
depend on whether they actively moved their heads or not. Since the maximum angular displacement of the targets was only  $\pm 35$  degrees, subjects were not forced to move their heads to track the target. Although this was a limitation of the equipment, it allowed the subjects to use whichever strategy they wanted.



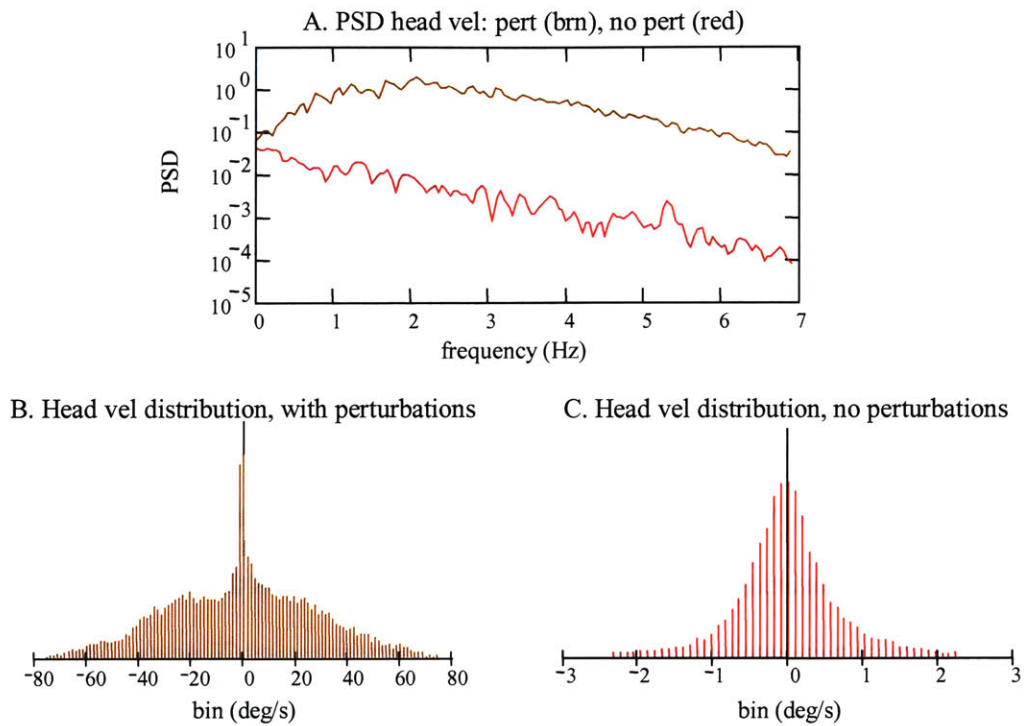
**Figure 7.11:** Representative gaze tracking response where subject moved the head through nearly the same range as the target (plot A). Plots C and D show the head (brown), eye (blue), and target (red) position sequences.

The spectra and amplitudes of the head's velocity as it was perturbed met and exceeded what was desired. The goal was to create horizontal plane disturbances that were similar to what is experienced by the head during locomotion, namely frequencies between 0.5 and 4.0 Hz, and velocities up to 60 deg/s. It was difficult to identify the disturbance to the head's motion by inspection alone. The head movements that were measured during the VOR experiments were a combination of active motions that were used to track the target and the disturbance motions caused by the torque. To identify accurately the motions caused by the torque, a cross-correlation

analysis would have to be conducted (Section 7.6). However, the effect of the perturbations on the subjects' heads can be understood by comparing the head velocities during a trial where torque perturbations were applied to those measured when there were no perturbations. An example of the head velocity PSD for two representative trials is shown in Figure 7.13. The PSD of the head velocity when perturbations were applied is greater at all frequencies. Near 0 Hz, the magnitudes of the PSDs are similar, but by 0.5 Hz, the magnitude of the PSD when torque was applied is an order of magnitude greater than when no perturbations were applied. The difference increases and remains steady at approximately two orders of magnitude difference at all frequencies above 2.0 Hz. When there were no perturbations, subjects tracked the target slowly with their head. The maximum velocity measured for the head was less than 3 deg/s. When the torque perturbations were applied the maximum head velocity exceeded 70 deg/s. The subject may, or may not, have moved the head differently to track the target when the perturbations were applied, but it can be assumed that most of the increase in the head velocity was caused by the applied torque.



**Figure 7.12:** Representative head (plot A) and eye (plot B) position for a subject that did not move the head with the target (red).



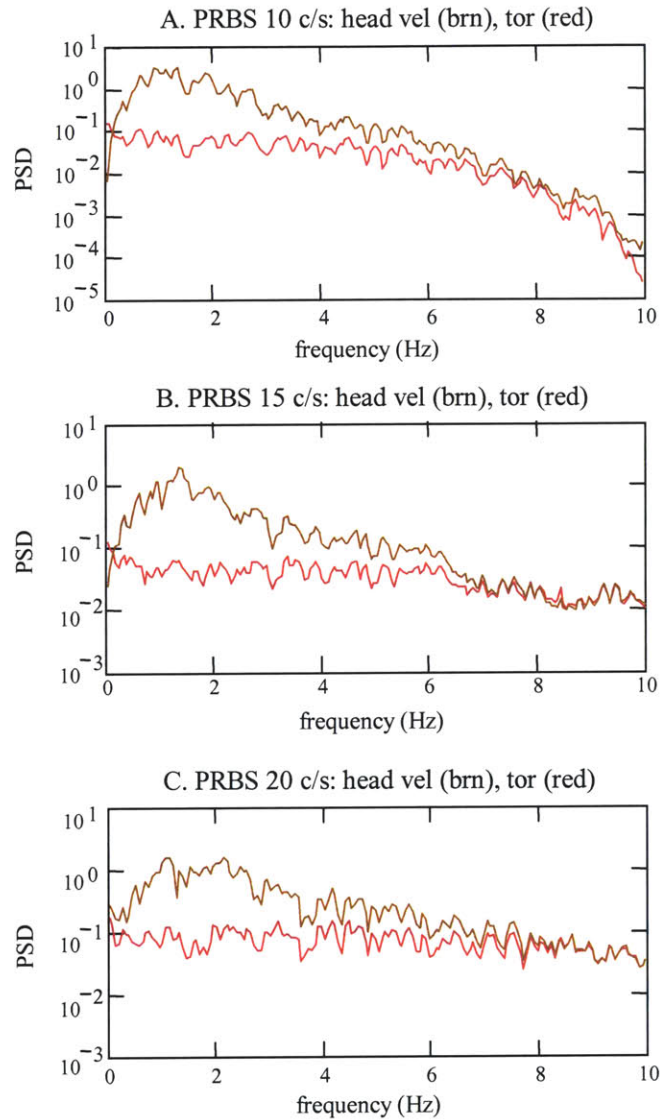
**Figure 7.13:** Comparison of the head velocity during a tracking trial when torque perturbations were applied (brown) and when no perturbations were applied (red). The power spectral densities of the sequences are shown in plot A, and the distributions of the velocity amplitudes are shown in plots B (with perturbations) and C (without perturbations).

The amplitude and spectrum of the head velocity response was surprisingly insensitive to the rate at which the PRB torque sequence was executed. Most of the trials were run at 10 changes/s, but some trials were run at 9, 15, and even 20 changes/s. The effective bandwidth of an ideal PRBS is one third of its execution frequency. The highest frequency of the effective bandwidth for PRB sequences executed at 9 or 20 changes/s occurred very near the upper frequency of the disturbance motions that the perturbations were designed to create. Figure 7.14 shows representative PSDs for the same subject when perturbations were applied at 10, 15, and 20 changes/s. It is obvious that the spectrum of the 10 changes/s PRBS decayed more rapidly than the 15 and 20 changes/s PRBS. The response of the head, however, did not vary significantly. The differences seen between the head velocity PSDs when the torque sequences were executed

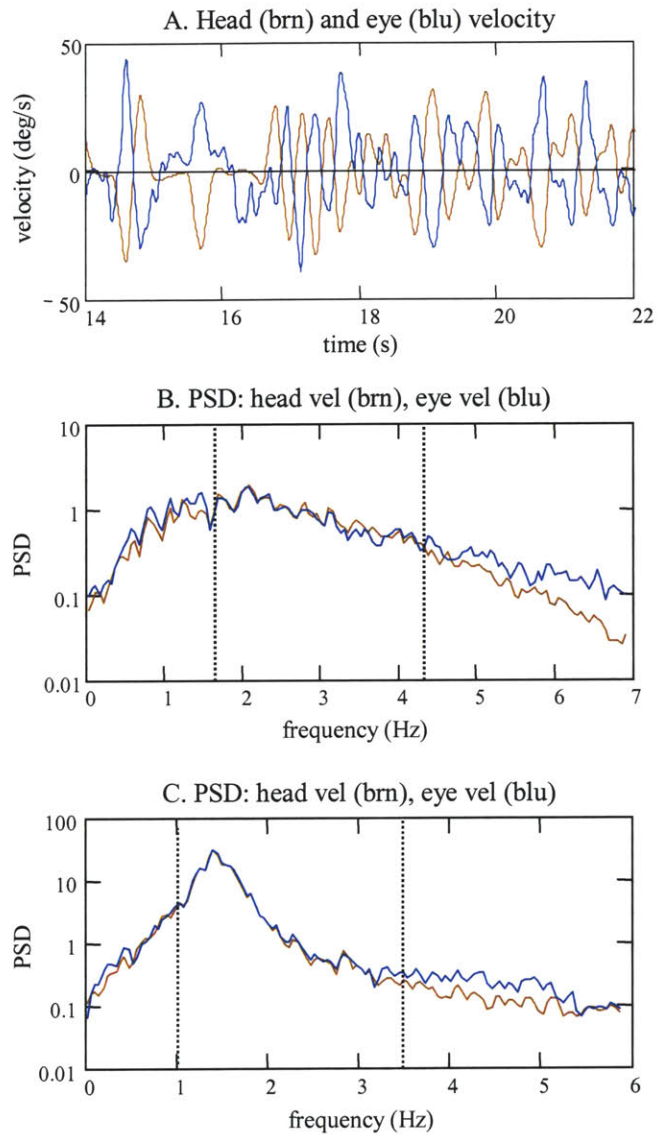


at the different rates were not much greater than the differences seen in head velocity PSDs when the same execution rate was repeated. In the range of frequencies that were of interest for the VOR analysis there was no advantage to using sequences executed at 15 and 20 changes/s as compared to sequences executed at 10 changes/s. The faster execution rates did not create a wider spectrum of disturbances, and in comments made between trials, the subjects reported that the 15 and 20 changes/s sequences felt faster than the 10 changes/s sequences, and were less comfortable. An advantage to using the slowest execution was that it concentrated the total power into a smaller range of frequencies and therefore stimulated more responses in the frequencies that were of interest.

The eye velocities were dominated by components that were nearly equal and opposite to the head velocities, as shown in Figure 7.15. The amplitudes of the head and eye responses were rarely equal, and when combined they summed to be nearly equal to the target's velocity. There was little evidence of saccadic eye motions in the eye response. The "equal and opposite" behavior of the eyes was evidence that the VOR had functioned during the head-free pursuit of the target. The VOR gain relationship could not be accurately determined by simply comparing the head and eye amplitudes since the eyes were also driven by the target, but based on inspections of subjects' data the VOR gain did appear to be close to unity. The plots of the eye velocity PSDs were unique to the subjects and to the tracking trial, but there was a common pattern in many of the trials. At the frequencies over which the target had been controlled, the head velocity and the eye velocity PSD had similar amplitudes and shape, but there were some differences. Although it would be expected that healthy subjects would have similar PSDs for head and the VOR-induced eye motions, the head and eyes were used differently to track the target and the PSD reflected these different motions. From just above the maximum frequency of the target to about 4 Hz, the PSD for the eye and head velocities were nearly identical. Above approximately 4 Hz, the eye velocity PSD deviated from the head velocity PSD, and this was probably due to noise corrupting the eye velocity calculation.



**Figure 7.14:** Representative power spectral densities for the head velocity (brown) and the torque perturbation (red) when the PRB torque sequence was executed at 10 changes/s (plot A), 15 changes/s (plot B), and 20 changes/s (plot C).



**Figure 7.15:** Head (brown) and eye (blue) velocity when torque perturbations were applied to the head (plot A). Plots B and C show representative power spectral densities for the eye and head velocity sequences for two subjects

## 7.4 Head response

The head was used “naturally” by subjects to track the targets. The motion of the head was disturbed by the torque perturbations, but the subjects controlled how the head tracked the visual target. The measured head motions could be decomposed into two components; a component that was driven by the target and that was under the voluntary control of the subject, and a component that resulted from the torque perturbations applied to the head.

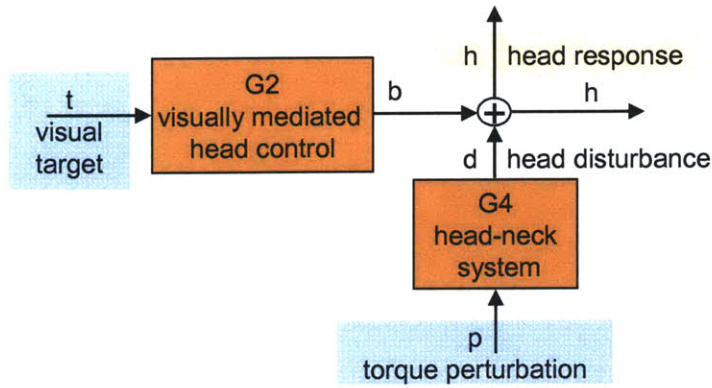
It was important to analyze the head response to assess the linearity of the head control systems. The linearity of the head control systems has a direct impact on the ability to assess the VOR with linear system identification. The eye motions caused by the VOR, and those caused by the visual system, can be distinguished only when they can be linearly related back to the torque and target inputs through the subsystems of the gaze control system. This relation is degraded when nonlinear relations exist.

### 7.4.1 Analysis of IRF and FRF for visually mediated ( $g_2$ ) and head-neck ( $g_4$ ) control systems

Head velocity can be considered the output of a two-input system (Figure 7.16). The torque perturbation ( $p$ ) stimulates a response from the head-neck system ( $g_4$ ) and the visual target ( $t$ ) stimulates a response from the visually mediated control system ( $g_2$ ). The responses of these systems, along with noise, sum to create the head velocity ( $h$ ). The linear convolution equation that relates the head velocity to the inputs is:

$$h = t \otimes g_2 + p \otimes g_4 + n, \quad \text{Eq. 7.1}$$

The calculation of the estimates of the  $g_2$  and  $g_4$  impulse response functions (IRF) is simplified because the torque and the target input sequence are uncorrelated.



**Figure 7.16:** Block diagram for the head control systems.

The IRF for  $g_4$  is calculated by taking the cross-correlation of the signals in Eq. 7.1 with the perturbation torque. The linear convolution equation becomes:

$$r_{PH} = r_{PT} \otimes g_2 + r_{PP} \otimes g_4 + r_{Pn}, \quad \text{Eq. 7.2}$$

where the linear convolution operator is represented by “ $\otimes$ ”,  $r_{PH}$  is the cross-correlation function between the torque and the head velocity,  $r_{PP}$  the auto-correlation function of the torque perturbation,  $r_{PT}$  the cross-correlation function between the torque and the target, and  $r_{Pn}$  the cross-correlation between the torque and the noise. Under the assumptions that the noise is uncorrelated with the input sequences, and the torque and target inputs are, by design, uncorrelated, Eq 7.2 reduces to:

$$r_{PH} = r_{PP} \otimes g_4. \quad \text{Eq. 7.3}$$

The impulse response function for the head-neck system could then be estimated by deconvolving  $r_{PP}$  from  $r_{PH}$  via a matrix inversion. This technique was fully described in Chapter 2.

A similar process can be used to calculate the  $g_2$  IRF. Instead of forming the cross-correlation of the signals in Eq. 7.1 with the torque input, the cross-correlation was formed with

the target velocity. The torque-target CCF and the noise-target CCF could be approximated as zero, and the resultant convolution relation was:

$$r_{TH} = r_{TT} \otimes g_2. \quad \text{Eq. 7.4}$$

The  $g_2$  IRF was estimated by deconvolving  $r_{TT}$  from  $r_{TH}$ .

An example of the analysis program that was used to estimate the impulse response functions for the head system is included in Appendix E.

#### 7.4.2 Head-neck ( $g_4$ ) and visually ( $g_2$ ) mediated control systems

##### *Head response to the torque perturbations ( $g_4$ )*

The dynamics of the head-neck systems in response to the torque perturbations resembled those of second-order, underdamped system. A representative example of one subject's data is shown in Figure 7.17. The impulse response functions and the gain and phase plots of the frequency response functions had the characteristics of second-order admittance functions (torque input, velocity output). The IRFs were fitted well by second-order parametric models. The gain of the FRF increased until it reached a resonance peak, and then decayed with nearly the same slope. The phase plot began at a phase near 90 deg, decayed and crossed 0 degrees at the same frequency as the resonance peak in the gain plot, and then asymptotically approached -90 deg, all of which are characteristic of second-order systems.

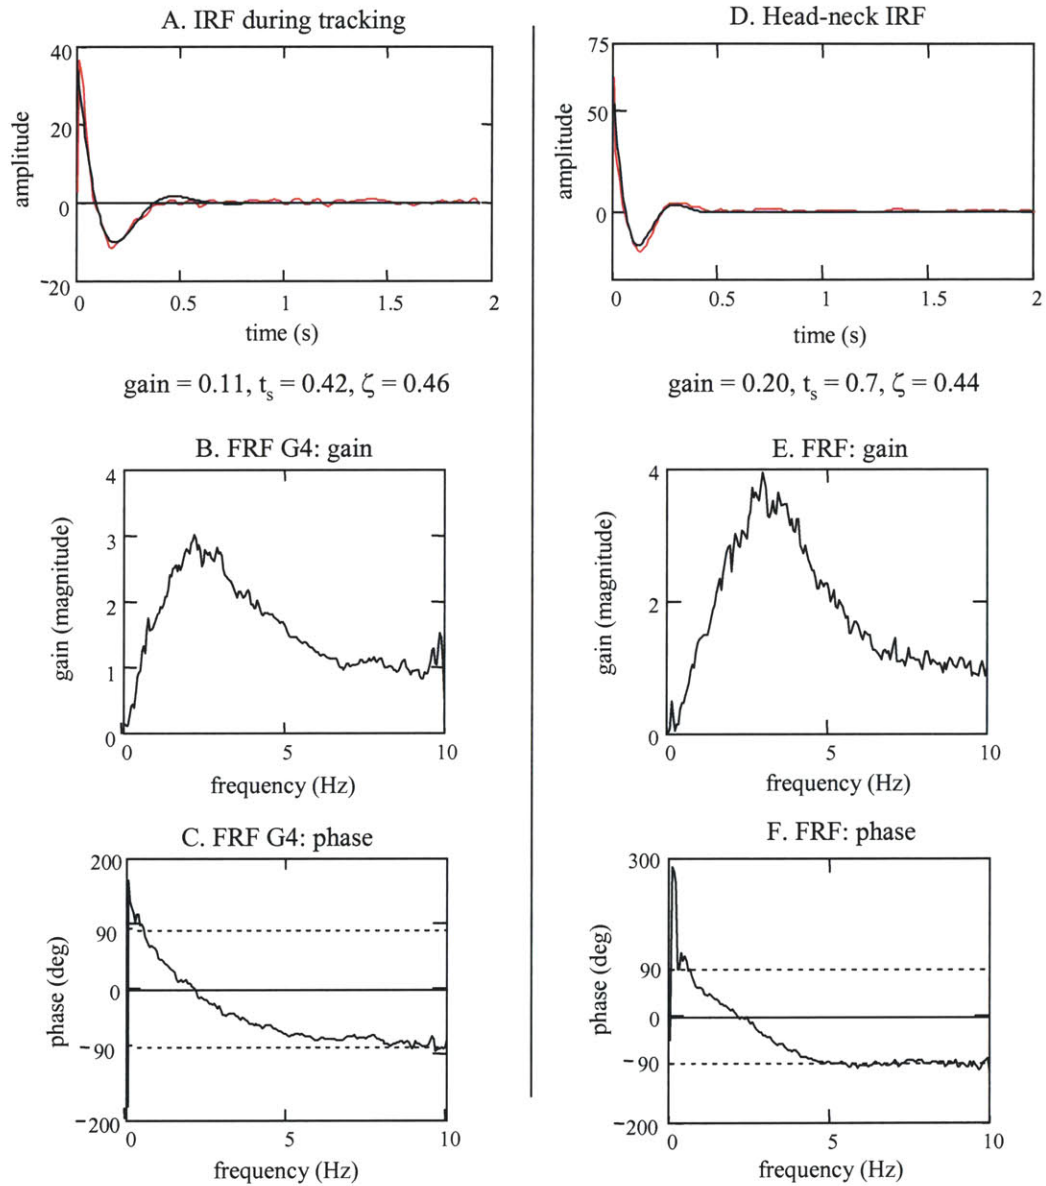
The IRFs that were calculated for the head-neck system during the VOR experiments were similar to those calculated for the same subjects during the head-neck system experiments (Chapter 6). The protocols followed during the VOR experiments and the "head-neck" experiments were very different, but similarities could be seen in the responses of the subjects who had participated in both.

In Figure 7.17, the IRF and FRF that were calculated for a subject during the VOR and head-neck experiments are shown. The response functions are not the same, but there are strong similarities. The IRFs are both fitted well by the parametric model. The parametric IRF to the left has a damping ratio of 0.46, and the IRF to the right has a damping ratio of 0.44. The peak gain in the FRF is larger in the plot at the right, but both responses peak at a frequency around 2.5 Hz, and have a similar decay as the frequency increases. Both phase plots start with a phase of 90 deg at 0 Hz, and by 5 Hz, both phase plots show a phase of -90 deg. This subject tended to have a very strong linear, second-order response during both the VOR and head-neck experiments. During the VOR experiment the non-parametric IRFs were able to predict the head output with a VAF of 96%, and the parametric IRF predicted the output with a VAF of 91%.

There are many reasons for there to be differences in the response functions calculated for the head during the VOR and head-neck experiments. During the head-neck experiments the protocol was controlled so that the head was used in the same manner in each trial. The protocol was designed to stimulate linear, time invariant behavior from the head-neck system, so that it could be compared at different equilibrium points. Subjects were required to keep the head directed forward (with help of visual feedback), and to correct for the disturbances to the head's position that were created by the applied torque perturbations. Generally, the displacements to the head were less than 10 deg. During the VOR experiments the subjects were free to move their heads as they wished. Some subjects moved the head with the target through displacements of more than 25 deg, and others kept the head centered. Considering how different the protocols were during the VOR and head-neck experiments, it is remarkable how close the response functions were.

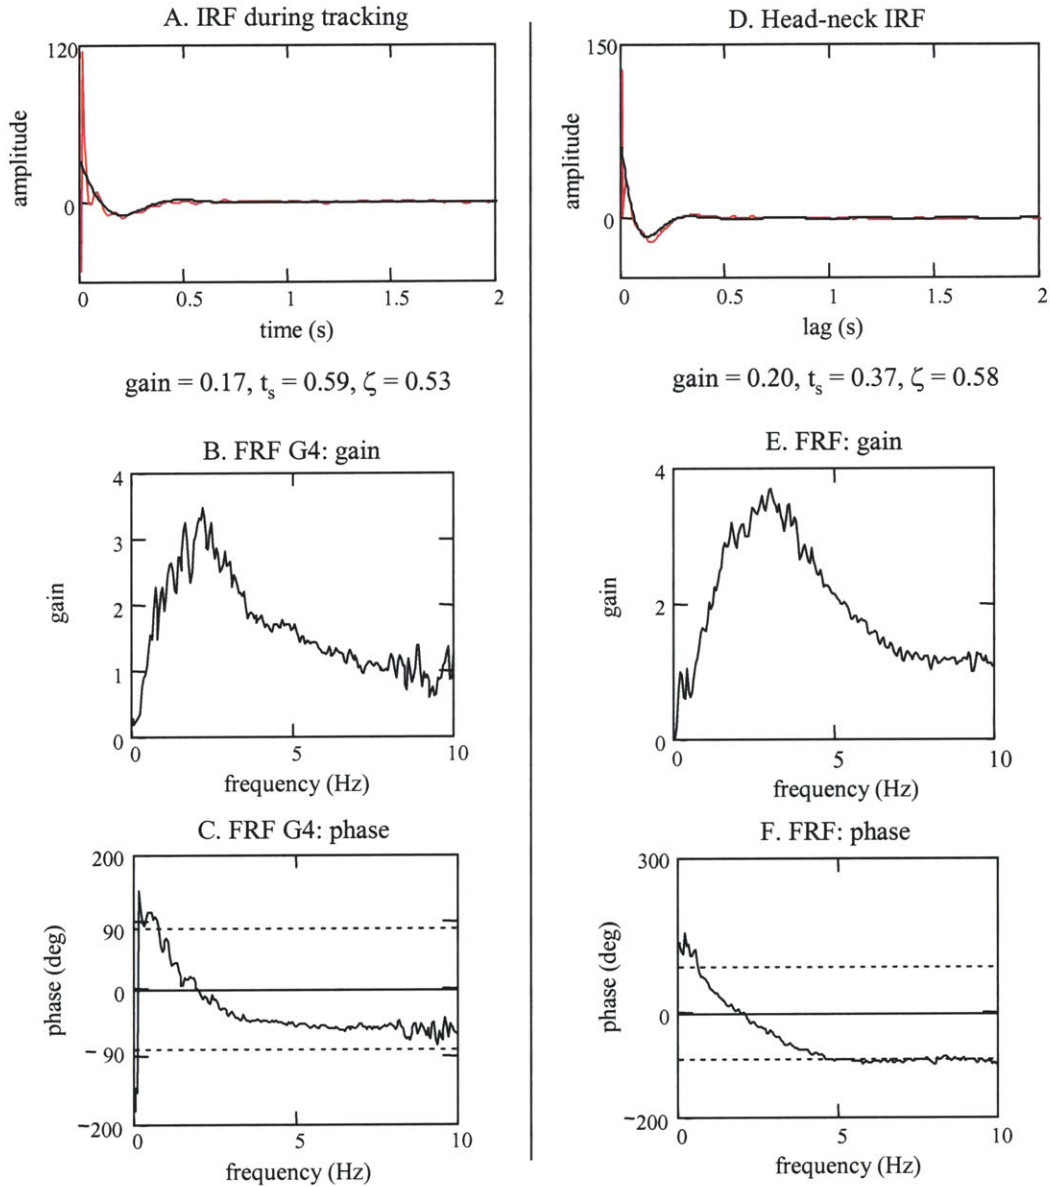
Not all subjects had head-neck ( $g_4$ ) IRFs that were so well modeled by second-order, linear systems. Figure 7.18 shows the IRF and FRF responses for another subject. The IRF and FRF are not as clean and not as closely matched by second-order models as they are for the subject shown in Figure 7.17. The responses still show characteristics of second-order systems.

There was some degree of variability in how closely the subject's response resembled an ideal, second-order system, but there was some consistency in the subject's head-neck responses during the VOR and head-neck experiments.



**Figure 7.17:** Representative impulse and frequency response functions for the head-neck system for one subject. Plots A, B and C show the results for the head-neck system during a VOR experiment. Plots D, E, and F show the results of a head-neck experiment when the mean torque that was applied to the subject's head was 0.1 N.m.





**Figure 7.18:** Representative impulse and frequency response functions for the head-neck system for a subject different from the subject in Figure 7.17. Plots A, B and C show the results for the head-neck system during a VOR experiment. Plots D, E, and F show the results from a head-neck experiment.

### *Voluntary, visually mediated response (g2)*


The g2 IRF represented the relation between the target and the subject's head movements. This relation was not as strong as that for the g4 IRF as subjects were free to decide how much head motion to use to follow the target. The target did not directly drive the head motions, but drove the gaze. How the gaze response was split between head and eye motions was up to the subject, and since the subject did not have to maintain a consistent strategy during the trial, how the head was used could vary. There tended to be three types of head uses among the subjects. Some subjects moved the head closely with the target (Figure 7.11), whereas others kept the head near 0 deg and followed the target mostly with eye motions (Figure 7.12), and others used a combination of the two.

The gain and phase plots of a G2 FRF are shown in Figure 7.19. The plots are representative of the class of subjects who tended to use a combination of head and eye movements to track the target. In this example the target trajectory had a frequency spectrum between 0 and 1.8 Hz. The gain plot shows that the subject's head motions had a gain of about 0.6 across throughout this frequency range. The subject moved the head with the target even at the higher velocities. On average the phase increased linearly with frequency, as is shown in the phase plot. Therefore the head response lagged the target by a nearly constant time lag across all frequencies. Since the targets were unpredictable, this suggests that the visually driven response was in fact driven by visual mechanisms and not by prediction. In most cases, independently of whether the gain was high or not, the phase indicated that there was a delay between the target's motion and the response of the head.

### *Predictions of the head response*

Although the head and neck are controlled and stabilized by nonlinear systems, the linear g2 and g4 IRFs were usually able to capture a great deal of the head's dynamic response during the tracking trials. Predictions of head velocity sequences were made by convolving the input

sequences with the g2 and g4 IRFs. Comparisons of the predicted head velocities with the measured sequences showed where the linear analysis captured the dynamics of the head during the head-free tracking, and where the head-neck system had nonlinearities that could not be modeled well by linear system identification.



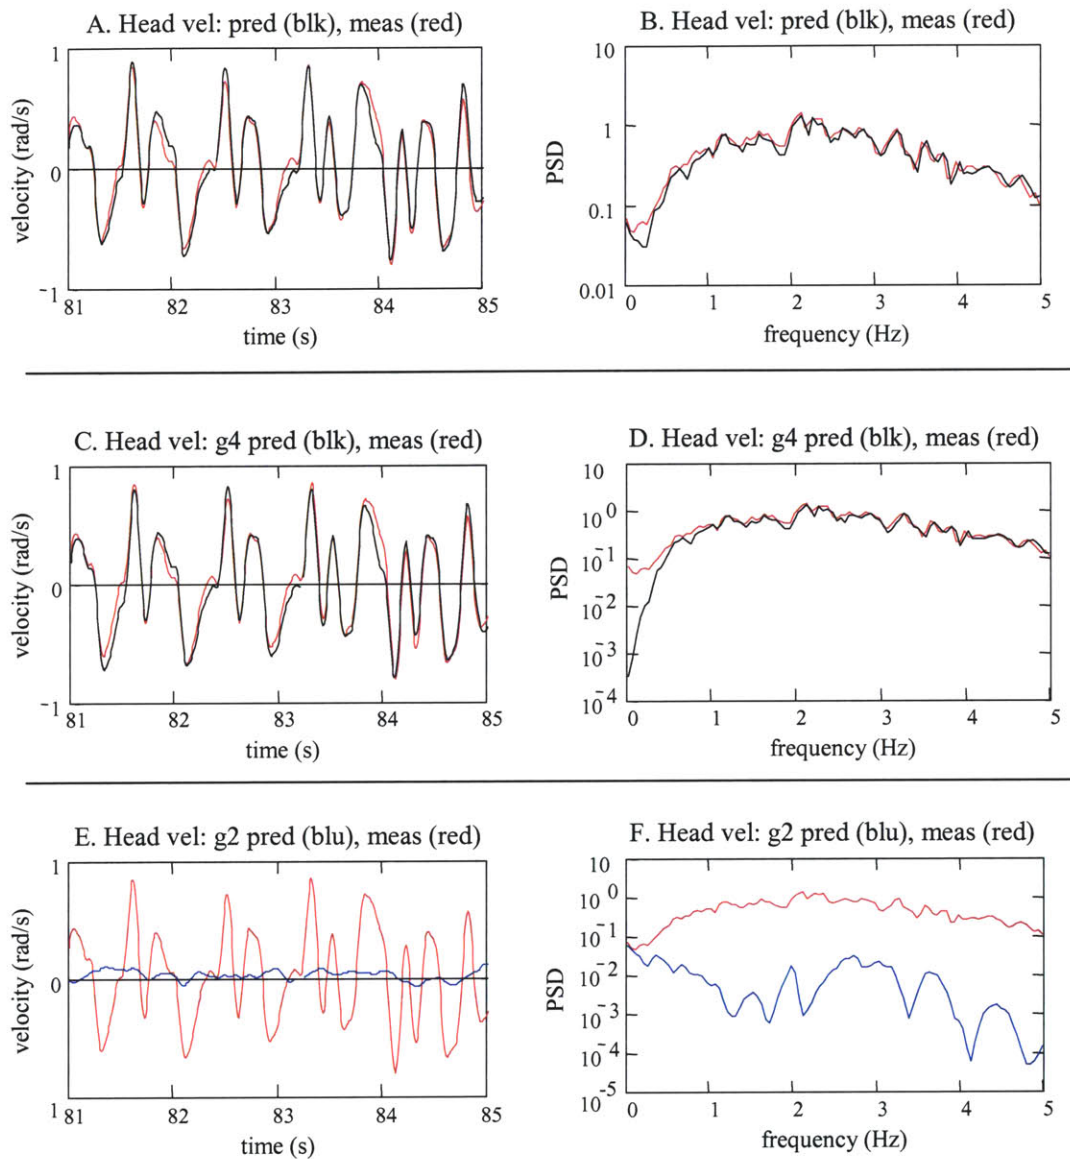
**Figure 7.19:** Representative gain (plot A) and phase (plot B) response plots for the visually mediated head control system during the tracking of a 0.0 to 2.0 Hz visual target.

Figure 7.20 shows a representative example of head velocity, its linear prediction, and the breakdown of the prediction into a visually driven component and a torque driven component. Overall, the linear IRFs were able to capture a great deal of the dynamics of the head-control systems. The prediction of the head velocity had a VAF of 92%. It can be seen from plots of the power spectral density functions that the predicted PSD matched the measured PSD at all frequencies, with one exception. Between 0.2 and 0.5 Hz, the linear models did not predict the head velocity well. Just below 0.5 Hz, the PSD of the predicted response deviated from the PSD of the measured head velocity. The two match again, however, near 0.0 Hz. These frequencies mark the two regions where active control dominates the head response, and where the head response is dominated by a reaction to the torque. The g3 IRF accounts for most of the head's dynamics at frequencies above 0.5 to 1.0 Hz. In plot D of Figure 7.20, the PSD of the torque


driven component matches the PSD of the head at frequencies above 0.5 Hz, but drops by over three orders of magnitude between 0.5 Hz and 0.0 Hz. Although the torque perturbation has power down towards 0.0 Hz, the active control of the head prevents the torque from disturbing the motion of the head at low frequencies, as had been observed in the head-neck experiments. In plots E and F the visually driven component of the head velocity, and its PSD, is shown. The target driven component matches the measured head velocity below 0.2 Hz. In this case, only a small part of the measured velocity of the head is predicted by the linear g2 system. More of the head motion may in fact be visually mediated, but its dynamics are not captured by the linear g2 IRF.

It can be seen from plots of the head position how the visually mediated and the torque driven components contributed to the head's overall motions. Figure 7.21 shows the predicted velocities integrated into position sequences. The total head response can easily be seen to be the sum of two components. The visually driven movement of the head follows the low frequency motion of the head, and the torque driven component captures the higher frequency motions. As was intended by the experimental protocol, the torque driven component appears as a disturbance to the head's motion as the head is used to track the target.

The characteristics of this example were common to most of the predicted head velocity PSDs. The predicted responses matched the measured head velocity well at all frequencies near 0.0 Hz and at all frequencies above 1.0 Hz. In some range between 0 and 1.0 Hz, the PSD of the predicted response did not match the PSD of the measured head velocity. The response in this area can probably be attributed to nonlinear active control. There are certainly head velocities that arise from the natural use of the head to track the target, and from the torque disturbing the position of the head, but these dynamics were not captured by the linear analysis of the g2 and g4 systems.



**Figure 7.20:** Linear predictions of the head velocity. The plots on the left show the measured head velocity (red) and the predictions (black or blue). The plots on the right show the corresponding power spectral densities. Plots A and B show the prediction of the head velocity. Plots C and D show the component of the head velocity that was driven by the perturbation torque. Plots E and F show the component of the head velocity that was driven by the visual target.



**Figure 7.21:** Prediction of the head position. The position was calculated by numerically integrating the predictions of the head velocity. The visually driven component is shown in blue, the component driven by the torque perturbation is in black, and the complete head prediction is shown in red.

#### 7.4.3 Coherence-squared functions: system nonlinearities

What is most important to the VOR analysis is that there is a strong, linear output-input relation through the head control systems, especially through the head-neck system ( $g_4$ ), not that the  $g_2$  and the  $g_4$  systems can be modeled by familiar parametric systems. The modeling of the  $g_4$  system as a second-order system shows that the head-neck system is relatively invariant, but is not essential to the identification of the VOR system. What is important is that the head responds linearly to the applied torque perturbations over a wide range of frequencies. This allows the eye motions that were created by the VOR to be linearly related to the torque perturbations through the head-neck system ( $g_4$ ). Since the torque and the target are the controlled experimental inputs, the measured eye responses can be separated into components stimulated by the torque, and those stimulated by the target, only if there is a linear relation between the inputs and the output. Nonlinearities in the head-neck system degrade the linear cascade of systems between the controlled torque input, and the VOR output, and limit how well the linear system identification technique can evaluate the VOR's performance.

The coherence-squared functions between the torque and target inputs, and the head velocity output, show where the measured head velocity was linearly related to the inputs, and

identify where the linear relation was corrupted by noise and nonlinear system behavior. The coherence-squared function essentially decomposes the “percent variance accounted for” that was calculated between the prediction of the head velocity and the measured head velocity by frequency. The measured head velocity is the output of a two-input, single-output system. The interpretation of the coherence-squared function for a two-input, single-output system is not as straightforward as the coherence-squared function for a single-input single-output system (Juang, 1994). The coherence-squared function for a two-input single-output system is the sum of three groups of terms (Chapter 2). There are coherence-squared terms between each input and the measured output. These are calculated in the same way as for a single input-single output system. There is a coherence-squared measure between the inputs, and two coherence-squared terms that simultaneously relate both inputs and the output (“triple coherence-squared”). The sum of the three groups of terms (coherence-squared sum) is bound by unity, as opposed to the case for a single-input single system, where the coherence between the single input and the output is bound by unity. The interpretation of the terms in the coherence-squared sum is simplified when the input signals are uncorrelated, as was the case in the VOR experiments.

Relative to the coherence-squared between either the torque or the target and the head velocity output, the coherence-squared between the inputs, and the “triple coherence squared” terms, are about zero. When these terms are ignored, the coherence-squared functions between the torque and the head velocity, and between the target and the head velocity, can be interpreted similarly to that for a single-input single-output system. Each coherence-squared function can be interpreted as the fraction of the head velocity’s total response that was linearly related to the particular input. In the single-input single-output system, the coherence-squared decreases with noise and nonlinearities. In the two-input, single-output system, noise and nonlinearities still degrade the coherence-squared value, but the coherence-squared between one of the inputs and the output is also reduced from unity by the linear contribution to the output from other inputs. The relationship between the input and its component of the output may be strictly linear, but if

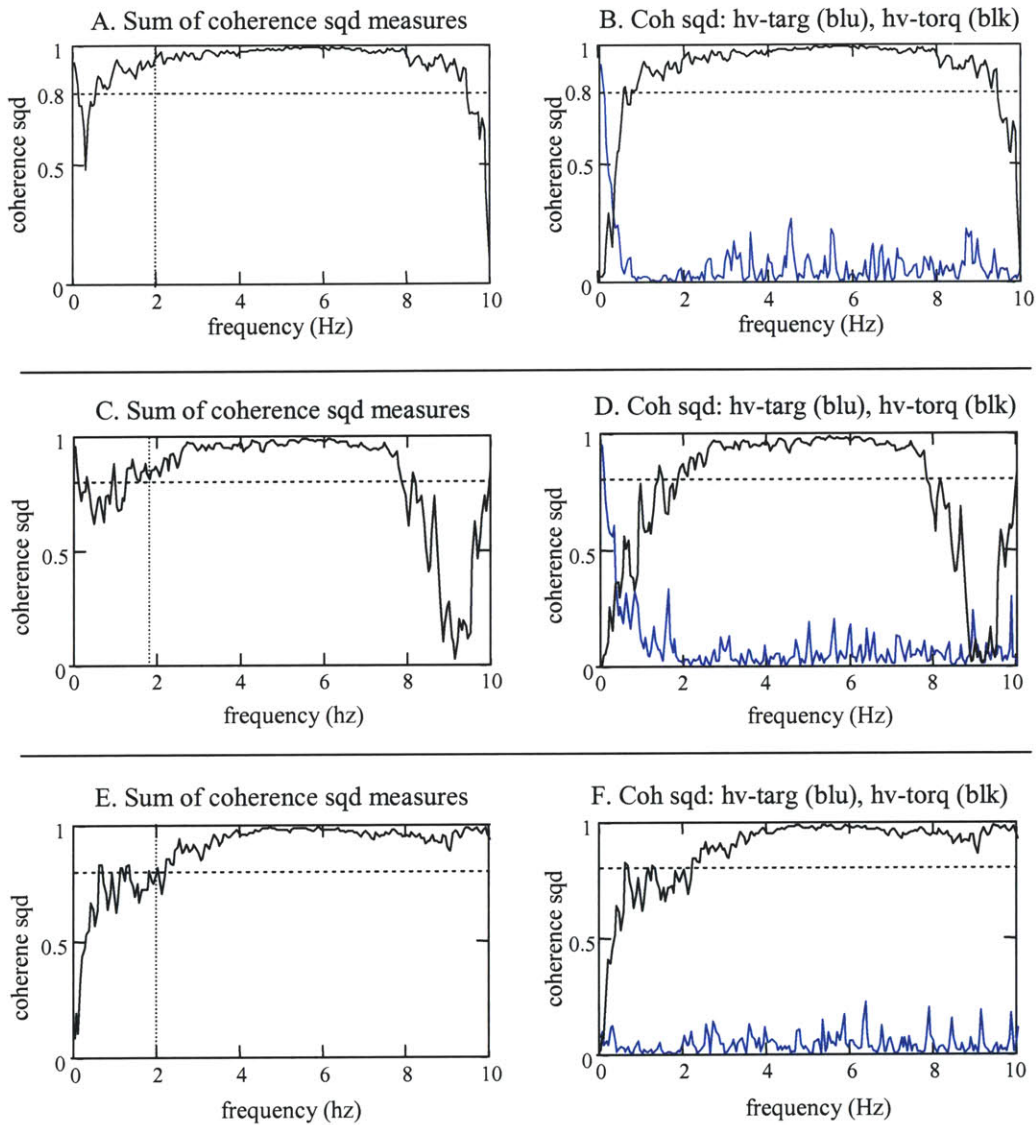
the measured output is also stimulated by other inputs, or degraded by noise or nonlinearities, the coherence-squared measure will be below unity.

The relations between the measured head velocity and the torque and target inputs can be divided into two parts: one component that involves frequencies above the target's spectrum, the other covers frequencies where the target was controlled. Coherence-squared functions that are representative of data from three subjects are shown in Figure 7.22. As a general rule, the frequency spectra of the visual targets used in the VOR experiments extended from 0.0 Hz to between 1.5 and 2.0 Hz. At frequencies above the target's spectrum, the coherence-squared plots indicated that the motion of the head was caused predominantly by the torque perturbation. This was consistent with the PSD of the predicted head velocity. As shown in Figure 7.20, the PSD of the head velocity component that was driven by the torque matched the PSD of the head velocity closely at all frequencies above 1.0 Hz. At these frequencies, the coherence-squared between the torque and the head velocity was the main contributor to the coherence-squared sum. The coherence-squared between the target and the torque was near zero, and the plot of the coherence-squared between the torque and the head velocity was nearly identical to the figure that showed the sum of the coherence-squared measures (Figure 7.22). The high coherence-squared values ( $>0.9$ ) indicated that there was a strong, linear relationship between the head velocity and the torque perturbation. The relationship was not degraded by noise, nonlinearities, or contributions to the head velocity from other inputs.

At the highest frequencies for which the visual targets were controlled, the coherence-squared was near, or greater than, 0.8. It improved as the frequency increased. By 4 Hz the coherence-squared was near unity, and it remained near unity until the frequency approached the rate at which the torque PRBS was executed. At this frequency the coherence-squared dropped rapidly since the power in the input torque dropped considerably, and the signal to noise ratio became low. In the top plot of Figure 7.22 the torque had been executed at 10 changes/s, in the



middle it was executed at 9 changes/s, and in the bottom plot it had been executed at 15 changes/s.



**Figure 7.22:** Coherence-squared functions for three subjects. The plots on the left show the complete coherence-squared function (as described in Appendix A) between the torque and target inputs and the head velocity output. The plots on the right show the coherence-squared functions between the torque and the head velocity (black) and the target velocity and the head velocity (blue). Plots A and B are representative of the data from one subject who had an exceptionally linear relation between the head response and the torque perturbations. Plots C and D are representative of the majority of the trials for all subjects, except for the subject shown in the plots A and B. Plots E and F are representative of trials where a subject tracked the target primarily with eye motions, and did not voluntarily move the head to follow the target. In these cases, the head was moved by perturbations alone. The upper frequency of the visual target’s spectrum is marked by a dotted vertical line

The coherence-squared sum was usually lower at the frequencies for which the visual target had power than it was at higher frequencies. In the majority of trials, the coherence-squared remained above 0.7 between 0.0 and 2.0 Hz. This indicated that there was a strong linear relationship between the head velocity and the torque and target inputs, although the linear relation was not as strong as it was at frequencies for which the head was driven only by the torque perturbation. A characteristic that was common to most trials for which the subject had used head motions to track the target, was that the lowest coherence-squared values occurred between 0.2 and 1.0 Hz. The coherence-squared was often near unity at 0.0 Hz. It would quickly drop to its lowest value(s) in the frequency range from approximately 0.2 to 1.0 Hz, and then would increase towards the end of the spectrum where the target had power. The plots of the coherence-squared values between the target and the head velocity and the torque and the head velocity show how the drop in the coherence-squared developed. Near 0.0 Hz the head velocity was strongly dependent on the target, and then this linear dependence quickly decayed toward zero. As this decayed, the coherence-squared between the torque and the head velocity increased. In the region from 0.2 to 1.0 Hz a portion of the head velocity was nonlinearly related to the inputs. This is the region where active control and stability of the head prevents the torque inputs from stimulating head disturbances in a linear fashion. These results are consistent with the results of the head-neck study of Chapter 6; below 0.5 Hz the head-neck response to horizontal plane torque perturbations was not modeled well by a linear IRF.

At the lowest frequencies, the torque perturbations did not stimulate head movement, but the head velocity was linearly related to the visual target's motion. As the frequency increased from 0.0 Hz, there was less of a linear response from the head to follow the target velocity. This was also seen in the PSD of the head velocity prediction (Figure 7.20). Between 0 and about 0.3 Hz, the PSD of the head velocity component that was caused by the visually-mediated system matched the PSD of the head velocity (Figure 7.20, plot F). Beyond 0.3 Hz, the visually driven component had almost no power.

The coherence-squared functions for the head velocity demonstrated that above 0.5 to 1.0 Hz the head is linearly related to the torque perturbation, and the dynamics of the head can be modeled well by a linear IRFs and FRFs. Below 0.5 to 1.0 Hz nonlinearities in the head control and stability systems degrade how well the systems can be represented by linear models.

The linear analysis of the VOR will be degraded by the same nonlinearities. The input to the subject's vestibular system that stimulates a response from the VOR is the head disturbance, but the stimuli to the VOR, within the framework of the linear analysis, are the components of the head disturbance that can be linearly related to the applied torque. The coherence-squared function that will relate the output of the VOR to the torque perturbation will be no better than the coherence-squared function that relates the torque to the head velocity.

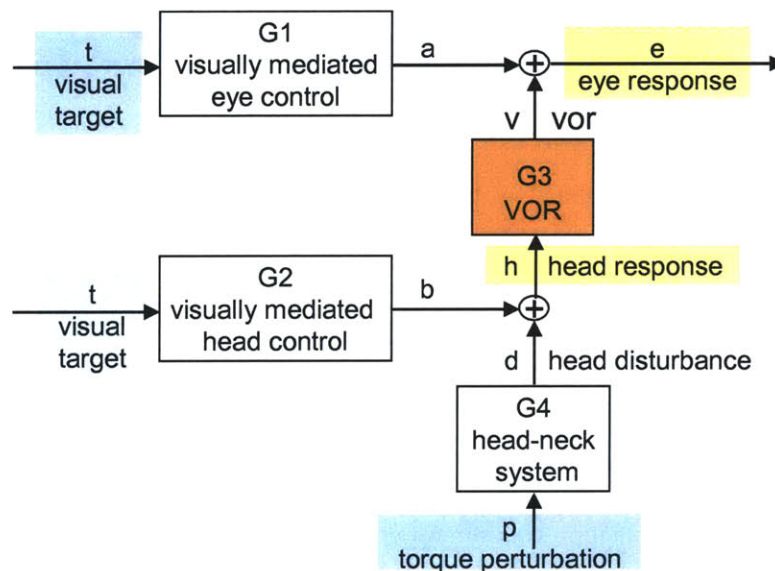
The nonlinearities of the head-neck system would not affect the linear system identification of the VOR if the VOR eye motions could be related directly to a measurement of the head's disturbance position. This is not possible unless the head's disturbance position becomes a controlled input, rather than it being a perturbation to the subject's natural head motions. If the disturbance's position was controlled, and the subject was prevented from using natural motions to track the target, the analysis would not assess the natural function of the VOR during head-free pursuit, and the normal interactions of subsystems in the gaze control system would be altered. The appropriate solution would be to use a nonlinear model to represent the dynamics of the head-neck system. That would require a very different analysis and a protocol that was more appropriate to nonlinear system identification.

## **7.5 VOR response**

One of the benefits of testing the gaze tracking system with random, broadband inputs was that the system-level response of the VOR could be estimated in both the frequency and time domains. The frequency response function (FRF) could be calculated over a wide range of

frequencies, and estimates of the impulse response function (IRF) could be made directly from the time domain data. The FRF and the IRF provided the same information about the VOR system, but presented the information differently. The system characteristics of the VOR were more easily interpreted from the gain and phase functions of the FRF.

### 7.5.1 Analysis of VOR IRF ( $g_3$ ) and FRF ( $G_3$ )



**Figure 7.23:** Block diagram of the inputs, outputs, and systems directly involved in the analysis of the VOR

The VOR system response function represented the VOR's response to the perturbations applied to the head. It was not the relation between the head velocity and eye velocity, even though those were the head and eye signals that could be directly measured. The impulse and frequency response functions that described the linear behavior of the VOR were found by decomposing the measured eye response into visually driven and VOR-driven eye motions, and relating the VOR eye motions to the movements of the head that were caused by the torque perturbations. The mathematics of the analysis was discussed fully in Chapter 2, but a summary

of the equations is presented in this section. An example of the analysis program that was used to analyze the VOR system is included in Appendix E.

The eye velocity sequence ( $e$ ) that was calculated for each tracking trial represented the sum of the eye motions from the visually mediated eye control system ( $a = t \otimes g1$ ) and from the VOR ( $v = h \otimes g3$ ) (Figure 7.23):

$$e = a + v = t \otimes g1 + h \otimes g3, \quad \text{Eq. 7.5}$$

where the linear convolution operator is represented by “ $\otimes$ ”.

The head response ( $h$ ) can be decomposed into head movements that resulted from the visually mediated head control system ( $b = t \otimes g2$ ) and disturbance motions that were caused by the torque applied by the head-perturber ( $d = p \otimes g4$ ):

$$h = b + d = t \otimes g2 + p \otimes g4. \quad \text{Eq. 7.6}$$

Since the convolution is a linear operator, Eq. 7.5 can be expanded into:

$$e = t \otimes g1 + (t \otimes g2) \otimes g3 + (p \otimes g4) \otimes g3. \quad \text{Eq. 7.7}$$

The eye motions that were in response to the visual target, either through voluntary eye control ( $g1$ ) or indirectly through voluntary head control ( $g2$ ), are represented by  $t \otimes g1 + (t \otimes g2) \otimes g3$ . The eye motions that were created by the VOR to stabilize vision during the perturbation of the head are represented by  $(p \otimes g4) \otimes g3$ .

Eq 7.7 can be changed from representing the relation between the total eye response and the torque and target inputs, to an equation that considers only eye motions that were driven by the torque perturbation, by forming the cross-correlation of the input and output sequences with the perturbation torque:

$$r_{PE} = r_{PT} \oplus g1 + r_{PT} \otimes g2 \otimes g3 + r_{PP} \otimes g4 \otimes g3. \quad \text{Eq. 7.8}$$

Since the torque perturbation sequence and the target velocity were made to be uncorrelated, their cross-correlation function can be set to zero ( $r_{PT} = 0$ ), and Eq. 7.7 is reduced to:

$$r_{PE} = r_{PP} \otimes g_4 \otimes g_3 \quad . \quad \text{Eq. 7.9}$$

Eq. 7.9 shows the linear relation between eye motions that were correlated with the torque stimulus and the torque perturbation that disturbed the head. Although the data were gathered under natural tracking conditions, during which the head and eyes interacted normally to track the target, the influence of the visual target has been removed from the analysis.

It was shown in Section 7.5 that:

$$r_{PH} = r_{PP} \otimes g_4, \quad \text{Eq. 7.10}$$

where  $r_{PH}$  represents the movement of the head that was caused by the torque perturbation,  $r_{PP}$ . Eq. 7.10 can be substituted into Eq 7.9 to create an equation that relates the VOR's impulse response function ( $g_3$ ) to the cross-correlation functions that represent the head and the eye motions that were driven by the torque perturbation:

$$r_{PE} = r_{PH} \otimes g_3, \quad \text{Eq. 7.11}$$

Eq 7.11 can be developed more quickly from Eq 7.5 by forming the cross-correlation functions between the torque perturbation and the head and target velocity input sequences, and the eye velocity output sequence. This is unorthodox, as it is rare, in the analysis of a two-input, single-output system to form the cross-correlation of the inputs and the output with a third sequence that is not also one of the inputs or outputs in the analysis equation. The procedure is justified, however, because it does not involve any physical relations that did not hold in the development of Eq. 7.11. The result is:

$$r_{PE} = r_{PT} \otimes g_1 + r_{PH} \otimes g_3. \quad \text{Eq. 7.12}$$

Since  $r_{PT}$  is approximated as 0, Eq. 7.12 simplifies to Eq. 7.11

The VOR impulse response function ( $g_3$ ) was estimated using Eq. 7.11 by de-convolving  $r_{PH}$  from  $r_{PE}$ . Eq. 7.11 was expanded into a series of equations and then represented in a matrix. The de-convolution was conducted via a matrix inversion. The process was fully described in Chapter 2.

The frequency domain equations extend immediately from the time domain equations. The linear convolution in the time domain becomes a multiplication in the frequency domain, and the time domain cross-correlation functions become the frequency domain cross-power spectral density functions. In the frequency domain Eq 7.11 becomes:

$$\bar{S}_{PE}(j\omega) = \bar{S}_{PH}(j\omega) \cdot G3(j\omega). \quad \text{Eq. 7.13}$$

The frequency response function of the VOR ( $G3$ ) was found by dividing the cross-PSD between the eye velocity and the perturbation torque by the cross-PSD between the head velocity and the torque perturbation. In the text *Applied System Identification* (Juang, 1994), Juang proposed that cross-power spectral density functions should not be calculated directly from a full data record, but from the average of smaller records. The data records were broken into smaller records, and the cross-PSDs for Eq. 7.13 calculated as the average of the cross-PSDs of the smaller records. This reduced the variance of the cross-PSD functions as estimators of the cross spectral density of the random process that created the data, and improved  $G3$  as an estimator of the system performance.

The VOR dynamics that can be solved for by Eq. 7.11, are those that are stimulated by the head disturbances that can be linearly related to the input torque, not those that are stimulated by all head disturbances. Eqs. 7.9 and 7.10 make it clear why this occurs. The perturbation torque was one of the two controlled inputs used in the experiment, with the visual target being the other input. The torque perturbation did not, however, stimulate the VOR. The VOR ( $g3$ ) was stimulated by head disturbances that were caused by the torque, and the linear equations only accounted for disturbances that could be linearly related to the torque perturbation ( $p \otimes g4$ ). Although the torque perturbation sequences had power at frequencies from 0 through 8 to 10 Hz, the linear stimulus to the VOR did not. The linear analysis could not therefore assess the performance of the VOR over all frequencies for which there was power in the torque sequence.

The VOR could only be evaluated over frequencies for which the head disturbance that was linearly related to the torque,  $p \otimes g_4$ , had power.


One of the fundamental assumptions made during the development of the equations that were used to identify the  $g_3$  response was that the torque and target sequences could be approximated as uncorrelated. This allowed  $r_{PT}$  to be set to 0 and to be removed from the analysis equations. Without this assumption,  $g_3$  could not be identified. When the input sequences were being designed, the torque and target sequences were accepted as uncorrelated when the magnitude of their cross-correlation was small relative to the auto-correlation function of the PRB torque sequence (Chapter 4). The ACF of the PRB torque sequence was an appropriate measure with which to judge the sequences as uncorrelated, but it was also the only measure available during the design process with which to judge the level of the torque-target correlation. Once the tracking trial had been run and the data collected, other measures became available that could be used to assess the influence of the torque-target CCF on the analysis and to justify setting the torque-target CCF to zero.

Eq. 7.12 shows the relation between the two eye-control systems,  $g_1$  and  $g_3$ , and the three signals that can be measured: target velocity, torque perturbation, and head velocity. When  $r_{PT}$  is set to zero the influence of the  $g_1$  system on the output measure,  $r_{PE}$ , is assumed to be insignificant. Without knowledge of the  $g_1$  system, it is not known how sensitive the system is to small values of  $r_{PT}$ . To be conservative in accepting the assumption that  $r_{PT} \otimes g_1$  does go to zero,  $r_{PT}$  should be small relative to the other cross-correlation values in the equation,  $r_{PE}$  and  $r_{PH}$ , not just small relative to the ACF of the torque, which is not explicitly part of the equation from which  $g_3$  is determined.

Typical cross-correlation functions are shown in Figure 7.24. The plots show that the cross-correlation between the target and the torque was much smaller than either the CCF between the eye velocity and the torque (the output measure) or the CCF between head velocity



and the torque (the input measure). In this example the two-norm of the eye velocity-torque CCF and the two-norm of the head velocity-torque CCF are 114 and 108 times the two-norm of the target velocity-torque CCF. Across all experiments, it was common for the two-norm of the target-torque CCF to be 40 to 120 times smaller than the two-norm of the CCF between either the eye velocity, or the head velocity, and the torque. These results supported setting  $r_{PT}$  to zero. The torque and target sequences that had been designed to be uncorrelated, and that had retained their qualities when they were converted into physical inputs and sampled, did in fact end up being insignificantly correlated when compared to the other sequences used in the analysis equations.



**Figure 7.24:** Cross-covariance function (CCVF) between the torque and target (red), the torque and eye velocity (blue), and the torque and head velocity (brown). The 2-norm of the torque-eye velocity CCVF is 50 times the 2-norm of the torque-target CCVF, and the 2-norm of the torque-head velocity CCVF is 45 times the 2-norm of the torque-target CCVF.

## 7.5.2 VOR FRF: gain, phase, coherence-squared, high pass models

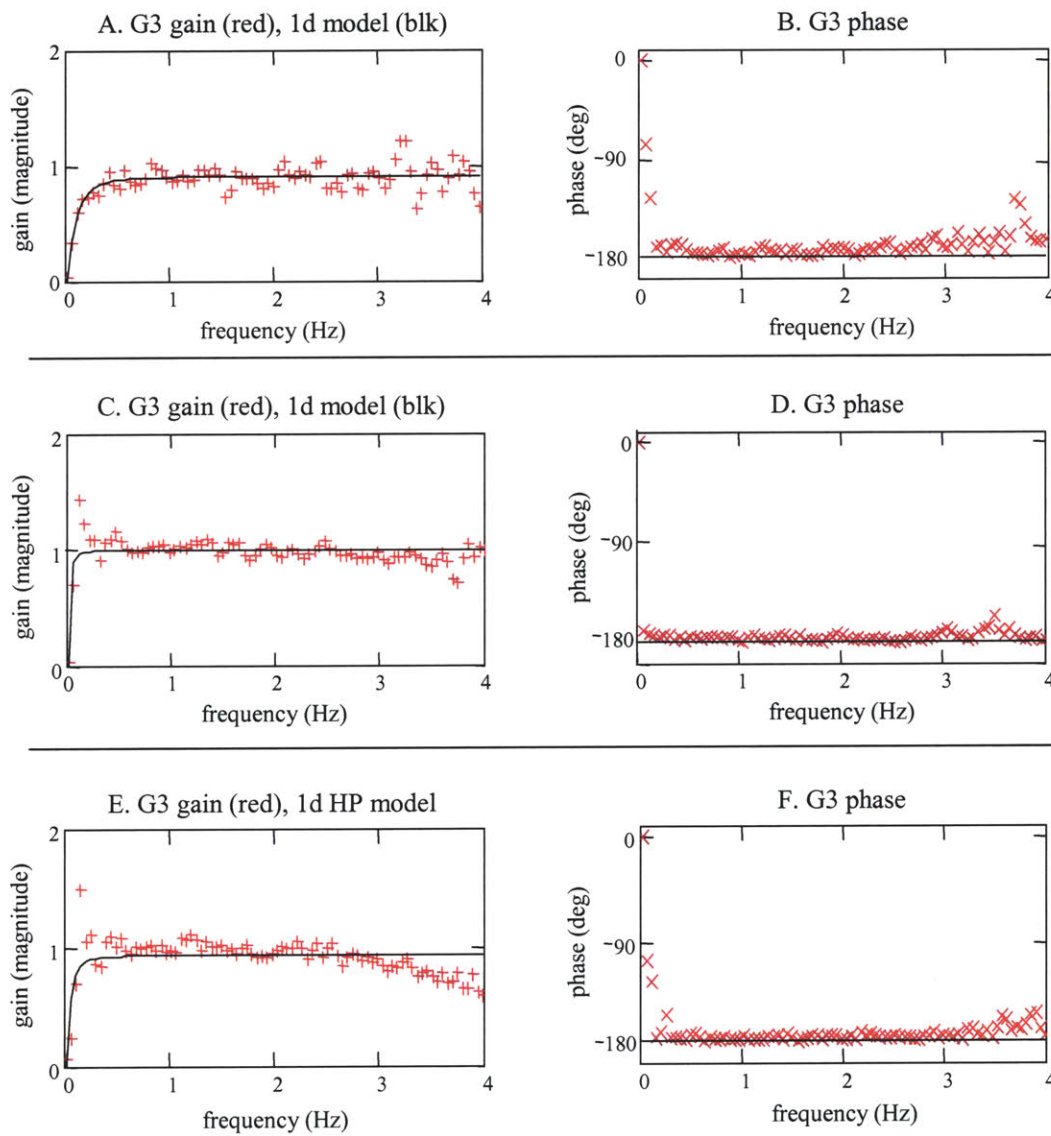
### *Gain and phase plots*

The frequency response functions that were calculated for the subject's VOR resembled high-pass systems. Gain and phase plots that are representative of the results from three of the subjects are shown in Figure 7.25. The FRFs usually began with both the gain and the phase near zero at 0.0 Hz. The gain values quickly rose as the frequency increased, and usually between 0.3 and 0.7 Hz settled near unity, which indicated that the VOR eye response was approximately

equal in magnitude to the head disturbance. The gain values remained close to unity until between 3 to 5 Hz. Above this, the difference between consecutive gain values increased, although the average of the gain values was usually near unity. An exception to this was in the results for the subject shown in plots E and F, whose gain consistently dropped between 3 and 5 Hz, but then returned to an average near unity at higher frequencies. This was the subject who had used very little head motion to track the targets (Section 7.4), and as was noted previously, his results tended to be slightly different from those of the other subjects. The phase plots usually began at 0 deg and by 0.1 to 0.3 Hz had settled to near 180 deg. Even at the lowest frequencies, the FRF indicated that the VOR moved the eye directly opposite to the head's disturbance motions, with little or no time lag between the head disturbance and the compensatory VOR motion. The phase remained close to 180 degrees as the frequency increased, but just as with the gain plot, there was an increase in the variance in the gain values at frequencies above 3 to 5 Hz.

The VOR is traditionally modeled as a high-pass system, so it was reassuring to see that the FRFs from this analysis also looked like high-pass systems. However, traditional modeling of the VOR usually shows that the corner of the high-pass VOR model is at a lower frequency than was found in these results (Leigh and Zee, 1999). The difference in testing conditions was probably responsible for the change in the corner frequency that was calculated for the high pass models. Traditional VOR testing evaluates the vestibular system in the dark or with opaque glasses over the subject's eyes, with the hope of removing the influence of the visual system on the measured eye response. The visual system is not controlled, and any influence of the visual system on the stabilization of the eyes cannot be distinguished from the VOR's response. Our testing protocol evaluated the VOR under natural conditions where visual and vestibular interactions were an important part of the protocol. The vestibular system was thus evaluated as it interacted with the visual system to stabilize the vision, but the contributions of the visual and vestibular system could be distinguished because the visual and vestibular inputs were uncorrelated. Because the visual and vestibular oculomotor control systems share common neural

pathways and control loops, it is not unreasonable to expect that the VOR component of a combined visual-vestibular response would indicate that the VOR's response had gain that improves toward unity at frequencies where the visual system is no longer effective at stabilizing the gaze. As the frequency of the head disturbance increased towards and beyond 1.0 Hz, the ability of the visual system to stabilize the vision degraded. It is thus at these frequencies that the VOR was needed to stabilize the gaze and its gain response approached unity.



**Figure 7.25:** Representative gain and phase responses for the VOR for three subjects.

The gain and phase plots had values that were within the anticipated range. The tracking performance of subjects during the tracking trials indicated that their VORs should have a gain near one and a phase near 180 degrees. The subjects had been able to track the targets well, and there was little difference in their gaze performance independently of whether perturbations were applied to the head or not. Obviously, their horizontal VORs had functioned well and had stabilized their vision in response to the head disturbances.

Although the low frequency gain and phase values were consistent with what had been expected for the VOR, the values had to be viewed as with caution. ("Low frequency" refers to the frequencies before the gain had settled near unity which was from 0.0 Hz to between 0.3 and 0.7 Hz, depending on the subject and the trial). The low frequency gain response was sensitive to how the cross-power spectral densities (cross PSD) that were used in the FRF calculation ratio were themselves calculated. The cross-PSDs were calculated from finite length, sampled data records, and were thus estimators of the true cross-PSDs for the processes that had produced the data. The shape of the low frequency gain response was often affected by small changes in the number of records that the data sequence was divided into. If the analysis had captured the system level response well, small changes in how the data were divided and averaged would have resulted in only small changes in the gain plot. As shown in plot A of Figure 7.25 the low frequency gain responses often looked like a first order system. When the number of records was changed the shape of the gain response sometimes became more characteristic of a second-order high-pass system and the gain would peak before it settled near unity, rather than climb steadily to its high frequency gain value.

The gain response at frequencies above the low frequency region was much less sensitive to how the data were broken into records. The gain values changed by a very small amount when a different number of records was used to calculate the cross spectral densities. The variance in the gain values was reduced when data from numerous trials were averaged, but the basic shape of the response and its average values were not significantly affected.

The gain response was sensitive to how the cross-power spectral densities were calculated which provides an indication of the frequencies over which the linear analysis of the VOR did and did not hold well. The gain response was affected most at the same low frequencies for which the linear analysis of the head-neck system had been degraded due to the nonlinear relationship between the torque perturbations and the head velocity. This was the region where the head was stabilized by active control, and where the drive to the head changed from coming primarily from the visual system to mostly from the torque input. The coherence-squared functions between the head velocity output and the target velocity and torque inputs were shown in Figure 7.5.8 for the same tracking trials that are shown in Figure 7.25. The torque perturbations that were applied to the head had power down to 0.0 Hz, but the coherence-squared plots show that the head did not have a strong linear response to the torque input below about 0.3 to 0.7 Hz. There was little power in the head disturbances that could be linearly related to the torque perturbations at these frequencies. The G3 FRF related eye velocities to the head disturbances that were created by the torque, but there was little input power from which to analyze the G3 response. Although gain and phase values could be calculated for the VOR at these low frequencies, the values do not necessarily reflect the dynamics of the system. Above 0.3 to 0.7 Hz the strength of the response between the head velocity and the torque perturbations implied that the linear analysis of the VOR held well. The performance of the VOR could therefore be captured by a linear system analysis over the frequencies for which the system was used to stabilize gaze. This linear analysis was limited at low frequencies by the nonlinear response of the head to the applied torque perturbations.

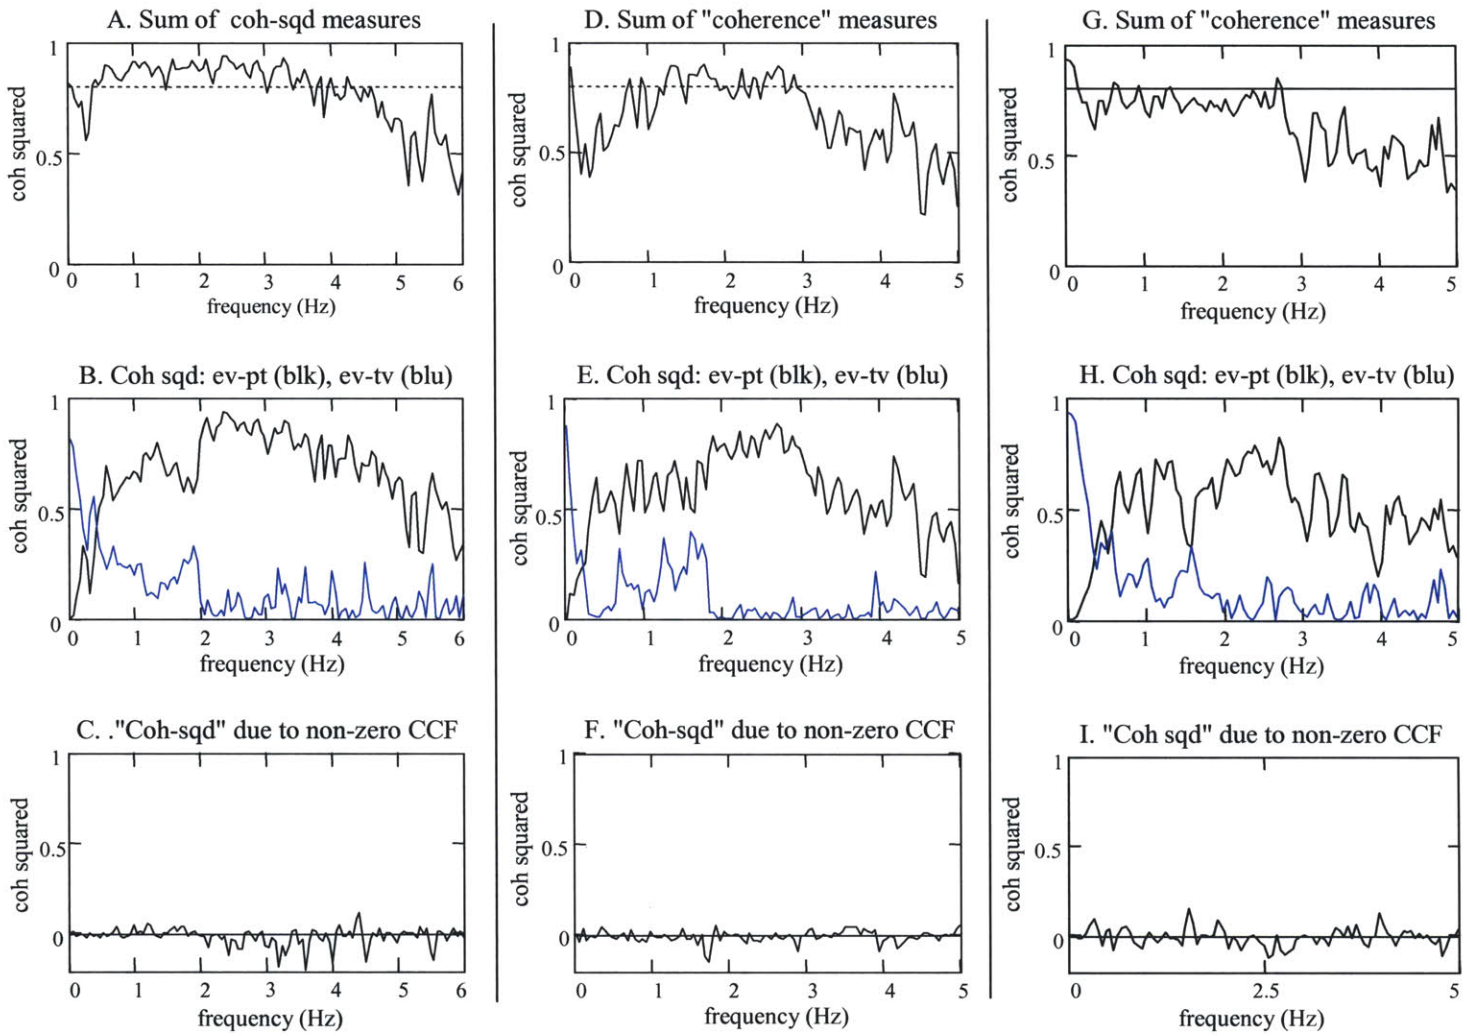
#### *Coherence-squared functions*

How well the linear relations between the torque and target inputs and the eye velocity output captured the dynamics of the VOR, and at what frequencies the analysis was degraded by nonlinearities and noise in the data, was more rigorously analyzed using coherence-squared

functions. The eye velocity was the output of a two input system (Figure 7.23). Similar to the analysis in Section 7.5 for the head velocity, coherence-squared functions were calculated to assess the linearity of the subsystems that transformed the torque and target inputs into eye movements, and to evaluate the quality of the data. The complete coherence-squared function was the sum of coherence-squared measures calculated between the torque and the eye velocity ( $\gamma_{EVPT}^2$ ), the target velocity and the eye velocity ( $\gamma_{EVTV}^2$ ), and three coherence-squared measures that were caused by the inputs not being completely uncorrelated. The sum of these three measures was always small when compared to the coherence-squared between the inputs and the output, and so for simplicity in the coherence-squared plots the three measures were lumped together into one term ( $\gamma_{COR}^2$ ).  $\gamma_{EVTV}^2$  represents the visual target's effect on the eye velocity. The visual target affected the eye velocity through two paths: directly through the g1 subsystem, and indirectly through the visually mediated head motion system (g2) when voluntary head movements stimulated a VOR response. As discussed in Chapter 2, the VOR should not have been sensitive to voluntarily controlled head motions.  $\gamma_{EVPT}^2$  represented the perturbation torque's effect on the eye velocity. The torque affected the eye velocity through the head-neck system (g4) and then through the VOR (g3). The output of the g4 subsystem was the input to the g3 subsystem. The coherence-squared between the eye velocity and the torque was therefore bound by the coherence-squared between the head velocity and the input torque.

The coherence-squared sum indicates where the analysis of the VOR has been degraded by nonlinearities and noise, but this degradation cannot be attributed to the VOR system. The sum of all the coherence-squared terms shows how well the eye velocity was linearly related to the torque and target inputs. Its deviation from unity was caused by noise in the measured signals and from nonlinear responses in all the subsystems through which the signals passed.

**Figure 7.26:** Representative coherence-squared functions between the torque and target input, and the eye velocity output, for the same three subjects as shown in Figure 7.25.



The coherence-squared functions for three experiments are shown in Figures 7.26. The coherence-squared functions for the three subjects generally followed a similar pattern, although the results from the subject who did not use many voluntary head motions to track the target were consistently different from the other subjects. The coherence-squared functions could be divided into three frequency regions: below 0.5 to 1.0 Hz, up to about 3 to 4.5 Hz, and then finally frequencies above 3 to 4.5 Hz. Between about 1 Hz and 3 Hz, and in the best cases between 0.5 and 4.5 Hz, the coherence-squared was consistently high, near or greater than 0.8. This implied that across these frequencies the eye velocity had a strong linear relation to the inputs, and that the dynamics of the G3 subsystem was captured well by the linear analysis. It can be seen in the plots that show the coherence-squared measures between each input and the eye velocity output that this region ranges from frequencies where both the visual target and the perturbation torque contributed to the eye velocity, to frequencies where the target was no longer controlled, and the eye velocity was stimulated predominantly by the torque input that disturbed the head. At these frequencies the head-neck system had behaved linearly and there was a strong linear relation between the applied torque perturbations and the disturbance motion of the head, and so it was not surprising that the linear analysis was not degraded by nonlinearities. The sum of coherence-squared measures for the eye velocity output is less than that for the head velocity output. Therefore there were nonlinearities, and/or noise, beyond that caused by the head-neck system that degraded the linear relations between the inputs and the eye velocity.

Below 0.5 to 1.0 Hz the coherence-squared functions exhibited a similar behavior to that of the coherence-squared functions that related the head-velocity output to the torque and target inputs (Figure 7.22). The coherence-squared functions were high near 0.0 Hz, and again near and above 1.0 Hz, but between 0 and 1.0 Hz were lower. In the best cases, the coherence-squared remained above 0.6 and returned to above 0.8 by 0.5 Hz, but in many trials the coherence-squared dropped below 0.5 and did not increase until the frequency was near 1.0 Hz. The coherence-squared functions between the eye velocity and the individual inputs were also similar




to those for the head velocity. The coherence-squared between the visual target and the eye velocity showed that the drive from the target was strongest near zero and then decreased with frequency. The drive from the torque perturbation started near zero and then rose with frequency. It was not surprising that the shapes of the coherence-squared functions for the eye velocity systems were similar to those for the head velocity since the torque passed through the head-neck system to create the head disturbance that was the stimulus for the VOR. The nonlinearities of the head-neck system degraded the eye velocity coherence-squared function just as they did the coherence-squared for the head-neck system.

The similarities between the coherence-squared functions for eye and head velocity imply that the VOR subsystem behaved linearly. The analysis of the VOR at frequencies below 3 to 4 Hz was limited by the nonlinearities of the head-neck system. There were no indications that the VOR had nonlinearities during natural head-free tracking that could not be captured by a linear system analysis.

At higher frequencies, above 3 to 4 Hz, the coherence-squared for the eye velocity dropped, and remained low as the frequency increased. This was different from the coherence-squared functions for the head velocity, which had remained high until the power in the input torque had dropped, which was around 10 Hz. The drop in the coherence function had to be attributed to either nonlinear behaviors in the g3 subsystem or to noise in the data. Although nonlinearities in the g3 system could not be ruled out, its response probably did not have a strong nonlinear component. Its low frequency response did not have a strong nonlinear component, and the VOR had stabilized the subjects' gaze during perturbations at the higher frequencies as well as at the low frequencies. The drop in the coherence-squared value is best explained by a low signal to noise ratio for the eye movement data. The head velocity data were directly measured by an angular rate sensor and there was little noise in its measurement relative to the signal (Chapter 3). Eye velocity was calculated from the measured eye position data, and its noise was relatively high at frequencies above 4 Hz. Figure 7.27 shows typical PSDs for head and eye velocity.

Between 0 and 4.0 Hz, the plots are very similar with both the head and the eye velocity PSDs peaking near 1.5 Hz, and then decaying as frequency increases. Above 4.0 Hz, the eye velocity PSD stops decreasing and remains relatively constant as frequency increases, whereas the head velocity PSD continues to decline. This was shown in Chapter 3 to be due to the noise created when the eye velocity was calculated from the position data. This means that as the strength of the stimulus to the VOR decreased, the noise in the eye velocity measurement increased and so the ratio of the input signal to the noise decreased, resulting in a decline in the coherence-squared function.



**Figure 7.27:** Power spectral densities for the head velocity (brown) and eye velocity (blue) while the subject's head was perturbed.

The coherence-squared functions showed that the linear systems analysis of the VOR using measures of the eye velocity output and the torque and target inputs was very good between at least 1 and 3 Hz, and in the best of cases was very good between 0.5 and 4.5 Hz. Below 0.5 to 1.0 Hz the analysis was degraded by nonlinearities caused by active control of the head-neck system, and at frequencies above 3 to 4 Hz, noise in the eye velocity data affected the signal to noise ratio and hence the linear analysis. These results were consistent with observations made when the G3 FRFs were calculated. The low frequency analysis was sensitive to how the calculations of the spectral densities were made, which was interpreted as indicating that the

linear functions had not captured the nonlinear response of the head-neck system well. At the higher frequencies, there was greater variance in the gain and phase values.

*High pass models (incomplete description of parameters)*

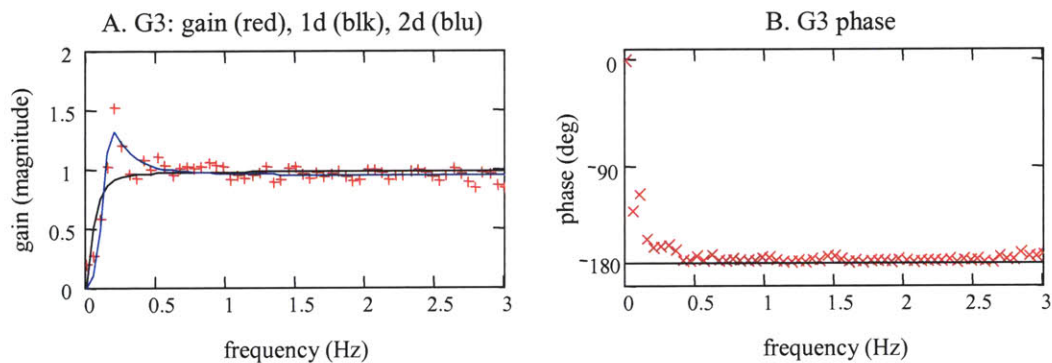
First- (Eq. 7.14) and second-order (Eq. 7.15) parametric models of the frequency response function for high pass systems were fitted to the VOR gain responses that had been calculated for each tracking trial (Figure 7.28). The models were fitted to the data from 0.0 Hz out to the frequency for which the coherence-squared values dropped, which was between 3 and 4 Hz. The models were developed using an electrical system analog with parameters for gain, resistance (R), capacitance (C), and inductance (L). In the first-order model the resistance and capacitance always appear as a product, and so were combined into a single RC parameter. The magnitudes of the models were fitted to the gain using a Levenberg-Marquardt minimization technique to minimize the sum of the squared difference between the model and the calculated gain values (Press *et al.*, 1997).

$$FRF_{1D} = \frac{gain \cdot RC \cdot j \cdot 2\pi f}{1 + RC \cdot j \cdot 2\pi f} \quad \text{Eq. 7.14}$$

$$FRF_{2D} = \frac{gain(j \cdot 2\pi f)^2}{(j \cdot 2\pi f)^2 + \frac{R}{L}(j \cdot 2\pi f) + \frac{1}{LC}} \quad \text{Eq. 7.15}$$

The VOR gain responses were fitted well by the high-pass models. As can be seen in Figure 7.28 if a first-order model did not fit the data well, then a second-order model often would. The values of the parameters for the first- and second-order models that were fitted to the gaze responses of one of the subjects are listed in Table 7.2. The values of the gain parameter were consistent from trial to trial, but the variance in the other parameters was large. For both the first- and second-order models the values of gain parameter were around 1, which is what would be

expected for a VOR that had stabilized the gaze well. The gain values for the first-order models had a mean of 0.96 and a standard deviation of 0.03. The gain values for second-order models had a mean of 0.94 and a standard deviation of 0.05. The low variance in the gain parameters indicates that the gain parameter was a good representative of the gain of the subject's system. The values for the other parameters in the models were much less consistent. The "RC" parameter of the first-order model spanned a very large range, from 1.39 to 8715. The standard deviation of the "RC" parameter was more than twice the value of the mean. In the second-order models there was large variance in the L, R, and C parameters. In all cases, the standard deviation of the values found for the parameters in the trials, was more than 50 % larger than the mean value. When the standard deviation, and thus the variance, of a parameter is this large, no single value could be said to be an accurate representation of the system in each of the trials.



**Figure 7.28:** Gain (plot A) and phase (plot B) response for the VOR. First-order (blue) and second-order (black) high-pass system model were fitted to the gain response.

Although the high pass models fit the VOR gain responses well, there could only be confidence in the gain parameter values as a measure of the subject's system response. The low frequency response of the high-pass models, that is the rise of the response from zero, the magnitude of the peak in the second-order system, and the corner frequency, are defined by the R, C, and L parameters. Since the linear system analysis did not capture the low frequency dynamics of the VOR well, the values of the parameters fitted to the low frequency gain response could not

be accepted as representative of the subject's VOR. The variance in these model parameters were a result of the limitations of the analysis.

Table 7.2: VOR high-pass model parameters						
	first-order model		second-order model			
trial	gain	RC	gain	L	R	C
1	0.96	1.39	0.96	1.53	1.35	3.57
2	0.93	5.48	0.92	2.23	0.78	2.03
3	0.94	1667	0.93	2.39	0.42	3.98
4	0.96	0.72	0.84	0.19	0.44	0.36
5	1.03	11.4	1.03	7.3	1.52	7.2
6	0.98	8715	0.95	6.08	0.27	1.24
7	0.98	1.6	0.97	0.96	0.79	1.78
8	0.93	2.5	0.94	2.92	1.48	10.5
<b>mean</b>	0.96	1.3x10 <sup>3</sup>	0.94	2.95	0.88	3.8
<b>std dev</b>	0.03	2.9x10 <sup>3</sup>	0.05	2.3	0.47	3.2

**Table 7.2:** Parameters of high-pass models fitted to VOR gain response.

The high frequency response of the models, which is the gain at which the response settled, is defined by the gain parameter. The coherence-squared functions demonstrated that the linear analysis of the VOR was very good in the 1 to 3 Hz frequency range, and since the models fit the data well, the gain parameter reflected the performance of the subject's VOR at these frequencies. The consistency in the gain values indicated that the analysis had captured a consistent system response from trial to trial.

The VOR is often thought to be suppressed during head-free gaze motions (Leigh and Zee, 1999). The peripheral vestibular organs sense the motion of the head. The canals, saccule, and utricle are stimulated similarly by voluntary head motions and perturbations applied to the head. However, the response to the signals must be different. A fully functional VOR that responded to all head motions would create eye motions that compensated for voluntary head


motions. If the head was driven to follow a target moving to the left, a fully functioning VOR would drive the eyes to the right to compensate for the leftward head motions. But this does not happen. When we track a target with head and eye motions, the head and eyes often move together. There must be some mechanism in the VOR that stops the eyes from compensating for voluntary head motions, but that allows the VOR to stabilize gaze during against disturbances. It has been proposed that a copy of the head motion, such as a corollary discharge, is used to inform the VOR of the intended movement, or that during head-free tracking the VOR is suppressed (Leigh and Zee, 1999). The results presented here indicate that the gain of the VOR, or at least the component that is stimulated by the perturbations applied to the head, was not affected during head-free tracking. When subjects moved their heads to follow the target the gain was near unity. Subjects did not necessarily move their heads quickly, and the head velocity that was attributed to the target was smaller than the head motion that was attributed to the torque. However, between 1 and 2 Hz, which was above the frequency range where the nonlinearities of the head-neck system had affected the analysis, but before the target was no longer controlled, the drive to the eyes was mixed between the torque and the target. At these frequencies, the gain of the VOR was similar to the gain measured at frequencies for which the visual target was not controlled. The VOR gain was consistent at frequencies for which the target was controlled and the head was free to move to pursue the target, and at frequencies for which the target was essentially stationary and the head moved by torque perturbations alone.

### 7.5.3 VOR impulse response functions

#### *VOR IRF estimate (g3)*

The VOR impulse response function estimates ( $g_3$ ) were calculated by de-convolving the cross-covariance function (CCVF) between the head velocity and the torque perturbation from the CCVF between the eye velocity and the torque perturbation. To improve the IRFs as models of the subjects' VOR systems, the CCVFs were calculated from data sequences that were

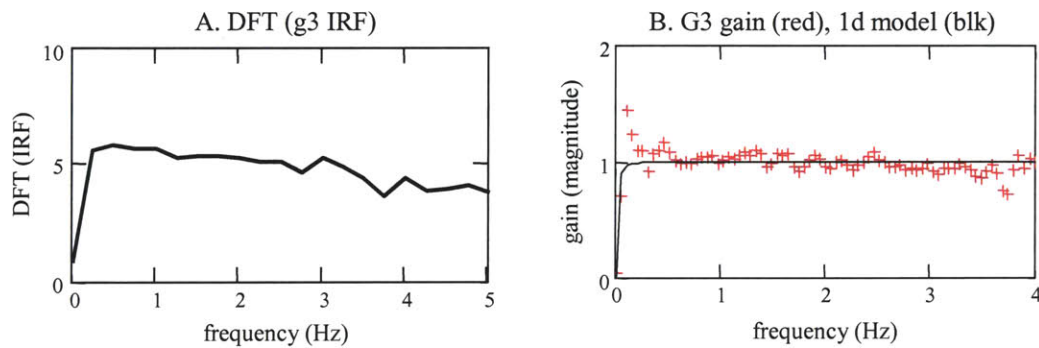
constructed by joining the experimental data records from several tracking trials. Two IRF estimates were calculated for each data sequence: a “single IRF” and an “average IRF”. The “single IRFs” were calculated by de-convolving CCVFs that had been calculated from the full length data sequences. The “average IRFs” were estimated by dividing the long data sequences into smaller, overlapping records, calculating an IRF for each record, and then averaging the IRFs of the records. The averaging technique tended to reduce the amount of noise that was modeled by the IRF estimates, and was similar to that used in the FRF analysis to reduce the variance of the cross-PSD estimators. Although the single IRF and the average IRF that were calculated for the same data set looked different, it was found that there was little difference in the dynamics of the VOR that they had captured.



**Figure 7.29:** Impulse response functions for the VOR. The IRF in plot A was calculated as a “single” IRF, and in Plot B as an “average” IRF.

The basic shapes of the single IRFs and of the average IRFs were much the same across subjects. There were consistent traits among the IRFs calculated from the full data records, and also among the IRFs calculated as averages of the IRFs from smaller records. A single IRF and an average IRF are shown for a representative data set in Figure 7.29. The single IRFs were dominated by spikes during the first few lag values. Most often the IRFs would start with a large pulse that went to a negative value, and then would quickly oscillate between a few positive and negative spikes that were close in magnitude to the initial pulse. After these initial spikes, the

amplitude of the IRF would drop in magnitude and would then oscillate more slowly, and smoothly. The oscillations slowly decayed towards zero. Plots of the average IRFs were simpler. They were dominated by a single, negative amplitude spike at a lag near zero. The initial pulse was followed by a smaller pulse with a positive amplitude, almost an overshoot, and then the amplitude decayed quickly to zero, with fewer and smaller oscillations than the single IRF.



**Figure 7.30:** Estimate of a VOR frequency response function (plot A) and the actual VOR FRF (plot B) calculated for the same tracking trial.

The VOR IRFs were not easily recognized as characteristic of any particular system type, but after taking the Fourier transform of the IRF, it was evident that most of the IRF estimates represented high-pass systems. This was consistent with the results of the frequency response functions. The IRF and the FRF are formally a Fourier transform pair, and so an estimate of the FRF is produced by taking the discrete Fourier transform (DFT) of the IRF. The DFT of the IRF in Figure 7.29 is shown in Figure 7.30, together with the FRF that was calculated for the same tracking trial. As was seen with the high frequency responses in the FRFs, the amplitudes of the DFT of the IRFs began near 0 at 0 Hz, rose to a higher magnitude, and then stayed near that level as frequency increased. The DFT of the IRF usually had a similar shape to the FRF, rising from a magnitude near zero to its high frequency gain value within the same frequency range, and like the FRFs, the low frequency values were sensitive to changes in the length of the IRF. The DFT of the IRFs did not have the same gain values as the frequency response functions, which was due



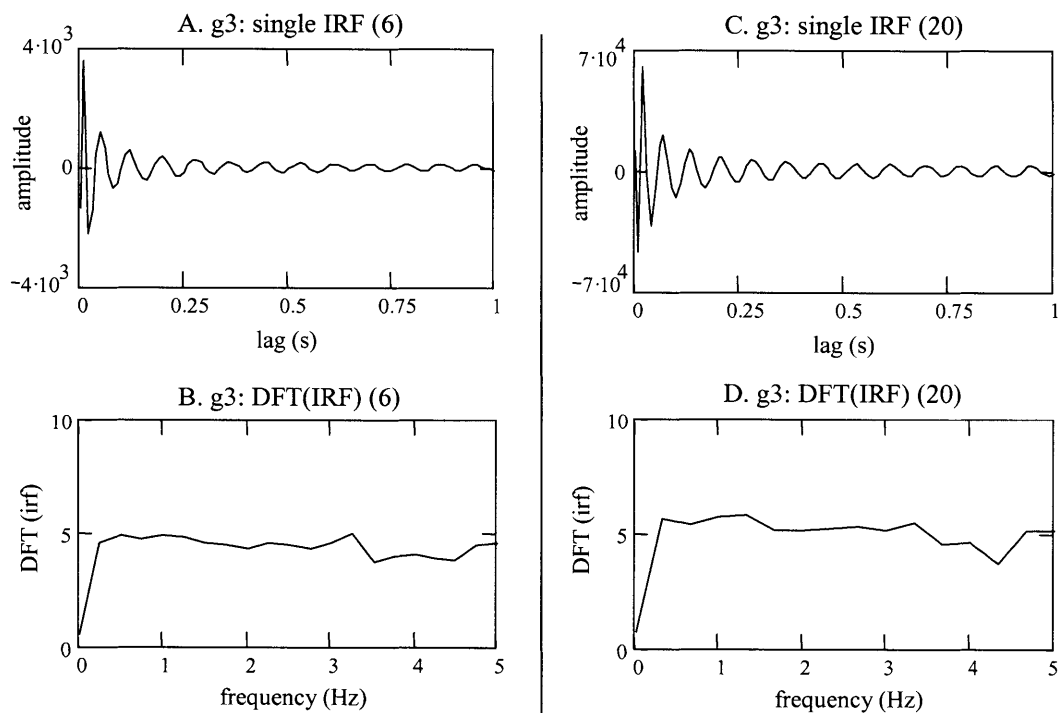
to the particular function forms that were used to calculate the DFT and the FRF. The DFT of the IRF had a cruder frequency interval than the FRF. The IRFs were only 2 to 4 seconds long which meant that the discrete time Fourier transforms had a large frequency increment. Despite these differences, the similarities between the FRFs and the DFT of the IRFs showed that the same VOR dynamics had been captured by the time and the frequency domain analyses.

#### *Similarities and variance between IRFs*

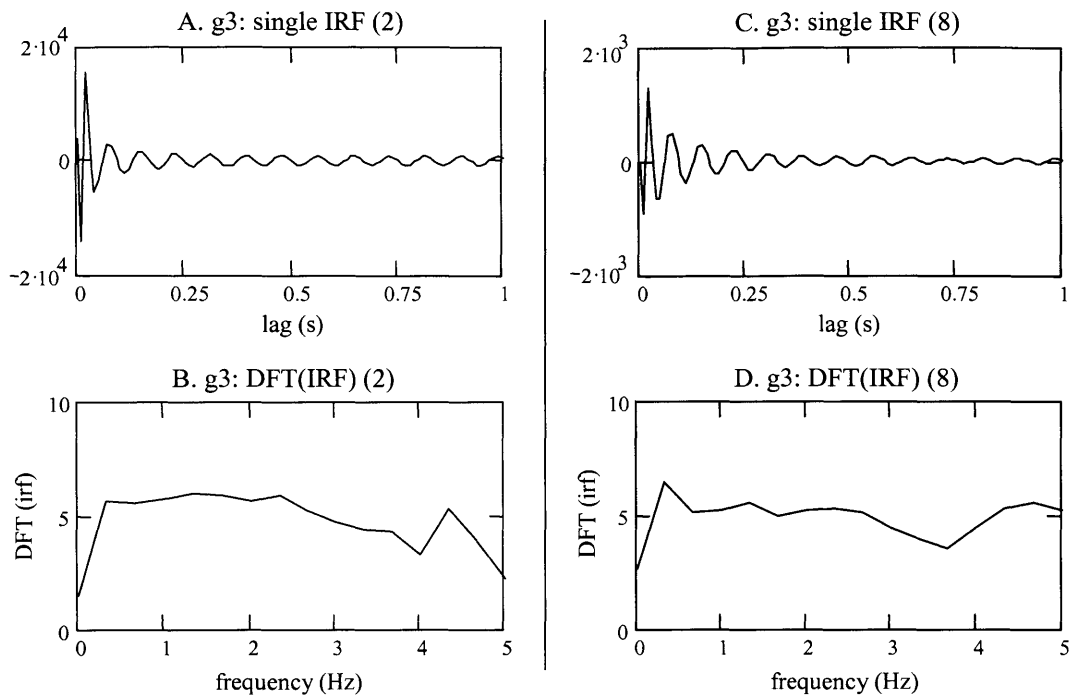
When the IRF calculations were first made, a separate VOR IRF was calculated for each trial. The data sequences for most trials were short, about 90 s in length, but because a rich range of head and eye movements had been stimulated during each experimental trial, the data were thought to be informative enough (Ljung, 1999) for the IRFs to capture the full dynamic response of the VOR.

The plots of the IRFs that were calculated from different trials, with the same subject were usually very similar. The IRFs from two subjects are shown in Figures 7.31 and 7.32. For both subjects, the only obvious difference between the IRFs from the different trials was in the amplitude scale. In fact, when the IRFs were overlaid, it could be seen that their time course and the decay of their amplitudes were nearly identical. The DFT of the IRF also showed that both analyses had captured the same frequency response for the VOR during the two experiments. The fact that nearly identical results were calculated from experimental trials that were conducted with different, rich spectrum, random inputs, was evidence that the IRFs had captured the underlying system performance of the VOR. Ideally, the IRF estimates that were calculated from the individual tracking trials would have represented only the dynamics of the subject's VOR and would have been identical. Realistically, however, noise variations in the voluntary responses of the subject, and other uncontrollable aspects of experimentation introduce some variance into the results for the IRF.

It was remarkable how similar some of the IRFs were from different trials, but the variability between the IRFs and in their ability to predict the output sequence, implied that the IRFs that were calculated from single tracking trials could be further improved as models of the subjects' VORs. When the IRFs were convolved with the input sequences to make predictions of the measured output signals, the IRFs were better at modeling the data from which they were analyzed, than the data sequences from other trials. This was not surprising, since the IRF is simply the solution to the set of equations established between the input and output data. It makes sense that an IRF would be a better solution for the equations from which it was solved, than for a different set of equations.




**Figure 7.31:** IRFs (plots A and C) and their Fourier transforms (plots B and D) for two data sets from the same subject.



**Figure 7.32:** IRFs (plots A and C) and their Fourier transforms (plots B and D) for two data sets from the same subject. The subject was different from the subject of Figure 7.31.

To improve the IRFs as estimators of a subject's VOR system, and make the IRFs less reflective of the results from a single tracking trial, data from several trials were joined into longer sequences from which the IRFs were estimated. The amount of data from which the IRFs were calculated increased from 90 s for a single trial, to 180 to 360 s for multiple trials. By joining the data from several trials, the sequence was more likely to have the necessary information from which to determine the full dynamic response of the VOR, and the IRF calculated from the data would be less likely to model the noise from any particular trial. Since the IRFs that had been calculated from the individual trials were so similar, it was not surprising that the IRFs that were calculated from longer data sequences were not significantly different from the IRFs of the single trials. The IRFs that were calculated by joining the data from the single trials shown in Figures 7.31 and 7.32, are shown in Figure 7.33. The IRFs from the multiple trials predicted the output data of each of the tracking trials well. The predictions were


not quite as good as when an IRF that was calculated from a single trial was used to predict the data from the same trial, but the predictions were better than when the IRF from the single trial was used to calculate data from a different trial.



**Figure 7.33:** IRFs calculated from data from multiple trials.

#### *Averaging IRF to reduce noise*

In some of the “single IRF” estimates, the oscillations that occurred after the initial spikes were much larger, relative to the magnitude of the pulses, than for other IRF estimates, as shown in Figure 7.34. The magnitude of the oscillations decayed slowly, if at all, and their magnitude relative to the amplitude of the spikes would change depending on the lag duration for which the IRF was calculated. In some cases the IRF was dominated by these oscillations and not by the dynamics of the VOR. When that occurred, the predictions of the VOR output made from the IRF were poor, and did not resemble the expected VOR response. Since it was known that the subject’s gaze had remained close to the target during the trial, the predicted VOR eye motions should have been similar to the predicted disturbance motions of the head. If the eye motions were not similar, they could not have stabilized vision during the disturbance and maintained the gaze close to the target.



**Figure 7.34:** VOR IRF with oscillations (plot A). The velocity of the head's disturbance motion (black) and the prediction of the VOR eye velocity (red) are shown in plot B.

It was initially thought that the oscillations were the result of an instability in the de-convolution operation that had been used to calculate the IRF from the cross-covariance functions. The de-convolution was conducted by inverting a matrix representation of the torque-head velocity CCVF. The matrix was of a Toeplitz form (constant along the negative diagonals), and inverting Toeplitz matrices can sometimes result in an instability that appears as an oscillation (Bottcher and Grudsky, 2000). The stability of the inversions was checked by inverting the Toeplitz matrix via singular value decomposition, which will usually not have the same instability. The same oscillations existed no matter how the matrix inverses were calculated, and so it appears that they are not the result of an instability in the Toeplitz matrix.

The most likely source of the oscillations in the IRF was noise in the data that had corrupted the IRF calculation. The IRFs possibly modeled the noise in the data more than the underlying VOR behavior. It was found that these IRFs were more sensitive to how the data were low pass filtered than IRFs that did not exhibit the large oscillations. The oscillations would sometimes not occur if the data were filtered at frequencies near 4 to 6 Hz, rather than at 10 Hz, which was the standard when the raw data from the tracking trails were processed (Section 7.3). In plot B of Figure 7.35, the IRF was calculated after the data were filtered with a 5 Hz low-pass filter. The data from which plot A was made were filtered at 10 Hz. The 5 Hz filter made a


considerable difference to the number of oscillations that occurred after the initial pulses. The IRF retained its pulses in the first few lag values, but the oscillatory response decayed quickly.

One advantage of calculating a system's IRF from the cross-correlation, or cross-covariance, is that it helps prevent the IRF from capturing and modeling noise in the data measurements. If measurement noise is uncorrelated with the input sequence from which the cross-correlation functions are formed (the torque perturbation in the case of the VOR analysis), the effect of the noise on the IRF is diminished. In these tracking experiments the noise in the data may not have been uncorrelated with the torque input. Much of the noise in the data was from the eye velocity sequence. The noise in the eye velocity sequence was not measurement noise, per se, but noise that resulted from calculating the eye velocity from eye position. The velocity calculation amplified any measurement noise that was in the eye position data. Since the eye motions were caused in part by the torque, it is possible that this noise was correlated with the torque. If that were the case, the cross-correlation between the noise and the torque would not have been zero, the effect of the noise would have remained in the set of equations that the IRF solved, and the noise would have been modeled by the IRF.

The frequency analysis was also affected by noise in the data. The noise was modeled by the FRF, but in the FRF relations are separated by frequency. The effects of the noise were evident at higher frequencies in the gain and phase plots where values had increased variance, and were quantified in the coherence-squared function when it dropped in value. The higher frequency noise content did not corrupt the lower frequency gain and phase relations that were attributed to the VOR. The IRFs are not separated by frequency and so the noise that was modeled by the IRF was spread across all lags.

The "average IRFs" were calculated to reduce the effects of noise on the IRFs. In many cases low-pass filtering eliminated the large magnitude oscillations in the IRF, but the filter was indiscriminant and removed data as well as noise. Averaging can sometimes "average out" the noise, but retain the dynamics from the system. The data sequences that had been formed by

joining data from several tracking trials were often 180 to 360 seconds long. The responses of the IRFs that had been calculated from the full length data sequences had settled within a few seconds. There was little advantage to calculating the IRFs to lags any greater than 2 to 4 s. It was therefore appropriate to divide the long data sequences into smaller records of 30 to 40 seconds, and still expect the average IRF to capture the dynamics of the system. One drawback to breaking the data into records was that it disturbed the continuity between the output and the input sequences. The output data at the beginning of a record had been affected by inputs that were at the end of a previous record. To reduce these effects, the records were overlapped, and the IRF was calculated only from data for which the full input-output convolution relation could be formed.



**Figure 7.35:** VOR IRFs. The IRFs in plots A and B were calculated from data that was low-pass filtered with a 10 Hz filter. The IRFs in plots C and D were calculated from data that was low-pass filtered with a 5 Hz filter.

The “average IRFs” looked more like impulses than models of systems with complex, or long duration, dynamics. The “average IRFs” did not have as much oscillatory behavior as the “single IRFs”. When there were oscillations, they would typically decay to almost zero by 0.50 s. The average IRFs were not very sensitive to the cutoff frequency at which the data were low pass filtered. As can be seen in plots C and D of Figure 7.35, there was often little difference between an average IRF estimator that was from data filtered at 10 Hz and from data that was filtered as low as 4 to 6 Hz. The average IRFs seemed to have reduced the amount of noise that was modeled by the IRF as effectively as the low-pass filters. It was reasonable that the average IRFs looked similar to a simple pulse function. If the average IRFs had been a pure pulse with a negative amplitude, then they would have indicated that the eye motions of the VOR were exactly equal and opposite to the head disturbance. This would be ideal for the VOR, as it represents perfect stabilization against the head disturbances. This is approximately what was observed in the subjects’ tracking data. Therefore, the VOR had effectively compensated for the head disturbances, which implies that the eye and head motions were equal and opposite. It should not be surprising that the average IRF for g3 would look like a negative impulse. The differences in the average IRFs from an ideal negative pulse did indicate that the VOR stabilization was not ideal, that is, the VOR eye motions could be delayed, that noise may have been captured by the IRFs, and that the VOR’s response to disturbances was not the same at all frequencies.

#### *Predictions of VOR eye motions*

Predictions of the VOR eye motions were made by convolving the estimates of the VOR IRF with predictions of the head’s disturbance motion. Head velocity had been measured, but the g3 IRF was calculated as the response of the eyes to the head’s disturbance motion, not its total motion. The best estimate of the head’s disturbance motion was found by convolving the IRF of the head-neck system (g4) with the measured torque perturbation. The predicted disturbance head velocity usually had a PSD that was similar to the PSD of the measured head velocity for

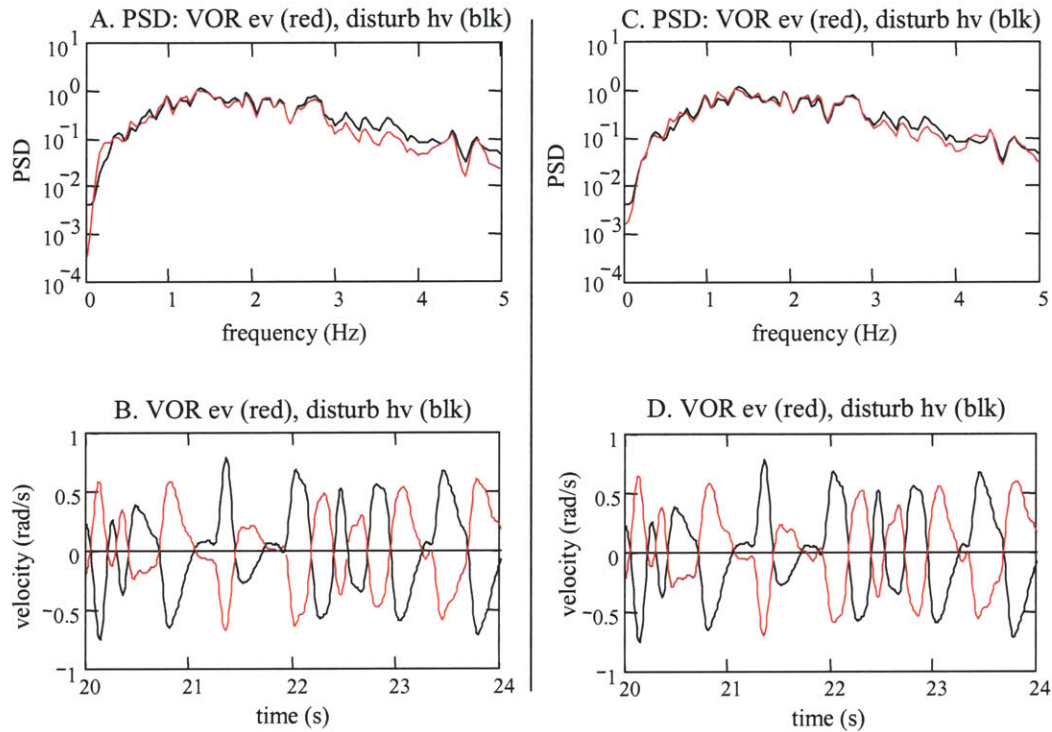


frequencies above 1.0 Hz, but deviated below 0.5 to 1.0 Hz (Figure 7.21). The lower frequency components of head velocity are controlled by the visually mediated system. During tracking, the head moved with velocities that had power across all frequencies, but the disturbance component rarely had much power below 0.5 Hz. Therefore the “linear g3” system was not stimulated well below 0.5 Hz. This meant that the IRFs could not have captured the behavior of the VOR at low frequencies (which was also seen in the FRF analysis). This was not a major concern, since the goal was to evaluate the normal, functional dynamics of the VOR. Grossman reported that below 0.6 Hz the head is disturbed little in the horizontal plane (Grossman *et al.*, 1988) and that the VOR does not have to stabilize the gaze often against disturbances below 0.5 Hz. The linear disturbance, therefore, matches the disturbance motions that are experienced by the horizontal VOR during normal function.

Figure 7.36 shows representative predictions of the VOR eye velocity, and the PSDs of the head disturbance and the VOR eye motions, that were made with single IRFs and average IRFs. In most cases there was usually very little difference between the predictions made using the single or average IRF. There were slight differences in the amplitudes of the PSDs at some frequencies, and subtle differences between the plots of the VOR eye velocities, but for the most part the predictions implied that the same system dynamics had been captured by both IRFs. This is interesting in light of the differences seen in the plots of the single and average IRFs for the same tracking trials (Figures 7.29, 7.35) and lent further support to the hypothesis that the oscillations in the single IRFs were from noise, and were not representative of dynamics of the VOR system.

The plots of the VOR eye motion against the head’s disturbance showed that the IRF analysis predicted that the VOR had stabilized the gaze against the head disturbance well. The amplitudes of the VOR eye velocities were nearly equal to the corresponding head disturbance velocities, and there was no noticeable lag between the head motion and the opposing eye motions. The %VAF between the VOR eye velocity and the disturbance head velocity were

typically above 95%, and the RMS difference was usually less than 4 deg/s (The VOR eye velocity was compared against -1 times the disturbance head velocity). These predictions agreed with the near unity gain and 180 deg phase difference that had been consistently found across most frequencies during the analysis of the VOR's frequency response function.



**Figure 7.36:** The velocities of the predicted head disturbance (black) and the VOR eye response (red). Plots A and B were calculated using a “single” VOR IRF, and plots C and D were calculated using an “average” IRF.

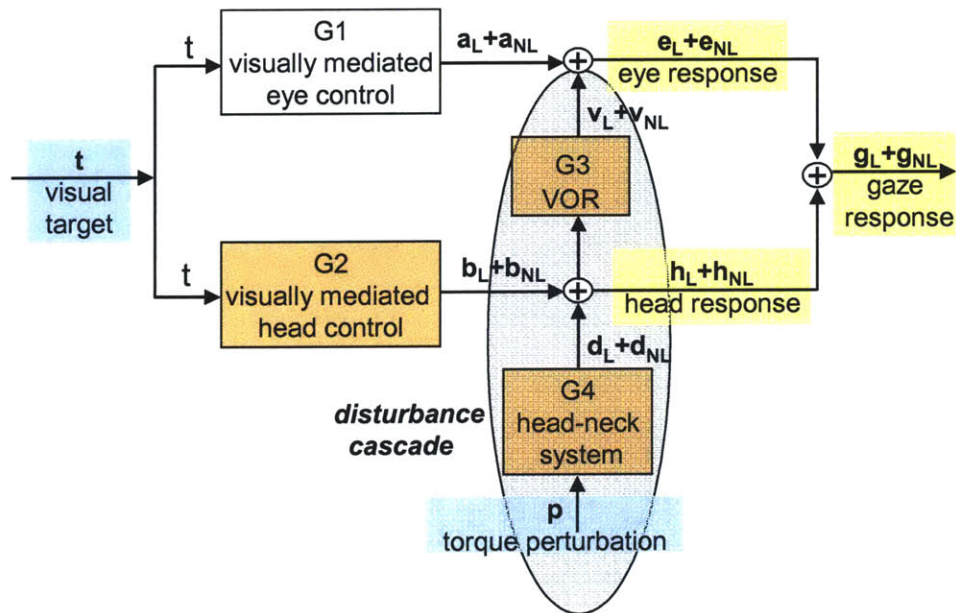
The response of the VOR is believed to lag the disturbance motion of the head by about 15 ms. It takes time from when the head is disturbed to when the stimulus from the semi-circular travels through the neural pathways to stimulate the neurons of the vestibular system’s “three neuron arc”, and causes the oculomotor muscles to move the eyes (Goldberg and Hudspeth, 2000). The IRFs, however, predicted that the VOR’s stabilizing motions occurred at the same time as the head’s motions. No consistent lag was found when the subjects’ VOR eye motions and the head disturbances were compared using cross-correlations, or even when one sequence was shifted to

minimize the difference between the two. The lag between the head movement and VOR reaction was too small for the IRFs to capture. Most of the IRFs were calculated from data sampled at 100 samples/s, and so the lag increment of the IRF and between predicted values of the eye response was 10 ms. The predicted VOR eye motions would have had to have been delayed relative to the head velocity by only 1 to 2 lag increments to show that a lag of 15 ms existed between the head and the eyes. This was too small a difference to be seen in this analysis. To record these delays, the noise in the measurements and predictions would have to be reduced, and the sampling frequency would have to be increased.

The main difference between the predicted VOR eye motions and the head disturbance motions was seen at frequencies below 1.0 Hz, which are the frequencies for which there can be little confidence in the analysis. Although their PSDs never matched exactly at high frequencies, the differences between the sequences were consistently greater at frequencies below 0.5 to 1.0 Hz. In the example shown in Figure 7.36, the VOR eye motion had a smaller power spectral density below 0.2 Hz than the head disturbance velocity. This gives the impression that the analysis showed that the VOR was a high-pass system, that is, the ratio of the eye PSD to the head PSD was near unity at high frequencies, but was lower below at low frequencies. This interpretation suffers from the same problems as the low frequency analysis of the frequency response function. The nonlinearities of the head-neck system prevented head disturbances (at least linearly related perturbations) from stimulating head motions at low frequencies. Even when the analysis did show that the low frequency response of the VOR decayed like that of a high pass system, that results cannot be accepted with confidence. Conclusions about the system level performance of the VOR at frequencies below 0.5 Hz can not be made from these findings.

## 7.6 Prediction of the gaze response

All the results from the IRFs indicated that the experimental protocol and the analytic techniques had successfully captured the system level dynamics of the VOR during head-free tracking. The IRFs captured consistent VOR dynamics for a subject from trial to trial, the dynamics modeled by the IRF agreed with those found from the frequency response functions, the predicted VOR outputs confirmed that the subject's gaze had been well stabilized by the vestibular system, and similar dynamics were found for different subjects. How well the predicted VOR eye motions compensated for head disturbances could be evaluated by comparing the predicted VOR output against the disturbance motions, unfortunately there were no measured (experimental) signals against which the predicted VOR eye motions could be directly compared in order to quantify how much of a subject's VOR dynamics had been accounted for by the linear models. This was similar to the situation where the frequency response functions had been verified. As with the  $g_3$  IRF, the FRF was calculated from the input and output signals of the gaze system. There were no direct measures of the VOR's input and output that could be used to create a coherence-squared function across the VOR system. Thus, the linearity of the VOR system, and the quality of the data from which the FRF was calculated, were judged from coherence-squared functions that were calculated from the controlled inputs to the gaze tracking system and the eye and head velocity outputs. In an analogous way, how well the linear IRFs had captured the system level dynamics of the subjects' VORs could be assessed by using the  $g_3$  IRF, along with the IRFs that were estimated for the head control systems ( $g_2$  and  $g_4$ ), to reconstruct the signals within the gaze control system. The linear predictions could be combined into a prediction of the gaze output and then compared with the measured output sequences, which were the result of the linear and nonlinear system dynamics.



**Figure 7.37:** Gaze tracking block diagram. Each signal is shown to have a linear (L) and a nonlinear (NL) component.

The intent of system identification is to determine the underlying properties of the system being tested. The goal is to develop a mathematical model of the system, such as an impulse response function, that behaves like the actual system when the model and system are stimulated by the same input. The capacity of an IRF to represent a system's dynamics comes not from the complexity of its calculation, but from the design and execution of the experiments. The IRF that is calculated from the cross-correlation relations is essentially the estimator, with the least squared error, that satisfies the linear convolution relations between the input and output sequences (Chapter 2). To get a meaningful IRF, the experiments must be designed so that the measured data, from which the IRF is calculated, is informative enough for the IRF to capture the full dynamics of the system (Ljung, 1999; Goodwin and Payne, 1977). The appropriate test of an IRF as a model of a system's dynamics is to evaluate the IRF's ability to predict outputs of the system from data that were not used to solve for the IRF. It is not as rigorous to validate an IRF as a model of a system's dynamics using the same data sequences from which the IRF was

calculated, because the IRF is already the least squares solution to the set of equations from the data. If an IRF is able to make good predictions of a system's response to rich, broad spectrum inputs, then the IRF probably represents the same dynamics as the system. If the IRF is only able to predict the data from which it was solved, then the IRF is not a good model of the system, but just the solution to one set of linear equations.

The ability of the IRFs to model the interactions of the subjects' gaze tracking systems was tested using data from trials that were not used to calculate the IRFs. As described in Section 7.5, the IRFs were calculated from data sequences that were formed by joining data from several tracking trials. This was done to make them better estimators of the subject's gaze tracking systems. These IRFs were then used to predict the output sequence from tracking trials that had not been used to calculate the IRF estimates. Since all the tracking trials were designed to stimulate a full range of head, eye, and VOR responses that a subject would experience during natural activities, any trials were appropriate for validating that the IRFs had indeed captured the underlying dynamics of the VOR and the gaze tracking systems. If the predictions from the linear IRFs matched the measured outputs well, then within the construct of the VOR experiment, two important conclusions could be drawn: the IRFs were good models of the systems they represented, which meant that the systems were primarily linear, and the IRFs represented the dynamics of the subject's vestibular systems, not merely a relationship between the output and input data from one experimental trial.

#### *Gaze response: signals and systems*

Consider the block diagram model for the gaze tracking system shown in Figure 7.37. The outputs from the physiological systems that are represented by the blocks are a mixture of linear (subscript L) and nonlinear (subscript NL) responses to the input. There is also noise in the measurements of the outputs (n). The output predicted from the IRF models of the systems modeled the linear output, including noise. The goal of the tracking system is to maintain the gaze

focused on the target, and to maintain the gaze velocity with the target's velocity. The top half of the block diagram represents the eye control systems. The bottom half represents the head control systems. The VOR connects the head and eye control systems by generating corrective eye motions in response to head disturbances. The cascade from the applied torque perturbation through the head-neck system, and through the VOR represents the path of disturbances through the system. The visually mediated systems, modeled by g1 and g2, drive the eyes (a) and head (b) to follow the target. Disturbance motions (d) move the head and the eyes away from the target but are compensated for by the VOR eye motions (v). When the systems are working properly, the gaze (g) is maintained close to the target (t). The VOR testing protocol was designed specifically to analyze the response of the VOR (g3), but it also allowed for IRF estimates to be made for the head control systems (g2 and g4), and the overall gaze tracking system. The one subsystem for which an IRF estimate could not be determined directly from the input and output signals measured during the tracking trials was the visually mediated eye control system (g1). An estimate of its IRF could be made, however once a prediction of the VOR eye motions was made, and the output of the g1 system (a) estimated as the measured eye motion (e) minus the VOR prediction.

### *Linear predictions of the gaze*

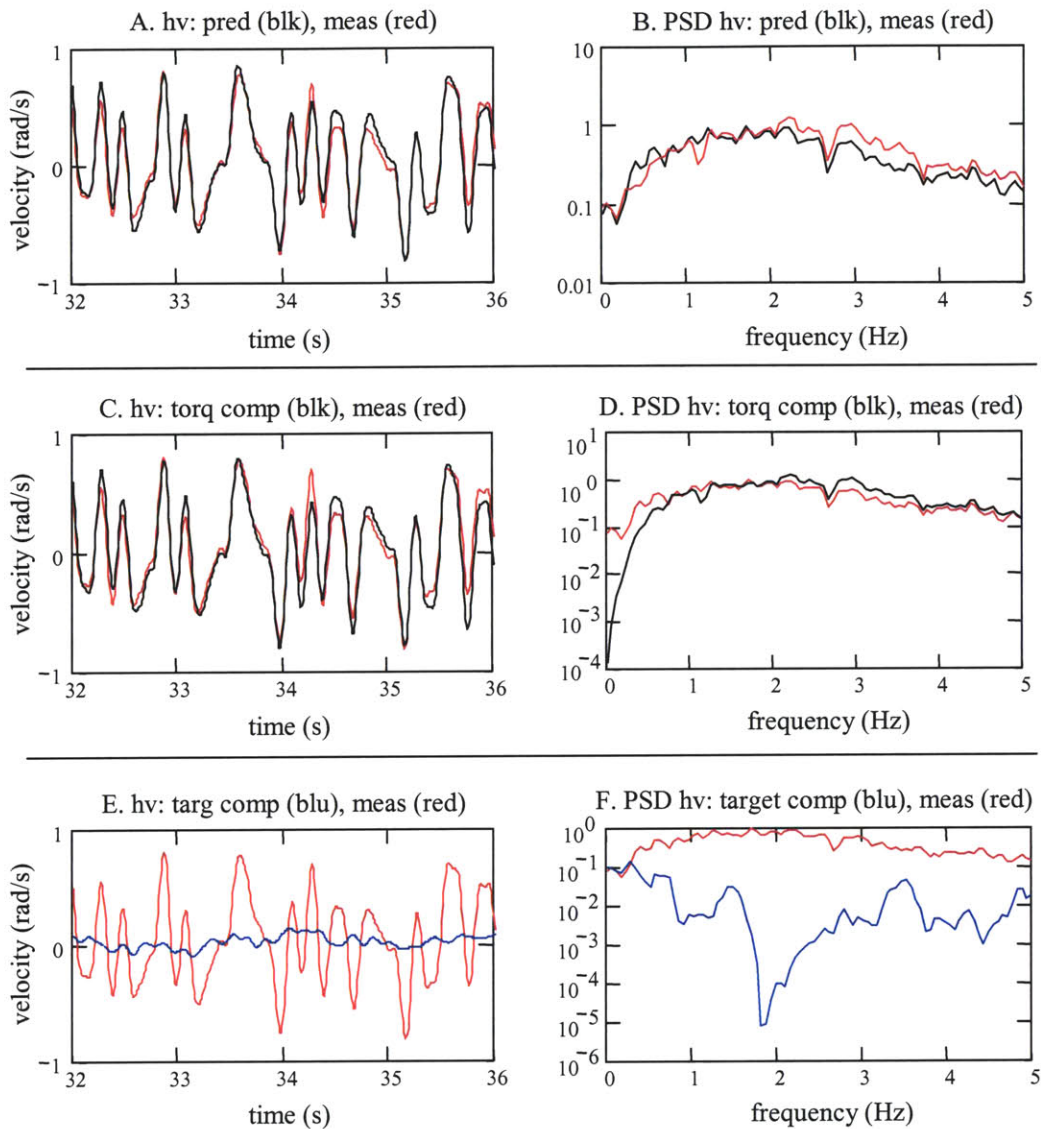
To illustrate how the linear predictions of the signals within the gaze system built upon each other to create a linear prediction of the gaze velocity, it is easiest to begin with the head velocity. The results from one subject will be used to demonstrate this analysis. The specific responses were unique to the subject, but in general the results are representative.

The measured head velocity sequence (h) could be divided into a component that was driven by the visual target (b), a disturbance component that was driven by the torque perturbation (d), and some noise (n). Each component could be further divided into linear and nonlinear portions. The breakdown of the head velocity into components is:

$$h = b_L + b_{NL} + d_L + d_{NL} + n = (t \otimes g_2 + p \otimes g_4) + (b_{NL} + d_{NL} + n). \quad \text{Eq. 7.16}$$

The linear prediction of a subject's head velocity sequence was the sum of a target driven component from the  $g_2$  IRF,  $b_L = t \otimes g_2$ , and a torque driven component from the head-neck IRF,  $d_L = p \otimes g_4$ . It was shown in Section 7.4 that the linear models of the head control system accounted for nearly all of the dynamics of the head and that it was typical for the linear model to account for nearly 90% of the variance in the measured output. Nonlinearities were significant only in the region between approximately 0.2 Hz and 0.5 to 1.0 Hz. In the example shown in Figure 7.38 the linear prediction of the head velocity had a "percent variance accounted for" of 86.5%, and the predicted head velocity clearly had the same motion as the measured velocity. It can be seen in Figure 7.38 that the visually driven component accounted for the very low frequency head velocities. The target had a frequency spectrum out to 1.8 Hz, yet the PSD of the head velocity that was linearly driven by the target was close to the PSD of the measured head velocity only below 0.5 Hz. The torque driven prediction accounted for much of the high frequency head motions. The power in the torque driven component of the head velocity dropped below 1 Hz. The vast majority of the head system's dynamic response was accounted for by the linear models. The 13.5 % of the variance in the measured head velocity that was not modeled by the predicted output was split between noise and nonlinear responses to the perturbations and target.






**Figure 7.38:** Linear predictions of the head velocity. The plots on the left show the measured head velocity (red) and the predictions (black or blue). The plots on the right show the power spectral densities. Plots A and B show the prediction of the head velocity. Plots C and D show the component of the head velocity that was driven by the perturbation torque. Plots E and F show the component driven by the visual target.

Although the subject's VOR did not discriminate between disturbances that were linearly ( $d_L$ ), or nonlinearly ( $d_{NL}$ ), related to the torque perturbations, the linear g3 IRF accounted for only the disturbance motions that were linearly related to the torque through the head-neck system. The

nonlinear VOR eye velocity component,  $v_{NL}$ , encompassed all VOR eye motions that were not accounted for by the linear analysis, including eye motions that were due to a nonlinear response of the vestibular system, and VOR eye motions that were stimulated by head disturbances that were related nonlinearly to the torque input. Both of these nonlinear components were probably small. The coherence-squared functions which were used to evaluate the VOR FRF had shown that the VOR did not have a strong nonlinear dynamic. As the linear models of the head control systems were able to predict 85 to 90% of the head's response, the nonlinear disturbance was also small. A prediction of the linear VOR eye motions is shown in Figure 7.39. The  $g_3$  IRF that was used to make the prediction was the IRF shown in plot A of Figure 7.33. The %VAF for the VOR eye motions was 98% in this example (calculated against  $-1 \times$  head velocity). The prediction for this trial had one of the higher %VAF measures, but it was common in other trials and for other subjects, that the prediction of the VOR eye motions would account for over 95% of the head disturbance. As discussed in Section 7.5, these predictions of the VOR driven eye velocity compensated very well for the linear head disturbances, and it was evident that the  $g_3$  IRFs were good models of the linear dynamics of the subjects' VOR systems.



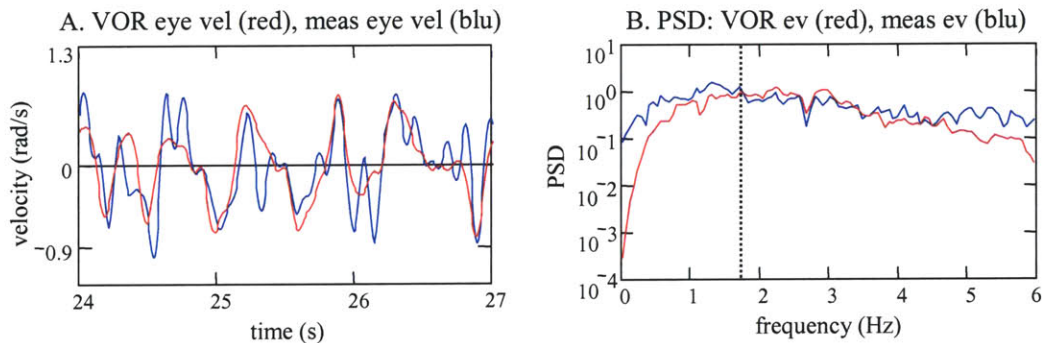
**Figure 7.39:** Prediction of the VOR eye velocity (red) and the velocity of the head's disturbance velocity (black). The prediction of the VOR accounted for 98% of the variance in the head's disturbance velocity.

The predicted VOR eye velocity is plotted against the measured eye velocity in Figure 7.40. The measured eye velocity was the sum of a component driven to follow the target (a), a

component from the VOR ( $v$ ), and noise ( $n$ ). The breakdown of the measured eye velocity into linear and nonlinear components is:

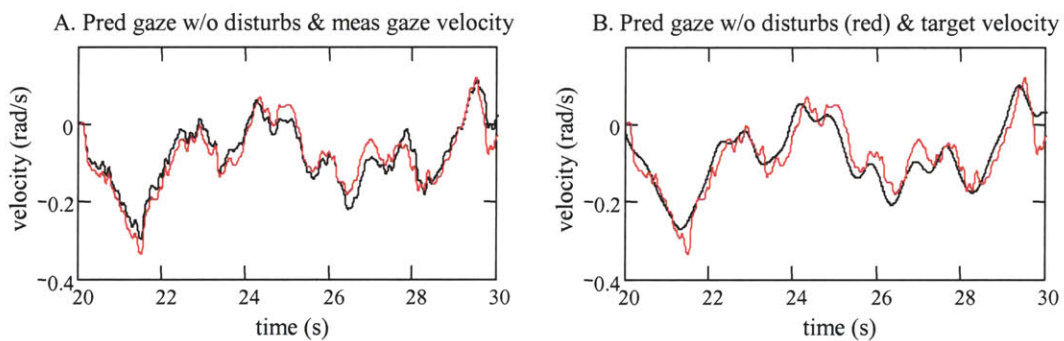
$$e = a + v + n = (a_L + p \otimes g_4 \otimes g_3) + (a_{NL} + v_{NL} + n). \quad \text{Eq. 7.17}$$

The predicted VOR eye motions,  $p \otimes g_4 \otimes g_3$ , comprised a large portion of the measured eye response, but they were not all of it. The two time-domain traces move with each other, but the measured eye velocity has a high frequency component that is not seen in the VOR eye motion. From the plots of the power spectral densities, it can be seen that the predicted VOR component matched the measured eye velocity closely from around 1.0 Hz to about 5.0 Hz. Above 5.0 Hz the power in the two sequences differed mainly because of the noise in the eye velocity resulting from calculating eye velocity from eye position data (Section 7.3). The difference between the two plots at low frequencies begins near 1.8 Hz, which is near the upper frequency of the target's spectrum. Below 1.8 Hz, the power in the VOR eye velocity dropped faster than that for the measured eye velocity. The power in the VOR eye component dropped just as the power in the disturbance head velocity dropped. The power in the total eye velocity remained higher, because at low frequencies it was driven by the visual target, as was head velocity.



**Figure 7.40:** The predicted VOR component of the eye velocity (red) and the measured eye velocity (blue). In this example, the VOR component of the eye velocity accounted for 40% of the variance in the measured eye velocity.

When the predicted VOR component was subtracted from the measured eye velocity, what remained was the portion of the eye velocity that was driven by the target, plus the nonlinear VOR component and noise, both of which were small relative to the eye motions driven by the target. When this portion of the eye velocity was added to the (predicted) component of the head velocity that was driven by the visual target ( $b_L = t \otimes g_2$ ), the outcome was very similar to the measured gaze response (Figure 7.41). This estimate of the gaze response reflected the responses of the visually mediated systems before the torque disturbance disrupted head and eye motions. The predicted gaze was, however, affected by noise and nonlinear VOR responses. The eye velocity component used in the estimate was not purely driven by the visual target, as it included measured noise and any nonlinear VOR eye motions.



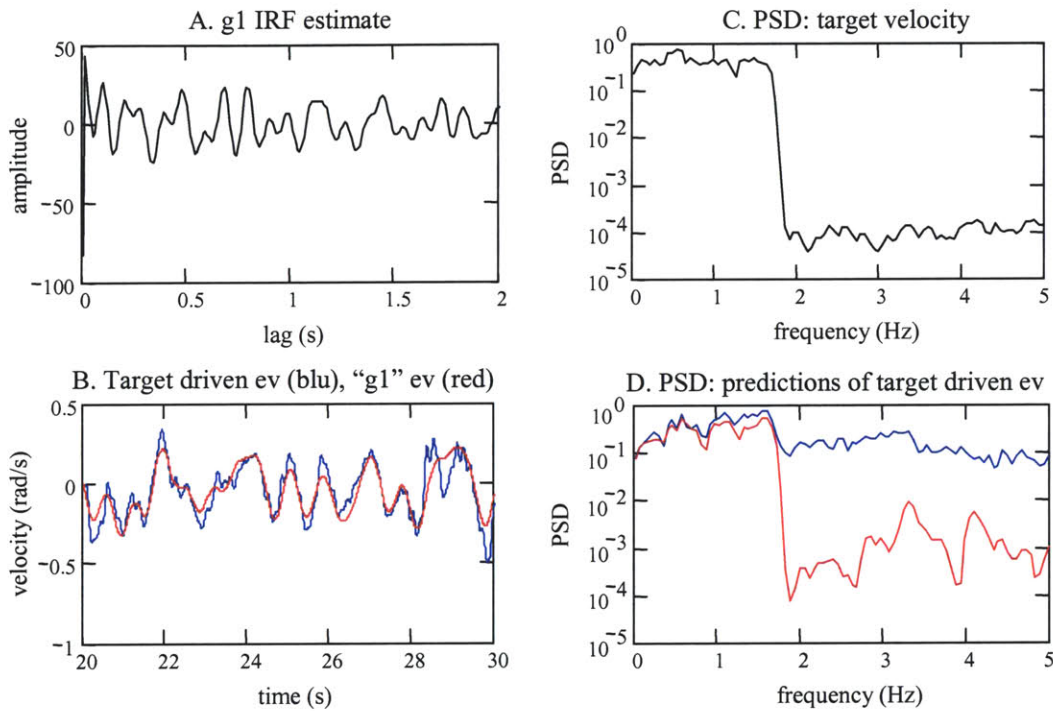
**Figure 7.41:** The prediction of the “gaze velocity without disturbances” (red) shown with the measured gaze velocity (plot A) and with the target velocity (plot B).

This eye velocity component can be used to create an estimate for the  $g_1$  IRF, so that a completely linear reconstruction of the gaze response can be made. After the linear VOR prediction is removed, the remaining portion of the eye velocity can be expressed as:

$$e_{REM} = t * g_1 + a_{NL} + v_{NL} + n \quad \text{Eq. 7.18}$$

The  $g_1$  IRF was estimated by de-convolving the auto-covariance of the target velocity (input) from the cross-covariance between the target and the component of the eye velocity,

$ev_{REM}$ . The effects of the nonlinear VOR response and the noise on the g1 IRF were minimized by using the cross-correlation functions. Neither the VOR response nor the noise was driven by the input target, so their cross-covariance with the target sequence was near zero. An example of a g1 IRF is shown in plot A of Figure 7.42.

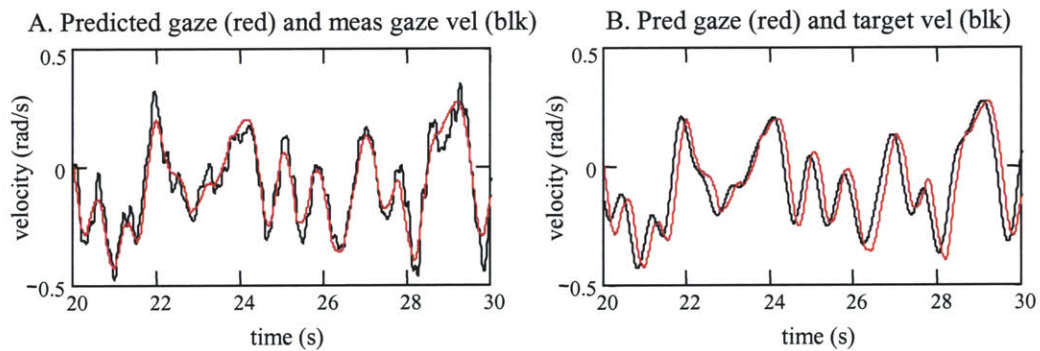


**Figure 7.42:** The visually mediated component of the eye velocity. A representative IRF for the visually mediated eye control system is shown in plot A. Two estimates of the visually driven eye velocity are shown in plot B. The estimate that was calculated by convolving the visual target sequence with the g1 IRF is shown in red, and the estimate that was calculated by subtracting the VOR eye velocity from the measured eye velocity is shown in blue. The PSDs of these estimates are shown in plot D.

An estimate of the linear output of the visually mediated eye control system was then made by convolving the g1 IRF with the target sequence ( $a_L = t \otimes g1$ ). The linear estimates of the eye velocity component resembled smoother versions of the eye velocity minus the predicted VOR (Figure 7.42). The estimates did not have the higher frequency motions that were in the

measured eye velocity. The output from the g1 IRF was linearly related to the input target, which had power only in the frequencies over which it was controlled, that is between 0 and an upper bound of 2.0 Hz (Figure 7.42). What was missing from the estimate was the noise, and if they existed, nonlinear outputs of the visually mediated systems and any nonlinear VOR responses.

Once this linear component of the eye velocity was found, all the components that were needed to make a linear estimate of the gaze response were available. In Figure 7.43, the linear prediction of the gaze response and the measured gaze response are plotted. The prediction was remarkably similar to the measured gaze response. It matched the gaze response closely, and in the example shown accounted for 87% of the gaze velocity's variance. The linear prediction does not account for noise or the nonlinear components of the eye velocity that were driven by the target, the head velocity driven by the target, the head disturbance driven by the torque, or the VOR, which indicates that these are small when compared to the linear responses of the systems.



**Figure 7.43:** Linear prediction of the gaze velocity (red). A prediction of the gaze velocity (red) and the measured gaze velocity (black) are shown in plot A. The prediction (red) is shown in plot B with the target velocity (black).

## Chapter 8

### Conclusions and future work

- 8.1 Conclusions of the VOR experiments
- 8.2 Other contributions
- 8.3 Future work
- 8.4 Thesis conclusion

#### 8.1 Conclusions of the VOR experiments

The results from the VOR experiments indicated that the dynamic performance of the VOR can be successfully evaluated while the VOR is used to stabilize vision during head-free tracking of a visual target. The testing protocol that was developed created the conditions from which the performance of the horizontal, rotational VOR could be evaluated across the full bandwidth for which the VOR naturally functions to stabilize gaze, that is from 0.5 to above 3.0 Hz. Impulse and frequency response functions that were consistent estimators of the VOR's performance could be calculated from as little as 90 s of data.

The visual and vestibular inputs that were generated using the modified Hunter-Kearney procedure stimulated the gaze tracking and VOR responses that were required for the VOR evaluation to be conducted. The PRB torque perturbation disturbed the head's motion with frequencies up to 8.0 Hz and velocities that exceeded 70 deg/s. Despite these strong disturbances, subjects tracked the visual targets well, and freely used their heads to direct their gaze. The subjects' gaze responses lagged the visual input with lags that exceeded 100 ms, which implied that the target was unpredictable and that the tracking response was driven by visually mediated

systems. The targets were tracked with a small position error, but large gaze saccades were not needed to maintain sustained tracking. The experimental torque and target sequences were very similar to the designed input signals, but because of noise they were slightly more correlated than the designed signals. However, they were still many times less correlated than the PRB torque sequence was with itself, and therefore could be approximated as uncorrelated, which was a requirement for the analysis to be conducted.

Impulse and frequency response functions for the VOR, the head-neck system, and the visually mediated head control system could be consistently estimated from data acquired in trials as short as 90 s. The head-neck system was well modeled as linear from below 0.5 Hz through to 8.0 Hz. The nonlinear dynamics of the head's stability and voluntary control mechanisms prevented the linear analysis from modeling the dynamics of the head-neck system accurately below 0.5 Hz. Low power in the torque perturbation input prevented the analysis from adequately stimulating the head-neck system at the frequencies above 8.0 Hz.

The linear systems analysis captured the dynamics of the VOR between 0.5 and 4.0 Hz. The dynamics captured by the impulse response functions agreed with that found for the system in the frequency domain. Response functions that were calculated from data acquired in one trial could be used to predict accurately the subject's response in another trial. The consistency of the results across domains, and from trial to trial, indicated that the analysis had successfully captured the underlying behavior of the subject's VOR. The VOR analysis was limited at low frequencies because of the nonlinear response of the head and neck, and at high frequencies because of noise that existed in the calculation of the eye velocity.

The IRFs for the head control and VOR systems were used to model the entire gaze tracking system. These models were often able to predict the gaze output from a tracking trial and accounted for over 80% of the variance. Although the gaze tracking system, like most physiological systems, is a nonlinear system, within the constraints of these stochastic VOR experiments, its behavior was very well modeled with linear systems.



Three observations can be made from the results of these VOR experiments that cannot be made from the results of traditional VOR tests. First, and most obvious, the impulse and frequency response functions experimentally verify that the VOR has the performance that is expected of it when it stabilizes gaze, that is, it has a gain near one and a phase of 180 deg. This is often not captured in traditional rotation tests, where it is common for the VOR gain to be calculated as less than 0.8, or for the response reported for the VOR to be corrupted by visually driven eye motions (Leigh, 1996; Collewijn, 1989). This new VOR testing protocol enables the performance of the VOR to be measured accurately, because it tests the VOR under conditions in which it functions naturally to stabilize vision, and in a manner that allows the eye motions driven by the vestibular system to be distinguished from those driven by the visual system. The gain and phase responses that are calculated for the VOR are not corrupted by the visual system, since the eye motions that are stimulated by the visual target can be mathematically removed from the analysis.

Second, there was no indication that the VOR was significantly suppressed during head-free tracking. Subjects did not make large head motions while tracking the target, which may have influenced the results, but with the motions that were used the VOR gain did not appear to be suppressed or cancelled. There was no difference in the VOR gain response at frequencies where the target had power and drove eye and head motions, or at frequencies where the target had little power and the eyes motions were driven primarily by the head disturbance. There was no indication that the smooth pursuit system had suppressed the VOR gain at any frequencies for which there were smooth pursuit tracking motions.

Third, the latency between the head disturbance and the compensatory eye motion was very short and indicated that the stabilization of the vision against the head disturbances was indeed accomplished by the VOR, and not by the visual system. The latency between the head disturbance and the corrective eye motion was so small that it could not be distinguished consistently from zero in data that had been sampled at 100 Hz. The compensatory eye motions

appeared to have occurred almost immediately after the head disturbance. The analysis also showed that the subjects required over 100 ms to create visually mediated eye and head movements. This implies that the visual system is too slow to be considered the cause of the eye motions that compensated for the head disturbance. Only the vestibular system is fast enough to have created the rapid compensatory eye motions.

## **8.2 Other contributions**

In addition to the development of a new technique to evaluate the VOR, this research produced several other important findings. First, the requirements that must be met by torque and target inputs to control the visual system and to disturb the head were identified. This will make it possible to develop a portable version of the testing equipment, using one of the systems discussed in Chapter 3. This, in turn, will enable system-level testing of a patient's VOR to be conducted outside of the laboratory, such as in a hospital, or clinic, or even in a patient's home. Second, it was shown that the Hunter-Kearney technique for creating random sequences with arbitrarily defined ACFs and PDFs could be used to create pseudo-random visual stimuli. These stimuli are considerably different from the pseudo-random stimuli that are traditionally used when testing the gaze tracking system. It was observed that these types of target are tracked differently from pseudo-random visual targets that are made by combining predictable sinusoids.

The two most important developments that allowed this research to be conducted successfully were the development of the modified Hunter-Kearney algorithm and the identification of the head-neck system as a quasi-linear system. The modified Hunter-Kearney technique can be used to create sequences with arbitrary spectral and cross-correlation characteristics. It is an extremely powerful tool for creating input signals for multi-input physiological systems, where each input may have to be of a different signal type in order for it to be tolerated well by the system under study.

The study of the head-neck dynamics demonstrated that although the head-neck system is stabilized by frequency dependent mechanisms, and is in general a nonlinear system, its response to stochastic perturbations can be well represented by a quasi-linear, second-order system. When the mean torque applied to the head was held constant, and the position disturbances that were applied were small, the viscoelastic properties of the system were not dependent on the stimulus frequency between 0.5 and 10 Hz. This allowed for the linear systems analysis to be conducted across the head-neck system, and for the torque perturbation to be a controlled experimental input. Had this not been the case, nonlinear techniques would have been required, which would have increased the complexity of the experiment design, analysis, and system modeling.

### **8.3 Future work**

The development of the equipment, testing protocol and analyses was undertaken using healthy adults as test subjects. The system should be tested next with people who have vestibular dysfunction to determine if different protocols are more appropriate for use with people whose natural gaze tracking may not be normal. The testing equipment can also be used for the analysis of systems other than the VOR. For example, the equipment could be used to study proprioceptive performance in the head-neck system, and to conduct more thorough tests of oculomotor performance.

In this research, the linear analysis of the VOR was shown to be limited to a bandwidth between approximately 0.5 and 4.0 Hz. The analysis could potentially be improved at low frequencies by modeling the head-neck system using nonlinear models, such as LNL, Wiener, and Hammerstein cascade systems (Hunter and Korenberg, 1986; Korenberg and Hunter, 1986), and at high frequencies by implementing an eye tracking device that can measure eye velocities with less noise. It is reasonable to expect that a nonlinear modeling of the head-neck system would lower the frequencies at which the response of VOR could be associated with the controlled

torque input, and thus would lower the frequencies at which the VOR response could be analyzed. The protocols used in this research were developed specifically to analyze the systems with linear techniques, and so it might be necessary to redesign the input signals and the experimental protocol so that the systems are stimulated appropriately for a nonlinear analysis. An immediate improvement in the analysis of the VOR at frequencies above 4.0 Hz could occur if eye position was tracked with a scleral contact, which measures eye position more accurately and with less noise than EOG. However, this would only be appropriate if the experimentation was conducted in a medical laboratory where an anesthetic could be applied to the eye, and care could be taken to avoid corneal abrasions from the equipment. Video based eye tracking devices offer more accurate tracking of the eye position than the EOG, but would only be appropriate for use with this protocol if they could be designed with a lower mass, and with a better means for correcting for relative position changes between the sensor and the head as the head is perturbed.

## **8.4 Conclusion**

This thesis has shown that the dynamic response of the VOR can be evaluated under natural functional conditions where the VOR is used to stabilize vision during head-free tracking of a visual target. Because of this, a more accurate analysis of the normal and dysfunctional behavior of the VOR can be made over the frequencies for which humans rely upon the VOR to maintain clear vision. It is believed that the equipment and techniques that were developed during the course of this research will serve as the foundation for laboratory and clinical techniques that will be used for a more thorough diagnosis of vestibular.

## References

- Abramowitz, M. and Stegun, I. A., *Handbook of Mathematical Functions*, Dover Publication, Inc., New York, 1972.
- Acoustical Society of America, American National Standard: Procedures for Testing Basic Vestibular Function (ANSI S3.45-1999). 1999. New York, Acoustical Society of America.
- ADInstruments, Human Electro-Oculography. 2003. Retrieved March 2003 from [www.adinstruments.com/teaching/tmats/tnotes/INB11b\\_Human\\_EOG.pdf](http://www.adinstruments.com/teaching/tmats/tnotes/INB11b_Human_EOG.pdf)
- American Academy of Otolaryngology, Balance: Dizziness and Motion Sickness. 2003. Retrieved March 2003 from [www.entnet.org](http://www.entnet.org).
- Bahill, A. T., Iandolo, M. J., and Troost, B. T., "Smooth Pursuit Eye Movements in Response to Unpredictable Target Waveforms", *Vision Research*, Vol. 20, 1980, pp. 923-931.
- Baloh, R. W., Langhoffer, L., Honrubia, V., and Yee, R. D., "On-Line Analysis of Eye Movements Using a Digital Computer", *Aviation, Space, and Environmental Medicine*, Vol. 51, 1980.
- Barnes, G. R. and Asselman, P. T., "The Mechanism of Prediction on Human Smooth Pursuit Eye Movements", *Journal of Physiology*, Vol. 439, 1991, pp. 439-461.
- Barnes, G. R., Donnelly, S. F., and Eason, R. D., "Predictive Velocity Estimation in the Pursuit Reflex Response to Pseudo-Random and Step Displacement Stimuli in Man", *Journal of Physiology*, Vol. 389, 1987, pp. 111-136.
- Barnes, G. R. and Grealy, M. A., "Predictive Mechanisms of Head-Eye Coordination and VOR Suppression in Humans", *Journal of Vestibular Research*, Vol. 2, 1992.
- Barnes, G. R. and Lawson, J. F., "Head-Free Pursuit in the Human of a Visual Target Moving in a Pseudo-Random Manner", *Journal of Physiology*, Vol. 410, 1989, pp. 137-155.
- Barnes, G. R. and Ruddock, C. J. S., "Factors Affecting the Predictability of Pseudo-Random Motion Stimuli in the Pursuit Reflex of Man", *Journal of Physiology*, Vol. 408, 1989, pp. 137-165.
- Beier, G., Schuller, E., and Schuck, M., "Center of Gravity and Moments of Inertia of Human Heads", Naval Biodynamics Laboratory, 1980.
- Bottcher, A. and Grudsky, S. M., *Toeplitz Matrices, Asymptotic Linear Algebra, and Functional Analysis*, Birkhauser, Boston, 2000.
- Carl, J. R. and Gellman, R. S., "Human Smooth Pursuit: Stimulus-Dependent Responses", *Journal of Neurophysiology*, Vol. 57, No. 5, 1987, pp. 1447-1463.
- Clark, J. W., "The Origin of Biopotentials", *Medical Instrumentation: Application and Design*, edited by J. G. Webster John Wiley and Sons, Inc., New York, 1995, pp. 150-226.
- Collewijn, H., "The vestibulo-ocular reflex: an outdated concept?", *Progressive Brain Research*, Vol. 80, 1989, pp. 197-209.

- Davies, W. D. T., "System Identification for Self Adaptive Control", *System Identification for Self-Adaptive Control*, edited by W. D. T. Davies Wiley-Interscience, London, 1970.
- Demer, J. L., "How Does the Visual System Interact with the Vestibulo-Ocular Reflex", *Disorders of the Vestibular System*, edited by R. W. Baloh and G. M. Halmagyi Oxford University Press, New York, 1996, pp. 73-92.
- Demer, J. L., Goldberg, J., and Porter, F. I., "Effect of telescopic spectacles on head stability in normal and low vision", *Journal of Vestibular Research*, Vol. 1, 1991, pp. 109-122.
- DiFabio, R. P., "Sensitivity and Specificity of Platform Posturography for Identifying Patients with Vestibular Dysfunction", *Physical Therapy*, Vol. 75, No. 4, 1995, pp. 290-305.
- Duchowski, A., *Eye-Based Interaction in Graphical Systems: Theory and Practice*. 2000. Retrieved March 2003 from [www.vr.clemson.edu/eyetracking/sigcourse/notes/c05-TOC.pdf](http://www.vr.clemson.edu/eyetracking/sigcourse/notes/c05-TOC.pdf).
- Eykhoff, P., *System Identification: Parameter and State Estimation*, John Wiley and Sons, London, 1974.
- Foreman, J., *Better Strategies for Coping with Vertigo and Dizziness*. Section E-1. 12 Sep 2000. Boston, Boston Globe.
- Foster, K. and Parker, G. A., *Fluidics: Components and Circuits*, Wiley-Interscience, New York, 1970.
- Furman, J. M. and Cass, S. P., "Laboratory Evaluation: Electronystagmography and Rotational Testing", *Disorders of the Vestibular System*, edited by R. W. Baloh and G. M. Halmagyi, Oxford University Press, New York, 1996, pp. 191-210.
- Furman, J. M. and Durrant, J. D., "Head-Only Rotational Testing in the Elderly", *Journal of Vestibular Research*, Vol. 8, No. 5, 1998, pp. 355-361.
- Gauthier, G. M. and Vercher, J. L., "Visual vestibular interaction: vestibulo-ocular reflex suppression with head fixed target fixation", *Experimental Brain Research*, Vol. 81, 1990, pp. 150-160.
- Genta, G., *Kinetic Energy Storage: Theory and Practice of Advanced Flywheel Systems*, edited by G. Genta, Butterworths and Co., Ltd., London, 1985.
- Godfrey, K., "Introduction to Perturbation Signals for Time-domain System Identification," *Perturbation Signals for System Identification*, edited by K. Godfrey, Prentice Hall International Series in Acoustics, Speech, and Signal Processing, Prentice Hall, New York, 1993, pp. 1-59.
- Goebel, J. A., Isipradit, P., and Hanson, J. M., "Manual Rotational Testing of the Vestibulo-ocular Reflex," *The Laryngoscope*, Vol. 110, 2000, pp. 517-535.
- Goldberg, M. and Hudspeth, A. J., "The Vestibular System", *Principles of Neural Science*, edited by E.R. Kandel, J. H. Schwartz, and T. M. Jessel, McGraw Hill, New York, 2000, pp. 801-815.
- Goldberg, M. E., "The Control of Gaze", *Principles of Neural Science*", edited by E. R. Kandel, J. H. Schwartz, and T. M. Jessel, McGraw Hill, New York, 2000, pp. 783-800.
- Goodwin, G. C. and Payne, R. L., *Dynamic System Identification: Experiment Design and Data Analysis*, Academic Press, New York, 1977.

- Grossman, G. E., Leigh, R. J., Abel, L. A., Lanska, D. J., and Thurston, S. E., "Frequency and velocity of rotational head perturbations during locomotion", *Experimental Brain Research*, Vol. 70, 1988, pp. 470-476.
- Grossman, G. E., Leigh, R. J., Bruce, E. M., Huebner, W. P., and Lanska, D. J., "Performance of the human vestibulo-ocular reflex during locomotion", *Neurophysiology*, Vol. 62, 1989, pp. 264-272.
- Hain, T. C., Vestibular Testing. 2002. Retrieved February 2003 from [www.tchain.com/otoneurology/testing/engrot.html](http://www.tchain.com/otoneurology/testing/engrot.html).
- Halliday, D. and Resnick, R., *Fundamentals of Physics*, 2 ed., John Wiley and Sons, New York, 1986.
- Hanson, J. M. and Goebel, J. A., "Comparison of manual whole-body and passive and active head-on-body rotational testing with conventional rotary chair testing", *Journal of Vestibular Research*, Vol. 8, No. 3, 1998, pp. 273-282.
- Highstein, S. M., "How Does the Vestibular Part of the Inner Ear Work (Ch. 1)", *Disorders of the Vestibular System*, edited by R. W. Baloh and G. M. Halmagyi Oxford University Press, New York, 1996, pp. 3-11.
- Hirvonen, T. P., Aalto, H., Pykko, I., and Juola, M., "Comparison of two head autorotation tests", *Journal of Vestibular Research*, Vol. 9, 1999, pp. 119-125.
- Hunter, I. W. and Kearney, R. E., "Generation of Random Sequences with Jointly Specified Probability Density and Autocorrelation Functions", *Biological Cybernetics*, Vol. 47, 1983, pp. 141-146.
- Hunter, I. W. and Kearney, R. E., "Generation of Stochastic Sequences for System Identification", *Proc. 35th Ann. Conf. Eng. Med. Biol.*, Vol. 24, 1982, pp. 72.
- Hunter, I. W. and Korenberg, M. J., "The Identification of Nonlinear Biological Systems: Wiener and Hammerstein Cascade Models", *Biological Cybernetics*, Vol. 55, 1986, pp. 135-144.
- Jenkins, G. M. and Watts, D. G., *Spectral Analysis and its Applications*, Emerson-Adams Press, Inc, Boca Raton, FL, 1968.
- Johns Hopkins University Center for Hearing and Balance, Research and Training Center for Hearing and Balance. 2003. Retrieved February 2003 from [www.bme.jhu.edu/labs/chb/index.html](http://www.bme.jhu.edu/labs/chb/index.html).
- Juang, J. N., *Applied System Identification*, Prentice Hall PTR, Englewood Cliffs, NJ, 1994.
- Kearney, R. E. and Hunter, I. W., "System Identification of Human Joint Dynamics", *Critical Reviews in Biomedical Engineering*, Vol. 18, No. 1, 1990, pp. 55-87.
- Keshner, E. A., Hain, T. C., and Chen, K. J., "Predicting Control Mechanisms for Human Head Stabilization by Altering the Passive Mechanics", *Journal of Vestibular Research*, Vol. 9, 1999, pp. 423-434.
- Keshner, E. A. and Peterson, B. W., "Mechanisms Controlling Human Head Stabilization. I. Head-Neck Dynamics During Random Rotations in the Horizontal Plane", *Journal of Neurophysiology*, Vol. 23, No. 6, 1995, pp. 2293-2301.

- Korenberg, M. J. and Hunter, I. W., "The Identification of Nonlinear Biological Systems: LNL Cascade Models", *Biological Cybernetics*, Vol. 55, 1986, pp. 125-134.
- Leigh, R. J., "What is the Vestibulo-Ocular Reflex and Why Do We Need It?", *Disorders of the Vestibular System*, edited by R. W. Baloh and G. M. Halmagyi, Oxford University Press, New York, 1996, pp. 12-19.
- Leigh, R. J. and Zee, D. S., *The Neurology of Eye Movements*, 3 ed., Oxford University Press, New York, 1999.
- Lindner, A. and Ilg, U. J., "Initiation of smooth-pursuit eye movements to first-order and second-order stimuli", *Experimental Brain Research*, Vol. 133, 2000, pp. 450-457.
- Ljung, L., *System Identification: Theory for the User*, 2 ed., Prentice Hall PTR, Upper Saddle River, NJ, 1999.
- Marmarelis, V. Z., "Coherence and Apparent Transfer Function Measurements for Nonlinear Physiological Systems", *Annals of Biomedical Engineering*, Vol. 17, 1988, pp. 143-157.
- Mayo Clinic, Testing for Balance or Dizziness Problems. 2003. Retrieved March 2003 from [www.mayoclinic.edu/balance-rst/testing/html](http://www.mayoclinic.edu/balance-rst/testing/html).
- Michael, J. A. and Melvill Jones, G., "Dependence of Visual Tracking Capability upon Stimulus Predictability", *Vision Research*, Vol. 6, 1966, pp. 707-716.
- O'Leary, D. P. and Davis, L. L., "Spectral analysis of low frequency, active head, VOR responses", *Journal of Vestibular Research*, Vol. 8, 1998, pp.313-324.
- Oppenheim, A. V., Schafer, R. W., and Buck, J. R., *Discrete-Time Signal Processing*, 2 ed., Prentice Hall, Upper Saddle River, NJ, 1999.
- Oppenheim, A. V., Willsky, A. S., and Nawab, S. H., *Signals and Systems*, 2 ed., Prentice Hall, Upper Saddle River, NJ, 1996.
- Pola, J. and Wyatt, H. J., *Target Position and Velocity: The Stimuli for Smooth Pursuit Eye Movements*, Vol. 20, 1980, pp. 523-534.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P., "Numerical Recipes in C: The Art of Scientific Computing", *Numerical Recipes in C: The Art of Scientific Computing*, edited by W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, Cambridge University Press, Cambridge, UK, 1997.
- Rashbass, C., "The relationship between saccadic and smooth tracking eye movements", *Journal of Physiology*, Vol. 159, 1961, pp. 326-338.
- Sawyer, R. N., Thurston, S. E., Becker, K. R., Ackley, C. V., Seidman., S. H., and Leigh, R. J., "The cervico-ocular reflex of normal human subjects in response to transient and sinusoidal trunk rotations", *Journal of Vestibular Research*, Vol. 2, 1994, pp. 245-249.
- Skalar Medical BV, Skalar Medical. 2003. Retrieved April 2003 from [www.skalar.nl](http://www.skalar.nl).
- SR Research Ltd., Eyelink. 2003. Retrieved March 2003 from [www.eyelinkinfo.com](http://www.eyelinkinfo.com).



- Strang, G., *Introduction to Linear Algebra*, 2 ed., Wellesley-Cambridge Press, Wellesley, MA, 1998.
- Thomas, G. B. and Finney, R. L., *Calculus and Analytic Geometry*, 6 ed., Addison-Wesley Publishing Company, Reading, MA, 1984.
- Van Gelder, P. and Lebedev, S., "Anticipatory saccades in smooth pursuit: Task effects and pursuit vector after saccades", *Vision Research*, Vol. 35, 1995, pp. 667-678.
- Vasavada, A. N., Peterson, B. W., and Delp, S. L., "Three-dimensional spatial tuning of neck muscle activation in humans", *Experimental Brain Research*, Vol. 147, 2002, pp. 437-448.
- Vestibular Disorders Association, Vestibular Disorders Association. 2003. Retrieved March 2003 from [www.vestibular.org](http://www.vestibular.org).
- Wall, C., Black, F. O., and O'Leary, D. P., "Clinical use of pseudorandom binary sequence white noise in assessment of the human vestibulo-ocular system", *Annals of Otolaryngology, Rhinology, and Laryngology*, Vol. 87, No. 6, 1978, pp. 845-852.
- Wall, C. and Furman, J. M., "Eyes Open Versus Eyes Closed: Effect on Human Rotational Responses", *Annals of Otolaryngology, Rhinology, and Laryngology*, Vol. 98, No. 8, 1989, pp. 625-629.
- Webster, J. G., "Amplifiers and Signal Processing", *Medical Instrumentation: Application and Design*, edited by J. G. Webster, John Wiley and Sons, Inc., New York, 1995, pp. 112-149.
- Wellenius, G. A. and Cullen, K. E., "A comparison of head-unrestrained and head-restrained pursuit: influence of eye position and target velocity on latency", *Experimental Brain Research*, Vol. 133, 2000, pp. 139-155.
- Wells, S. G. and Barnes, G. R., "Predictive smooth pursuit eye movements during identification of moving acuity targets", *Vision Research*, Vol. 39, 1999, pp. 2767-2775.
- Williams, R. A. and Fender, D. H., "Velocity Precision in Smooth Pursuit", *Vision Research*, Vol. 19, No. 3, 1979, pp. 343-348.
- Xu, Y., Hunter, I. W., Hollerbach, J. M., and Bennett, D. J., "An Airjet Actuator System for Identification of the Human Arm Joint Mechanical Properties", *IEEE Transactions on Biomedical Engineering*, Vol. 38, No. 11, 1991, pp. 1111-1122.
- Young, L. R., "Pursuit Eye Movement - What is Being Pursued?", *Control of Gaze by Brain Stem Neurons*, Vol. 1, Developments in Neuroscience, Elsevier/North Holland Biomedical Press, 1977, pp. 29-36.
- Young, L. R. and Stark, L., "Variable Feedback Experiments Testing a Sampled Data Model for Eye Tracking Movements," *IEEE Transactions on Human Factors in Electronics*, Vol. HFE-4, 1963, pp. 38-51.



# Appendix A

## Additional mathematics

- A.1 Least squares solution of the VOR IRF
- A.2 Coherence-squared function for a two-input, single-output system

### A.1 Least squares solution of the VOR IRF

The variables that are used in development of the correlation equations from which the VOR IRF is solved are defined in the following block diagram:

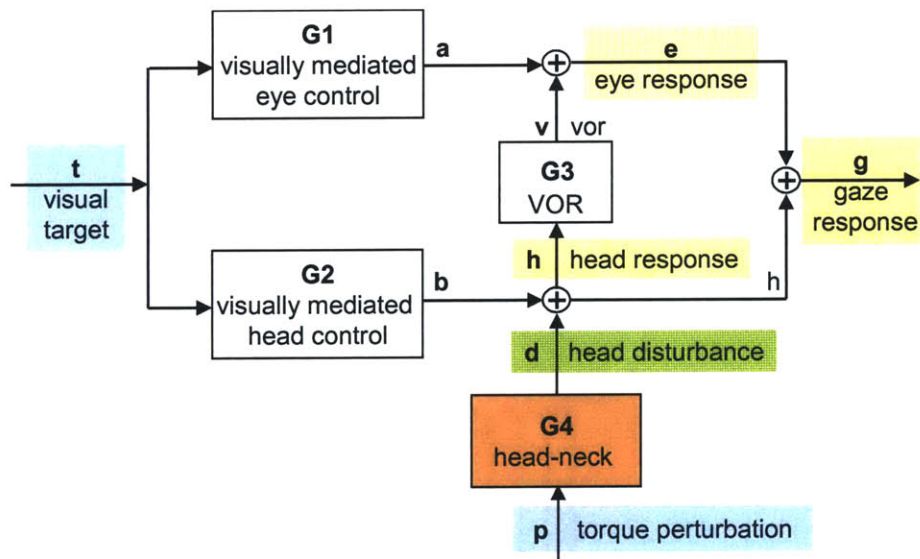


Figure A.1: Block diagram of the gaze tracking system.

In Chapter 2 the impulse response function (IRF) for the VOR ( $g_3$ ) was estimated by solving a set of convolution equations that were expressed with correlation functions. The linear convolution equations that related the torque and target inputs to the eye and head velocity

outputs were established first, and were converted so that the input and output sequences were replaced by auto- (ACF) and cross-correlation functions (CCF). The correlation based convolution equation from which the VOR IRF was solved was:

$$r_{PE} = r_{PT} \otimes g1 + r_{PH} \otimes g3 \quad \text{Eq. A.1}$$

which, because the torque and the target inputs were uncorrelated, reduced to:

$$r_{PE} = r_{PH} \otimes g3. \quad \text{Eq. 2.10}$$

The VOR IRF,  $g3$ , was found by deconvolving the CCF between the torque and the head velocity,  $r_{PH}$ , from the CCF between the torque and the eye velocity,  $r_{PE}$ , using a Toeplitz matrix inversion.

The estimation of the VOR IRF using the correlation based convolution equations is similar to finding the least squares estimate of the IRF directly from convolution equations that relate the input and output sequences. This, of course, is under the assumption that the input and output sequences are outputs of random processes and are ergodic.

The following develops the least squares solution for the VOR IRF ( $g3$ ) directly from the linear convolution equations, and shows that the solution to the equation is similar to the correlation based convolution equations.

The eye velocity ( $e$ ) can be considered the output of a two input system, where the target velocity ( $t$ ) and the head velocity ( $h$ ) are the inputs. Unlike the target velocity, the head velocity is not a controlled experimental input, but it can be measured. This convolution relation between the target velocity, head velocity, and eye velocity is:

$$e = t \otimes g1 + h \otimes g3, \quad \text{Eq. A.2}$$

which can be expanded into:

$$e_n = \Delta t \sum_{i=0}^M g1_i t_{n-i} + \Delta t \sum_{k=0}^M g3_k h_{n-k}, \quad \text{Eq. A.3}$$

where  $M+1$  is the length of the impulse response function.

Equations A.2 and A.3 were found more directly than the equations that were developed in Chapter 2. The development of these equations, and how they are no different than that developed in Chapter 2, is described in Chapter 7.

Equations A.2 and A.3 can be expressed as a series of linear equations which can be put in matrix form. For values of the eye response at times less than M, the causal convolution equation does not use the full length of the IRF, but for times n that are greater than or equal to M, a set of equations can be formed where all the terms of the response functions g1 and g3 are present. This linear set of equations is:

$$\begin{bmatrix} e_M \\ e_{M+1} \\ \dots \\ e_{N-1} \end{bmatrix} = \Delta T \cdot T \cdot G1 + \Delta T \cdot \begin{bmatrix} h_M & h_{M-1} & \dots & h_0 \\ h_{M+1} & h_M & \dots & h_1 \\ \dots & \dots & \dots & \dots \\ h_{N-1} & h_{N-2} & \dots & h_{N-1-M} \end{bmatrix} \cdot \begin{bmatrix} g3_0 \\ g3_1 \\ \dots \\ g3_M \end{bmatrix}, \text{ Eq. A.4}$$

where only the terms of the VOR IRF (g3) and the matrix for the head velocity (h) are shown in an expanded form. Eq. A.3 can be expressed as:

$$E = \Delta T \cdot T \cdot G1 + \Delta T \cdot H \cdot G3 \quad \text{Eq. A.5}$$

where E is an (N-M × 1) vector that represents the eye velocity at times M through N-1, ΔT is the sampling time times an (N-M × N-M) identity matrix, T and H are (N-1 × M+1) Toeplitz matrices with terms from the torque input and the head velocity, and G1 and G3 are (M+1 × 1) vectors with the unknown terms of the system response functions.

This is an over-determined set of equations. Since the number of data points is much greater than the length of the system response functions, there are far more equations (N-M) than there are unknowns (M+1 for each response function), and therefore it is unlikely that an exact solution can be found. Also, there is noise in the data, which makes it even more unlikely that there is an exact solution for either vector G1 or G3. The least squares solution to such a set of equations may be found by transforming the matrices into invertible, full-rank square matrices by

multiplying through by the matrix transpose (Strang, 1998). The solution is found by inverting the new full-rank matrix, which is often called the pseudo-inverse

A slightly unorthodox approach has to be used to solve for G3. Rather than multiplying Eq. A.5 through by the transpose of the T or the H matrix, the torque perturbation sequence can be used to remove the influence of the visual target (T) so that the terms of the G3 matrix can be solved. The solution for G1 can be found once the solution for G3 is known.

A matrix is made for the perturbation torque, P, in the same form as the target input matrix, T, and the head velocity matrix, H.

$$P = \begin{bmatrix} P_M & P_{M-1} & \cdots & P_0 \\ P_{M+1} & P_M & \cdots & P_1 \\ \cdots & \cdots & \cdots & \cdots \\ P_{N-1} & P_{N-2} & \cdots & P_{N-M} \end{bmatrix}. \quad \text{Eq. A.6}$$

The transpose of P,  $P^T$ , is multiplied through Eq. A.5 to form:

$$P^T E = \Delta T \cdot P^T T \cdot G1 + \Delta T \cdot P^T H \cdot G3, \quad \text{Eq. A.7}$$

where the  $\Delta T$  matrix, which represents a scalar multiplication, has had its dimensions and its position in the equation changed. To simplify the notation, the  $\Delta T$  matrix will be dropped, but it must be understood that it still exists in the equations. Eq. A.7 becomes:

$$P^T E = P^T T \cdot G1 + P^T H \cdot G3, \quad \text{Eq. A.8}$$

*The matrix equations are approximately the same as the correlation equations*

The set of matrix equations in Eq. A.8 can be shown to be approximately the same as the correlation based convolution equations, which were:

$$r_{PE} = r_{PT} \otimes g1 + r_{PH} \otimes g3 \quad \text{Eq. A.1}$$

It is obvious that the sequences that are used to form the components of each equation are the same, and it will be seen that summations of products represented by both the matrix multiplications and by the cross-correlation functions are approximations of each other.

*Output cross-vector  $P^T E$*

First consider  $P^T E$ , which will be called the output cross-vector. The product results in a vector of the order  $(M+1 \times N-M)(N-M \times 1) = (M \times 1)$ .

$$P^T E = \begin{bmatrix} p_M & p_{M+1} & \cdots & p_{N-1} \\ p_{M-1} & p_M & \cdots & p_{N-2} \\ \cdots & \cdots & \cdots & \cdots \\ p_0 & p_1 & \cdots & p_{N-M} \end{bmatrix} \begin{bmatrix} e_M \\ e_{M+1} \\ \cdots \\ e_{N-1} \end{bmatrix} = \begin{bmatrix} p_M e_M + p_{M+1} e_{M+1} + \cdots + p_{N-1} e_{N-1} \\ p_{M-1} e_M + p_M e_{M+1} + \cdots + p_{N-2} e_{N-1} \\ \cdots \\ p_0 e_M + p_1 e_{M+1} + \cdots + p_{N-M} e_{N-1} \end{bmatrix}$$

Eq. A.9

The terms in the output cross-vector are not quite the terms of the cross-correlation between the perturbation torque and the eye velocity, but they are very similar. The biased cross correlation can be defined as:

$$r_{PE}[j] = \frac{1}{N} \sum_{i=j}^{N-1} p_{i-j} e_i \quad \text{Eq. A.10}$$

where the lag  $j$  can range from 0 to  $N-1$ . If the first  $M$  lags of the cross-correlation are put into a vector, the result is:

$$\begin{bmatrix} r_{PE}[0] \\ r_{PE}[1] \\ \cdots \\ r_{PE}[M] \end{bmatrix} = \frac{1}{N} \begin{bmatrix} p_0 e_0 + p_1 e_1 + \cdots + p_M e_M + p_{M+1} e_{M+1} + \cdots + p_{N-1} e_{N-1} \\ p_0 e_1 + p_1 e_2 + \cdots + p_{M-1} e_M + p_M e_{M+1} + \cdots + p_{N-2} e_{N-1} \\ \cdots \\ p_0 e_M + p_1 e_{M+1} + \cdots + p_{N-M} e_{N-1} \end{bmatrix} \quad \text{Eq. A.11}$$

The matrix in Eq. A.11 can be made the same as the matrix in Eq. A.9 simply by scaling A.11 by  $N$  ( the length of the sequences), and by setting the first  $M-1$  terms of the eye velocity to zero (which means that only the eye velocities that are influenced by a full length system response function are considered). It is legitimate to substitute the cross-correlation terms in Eq A.11 for the matrix terms in Eq. A.9 if the terms in the matrix are summations of enough products so that they contain all the statistical information about the random process (ergodic) (Juang, 1994).

### *Input cross-matrices*

The matrix products  $P^T T$  and  $P^T H$  will be called the input cross-matrices, and will be shown to be similar to the cross-correlation between the torque and the target, and the torque and the head velocity. The general form of  $P^T T$  and  $P^T H$  is the expansion of  $P^T X$ , which is a  $(M+1 \times N-M)$   $(N-M \times M+1) = (M+1 \times M+1)$  matrix where:

$$P^T X = \begin{bmatrix} p_M & p_{M+1} & \cdots & p_{N-1} \\ p_{M-1} & p_M & \cdots & p_{N-2} \\ \cdots & \cdots & \cdots & \cdots \\ p_0 & p_1 & \cdots & p_{N-M} \end{bmatrix} \begin{bmatrix} x_M & x_{M-1} & \cdots & x_0 \\ x_{M+1} & x_M & \cdots & x_1 \\ \cdots & \cdots & \cdots & \cdots \\ x_{N-1} & x_{N-2} & \cdots & x_{N-1-M} \end{bmatrix}. \quad A.12$$

When the matrix multiplication is conducted it can be seen that each term in the resulting matrix, at row  $R$  and column  $C$ , is of the form:

$$P^T X_{R,C} = \sum_{i=M}^{N-1} p_{i-R} x_{i-C}, \quad A.13$$

where the rows and columns range from 0 to  $M$ . The terms of the  $(M+1 \times M+1)$  matrix are approximately the terms of the cross-correlation function  $r_{pX}$  at a lag  $(R-C)$ . As before it is legitimate to substitute the matrix terms in Eq. A.12 with the terms of the cross-correlation function if the summations contain enough products so that they contain all the statistical information about the random process (ergodic) (Juang, 1994).



### *Solution for G3*

The solution for the terms of the VOR's response function ( $G3$ ) in Eq. A.8 is found by inverting  $P^T H$  so that:

$$G3 = (P^T H)^{-1} (P^T E) - (P^T H)^{-1} (P^T T) \cdot G1 \quad \text{Eq. A.14}$$

The product of the matrices  $(P^T H)^{-1} (P^T E)$  approximates the deconvolution of the torque-head velocity cross-correlation,  $r_{PH}$ , from the torque-eye velocity cross-correlation,  $r_{ET}$ .

The product of the matrices  $(P^T H)^{-1} (P^T T)$  approximates the deconvolution of the torque-head velocity cross-correlation,  $r_{PH}$ , from the torque-target cross-correlation,  $r_{PT}$ . Since the torque and the target are designed to be uncorrelated,  $P^T T$  can be approximated as zero. This can also be seen simply by evaluating the terms of  $P^T T$ .

Eq A.11 can be simplified to:

$$G3 = (P^T H)^{-1} (P^T E), \quad \text{Eq. A.15}$$

which approximates the deconvolution of the torque-head velocity CCF from the torque-eye velocity CCF that must be done to solve for the  $g3$  IRF in Eq. 2.10.

Although the VOR response function can be solved using these matrix equations, there is a very good practical reason to use the cross-correlations rather than the matrix equations. As can be seen from the expanded form of  $P^T X$  in Eq. A.13, each term of the matrix must be calculated individually. It is much less time consuming to use the cross-correlation function, for example with a lag of 2, for all the terms where the row minus the column equals 2, than it is to calculate each term of the matrix individually. When there are a sufficient number of summations, the difference in the individual terms has an insignificant effect on the outcome of the solution. When the correlation functions are used, the terms are substituted so that the matrices are in a Toeplitz

form, which means that the matrix has constant values diagonals. It is thus extremely time consuming to calculate the terms required to find the matrix solution.

## A.2 Coherence-squared function for a two-input, single-output system

The coherence-squared function is typically developed for the single-input, single-output case. It is not commonly seen developed for multi-input cases. The development of the coherence-squared function for the two-input, single-output case follows. The development is done using the same procedure used by Juang (1994) for the single-input, single-output case.

The coherence-squared between the input  $y_i$  and the output  $u_j$  is defined as:

$$\gamma_{y_i u_j}^2 = \frac{\overline{S_{y_i u_j}} \overline{S_{y_i u_j}}^*}{\overline{S_{y_i y_i}} \overline{S_{u_j u_j}}}$$

where “S” denotes the spectral density which must be averaged from multiple records, and “ $\otimes$ ” denotes the complex conjugate.

The two inputs sequences are  $u_1(n)$  and  $u_2(n)$ , and the output is  $y(k)$ . Consider  $U_1(k)$ ,  $U_2(k)$ , and  $Y(k)$  as the discrete Fourier transform of the input and output sequences. Form a column vector from the DFT of the inputs and output, multiply against the complex conjugate transpose, and take the expected value. Since the product of a complex column vector and its complex conjugate transpose is a positive semi definite matrix, the expected value must be greater than 0.

$$E \left[ \begin{bmatrix} Y \\ U_1 \\ U_2 \end{bmatrix} \begin{bmatrix} Y^* & U_1^* & U_2^* \end{bmatrix} \right] = \begin{bmatrix} E[YY^*] & E[YU_1^*] & E[YU_2^*] \\ E[U_1Y^*] & E[U_1U_1^*] & E[U_1U_2^*] \\ E[U_2Y^*] & E[U_2U_1^*] & E[U_2U_2^*] \end{bmatrix} > 0 \quad \text{Eq. 5}$$

From the correlation theorem,  $E[YU^*]$  equals the spectral density  $S_{YU}$ , and so Eq. 5 can be expressed using cross spectral densities.

$$\begin{bmatrix} \bar{S}_{YY} & \bar{S}_{YU_1} & \bar{S}_{YU_2} \\ \bar{S}_{YU_1}^* & \bar{S}_{U_1U_1} & \bar{S}_{U_1U_2} \\ \bar{S}_{YU_2}^* & \bar{S}_{U_1U_2}^* & \bar{S}_{U_2U_2} \end{bmatrix} > 0 \quad \text{Eq. 6}$$

The determinant of a positive semi-definite matrix is greater or equal to zero, and so:

$$\bar{S}_{YY} \bar{S}_{U_1U_1} \bar{S}_{U_2U_2} - \bar{S}_{YU_2} \bar{S}_{YU_2}^* \bar{S}_{U_1U_1} - \bar{S}_{YU_1} \bar{S}_{YU_1}^* \bar{S}_{U_2U_2} + \bar{S}_{YU_1}^* \bar{S}_{YU_2} \bar{S}_{U_1U_2} + \bar{S}_{YU_1} \bar{S}_{YU_2}^* \bar{S}_{U_1U_2} - \bar{S}_{YY} \bar{S}_{U_1U_2} \bar{S}_{U_1U_2}^* \geq 0$$

The goal is to express this equation in the form of the coherence-squared function of Eq.

3. When the inputs  $u_1$  and  $u_2$  are uncorrelated, the cross spectral densities between  $u_1$  and  $u_2$  can be set to zero. When the inputs are not perfectly uncorrelated all terms must be considered. Eq. 4 can be rewritten as:

$$\bar{S}_{YY} \bar{S}_{U_1U_1} \bar{S}_{U_2U_2} \geq \bar{S}_{YU_2} \bar{S}_{YU_2}^* \bar{S}_{U_1U_1} + \bar{S}_{YU_1} \bar{S}_{YU_1}^* \bar{S}_{U_2U_2} - \bar{S}_{YU_1}^* \bar{S}_{YU_2} \bar{S}_{U_1U_2} - \bar{S}_{YU_1} \bar{S}_{YU_2}^* \bar{S}_{U_1U_2} + \bar{S}_{YY} \bar{S}_{U_1U_2} \bar{S}_{U_1U_2}^*$$

which leads to:

$$1 \geq \gamma_{YU_1}^2 + \gamma_{YU_2}^2 + \gamma_{U_1U_2}^2 - \frac{\bar{S}_{YU_1}^* \bar{S}_{YU_2} \bar{S}_{U_1U_2}}{\bar{S}_{YY} \bar{S}_{U_1U_1} \bar{S}_{U_2U_2}} - \frac{\bar{S}_{YU_1} \bar{S}_{YU_2}^* \bar{S}_{U_1U_2}}{\bar{S}_{YY} \bar{S}_{U_1U_1} \bar{S}_{U_2U_2}}$$

Eq. 8 expresses a relationship between coherence squared functions between each input and the output, between both inputs, and “triple” products between both inputs and the output. This quantity is bound by one. IF the inputs are strictly uncorrelated, then it is the sum of the coherence-square between each input and the output that is bound by one.



## **Appendix B**

### **Equipment specification sheets**

- B.1 Kollmorgen Goldline XT MT304A1 servomotor
- B.2 Kollmorgen Servostar CD CR03200 controller
- B.3 Assurance Technologies Inc. Mini force/torque transducer
- B.4 ATA Sensors ARS-04E magneto-hydrodynamic sensor
- B.5 ATA Sensors digital compensation filter
- B.6 Watson Industries ARS-C141 1AR8P angular rate sensor

## B.1 Kollmorgen Goldline XT MT304A1 servomotor

### Kollmorgen GOLDLINE XT

#### 30x PERFORMANCE DATA

Parameters	Symbols	Units	MT302A	MT302B	MT304A	MT304B
Horsepower	HP Rated	HP	0.83	1.02	1.00	1.85
Kilowatts	kW Rated	kW	0.62	0.76	0.75	1.38
Speed at Rated Power	N Rated	RPM	4500	6000	3000	6000
Max Operating Speed	N Max	RPM	4500	6000	3000	6000
Continuous Torque (Stall) at 40°C	Tc	lb-ft	1.00	1.00	1.95	1.90
		N-m	1.36	1.36	2.64	2.58
Continuous Torque (Stall) at 25°C	Tc	lb-ft	1.06	1.05	2.07	2.01
		N-m	1.44	1.42	2.80	2.73
Continuous Line Current	Ic	Amps RMS	2.57	4.67	2.72	5.53
Peak Torque	Tp	lb-ft	2.38	2.38	5.46	5.46
		N-m	3.23	3.22	7.40	7.41
Peak Line Current	Ip	Amps RMS	9.0	16.5	10.0	20.9
Max Theoretical Acceleration	Z	rad / sec <sup>2</sup>	69,364	69,346	91,544	91,676
Torque Sensitivity (Stall) ±10%	Kt	lb-ft / Amp RMS	0.398	0.217	0.727	0.349
		N-m / Amp RMS	0.540	0.294	0.986	0.473
Back EMF (Line-to-Line) ±10%	Kb	VRMS / kRPM	32.6	17.8	59.6	28.6
Max Line-to-Line Volts	VMax	Volts RMS	250	250	250	250
DC Res at 25°C (Line-to-Line) ±10%	Rm	Ohms	7.77	2.24	7.93	1.79
Inductance (Line-to-Line) ±30%	Lm	mh	31	10	51	12
Rotor Inertia	Jm	lb-ft-sec <sup>2</sup>	0.0000343	0.0000343	0.0000596	0.0000596
		kg-m <sup>2</sup>	0.0000465	0.0000465	0.0000808	0.0000808
Weight (without brake)	Wt	lb	5.6	5.6	7.3	7.3
		kg	2.5	2.5	3.3	3.3
Weight (with brake)	Wt	lb	8.0	8.0	9.7	9.7
		kg	3.6	3.6	4.3	4.3
Static Friction	Tf	lb-ft	0.022	0.022	0.026	0.026
		N-m	0.030	0.030	0.035	0.035
Thermal Time Constant	TCT	Minutes	22	22	25	25
Viscous Damping Z Source	Fi	lb-ft / kRPM	0.0057	0.0057	0.0083	0.0083
		N-m / kRPM	0.0078	0.0078	0.0112	0.0112
Motor Constant at 25°C	Km	lb-ft / √Watts	0.124	0.126	0.224	0.226
		N-m / √Watts	0.168	0.171	0.304	0.306
Thermal Resistance at Stall	Rth	°C / Watt	0.93	0.93	0.81	0.87
Number of Poles			8	8	8	8

Notes:

1. Please refer to the Kollmorgen SERVOSTAR™ CD Series SYSTEMS TECHNICAL PUBLICATION for Kollmorgen GOLDLINE™ XT Performance Curves.
2. For models with shaft seals, derate torque at all speeds by 0.059 lb-ft (0.08 N-m)
3. Continuous duty operation is based on using 8"x 12"x 3/8" aluminum plate.

2.1.3 Electrical Specifications

Product Model		Cx03	Cx06	Sx03	Sx06	Sx10	Sx20	Sx30	Sx55	Sx85	
<b>Main Input (BUS+ / BUS-)</b>	BUS (Vdc)	See BUS Module specs for Cx main input data		125 to 360	125 to 360	125 to 360	260 to 360	260 to 360	260 to 360	260 to 360	
	Rated Power @ dc (kW)			0.63-1.4	1.26-2.79	1.96-4.34	8.68	13.33	24.45	37.20	
<b>Rated Main Output (MA, MB, MC)</b>	Continuous Power (KVA) @ 115Vac 1φ Line Input (45° Ambient)	.55	1.1	0.55	1.1	1.6	n/a	n/a	n/a	n/a	
	Continuous Power (KVA) @ 230Vac 1φ Line Input (45°C Ambient)	1.1	2.2	1.1	2.2	3.2	3.2	n/a	n/a	n/a	
	Continuous Power (KVA) @ 230Vac 3φ Line Input(45° Ambient)			1.1	2.2	3.6	7.3	11	20	30	
	Continuous Current (Arms / ADC)	3	6	3	6	10	20	30	55	85	
	Peak Current (Arms / Adc) for 500mSec	9	18								
	Peak Current (Arms / Adc) for 2Sec			6	12	20	40	60	110	170	
	PWM Frequency (kHz)	16	16	16	16	16	8	8	8	8	
	PWM Motor Current Ripple (kHz)	32	32	32	32	32	16	16	16	16	
	Form Factor (rms/avg)	≤1.01	≤1.01	≤1.01	≤1.01	≤1.01	≤1.01	≤1.01	≤1.01	≤1.01	
<b>Control Input</b>	+8Vdc Supply Voltage	Internally Supplied		7.3-8.5	7.3-8.5	7.3-8.5	7.3-8.5	7.3-8.5	7.3-8.5	7.3-8.5	
	+8Vdc Supply Current (amps)			1.1	1.1	1.1	1.1	1.1	1.1	1.1	1.1
	±15Vdc Supply Voltage			14.3-15.5	14.3-15.5	14.3-15.5	14.3-15.5	14.3-15.5	14.3-15.5	14.3-15.5	14.3-15.5
	±15Vdc Supply Current (amps)	.37	.37	.38	.5	.47	.66	.87			
	+24Vdc Ext. Logic Voltage (volts)	This option is not currently available									
	+24Vdc Ext. Logic Current (amps sink)										
<b>Protective Functions</b>	UnderVoltage Trip (on power-up)	90 Vdc	90 Vdc	90 Vdc	90 Vdc	90 Vdc	255 Vdc	255 Vdc	255 Vdc	255 Vdc	
	UnderVoltage Trip (nominal)	90 Vdc	90 Vdc	90 Vdc	90 Vdc	90 Vdc	125 Vdc	125 Vdc	125 Vdc	125 Vdc	
	OverVoltage Trip	430 Vdc	430 Vdc	430 Vdc	430 Vdc	430 Vdc	430 Vdc	430 Vdc	430 Vdc	430 Vdc	
	OverTemperature Trip	80°C	80°C	118°C	118°C	118°C	90°C	90°C	90°C	90°C	
<b>Environment</b>	Internal heat dissipation in 45° ambient @ continuous current (not including regen)	84w	138w	37 w	84 w	120 w	240 w	254 w	465 w	675 w	
	Operation temperature (°C)	0 to 45°									
	Storage temperature (°C)	-20 to 70°									
	Humidity (non-condensing)	10% to 90%									
	Atmosphere	w/o corrosive gasses or dust									
	Altitude	Derate 5% per 1000ft above 3300ft									
	Vibration	0.5 g									

Table 2-1 SERVOSTAR® Electrical Specifications

Product Model		Cx03	Cx06	Sx03	Sx06	Sx10	Sx20	Sx30	Sx55	Sx85
<b>Current Loop</b>	Update Rate	62.5 $\mu$ S (16kHz)								
	Bandwidth	<2000Hz								
<b>Commutation Loop</b>	Update Rate	62.5 $\mu$ S (16kHz)								
	Output Waveform	Sinusoidal								
<b>Velocity Loop</b>	Update Rate	250 $\mu$ S (4kHz)								
	Bandwidth	<400Hz								
	Speed Command Resolution	Serial/Analog: 12Bit (Optional 14)								
	Maximum / Minimum Speed	500Hz / [Serial: 1 RPM or Analog: (1/2048) x VMAX]								
<b>Position Loop</b>	Update Rate	.01% ( $\mu$ P clock tolerance)								
<b>I/O Connector (C3 by pinout)</b>										
Analog Input (2,3)	Maximum Voltage	13 V differential								
	Input Resolution	12 Bit (Optional 14 Bit)								
	Sensitivity	4.9mV @12Bit								
	Voltage Range	-8 to +8V (Standard) / -10 to +10 V (Rescaled)								
	Input Impedance / CMR	> 10 K ohms / 50 dB								
	Long-term Drift	100 ppm (0.075%/°C)								
Fault Output Relay (5,6)	Max Capacity	1 A @ 24 Vdc								
Remote Enable (7,8) & Configurable Inputs (7,9,10,11)	Input Frequency	2.5kHz (Opto-isolated)								
	Input Voltage Range Min. On / Max. Off	12 to 24 V Nominal 10 V / 1 V								
	Current Demand per Input	20mA								
Configurable Digital Output (7,12) Open Collector	Output Voltage (max.) (Min. On)	0 to 48 V Nominal - Bi-directional 1V								
	Max. Output Current	60mA								
Configurable Analog Output (13,4)	Max. Output Current	1mA (1K $\Omega$ internal series resistance)								
	Sensitivity / Resolution	4.9mV / 12 Bit								
	Voltage Range	-10 to +10V								
<b>Encoder Equivalent Output (C4 by pinout)</b>										
A/B/I & Complements (1,2,4,5,7,8)	Output Voltage (high level) @25°C	2.5V min @ 20mA								
	Output Voltage (low level) @25°C	.5 max @ 20mA								
	RS 485 Line Drive Type	DS26C31TM								
<b>Remote Encoder Input (C8 by pinout)</b>										
A/B/I & Complements (1,2,4,5,7,8)	Input Voltage (high/low level) @25°C	5V / 0V nominal (+.2V / -.2V)								
	Input Impedance	100 $\Omega$								
	RS 485 Line Receiver Type	SN75173								
See section 6.4.2.4 Position Loop on page 6-6 for features using this input										



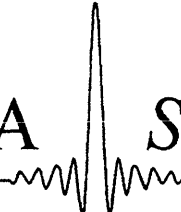
### B.3 Assurance Technologies Inc. Mini force/torque transducer

#### Metric Calibrations for Nano and Mini



Transducer models	Nano			Mini		
	50/500	25/250	12/125	80/4	40/2	20/1
<b>Standard Metric Calibrations</b>						
<b>Sensing Ranges</b>						
±Force x and y range, (N)	50	25	12.5	80	40	20
±Force z range, (N)	90	45	22.5	240	120	60
±Torque range, Nano: N-mm, Mini: N-m	500	250	125	4	2	1
<b>Resolution</b>						
Fx, Fy, (N)	0.05	0.025	0.0125	0.08	0.04	0.02
Fz, (N)	0.10	0.05	0.025	0.24	0.12	0.06
Tx, Ty, Tz, (N-mm)	0.25	0.125	0.0625	2.0	1.0	0.5
<b>Counts measurement value</b>						
Fx, Fy, Fz, (counts/N)	20	40	80	12.5	25	50
Tx, Ty, Tz, Nano: counts/N-mm, Mini: counts/N-m	4	8	16	500	1000	2000
<b>Stiffness</b>						
Kx, Ky, (N/m)	9x10 <sup>6</sup>			11x10 <sup>6</sup>		
Kz, (N/m)	12x10 <sup>6</sup>			24x10 <sup>6</sup>		
Ktx, Kty, (N-m/rad)	250			3300		
Ktz, (N-m/rad)	390			4300		
<b>Weight, (grams)</b>						
Transducer, tool & mount adaptor	9.4			54		
<b>Material</b>						
Transducer	Hardened Stainless Steel			Hardened Stainless Steel		
Tool adaptor	Aircraft Aluminum			Aircraft Aluminum		
Mounting adaptor	Aircraft Aluminum			Aircraft Aluminum		
Mounting ring-plug adaptor	Not available			Not available		
Tool ring-plug adaptor	Not available			Not available		
<b>Overload Protection</b>						
Max. Fxy allowable, (±N)	350			1300		
Max. Fz allowable, (±N)	800			2900		
Max. Txyz allowable, (±N-m)	2.5			25		

## B.4 ATA Sensors ARS-04E magneto-hydrodynamic sensor

ATA  Sensors

P.O. Box 92303, ABQ, NM 87199-2303  
ph: (505) 823-1320, fx: (505) 823-1560  
e-mail: [atasensors@aptec.com](mailto:atasensors@aptec.com)

**Doc. No. S001001-**  
**MHD ANGULAR RATE SENSOR**  
**Model: ARS-04E SN: 511**  
**Date Calibrated: 09/08/00**  
Test Data

Scale Factor @ 10 Hz $K\omega$	=	156.6 mV/(rad/s)
Low Frequency -3dB Point	=	0.378 Hz
Low Frequency +45° Point	=	0.466 Hz
Pole Frequency f1	=	0.37 Hz
Pole Frequency f2	=	0.05 Hz
High Frequency -3dB Point	=	1270 Hz
Input Power @ +/- 15 VDC	Positive Side @	15.10 mA x 15VDC = 226.5 mW
	Negative Side @	15.15 mA x 15VDC = 227.3 mW
		Total Power = 453.8 mW
Output RMS Voltage Noise	=	9.9 mV
Output Offset (After 10 Minute Warming)	=	-23.7 mV
Output Impedance	=	<100 $\Omega$

### Wiring Instructions\*

Red Lead: +15 VDC  
White Lead: -15 VDC  
Black Lead: Power and Signal Common  
Yellow Lead: Output Signal

\*If used with an ILC, follow the ILC wiring instructions.  
**CAUTION: Improper hookup can result in transducer failure.**

### **WARNING:**

**Sharp Impacts May Cause Damage to This Sensor.**  
**This Sensor Contains Mercury and is Not Designed**  
**for Safe Operation or Use Outside the Specifications.**

## B.5 ATA Sensors digital compensation filter

### Digital compensation filter for MHD sensor, s/n 496

The digital compensation filter is implemented as a linear difference equation.

1. Set vectors  $y$  and  $mhd$  to the measured head velocity signal.  $y := head$   
The sequence  $vel$  will become the compensated vector.  $mhd := head$
2. Set desired low frequency pole frequencies:

<u>mhd sensor 496 has 2 poles:</u>	<u>physical</u>	<u>electronic</u>
original LF pole freqs (Hz):	$f_1 := 0.37$	$f_2 := 0.08$
desired LF pole freqs (Hz)	$f_{d1} := 0.05$	$f_{d2} := 0.01$

calculate DT frequency equivalents (rad/sec)      sampling period:  $\tau := 0.01$

original corners

desired corners

$$\begin{array}{llll} \Omega_1 := 2 \cdot \pi \cdot f_1 \cdot \tau & \Omega_1 = 0.023 & \Omega_{d1} := 2 \cdot \pi \cdot f_{d1} \cdot \tau & \Omega_{d1} = 3.142 \times 10^{-3} \\ \Omega_2 := 2 \cdot \pi \cdot f_2 \cdot \tau & \Omega_2 = 5.027 \times 10^{-3} & \Omega_{d2} := 2 \cdot \pi \cdot f_{d2} \cdot \tau & \Omega_{d2} = 6.283 \times 10^{-4} \end{array}$$

3. Create filter difference equation with scale factor that converts measurement to rad/sec

equation coefficients

$$C_1 := \frac{\left( e^{-\Omega_1} \cdot e^{-\Omega_2} \right)}{e^{-\Omega_{d1}} \cdot e^{-\Omega_{d2}}} \quad C_1 = 1.025$$

$$C_2 := e^{-\Omega_1} + e^{-\Omega_2} \quad C_2 = 1.972 \quad C_3 := e^{-\Omega_1} \cdot e^{-\Omega_2} \quad C_3 = 0.972$$

$$C_4 := e^{-\Omega_{d1}} + e^{-\Omega_{d2}} \quad C_4 = 1.996 \quad C_5 := e^{-\Omega_{d1}} \cdot e^{-\Omega_{d2}} \quad C_5 = 0.996$$

4. Apply difference equation to head velocity sequence

$$vel := \left| \begin{array}{l} \text{for } k \in 2..length(y) \\ \quad \left| \begin{array}{l} temp \leftarrow C_1 \cdot (mhd_k - C_2 \cdot mhd_{k-1} + C_3 \cdot mhd_{k-2}) + C_4 \cdot y_{k-1} - C_5 \cdot y_{k-2} \\ y_k \leftarrow temp \end{array} \right. \\ \quad \left. y \right. \end{array} \right.$$



**WATSON INDUSTRIES, INC.**

3041 Melby Road Eau Claire, WI 54703

Phone (715) 839-0628 \* FAX (715) 839-8248

www.watson-gyro.com \* email: support@watson-gyro.com

**Certificate of Conformance**

Product P/N - ARS-C141-1AR8P

S/N 14150

Rate %/Second	Ideal (VDC)	X Axis Data (VDC)	X Axis Error (%/Second)
-300	-10.00	-10.00	0.00
-225	-7.50	-7.49	-0.30
-150	-5.00	-4.99	-0.30
-75	-2.50	-2.50	0.00
0	0.00	0.01	0.30
75	2.50	2.51	0.30
150	5.00	5.01	0.30
225	7.50	7.51	0.30
300	10.00	10.01	0.30
			<b>PASS</b>

	X Axis Thermal Bias	X Axis 10 Min. Drift
+20° C	0.00V	26 mV
-20° C	0.28V	<b>PASS</b>
+50° C	0.16V	<b>Scale Factor</b>
+20° C	-0.10V	29.99°/s/V
		<b>PASS</b>

*NEW S/N 14150*

$2.6 \frac{mV}{MIN} = 0.0780 \frac{\%}{MIN}$   
 $0.0018 \frac{\%}{s} \text{ DRIFT}$


340

This product has been tested with equipment traceable to N.I.S.T. Standards and is hereby certified to meet or exceed manufacturer's specifications.

*[Signature]*  
 Signature

1/10/02  
 Date

**WARNING**  
 THIS UNIT IS MADE TO WITHSTAND OVER 200 g's. A DROP OF 10 INCHES MAY EXCEED THIS LIMIT. DRILLING AND SAWING OPERATIONS ON THE CASE ARE LIKELY TO DAMAGE THE UNIT.

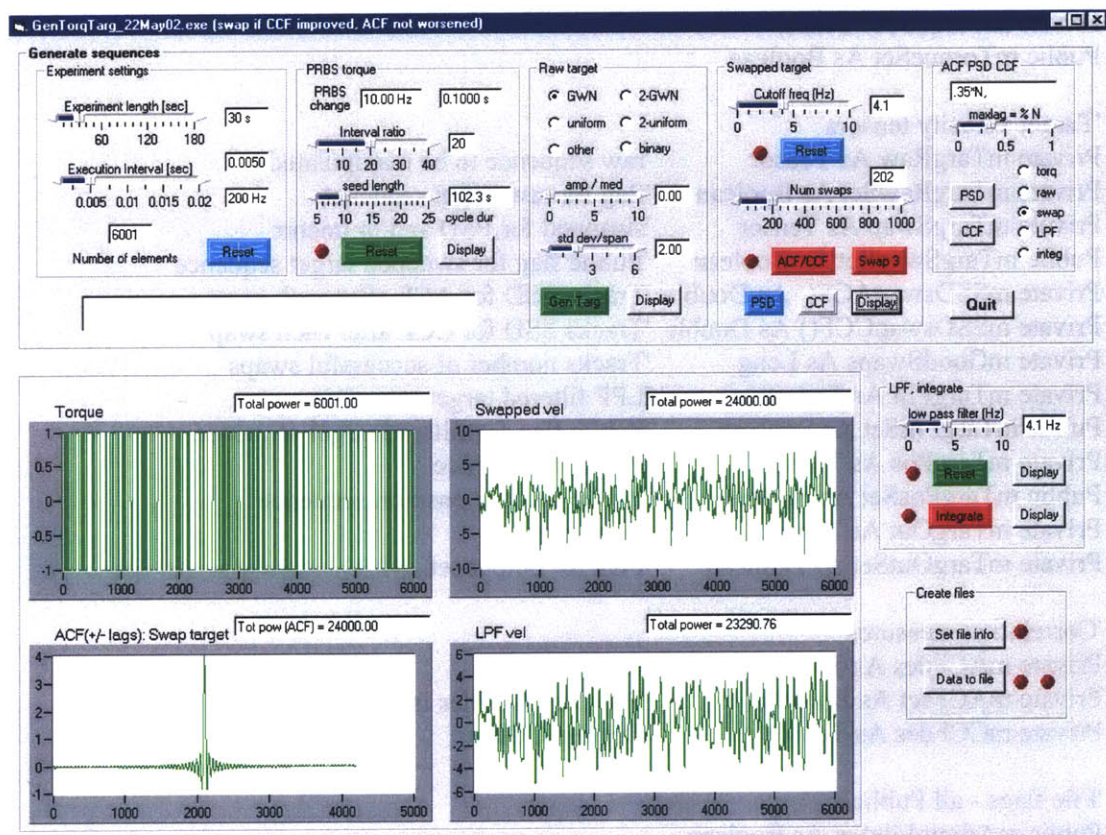


B.6 Watson Industries ARS-C141 1AR8P angular rate sensor

## Appendix C

### Code to generate the uncorrelated torque and target inputs

Appendix C includes only the portions of the program's code that were essential to generating the uncorrelated torque and target sequences. Much of the code that was required to execute the program, such as the code for of the GUIs that allowed the user to interact with the program, has been removed. The code was written in Visual Basic 6.0.



**Figure C.1:** Screen capture of the interface from the program that was used to generate the uncorrelated torque and target sequences.

## Excerpts of the code for “GenUncorTorqTarg”

'Module level variable declarations

'Requires the Plexus tensor PlexusTensor.dll  
'National Instruments Component works UI

'Set option explicit so VB verifies all variables declared  
Option Explicit

'Module level variables

'Common experiment settings

Private mExpLen As Integer 'experiment length in seconds  
Private mExecInt As Double 'interval at which to execute thru sequence  
Private mN As Long 'number of terms in sequence  
Private mSettingsSet As Boolean 'flag for experimental settings

'Torque sequence

Private mTorque As Tensor  
Public mTorqueSet As Boolean

'Target velocity tensors

Private mTargRaw As Tensor 'raw sequence to be manipulated  
Private mTargRawSet As Boolean 'flag for raw target sequence  
Private mTargSwap As Tensor 'swapped for PSD and/or uncorr  
Public mTargSwapSet As Boolean 'Public flag for swapped target sequence  
Private mSSDswapACF() As Double 'Tracks SSD for ACF after each swap  
Private mSSDswapCCF() As Double 'Tracks SSD for CCF after each swap  
Private mGoodSwaps As Long 'Tracks number of successful swaps  
Private mTargFilt As Tensor 'LPF filtered target sequence  
Public mTargFiltSet As Boolean 'Public flag for filtered target sequence  
Private mTargPos As Tensor 'Position sequence  
Public mTargPosSet As Boolean 'Public flag for position sequence  
Private mTargOut As Tensor 'Output tensor  
Private mTargOutSet As Boolean 'Flag for output set

'Correlation measures

Private mACFdes As Tensor 'Desired ACF  
Private mACFset As Boolean 'Flag that ACFdes is set  
Private mCCFdes As Tensor 'Desired CCF

'File flags - all Public

Public mAdminFileSet As Boolean  
Public mTorqueFileSet As Boolean  
Public mSwapFileSet As Boolean  
Public mACFssdFileSet As Boolean  
Public mCCFssdFileSet As Boolean  
Public mFiltFileSet As Boolean  
Public mPosFileSet As Boolean

Private Sub Form\_Load()

' Description:

' This Program generates PRBS torque and random target sequences.

' The spectrum of each sequence is chosen, and the target velocity sequence

' can be made so that the torque and target velocity sequences are uncorrelated.

' The position sequence is created by integrating the velocity sequence.

' The position sequence is manipulated after the integration so that it begins

' and ends at zero position.

' The sequences are made to be output at a certain frequency. Above 100 Hz is beyond

' the capability of the MS Timer. The WinTimer DLL, by Serge LaFontain, must be used.

'Module level variables

'Common experiment settings

'Private mExpLen As Integer 'experiment length in seconds

'Private mExecInt As Double 'interval at which to execute thru sequence

'Private mN As Long 'number of terms in sequence

'Public mSettingsSet As Boolean 'flag for experimental settings

'Torque sequence

'Private mTorque as Tensor

'Public mTorqueSet As Boolean

'Target velocity tensors

'Private mTargRaw As Tensor 'raw sequence to be manipulated

'Public mTargRawSet As Boolean 'flag for mTargRaw generated

'Private mTargSwap As Tensor 'swapped for PSD and/or uncorr

'Public mTargSwapSet As Boolean 'flag for mTargSwap generated

'Private mSSDswapACF() As Double 'Tracks SSD for ACF after each swap

'Private mSSDswapCCF() As Double 'Tracks SSD for CCF after each swap

'Private mGoodSwaps As Long 'Tracks number of successful swaps

'Private mTargFilt As Tensor 'LPF filtered target sequence

'Public mTargFiltSet As Boolean 'Flag for filtered Target sequence

'Public mTargPos as tensor

'Correlation measures

'Private mPSDdes As Tensor 'Desired PSD

'Private mACFdes As Tensor 'Desired ACF

'Private mACFset As Boolean 'Flag that ACFdes is set

'Private mCCFdes As Tensor 'Desired CCF

'File flags - all Public

'Public mAdminFileSet as Boolean

'Public mTorqueFileSet as Boolean

'Public mSwapFileSet As Boolean

'Public mACFssdFileSet as Boolean

'Public mCCFssdFileSet as Boolean

'Public mFiltFileSet As Boolean

'Public mPosFileSet As Boolean

```

'Boolean flags all set false
'Common experimental setting
mSettingsSet = False

'Torque sequence
mTorqueSet = False

'Target sequence
mTargRawSet = False
mTargSwapSet = False
mTargFiltSet = False
mTargPosSet = False
mTargOutSet = False
mACFset = False

'Datafiles
mAdminFileSet = False
mTorqueFileSet = False
mSwapFileSet = False
mACFssdFileSet = False
mCCFssdFileSet = False
mFiltFileSet = False
mPosFileSet = False

'Update experiment length text box
txtExpLen.Text = Format(Int(sldExpLen.Value), "0") + " s"

'Update execution interval text box
txtExecInt.Text = Format(sldExecInt.Value, "0.0000")
txtExecFreq.Text = Format(Int(1 / sldExecInt.Value), "0") + " Hz"
txtNumElems.Text = Int(1 + sldExpLen.Value / sldExecInt.Value)

'Set the PRBS text windows.
'Update after experiment settings are set
txtPRBSFreq.Text = ""
txtPRBSint.Text = ""
txtIntRatio.Text = sldIntRatio.Value
txtPRBSdur.Text = ""

'Update amplitude/mean and standar deviation in text boxes
txtAmpMed.Text = Format(sldAmpMed.Value, "0.#0")
txtStdev.Text = Format(sldStdev.Value, "0.#0")

'Update maxlag percentage of sequence length
Dim percent As Double
percent = sldMaxLag.Value
txtMaxLag.Text = Format(percent, "#.#0") + " * N"

'Update number of swaps for interchange
txtNumSwaps.Text = Format(sldNumSwaps.Value, "0")

```



```
'Update low and high frequency for desired PSD
  txtHiFreq.Text = Format(sldHiFreq.Value, "0.0")
```

```
'LPF filter frequency
  txtLPF.Text = Format(sldLPF.Value, "0.0") + " Hz"
```

```
End Sub
```

```
Private Sub butGenmTargRaw_ValueChanged(ByVal Value As Boolean)
'Generates a raw sequence that will be manipulated into the Target sequence
```

```
'Module variables
```

```
'Common experiment settings
```

```
  'Private mExecInt As Double      'interval at which to execute thru sequence
```

```
  'Private mN As Long             'number of terms in sequence
```

```
  'Private mSettingsSet As Boolean 'flag for experimental settings
```

```
'Target velocity tensors
```

```
  'Private mTargRaw As Tensor     'raw sequence to be manipulated
```

```
  'Private mTargRawSet As Boolean 'flag for mTargRaw generated
```

```
'Local variables
```

```
  Dim N As Long                  'sequence length set by sldSeqLength.Value
```

```
  Dim SampInt As Double          'sampling interval set by sldSampFreq.Value
```

```
  Dim SigType As String
```

```
  Dim AmpMed As Double
```

```
  Dim stdev As Double
```

```
  Dim y1min As Double            'y bounds for ramp sequence
```

```
  Dim y1max As Double
```

```
  Dim y2min As Double
```

```
  Dim y2max As Double
```

```
  Dim rmin As Double
```

```
  Dim rmax As Double
```

```
  Dim imin As Long
```

```
  Dim imax As Long
```

```
  Dim i As Long
```

```
  Dim HalfSeq As Long
```

```
  Dim TotPow As Double
```

```
'Generate raw target sequence
```

```
If butGenmTargRaw.Value = True Then
```

```
  If mSettingsSet = True Then
```

```
    'Set flags of other target sequences to false
```

```
    mTargSwapSet = False
```

```
mTargFiltSet = False
mTargPosSet = False
```

```
'Sequence parameters
```

```
  N = mN                                'length set in Experimental settings form
  SampInt = mExecInt
  AmpMed = sldAmpMed.Value              'sets peak amplitude or median
  stdev = sldStdev.Value                 'sets standard deviation
```

```
'Generate target sequence mY
```

```
  If optGWN.Value = True Then           'generate Gaussian white sequence
    mTargRaw = GenGaussNoise(0, N - 1, SampInt, AmpMed, stdev)
    SigType = "Raw vel: GWN"
```

```
  ElseIf opt2GWN.Value = True Then
    mTargRaw = GenDualGWN(0, N - 1, SampInt, AmpMed, stdev)
    SigType = "Raw vel: Dual GWN"
```

```
  ElseIf optRamp.Value = True Then     'generate ramp
    imin = 0
    imax = N - 1
    mTargRaw = GenUniform(imin, imax, SampInt, -AmpMed, AmpMed)
    SigType = "Raw vel: ramp"
```

```
  ElseIf opt2Ramps.Value = True Then   'generate two ramps
    imin = 0
    imax = N - 1
    y1min = -AmpMed - stdev / 2
    y1max = -AmpMed + stdev / 2
    y2min = AmpMed - stdev / 2
    y2max = AmpMed + stdev / 2
    mTargRaw = Gen2Uniform(imin, imax, SampInt, y1min, y1max, y2min, y2max)
    SigType = "Raw vel: two ramps"
```

```
  ElseIf optBinary.Value = True Then   'generate ordered binary sequence
    mTargRaw = GenBinary(0, N - 1, SampInt, -AmpMed, AmpMed)
    SigType = "Raw vel: Binary"
```

```
  "ElseIf optPRBS.Value = True Then    'generate PRBS
    "mTargRaw = GenPRBS(12, 0, N - 1, SampInt)
    "SigType = "Raw vel: PRBS"
    'change amplitude to +/- 1, and then scale with AmpMed
    "imin = LBound(mTargRaw.Val)
    "imax = UBound(mTargRaw.Val)
    "For i = imin To imax
      "mTargRaw.Val(i) = AmpMed * (2 * mTargRaw.Val(i) - 1)
    "Next i
```

```
  ElseIf optOther.Value = True Then    'generate other
    rmin = 0
    rmax = (N - 1) * SampInt
```

```

    mTargRaw = GenSine(rmin, rmax, SampInt, AmpMed, 4, 0)
    SigType = "Raw vel: 4 Hz sine"

Else
    mTargRaw = GenGWN(0, N - 1, SampInt, AmpMed, stdev)
    SigType = "Raw vel: GWN"

End If

imin = LBound(mTargRaw.Val)
imax = UBound(mTargRaw.Val)

'Calculate total power from sequence values
TotPow = 0
For i = imin To imax
    TotPow = TotPow + mTargRaw.Val(i) * mTargRaw.Val(i)
Next i

'Set flag indicating sequence created
mTargRawSet = True

'Display sequences, total power
plot_D.PlotY mTargRaw.Val
lab_D.Caption = SigType
txt_D.Text = "Total power = " + Format(TotPow, "#.#0")

Else 'mSettingsSet flag not set
    txtComment.Text = "Experiment settings not set."
    butGenmTargRaw.Value = False
    mTargRawSet = False
    'Set flags of other target sequences to false
    mTargSwapSet = False
    mTargFiltSet = False
    mTargPosSet = False

End If

Else 'butGenmTargRaw turned off

    mTargRawSet = False 'Set flag indicating sequence not created
    'Set buttons for other target sequences to false
    'This will set flags to zero in their routines
    mTargSwapSet = False
    butPSDuncorr = False
    butLPF = False
    butIntegrate = False

End If

End Sub

```

Private Sub butGenPSDdes\_ValueChanged(ByVal Value As Boolean)

' Create desired two sided PSD for the target signal  
' Generates tensor mACFdes to be used in swap algorithms  
' Algorithm set up for pos and neg lag ACF and a PSD defined  
' from -0.5 to 0.5 dt cps

'Module level variables

'Target velocity tensors  
'Private mTargRaw As Tensor 'raw sequence to be manipulated  
'Private mTargRawSet As Boolean 'flag for raw target sequence

'Correlation measures

'Private mACFdes As Tensor 'Desired ACF  
'Private mACFset As Boolean 'Flag that ACFdes is set

'Local variable declarations

Dim N As Long 'length of mTargRaw  
Dim percentLen As Double 'percent of N for Maxlag  
Dim maxlag As Long 'Max lag set by sldMaxLag.value  
Dim minlag As Long  
Dim SampInt As Double 'sampling interval of sequence  
Dim maxFreq As Double '1/2 the sampling frequency

Dim PSDdes As Tensor 'Desired PSD

Dim ACFFraw As Tensor 'ACF of original sequence  
Dim PSDraw As Tensor 'PSD of original sequence

Dim TotPowACFFraw As Double 'N \* ACF(0)  
Dim SumElemsPSDraw As Double 'Sum of elements in PSD  
Dim TotPowPSDraw As Double 'Sum of elements scaled  
Dim imin As Long  
Dim imax As Long  
Dim i As Long  
Dim M As Long 'length of ACF, set by maxLag  
Dim PSDdisp As Tensor 'Display 0 freq at 0 element

Dim hiFreq As Long 'Freq content of PSDdes exists  
Dim loFreq As Long 'between loFreq and hiFreq  
Dim hiElem As Long 'element equivalent to hiFreq  
Dim loElem As Long

Dim SumElemsPSDdes As Double 'Sum of PSDdes elements  
Dim TotPowPSDdes As Double 'PSDdesSum scaled by ratio of sequence length to maxlag  
Dim TotPowACFdes As Double 'PSDmYSum scaled by ratio seq length to maxlag

Dim k As Long

'Initiate new PSDdes when button turned on  
If butGenPSDdes.Value = True Then

```

'Execute only if mTargRaw has been created - want to match sequence power
  If mTargRawSet = True Then

    'initiate sequence
      butPSDinit = True
    'length of sequence mTargRaw
      N = UBound(mTargRaw.Val) - LBound(mTargRaw.Val) + 1

    'Calculate ACF, PSD, and total power of raw sequence

      'min and max lags
      percentLen = sldMaxLag.Value
      maxlag = Int(percentLen * (N - 1))
      minlag = -maxlag
      txtMaxLag.Text = Format(percentLen, "#.#0") + "*N, " + Format(maxlag, "#")

      'ACF and PSD
      ACFraw = ACF(mTargRaw, minlag, maxlag)
      PSDraw = PSDfromACF(ACFraw, None)

    'Total power from ACFraw with MaxLag
      TotPowACFraw = N * ACFraw.Val(0) 'sequence length * zero lag correlation

    'Total power from PSDraw with MaxLag
      imin = LBound(PSDraw.Val)
      imax = UBound(PSDraw.Val)
      M = imax - imin + 1 'length of ACF and PSD
      SumElemsPSDraw = 0
      For i = imin To imax
        SumElemsPSDraw = SumElemsPSDraw + PSDraw.Val(i)
      Next i
      'PSD power is the element sum scaled by ratio of sequence length to PSD length
      TotPowPSDraw = (N / M) * SumElemsPSDraw

    'Define desired PSD with same sum of elements as for PSDmY
      'ACF should be based on PSD so TotPowPSD = (N/M)*SumElems
      'Defining the positive frequencies, which must also be mirrored by neg freqs

      PSDdes = PSDraw

      SampInt = 1# / ((M - 1) * PSDraw.Inc(1)) 'Gives sampling interval
      maxFreq = 1# / (2 * SampInt)

      'loFreq = sldLoFreq.Value
      loFreq = 0
      hiFreq = sldHiFreq.Value

      'loElem = Int(loFreq * SampInt * (m - 1))
      loFreq = 0
      hiElem = Int(hiFreq * SampInt * (M - 1))

```

```

'initialize all elements to zero
For i = imin To imax
    PSDdes.Val(i) = 0
Next i

'scale elements in desired freq range to have amplitude of 1
For i = loElem To hiElem
    PSDdes.Val(i) = 1
    PSDdes.Val(-i) = 1      'negative freq elements
    'PSDdes.val(imin+imax-i) 'more general relation that allows PSDs
    'defined for -0.5 to 0.5 or 0.0 to 1.0
Next i

'scale PSD elements so their sum has same sum as PSDmY

'original sum
SumElemsPSDdes = 0
For i = imin To imax
    SumElemsPSDdes = SumElemsPSDdes + PSDdes.Val(i)
Next i

'scale elements so that SumElemsPSDdes equals SumElemsPSDmX
'PSDmX elements
For i = imin To imax
    PSDdes.Val(i) = (SumElemsPSDraw / SumElemsPSDdes) * PSDdes.Val(i)
Next i

'recalculate power to ensure match
SumElemsPSDdes = 0
For i = imin To imax
    SumElemsPSDdes = SumElemsPSDdes + PSDdes.Val(i)
Next i
TotPowPSDdes = (N / M) * SumElemsPSDdes

'desired ACF
mACFdes = ACFfromPSD(PSDdes)
TotPowACFdes = N * mACFdes.Val(0)
mACFset = True      'Flag that desired ACF created

'clear text box
txtComment.Text = ""

'Turn button off when calculation done
butPSDinit = False

'Plots
plot_A.PlotY mACFdes.Val
lab_A.Caption = "Desired Ctt"
txt_A.Text = "Total power = " + Format(TotPowACFdes, "#.#0")

```

```

    PSDdisp = DisplayTensor(PSDdes)
    plot_D.PlotY PSDdisp.Val
    lab_D.Caption = "Desired PSD (disp)"
    txt_D.Text = "Total power = " + Format(TotPowPSDdes, "#.#0")

Else 'mTargRaw had not been created
    txtComment.Text = "mTargRaw has not been generated."
    butGenPSDdes.Value = False 'turn button off
End If

Else 'butGenPSDdes turned off
    mACFset = False 'Flag states PSDdes not created
    butPSDinit = False

    plot_A.PlotY 0
    lab_A.Caption = "Desired ACF"
    txt_A.Text = ""

    plot_D.PlotY 0
    lab_D.Caption = "Desired PSD (disp)"
    txt_D.Text = ""

End If
End Sub

Private Sub butGenTorque_ValueChanged(ByVal Value As Boolean)
' Description
' Create Torque sequence that executes at interval mExecInt, but appears to be
' a PRBS being executed at PRBSint. The sequence has mN elements. It takes
' mExpLen secs to execute through entire sequence

'Module variables
'Private mExpLen As Integer          'experiment length in seconds
'Private mExecInt As Double         'interval at which to execute thru sequence
'Private mN As Long                'number of terms in sequence
'Private mSettingsSet As Boolean    'flag for experimental settings

'Torque sequence
'Private mTorque As Tensor
'Public mTorqueSet As Boolean

'Local variables
Dim reqLen As Long                'required length of PRBS sequence
Dim R As Integer                  'ratio of PRBS to execution intervals
Dim PRBSint As Double             'effective PRBS interval
Dim M As Long                    'PRBS seedlength
Dim G As Long                    'number of complete groups of terms
Dim gr As Integer                 'group counter
Dim L As Integer                  '# elements in last group, no including last PRBS element

```

```

Dim initPRBS As Tensor      'output of Plexus PRBS algorithm
Dim TorSeq() As Double     'array that holds the elements of mTorque tensor
Dim i As Integer
Dim j As Integer

Dim TotPow As Double

'Generate Torque sequence
If butGenTorque = True Then

'Generate torque PRBS only if experiment settings are set
If mSettingsSet = True Then

'Turn butTorInit on. Signals process underway.
    butTorInit.Value = True

'Sequence parameters
    R = sldIntRatio.Value      'ratio of PRBS and execution intervals
    PRBSint = R * mExecInt     'interval at which PRBS appears to be executed
    M = sldSeedlength.Value    'seedlength
    reqLen = 1 + Int((mN - 1) / R) 'number of PRBS elements required to last for mExpLen
    G = reqLen - 1             'number of full groups that will be filled
    L = (mN - 1) - (G * R)     'number of elements in final group
                                'does not include final PRBS element

'Generate basic PRBS with amp (+/- 1) and dt = 1
    initPRBS = GenPRBS(M, 0, G, 1)

'change amplitude to +/- 1
    For i = LBound(initPRBS.Val) To UBound(initPRBS.Val)
        initPRBS.Val(i) = 1 * (2 * initPRBS.Val(i) - 1)
    Next i

'plot initial PRBS
    'plot_A.PlotY initPRBS.Val
    'lab_A.Caption = "PRBS"

'Create mN element array by "filling" in PRBS to execute at mExecInt
'Each PRBS element is first element of a group of the same values

ReDim TorSeq(0 To mN - 1)    'Array for sequence. Will become a tensor.

'Fill complete "groups"
    For gr = 0 To G - 1      'cycle through each PRBS element
        For i = 0 To R - 1  'elements of each group
            j = i + gr * R   'array index
            TorSeq(j) = initPRBS.Val(gr)
        Next i
    Next gr

'Fill last group. May be a partial group.

```



```

For j = (G * R) To (G * R) + L
    TorSeq(j) = initPRBS.Val(G)
Next j

'Convert array to a tensor
mTorque = GenTensor(TorSeq, mExecInt, "Amplitude", "", "Time", "s")

'Calculate total power from sequence values
TotPow = 0
For i = LBound(mTorque.Val) To UBound(mTorque.Val)
    TotPow = TotPow + mTorque.Val(i) * mTorque.Val(i)
Next i

'plot PRBS
plot_A.PlotY mTorque.Val
lab_A.Caption = "Torque"
txt_A.Text = "Total power = " + Format(TotPow, "#.##0")

'Set flag that torque sequence has been created
mTorqueSet = True
txtComment.Text = "Torque PRBS generated."

'Turn butTorInit off when calculation complete.
butTorInit.Value = False

Else 'mSettingsSet = False
    butGenTorque = False 'Will execute just as if button manually turned off
    txtComment.Text = "Experiment settings not set."

End If

Else 'if butGenTorque = false (turned off), set flag so settings can be reset
    mTorqueSet = False
    txtComment.Text = "Reset PRBS parameters."
    butTorInit.Value = False 'Turn butTorInit off.

End If
End Sub

Private Sub butLPF_ValueChanged(ByVal Value As Boolean)
' Description:
' Convolves swapped sequence with 2 sided LPF.
' Time domain LPF is created by invDFT of the ideal LPF spectrum

'Low pass filter mTargSwap, create mTargFilt

'Module level variables
'Target velocity tensors
'Private mTargSwap As Tensor 'swapped for PSD and/or uncorr

```

```

'Public mTargSwapSet As Boolean    'Public flag for swapped target sequence
'Private mTargFilt As Tensor      'LPF filtered target sequence
'Public mTargFiltSet As Boolean    'Flag for filtered target sequence

'Local variable declarations
Dim N As Long                    'length of mTargSwap
Dim maxlag As Long               'Max lag set by sldMaxLag.value
Dim minlag As Long
Dim ACForig As Tensor            'ACF of original sequence
Dim PSDorig As Tensor            'PSD of original sequence
Dim TotPowACForig As Double      'N * ACF(0)
Dim SumElems As Double           'Sum of elements in PSD
Dim TotPowPSDorig As Double      'Sum of elements scaled

Dim coFreq As Double             'LP cutoff freq
Dim ACFfilt As Tensor
Dim PSDfilt As Tensor
Dim TotPowACFfilt As Double      'N * ACF(0)
Dim TotPowPSDfilt As Double      'Sum of elements scaled

Dim PSDdisp As Tensor           'Display 0 freq at 0 element

Dim imin As Long
Dim imax As Long
Dim i As Long
Dim M As Long                   'length of ACF, set by maxLag

'Start with butLPFinitiated off
butLPFinitiated.Value = False

'Execute when button turned on
If butLPF = True Then

    'Execute if mTargSwap has been created
    If mTargSwapSet = True Then

        'turn butLPFinitiated on
        butLPFinitiated.Value = True

        'length of sequence mTargSwap
        N = UBound(mTargSwap.Val) - LBound(mTargSwap.Val) + 1

        'Calculate ACF, PSD, and total power of original mTargSwap sequence
        'min and max lags
        maxlag = N - 1
        minlag = -maxlag
        M = maxlag - minlag + 1

        'Generate filtered Target sequence
        coFreq = sldLPF.Value
        mTargFilt = LPF(mTargSwap, coFreq) 'Convolve sequence with ideal LPF in time domain

```

```

mTargFiltSet = True

'ACF and PSD
ACFfilt = ACF(mTargFilt, minlag, maxlag)
PSDfilt = PSDfromACF(ACFfilt, None)

'Total power from ACFfiltT
TotPowACFfilt = N * ACFfilt.Val(0) 'sequence length * zero lag correlation

'Total power from PSDfiltT
SumElems = 0
For i = LBound(PSDfilt.Val) To UBound(PSDfilt.Val)
    SumElems = SumElems + PSDfilt.Val(i)
Next i
'PSD power is the element sum scaled by ratio of sequence length to PSD length
TotPowPSDfilt = (N / M) * SumElems

'turn butLPFinitiated off
butLPFinitiated.Value = False

'plots of filtered sequence and PSD

'filtered
plot_A.PlotY mTargFilt.Val
lab_A.Caption = "LPF vel"
txt_A.Text = "Total power(LPF)= " + Format(TotPowACFfilt, "#.#0")

PSDdisp = DisplayTensor(PSDfilt)
plot_B.PlotY PSDdisp.Val
lab_B.Caption = "PSD(dis): LPF vel"
txt_B.Text = "Total power(LPF)= " + Format(TotPowPSDfilt, "#.#0")

'Swapped sequence not created
Else
    txtComment.Text = "Swapped target has not been generated."
    butLPF = False
End If

Else 'button turned off
    mTargFiltSet = False
    mTargPosSet = False
    butIntegrate = False
End If
End Sub

```

Private Sub butPSDuncorr\_ValueChanged(ByVal Value As Boolean)

' Subroutine uses a two element stochastic swap to order the target sequence  
' to have a desired ACF and to be statistically uncorrelated from the torque sequence.  
' Cross-correlation is evaluated symmetrically, for equal numbers of neg and pos lags  
' Output is the tensor mTargSwap.  
' Developed for correlation measures with negative and positive lags

'Module variables

'Common experiment settings

'Private mExpLen As Integer           'experiment length in seconds  
'Private mExecInt As Double         'interval at which to execute thru sequence  
'Private mN As Long                 'number of terms in sequence  
'Private mSettingsSet As Boolean     'flag for experimental settings

'Torque sequence

'Private mTorque As Tensor  
'Public mTorqueSet As Boolean

'Target velocity tensors

'Private mTargRaw As Tensor           'raw sequence to be manipulated  
'Private mTargRawSet As Boolean       'flag for raw target sequence  
'Private mTargSwap As Tensor         'swapped for PSD and/or uncorr  
'Public mTargSwapSet As Boolean       'Public flag for swapped target sequence  
'Private mSSDswapACF() As Double     'Tracks SSD for ACF after each swap  
'Private mSSDswapCCF() As Double     'Tracks SSD for CCF after each swap  
'Private mGoodSwaps As Long           'Tracks number of successful swaps

'Correlation measures

'Private mACFdes As Tensor           'Desired ACF  
'Private mACFset As Boolean           'Flag that ACFdes is set  
'Private mCCFdes As Tensor           'Desired CCF

'Local variables

Dim X As Tensor           'Set to torque sequence, mTorque  
Dim Y As Tensor           'Set to target sequence which will be swapped, mTargRaw

Dim imin As Long  
Dim imax As Long  
Dim i As Long  
Dim N As Long             'number elements in torque and target sequences  
Dim Nt As Long

Dim maxlag As Long       'set to the length of mACFdes  
Dim minlag As Long  
Dim M As Long             'number elements in ACF and PSD

Dim ACForig As Tensor     'measures for the sequence before stochastic swap  
Dim PSDorig As Tensor  
Dim CCForig As Tensor

```

Dim CCFdes As Tensor      'desired CCF, set to 0 for uncorr sequences
Dim cmin As Long
Dim cmax As Long

Dim ACFy As Tensor       'current measures
Dim PSDy As Tensor
Dim CCFy As Tensor

Dim ACFswap As Tensor    'efficiently updated measures
Dim PSDswap As Tensor
Dim CCFswap As Tensor

Dim ACFcalc As Tensor    'calculated directly from swapped sequence
Dim PSDcalc As Tensor
Dim CCFcalc As Tensor

Dim ACFdif As Tensor     'compare ACFcalc to final updated ACFold
Dim CCFdif As Tensor

Dim PSDdisp As Tensor    'use with DisplayTensor( ) to plot 0 freq 0 position

Dim TotPowSeq As Double
Dim TotPowACF As Double
Dim SumElemsPSD As Double
Dim TotPowPSD As Double

Dim ssdACFy As Double
Dim ssdCCFy As Double
Dim ssdACFswap As Double
Dim ssdCCFswap As Double
Dim MSEorig As Double    'MSE of original CCF
Dim MSEswap As Double    'MSE of swapped CCF

Dim j1 As Long           'randomly selected indices to swap
Dim j2 As Long
Dim MaxSwaps As Long
Dim PossSwap As Long
Dim k As Long
Dim temp As Double

' Set indicators and tensor flags

'Ensure led indicator turned off
butInitiated.Value = False

'Set flags for the LPF and position (target) sequences to false.
'This should be done whether butPSDuncorr is turned on, or reset to off.
mTargFiltSet = False
mTargPosSet = False

' Initiate routine if butPSDuncorr turned on, and other swapping routines turned off

```

```
If (butPSDuncorr = True And butSwap3 = False And butPSDswap = False And butUncorrSwap = False) Then
```

```
' Execute only if both torque and raw target tensors have been created,  
' are of same length, and desired ACF previously created
```

```
'Set N and Nt to sequence lengths  
If (mTorqueSet = True And mTargRawSet = True) Then  
    N = UBound(mTorque.Val) - LBound(mTorque.Val) + 1  
    Nt = UBound(mTargRaw.Val) - LBound(mTargRaw.Val) + 1  
End If
```

```
'Execute only if mTorque and mTargRaw created, and of same length  
If (mTorqueSet = True And mTargRawSet = True And N = Nt) Then
```

```
'Verify that desired ACF has been created  
If mACFset = True Then
```

```
'Turn butInitiated just before calculation begins  
butInitiated.Value = True
```

```
'1. Set algorithm tensors to torque and target sequences
```

```
X = mTorque      'X is set to torque sequence  
Y = mTargRaw     'Y is sequence to be manipulated for target motion  
imin = LBound(Y.Val)  
imax = UBound(Y.Val)  
N = imax - imin + 1 'sequence length
```

```
'2. Randomize sequence that will be swapped; seems to quicken approach to desired ACF
```

```
For i = 1 To (5 * N)  
    j1 = imin + (imax - imin) * Rnd() 'selects number between imin and imax  
    j2 = imin + (imax - imin) * Rnd()  
    temp = Y.Val(j1)  
    Y.Val(j1) = Y.Val(j2)  
    Y.Val(j2) = temp  
Next i
```

```
'3. Calculate original ACF, PSD, CCF
```

```
maxlag = UBound(mACFdes.Val) 'maxlag set same as that in desired ACF  
minlag = -maxlag  
M = maxlag - minlag + 1      'number of elements in ACF and PSD  
ACForig = ACF(Y, minlag, maxlag)  
PSDorig = PSDfromACF(ACForig, None)  
CCForig = CCF(X, Y, minlag, maxlag)
```

```
'4. Create desired CCF, CCF = 0
```

```
CCFdes = CCForig 'sets information in tensor  
cmin = LBound(CCForig.Val)  
cmax = UBound(CCForig.Val)  
For i = cmin To cmax  
    CCFdes.Val(i) = 0 'want sequences to be uncorrelated, CCF(j)= 0
```

```

Next i
mCCFdes = CCFdes                                'module level tensor

'5. Stochastic interchange

' Set initial ACF and CCF to original
ACFy = ACForig
ssdACFy = SSD(ACFy, mACFdes)
CCFy = CCForig
ssdCCFy = SSD(CCFy, mCCFdes)
MSEorig = (1 / M) * ssdCCFy                        'M is length of ACF & CCF

' Set #swaps, array to track change in SSD
MaxSwaps = N * sldNumSwaps.Value                  'set max number of swaps
ReDim mSSDswapACF(MaxSwaps)                       'SSD of ACF at each swap
mSSDswapACF(0) = ssdACFy                          'initial SSD, at swap 0
ReDim mSSDswapCCF(MaxSwaps)
mSSDswapCCF(0) = ssdCCFy
mGoodSwaps = 0                                    'counts swaps that are effective

'Ensure bounds set for sequence being swapped
imin = LBound(Y.Val)
imax = UBound(Y.Val)

' Two element stochastic interchange with efficient update of ACF & CCF
For PossSwap = 1 To MaxSwaps

    'Randomly choose two unequal elements to be swapped
    j1 = j2
    Do While j1 = j2
        j1 = imin + (imax - imin) * Rnd()
        j2 = imin + (imax - imin) * Rnd()
    Loop

    'ACF after swap
    ACFswap = ACFAfterSwap(Y, ACFy, minlag, maxlag, j1, j2)

    'CCF after swap
    CCFswap = CCFAfterSwap(Y, X, CCFy, minlag, maxlag, j1, j2)

    'Calculate SSD for ACF and CCF after swap
    ssdACFswap = SSD(ACFswap, mACFdes)
    ssdCCFswap = SSD(CCFswap, mCCFdes)

    'Exchange elements if swap improves ACF and CCF
    ' If SSD for ACF and CCF are lower after terms swapped, then swap
    ' Y.val(j1) and Y.val(j2) and update ACFold and CCFold to new
    If (ssdCCFswap < ssdCCFy) And (ssdACFswap <= ssdACFy) Then
        temp = Y.Val(j1)
        Y.Val(j1) = Y.Val(j2)
        Y.Val(j2) = temp
    End If

```

```

ACFy = ACFswap
ssdACFy = ssdACFswap
CCFy = CCFswap
ssdCCFy = ssdCCFswap
mGoodSwaps = mGoodSwaps + 1
End If

```

```

'Set SSD for swap # PossSwap
mSSDswapACF(PossSwap) = ssdACFy
mSSDswapCCF(PossSwap) = ssdCCFy

```

Next PossSwap

'6. Verify outcome

```

'Explicitly calculate ACF, PSD, & CCF from swapped sequence
ACFcalc = ACF(Y, minlag, maxlag)
PSDcalc = PSDfromACF(ACFcalc, None)
CCFcalc = CCF(X, Y, minlag, maxlag)
MSEswap = (1 / M) * SSD(CCFcalc, mCCFdes)

```

```

'Compare updated ACF to ACF calculated for entire swapped array
"imin = LBound(ACFcalc.Val)
"imax = UBound(ACFcalc.Val)
"ACFdif = ACFcalc
"For i = imin To imax
"ACFdif.Val(i) = ACFcalc.Val(i) - ACFy.Val(i)
"Next i

```

```

'Compare updated CCF to CCF calculated for entire swapped array
"imin = LBound(CCFcalc.Val)
"imax = UBound(CCFcalc.Val)
"CCFdif = CCFcalc
"For i = imin To imax
"CCFdif.Val(i) = CCFcalc.Val(i) - CCFy.Val(i)
"Next i

```

'7. Total power of swapped sequence

```

'from ACF
TotPowACF = N * ACFcalc.Val(0)

```

```

'from PSD
SumElemsPSD = 0
M = UBound(PSDcalc.Val) - LBound(PSDcalc.Val) + 1
For i = LBound(PSDcalc.Val) To UBound(PSDcalc.Val)
SumElemsPSD = SumElemsPSD + PSDcalc.Val(i)
Next i
TotPowPSD = (N / M) * SumElemsPSD

```

'8. Set the swapped target tensor



```

    mTargSwap = Y
    mTargSwapSet = True

    'Turn butInitiated off after calculation
    butInitiated.Value = False

    'plots
    plot_A.PlotY CCFcalc.Val
    lab_A.Caption = "CCF after swap"
    txt_A.Text = "MSE 1: " + Format(MSEorig, "0.#000000") + " 2: " +
Format(MSEswap, "0.#000000")

    "plot_A.PlotY ACFdif.Val
    "lab_A.Caption = "ACF difference at end"

    "plot_B.PlotY CCFdif.Val
    "lab_B.Caption = "CCF difference at end"

    plot_B.PlotY ACFcalc.Val
    lab_B.Caption = "ACF mTargSwap"
    txt_B.Text = "Total power = " + Format(TotPowACF, "#.#0")

    plot_D.PlotY Y.Val
    lab_D.Caption = "mTargSwap: PSD, CCF"
    txt_D.Text = ""

    PSDdisp = DisplayTensor(PSDcalc)
    plot_E.PlotY PSDdisp.Val
    lab_E.Caption = "PSD mTargSwap"
    txt_E.Text = "Total power = " + Format(TotPowPSD, "#.#0")

Else 'ACF not previously created
    txtComment.Text = "Desired ACF not created."
    butPSDuncorr = False
End If

Else 'Torque and target sequences not previously created properly
    txtComment.Text = "Both tensors not created, or not same length."
    butPSDuncorr = False
End If

Else 'butPSDuncorr turned off or either butPSDswap or butUncorrSwap already on
    If (butPSDuncorr = False And butPSDswap = False And butUncorrSwap = False) Then
        mTargSwapSet = False
        'Turn off LPF and Integrated buttons
        butLPF = False
        butIntegrate = False
    End If
End If
End Sub

```

Private Sub butSwap3\_ValueChanged(ByVal Value As Boolean)

' Subroutine uses a three element stochastic swap to order the target sequence  
' to have a desired ACF and to be statistically uncorrelated from the torque sequence.  
' Cross-correlation is evaluated symmetrically, for equal numbers of neg and pos lags  
' Output is the tensor mTargSwap.  
' Developed for correlation measures with negative and positive lags

'Module variables

'Common experiment settings

'Private mExpLen As Integer           'experiment length in seconds  
'Private mExecInt As Double         'interval at which to execute thru sequence  
'Private mN As Long                 'number of terms in sequence  
'Private mSettingsSet As Boolean     'flag for experimental settings

'Torque sequence

'Private mTorque As Tensor  
'Public mTorqueSet As Boolean

'Target velocity tensors

'Private mTargRaw As Tensor           'raw sequence to be manipulated  
'Private mTargRawSet As Boolean       'flag for raw target sequence  
'Private mTargSwap As Tensor         'swapped for PSD and/or uncorr  
'Public mTargSwapSet As Boolean       'Public flag for swapped target sequence  
'Private mSSDswapACF() As Double     'Tracks SSD for ACF after each swap  
'Private mSSDswapCCF() As Double     'Tracks SSD for CCF after each swap  
'Private mGoodSwaps As Long           'Tracks number of successful swaps

'Correlation measures

'Private mACFdes As Tensor           'Desired ACF  
'Private mACFset As Boolean           'Flag that ACFdes is set  
'Private mCCFdes As Tensor           'Desired CCF

'Local variables

Dim X As Tensor       'Set to torque sequence, mTorque  
Dim Y As Tensor       'Set to target sequence which will be swapped, mTargRaw  
  
Dim Yt As Tensor      'temporary tensor that holds order after two elements swapped  
Dim Ya As Tensor      '3 elements swapped, order a  
Dim Yb As Tensor      '3 elements swapped, order b  
  
Dim imin As Long  
Dim imax As Long  
Dim i As Long  
Dim N As Long         'number elements in torque and target sequences  
Dim Nt As Long  
  
Dim maxlag As Long    'set to the length of mACFdes  
Dim minlag As Long  
Dim M As Long         'number elements in ACF and PSD

Dim ACForig As Tensor 'measures for the sequence before stochastic swap  
 Dim PSDorig As Tensor  
 Dim CCForig As Tensor

Dim CCFdes As Tensor 'desired CCF, set to 0 for uncorr sequences  
 Dim cmin As Long  
 Dim cmax As Long

Dim ACfy As Tensor 'current measures for tensor Y  
 Dim ssdACfy As Double  
 Dim PSDy As Tensor  
 Dim CCFy As Tensor  
 Dim ssdCCfy As Double

Dim ACft As Tensor  
 Dim ssdACft As Double  
 Dim CCFt As Tensor  
 Dim ssdCCft As Double

Dim ACfa As Tensor  
 Dim ssdACfa As Double  
 Dim CCFa As Tensor  
 Dim ssdCCfa As Double

Dim ACfb As Tensor  
 Dim ssdACfb As Double  
 Dim CCFb As Tensor  
 Dim ssdCCfb As Double

Dim ACfcalc As Tensor 'calculated directly from swapped sequence  
 Dim PSDcalc As Tensor  
 Dim CCFcalc As Tensor

Dim ACfdif As Tensor 'compare ACfcalc to final updated ACfOld  
 Dim CCFdif As Tensor

Dim PSDdisp As Tensor 'use with DisplayTensor( ) to plot 0 freq 0 position

Dim TotPowSeq As Double  
 Dim TotPowACF As Double  
 Dim SumElemsPSD As Double  
 Dim TotPowPSD As Double

Dim MSEorig As Double 'MSE of original CCF  
 Dim MSEswap As Double 'MSE of swapped CCF

Dim j1 As Long 'randomly selected indices to swap  
 Dim j2 As Long  
 Dim j3 As Long

Dim n1 As Long

```

Dim n2 As Long

Dim MaxSwaps As Long
Dim PossSwap As Long
Dim k As Long
Dim temp As Double

' Set indicators and tensor flags

'Ensure led indicator turned off
butInitiated.Value = False

'Set flags for the LPF and position (target) sequences to false.
'This should be done whether butPSDuncorr is turned on, or reset to off.
mTargFiltSet = False
mTargPosSet = False

' Initiate routine if butSwap3 turned on, and other swapping routines turned off
If (butSwap3 = True And butPSDuncorr = False And butPSDswap = False And butUncorrSwap =
False) Then

' Execute only if both torque and raw target tensors have been created,
' are of same length, and desired ACF previously created

' Set N and Nt to sequence lengths
If (mTorqueSet = True And mTargRawSet = True) Then
    N = UBound(mTorque.Val) - LBound(mTorque.Val) + 1
    Nt = UBound(mTargRaw.Val) - LBound(mTargRaw.Val) + 1
End If

'Execute only if mTorque and mTargRaw created, and of same length
If (mTorqueSet = True And mTargRawSet = True And N = Nt) Then

'Verify that desired ACF has been created
If mACFset = True Then

'Turn butInitiated just before calculation begins
butInitiated.Value = True

'1. Set algorithm tensors to torque and target sequences
X = mTorque      'X is set to torque sequence
Y = mTargRaw     'Y is sequence to be manipulated for target motion
imin = LBound(Y.Val)
imax = UBound(Y.Val)
N = imax - imin + 1 'sequence length

'2. Randomize sequence that is to be ordered;
' seems to quicken approach to desired ACF&CCF
For i = 1 To (5 * N)
    j1 = imin + (imax - imin) * Rnd() 'selects number between imin and imax
    j2 = imin + (imax - imin) * Rnd()

```

```

    temp = Y.Val(j1)
    Y.Val(j1) = Y.Val(j2)
    Y.Val(j2) = temp
Next i

'3. Explicitly calculate original ACF, PSD, CCF
maxlag = UBound(mACFdes.Val) 'maxlag set same as that in desired ACF
minlag = -maxlag
M = maxlag - minlag + 1 'number of elements in ACF and PSD
ACForig = ACF(Y, minlag, maxlag)
PSDorig = PSDfromACF(ACForig, None)
CCForig = CCF(X, Y, minlag, maxlag)

'4. Create desired CCF; CCF = 0 for uncorrelated sequences
CCFdes = CCForig 'sets tensor information
cmin = LBound(CCForig.Val)
cmax = UBound(CCForig.Val)
For i = cmin To cmax
    CCFdes.Val(i) = 0 'uncorrelated sequences, CCF(j)= 0
Next i
mCCFdes = CCFdes 'module level tensor

'5. Three element stochastic interchange
'Compare result of swap to Y, ACFy, ssdACFy, CCFy, ssdCCFy

'5a. Set up

'Set initial ACF, CCF values to original
ACFy = ACForig
ssdACFy = SSD(ACForig, mACFdes)
CCFy = CCForig
ssdCCFy = SSD(CCForig, mCCFdes)
MSEorig = (1 / M) * ssdCCFy 'M is length of ACF & CCF

'Set #swaps, array to track change in SSD
MaxSwaps = N * sldNumSwaps.Value 'set max number of swaps
ReDim mSSDswapACF(MaxSwaps) 'SSD of ACF at each swap
mSSDswapACF(0) = ssdACFy 'initial SSD, at swap 0
ReDim mSSDswapCCF(MaxSwaps)
mSSDswapCCF(0) = ssdCCFy
mGoodSwaps = 0 'counts swaps that are effective

' Ensure bounds set to Y
imin = LBound(Y.Val)
imax = UBound(Y.Val)

'5b. Stochastic swap
For PossSwap = 1 To MaxSwaps

    '1. Randomly choose three elements to be swapped - ensure they are different
    j1 = 0

```

```

j2 = 0
j3 = 0
Do While j1 = j2 Or j2 = j3 Or j1 = j3
  j1 = imin + (imax - imin) * Rnd()
  j2 = imin + (imax - imin) * Rnd()
  j3 = imin + (imax - imin) * Rnd()
Loop

```

'2. Create temporary sequence by swapping j1 and j2

```

'swap elements of Y to create Yt
Yt = Y
temp = Yt.Val(j1)
Yt.Val(j1) = Yt.Val(j2)
Yt.Val(j2) = temp

'update ACF and CCF from original Y values
ACFt = ACFAfterSwap(Y, ACFy, minlag, maxlag, j1, j2)
CCFt = CCFAfterSwap(Y, X, CCFy, minlag, maxlag, j1, j2)
ssdACFt = SSD(ACFt, mACFdes)
ssdCCFt = SSD(CCFt, mCCFdes)

```

'3. Two options for the three element swap

'3a. Option a: swap elements j1 and j3

```

'swap elements of Yt to create Ya
Ya = Yt
temp = Ya.Val(j1)
Ya.Val(j1) = Ya.Val(j3)
Ya.Val(j3) = temp

'update ACF and CCF from Yt measures
ACFa = ACFAfterSwap(Yt, ACFt, minlag, maxlag, j1, j3)
ssdACFa = SSD(ACFa, mACFdes)
CCFa = CCFAfterSwap(Yt, X, CCFt, minlag, maxlag, j1, j3)
ssdCCFa = SSD(CCFa, mCCFdes)

```

'3b. Option b: swap elements j2 and j3

```

'swap elements of Yt to create Yb
Yb = Yt
temp = Yb.Val(j2)
Yb.Val(j2) = Yb.Val(j3)
Yb.Val(j3) = temp

'update ACF and CCF from Yt measures
ACFb = ACFAfterSwap(Yt, ACFt, minlag, maxlag, j2, j3)
ssdACFb = SSD(ACFb, mACFdes)
CCFb = CCFAfterSwap(Yt, X, CCFt, minlag, maxlag, j2, j3)

```

ssdCCFb = SSD(CCFb, mCCFdes)

- '4. Select which order should be used; a, b, or y
- ' Change sequence order when the swap improved the CCF and the ACF is improved
- ' or remains the same. Of the 2 orders, choose the order with the lower CCF

'Select order a

If (ssdCCFa < ssdCCFb) And (ssdCCFa < ssdCCFy) And (ssdACFa <= ssdACFy)

Then

Y = Ya  
ACFy = ACFa  
ssdACFy = ssdACFa  
CCFy = CCFa  
ssdCCFy = ssdCCFa

'Select order b

ElseIf (ssdCCFb < ssdCCFa) And (ssdCCFb < ssdCCFy) And (ssdACFb <=

ssdACFy) Then

Y = Yb  
ACFy = ACFb  
ssdACFy = ssdACFb  
CCFy = CCFb  
ssdCCFy = ssdCCFb

'Decide between a and b if CCF equal and less than old

ElseIf (ssdCCFb = ssdCCFa) And (ssdCCFa < ssdCCFy) Then

If (ssdACFa < ssdACFb) And (ssdACFa < ssdACFy) Then

Y = Ya  
ACFy = ACFa  
ssdACFy = ssdACFa  
CCFy = CCFa  
ssdCCFy = ssdCCFa

ElseIf (ssdACFb < ssdACFa) And (ssdACFb < ssdACFy) Then

Y = Yb  
ACFy = ACFb  
ssdACFy = ssdACFb  
CCFy = CCFb  
ssdCCFy = ssdCCFb

End If

Else 'No change

End If

'5c. Set SSD for swap # PossSwap

mSSDswapACF(PossSwap) = ssdACFy

mSSDswapCCF(PossSwap) = ssdCCFy

Next PossSwap

'6. Verify outcome

'Explicitly calculate ACF, PSD, & CCF from swapped sequence

```
ACFcalc = ACF(Y, minlag, maxlag)
PSDcalc = PSDfromACF(ACFcalc, None)
CCFcalc = CCF(X, Y, minlag, maxlag)
MSEswap = (1 / M) * SSD(CCFcalc, mCCFdes)
```

'Compare updated ACF to ACF calculated for entire swapped array

```
"imin = LBound(ACFcalc.Val)
"imax = UBound(ACFcalc.Val)
"ACFdif = ACFcalc
"For i = imin To imax
" ACFdif.Val(i) = ACFcalc.Val(i) - ACFy.Val(i)
"Next i
```

'Compare updated CCF to CCF calculated for entire swapped array

```
"imin = LBound(CCFcalc.Val)
"imax = UBound(CCFcalc.Val)
"CCFdif = CCFcalc
"For i = imin To imax
" CCFdif.Val(i) = CCFcalc.Val(i) - CCFy.Val(i)
"Next i
```

'7. Total power of swapped sequence

'from ACF

```
TotPowACF = N * ACFcalc.Val(0)
```

'from PSD

```
SumElemsPSD = 0
```

```
M = UBound(PSDcalc.Val) - LBound(PSDcalc.Val) + 1
```

```
For i = LBound(PSDcalc.Val) To UBound(PSDcalc.Val)
```

```
SumElemsPSD = SumElemsPSD + PSDcalc.Val(i)
```

```
Next i
```

```
TotPowPSD = (N / M) * SumElemsPSD
```

'8. Set the swapped target velocity to Y

```
mTargSwap = Y
```

```
mTargSwapSet = True
```

'Turn butInitiated off after calculation

```
butInitiated.Value = False
```

'plots

```
plot_A.PlotY CCFcalc.Val
```

```
lab_A.Caption = "CCF after swap"
```



```
txt_A.Text = "MSE 1: " + Format(MSEorig, "0.#000000") + " 2: " +  
Format(MSEswap, "0.#000000")
```

```
"plot_A.PlotY ACFdif.Val  
"lab_A.Caption = "ACF difference at end"
```

```
"plot_B.PlotY CCFdif.Val  
"lab_B.Caption = "CCF difference at end"
```

```
plot_B.PlotY ACFcalc.Val  
lab_B.Caption = "ACF mTargSwap"  
txt_B.Text = "Total power = " + Format(TotPowACF, "#.#0")
```

```
plot_D.PlotY Y.Val  
lab_D.Caption = "mTargSwap: PSD, CCF"  
txt_D.Text = ""
```

```
PSDdisp = DisplayTensor(PSDcalc)  
plot_E.PlotY PSDdisp.Val  
lab_E.Caption = "PSD mTargSwap"  
txt_E.Text = "Total power = " + Format(TotPowPSD, "#.#0")
```

```
Else 'ACF not previously created  
txtComment.Text = "Desired ACF not created."  
butPSDuncorr = False  
End If
```

```
Else 'Torque and target sequences not previously created properly  
txtComment.Text = "Both tensors not created, or not same length."  
butPSDuncorr = False  
End If
```

```
Else 'butPSDuncorr turned off or either butPSDswap or butUncorrSwap already on  
If (butPSDuncorr = False And butPSDswap = False And butUncorrSwap = False) Then  
mTargSwapSet = False  
'Turn off LPF and Integrated buttons  
butLPF = False  
butIntegrate = False  
End If  
End If  
End Sub
```

```

Private Sub cmdCCF_Click()
'Calculates the neg and pos lags of cross correlation

'Module variables
'Common experiment settings
  'Private mExpLen As Integer      'experiment length in seconds
  'Private mExecInt As Double     'interval at which to execute thru sequence
  'Private mN As Long             'number of terms in sequence

'Torque sequence
  'Private mTorque As Tensor
  'Private mTorqueSet As Boolean

'Raw Target sequence
  'Target velocity tensors
  'Private mTargRaw As Tensor     'raw sequence to be manipulated
  'Private mTargRawSet As Boolean 'flag for raw target sequence

'Local variables
  Dim Y As Tensor
  Dim Yset As Boolean
  Dim SeqName As String
  Dim N As Long
  Dim Ny As Long
  Dim percentLen As Double
  Dim maxlag As Long
  Dim minlag As Long
  Dim Cty As Tensor
  Dim Uncorr As Tensor
  Dim j As Long
  Dim MSE As Double
  Dim AmpSpec As Tensor
  Dim PhaseSpec As Tensor
  Dim PlotSpec As Tensor

' Decide which sequence is to be cross correlated with the Torque
' Even allow the torque sequence

  'Begin with sequence flag set false, set true once decide which seq to use
  Yset = False

  If (optTor = True And mTorqueSet = True) Then
    Y = mTorque
    SeqName = "torque"
    Yset = True
  ElseIf (optTargRaw = True And mTargRawSet = True) Then
    Y = mTargRaw
    SeqName = "Raw target"
    Yset = True
  ElseIf (optTargSwap = True And mTargSwapSet = True) Then
    Y = mTargSwap

```

```

    SeqName = "Swap target"
    Yset = True
ElseIf (optTargFilt = True And mTargFiltSet = True) Then
    Y = mTargFilt
    SeqName = "LPF target"
    Yset = True
ElseIf (optTargInt = True And mTargPosSet = True) Then
    Y = mTargPos
    SeqName = "Int target"
    Yset = True
Else
    txtComment.Text = "Specified sequence not created."
End If

'Check lengths of sequences
If (mTorqueSet = True And Yset = True) Then
    N = UBound(mTorque.Val) - LBound(mTorque.Val) + 1
    Ny = UBound(Y.Val) - LBound(Y.Val) + 1
End If

'Execute only if both tensors have been created and same length
If (mTorqueSet = True And Yset = True And N = Ny) Then

    'lags
    percentLen = sldMaxLag.Value
    maxlag = Int(percentLen * (N - 1))
    minlag = -maxlag
    txtMaxLag.Text = Format(percentLen, "#.##") + "*N, " + Format(maxlag, "#")

    'CCF
    Cty = CCF(mTorque, Y, minlag, maxlag)

    'Create zero tensor to represent uncorr ccf
    Uncorr = Cty
    For j = minlag To maxlag
        Uncorr.Val(j) = 0
    Next j

    'Mean square error between Cxy and uncorr
    MSE = (1 / (maxlag - minlag + 1)) * SSD(Cty, Uncorr)

    'Cross amplitude and phase spectra
    AmpSpec = CrossAmpSpectrum(mTorque, Y, maxlag)
    PhaseSpec = PhaseSpectrum(mTorque, Y, maxlag)

    'plots
    plot_A.PlotY Y.Val
    lab_A.Caption = SeqName
    txt_A.Text = "Crossed with torque"

    plot_B.PlotY Cty.Val

```

```
lab_B.Caption = "Tor-" + SeqName + " CCF"
txt_B.Text = "MSE CCF = " + Format(MSE, "0.###00000")
```

```
PlotSpec = DisplayTensor(PhaseSpec)
plot_D.PlotY PlotSpec.Val
lab_D.Caption = "Phase (deg)"
txt_D.Text = "Tor-" + SeqName
```

```
PlotSpec = DisplayTensor(AmpSpec)
plot_E.PlotY PlotSpec.Val
lab_E.Caption = "Cross-amp (disp)"
txt_E.Text = "Tor-" + SeqName
```

```
'clear coment txt
txtComment.Text = ""
```

```
Else: txtComment.Text = "Both sequences not created."
```

```
End If
End Sub
```

```
Private Sub cmdPSD_Click()
```

```
' ACF and PSD of mTorque and mTargRaw sequences
' Uses Plexus function ACF() to calculate negative and positive lags,
' and Plexis function
```

```
'Module variables
```

```
'Common experiment settings
'Private mExpLen As Integer          'experiment length in seconds
'Private mExecInt As Double         'interval at which to execute thru sequence
'Private mN As Long                 'number of terms in sequence
```

```
'Torque sequence
```

```
'Private mTorque As Tensor
'Private mTorqueSet As Boolean
```

```
'Target velocity tensors
```

```
'Private mTargRaw As Tensor         'raw sequence to be manipulated
'Private mTargRawSet As Boolean     'flag for raw target sequence
'Private mTargSwap As Tensor       'swapped for PSD and/or uncorr
'Public mTargSwapSet As Boolean     'Public flag for swapped target sequence
'Private mTargFilt As Tensor       'LPF filtered target sequence
'Public mTargFiltSet As Boolean     'Public flag for filtered target sequence
'Private mTargPos As Tensor        'Position sequence
'Public mTargPosSet As Boolean      'Public flag for position sequence
```

```
'Local variables
```

```
Dim Target As Tensor              'Assiged to hold desired target sequence
Dim TargetSet As Boolean          'Flag that tensor was set
```

```

Dim N As Long                'length of sequences
Dim percentLen As Double    'percent of length for lag
Dim maxlag As Long          'max lag for ACF
Dim minlag As Long
Dim SeqName As String

Dim Ctt As Tensor          'pos and neg lags
Dim TotPowACF As Double
Dim PSD As Tensor
Dim lenPSD As Long
Dim i As Long

Dim SumElems As Double
Dim TotPowPSD As Double
Dim PSDdisp As Tensor

' Execute only if flags for sequences are set true
' else report tensors not created.

'optTor selected
If optTor = True Then

    'PSD of Torque
    If mTorqueSet = True Then

        'length of mTorque
        N = UBound(mTorque.Val) - LBound(mTorque.Val) + 1

        'set min and max lag
        percentLen = GenSeqs.sldMaxLag.Value
        txtMaxLag.Text = Format(percentLen, "#.#0") + "*N, " + Format(maxlag, "#")
        maxlag = Int(percentLen * (N - 1))
        minlag = -maxlag

        'Torque ACF
        Ctt = ACF(mTorque, minlag, maxlag)
        'total power from ACF
        TotPowACF = N * Ctt.Val(0)

        'Torque PSD
        PSD = PSDfromACF(Ctt, None)
        lenPSD = UBound(PSD.Val) - LBound(PSD.Val) + 1
        SumElems = 0
        For i = LBound(PSD.Val) To UBound(PSD.Val)
            SumElems = SumElems + PSD.Val(i)
        Next i
        TotPowPSD = (N / lenPSD) * SumElems

        'Torque plots
        plot_A.PlotY mTorque.Val
        lab_A.Caption = "Torque"

```

```

txt_A.Text = ""

plot_B.PlotY Ctt.Val
lab_B.Caption = "ACF(+/- lags): Torque"
txt_B.Text = "Tot pow (ACF) = " + Format(TotPowACF, "#.#0")

plot_D.PlotY 0
lab_D.Caption = ""
txt_D.Text = ""

PSDdisp = DisplayTensor(PSD)
plot_E.PlotY PSDdisp.Val
lab_E.Caption = "PSD(dis): Torque"
txt_E.Text = "Tot pow (PSD) = " + Format(TotPowPSD, "#.#0")

Else: txtComment.Text = "Torque sequence not created."

End If

'Option selected for one of the target sequences
Else

'PSD of mTargRaw

'Decide which target sequence to use
'Begin with target flag set false, set true once decide which seq to use
TargetSet = False

If (optTargRaw = True And mTargRawSet = True) Then
    Target = mTargRaw
    SeqName = "Raw target"
    TargetSet = True
ElseIf (optTargSwap = True And mTargSwapSet = True) Then
    Target = mTargSwap
    SeqName = "Swap target"
    TargetSet = True
ElseIf (optTargFilt = True And mTargFiltSet = True) Then
    Target = mTargFilt
    SeqName = "LPF target"
    TargetSet = True
ElseIf (optTargInt = True And mTargPosSet = True) Then
    Target = mTargPos
    SeqName = "Int target"
    TargetSet = True
Else
    TargetSet = False
    txtComment.Text = "Specified target sequence not created."
End If

If TargetSet = True Then

```

```

'length of mTargRaw
  N = UBound(Target.Val) - LBound(Target.Val) + 1

'set min and max lag
  percentLen = GenSeqs.sldMaxLag.Value
  txtMaxLag.Text = Format(percentLen, "#.#0") + "*N, " + Format(maxlag, "#")
  maxlag = Int(percentLen * (N - 1))
  minlag = -maxlag

'ACF
  Ctt = ACF(Target, minlag, maxlag)
'total power from ACF
  TotPowACF = N * Ctt.Val(0)

'PSD
  PSD = PSDfromACF(Ctt, None)
  lenPSD = UBound(PSD.Val) - LBound(PSD.Val) + 1
  SumElems = 0
  For i = LBound(PSD.Val) To UBound(PSD.Val)
    SumElems = SumElems + PSD.Val(i)
  Next i
  TotPowPSD = (N / lenPSD) * SumElems

'clear txtComment
  txtComment.Text = ""
'plots
  plot_A.PlotY Target.Val
  lab_A.Caption = SeqName
  txt_A.Text = ""

  plot_B.PlotY Ctt.Val
  lab_B.Caption = "ACF(+/- lags): " + SeqName
  txt_B.Text = "Tot pow (ACF) = " + Format(TotPowACF, "#.#0")

  plot_D.PlotY 0
  lab_D.Caption = ""
  txt_D.Text = ""

  PSDdisp = DisplayTensor(PSD)
  plot_E.PlotY PSDdisp.Val
  lab_E.Caption = "PSD(disp): " + SeqName
  txt_E.Text = "Tot pow (PSD) = " + Format(TotPowPSD, "#.#0")

End If

End If
End Sub

```

```

Private Sub sldLPF_PointerValueChanged(ByVal Pointer As Long, Value As Variant)
'LPF filter frequency
    txtLPF.Text = Format(sldLPF.Value, "0.0") + " Hz"

End Sub

Private Sub sldMaxLag_PointerValueChanged(ByVal Pointer As Long, Value As Variant)
'Calculates max number of ACF lags as a % of sequence length

'Module variables
'Common experiment settings
    Private mExpLen As Integer        'experiment length in seconds
    Private mExecInt As Double       'interval at which to execute thru sequence
    Private mN As Long               'number of terms in sequence
    Private mSettingsSet As Boolean   'flag for experimental settings

'Local variables
    Dim i As Long
    Dim N As Long
    Dim maxlag As Long
    Dim percent As Double

'Execute only if mHvel created
    If mSettingsSet = True Then
        N = mN
        percent = sldMaxLag.Value
        maxlag = Int(percent * N)
        txtMaxLag.Text = Format(percent, "#.#0") + "*N, " + Format(maxlag, "#")
        txtComment.Text = ""

    Else
        txtComment.Text = "Experiment settings must be set."
        sldMaxLag.Value = 1

    End If
End Sub

Private Sub sldNumSwaps_PointerValueChanged(ByVal Pointer As Long, Value As Variant)
'Update text window
    txtNumSwaps.Text = Format(sldNumSwaps.Value, "0")
End Sub

Private Sub sldSeedlength_PointerValueChanged(ByVal Pointer As Long, Value As Variant)
'Local variables
    Dim P As Long
    Dim PRBSint As Double
    Dim PRBSdur As Double

```



```
'Turn butGenTorque off, ensures values are not changed after sequence set
  butGenTorque.Value = False
  mTorqueSet = False
```

```
'Can only calculate duration once experimental settings are set
If mSettingsSet = True Then
  P = (2 ^ sldSeedlength.Value) - 1      'PRBS length
  PRBSint = sldIntRatio.Value * mExecInt 'effective PRBS interval
  PRBSdur = P * PRBSint                  'duration of PRBS
  txtPRBSdur.Text = Format(PRBSdur, "0.0") + " s"
Else

  End If
End Sub
```

## User defined functions

The user defined functions were programmed using the Plexus programming format developed by Prof. Ian W. Hunter.

```
.....
'           Plexus functions added by J. L. Tangorra
'           Last update Apr 2002
'
'.....
'
' GenComplexArray
'
' GenPRBS
' GenGWN
' GenDualGWN
' GenBinary
' GenUniform
' Gen2Uniform
'
' Convolve2          Conducts convolution h*x - causal or noncausal
' dtConvolve        Convolves h(n)*x(n). Scales by 1/dT.
'
' ReDFT              Calculates the cosine transform of a real valued tensor
'                   A of DFT = A - jB
' ImDFT              Calculates the sine transform of a real valued tensor
'                   B of DFT = A - jB
' InvDFT             Calculates the inverse DFT from the real and imaginary
'                   components of the DFT = A - jB
' DisplayTensor      Rearranges elements of tensor so elements with negative
'                   indices are moved after the positive indices for graphing.
' AutoCorrelation
' CrossCorrelation
' ACF                Calculates (-/+ )lag values of auto-correlation
' CCF
' TwoSidedPSD        PSD calculated from (+) lags of ACF
' PSDfromACF         -0.5 to 0.5 cps PSD calculated from (-/+) lags of ACF
' ACFfromPSD         (-/+) lags of ACF from PSD -0.5 to 0.5 Hz
'
' ReCrossPSD         Calculates the real part of the cross spectral density
'                   Sometimes called the co-spectral estimator.
' ImCrossPSD         Calculates the imaginary part of the cross spectral density
'                   Sometimes called the quadrature estimator
'
' CrossAmpSpectrum   SqRoot( ReCrossPSD^2 + ImCrossPSD^2)
' PhaseSpectrum      ArcTan( - ImCrossPSD/ReCrossPSD)
'
' SSD                Calculates sum of square difference between two tensors
```

```

'
' IntegrateTR          Integrates the elements of a tensor using a simple
'                    Trapezoidal Rule
'
' ScaleTensor         Scales tensor by a constant
' ACFAfterSwap
' CCFAfterSwap
'
' LPF                 Convolves sequence with non-causal LPF irf found from invDFT of H(jw)
' LPFwind             Windowed LPF
' ModBessI0()         Evaluates zeroth-order modified Bessel function of first kind
' LPFkaiser           Creates LPF from sinc with a Kaiser window. Requires ModBessI0()
'

```

```

.....

```

```

Public Function GenComplexArray(Y() As Double) As Complex()
'

```

```

' Converts an array of real, double, values into an array of type Complex.
' Values in complex arrays have real and imag parts, x(i).re and x(i).im.
' The motivation is to have a function to convert real arrays in complex() for
' the FFT functions

```

```

Dim imin As Double
Dim imax As Double
Dim i As Double
Dim X() As Complex

```

```

imin = LBound(Y)
imax = UBound(Y)

```

```

ReDim X(imin To imax)

```

```

For i = imin To imax
    X(i).re = Y(i)
    X(i).im = 0
Next i

```

```

GenComplexArray = X

```

```

End Function

```

```

Public Function GenPRBS( _
    Optional SeedLength As Long = 12, _
    Optional imin As Long = 0, _
    Optional imax As Long = 1000, _
    Optional dT As Double = 0.05 _
) As Tensor

```

```

'Generates a Plexus Tensor Vector
' prbs(i) i=imin, imin+1,...,imax

```

```
' imin (default = 0)
' imax (default = 1000)
' dt domain increment (default 0.05)
'
```

'The PRBS repeats after  $(2^{\text{SeedLength}})-1$  elements

'This function was created by following the structure used  
'by Ian Hunter for the Plexus functions.

'It was written by James Tangorra

'Last update 03 Aug 2001

'Local variable declarations

Dim prbs() As Double

Dim tap1 As Integer

Dim tap2 As Integer

Dim i As Long

Dim N As Long

Dim SeedSum As Long

'Select XOR tap values based on SeedLength

Select Case SeedLength

Case 6

tap1 = 5

tap2 = 6

Case 7

tap1 = 6

tap2 = 7

Case 8

tap1 = 6

tap2 = 8

Case 9

tap1 = 5

tap2 = 9

Case 10

tap1 = 7

tap2 = 10

Case 11

tap1 = 9

tap2 = 11

Case 12

tap1 = 11

tap2 = 12

Case 13

tap1 = 12

tap2 = 13

Case 14

tap1 = 13

tap2 = 14

Case 15

tap1 = 14

```

        tap2 = 15
    Case Else
        'default to taps for SeedLength of 12, SeedLength default is 12
        tap1 = 11
        tap2 = 12
    End Select

'size prbs() and initialize elements to zero
ReDim prbs(imin To imax)
For i = imin To imax
    prbs(i) = 0
Next i

'Generate random seeds of 0 or 1, place at rightmost end of array
'If all seed values equal zero, change last seed value to 1
SeedSum = 0
For i = 0 To (SeedLength - 1)
    prbs(imax - i) = Int(Rnd + 0.5)
    SeedSum = SeedSum + prbs(imax - i)
Next i
If SeedSum = 0 Then prbs(imax) = 1

'Generate PRBS sequence
'place XOR value to left of shift register
'begin at value n = imax - SeedLength
'end at imin
For i = (imax - SeedLength) To imin Step -1
    prbs(i) = prbs(i + tap1) Xor prbs(i + tap2)
Next i

ReDim GenPRBS.Rname(1 To 1)
ReDim GenPRBS.Runit(1 To 1)
ReDim GenPRBS.Dname(1 To 1)
ReDim GenPRBS.Dunit(1 To 1)
ReDim GenPRBS.Inc(1 To 1)
GenPRBS.Note = "Pseudo-random binary sequence"
GenPRBS.Class = "Vector"
GenPRBS.Type = "Signal"
GenPRBS.Rname(1) = "Amplitude"
GenPRBS.Dname(1) = "Time"
GenPRBS.Dunit(1) = "s"
GenPRBS.Inc(1) = dT
GenPRBS.Val = prbs

End Function

```

```

Public Function GenGWN( _
    Optional imin As Long = 0, _
    Optional imax As Long = 1000, _
    Optional dT As Double = 0.05, _
    Optional Mean As Double = 0#, _
    Optional SD As Double = 1# _
) As Tensor

'Generates a Plexus Tensor Vector of N Gaussian white noise samples,
' x(i) i=imin, imin+1,...,imax
' by adding 20 uniformly distributed (0 to 1) random numbers)
'GenGaussNoise function sets x to the specified
' imin (default 0),
' imax (default 1000),
' dt domain increment (default 0.5),
' mean (default 0.0) and,
' standard deviation (default 1.0)
,

'This function was modified from GenGaussNoise.
'It was written by Ian Hunter Jan 14 1999

Dim X() As Double
ReDim X(imin To imax)
Dim i As Long
Dim j As Long
Dim gwn As Double
Dim s1 As Double
Dim s2 As Double

s1 = 0#

For i = imin To imax 'Adds 20 elements, each between 0 and 1
    gwn = -10# 'Puts sum from -10 to +10
    For j = 1 To 20
        gwn = gwn + Rnd()
    Next j
    X(i) = gwn
    s1 = s1 + gwn
Next i
s1 = s1 / (imax - imin + 1) 'mean of sequence

s2 = 0#
For i = imin To imax
    X(i) = X(i) - s1 'Makes sequence zero mean
    s2 = s2 + X(i) ^ 2
Next i
s2 = Sqr(s2 / (imax - imin)) 'Standard deviation of zero mean sequence

For i = imin To imax
    X(i) = Mean + SD * X(i) / s2 'Give sequence desired mean and std deviation
Next i

```

```

ReDim GenGWN.Rname(1 To 1)
ReDim GenGWN.Runit(1 To 1)
ReDim GenGWN.Dname(1 To 1)
ReDim GenGWN.Dunit(1 To 1)
ReDim GenGWN.Inc(1 To 1)

```

```

GenGWN.Note = "Gaussian white noise"
GenGWN.Class = "Vector"
GenGWN.Type = "Signal"
GenGWN.Rname(1) = "Amplitude"
GenGWN.Dname(1) = "Time"
GenGWN.Dunit(1) = "s"
GenGWN.Inc(1) = dT
GenGWN.Val = X
End Function

```

```

Public Function GenDualGWN( _
    Optional imin As Long = 0, _
    Optional imax As Long = 1000, _
    Optional dT As Double = 0.05, _
    Optional Mean As Double = 1#, _
    Optional SD As Double = 1# _
) As Tensor

```

```

' Generates a Tensor from two Gaussian distributions.
' The Gaussian distributions are made by adding 20 uniformly
' distributed random numbers, between 0 and 1
'
' This function was modified from GenGWN.

```

```

Dim X() As Double
ReDim X(imin To imax)
Dim i As Long
Dim j As Long
Dim N As Long
Dim jmin As Long
Dim jmax As Long
Dim gwn As Double
Dim s1 As Double
Dim s2 As Double

```

```

'Split the elements in half
N = imax - imin + 1

```

```

'First half, give the negative mean
jmin = imin
jmax = imin + (N / 2) - 1

```

```

s1 = 0#
For i = jmin To jmax
  gwn = -10#
  For j = 1 To 20
    gwn = gwn + Rnd()
  Next j
  X(i) = gwn
  s1 = s1 + gwn          'summing elements
Next i
s1 = s1 / (jmax - jmin + 1)  'mean of sequence

s2 = 0#
For i = jmin To jmax
  X(i) = X(i) - s1      'give elements zero mean
  s2 = s2 + X(i) ^ 2   'summation for variance of zero mean sequence
Next i
s2 = Sqr(s2 / (jmax - jmin)) 'standard deviation

For i = jmin To jmax
  X(i) = -Mean + SD * X(i) / s2
Next i

'Second half, give the positive mean
jmin = imin + (N / 2)
jmax = imax

s1 = 0#
For i = jmin To jmax
  gwn = -10#
  For j = 1 To 20
    gwn = gwn + Rnd()
  Next j
  X(i) = gwn
  s1 = s1 + gwn          'summing elements
Next i
s1 = s1 / (jmax - jmin + 1)  'mean of sequence

s2 = 0#
For i = jmin To jmax
  X(i) = X(i) - s1      'give elements zero mean
  s2 = s2 + X(i) ^ 2   'summation for variance of zero mean sequence
Next i
s2 = Sqr(s2 / (jmax - jmin)) 'standard deviation

For i = jmin To jmax
  X(i) = Mean + SD * X(i) / s2
Next i

ReDim GenDualGWN.Rname(1 To 1)
ReDim GenDualGWN.Runit(1 To 1)

```



```

ReDim GenDualGWN.Dname(1 To 1)
ReDim GenDualGWN.Dunit(1 To 1)
ReDim GenDualGWN.Inc(1 To 1)
GenDualGWN.Note = "Gaussian white noise"
GenDualGWN.Class = "Vector"
GenDualGWN.Type = "Signal"
GenDualGWN.Rname(1) = "Amplitude"
GenDualGWN.Dname(1) = "Time"
GenDualGWN.Dunit(1) = "s"
GenDualGWN.Inc(1) = dT
GenDualGWN.Val = X
End Function

```

```

Public Function GenBinary( _
    Optional imin As Long = 0, _
    Optional imax As Long = 1000, _
    Optional dT As Double = 0.05, _
    Optional Amp1 As Double = 1#, _
    Optional Amp2 As Double = -1# _
) As Tensor

```

'Generates a binary sequence with amplitudes amp1 and amp2.

```

' imin (default 0),
' imax (default 1000),
' dt domain increment (default 0.05),
' amp1 (default 1.0) and,
' amp2 (default -1.0)

```

```

Dim X() As Double
ReDim X(imin To imax)
Dim j As Long
Dim N As Long
Dim jmin As Long
Dim jmax As Long

```

```

'Length of the sequence
N = imax - imin + 1

```

```

'Give first half an amplitude of amp1
jmin = imin
jmax = imin + (N / 2) - 1
For j = jmin To jmax
    X(j) = Amp1
Next j

```

```

'Give second half an amplitude of amp2
jmin = imin + (N / 2)

```

```

jmax = imax
For j = jmin To jmax
    X(j) = Amp2
Next j

```

```

ReDim GenBinary.Rname(1 To 1)
ReDim GenBinary.Runit(1 To 1)
ReDim GenBinary.Dname(1 To 1)
ReDim GenBinary.Dunit(1 To 1)
ReDim GenBinary.Inc(1 To 1)

```

```

GenBinary.Note = "Binary sequence"
GenBinary.Class = "Vector"
GenBinary.Type = "Signal"
GenBinary.Rname(1) = "Amplitude"
GenBinary.Dname(1) = "Time"
GenBinary.Dunit(1) = "s"
GenBinary.Inc(1) = dT
GenBinary.Val = X
End Function

```

```

Public Function GenUniform( _
    Optional imin As Long = 0#, _
    Optional imax As Long = 1000#, _
    Optional dT As Double = 0.05, _
    Optional ymin As Double = 0#, _
    Optional ymax As Double = 1# _
) As Tensor

```

'Program created by modifying GenRamp( ) written by I.W. Hunter

```

Dim i As Long
Dim N As Long
Dim yinc As Double

```

```

'Number of elements
N = imax - imin + 1

```

```

ReDim GenUniform.Val(imin To imax)
yinc = (ymax - ymin) / (N - 1#)
For i = imin To imax
    GenUniform.Val(i) = ymin + i * yinc
Next i

```

```

ReDim GenUniform.Rname(1 To 1)
ReDim GenUniform.Runit(1 To 1)
ReDim GenUniform.Dname(1 To 1)
ReDim GenUniform.Dunit(1 To 1)

```

```

ReDim GenUniform.Inc(1 To 1)
GenUniform.Note = "Ramp"
GenUniform.Rname(1) = "Amplitude"
GenUniform.Dname(1) = "Time"
GenUniform.Dunit(1) = "s"
GenUniform.Inc(1) = dT

```

End Function

```

Public Function Gen2Uniform( _
    Optional imin As Long = 0#, _
    Optional imax As Long = 1000#, _
    Optional dT As Double = 0.05, _
    Optional y1min As Double = -2#, _
    Optional y1max As Double = -1#, _
    Optional y2min As Double = 1#, _
    Optional y2max As Double = 2# _
) As Tensor

```

'Program created by modifying GenRamp( ) written by I.W. Hunter

```

Dim N As Long
Dim Rec As Long
Dim j As Long
Dim jmin As Long
Dim jmax As Long
Dim yinc As Double

```

```

ReDim Gen2Uniform.Val(imin To imax)

```

```

'Number of elements
N = imax - imin + 1

```

```

'First record (half sequence)
jmin = imin
jmax = imin + (N / 2) - 1

```

```

'Record length
Rec = jmax - jmin + 1

```

```

'Increment each value
yinc = (y1max - y1min) / (Rec - 1#)
For j = jmin To jmax
    Gen2Uniform.Val(j) = y1min + (j - jmin) * yinc
Next j

```

```

'Second record (half sequence)

```

```

jmin = imin + (N / 2)
jmax = imax

'Record length
Rec = jmax - jmin + 1

'Increment each value
yinc = (y2max - y2min) / (Rec - 1#)
For j = jmin To jmax
    Gen2Uniform.Val(j) = y2min + (j - jmin) * yinc
Next j

ReDim Gen2Uniform.Rname(1 To 1)
ReDim Gen2Uniform.Runit(1 To 1)
ReDim Gen2Uniform.Dname(1 To 1)
ReDim Gen2Uniform.Dunit(1 To 1)
ReDim Gen2Uniform.Inc(1 To 1)

Gen2Uniform.Note = "Two ramps"
Gen2Uniform.Rname(1) = "Amplitude"
Gen2Uniform.Dname(1) = "Time"
Gen2Uniform.Dunit(1) = "s"
Gen2Uniform.Inc(1) = dT
End Function

Public Function Convolve2( _
    h As Tensor, _
    X As Tensor _
) As Tensor

' Computes the convolution between IRF h and sequence x
' Does not care if h is causal or acausal
'
' y(i) = summation(j = jmin to jmax) h(j)x(i-j)
'
' Example: h is the irf, h.val(j), where j=jmin,...,jmax
' x is the input signal, x.val(i) where i=imin,...,imax
' h and x must have the same sampling interval, x.inc(1)
' y is the result, Y(i) where i=imin,...,imax

Dim i As Long
Dim j As Long
Dim imin As Long
Dim imax As Long
Dim jmin As Long
Dim jmax As Long
Dim sum As Double

```

```

'bounds of sequence x
  imin = LBound(X.Val)
  imax = UBound(X.Val)

'bounds of irf h
  jmin = LBound(h.Val)
  jmax = UBound(h.Val)

Convolve2 = X

For i = imin To imax
  sum = 0#
  'conduct sum across bounds of irf, include only terms within bounds of x
  For j = jmin To jmax
    If (i - j) >= imin And (i - j) <= imax Then
      sum = sum + h.Val(j) * X.Val(i - j)
    End If
  Next j
  Convolve2.Val(i) = X.Inc(1) * sum
Next i

End Function

Public Function dtConvolve( _
  h As Tensor, _
  X As Tensor _
) As Tensor

' Convolves IRF h with sequence x.
' Does not care if h is causal or noncausal.
'
' y(i) = (dT)*summation(j = jmin to jmax) h(j)*x(i-j)
'
' Example: h is the irf, h.val(j), where j=jmin,...,jmax
'         x is the input signal, x.val(i) where i=imin,...,imax
'         h and x must have the same sampling interval, x.inc(1)
'         y is the result, Y(i) where i=imin,...,imax

' Local variables
  Dim i As Long
  Dim j As Long
  Dim imin As Long
  Dim imax As Long
  Dim jmin As Long
  Dim jmax As Long
  Dim sum As Double

' Bounds of sequence x
  imin = LBound(X.Val)
  imax = UBound(X.Val)

```

```

' Bounds of irf h
  jmin = LBound(h.Val)
  jmax = UBound(h.Val)

' Assign dtConvolve tensor info from tensor x
' dT should be the same.
  dtConvolve = X

' Convolution sum
  For i = imin To imax ' do for all terms
    sum = 0#
    'conduct sum across bounds of IRF, include only terms within bounds of x
    For j = jmin To jmax
      If (i - j) >= imin And (i - j) <= imax Then
        sum = sum + h.Val(j) * X.Val(i - j)
      End If
    Next j
    dtConvolve.Val(i) = X.Inc(1) * sum
  Next i
End Function

```

```

Public Function ReDFT( _
  X As Tensor, _
  Optional WindowFunction As WindowFn = Tukey _
) As Tensor

```

' Returns the real part of DFT(x), where x is a real valued sequence.

'  $DFT(n) = A(n) - j * B(n)$

'  $A(n) = \text{summation} \{x(k) * \cos(2 * \pi * k * n / N)\}$  for k = 0 to N-1

'  $B(n) = \text{summation} \{x(k) * \sin(2 * \pi * k * n / N)\}$  for k = 0 to N-1

' A(n) is the real part, and B(n) the imaginary part of the DFT

,

' Function ReDFT() conducts the N to N cosine transformation for the real part  
' of the DFT

'  $ReDFT(x) = A(n)$

,

' Written by JLT, Aug 2001

'Local variable declarations

Dim i As Long

Dim imin As Long

Dim imax As Long

Dim N As Long 'number of elements in sequence

Dim k As Long

Dim freqInc As Double 'ct frequency increment

Dim wx() As Double 'windowed sequence

Dim cosf As Double

```

Dim cos0 As Double
Dim cos1 As Double
Dim cos2 As Double
Dim Pi As Double

Pi = 4# * Atn(1#)

'find bounds of sequence X.val
imin = 0
imax = UBound(X.Val)
N = imax - imin + 1 'number of elements in sequence

'redim the array to hold windowed elements
ReDim wx(0 To imax)

'frequency interval, ct Hz equivalent
freqInc = 1# / (imax * X.Inc(1))

'set output tensor to x to set sequence information
ReDFT = X
ReDim ReDFT.Val(0 To imax)

'Apply a window to sequence
If WindowFunction = Tukey Then 'default
  For i = 0 To imax
    wx(i) = 0.5 * (1# + Cos(Pi * i / imax)) * X.Val(i)
  Next i

ElseIf WindowFunction = Triangular Then '(i.e. Bartlett)
  For i = 0 To imax
    wx(i) = (1# - i / imax) * X.Val(i)
  Next i

Else 'If WindowFunction = None (i.e. rectangular)
  For i = 0 To imax
    wx(i) = X.Val(i)
  Next i
End If

'Calculate ReDFT - a cosine transformation of x

'Do for all frequencies, from 0/N to (N-1)/N
For i = 0 To imax

  'initialize cos constants for freq
  cosf = Cos(2 * Pi * i / N)
  cos0 = 0#
  cos1 = 0#

  'recursive calculation of summation {x(k)cos(2*pi*k*j/N)}
  For k = imax To 1 Step -1

```

```

        cos2 = 2# * cosf * cos1 - cos0 + wx(k)
        cos0 = cos1
        cos1 = cos2
    Next k

    ReDFT.Val(i) = wx(0) + cos1 * cosf - cos0

Next i

ReDFT.Rname(1) = "Power"
ReDFT.Dname(1) = "Frequency"
ReDFT.Dunit(1) = "Hz"
ReDFT.Inc(1) = freqInc
End Function

Public Function ImDFT( _
    X As Tensor, _
    Optional WindowFunction As WindowFn = Tukey _
) As Tensor

' Returns the imaginary part of DFT(x), where x is a real valued sequence
' DFT(n)= A(n) - j*B(n)
'   A(n) = summation {x(k)*cos(2*pi*k*n/N)} for k = 0 to N-1
'   B(n) = summation {x(k)*cos(2*pi*k*n/N)} for k = 0 to N-1
' A(n) is the real part, and B(n) the imaginary part of the DFT

' Function ImDFT( ) conducts the N to N sine transformation for the imaginary
' part of the DFT
' ImDFT(x) = B(n)

'Written by JLT, Aug 2001

'Local variable declarations
    Dim i As Long
    Dim imin As Long
    Dim imax As Long
    Dim N As Long
    Dim k As Long
    Dim freqInc As Double
    Dim wx() As Double 'windowed sequence
    Dim cosf As Double
    Dim sinf As Double
    Dim sin0 As Double
    Dim sin1 As Double
    Dim sin2 As Double
    Dim Pi As Double

```



```

Pi = 4# * Atn(1#)

'find bounds of sequence
imin = 0          'calculate from zero to end of sequence
imax = UBound(X.Val)
N = imax - imin + 1 'sequence length

'redim the array to hold windowed elements
ReDim wx(0 To imax)

'frequency interval, ct Hz equivalent
freqInc = 1# / (imax * X.Inc(1))

'set output tensor to x to set sequence information
ImDFT = X
ReDim ImDFT.Val(0 To imax)

'Apply a window to sequence
If WindowFunction = Tukey Then 'default
  For i = 0 To imax
    wx(i) = 0.5 * (1# + Cos(Pi * i / imax)) * X.Val(i)
  Next i

ElseIf WindowFunction = Triangular Then '(i.e. Bartlett)
  For i = 0 To imax
    wx(i) = (1# - i / imax) * X.Val(i)
  Next i

Else 'If WindowFunction = None (i.e. rectangular)
  For i = 0 To imax
    wx(i) = X.Val(i)
  Next i
End If

'Calculate imDFT - a sine transformation of x

'Do for all frequencies, from 0/N to (N-1)/N
For i = 0 To imax

  'initialize cos and sin constants for freq
  cosf = Cos(2 * Pi * i / N) 'cos term needed to estimate the recursive sine
  sinf = Sin(2 * Pi * i / N)
  sin0 = 0#
  sin1 = 0#

  'recursive calculation of summation {x(k)sin(2*pi*k*i/N)}
  For k = imax To 1 Step -1
    sin2 = 2# * cosf * sin1 - sin0 + wx(k) 'J&W shows cos, not sin, term here!
    sin0 = sin1
    sin1 = sin2
  Next k

```

```

    ImDFT.Val(i) = sin1 * sinf

Next i

ImDFT.Rname(1) = "Power"
ImDFT.Dname(1) = "Frequency"
ImDFT.Dunit(1) = "Hz"
ImDFT.Inc(1) = freqInc
End Function

Public Function InvDFT( _
    realDFT As Tensor, _ imagDFT As Tensor _
) As Tensor

' Call: InvDFT(A as real part of DFT, B as imaginary part of DFT)
' summation(1/N){A(k)cos(2*pi*k*j/N)+ B(k)sin(2*pi*k*j/N)}
'
' Returns the real part of the inverse DFT from the real
' part of the DFT, A(n), and the imaginary part of the DFT, B(n).
' Where DFT(n)= A(n) - j*B(n)
'   A(n) = summation {x(k)*cos(2*pi*k*n/N)} for k = 0 to N-1
'   B(n) = summation {x(k)*sin(2*pi*k*n/N)} for k = 0 to N-1
' If the original sequence was real (before DFT) then this will return
' the sequence
'
'Written by JLT, Aug 2001

'Local variable declarations
    Dim i As Long
    Dim imin As Long
    Dim imax As Long
    Dim N As Long
    Dim k As Long
    Dim freqInc As Double
    Dim A() As Double 'realDFT
    Dim B() As Double 'imagDFT
    Dim cosf As Double
    Dim cos0 As Double
    Dim cos1 As Double
    Dim cos2 As Double
    Dim sinf As Double
    Dim sin0 As Double
    Dim sin1 As Double
    Dim sin2 As Double
    Dim Pi As Double
    Dim SampPeriod As Double

```

```

Pi = 4# * Atn(1#)

'find bounds of sequence
imin = 0      'calculate from zero to end of sequence
imax = UBound(realDFT.Val)
N = imax - imin + 1 'sequence length

'redim the array to hold real and imaginary
ReDim A(0 To imax)
ReDim B(0 To imax)

'Sampling period
SampPeriod = 1# / (imax * realDFT.Inc(1))

'set output tensor, InvDFT, to input realDFT to set sequence information
InvDFT = realDFT
ReDim InvDFT.Val(0 To imax)

'Set real and imaginary values
For i = 0 To imax
    A(i) = realDFT.Val(i)
    B(i) = imagDFT.Val(i)
Next i

'Calculate InvDFT

'Do for all frequencies, from 0/N to (N-1)/N
For i = 0 To imax

    'initialize cos and sin constants for freq
    cosf = Cos(2 * Pi * i / N)
    cos0 = 0#
    cos1 = 0#
    sinf = Sin(2 * Pi * i / N)
    sin0 = 0#
    sin1 = 0#

    'recursive calculation of summation(1/N){A(k)cos(2*pi*k*j/N)+ B(k)sin(2*pi*k*j/N)}
    For k = imax To 1 Step -1
        cos2 = 2# * cosf * cos1 - cos0 + A(k)
        cos0 = cos1
        cos1 = cos2
        sin2 = 2# * cosf * sin1 - sin0 + B(k) 'J&W shows cosf, not sinf, term here
        sin0 = sin1
        sin1 = sin2
    Next k

    InvDFT.Val(i) = (1 / N) * (A(0) + cos1 * cosf - cos0 + sin1 * sinf)

Next i

```

```

InvDFT.Rname(1) = "sequence"
InvDFT.Dname(1) = "Time"
InvDFT.Dunit(1) = "s"
InvDFT.Inc(1) = SampPeriod

```

End Function

```

Public Function DisplayTensor( _
    X As Tensor _
) As Tensor

```

```

' Rearranges the order of a tensor X when there are elements with negative indices.
' The elements with positive indices remain in the same position.
' The elements with negative indices, x.val(k) where k < 0, are moved
' to DisplayTensor.val(len+k)
' Elements 0 to kmax will not change. The negative values will be moved en masse
' after kmax

```

```

'Variable declaration
Dim imin As Long
Dim imax As Long
Dim i As Long
Dim N As Long

```

```

'set the tensor equal to the argument
    DisplayTensor = X

```

```

'get bounds of sequence
    imin = LBound(X.Val)
    imax = UBound(X.Val)
    N = imax - imin + 1 'Total number of elements

```

```

    ReDim DisplayTensor.Val(0 To (N - 1)) 'Redim so there are only positive indices

```

```

'positive indices
    For i = 0 To imax
        DisplayTensor.Val(i) = X.Val(i)
    Next i

```

```

'negative values, if there are any
    If imin < 0 Then
        For i = imin To -1
            DisplayTensor.Val(N + i) = X.Val(i)
        Next i
    End If

```

End Function

```

Public Function AutoCorrelation( _
  X As Tensor, _
  Optional jmin As Long = 0, _
  Optional jmax As Long = 100 _
) As Tensor

' Computes positive lags of the biased auto-correlation function of vector X

' The auto-correlation function returned is Cxx(j) j=jmin,...jmax
' where jmin is the specified minimum correlation lag (default 0)
'   jmax is the specified maximum correlation lag (default 100)
'
' The auto-correlation function is symmetric about j=0 and so
' Cxx(j) = Cxx(-j). Normally jmin will be set = 0 for auto-correlation
' function determination.
'
' Note that this function determines the biased estimator which reduces
' the variance of the longer lag auto-correlation function estimates.
'
' This function was adapted by J.Tangorra from the auto-covariance function
' written by Ian Hunter in the System Id section
' Last update 31 Jul 2001

Dim i As Long
Dim j As Long
Dim imin As Long
Dim imax As Long
Dim sum As Double
imin = LBound(X.Val)
imax = UBound(X.Val)

'Set AutoCorrelation equal to tensor x to set all increments the same
AutoCorrelation = X
AutoCorrelation.Type = "Auto-Correlation Function"
AutoCorrelation.Rname(1) = "Correlation"
AutoCorrelation.Dname(1) = "Lag"
ReDim AutoCorrelation.Val(jmin To jmax)

For j = jmin To jmax
  sum = 0#
  For i = imin + j To imax + jmin
    If i - j >= imin And i >= imin And i - j <= imax And i <= imax Then
      sum = sum + X.Val(i - j) * X.Val(i)
    End If
  Next i
  AutoCorrelation.Val(j) = sum / (imax - imin + 1) 'biased estimator
Next j
End Function

```

```

Public Function CrossCorrelation( _
  X As Tensor, _
  Y As Tensor, _
  Optional jmin As Long = 0, _
  Optional jmax As Long = 100 _
) As Tensor

'Computes the biased Cross-Covariance function between vectors X and Y
'which must have the same lengths.
'The CrossCorrelation function returned is Cxy(j) j=jmin,...,jmax
'where jmin is the specified minimum CrossCorrelation lag (default 0)
'jmax is the specified maximum CrossCorrelation lag (default 100)
,

'The CrossCorrelation function is not symmetric and so
'Cxy(j) is not equal to Cxy(-j).
,

'Note that this function determines the biased estimator which reduces
' the variance of the longer lag cross-covariance function estimates.
,

'This function was adapted by JLT from the CrossCovariance function
'written by Ian Hunter.
'Last update Jul 31 2001

Dim i As Long
Dim j As Long
Dim imin As Long
Dim imax As Long
Dim sum As Double
Dim MeanX As Double
Dim MeanY As Double
imin = LBound(X.Val)
imax = UBound(X.Val)
'note that we do not check that y has the same imin and imax

CrossCorrelation = X
CrossCorrelation.Type = "Cross-Correlation Function"
CrossCorrelation.Rname(1) = "Correlation"
CrossCorrelation.Dname(1) = "Lag"
ReDim CrossCorrelation.Val(jmin To jmax)

For j = jmin To jmax
  sum = 0#
  For i = imin + j To imax + jmin
    If i - j >= imin And i >= imin And i - j <= imax And i <= imax Then
      sum = sum + X.Val(i - j) * Y.Val(i)
    End If
  Next i
  CrossCorrelation.Val(j) = sum / (imax - imin + 1) 'biased estimator
Next j
End Function

```

```

Public Function ACF( _
  X As Tensor, _
  Optional jmin As Long = -100, _
  Optional jmax As Long = 100 _
) As Tensor

' Computes the (-/+ ) lags biased auto correlation function of vector X
'
' The ACF function returned is Cxx(j) j=jmin,...jmax
' where jmin is the specified minimum auto correlation lag (default -100)
' jmax is the specified maximum cross correlation lag (default 100)
'
' The auto-correlation function is symmetric; Cxx(j)= Cxx(-j).
'
' Note that this function determines the biased estimator which reduces
' the variance of the longer lag auto correlation function estimates.
'
' This function was adapted by JLT from the AutoCovariance function
' written by Ian Hunter.
' Last update Sep 17 2001

'Variable declarations
  Dim N As Long      'sequence length
  Dim i As Long      'sequence element
  Dim imin As Long   'lower bound for correlation sum
  Dim imax As Long   'upper bound for correlation sum
  Dim j As Long      'lag values
  Dim sum As Double

'Length of tensors
  N = UBound(X.Val) - LBound(X.Val) + 1
  'note that we do not check that y has the same bounds

'Set ACF info, redim to hold lags
  ACF = X
  ACF.Type = "Cross Correlation Function"
  ACF.Rname(1) = "Correlation"
  ACF.Dname(1) = "Lag"
  ReDim ACF.Val(jmin To jmax)

'ACF
  For j = jmin To jmax
    sum = 0
    'set bounds for correlation sum
    If j < 0 Then
      imin = -j
      imax = N - 1
    ElseIf j >= 0 Then
      imin = 0
      imax = N - 1 - j
    End If

```

```

    'sum elements
    For i = imin To imax
        sum = sum + (X.Val(i) * X.Val(i + j))
    Next i

    ACF.Val(j) = (1 / N) * sum
Next j

End Function

Public Function CCF( _
    X As Tensor, _
    Y As Tensor, _
    Optional jmin As Long = -100, _
    Optional jmax As Long = 100 _
) As Tensor

' Computes the biased (-/+ ) lags of cross correlation function between vectors X and Y
' which must have the same lengths.
'
' The CCF function returned is Cxy(j) j=jmin,...,jmax
' where jmin is the specified minimum cross correlation lag (default -100)
' jmax is the specified maximum cross correlation lag (default 100)
'
' The cross correlation function is not symmetric and so
' Cxy(j) is not equal to Cxy(-j), but Cxy(j) does equal Cyx(-j)
'
' Note that this function determines the biased estimator which reduces
' the variance of the longer lag cross correlation function estimates.
'
' This function was adapted by JLT from the CrossCovariance function
' written by Ian Hunter.
' Last update Sep 17 2001

'Variable declarations
    Dim N As Long      'sequence length
    Dim i As Long      'sequence element
    Dim imin As Long   'lower bound for correlation sum
    Dim imax As Long   'upper bound for correlation sum
    Dim j As Long      'lag values
    Dim sum As Double

'Length of tensors
    N = UBound(X.Val) - LBound(X.Val) + 1
    'note that we do not check that y has the same bounds

'Set CrossCorr info, redim to hold lags
    CCF = X
    CCF.Type = "Cross Correlation Function"

```



```

CCF.Rname(1) = "Correlation"
CCF.Dname(1) = "Lag"
ReDim CCF.Val(jmin To jmax)

'CCF
  For j = jmin To jmax
    sum = 0
    'set bounds for correlation sum
    If j < 0 Then
      imin = -j
      imax = N - 1
    ElseIf j >= 0 Then
      imin = 0
      imax = N - 1 - j
    End If
    'sum elements
    For i = imin To imax
      sum = sum + (X.Val(i) * Y.Val(i + j))
    Next i

    CCF.Val(j) = (1 / N) * sum
  Next j
End Function

```

```

Public Function TwoSidedPSD( _
  Cxx As Tensor, _
  Optional WindowFunction As WindowFn = Tukey _
) As Tensor

' Created by modifying AutoSpectrum function
'
' TwoSidedPSD( ) calculates the two sided PSD by taking the cosine
' transformation of an ACF with positive lags only. The two sided PSD calculates
' the transform from 0 to on increment below 1.0 dtHz. The output is symmetric
' about the 0.5 dtHz value, with the exception of the 0.0 Hz value.
' The ACF argument has positive lags only, so is not actually even.
' When the ACF is not even, the Fourier transform is not real. The cosine
' transform ignores the imaginary content.
' The ACF function is multiplied by the specified window weighting
' function which can be:
' Tukey (default),
' Triangular (i.e. Bartlett), or
' None (i.e. rectangular).
' If the ACF function has already been windowed using the window()
' function then WindowFunction should be set = None.
' Note that for the Tukey and Triangular windows a sinusoid of amplitude 1.0 will
' have a power of 1.0, and a sinusoid of amplitude 2.0 will have a power of 4.0.

```

```

Dim i As Long
Dim j As Long
Dim k As Long
Dim imin As Long
Dim imax As Long
Dim N As Long           'sequence length
Dim freqInc As Double  'frequency increment in ctHz
Dim wCxx() As Double   'windowed ACF
Dim cos0 As Double
Dim cos1 As Double
Dim cos2 As Double
Dim cosf As Double
Dim Pi As Double

```

```
Pi = 4# * Atn(1#)
```

```
imin = LBound(Cxx.Val) 'not used (start at 0)
```

```
imin = 0
```

```
imax = UBound(Cxx.Val)
```

```
N = imax - imin + 1           'length of ACF
```

```
ReDim wCxx(0 To imax)       'array of windowed ACF values
```

```
freqInc = 1# / ((N - 1) * Cxx.Inc(1)) 'Cxx.Inc(1) is the sampling period
```

```
TwoSidedPSD = Cxx           'set equal for information values
```

```
ReDim TwoSidedPSD.Val(0 To imax)
```

```
'Window ACF values
```

```
If WindowFunction = Tukey Then 'default
```

```
  For i = 0 To imax
```

```
    wCxx(i) = 0.5 * (1# + Cos(Pi * i / imax)) * Cxx.Val(i)
```

```
  Next i
```

```
ElseIf WindowFunction = Triangular Then '(i.e. Bartlett)
```

```
  For i = 0 To imax
```

```
    wCxx(i) = (1# - i / imax) * Cxx.Val(i)
```

```
  Next i
```

```
Else If WindowFunction = None (i.e. rectangular)
```

```
  For i = 0 To imax
```

```
    wCxx(i) = Cxx.Val(i)
```

```
  Next i
```

```
End If
```

```
'Perform cosine (real part of Fourier) transform
```

```
'Calculate for 0 to 1 freq increment below 1.0 dtHz, else 0 Hz would be calculated twice
```

```
For j = 0 To imax
```

```
  'initial cos terms for freq j/N
```

```
  cosf = Cos(2 * Pi * j / N)
```

```
  cos0 = 0#
```

```

cos1 = 0#

For k = imax To 1 Step -1 'recursive calculation of summation {cos(2*pi*k*j/N)}
  cos2 = 2# * cosf * cos1 - cos0 + wCxx(k)
  cos0 = cos1
  cos1 = cos2
Next k

'using one sided ACF, so amplitude should be doubled
'but calculating for 0 thru 1.0 dtHz so don't double value of entire sequence
TwoSidedPSD.Val(j) = wCxx(0) + 2# * (cos1 * cosf - cos0)

Next j

TwoSidedPSD.Rname(1) = "Power"
TwoSidedPSD.Dname(1) = "Frequency"
TwoSidedPSD.Dunit(1) = "Hz"
TwoSidedPSD.Inc(1) = freqInc

End Function

```

```

Public Function PSDfromACF( _
  Cxx As Tensor, _
  Optional WindowFunction As WindowFn = Tukey _
) As Tensor

' Works in conjunction with ACF()
' Created by modifying TwoSidedPSD function

Dim i As Long
Dim j As Long
Dim k As Long
Dim imin As Long
Dim imax As Long
Dim N As Long          'sequence length
Dim M As Long
Dim freqInc As Double  'frequency increment in ctHz
Dim wCxx() As Double   'windowed ACF
Dim cos0 As Double
Dim cos1 As Double
Dim cos2 As Double
Dim cosf As Double
Dim Pi As Double

Pi = 4# * Atn(1#)

'Bounds of ACF argument, pos and neg lags should be symmetric
imin = LBound(Cxx.Val)

```

```

imax = UBound(Cxx.Val)
If imin <> (-imax) Then imin = 0      'verify that function is symmetric else just use pos lags
N = imax + 1                        'represents 0 through poslags
M = imax - imin + 1                 'length of ACF

'some terms
freqInc = 1# / ((M - 1) * Cxx.Inc(1)) 'Cxx.Inc(1) is the sampling period
PSDfromACF = Cxx                     'set equal to establish information

'Window ACF values
ReDim wCxx(imin To imax)             'array for windowed ACF values

If WindowFunction = Tukey Then      'default
  For i = imin To imax
    wCxx(i) = 0.5 * (1# + Cos(Pi * i / (M - 1))) * Cxx.Val(i)
  Next i
ElseIf WindowFunction = Triangular Then '(i.e. Bartlett)
  For i = imin To imax
    wCxx(i) = (1# - i / (M - 1)) * Cxx.Val(i)
  Next i
Else 'If WindowFunction = None (i.e. rectangular)
  For i = imin To imax
    wCxx(i) = Cxx.Val(i)
  Next i
End If

'Perform cosine (real part of Fourier) transform

'Calculate for -0.5 to 0.5
For j = imin To imax
  'initial cos terms for freq j/M
  cosf = Cos(2 * Pi * j / M)
  cos0 = 0#
  cos1 = 0#

  'Recursive calculation of summation{cos(2*pi*k*j/N)}
  ' Only necessary to sum from k = imax to 1, since the cos of the
  ' negative k indices is equal to the cos of the positive indices.
  ' But, must double the amplitude of the sum.
  For k = imax To 1 Step -1
    cos2 = 2# * cosf * cos1 - cos0 + wCxx(k)
    cos0 = cos1
    cos1 = cos2
  Next k

  'Only summed from imax to 1, so double amplitude of sum to account for -1 to imin
  PSDfromACF.Val(j) = wCxx(0) + 2# * (cos1 * cosf - cos0)

Next j

PSDfromACF.Rname(1) = "Power"

```

```

PSDfromACF.Dname(1) = "Frequency"
PSDfromACF.Dunit(1) = "Hz"
PSDfromACF.Inc(1) = freqInc
End Function

```

```

Public Function ACFfromPSD( _
    PSD As Tensor _
) As Tensor

```

```

' ACFfromPSD( ) transforms a two sided PSD into an ACF with a
' cosine transform (real part of the inverse Fourier transform).
' The PSD argument is real and even, so the invDFT just has a real part
' which is the cosine transformation.
' The ACF found from this algorithm is symmetric.

```

```

' JLT, Sep 2001

```

```

Dim i As Long
Dim j As Long
Dim imin As Long
Dim imax As Long
Dim N As Long 'length of PSD
Dim M As Long
Dim sampPer As Double
Dim cos0 As Double
Dim cos1 As Double
Dim cos2 As Double
Dim cosf As Double
Dim Pi As Double

```

```

Pi = 4# * Atn(1#)

```

```

'Bounds of input PSD
imin = LBound(PSD.Val) 'not used (start at 0)
imax = UBound(PSD.Val)
If imin <> -imax Then imin = 0
N = imax + 1 'length of positive lags, and original sequence
M = imax - imin + 1 'length of entire PSD

```

```

'Sampling period
sampPer = 1# / (imax * PSD.Inc(1))

```

```

'Set aCF tensor equal to PSD for sequence info
ACFfromPSD = PSD
ReDim ACFfromPSD.Val(imin To imax)

```

```

'Perform real part of the invDFT.

```

'This calculation is the same as the real Fourier transform (cosine xform)  
'scaled by 1/M

For j = imin To imax

cosf = Cos(2 \* Pi \* j / M)  
cos0 = 0#  
cos1 = 0#

For i = imax To 1 Step -1  
cos2 = 2# \* cosf \* cos1 - cos0 + PSD.Val(i)  
cos0 = cos1  
cos1 = cos2  
Next i

ACFfromPSD.Val(j) = (1 / M) \* (PSD.Val(0) + 2 \* (cos1 \* cosf - cos0))

Next j

ACFfromPSD.Rname(1) = "ACF"  
ACFfromPSD.Dname(1) = "Time"  
ACFfromPSD.Dunit(1) = "s"  
ACFfromPSD.Inc(1) = sampPer

End Function

Public Function ReCrossPSD( \_  
X As Tensor, \_  
Y As Tensor, \_  
maxlag As Long \_  
) As Tensor

' Calculates the real part of the cross spectral density by taking the DFT of the  
' cross-correlation of tensor X and Y.

' The two tensors should be the same length, and should have been  
' sampled at the same rate.

' MaxLag sets the largest lag used in the calculation of the CCF

'  $CrossPSD(n) = \sum \{ C_{xy}(k) * \cos(2 * \pi * n * k / M) \} - j \sum \{ C_{xy}(k) * \sin(2 * \pi * n * k / M) \}$   
'  $= ReCrossPSD(n) - j ImPSD(n)$

' Algorithm exploits  $C_{xy}(-j) = C_{yx}(j)$ . By using only the positive lags  
' of  $C_{xy}$  and  $C_{yx}$ , the cosine/sine terms in the summation can be  
' calculated with a recursive method from Jenkins and Watts

'Local variables

Dim Pi As Double  
Dim Nx As Long 'length of X  
Dim Ny As Long 'length of Y

```

Dim N As Long      'min length of X and Y
Dim Cxy As Tensor
Dim Cyx As Tensor
Dim jmin As Long
Dim jmax As Long
Dim M As Long      'equivalent length of - and + lags for correlation
Dim freqInc As Double

Dim j As Long
Dim k As Long

Dim cosf As Double
Dim cos0 As Double
Dim cos1 As Double
Dim cos2 As Double

'define Pi
Pi = 4# * Atn(1#)

'lengths of X and Y
Nx = UBound(X.Val) - LBound(X.Val) + 1
Ny = UBound(Y.Val) - LBound(Y.Val) + 1

'set N to smallest tensor length
If Nx < Ny Then
    N = Nx
Else: N = Ny
End If

'Calculate pos lags of both Cxy and Cyx
'Ensure maxlag is less than min length
If maxlag > N - 1 Then
    maxlag = N - 1
Else: maxlag = maxlag
End If

Cxy = CCF(X, Y, 0, maxlag)
Cyx = CCF(Y, X, 0, maxlag)

'some terms
jmax = UBound(Cxy.Val)
jmin = -jmax
M = jmax - jmin + 1      'equivalent length of neg and pos lags values for one CCF
freqInc = 1 / ((M - 1) * Cxy.Inc(1))

'Set bounds of Cross PSD
ReCrossPSD = Cxy
ReDim ReCrossPSD.Val(jmin To jmax)

'Windowing - do later

```

```

'Perform cosine transformation.
'Split into two sums, one with Cxy and other with Cyx.
'Realize that the summation indices can be combined to
'get a single summation that uses the sum of Cxy and Cyx

For j = jmin To jmax

'Initialize
  cosf = Cos(2 * Pi * j / M)
  cos0 = 0
  cos1 = 0

'recursive calculation
  For k = maxlag To 1 Step -1
    cos2 = 2 * cosf * cos1 - cos0 + (Cxy.Val(k) + Cyx.Val(k))
    cos0 = cos1
    cos1 = cos2
  Next k

'set Re(spectrum) value
  ReCrossPSD.Val(j) = Cxy.Val(0) + (cos1 * cosf - cos0)

Next j

ReCrossPSD.Rname(1) = "Power"
ReCrossPSD.Dname(1) = "Frequency"
ReCrossPSD.Dunit(1) = "Hz"
ReCrossPSD.Inc(1) = freqInc

End Function

```

```

Public Function ImCrossPSD( _
  X As Tensor, _
  Y As Tensor, _
  maxlag As Long _
) As Tensor

' Calculates the imaginary part of the cross spectral density by taking the DFT of the
' cross-correlation of tensor X and Y.
' The two tensors should be the same length, and should have been
' sampled at the same rate.
' MaxLag sets the largest lag used in the calculation of the CCF
'
' CrossPSD(n) = sum{ Cxy(k)*cos(2*pi*n*k/M) } - j sum{ Cxy(k)*sin(2*pi*n*k/M) }
'             = ReCrossPSD(n) - j ImPSD(n)
' Algorithm exploits Cxy(-j)=Cyx(j). By using only the positive lags

```



' of Cxy and Cyx, the cosine/sine terms in the summation can be  
' calculated with a recursive method from Jenkins and Watts

'Local variables

```
Dim Pi As Double
Dim Nx As Long   'length of X
Dim Ny As Long   'length of Y
Dim N As Long    'min length of X and Y
Dim Cxy As Tensor
Dim Cyx As Tensor
Dim jmin As Long
Dim jmax As Long
Dim M As Long    'equivalent length of - and + lags for correlation
Dim freqInc As Double
```

```
Dim j As Long
Dim k As Long
```

```
Dim cosf As Double
Dim sinf As Double
Dim sin0 As Double
Dim sin1 As Double
Dim sin2 As Double
```

'define Pi

```
Pi = 4# * Atn(1#)
```

'lengths of X and Y

```
Nx = UBound(X.Val) - LBound(X.Val) + 1
Ny = UBound(Y.Val) - LBound(Y.Val) + 1
```

'set N to smallest tensor length

```
If Nx < Ny Then
    N = Nx
Else: N = Ny
End If
```

'Calculate pos lags of both Cxy and Cyx

```
'Ensure maxlag is less than min length
If maxlag > N - 1 Then
    maxlag = N - 1
Else: maxlag = maxlag
End If
Cxy = CCF(X, Y, 0, maxlag)
Cyx = CCF(Y, X, 0, maxlag)
```

'some terms

```
jmax = UBound(Cxy.Val)
jmin = -jmax
M = jmax - jmin + 1  'equivalent length of neg and pos lags values for one CCF
```

```

    freqInc = 1 / ((M - 1) * Cxy.Inc(1))

'Initialize and set bounds of Cross PSD
    ImCrossPSD = Cxy
    ReDim ImCrossPSD.Val(jmin To jmax)

'Windowing - do later

'Perform sine transformation.
    'Split into two sums, one with Cxy and other with Cyx.
    'Realize that the summation indices can be combined to
    'get a single summation that uses the difference of Cxy and Cyx

For j = jmin To jmax

'Initialize
    cosf = Cos(2 * Pi * j / M)
    sinf = Sin(2 * Pi * j / M)
    sin0 = 0
    sin1 = 0

'recursive calculation
    For k = maxlag To 1 Step -1
        sin2 = 2 * cosf * sin1 - sin0 + (Cxy.Val(k) - Cyx.Val(k))
        sin0 = sin1
        sin1 = sin2
    Next k

'set Im(spectrum) value
    ImCrossPSD.Val(j) = (sin1 * sinf - sin0)

Next j

    ImCrossPSD.Rname(1) = "Power"
    ImCrossPSD.Dname(1) = "Frequency"
    ImCrossPSD.Dunit(1) = "Hz"
    ImCrossPSD.Inc(1) = freqInc

End Function

Public Function CrossAmpSpectrum( _
    X As Tensor, _
    Y As Tensor, _
    maxlag As Long _
) As Tensor

' Calculates the Cross amplitude spectrum of tensor x and y,
' By taking the DFT of the cross correlation function.
' Uses Plexus functions ReCrossPSD and ImCrossPSD.

```

```
' Want X and Y to be of same length and to have been  
' sampled at same frequency
```

```
' MaxLag sets the largest lag used in the calculation of the CCF
```

```
'The output tensor will have neg and pos frequencies
```

```
'Know that before displaying the tensor values
```

```
'May want to use DisplayTensor function
```

```
'Local variables
```

```
Dim ReCross As Tensor
```

```
Dim ImCross As Tensor
```

```
Dim i As Long
```

```
'Find real and imaginary parts of spectrum
```

```
ReCross = ReCrossPSD(X, Y, maxlag)
```

```
ImCross = ImCrossPSD(X, Y, maxlag)
```

```
' Initialize tensor information
```

```
CrossAmpSpectrum = ReCross
```

```
For i = LBound(CrossAmpSpectrum.Val) To UBound(CrossAmpSpectrum.Val)
```

```
    CrossAmpSpectrum.Val(i) = Sqr(ReCross.Val(i) ^ 2 + ImCross.Val(i) ^ 2)
```

```
Next i
```

```
CrossAmpSpectrum.Rname(1) = "Power"
```

```
CrossAmpSpectrum.Dname(1) = "Frequency"
```

```
CrossAmpSpectrum.Dunit(1) = "Hz"
```

```
'CrossAmpSpectrum.Inc(1) = freqInc 'already set by ReCross
```

```
End Function
```

```
Public Function PhaseSpectrum( _
```

```
    X As Tensor, _
```

```
    Y As Tensor, _
```

```
    maxlag As Long _
```

```
) As Tensor
```

```
' Calculates the Phase spectrum of tensors x and y.
```

```
' Output in degrees.
```

```
' Uses Plexus functions ReCrossPSD and ImCrossPSD.
```

```
' Want X and Y to be of same length and to have been
```

```
' sampled at same frequency
```

```
' MaxLag sets the largest lag used in the calculation of the CCF
```

```
' The output tensor will have neg and pos frequencies
```

```
' Know that before displaying the tensor values
```

```
' May want to use DisplayTensor function
```

```

'Local variables
  Dim Pi As Double
  Dim ReCross As Tensor
  Dim ImCross As Tensor
  Dim i As Long

'define Pi
  Pi = 4# * Atn(1#)

'Find real and imaginary parts of spectrum
  ReCross = ReCrossPSD(X, Y, maxlag)
  ImCross = ImCrossPSD(X, Y, maxlag)

' Initialize tensor information
  PhaseSpectrum = ReCross

  For i = LBound(PhaseSpectrum.Val) To UBound(PhaseSpectrum.Val)
    PhaseSpectrum.Val(i) = Atn(-ImCross.Val(i) / ReCross.Val(i)) * 180 / Pi
  Next i

  PhaseSpectrum.Rname(1) = "Phase"
  PhaseSpectrum.Dname(1) = "Frequency"
  PhaseSpectrum.Dunit(1) = "Deg"
  'PhaseSpectrum.Inc(1) = freqInc 'already set by ReCross
End Function

```

```

Public Function SSD(X1 As Tensor, X2 As Tensor) As Double

```

```

' Calculates the sum of the square difference between the elements of
' two tensors of the same length.
' SSD = 1010101010101010 returned if tensors not of same length.
' Written originally to be used to compute the SSD between the ACF of a sequence
' and a desired ACF in the Swap4PSD function

```

```

  Dim kmin As Long
  Dim kmax As Long
  Dim k As Long

```

```

  kmin = LBound(X1.Val)
  kmax = UBound(X1.Val)

```

```

  If LBound(X1.Val) = LBound(X2.Val) And UBound(X1.Val) = UBound(X2.Val) Then
    'conduct SSD calculation
    SSD = 0
    For k = kmin To kmax
      SSD = SSD + (X1.Val(k) - X2.Val(k)) ^ 2
    Next k
  End If

```

```

Else
  'Give error message
  SSD = 1.01010101010101E+17
End If
End Function

```

```

Public Function IntegrateTR(X As Tensor) As Tensor
'Integrates X using a simple Trapezoidal Rule
'Similar procedure in Sys ID section "Integrate( )"

```

```

'Local variable declarations

```

```

  Dim N As Long          'length of X

```

```

  Dim SampInt As Double  'sampling interval of sequence

```

```

  Dim imin As Long

```

```

  Dim imax As Long

```

```

  Dim i As Long

```

```

'set output tensor to X

```

```

  IntegrateTR = X

```

```

'Sampling interval

```

```

  SampInt = X.Inc(1)

```

```

'length of sequence

```

```

  imin = LBound(X.Val)

```

```

  imax = UBound(X.Val)

```

```

  N = imax - imin + 1

```

```

'initial integral value is 0

```

```

  IntegrateTR.Val(imin) = 0

```

```

'add update to previous val

```

```

  For i = imin + 1 To imax

```

```

    IntegrateTR.Val(i) = IntegrateTR.Val(i - 1) + (1 / 2) * SampInt * (X.Val(i - 1) + X.Val(i))

```

```

  Next i

```

```

End Function

```

```

Public Function ScaleTensor(X As Tensor, A As Double) As Tensor

```

```

'Returns X.val(i) = A*X.val(i)

```

```

,

```

```

'This function has not been copied or adapted from any software.

```

```

'It was written by JLT

```

```

'Last update 15 Oct 2001

```

Dim i As Long

'set new tensor info to old tensor info  
ScaleTensor = X

'Scale amplitude  
For i = LBound(X.Val) To UBound(X.Val)  
  ScaleTensor.Val(i) = A \* X.Val(i)  
Next i  
End Function

Public Function ACFAfterSwap( \_  
  Y As Tensor, \_  
  ACFold As Tensor, \_  
  minlag As Long, \_  
  maxlag As Long, \_  
  j1 As Long, \_  
  j2 As Long \_  
  ) As Tensor

' Procedure updates the ACF of a sequence after a swap of j1 and j2  
' Arguments  
' Y - sequence to be swapped  
' ACFold - ACF of sequence before swap  
' minlag, maxlag - lags to be considered fo update  
' j1, j2 - elements of sequence to be swapped

Dim k As Long  
Dim imin As Long  
Dim imax As Long  
Dim N As Long  
Dim Dacf As Double

'Begin with update set to old  
ACFAfterSwap = ACFold

'Bounds of the sequence being swapped  
imin = LBound(Y.Val)  
imax = UBound(Y.Val)  
N = imax - imin + 1

For k = minlag To maxlag  
' Calculate effect of swap on ACFold, CCFold at every lag value.  
' Define update as ACFnew, CCFnew.

'ACF update

```

'Developed from perspective of  $C_{yy}(k) = (1/N) * \text{summation}[y(i+k)*y(i)]$ 
Dacf = 0
If ((j1 + k) >= imin And (j1 + k) <= imax And (j1 + k) <> j2) Then
    Dacf = Dacf - Y.Val(j1) * Y.Val(j1 + k) + Y.Val(j2) * Y.Val(j1 + k)
End If

If ((j1 - k) >= imin And (j1 - k) <= imax And (j1 - k) <> j2) Then
    Dacf = Dacf - Y.Val(j1) * Y.Val(j1 - k) + Y.Val(j2) * Y.Val(j1 - k)
End If

If ((j2 + k) >= imin And (j2 + k) <= imax And (j2 + k) <> j1) Then
    Dacf = Dacf - Y.Val(j2) * Y.Val(j2 + k) + Y.Val(j1) * Y.Val(j2 + k)
End If

If ((j2 - k) >= imin And (j2 - k) <= imax And (j2 - k) <> j1) Then
    Dacf = Dacf - Y.Val(j2) * Y.Val(j2 - k) + Y.Val(j1) * Y.Val(j2 - k)
End If

ACAfterSwap.Val(k) = ACFold.Val(k) + (Dacf / N)

Next k
End Function

```

```

Public Function CCAfterSwap( _
    Y As Tensor, _
    X As Tensor, _
    CCFold As Tensor, _
    minlag As Long, _
    maxlag As Long, _
    j1 As Long, _
    j2 As Long _
) As Tensor

```

```

' Procedure updates the CCF of a sequence after a swap of j1 and j2
' Arguments
' Y - sequence to be swapped
' X - remains the same
' CCFold - CCF of sequence before swap
' minlag, maxlag - lags to be considered fo update
' j1, j2 - elements of sequence to be swapped

```

```

Dim k As Long
Dim imin As Long
Dim imax As Long
Dim N As Long
Dim Dccf As Double

```

```

'Begin with update set to old

```

CCFafterSwap = CCFold

'Bounds of the sequence being swapped

imin = LBound(Y.Val)

imax = UBound(Y.Val)

N = imax - imin + 1

For k = minlag To maxlag

'CCF update

'Developed from perspective of  $Cyx(k) = (1/N) * \text{summation}[y(i) * x(i-k)]$

'where y is sequence being swapped

Dccf = 0

If ((j1 - k) >= imin And (j1 - k) <= imax And j1 <> j2) Then

    Dccf = Dccf - X.Val(j1 - k) \* Y.Val(j1) + X.Val(j1 - k) \* Y.Val(j2)

End If

If ((j2 - k) >= imin And (j2 - k) <= imax And j2 <> j1) Then

    Dccf = Dccf - X.Val(j2 - k) \* Y.Val(j2) + X.Val(j2 - k) \* Y.Val(j1)

End If

CCFafterSwap.Val(k) = CCFold.Val(k) + (Dccf / N)

Next k

End Function

Public Function LPF(X As Tensor, CutOffFreq As Double) As Tensor

'Convolve X with the IRF of a finite impulse response (non-recursive) low pass filter

'Filter IRF approximated as the inverse Fourier transform of the filter FRF. This algorithm does not apply a tapering window the IRF, so will be affected by Gibbs phenomenon and other problems of "rectangular"

'windowed sequences.

'Local variable declarations

Dim N As Long           'length of X

Dim maxlag As Long

Dim ACForig As Tensor   'ACF of original sequence

Dim PSDorig As Tensor   'PSD of original sequence

Dim M As Long           'length of ACF, set by maxLag

Dim LPFfreq As Tensor   'Filter's spectrum

Dim LPFtime As Tensor   'Time domain filter elements

Dim SampInt As Double   'sampling interval of sequence

Dim CutOffElem As Long



```

Dim imin As Long
Dim imax As Long
Dim i As Long

'set output tensor to X
LPF = X

'length of original sequence
N = UBound(X.Val) - LBound(X.Val) + 1

'ACF & PSD original sequence
maxlag = N - 1
ACForig = ACF(X, -maxlag, maxlag)
PSDorig = PSDfromACF(ACForig, None)
imin = LBound(PSDorig.Val)
imax = UBound(PSDorig.Val)
M = imax - imin + 1

'Define filter spectrum with same num of elements as PSDorig
LPFfreq = PSDorig 'set equal for info
SampInt = X.Inc(1)

'initialize all filter PSD elements to zero
For i = imin To imax
    LPFfreq.Val(i) = 0
Next i

'Scale elements in desired freq range to have amplitude of (1/sampInt)
'accounts for sampling interval necessary in convolution
CutOffElem = Int(CutOffFreq * SampInt * (M - 1))
For i = 0 To CutOffElem
    LPFfreq.Val(i) = 1 / SampInt 'positive freq elements
    LPFfreq.Val(-i) = 1 / SampInt 'negative freq elements
Next i

'Time domain filter IRF - no windowing applied
LPFtime = ACFfromPSD(LPFfreq) 'essentially the inverse Fourier transformation to get the
Filter IRF

'Convolve filter IRF with sequence to be filtered
LPF = dtConvolve(LPFtime, X)

End Function

```

```

Public Function LPFwind(X As Tensor, CutOffFreq As Double, win As Integer) As Tensor
'
' Convolves sequence X with an acausal IRF of a finite impulse response (non-recursive)
' low pass filter. By default the IRF is 2x the length of the sequence. This means that
' much of the IRF is ~ zero, and significantly, and unnecessarily, lengthens the calculation.
'
' The ideal LPF filter IRF is approximated by taking the inverse Fourier transform of the
' filter FRF. An alternate method would be to specify the IRF with the sinc function. Using
' the invDFT allows the desired freqs and amplitudes to be specified.
'
' This algorithm applies a window to the IRF
' win = 0 (rectangular)
' win = 1 (Hanning)
' win = 2 (Hamming)
' win = 3 (Blackman)

'Local variable declarations
Dim N As Long          'length of sequence X
Dim maxlag As Long

Dim imin As Long
Dim imax As Long
Dim i As Long

Dim ACForig As Tensor  'ACF of original sequence
Dim PSDorig As Tensor  'PSD of original sequence
Dim M As Long          'length of ACF, set by maxLag

Dim LPFfreq As Tensor  'Filter's spectrum
Dim LPFfirf As Tensor  'IRF: time domain filter elements
Dim winLPFfirf() As Double 'Windowed LPF IRF

Dim SampInt As Double  'sampling interval of sequence
Dim CutOffElem As Long

' 0. Set output tensor to X to give it the appropriate tensor information
LPFwind = X

' 1. Create unwindowed LPF IRF tensor from invDFT of desired filter PSD - LPFfirf

'ACF & PSD unfiltered sequence
N = UBound(X.Val) - LBound(X.Val) + 1 'length of original sequence
maxlag = N - 1 'set maxlag to full sequence length
ACForig = ACF(X, -maxlag, maxlag)
PSDorig = PSDfromACF(ACForig, None)
imin = LBound(PSDorig.Val)
imax = UBound(PSDorig.Val)
M = imax - imin + 1 'length of ACF and PSD

```

```

'Define filter spectrum with same num of elements as PSDorig
'Initially set all elements to zero
  LPFfreq = PSDorig 'set equal for info
  SampInt = X.Inc(1)
  For i = imin To imax
    LPFfreq.Val(i) = 0
  Next i

'Scale elements in the desired freq range to have amplitude of (1/sampInt).
'This amplitude accounts for sampling interval necessary in convolution.
  CutOffElem = Int(CutOffFreq * SampInt * M)
  For i = 0 To CutOffElem
    LPFfreq.Val(i) = 1 / SampInt 'positive freq elements
    LPFfreq.Val(-i) = 1 / SampInt 'negative freq elements
  Next i

'Estimate LPF IRF as the invDFT of the desired 2 sided PSD.
  LPFfirf = ACFfromPSD(LPFfreq)

' 2. Create windowed LPF IRF. Window is symmetric and the same length as the IRF
'   win = 0 (rectangular)
'   win = 1 (Hanning)
'   win = 2 (Hamming)
'   win = 3 (Blackman)

  imin = LBound(LPFfirf.Val)
  imax = UBound(LPFfirf.Val)
  M = imax - imin + 1
  ReDim winLPFfirf(imin To imax) 'array to hold windowed ACF values

  If win = 1 Then
    'Hanning window
    For i = imin To imax
      winLPFfirf(i) = 0.5 * (1# + Cos(2 * Pi * i / (M - 1))) * LPFfirf.Val(i)
    Next i

  ElseIf win = 2 Then
    'Hamming window
    For i = imin To imax
      winLPFfirf(i) = (0.54 + 0.46 * Cos(2 * Pi * i / (M - 1))) * LPFfirf.Val(i)
    Next i

  ElseIf win = 3 Then
    'Blackman window
    For i = imin To imax
      winLPFfirf(i) = (0.42 + 0.5 * Cos(2 * Pi * i / (M - 2)) + 0.08 * Cos(4 * Pi * i / (M - 2)))
    * LPFfirf.Val(i)
    Next i

  Else
    'Rectangular window

```

```

    For i = imin To imax
        winLPFfir(i) = LPFfir.Val(i)
    Next i

End If

LPFfir.Val = winLPFfir

' 3. Convolve filter IRF with sequence to be filtered
    LPFwind = dtConvolve(LPfFir, X)

End Function

Public Function ModBessI0(X As Double) As Double
' Returns the modified Bessel function I0(x) for any real x
' Uses polynomial approximation from "Numerical Recipes in C" which
' is based on the algorithm in Abramowitz and Stegun.

'Local variable
    Dim ax As Double
    Dim sqrtax As Double
    Dim ans As Double
    Dim Y As Double
    Dim p1, p2 As Double

'Algorithm

    ax = Abs(X)
    If ax < 3.75 Then
        Y = X / 3.75
        Y = Y * Y
        'ans = 1# + y * (3.5156229 + y * (3.0899424 + y * (1.2067492 + y * (0.2659732 + y *
(0.0360768 + y * 0.0045813))))))
        'full expression too complex, so breakdown
        p1 = (0.2659732 + Y * (0.0360768 + Y * 0.0045813))
        p2 = (3.0899424 + Y * (1.2067492 + Y * p1))
        ans = 1# + Y * (3.5156229 + Y * p2)

    Else
        Y = 3.75 / ax
        sqrtax = ax ^ (1 / 2)
        'ans = (Exp(ax) / sqrtax) * (0.39894228 + y * (0.01328592 + y * (0.00225319 + y * (-
0.00157565 + y * (0.00916281 + y * (-0.02057706 + y * (0.02635537 + y * (-0.01647633 + y *
0.00392377))))))))))
        p1 = (-0.02057706 + Y * (0.02635537 + Y * (-0.01647633 + Y * 0.00392377)))
        p2 = (0.00225319 + Y * (-0.00157565 + Y * (0.00916281 + Y * p1)))

```

```
ans = (Exp(ax) / sqrtax) * (0.39894228 + Y * (0.01328592 + Y * p2))
```

```
End If
```

```
ModBessI0 = ans
```

```
End Function
```

```
Public Function LPFkaiser( _  
    Optional Wpassband As Double = 0.15 * 3.141592654, _  
    Optional Wstopband As Double = 0.25 * 3.141592654, _  
    Optional derr As Double = 0.0001, _  
    Optional dT As Double = 0.01 _  
    ) As Tensor
```

```
' Creates a LPF with a Kaiser window. The filter IRF is based on an FRF with an  
' amplitude of 1. When convolved with a sequence, the convolution must be scaled by  
' 1/dT to produce an estimate of a sampled sequence (see Plexus dtConvolve( ) function )  
,
```

```
' The window is created in terms of the normalized frequencies. Passband cutoff of  
' the LPF is the highest freq such that the freq response |H(jw)| > 1-derr, and the  
' stopband freq is defined as the lowest freq such that |H(jw)| < derr.  
,
```

```
' *** It is important to inspect the window before using. It can sometimes oscillate  
' quickly between zero and non-zero values.
```

```
' Local variables
```

```
    Dim Wtrans As Double      'transition region  
    Dim Wco As Double        'cutoff freq of ideal LPF  
    Dim A As Double  
    Dim B As Double          'shape parameter  
    Dim Mtemp As Single  
    Dim M As Integer         'integer length of filter  
    Dim Mhalf As Single  
    Dim Mhalfint As Integer  
    Dim Rmndr As Single  
    Dim N, nmin, nmax As Integer  
    Dim Num As Double  
    Dim NumArg As Double  
    Dim Den As Double
```

```
    Dim ncLPFideal() As Double 'non-causal, ideal LPF IRF  
    Dim LPFwind() As Double   'windowed filter IRF  
    Dim Pi As Double  
    Pi = 4# * Atn(1#)
```

'1. Calculate necessary filter length (M) and shape parameter (B) to achieve desired specs

' Kaiser intermediary terms

Wtrans = Wstopband - Wpassband

A = -20 \* Log(derr) / Log(10) 'Log base 10

' Shape parameter (B)

If A > 50 Then

B = 0.1102 \* (A - 8.7)

ElseIf (A >= 21) And (A <= 50) Then

B = 0.5842 \* (A - 21) ^ 0.4 + 0.07886 \* (A - 21)

Else: B = 0# 'equivalent to a rectangular window

End If

' Filter length (M)

Mtemp = (A - 8) / (2.285 \* Wtrans)

M = Mtemp + 0.6

'If M is even, add term to make odd so IRF is even about n=0

Mhalf = M / 2#

Mhalfint = M / 2#

Rmndr = Mhalf - Mhalfint

If Abs(Rmndr) < 0.01 Then M = M + 1

'2. Define Kaiser window elements from n = -(M-1)/2 to (M-1)/2

' Evaluates modified Bessel with function ModBessI0()

nmin = -(M - 1) / 2

nmax = (M - 1) / 2

ReDim w(nmin To nmax)

If B <> 0 Then ' If B = 0 then window is equivalent to a rectangular window

For N = nmin To nmax

NumArg = B \* (1 - (2 \* N / (M - 1)) ^ 2) ^ (1 / 2)

Num = ModBessI0(NumArg)

Den = ModBessI0(B)

w(N) = Num / Den

Next N

Else 'Rectangular window

For N = nmin To nmax

w(N) = 1

Next N

End If

'3. Define ideal and windowed IRF. Both are created only so a comparison can be made.

'a. Ideal LPF IRF

ReDim ncLPFideal(nmin To nmax)

'ideal cutoff freq

Wco = (Wpassband + Wstopband) / 2

'n = 0

ncLPFideal(0) = Wco / Pi

'n > 0

For N = 1 To nmax

```

ncLPFideal(N) = Sin(N * Wco) / (N * Pi)
Next N
'n < 0
For N = nmin To -1
ncLPFideal(N) = Sin(N * Wco) / (N * Pi)
Next N

```

'b. Windowed LPF IRF

```

ReDim LPFwind(nmin To nmax)
For N = nmin To nmax
LPFwind(N) = ncLPFideal(N) * w(N)
Next N

```

'4. Set output tensor to windowed LPF IRF

```

ReDim LPFkaiser.Val(nmin To nmax)
For N = nmin To nmax
LPFkaiser.Val(N) = LPFwind(N)
Next N
'Scale by 1/dt so summation can be done with dtConvolve() function
LPFkaiser = ScaleTensor(LPFkaiser, (1 / dT))

```

'5. Define LPFkaiser tensor information

```

ReDim LPFkaiser.Rname(1 To 1)
ReDim LPFkaiser.Runit(1 To 1)
ReDim LPFkaiser.Dname(1 To 1)
ReDim LPFkaiser.Dunit(1 To 1)
ReDim LPFkaiser.Inc(1 To 1)
LPFkaiser.Rname(1) = "Amplitude"
LPFkaiser.Runit(1) = ""
LPFkaiser.Dname(1) = "time"
LPFkaiser.Dunit(1) = "s"
LPFkaiser.Inc(1) = dT 'value set in function call
LPFkaiser.Dim = 1

```

End Function





## Appendix D

### VOR experiment conditions

Experiment date: 28 Nov 01								
Experiment			Visual target			Torque perturbation		
trial	len (s)	samp (Hz)	PDF	deg	PSD (Hz)	PRBS (c/s)	amp (N.m)	seed
1	60	100	GWN	25	3.0	14.3	0.20	10
2	60	100	GWN	25	3.0	14.3	0.20	10
3	60	100	PRBS	30	3.0	14.3	0.20	10
4	60	100	uniform	30	3.0	14.3	0.20	10
5	60	100	GWN	30	4.0	14.3	0.20	10
6	60	100	uniform	30	4.0	14.3	0.20	10
7	60	100	uniform	30	4.0	12.5	0.20	10
8	60	100	uniform	35	4.0	20.0	0.20	11
9	60	100	GWN	25	3.0	14.3	0.20	10
10	60	100	GWN	25	3.0	14.3	0.20	10
11	60	100	PRBS	30	3.0	14.3	0.20	10
12	60	100	uniform	30	3.0	14.3	0.20	10
13	60	100	GWN	30	4.0	14.3	0.20	10
14	60	100	uniform	30	4.0	14.3	0.20	10
15	60	100	uniform	30	4.0	12.5	0.20	10
16	60	100	uniform	35	4.0	20.0	0.20	11
note 1: EOG gain 600, subject distance 1550 mm.								
note 2: Calibration trials were run between each tracking trial.								
note 3: Second 8 trials were a repeat of the first 8 trials.								

**Experiment date: 10 Dec 01**

	Experiment		Target		Torque	
trial	len (s)	samp (Hz)	seq #	PSD (Hz)	amp (N.m)	PRBS (c/s)
0	67	100	cal	cal		
1	67	100	cal	cal		
2	60	100	9	1.5	0.00	0.0
3	30	100	12	2	0.25	15.4
4	67	100	cal	cal		
5	60	100	10	1.5	0.25	15.4
6	45	100	13	2	0.00	0.0
7	67	100	cal	cal		

note 1: EOG gain set to 600, subject distance 1530 mm.

note 2: "cal" means calibration trial.

note 3: "seq #" refers to sequence from "sequences" data folder.

**Experiment date: 11 Dec 01 - no torque perturbations**

	Experiment		Target		Torque	
trial	length (s)	samp (Hz)	seq #	PSD (Hz)	amp (N.m)	PRBS (c/s)
0	67	100	cal	cal	0.00	
1	67	100	cal	cal	0.00	
2	120	100	9	1.5	0.00	
3	60	100	12	2.0	0.00	
4	67	100	cal	cal	0.00	
5	120	100	10	1.5	0.00	
6	90	100	13	2.0	0.00	
7	67	100	cal	cal	0.00	
8	60	200	15	3.5	0.00	
9	60	200	16	3.5	0.00	
10	67	200	cal	cal	0.00	

note 1: EOG gain set to 600, subject distance 1560

note 2: sequences from "sequences" data folder

**Experiment date: 16 Dec 01 - no torque perturbations**

trial	Experiment		Target		Torque	
	length (s)	samp (Hz)	seq #	PSD (Hz)	amp (N.m)	PRBS (c/s)
1	65	100	cal	cal	0.00	
2	60	100	12	2.0	0.00	
3	120	100	9	1.5	0.00	
4	65	100	cal	cal	0.00	
5	60	100	16	3.5	0.00	
6	90	100	14	3.0	0.00	
7	65	100	cal	cal	0.00	
8	90	100	13	2	0.00	
9	60	100	15	3.5	0.00	
11	120	100	11	2.5	0.00	
12	65	100	cal	cal	0.00	
13	90	100	17	3	0.00	

note 1: trial #10 was not sampled properly

note 2: EOG gain 600, subject distance 1550 mm

note 3: sequences from "sequences" data folder

**Experiment date: 18 Jan 02**

trial	Experiment		Target		Torque	
	length (s)	samp (Hz)	seq #	PSD (Hz)	amp (N.m)	PRBS (c/s)
1	60	100	cal	cal		
3	60	100	12	2.0	0.00	0.0
4	60	100	cal	cal		
5	60	100	12	2.0	0.20	14.3
6	60	100	cal	cal		
7	60	100	16	3.5	0.20	14.3
8	60	100	cal	cal		
9	60	100	16	3.5	0.20	14.3
10	60	100	16	3.5	0.20	14.3
12	95	100	13	2.0	0.00	0.0
13	60	100	cal	cal		
14	95	100	13	2.0	0.25	14.3

note 1: trial #2, #11 were not sampled properly

note 2: EOG gain 200, subject distance 1550 mm

note 3: sequences from "sequences" data folder

**Experiment date: 24 Jan 02**

	Experiment		Target		Torque	
trial	length (s)	samp (Hz)	seq #	PSD (Hz)	amp (N.m)	PRBS (c/s)
1	50	200	cal	cal		
2	60	200	12	2.0	0.00	0.0
3	50	200	cal	cal		
4	60	200	12	2.0	0.24	15.4
5	50	200	cal	cal		
6	60	200	16	3.5	0.00	0.0
7	50	200	cal	cal		
8	60	200	16	3.5	0.20	15.4
9	50	200	cal	cal		
10	90	200	13	2.0	0.20	15.4
11	50	200	cal	cal		
12	90	200	13	2.0	0.00	0.0
13	50	200	cal	cal		

note 1: EOG gain 600, subject distance 1500 mm

note 2: sequences from "sequences" data folder

**Experiment date: 12 Mar 02**

trial	Experiment			Target		Torque	
	length (s)	samp (Hz)	EOG gain	seq #	PSD (Hz)	amp (N.m)	PRBS (c/s)
1	40	200	600	cal	cal		
2	45	200	600	30	2.0	0.00	0.0
3	40	200	600	cal	cal		
4	45	200	600	30	2.0	0.25	15.4
5	40	200	600	cal	cal		
6	45	200	600	31	2.0	0.00	0.0
7	40	200	600	cal	cal		
8	45	200	600	31	2.0	0.20	20.0
9	40	200	200	cal	cal		
10	45	200	200	6	2.5	0.00	0.0
11	40	200	200	cal	cal		
12	45	200	200	6	2.5	0.25	15.4
13	40	200	200	cal	cal		
14	45	200	200	30	2.0	0.00	0.0
15	40	200	200	cal	cal		
16	45	200	200	30	2.0	0.25	15.4
17	40	200	200	cal	cal		
18	30	200	200	sin	0.5 Hz	0.00	0.0
19	30	200	200	sin	0.5 Hz	0.25	15.4

note 1: subject distance 1530 mm

note 2: sequences from "seq45" data folder

**Experiment date: 15 Mar 02**

trial	Experiment			Target		Torque	
	length (s)	samp (Hz)	EOG gain	seq #	PSD (Hz)	amp (N.m)	PRBS (c/s)
0	40	100	200	cal	cal	0.20	10.0
1	40	100	200	cal	cal		
2	45	100	200	30	2.0	0.00	0.0
3	40	100	200	cal	cal		
4	45	100	200	30	2.0	0.25	15.4
5	40	100	200	cal	cal		
6	45	100	200	32	2.0	0.25	10.0
7	40	100	200	cal	cal		
10	40	100	200	cal	cal		
11	45	100	200	36	2.0	0.00	0.0
12	40	100	200	cal	cal		
13	45	100	200	36	2.0	0.25	10.0
14	40	100	200	cal	cal		
15	45	100	200	2	2.0	0.25	10.0
16	40	100	200	cal	cal		
17	45	100	200	2	2.0	0.00	0.0
18	30	100	200	sin	0.5 Hz	0.00	0.0
19	30	100	200	sin	0.5 Hz	0.25	10.0

note 1: subject distance 1530 mm

note 2: sequences from "seq45" data folder

**Experiment date: 18 Apr 02**

trial	Experiment			Target		Torque	
	length (s)	samp (Hz)	EOG gain	seq #	PSD (Hz)	amp (N.m)	PRBS (c/s)
1	50	100	800	cal	cal		
2	100	100	800	30	1.0	0.00	0.0
3	50	100	800	cal	cal		
4	100	100	800	6	1.0	0.20	15.4
5	50	100	800	cal	cal		
6	100	100	800	3b	1.0	0.22	15.4
7	50	100	800	cal	cal		
8	100	100	800	6	1.0	0.00	0.0
9	50	100	800	cal	cal		
10	100	100	800	1	1.0	0.22	15.4
11	50	100	800	cal	cal		
12	100	100	800	3b	1.0	0.00	0.0

note 1: subject distance 1550 mm

note 2: sequences from "seq90" data folder

**Experiment date: 31 May 02**

	Experiment			Target		Torque	
trial	length (s)	samp (Hz)	EOG gain	seq #	PSD (Hz)	amp (N.m)	PRBS (c/s)
3	45	100	600	cal	cal		
4	95	100	200	22	1.7	0.22	10.0
5	45	100	200	cal	cal		
6	95	100	200	6	1.0	0.00	0.0
7	45	100	200	cal	cal		
8	95	100	200	15	2.0	0.22	15.4
9	45	100	200	cal	cal		
10	95	100	200	16	1.2	0.00	0.0
11	45	100	200	cal	cal		
12	95	100	600	20	1.4	0.22	15.4
13	45	100	600	cal	cal		
14	95	100	600	22	3.0	0.00	0.0
15	45	100	600	cal	cal		
16	95	100	600	6	1.0	0.22	15.4
17	45	100	600	cal	cal		
18	95	100	600	15	2.0	0.00	0.0
19	45	100	600	cal	cal		
20	95	100	600	16	1.2	0.22	15.4
21	45	100	600	cal	cal		
22	30	100	600	stat	0.0	0.22	15.4
23	30	100	600	stat	0.0	0.22	10.0
25	45	100	600	cal	cal		
26	30	100	600	20	1.4	0.24	15.4

note 1: subject distance 1550 mm

note 2: sequences from "seq90" data folder

note 3: "stat" means that target was stationary

**Experiment date: 14 Jun 02**

trial	Experiment			Target		Torque	
	length (s)	samp (Hz)	EOG gain	seq #	PSD (Hz)	amp (N.m)	PRBS (c/s)
1	45	100	600	cal	cal		
2	95	100	600	29	1.5	0.00	0.0
3	45	100	600	cal	cal		
4	95	100	600	30	1.8	0.30	10.0
5	45	100	600	cal	cal		
6	95	100	600	26	2.0	0.30	10.0
7	45	100	600	cal	cal		
8	95	100	600	24	1.5	0.00	0.0
9	45	100	600	cal	cal		
10	95	100	600	25	1.7	0.30	10.0
11	45	100	600	cal	cal		
12	95	100	600	23	1.2	0.30	10.0
13	45	100	600	cal	cal		
14	95	100	600	29	1.5	0.30	10.0
15	45	100	200	cal	cal		
16	95	100	200	30	1.8	0.00	0.0
17	45	100	200	cal	cal		
18	95	100	200	24	1.5	0.30	10.0
19	45	100	200	cal	cal		
20	95	100	200	26	2.0	0.30	10.0
21	45	100	200	cal	cal		
22	95	100	200	25	1.7	0.00	0.0
23	45	100	200	cal	cal		
24	95	100	200	24	1.5	0.30	10.0
25	45	100	200	cal	cal		
26	30	100	200	stat	0	0.30	10.0
27	45	100	200	cal	cal		
28	30	100	200	stat	0.0	0.30	10.0
29	45	100	200	cal	cal		
30	30	100	200	stat	0.0	0.25	10.0

note 1: subject distance 1520 mm

note 2: sequences from "seq90" data folder

note 3: "stat" means that target was stationary



**Experiment date: 04 Mar 03**

trial	Experiment			Target		Torque	
	length (s)	samp (Hz)	EOG gain	seq #	PSD (Hz)	amp (N.m)	PRBS (c/s)
1	45	100	600	cal	cal		
2	95	100	600	34	1.8	0.35	15.4
3	45	100	600	cal	cal		
4	95	100	600	26	2.0	0.32	10.0
5	45	100	600	cal	cal		
6	95	100	600	29	1.5	0.32	10.0
7	45	100	600	cal	cal		
8	125	100	600	2	1.8	0.31	9.1
9	45	100	600	cal	cal		
10	95	100	600	31	2.0	0.32	15.4
11	45	100	600	cal	cal		
12	95	100	600	30	1.8	0.33	10.0
13	45	100	600	cal	cal		
14	125	100	600	1	1.8	0.33	9.1
15	45	100	600	cal	cal		
16	95	100	600	24	1.5	0.30	10.0
17	45	100	600	cal	cal		
18	95	100	600	stat	0.0	0.30	10.0

note 1: subject distance 1540 mm

note 2: seqs 1 & 2 from "seq120" data folder, others from "seqs90"



## **Appendix E**

### **Analysis programs**

- E.1 MakeFiles
- E.2 G2 G4 Velocity
- E.3 G3 FRF IRF
- E.4 UserDefinedFunctions

The analysis programs were written in Mathcad 2001i (MathSoft, Inc., Cambridge, MA).



## MakeFiles

Program creates processed position and velocity data files for further analysis in Matlab and Mathcad.

filename: MakeFiles 04Mar03 16b (seq 24)

---

### Description

1. Process raw data into sequences for analysis
2. Inspect sequences, spectra, correlations (use radians)
3. Create Mathcad .mcd files

### A. User defined functions

5 point equal-abscissa derivative function:  $\text{Deriv5pt}(X, \Delta t)$



Trapezoidal integration:  $\text{Pos}(\text{vel}, \Delta t)$



Watson sensor 14150 calibration correction:  $\text{HVcorr14150}(hv)$



Noncausal LPF with Kaiser window:  $\text{ncLPFkaiser}(\Delta t, f_{pb}, f_{sb}, \delta)$



Noncausal convolution with IRF symmetric about 0:  $\text{ncSymDtConvolve}(h, x, \Delta t)$



Mean square difference:  $\text{MSD}(x, y)$



B. Raw data signals - tp, ep, hv

1. Raw data, DAQ gains, sensor scales

folder: 04Mar03

data set: 16

target: Seqs90 / 24 ; PSD 1.5 Hz,  $\tau_{\max}$  9.0 s, length 90s, exec 200 Hz

torque: PRBS, 0.3 v

DAQ @ 100 Hz for 95 sec

*Select files, set EOG calibration, apply median and low pass filters, remove transients.*

**admin file (adm)**

adm :=

	0
0	96
1	100
2	20
3	5
4	20
5	2
6	0
7	0

adm<sub>0</sub> = exper length

adm<sub>1</sub> = f<sub>samp</sub>

adm<sub>2</sub> = ch 0 gain

adm<sub>3</sub> = ch 1 gain

adm<sub>4</sub> = ch 2 gain

adm<sub>5</sub> = ch 3 gain

adm<sub>6</sub> = DC voltage

adm<sub>7</sub> = PRBS voltage

**eye position (ep)**

ep :=

  
C:\..\EOG16fix.dat

**head velocity (hv)**

hv :=

  
C:\..\Head16fix.dat

**target position (tp)**

tp :=

  
C:\..\Targ16fix.dat

**torque (tor)**

tor :=


  
C:\..\Torq16fix.dat

file length: N := length(ep)    N = 9.601 × 10<sup>3</sup>    indices    i := 0..N - 1

sampling frequency (Hz)    f<sub>samp</sub> := adm<sub>1</sub>    f<sub>samp</sub> = 100


sampling period     $\Delta t := \frac{1}{f_{\text{samp}}}$      $\Delta t = 0.01$

2. Inspect raw signals before scaling by gains




$$\text{mean}(ep) = 0.08623$$

$$\text{var}(ep) = 1.72352$$




$$\text{mean}(hv) = -2.59273$$

$$\text{var}(hv) = 13.22132$$



$$\text{mean}(tp) = -0.02673$$

$$\text{var}(tp) = 6.74961$$



$$\text{mean}(tor) = -0.39098$$

$$\text{var}(tor) = 26.92009$$





Target position: initially make radians  $tp := \left( \frac{g_{\text{targ}}}{k_{\text{targ}}} \right) \cdot tp$

convert target angle (radians) into angle (degrees) viewed by subject

mirror distance  $md := 1690$  subject distance  $sd := 1540$

radians  $tp := \text{atan}\left(\frac{md}{sd} \cdot \tan(tp)\right)$  degrees  $tp := \frac{180}{\pi} \cdot tp$

Torque: convert to N.m  $tor := \left( \frac{g_{\text{tor}}}{k_{\text{tor}}} \right) \cdot tor$

#### 4. Median filter

Use median filter to remove aberrant data, not as a substitute for a LPF. Too large a neighborhood changes the sequence shape. Median filters are nonlinear, so do not want to abuse.

set median filter neighborhood:  $nh := 5$  equates to  $nh \cdot \Delta t = 0.05$  seconds of data

$ep_{\text{fil}} := \text{medsmooth}(ep, nh)$   $hv_{\text{fil}} := \text{medsmooth}(hv, nh)$

$tp_{\text{fil}} := \text{medsmooth}(tp, nh)$   $tor_{\text{fil}} := \text{medsmooth}(tor, nh)$

*How does the filter change the PSD?*  $N = 9.601 \times 10^3$   $nr := 10$   $ol := 0.5$   $win := 6$

$S_{epf} := \text{pspectrum}(ep_{\text{fil}}, nr, ol, win)$   $S_{ep} := \text{pspectrum}(ep, nr, ol, win)$   $K := \text{last}(S_{epf})$   $k := 0.. \frac{K}{5}$





$S_{hvf} := \text{pspectrum}(hv_{\text{fil}}, nr, ol, win)$   $S_{hv} := \text{pspectrum}(hv, nr, ol, win)$





$S_{tpf} := \text{pspectrum}(tp_{\text{fil}}, nr, ol, win)$   $S_{tp} := \text{pspectrum}(tp, nr, ol, win)$

$S_{torf} := \text{pspectrum}(tor_{\text{fil}}, nr, ol, win)$   $S_{tor} := \text{pspectrum}(tor, nr, ol, win)$

$n := 1600..2000$

Effect of median filter:  $nh = 5$  ,  $nh \cdot \Delta t = 0.05$





Set sequences to median filtered version       $ep := ep_{fil}$        $hv := hv_{fil}$        $tp := tp_{fil}$        $tor := tor_{fil}$


### 5. LPF with Kaiser window

It is especially important to filter the eye position before differentiating the signal into eye velocity. Since linear relations are being made, all the signals must be filtered.

set Kaiser window LP filter settings:  $f_{pb} := 10$        $f_{sb} := 12$        $\delta := 0.000001$

generate filter IRF:  $h_{lpf} := \text{ncLPFkaiser}(\Delta t, f_{pb}, f_{sb}, \delta)$

$M := \text{length}(h_{lpf})$        $M = 196$        $m := 0.. \frac{M}{2}$



convolve acausal filter with sequences

$$ep_{fil} := \frac{1}{\Delta t} \cdot \text{ncSymDtConvolve}(h_{lpf}, ep, \Delta t) \qquad hv_{fil} := \frac{1}{\Delta t} \cdot \text{ncSymDtConvolve}(h_{lpf}, hv, \Delta t)$$

$$tp_{fil} := \frac{1}{\Delta t} \cdot \text{ncSymDtConvolve}(h_{lpf}, tp, \Delta t) \qquad tor_{fil} := \frac{1}{\Delta t} \cdot \text{ncSymDtConvolve}(h_{lpf}, tor, \Delta t)$$

calculate eye velocities:  $eva_{fil} := \text{Deriv5pt}(ep_{fil}, \Delta t)$

$eva := \text{Deriv5pt}(ep, \Delta t)$

How does the LPF change the PSD?  $N = 9.601 \times 10^3$   $nr := 10$   $ol := 0.5$   $win := 6$

Spectral densities

$S_{epf} := \text{pspectrum}(ep_{fil}, nr, ol, win)$   $S_{ep} := \text{pspectrum}(ep, nr, ol, win)$   $K := \text{last}(S_{epf})$   $k := 0.. \frac{K}{5}$

$S_{evf} := \text{pspectrum}(eva_{fil}, nr, ol, win)$   $S_{ev} := \text{pspectrum}(eva, nr, ol, win)$



$S_{hvf} := \text{pspectrum}(hv_{fil}, nr, ol, win)$   $S_{hv} := \text{pspectrum}(hv, nr, ol, win)$





$S_{tpf} := \text{pspectrum}(tp_{fil}, nr, ol, win)$   $S_{tp} := \text{pspectrum}(tp, nr, ol, win)$





$S_{torf} := \text{pspectrum}(tor_{fil}, nr, ol, win)$   $S_{tor} := \text{pspectrum}(tor, nr, ol, win)$

set plot indices:  $A1 := 1800$   $n := A1.. A1 + 150$

Effect of LPF:  $f_{pb} = 10$  ,  $f_{sb} = 12$  ,  $\delta = 1 \times 10^{-6}$







Set EP, HV, & TP to the LPF version

$ep := epfil$

$hv := hvfil$

$tp := tpfil$

$tor := torfil$

6. Remove transient data points at beginning and end of sequence that may be due to starting/stopping the trial

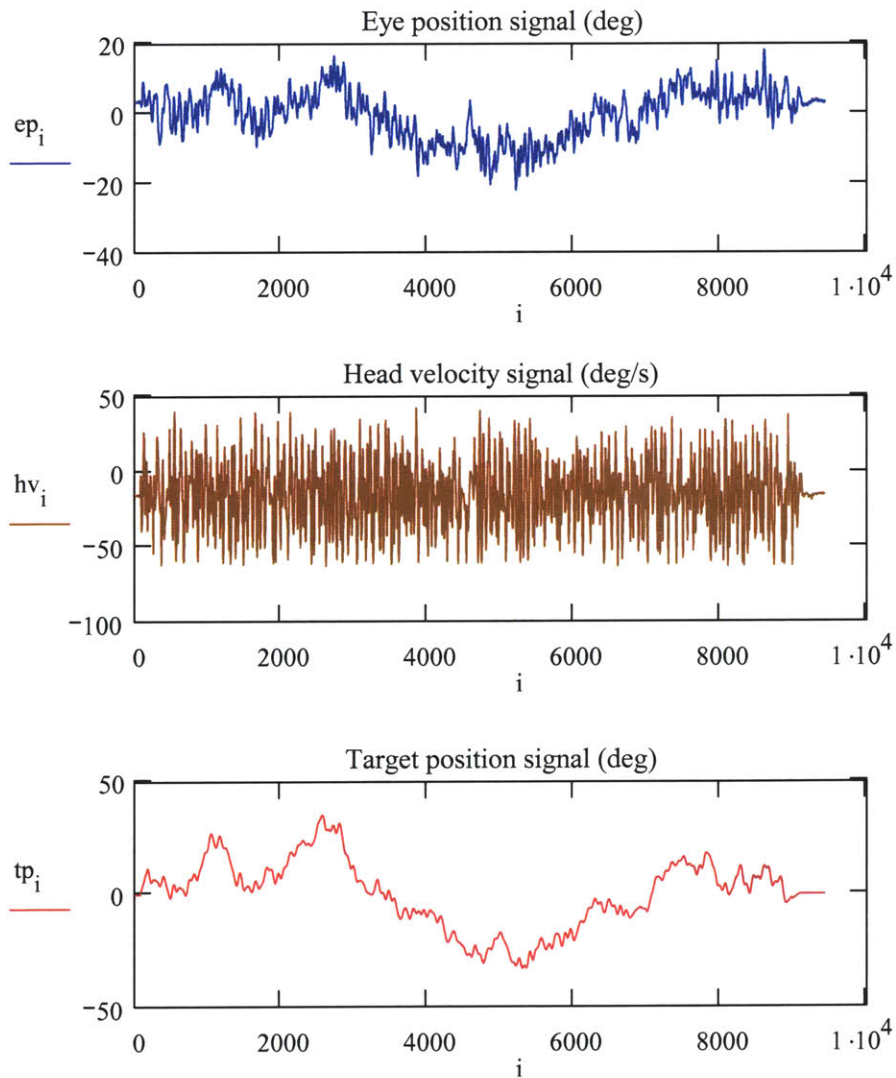
identify data points to be removed:      beg := 10      end := 160

ep := submatrix(ep, beg, N - 1 - end, 0, 0)      hv := submatrix(hv, beg, N - 1 - end, 0, 0)

tp := submatrix(tp, beg, N - 1 - end, 0, 0)      tor := submatrix(tor, beg, N - 1 - end, 0, 0)

reset indices      N := length(hv)      N =  $9.431 \times 10^3$       i := 0..N - 1

Sequences with proper units; offset and drift not yet removed






C. Remove drift and offset from hv and ep, then calculate head and gaze positions

The target signal is position feedback from a galvanometer and should not have appreciable drift, but may have an offset. This offset can be removed by identifying zero position regions and averaging the signal in these regions.

Although the Watson sensor claims to be a "dc" sensor, it does drift. Watson reports a drift of 0.001 deg/s per sec. The variance of the signal when measured at rest was 0.032, and the sensor appears to have a slight linear drift. The protocol begins and ends at zero position, so the mean head velocity should be zero in these regions. Consider fitting line to, and removing from the velocity, assuming a linear drift between beginning and end regions, or setting mean to zero.

Can expect the EOG sensor to drift, hopefully only a little, during the experiment. Despite careful referencing, using a differential measurement, reducing external light, and letting the subject's eyes acclimate to the environment, the signal will drift unless processed by a HPF, which would affect the dc data. Ideally the drift should be calculated and removed from the measurements. However, it is difficult during experimentation to force the eyes and head to a reference position, from which the drift can be calculated, without interfering with the subjects natural use of the head and eyes. Testing with the eyes in one place has shown the drift to be slow and small.

1. Target position: Estimate the bias as the mean of the reference regions.



identify regions of zero position - use "track data" of trace

region 1: a1 := 0            b1 := 57

region 2: a2 := 9134        b2 := N - 1

zp1 := submatrix(tp, a1, b1, 0, 0)            mean(zp1) = -0.08921

zp2 := submatrix(tp, a2, b2, 0, 0)            mean(zp2) = -0.08163

define bias as ave of regions:

Targbias := mean(zp1, zp2)    Targbias = -0.08287

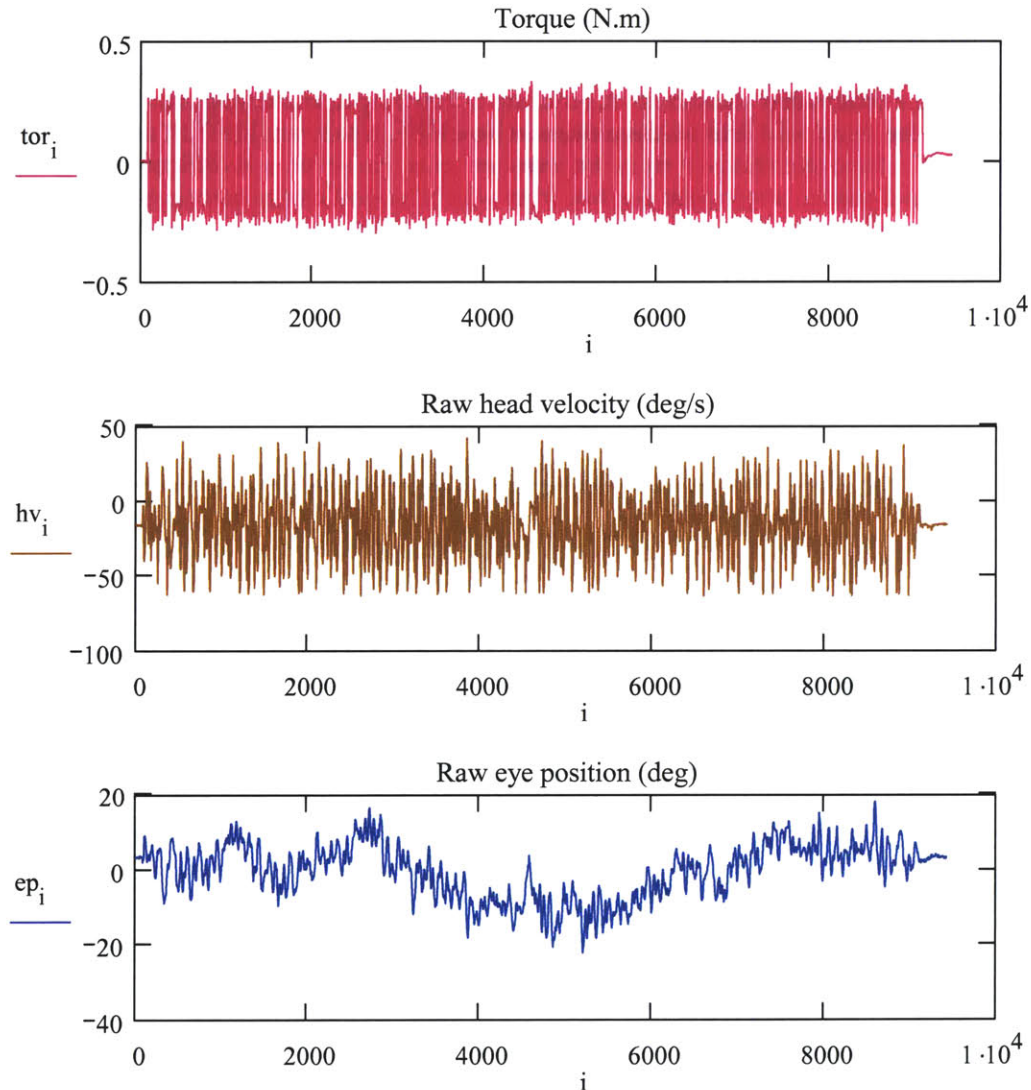
remove bias from target:

tp := tp - Targbias

## 2. Head velocity and eye position; calculate head position and gaze position

The subject is instructed to bring head and eyes to rest with zero displacement and zero velocity during the reference regions. Decide on drift and offset under the assumption that the head is at rest and the gaze is at zero during the reference regions.

Inspect the torque, head, and eye sequences



a. Consider the regions of head and eye data that represent when gaze was upon the target zero

regions set from target position       $b1 = 57$        $a2 = 9.134 \times 10^3$

reset b1 and a2:       $b1 := 57$        $a2 := 9200$

Do the head and eye signals drift from beginning to end ?

Head velocity    hv1 := submatrix(hv, a1, b1, 0, 0)     $\mu_{hv1} := \text{mean}(hv1)$      $\mu_{hv1} = -16.00356$   
                      hv2 := submatrix(hv, a2, b2, 0, 0)     $\mu_{hv2} := \text{mean}(hv2)$      $\mu_{hv2} = -16.47485$   
                      hv ave "drift" per second:  $\frac{\mu_{hv2} - \mu_{hv1}}{N \cdot \Delta t} = -4.99721 \times 10^{-3}$

Probably safe to assume that the head velocity is nearly zero at the beginning and the end of the trial. Harder to make the same assumption about head position. In this case there seems to be a constant decrease in the velocity, which indicates a constant acceleration, which probably is not the case.

Eye position       ep1 := submatrix(ep, a1, b1, 0, 0)        $\mu_{ep1} := \text{mean}(ep1)$         $\mu_{ep1} = 3.36547$   
                      ep2 := submatrix(ep, a2, b2, 0, 0)        $\mu_{ep2} := \text{mean}(ep2)$         $\mu_{ep2} = 3.20756$   
                      ep ave "drift" per second:  $\frac{\mu_{ep2} - \mu_{ep1}}{N \cdot \Delta t} = -1.67442 \times 10^{-3}$

b. Torque - estimate bias from the zero reference regions set for the head and eye zero positions.

a1 = 0

tor1 := submatrix(tor, a1, b1, 0, 0)        $\text{mean}(tor1) = 2.08451 \times 10^{-3}$   
 tor2 := submatrix(tor, a2, b2, 0, 0)        $\text{mean}(tor2) = 0.02735$

define bias as ave of regions:     $\text{tor}_{\text{bias}} := \text{mean}(tor1, tor2)$      $\text{tor}_{\text{bias}} = 0.02228$

remove bias from target:        $\text{tor} := \text{tor} - \text{tor}_{\text{bias}}$      $\mu_{\text{tor}} := \text{mean}(\text{tor})$      $\mu_{\text{tor}} = -6.78183 \times 10^{-3}$

c. Remove head velocity drift

- a. Fit line between the offsets at beginning and end; remove. Does not give data a zero mean
- b. Fit line to all data. This essentially makes the data zero mean which may not be the case, but is not a bad approximation because the head ends the experiment at nearly the same position as it began.
- c. Zero mean velocity data




c. (b). Fit line to head velocity data, remove.

define vector of x points:  $x_i := i$

least sq line :  $\phi_{hv} := \text{line}(x, hv)$   $\phi_{hv} = \begin{pmatrix} -15.6584 \\ -2.4389 \times 10^{-5} \end{pmatrix}$   $hvl_{i} := \phi_{hv_1} \cdot x_i + \phi_{hv_0}$


remove line from data:  $hvb := hv - hvl_{i}$   $\mu_{hvb} := \text{mean}(hvb)$   $\lambda_{hvb} := \text{var}(hvb)$

calculate head position:  $hpb := \text{Pos}(hvb, \Delta t)$



$$\mu_{hvb} = 1.07681 \times 10^{-15}$$

$$\lambda_{hvb} = 474.13594$$



Resulting head position (b):

1. Does it move in synch with the target ?
2. Do the positions at the beginning and end make sense?


$$hp1 := \text{submatrix}(hpb, a1, b1, 0, 0) \quad \mu_{hpb1} := \text{mean}(hp1) \quad \mu_{hpb1} = -0.09596$$

$$hp2 := \text{submatrix}(hpb, a2, b2, 0, 0) \quad \mu_{hpb2} := \text{mean}(hp2) \quad \mu_{hpb2} = 0.20958$$

c. (c). Zero mean head velocity data


remove mean from data:  $\mu_{hv} := \text{mean}(hv)$     $hv_c := hv - \mu_{hv}$     $\mu_{hvc} := \text{mean}(hvc)$     $\lambda_{hvc} := \text{var}(hvc)$

calculate head position:  $hpc := \text{Pos}(hvc, \Delta t)$



$$\mu_{hvc} = 0$$

$$\lambda_{hvc} = 474.14035$$



Resulting head position (c):

1. Does it move in synch with the target ?
2. Do the positions at the beginning and end make sense?


$$hpc1 := \text{submatrix}(hpc, a1, b1, 0, 0) \quad \mu_{hpc1} := \text{mean}(hpc1) \quad \mu_{hpc1} = -0.06332$$

$$hpc2 := \text{submatrix}(hpc, a2, b2, 0, 0) \quad \mu_{hpc2} := \text{mean}(hpc2) \quad \mu_{hpc2} = 0.33966$$






d (b). Remove first offset, assume no drift.


Remove initial offset from data:  $ep_b := ep - \text{mean}(\mu_{ep1}, \mu_{ep2})$



3. Create gaze position - decide which eye-head combination makes sense







### Sequences

$hp_a$  - assumed linear drift between offsets at beginning and end of velocity sequence



$hp_b$  - line fitted to, then removed from velocity data

$hp_c$  - velocity given zero mean

$ep_a$  - drift estimated as line fit between offsets at zero regions


$ep_b$  - offset estimated as mean of zero regions. Assumes no moving drift.

**Create gaze:**     $gp := hp_c + ep_b$              $MSD_{gp} := MSD(gp, tp)$              $MSD_{gp} = 4.71267$   
                           $ge := tp - gp$              $mean(ge) = 1.2222$              $stdev(ge) = 1.79413$




4. Define the sequences to be used


measured head velocity	$hv := hv_c$	
head position:	$hp := hp_c$	
measured eye position:	$ep := ep_b$	
gaze position:	$gp := ep + hp$	
gaze position error:	$ge := tp - gp$	
	$\mu_{ge} := \text{mean}(ge)$	$\mu_{ge} = 1.2222$
	$\sigma_{ge} := \text{stdev}(ge)$	$\sigma_{ge} = 1.79413$
head velocity statistics	$\mu_{hv} := \text{mean}(hv)$	$\mu_{hv} = 0$
	$\sigma_{hv} := \text{stdev}(hv)$	$\sigma_{hv} = 30.48467$
	$hv_{\max} := \text{max}(hv)$	$hv_{\max} = 80.74318$
	$hv_{\min} := \text{min}(hv)$	$hv_{\min} = -67.10871$
head position statistics	$\mu_{hp} := \text{mean}(hp)$	$\mu_{hp} = 2.52671$
	$\sigma_{hp} := \text{stdev}(hp)$	$\sigma_{hp} = 10.25277$
	$hp_{\max} := \text{max}(hp)$	$hp_{\max} = 28.32808$
	$hp_{\min} := \text{min}(hp)$	$hp_{\min} = -26.35789$
eye position statistics	$\mu_{ep} := \text{mean}(ep)$	$\mu_{ep} = -3.8115$
	$\sigma_{ep} := \text{stdev}(ep)$	$\sigma_{ep} = 7.18542$
	$ep_{\max} := \text{max}(ep)$	$ep_{\max} = 14.79346$
	$ep_{\min} := \text{min}(ep)$	$ep_{\min} = -25.5895$




$\mu_{hv} = 0$                        $\sigma_{hv} = 30.48467$   
 $\mu_{hp} = 2.52671$                  $\sigma_{hp} = 10.25277$




$hv_{max} = 80.74318$              $hv_{min} = -67.10871$   
 $hp_{max} = 28.32808$              $hp_{min} = -26.35789$



$\mu_{ep} = -3.8115$                        $\sigma_{ep} = 7.18542$





$ep_{max} = 14.79346$                  $ep_{min} = -25.5895$





$MSD_{gp} = 4.71267$   
 $\mu_{ge} = 1.2222$   
 $\sigma_{ge} = 1.79413$

D. Calculate velocities - 5 pt derivative followed by median filter. Also consider effects of mean filter.


Set elements of mov ave    elem := 10    elem·Δt = 0.1    j := 1400..2000



$\sigma_{tv} = 9.1861$      $tv_{\max} = 23.74911$      $tv_{\min} = -26.6056$   
 $\sigma_{gv} = 18.93931$      $gv_{\max} = 115.7328$      $gv_{\min} = -95.19488$





$\mu_{gs} = -1.10602 \times 10^{-3}$      $\sigma_{gs} = 17.58891$      $gs_{\max} = 92.4133$      $gs_{\min} = -115.59538$





$\sigma_{hv} = 30.48467$   
 $hv_{\max} = 80.74318$   
 $hv_{\min} = -67.10871$   
 $\sigma_{ev} = 30.66935$   
 $ev_{\max} = 114.25608$   
 $ev_{\min} = -116.53452$

Velocity after running through moving average filter;  $\text{elem} \cdot \Delta t = 0.1$




$\sigma_{\text{tvm}} = 9.06255$        $\text{tvm}_{\text{max}} = 23.36072$        $\text{tvm}_{\text{min}} = -26.09157$

$\sigma_{\text{gvm}} = 13.65113$        $\text{gvm}_{\text{max}} = 74.0963$        $\text{gvm}_{\text{min}} = -54.30527$



$\mu_{\text{gsm}} = 1.83752 \times 10^{-3}$        $\sigma_{\text{gsm}} = 11.74658$        $\text{gsm}_{\text{max}} = 73.34833$        $\text{gsm}_{\text{min}} = -48.45093$



$\sigma_{\text{hvm}} = 28.0186$   
 $\text{hvm}_{\text{max}} = 73.85167$   
 $\text{hvm}_{\text{min}} = -62.3822$









$\sigma_{\text{evm}} = 26.19748$   
 $\text{evm}_{\text{max}} = 92.41708$   
 $\text{evm}_{\text{min}} = -84.56116$

E. Write files for urther analysis in Mathcad

Sequence information

sampling period	SeqInfo <sub>0</sub> := $\Delta t$		
LP filter cutoff	SeqInfo <sub>1</sub> := $\frac{f_{pb} + f_{sb}}{2}$	SeqInfo <sub>1</sub> = 11	
	SeqInfo <sub>2</sub> := $f_{pb}$	SeqInfo <sub>3</sub> := $f_{sb}$	SeqInfo <sub>4</sub> := $\delta$

Mathcad files (.csv) - folder d:/ VOR / Proc data

 C:\..\tor_16b.csv	 C:\..\tp_16b.csv	 C:\..\lep_16b.csv	 C:\..\hp_16b.csv
tor	tp	ep	hp
 C:\..\SeqInfo_16b.csv	 C:\..\tv_16b.csv	 C:\..\lev_16b.csv	 C:\..\hv_16b.csv
SeqInfo	tv	ev	hv

## G2 G4 Velocity

filename: G2 G4 Vel 04May03 2b 8b  
data filtered with 10 Hz LPF

last update: 01Apr 03

---

Description: The worksheet estimates the IRF and FRF for the head's visually mediated response (g2) and response of the head-neck to perturbations (g4).

### A. User defined functions











### B. Read processed data files, create gaze position and velocity

Read from file: c:\VOR\Proc data\04May03\2b (1.8 Hz) 8b (1.8 Hz)



sequence information: SeqInfo:  $\Delta t$  (s);  $f_{co}$ ,  $f_{pb}$ ,  $f_{sb}$  (Hz)



#### 1. Read processed data from .csv file. Data created by MakeFiles\_04May03 2b & 8b



##### Trial 2b



ep :=	 C:\..\lep_2b.csv	ev :=	 C:\..\lev_2b.csv	last(ep) = $9.46 \times 10^3$
hp :=	 C:\..\hlp_2b.csv	hv :=	 C:\..\hvv_2b.csv	last(hv) = $9.46 \times 10^3$
SeqInfo :=	 C:\..\SeqInfo_2b.csv	pt :=	 C:\..\tpor_2b.csv	last(pt) = $9.46 \times 10^3$
tp :=	 C:\..\tpp_2b.csv	tv :=	 C:\..\tvv_2b.csv	last(tv) = $9.46 \times 10^3$

Trial 8b

ep1 :=  C:\..\lep\_8b.csv      ev1 :=  C:\..\lev\_8b.csv      last(ep1) =  $1.247 \times 10^4$

hp1 :=  C:\..\hlp\_8b.csv      hv1 :=  C:\..\hvv\_8b.csv      last(hv1) =  $1.247 \times 10^4$

SeqInfo1 :=  C:\..\SeqInfo\_8b.csv      pt1 :=  C:\..\tpor\_8b.csv      last(pt1) =  $1.247 \times 10^4$

tp1 :=  C:\..\tpp\_8b.csv      tv1 :=  C:\..\tvv\_8b.csv      last(tv1) =  $1.247 \times 10^4$

stack data into single array:      ep := stack(ep, ep1)      ev := stack(ev, ev1)  
      hp := stack(hp, hp1)      hv := stack(hv, hv1)  
      tp := stack(tp, tp1)      tv := stack(tv, tv1)  
      pt := stack(pt, pt1)

clear out unnecessary arrays:      ep1 := 0      tp1 := 0      hp1 := 0      pt1 := 0  
      ev1 := 0      tv1 := 0      hv1 := 0

sequence length and indices:      N := length(tp)      N = 2.193 × 10<sup>4</sup>      i := 0..N - 1



### sequence sampling and LP filtering

sampling period:	$\Delta t := \text{SeqInfo0}$	$\Delta t = 0.01$	$f_{\text{samp}} := \frac{1}{\Delta t}$
LPF cutoff:	$f_{\text{co}} := \text{SeqInfo1}$	$f_{\text{co}} = 11$	
LPF passband:	$f_{\text{pb}} := \text{SeqInfo2}$	$f_{\text{pb}} = 10$	
LPF stopband:	$f_{\text{sb}} := \text{SeqInfo3}$	$f_{\text{sb}} = 12$	
LPF overshoot:	$\delta := \text{SeqInfo4}$	$\delta = 1 \times 10^{-6}$	

### 2. Create gaze position, gaze error, gaze velocity, gaze slip

gaze position and gaze error:	$gp := ep + hp$	$ge := tp - gp$
gaze velocity and slip:	$gv := ev + hv$	$gs := tv - gv$

### 3. Means and standard deviations

$\sigma_{pt} := \text{stdev}(pt)$	$\mu_{pt} := \text{mean}(pt)$		
$\sigma_{tp} := \text{stdev}(tp)$	$\mu_{tp} := \text{mean}(tp)$	$\sigma_{tv} := \text{stdev}(tv)$	$\mu_{tv} := \text{mean}(tv)$
$\sigma_{gp} := \text{stdev}(gp)$	$\mu_{gp} := \text{mean}(gp)$	$\sigma_{gv} := \text{stdev}(gv)$	$\mu_{gv} := \text{mean}(gv)$
$\sigma_{ge} := \text{stdev}(ge)$	$\mu_{ge} := \text{mean}(ge)$	$\sigma_{gs} := \text{stdev}(gs)$	$\mu_{gs} := \text{mean}(gs)$
$\sigma_{hp} := \text{stdev}(hp)$	$\mu_{hp} := \text{mean}(hp)$	$\sigma_{hv} := \text{stdev}(hv)$	$\mu_{hv} := \text{mean}(hv)$
$\sigma_{ep} := \text{stdev}(ep)$	$\mu_{ep} := \text{mean}(ep)$	$\sigma_{ev} := \text{stdev}(ev)$	$\mu_{ev} := \text{mean}(ev)$

### 4. Mean filter velocity data for comparison

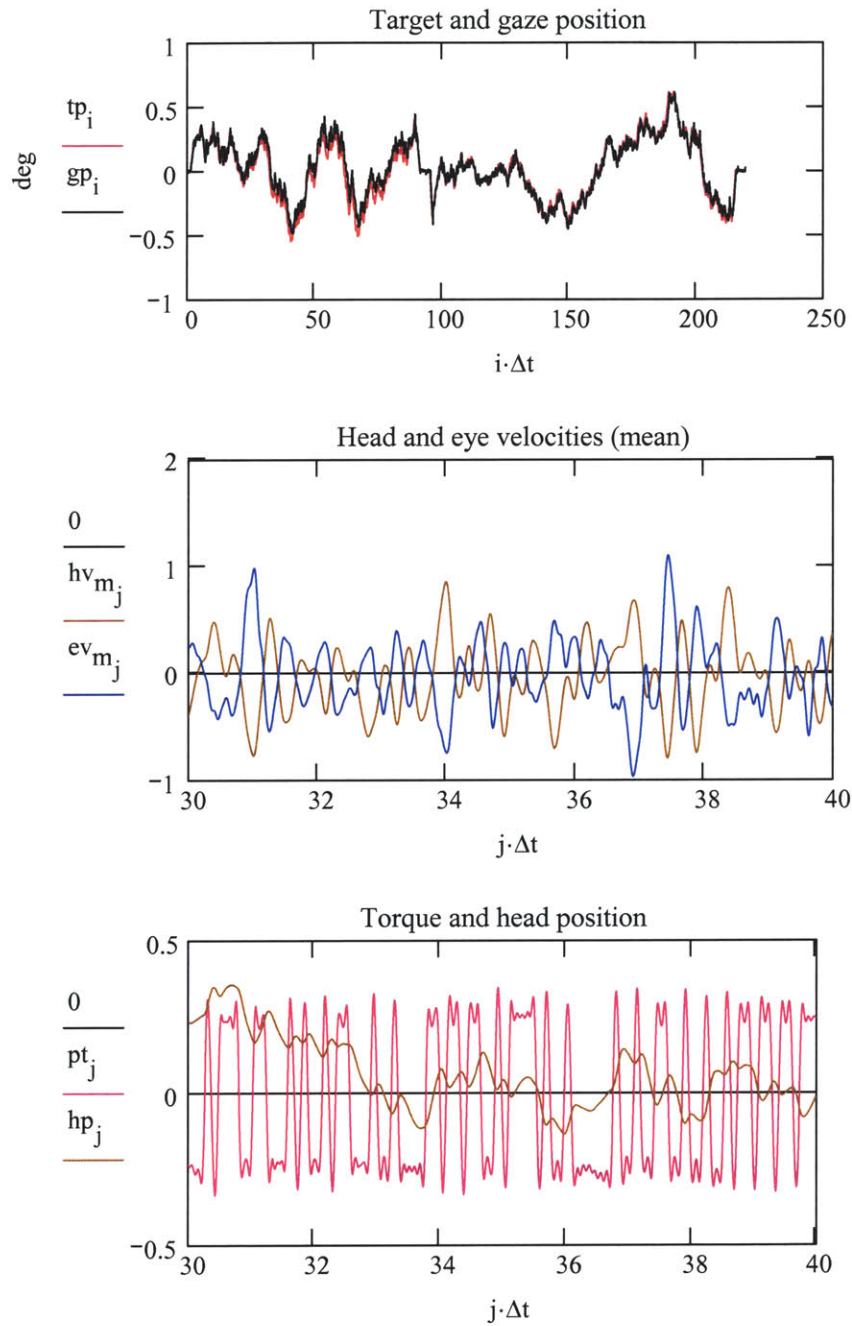
Set number of elements for mean filter:  $\text{elem} := 20$

$tv_m := \text{movavg}(tv, \text{elem})$        $hv_m := \text{movavg}(hv, \text{elem})$

$ev_m := \text{movavg}(ev, \text{elem})$        $gv_m := ev_m + hv_m$

C. Sequences, spectral measures, covariances

plot range:       $a1 := 3000$        $j := a1 .. a1 + 1000$



Set parameters for spectral calculations

$$N = 2.193 \times 10^4 \quad N \cdot \Delta t = 219.32$$

$$nr := 21 \quad \frac{N}{nr} = 1.044 \times 10^3$$

$$ol := 0.5 \quad win := 6$$

### Power spectral densities

$$S_{pt} := \text{pspectrum}(pt, nr, ol, win)$$

$$K := \text{length}(S_{pt})$$

$$S_{tv} := \text{pspectrum}(tv, nr, ol, win)$$

$$S_{hv} := \text{pspectrum}(hv, nr, ol, win)$$

### Non-normalized cross correlations and input PSD


$$C_{pttv} := \sigma_{pt} \cdot \sigma_{tv} \cdot \text{lcorr}(pt, tv)$$

$$C_{tvpt} := \sigma_{pt} \cdot \sigma_{tv} \cdot \text{lcorr}(tv, pt)$$


$$C_{pt} := \sigma_{pt} \cdot \sigma_{pt} \cdot \text{lcorr}(pt, pt)$$

$$C_{tv} := \sigma_{tv} \cdot \sigma_{tv} \cdot \text{lcorr}(tv, tv)$$

frequency increment (Hz):  $\frac{1}{K \cdot \Delta t} = 0.05$        $k := 0.. \frac{K}{10}$



$n := 0..1200$        $\text{amp} := 0.03$



D. Coherence and FRF for multi-input, single-output system

The coherence squared function between input  $u$  and output  $y$  is defined as  $coher_k := \frac{S_{yu_k} \cdot S_{yu_k}}{S_{yy_k} S_{uu_k}}$  ■

For a two input, single output case, where the inputs are perfectly uncorrelated, the sum of the coherences between each input and the output is bounded by 1. When the inputs are not perfectly uncorrelated, the coherence sum that is bounded by unity includes additional terms; the coherence between the two inputs, and two "triple term" coherences between both inputs and the output. The sum of these five coherences is reduced from unity by noise, nonlinearities, and from the correlation between the inputs.

Set params for spectra calcs:  $N = 2.193 \times 10^4$      $nr = 21$      $ol = 0.5$      $win = 6$

$$K = 1.992 \times 10^3 \quad \frac{1}{K \cdot \Delta t} = 0.05$$

1. Calculate coherence squared functions

$$\gamma_{hvpt} := CohSqdMC(hv, pt, nr, ol, win) \quad H := last(\gamma_{hvpt}) \quad h := 0..H$$

$$\gamma_{hvtv} := CohSqdMC(hv, tv, nr, ol, win)$$

$$\gamma_{tvpt} := CohSqdMC(tv, pt, nr, ol, win) \quad \gamma_{pttv} := CohSqdMC(pt, tv, nr, ol, win)$$

head vel - torque             $S_{hvpt} := cspectrum(hv, pt, nr, ol, win)$

head vel - target vel         $S_{hvtv} := cspectrum(hv, tv, nr, ol, win)$

target - torque                 $S_{tvpt} := cspectrum(tv, pt, nr, ol, win)$

torque - torque                 $S_{pt} := cspectrum(pt, pt, nr, ol, win)$

target - target                 $S_{tv} := cspectrum(tv, tv, nr, ol, win)$



head - head                     $S_{hv} := cspectrum(hv, hv, nr, ol, win)$


$$\gamma^{1\text{trip}_h} := \frac{\overline{S_{hvtv_h} \cdot S_{hvp_t_h} \cdot S_{tvpt_h}}}{S_{hv_h} S_{tv_h} \cdot S_{pt_h}} \quad \gamma^{1\text{trip}} := \text{Re}(\gamma^{1\text{trip}_h})$$

$$\gamma^{2\text{trip}_h} := \frac{S_{hvtv_h} \cdot \overline{S_{hvp_t_h} \cdot S_{tvpt_h}}}{S_{hv_h} S_{tv_h} \cdot S_{pt_h}} \quad \gamma^{2\text{trip}} := \text{Re}(\gamma^{2\text{trip}_h})$$

$$\gamma_{\text{sum}_h} := \gamma_{hvp_t_h} + \gamma_{hvtv_h} + \gamma_{tvpt_h} - \gamma^{1\text{trip}_h} - \gamma^{2\text{trip}_h}$$

$$H := \text{length}(\gamma_{hvp_t}) \quad h := 0.. \frac{H}{10}$$





2. Calculate frequency response functions (apparent transfer function)

G2 = head response to visual target

G4 = head response to torque perturbations

$$\text{visually mediated head response (G2)} \quad G2_{\text{FRF}} := \frac{\overrightarrow{S_{\text{hvtv}}}}{S_{\text{tv}}}$$





$$\text{head-neck response (G4)} \quad G4_{\text{FRF}} := \frac{\overrightarrow{S_{\text{hvppt}}}}{S_{\text{pt}}}$$

$$\text{magnitude: } \text{magG2} := \left| \overrightarrow{G2_{\text{FRF}}} \right| \quad P := \text{last}(\text{magG2})$$

$$\text{magG4} := \left| \overrightarrow{G4_{\text{FRF}}} \right| \quad P := \text{last}(\text{magG4})$$

phase (radians):  $\phi_{G2} := \arg(\overrightarrow{G2_{FRF}})$     degrees:  $\phi_{G2} := -\left(\frac{180}{\pi} \cdot \phi_{G2}\right)$

$\phi_{G4} := \arg(\overrightarrow{G4_{FRF}})$     degrees:  $\phi_{G4} := -\left(\frac{180}{\pi} \cdot \phi_{G4}\right)$





E. Time domain head response to torque perturbations,  $g_4$ , and to target velocity,  $g_2$

1. Analysis description

The head has a reflexive response to the torque perturbations,  $p$ , and a voluntary response to track the visual target,  $t$ . The causal input-output convolution relation is:

$$h_n := \Delta t \cdot \sum_{k=0}^{M-1} g_{4k} \cdot p_{n-k} + \Delta t \cdot \sum_{k=0}^{M-1} g_{2k} \cdot t_{n-k} \quad \text{for } n \geq M-1$$

a. Correlation based equations

The correlation based relation, with respect to the input torque,  $p$ , is:

$$C_{ph_j} := \Delta t \cdot \sum_{k=0}^{M-1} g_{4k} \cdot C_{pp_{j-k}} + \Delta t \cdot \sum_{k=0}^{M-1} g_{2k} \cdot C_{pt_{j-k}}$$

where  $C_{ph}$  is the input-output cross-correlation with the input perturbation,  $p$ , lagged wrt the head velocity output,  $h$ .

If the torque-target cross-covariance is assumed equal to zero for all relevant lags, then the relationship simplifies to:

$$C_{ph_j} := \Delta t \cdot \sum_{k=0}^{M-1} g_{4k} \cdot C_{pp_{j-k}}$$

and  $g_4$  can be found by deconvolving  $C_{pp}$  from  $C_{ph}$ .

A similar analysis can be done for the head's response to the target's velocity. The voluntary response to the visual target should be expected to be time variant, since the head does not have to match the target's motions; the gaze must match the target's motion. Expect that this is much less significant, and that the head is sensitive to the target's position.

b. Matrix manipulations

The convolution can be expressed as a series of equations:  $H = P \Delta T G_4 + T \Delta T G_2$ . All elements of the IRF are in each equation for elements  $n = M-1$  to  $N-1$ . The equations are overdetermined so should be solved with the "pseudo" inverse. Relations very similar to the correlation equation are created when multiply the equations by the transverse of the perturbation matrix,  $P^T$ , to form  $P^T H = P^T P \Delta T G_4 + P^T T \Delta T G_2$ .

The matrix relations are slightly different mathematically than the correlation relations. However, if the sequences are considered ergodic and wss,  $P^T H$ ,  $P^T P$ , and  $P^T T$  provide essentially the same estimate of the second order statistical properties of the system as the correlations estimators.

2. Calculate IRF relating torque (pt) and target (tv) to head velocity (hv)

a. Consider making sequences for analysis zero mean

Use lcorr to estimate the correlations. Use zero mean sequences. The sequences can therefore be considered perturbations. Must account for the offsets afterwards.

$$\begin{array}{lll} \text{mean}(pt) = -0.011 & \text{mean}(tv) = -5.678 \times 10^{-7} & \text{mean}(hv) = 0 \\ pt := pt - \text{mean}(pt) & tv := tv - \text{mean}(tv) & hv := hv - \text{mean}(hv) \\ \sigma_{pt} := \text{stdev}(pt) & \sigma_{tv} := \text{stdev}(tv) & \sigma_{hv} := \text{stdev}(hv) \end{array}$$

b. Set length of IRF (M):

$$\begin{array}{lll} N = 2.193 \times 10^4 & N \cdot \Delta t = 219.32 \\ M := 300 & M \cdot \Delta t = 3 & j := 0..M - 1 \\ \text{recLen} := 3000 & \text{recLen} \cdot \Delta t = 30 \\ \text{ovlap} := 1.5 \cdot M & \text{ovlap} \cdot \Delta t = 4.5 \end{array}$$

c. Solve for G2 and G4 using user-defined functions

$$\begin{array}{ll} \text{correlation IRFs} & g2_{\text{cor}} := \text{IRF}(tv_m, hv_m, \Delta t, M) \\ & g4_{\text{cor}} := \text{IRF}(pt, hv, \Delta t, M) \end{array}$$

$$\begin{array}{ll} \text{multiple record IRF} & g2_{\text{ave}} := \text{aveIRF}(tv, hv, \Delta t, M, \text{recLen}, \text{ovlap}) \\ & g4_{\text{ave}} := \text{aveIRF}(pt, hv, \Delta t, M, \text{recLen}, \text{ovlap}) \end{array}$$

F. Fit 2nd order parametric model to the g4 covariance IRF

choose IRF and length to be fitted:  $h_{est} := g4_{ave}$      $V := M$      $v := 0.. V - 1$   
 $g^2_{par} := g^2_{ave}$



set initial parameters     $gain := 0.25$      $\omega_n := 11$      $\zeta := 0.6$

calculate model system:  $h_{guess_v} := h2dmod(v \cdot \Delta t, gain, \omega_n, \zeta)$

estimate FRF from the DTFT of the position IRF

$H4_{vel} := cfft(h_{est})$      $H4_{vel} := magVect(H4_{vel})$      $LH := length(H4_{vel})$      $lh := 0.. \frac{LH}{10}$

$H4_{guess} := cfft(h_{guess})$      $H4_{guess} := magVect(H4_{guess})$



Solve for parameters that minimize mean square error

Define objective function to be minimized with parameter choice

set starting data point:  $U := 0$     end point:  $V = 300$

$$\text{sum of square differences: } SS_{\text{dif}}(\text{gain}, \omega_n, \zeta) := \sum_{j=U}^{V-1} \left( h_{\text{est}_j} - h_{2\text{dmod}}(j \cdot \Delta t, \text{gain}, \omega_n, \zeta) \right)^2$$

Minimize objective function with non-linear method (Lev-Marquardt, conj grad, or quasi-Newton)

Given

$$SS_{\text{dif}}(\text{gain}, \omega_n, \zeta) = 0$$

$1 = 1$   
 $1 = 1$

*Using Minerr in a solve block manner. Minerr requires as many equations after the "Given" command as there are parameters for which to solve. The "equals" sign is ctrl =, the "constrained to be equal to" operator.*

$$\begin{pmatrix} \text{gain} \\ \omega_n \\ \zeta \end{pmatrix} := \text{Minerr}(\text{gain}, \omega_n, \zeta)$$

2nd order model      parameter values:    gain = 0.197       $\omega_n = 11.92$        $\zeta = 0.51$

settling time:       $t_s := \frac{4}{\zeta \cdot \omega_n}$        $t_s = 0.657$

rotational inertia:     $J := \frac{1}{\text{gain} \cdot \omega_n^2}$        $J = 0.0356$

damping:       $B := \frac{2\zeta}{\text{gain} \cdot \omega_n}$        $B = 0.434$

spring stiffness:     $K := \frac{1}{\text{gain}}$        $K = 5.064$

parametric model response function       $g_{4\text{par}_v} := h_{2\text{dmod}}(v \cdot \Delta t, \text{gain}, \omega_n, \zeta)$

FRF from the DTFT of the position IRF       $H_{4\text{par}} := \text{cfft}(g_{4\text{par}})$        $H_{4\text{par}} := \text{magVect}(H_{4\text{par}})$

G. How well do the IRFs predict the output?

*hv predictions, %VAF, and mean square difference*                      k := 8000..9000

truncate IRFs to the same length

$$g2_{\text{cor}} := \text{submatrix}(g2_{\text{cor}}, 0, M - 1, 0, 0) \quad g4_{\text{cor}} := \text{submatrix}(g4_{\text{cor}}, 0, M - 1, 0, 0)$$

$$g2_{\text{ave}} := \text{submatrix}(g2_{\text{ave}}, 0, M - 1, 0, 0) \quad g4_{\text{ave}} := \text{submatrix}(g4_{\text{ave}}, 0, M - 1, 0, 0)$$

$$g2_{\text{par}} := \text{submatrix}(g2_{\text{par}}, 0, M - 1, 0, 0) \quad g4_{\text{par}} := \text{submatrix}(g4_{\text{par}}, 0, M - 1, 0, 0)$$

head velocity predictions - calculate components stimulated by torque (pt) and by target (tv)

from single correlation IRF

$$hv_{\text{corpt}} := \Delta t \cdot \text{convol}(g4_{\text{cor}}, \text{pt}) \quad hv_{\text{cortv}} := \Delta t \cdot \text{convol}(g2_{\text{cor}}, \text{tv})$$

$$hv_{\text{cor}} := hv_{\text{corpt}} + hv_{\text{cortv}}$$

from avg correlation IRF

$$hv_{\text{avept}} := \Delta t \cdot \text{convol}(g4_{\text{ave}}, \text{pt}) \quad hv_{\text{avetv}} := \Delta t \cdot \text{convol}(g2_{\text{ave}}, \text{tv})$$

$$hv_{\text{ave}} := hv_{\text{avept}} + hv_{\text{avetv}}$$

from parametric IRF

$$hv_{\text{parpt}} := \Delta t \cdot \text{convol}(g4_{\text{par}}, \text{pt}) \quad hv_{\text{partv}} := \Delta t \cdot \text{convol}(g2_{\text{par}}, \text{tv})$$

$$hv_{\text{par}} := hv_{\text{parpt}} + hv_{\text{partv}}$$

% VAF and MSD

Remove the same number of elements from end as there are in the impulse response function.

measured hv

$$hv2 := \text{submatrix}(hv, M, \text{last}(hv), 0, 0) \quad \text{last}(hv2) = 2.163 \times 10^4$$

$$\sigma_{hv2} := \text{stdev}(hv2) \quad \sigma_{hv2} = 0.452$$

from single correlation IRF

$$hv2_{\text{cor}} := \text{submatrix}(hv_{\text{cor}}, M, \text{last}(hv), 0, 0) \quad \text{last}(hv2_{\text{cor}}) = 2.163 \times 10^4$$

$$hv2_{\text{corpt}} := \text{submatrix}(hv_{\text{corpt}}, M, \text{last}(hv), 0, 0) \quad \text{last}(hv2_{\text{corpt}}) = 2.163 \times 10^4$$

$$hv2_{\text{cortv}} := \text{submatrix}(hv_{\text{cortv}}, M, \text{last}(hv), 0, 0) \quad \text{last}(hv2_{\text{cortv}}) = 2.163 \times 10^4$$

from avg correlation IRF

$$hv2_{\text{ave}} := \text{submatrix}(hv_{\text{ave}}, M, \text{last}(hv), 0, 0) \quad \text{last}(hv2_{\text{ave}}) = 2.163 \times 10^4$$

$$hv2_{\text{avept}} := \text{submatrix}(hv_{\text{avept}}, M, \text{last}(hv), 0, 0) \quad \text{last}(hv2_{\text{avept}}) = 2.163 \times 10^4$$

$$hv2_{\text{avetv}} := \text{submatrix}(hv_{\text{avetv}}, M, \text{last}(hv), 0, 0) \quad \text{last}(hv2_{\text{avetv}}) = 2.163 \times 10^4$$

from parametric IRF

$$hv2_{\text{par}} := \text{submatrix}(hv_{\text{par}}, M, \text{last}(hv), 0, 0) \quad \text{last}(hv2_{\text{par}}) = 2.163 \times 10^4$$

correlation estimate head velocity

$$hv_{\text{err}} := hv2_{\text{cor}} - hv2 \quad \sigma_{\text{err}} := \text{stdev}(hv_{\text{err}}) \quad \sigma_{\text{err}} = 0.187$$

$$\%VAF \quad VAF_{\text{hvc}} := 100 \left( 1 - \frac{\sigma_{\text{err}}^2}{\sigma_{hv2}^2} \right) \quad MSD_{\text{hvc}} := \text{MSD}(hv2, hv2_{\text{cor}})$$

Of the VAF, what portion is due to which component of the hv

$$hv_{err} := hv2_{corpt} - hv2 \quad \sigma_{err} := stdev(hv_{err}) \quad VAF_{hv_{cpt}} := 100 \left( 1 - \frac{\sigma_{err}^2}{\sigma_{hv2}^2} \right)$$

$$hv_{err} := hv2_{cortv} - hv2 \quad \sigma_{err} := stdev(hv_{err}) \quad VAF_{hv_{ctv}} := 100 \left( 1 - \frac{\sigma_{err}^2}{\sigma_{hv2}^2} \right)$$

$$S_{hvcor} := pspectrum(hv_{cor}, nr, ol, win)$$

average estimate

$$hv_{err} := hv2_{ave} - hv2 \quad \sigma_{err} := stdev(hv_{err}) \quad \sigma_{err} = 0.192$$

$$\%VAF \quad VAF_{hva} := 100 \left( 1 - \frac{\sigma_{err}^2}{\sigma_{hv2}^2} \right) \quad MSD_{hva} := MSD(hv2, hv2_{ave})$$

Of the VAF, what portion is due to which component of the hv

$$hv_{err} := hv2_{avept} - hv2 \quad \sigma_{err} := stdev(hv_{err}) \quad VAF_{hv_{apt}} := 100 \left( 1 - \frac{\sigma_{err}^2}{\sigma_{hv2}^2} \right)$$

$$hv_{err} := hv2_{avetv} - hv2 \quad \sigma_{err} := stdev(hv_{err}) \quad VAF_{hv_{atv}} := 100 \left( 1 - \frac{\sigma_{err}^2}{\sigma_{hv2}^2} \right)$$


parametric estimate of head velocity

$$hv_{err} := hv2_{par} - hv2 \quad \sigma_{err} := stdev(hv_{err}) \quad \sigma_{err} = 0.206$$


$$\%VAF \quad VAF_{hvp} := 100 \left( 1 - \frac{\sigma_{err}^2}{\sigma_{hv2}^2} \right) \quad MSD_{hvp} := MSD(hv2, hv2_{cor})$$

$$S_{hvpar} := pspectrum(hv_{par}, nr, ol, win)$$

g2: head driven by target



g4: head driven by torque



2d order params

gain = 0.197

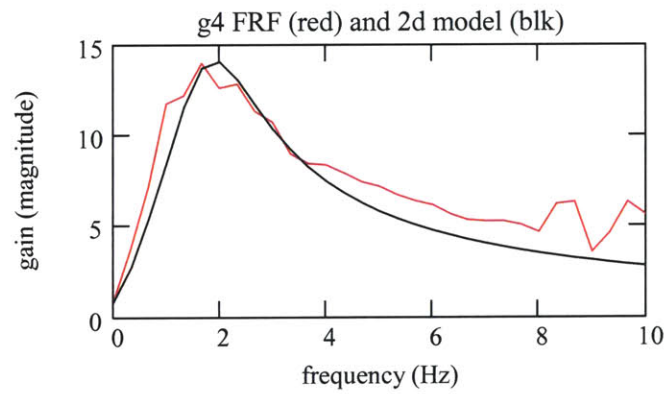
$\omega_n = 11.92$        $t_s = 0.657$

$\zeta = 0.51$

J = 0.036

B = 0.434

K = 5.064





g4: head driven by torque

2d order params

gain = 0.197



$\omega_n = 11.92$        $t_s = 0.657$

$\zeta = 0.51$

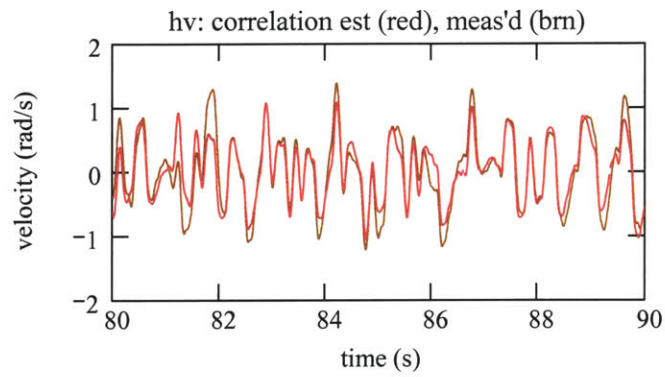
J = 0.036

B = 0.434

K = 5.064



How well do correlation or parametric IRFs model the measured head velocity?




$$VAF_{hvc} = 82.872$$

$$VAF_{hvcpt} = 75.555$$


$$VAF_{hvcvtv} = 7.662$$

$$MSD_{hvc} = 0.035$$



$$VAF_{hva} = 81.943$$

$$MSD_{hva} = 0.037$$



$$VAF_{hvp} = 79.183$$

$$MSD_{hvp} = 0.035$$

$$\frac{180}{\pi} \cdot MSD_{hvp} = 2.004$$

## G3 FRF IRF

Program calculates the impulse response function and frequency response function for the VOR.

filename: G3 FRF 14 Jun 02 6b 20b      last update: 07Apr03  
10 Hz LPF

---

### A. User defined functions



### B. Read processed data files, create gaze position and velocity

*Read processed data from .csv file. Data created by MakeFiles*



C. Is the torque-target CCF zero?

The analysis of the VOR relies upon the tor-targ CCF being approximated as zero. The CCF is measured relative to the off-impulse level of the torque ACF. During the calculation of g3 IRF, Cpttv is assumed zero, Cpthv is treated as the input, and Cptev is treated as the output.

$$\text{calculate standard devs: } \sigma_{pt} := \text{stdev}(pt) \quad \sigma_{tv} := \text{stdev}(tv) \quad \sigma_{tp} := \text{stdev}(tp)$$

$$\sigma_{hv} := \text{stdev}(hv) \quad \sigma_{ev} := \text{stdev}(ev)$$

$$\text{perturbation torque ACF: } C_{pt} := \sigma_{pt}^2 \cdot \text{lcorr}(pt, pt)$$

$$\text{head velocity ACF: } C_{hv} := \sigma_{hv}^2 \cdot \text{lcorr}(hv, hv)$$

$$\text{torque - target velocity CCF: } C_{pttv} := \sigma_{pt} \cdot \sigma_{tv} \cdot \text{lcorr}(pt, tv)$$

$$\text{torque - target position CCF: } C_{pttp} := \sigma_{pt} \cdot \sigma_{tp} \cdot \text{lcorr}(pt, tp)$$

$$\text{torque - head velocity CCF: } C_{pthv} := \sigma_{pt} \cdot \sigma_{hv} \cdot \text{lcorr}(pt, hv)$$

$$\text{torque - eye velocity CCF: } C_{ptev} := \sigma_{pt} \cdot \sigma_{ev} \cdot \text{lcorr}(pt, ev)$$

$$\text{head vel - target vel CCF: } C_{hvtv} := \sigma_{hv} \cdot \sigma_{tv} \cdot \text{lcorr}(hv, tv)$$

$$\text{head vel - eye vel CCF: } C_{hvev} := \sigma_{hv} \cdot \sigma_{ev} \cdot \text{lcorr}(hv, ev)$$


*Non-normalized covariances*

$$\text{plot ranges: } j := 0..2000$$

$$\text{set range for 2-norm: } t1 := 0.2 \quad x1 := \frac{t1}{\Delta t} \quad x1 = 20$$


$$t2 := 8 \quad x2 := \frac{t2}{\Delta t} \quad x2 = 800$$

Torque-target CCF vs. ACF of PRBS torque - important for design, but these are not the actual correlations that are used in the calculation of the VOR IRF.



$$\begin{aligned} \text{Norm2}(C_{pt}, x1, x2) &= 0.016 & \frac{\text{Norm2}(C_{pttv}, x1, x2)}{\text{Norm2}(C_{pt}, x1, x2)} &= 0.21 \\ \text{Norm2}(C_{pttv}, x1, x2) &= 3.406 \times 10^{-3} \end{aligned}$$

The analysis uses the CCF between the torque and the head velocity, and the torque and the eye velocity. These can be compared against the torque target CCF.



$$\begin{aligned} \frac{\text{Norm2}(C_{pthv}, 0, x2)}{\text{Norm2}(C_{pttv}, 0, x2)} &= 37.939 & \frac{\text{Norm2}(C_{ptev}, 0, x2)}{\text{Norm2}(C_{pttv}, 0, x2)} &= 38.087 \end{aligned}$$

#### D. FRF and coherence calculations

set spectral parameters:  $N = 1.857 \times 10^4$      $nr := 20$      $ol := 0.5$      $win := 6$

calculate G3 FRF

"output measure"     $S_{ptev} := cspectrum(pt, ev, nr, ol, win)$      $mS_{ptev} := mag(S_{ptev})$

"input measure"     $S_{pthv} := cspectrum(pt, hv, nr, ol, win)$      $mS_{pthv} := mag(S_{pthv})$



$L := length(S_{ptev})$      $freqIncr := \frac{1}{L \cdot \Delta t}$      $freqIncr = 0.057$

complex FRF:     $G3_{FRF} := \frac{S_{ptev}}{S_{pthv}}$

magnitude:  $mG3 := |G3_{FRF}|$      $P := last(mG3)$

phase (rad):  $\phi G3 := \arg(G3_{FRF})$     deg:  $\phi G3 := -\left(\frac{180}{\pi} \cdot \phi G3\right)$      $\phi G3 := Unwrap(\phi G3)$

$freqIncr = 0.057$      $p := 0.. \frac{P}{25}$



1. Coherence squared: two-input, single-output system

The coherence squared function between input  $u$  and output  $y$  is defined as  $\text{coher}_k := \frac{\overline{S_{yu_k} \cdot S_{yu_k}}}{S_{yy_k} S_{uu_k}}$  ■

For a two input, single output case, where the inputs are perfectly uncorrelated, the sum of the coherences between each input and the output is bounded by 1. When the inputs are not perfectly uncorrelated, the coherence sum that is bounded by unity includes additional terms; the coherence between the two inputs, and two "triple term" coherences between both inputs and the output. The sum of these five coherences is reduced from unity by noise, nonlinearities, and from the correlation between the inputs.

The single output is the eye velocity. The inputs are the target velocity and the head velocity.

pt input, ev output  $\gamma_{\text{evpt}} := \text{CohSqdMC}(\text{ev}, \text{pt}, \text{nr}, \text{ol}, \text{win})$   $H := \text{last}(\gamma_{\text{evpt}})$   $h := 0..H$

tv input, ev output  $\gamma_{\text{evtv}} := \text{CohSqdMC}(\text{ev}, \text{tv}, \text{nr}, \text{ol}, \text{win})$

between the inputs  $\gamma_{\text{pttv}} := \text{CohSqdMC}(\text{pt}, \text{tv}, \text{nr}, \text{ol}, \text{win})$

$\gamma_{\text{tvpt}} := \text{CohSqdMC}(\text{tv}, \text{pt}, \text{nr}, \text{ol}, \text{win})$

eye vel - torque  $S_{\text{evpt}} := \text{cspectrum}(\text{ev}, \text{pt}, \text{nr}, \text{ol}, \text{win})$

eye vel - target vel  $S_{\text{evtv}} := \text{cspectrum}(\text{ev}, \text{tv}, \text{nr}, \text{ol}, \text{win})$

target - tor  $S_{\text{tvpt}} := \text{cspectrum}(\text{tv}, \text{pt}, \text{nr}, \text{ol}, \text{win})$

tor - tor  $S_{\text{pt}} := \text{cspectrum}(\text{pt}, \text{pt}, \text{nr}, \text{ol}, \text{win})$

target - target  $S_{\text{tv}} := \text{cspectrum}(\text{tv}, \text{tv}, \text{nr}, \text{ol}, \text{win})$

eye - eye  $S_{\text{ev}} := \text{cspectrum}(\text{ev}, \text{ev}, \text{nr}, \text{ol}, \text{win})$




$$\gamma_{1\text{trip}_h} := \frac{\overline{S_{\text{evtv}_h} \cdot S_{\text{evpt}_h} \cdot S_{\text{tvpt}_h}}}{S_{\text{ev}_h} S_{\text{tv}_h} \cdot S_{\text{pt}_h}} \quad \gamma_{1\text{trip}} := \text{Re}(\gamma_{1\text{trip}})$$

$$\gamma_{2\text{trip}_h} := \frac{\overline{S_{\text{evtv}_h} \cdot S_{\text{evpt}_h} \cdot S_{\text{tvpt}_h}}}{S_{\text{ev}_h} S_{\text{tv}_h} \cdot S_{\text{pt}_h}} \quad \gamma_{2\text{trip}} := \text{Re}(\gamma_{2\text{trip}})$$



sum of the coherence-square functions  $\gamma_{\text{sum}_h} := \gamma_{\text{evpt}_h} + \gamma_{\text{evtv}_h} + \gamma_{\text{tvpt}_h} - \gamma^1_{\text{trip}_h} - \gamma^2_{\text{trip}_h}$

coh-squared due to inputs being correlated  $\gamma_{\text{cor}} := \gamma_{\text{pttv}} - \gamma^1_{\text{trip}} - \gamma^2_{\text{trip}}$

$H := \text{length}(\gamma_{\text{evpt}})$        $h := 0.. \frac{H}{16.6}$







## 2. Model FRF as high pass system

### a. Second order high pass system

set initial parameter guess     $\text{gain}_{\text{hf2}} := 1.0$      $L := 0.5$      $R := 0.5$      $C := 1$

define frequency range over which to fit the model     $f_{\text{min}} := 0$      $f_{\text{max}} := 4$

Calculate 2d HP system models and initial guess


$$\text{2d FRF: } G_{2\text{dhps}}(\text{gain}_{\text{hf2}}, L, R, C, f) := \frac{\text{gain}_{\text{hf2}} \cdot (j \cdot 2\pi f)^2}{(j \cdot 2\pi f)^2 + \frac{R}{L} \cdot (j \cdot 2\pi f) + \frac{1}{L \cdot C}} \quad \text{freqIncr} = 0.057$$

$$\text{magnitude: } mG_{2\text{dhps}}(\text{gain}_{\text{hf2}}, L, R, C, f) := |G_{2\text{dhps}}(\text{gain}_{\text{hf2}}, L, R, C, f)|$$

$$\text{phase: } \phi G_{2\text{dhps}}(\text{gain}_{\text{hf2}}, L, R, C, f) := \frac{180}{\pi} \cdot \arg(G_{2\text{dhps}}(\text{gain}_{\text{hf2}}, L, R, C, f))$$

$$j_{\text{min}} := \text{floor}\left(\frac{f_{\text{min}}}{\text{freqIncr}}\right) \quad j_{\text{max}} := \text{floor}\left(\frac{f_{\text{max}}}{\text{freqIncr}}\right) \quad j := j_{\text{min}} \cdot j_{\text{max}}$$

Calculate initial model:  $mG_{2\text{dmod}_j} := mG_{2\text{dhps}}(\text{gain}_{\text{hf2}}, L, R, C, j \cdot \text{freqIncr})$



Calculate parameters to minimize the SSD fit of 2d HP system from  $f1$  to  $f2$

Define objective function to be minimized

$$\text{ssdG3}(\text{gain}_{\text{hf2}}, L, R, C) := \sum_{j = j_{\text{min}}}^{j_{\text{max}}} \left( mG_{3j} - mG_{2\text{dhps}}(\text{gain}_{\text{hf2}}, L, R, C, j \cdot \text{freqIncr}) \right)^2$$

$$\text{ssdG3}(\text{gain}_{\text{hf2}}, L, R, C) = 2.726$$

Minimization of 2d order  
HF system parameters:

Given

$$\text{ssdG3}(\text{gain}_{\text{hf2}}, L, R, C) = 0$$

$$1 = 1$$

$$1 = 1$$


$$1 = 1$$

$$\begin{pmatrix} \text{gain}_{\text{hf2}} \\ L \\ R \\ C \end{pmatrix} := \text{Minerr}(\text{gain}_{\text{hf2}}, L, R, C)$$

"Min" values:  $\text{gain}_{\text{hf2}} = 0.969$     $L = 2.203$     $R = 0.761$     $C = 1.755$     $\frac{1}{R \cdot C} = 0.749$

Calculate model system:  $mG_{2\text{dmod},j} := mG_{2\text{dhps}}(\text{gain}_{\text{hf2}}, L, R, C, j \cdot \text{freqIncr})$

$$\text{ssdG3}(\text{gain}_{\text{hf2}}, L, R, C) = 0.31$$



b. First order high pass system

Set initial parameter guess     $rc := 1.2$              $gain_{hf1} := 1.0$

*Calculate 1d HP system models and initial guess*


Define 1st order HF system and initial models

FRF:             $G_{1dhps}(gain_{hf1}, rc, f) := \frac{gain_{hf1} \cdot rc \cdot j \cdot 2\pi f}{1 + rc \cdot j \cdot 2\pi f}$      $freqIncr = 0.057$

magnitude:     $mG_{1dhps}(gain_{hf1}, rc, f) := |G_{1dhps}(gain_{hf1}, rc, f)|$

phase:             $\phi_{G_{1dhps}}(gain_{hf1}, rc, f) := \frac{180}{\pi} \cdot \arg(G_{1dhps}(gain_{hf1}, rc, f))$

Calculate initial model:     $mG_{1dmod_j} := mG_{1dhps}(gain_{hf1}, rc, j \cdot freqIncr)$



Calculate parameters to minimize SSD fit of 1st order HP system

Define objective function to be minimized

$$\text{ssdG3}_{1d}(\text{gain}_{\text{hf1}}, \text{rc}) := \sum_{j = j_{\text{min}}}^{j_{\text{max}}} \left( mG3_j - mG_{1d\text{hps}}(\text{gain}_{\text{hf1}}, \text{rc}, j \cdot \text{freqIncr}) \right)^2$$

$$\text{ssdG3}_{1d}(\text{gain}_{\text{hf1}}, \text{rc}) = 1.251$$

Minimization technique

Given


$$\text{ssdG3}_{1d}(\text{gain}_{\text{hf1}}, \text{rc}) = 0$$

$$1 = 1$$

$$\begin{pmatrix} \text{gain}_{\text{hf1}} \\ \text{rc} \end{pmatrix} := \text{Minerr}(\text{gain}_{\text{hf1}}, \text{rc})$$

Calculate model with least squares fit system:  $mG_{1d\text{mod}_j} := mG_{1d\text{hps}}(\text{gain}_{\text{hf1}}, \text{rc}, j \cdot \text{freqIncr})$

$$\text{ssdG3}_{1d}(\text{gain}_{\text{hf1}}, \text{rc}) = 0.621$$



## E. IRF calculations

### Analysis description

The convolution relation for the eye motion is

$$e_n := \Delta t \cdot \sum_{i=0}^{M-1} g1_i \cdot t_{n-i} + \Delta t \cdot \sum_{i=0}^{M-1} g3_i \cdot h_{n-i} + \varepsilon_n$$

The target and the head inputs are not uncorrelated, so  $g3$  cannot be found from a straightforward cross-correlation analysis between  $h$  and  $t$ . The experiment was designed so the target velocity would be uncorrelated with the perturbation torque. To find  $g3$ , the input and output sequences must be cross-correlated with the perturbation torque.

Consider the cross-correlation with the perturbation torque:

$$C_{ptev_n} := \Delta t \cdot \sum_{i=0}^{M-1} g1_i \cdot C_{pttv_{n-i}} + \Delta t \cdot \sum_{i=0}^{M-1} g3_i \cdot C_{pthv_{n-i}} + C_{pte_n}$$

Under the assumption that  $C_{pttv}$  and  $C_{pte}$  are zero,  $g3$  is estimated from:

$$C_{ptev_n} := \Delta t \cdot \sum_{i=0}^{M-1} g3_i \cdot C_{pthv_{n-i}}$$

### Linear algebra solution

The convolution equation for elements  $M-1$  to  $N-1$  can be put in a matrix form:

$$E := T \cdot \Delta T \cdot G1 + H \cdot \Delta T \cdot G3 + \varepsilon$$

Following the same argument as for the correlation solution, multiply through by  $P^T$  to form:

$$P^T E := P^T H \cdot \Delta T \cdot G3$$

### Methods

1. Correlation method - set first  $M-1$  elements of output to 0
2. Correlation method - ave IRF from multiple records, set elements to 0

Sequence lengths:  $N = 1.857 \times 10^4$        $N \cdot \Delta t = 185.72$

Set desired IRF length       $M := 300$        $M \cdot \Delta t = 3$       lowest FRF frequency:  $\frac{1}{M \cdot \Delta t} = 0.333$

Set record length and overlap:       $lenIRF := M$        $lenOL := 1.5 \cdot M$

$$lenRec := \frac{N}{5} \quad lenRec \cdot \Delta t = 37.144$$

### IRF calculations

a. IRF using covariance - first M-1 output elements set to zero

$$\text{g3 IRF: } \quad \text{g3}_{\text{CZ}} := \text{g3}_{\text{corZ}}(\text{M}, \Delta\text{t}, \text{pt}, \text{hv}, \text{ev})$$

$$\text{G3 FRF = DFT(irf): } \quad \text{G3}_{\text{CZ}} := \text{cfft}(\text{g3}_{\text{CZ}}) \quad \text{mG3}_{\text{CZ}} := \text{mag}(\text{G3}_{\text{CZ}})$$

Estimate vor component of eye velocity

$$\text{ev}_{\text{CZ}} := \Delta\text{t} \cdot \text{convol}(\text{g3}_{\text{CZ}}, \text{hv})$$

$$\text{ev}_{\text{CZ}} := \text{submatrix}(\text{ev}_{\text{CZ}}, 0, \text{last}(\text{ev}), 0, 0) \quad \text{ev}_{\text{CZ}} := \text{Re}(\text{ev}_{\text{CZ}})$$

$$\text{S}_{\text{evCZ}} := \text{pspectrum}(\text{ev}_{\text{CZ}}, \text{nr}, \text{ol}, \text{win})$$

$$\text{VAF}_{\text{CZev}} := \text{VAF}\%(\text{ev}_{\text{CZ}}, \text{ev}, \text{M})$$

$$\text{VAF}_{\text{CZhv}} := \text{VAF}\%(\text{ev}_{\text{CZ}}, -\text{hv}, \text{M})$$

What remains of eye velocity must have been driven by target or other  
Mean filter to get rid of noise etc..

$$\text{ev}_{\text{CZtv}} := \text{ev} - \text{ev}_{\text{CZ}} \quad \text{ev}_{\text{CZtv}} := \text{movavg}(\text{ev}_{\text{CZtv}}, \text{elem})$$

b. IRF averaged from records

$$\text{ave IRF: } \quad \text{g3}_{\text{avCZ}} := \text{g3}_{\text{avcorZ}}(\text{pt}, \text{hv}, \text{ev}, \Delta\text{t}, \text{lenIRF}, \text{lenRec}, \text{lenOL})$$

$$\text{G3 FRF = DFT(irf): } \quad \text{G3}_{\text{avCZ}} := \text{cfft}(\text{g3}_{\text{avCZ}}) \quad \text{mG3}_{\text{avCZ}} := \text{mag}(\text{G3}_{\text{avCZ}})$$

Estimate vor component of eye velocity

$$\text{ev}_{\text{avCZ}} := \Delta\text{t} \cdot \text{convol}(\text{g3}_{\text{avCZ}}, \text{hv})$$

$$\text{ev}_{\text{avCZ}} := \text{submatrix}(\text{ev}_{\text{avCZ}}, 0, \text{last}(\text{ev}), 0, 0) \quad \text{ev}_{\text{avCZ}} := \text{Re}(\text{ev}_{\text{avCZ}})$$

$$S_{ev_{avcz}} := \text{pspectrum}(ev_{avcz}, nr, ol, win)$$

$$VAF_{avczev} := VAF\%(ev_{avcz}, ev, M)$$

$$VAF_{avczhv} := VAF\%(ev_{avcz}, -hv, M)$$

What remains of eye velocity must have been driven by target or other  
Mean filter to get rid of noise etc..

$$ev_{avcztv} := ev - ev_{avcz} \quad ev_{avcztv} := \text{movavg}(ev_{avcztv}, elem)$$

plot ranges

$$\text{IRF plot range:} \quad j := 0..100$$

$$\text{FRF plot range:} \quad H := \text{length}(mG3_{cz}) \quad h := 0.. \frac{H}{20}$$


$$\text{estimate plot:} \quad n := 1200..1600$$

$$\text{PSD of estimates} \quad K := \text{length}(S_{evcz}) \quad k := 0.. \frac{K}{50}$$


$$\text{target velocity plot} \quad r := 2000..6000$$



IRF estimates from covariance, zeros set in first M elements




IRF estimates from IRF averaged from multiple records




IRF estimates from covariance, zeros set in first M elements

M = 300




$VAF_{czev} = 67.234$



$VAF_{czhv} = 88.519$

The portion of the ev that is not driven by the VOR is driven by the target. However, the target drives the gaze, so the eye velocity may not look like the target



$elem \cdot \Delta t = 1$

## User Defined Functions

### A. Derivatives and integration

#### 5 point equal-abscissa derivative function: Deriv5pt(X, Δt)

Differentiation for sampled sequence with equally spaced abscissas. There are N elements, 0 through N-1. Algorithm from Abramowitz and Stegun, Handbook of Mathematical Functions, 25.3

The algorithm approximates the derivative p time steps from center value,  $x_0$ .  $f(x_{0-1})$  indicates the function evaluated at  $t_0 - 1Dt$ .

The derivative is found with the expansion:

$$\text{fdot}(x_0 + p \cdot \Delta t) := \frac{1}{\Delta t} (a(p) \cdot f(x_{0-2}) + b(p) \cdot f(x_{0-1}) + c(p) \cdot f(x_0) + d(p) \cdot f(x_1) + e(p) \cdot f(x_2))$$

where the coefficients:

$$a(p) := \frac{2p^3 - 3p^2 - p + 1}{12} \quad b(p) := -1 \cdot \frac{4 \cdot p^3 - 3p^2 - 8p + 4}{6} \quad c(p) := \frac{2p^3 - 5p}{2}$$

$$e(p) := \frac{2p^3 + 3p^2 - p - 1}{12} \quad d(p) := -1 \cdot \frac{4p^3 + 3p^2 - 8p - 4}{6}$$

$$p := 2$$

$$a(p) \cdot 12 = \blacksquare \quad b(p) \cdot 6 = \blacksquare \quad c(p) \cdot 2 = \blacksquare \quad d(p) \cdot 6 = \blacksquare \quad e(p) \cdot 12 = \blacksquare$$

This particular form of the algorithm calculates the deriv at elements 2 through N-3 as the center point of the 5 point expansion ( $p = 0$ ). Elements 0 and 1, and N-2 and N-1, cannot be calculated with  $p = 0$ .  $p = -2$  for element 0,  $p = -1$  for element 1,  $p = 1$  for element N-2, and  $p = 2$  for element N-1.

```

Deriv5pt(X, Δt) :=
  N ← length(X)
  for i ∈ 0
    Xdoti ←  $\frac{1}{\Delta t} \cdot \left( \frac{-25}{12} \cdot X_i + 4 \cdot X_{i+1} - 3 \cdot X_{i+2} + \frac{4}{3} \cdot X_{i+3} - \frac{1}{4} \cdot X_{i+4} \right)$ 
  for i ∈ 1
    Xdoti ←  $\frac{1}{\Delta t} \cdot \left( \frac{-1}{4} \cdot X_{i-1} - \frac{5}{6} X_i + \frac{3}{2} X_{i+1} - \frac{1}{2} X_{i+2} + \frac{1}{12} X_{i+3} \right)$ 
  for i ∈ 2..(N-3)
    Xdoti ←  $\frac{1}{\Delta t} \cdot \left( \frac{1}{12} X_{i-2} - \frac{2}{3} X_{i-1} + \frac{2}{3} X_{i+1} - \frac{1}{12} X_{i+2} \right)$ 
  for i ∈ N-2
    Xdoti ←  $\frac{1}{\Delta t} \cdot \left( \frac{-1}{12} X_{i-3} + \frac{1}{2} X_{i-2} - \frac{3}{2} X_{i-1} + \frac{5}{6} X_i + \frac{1}{4} X_{i+1} \right)$ 
  for i ∈ N-1
    Xdoti ←  $\frac{1}{\Delta t} \cdot \left( \frac{1}{4} X_{i-4} - \frac{4}{3} X_{i-3} + 3 X_{i-2} - 4 X_{i-1} + \frac{25}{12} X_i \right)$ 
  Xdot

```

Trapezoidal integration: Pos(vel, Δt)

```

Pos(vel, Δt) :=
  N ← length(vel)
  pos0 ← 0
  sum ← 0
  for i ∈ 1..N-1
    posi ← posi-1 + 0.5 · Δt · (veli-1 + veli)
  pos

```

## B. Correlations and covariances

### Biased auto-correlation estimator: AutoCorr(x, maxlag)

This estimator is not normalized by the sequence variance, so the zero lag value is equal to the sum of the squares of all terms - a measure of total sequence power. Though much slower than Mathcad compiled function, this function is explicitly defined. Similar functions are Mathcad's lcorr( ), the normalized covariance, and correl( ), the non-averaged correlation.

AutoCorr arguments: x, the sequence; maxlag, maximum lag desired. If maxlag is less than or equal to zero, a default maxlag is set to the number of terms in the sequence.

$$\text{AutoCorr}(x, \text{maxlag}) := \left\{ \begin{array}{l} M \leftarrow \text{length}(x) \\ J \leftarrow \text{last}(x) \text{ if } (\text{maxlag} = 0) \vee (\text{maxlag} \geq \text{last}(x)) \\ J \leftarrow \text{maxlag} \text{ otherwise} \\ \text{for } j \in 0..J \\ R_{xxj} \leftarrow \frac{1}{M} \cdot \sum_{i=0}^{M-1-j} x_i \cdot x_{i+j} \\ R_{xx} \end{array} \right.$$

### Biased auto-covariance estimator: AutoCovar(x, maxlag)

Covariance removes the sequence mean.

This estimator is not normalized by the sequence variance, so the zero lag value is equal to the sum of the squares of all terms - a measure of total sequence power. Though much slower than Mathcad compiled function, this function is explicitly defined. Similar functions are Mathcad's lcorr( ), the normalized covariance, and correl( ), the non-averaged correlation.

AutoCovar arguments: x, the sequence; maxlag, maximum lag desired. If maxlag is less than or equal to zero, a default maxlag is set to the number of terms in the sequence.

```

AutoCovar(x, maxlag) :=
  M ← length(x)
  μ ← mean(x)
  z ← x - μ
  J ← last(z) if (maxlag = 0) ∨ (maxlag ≥ last(z))
  J ← maxlag otherwise
  for j ∈ 0..J
    Cxxj ←  $\frac{1}{M} \cdot \sum_{i=0}^{M-1-j} z_i \cdot z_{i+j}$ 
  Cxx

```

### C. Gaussian related functions

Gaussian PDF:  $\text{PDF}_{\text{gauss}}(x, \mu, \sigma)$

$$\text{PDF}_{\text{gauss}}(a, \mu, \sigma) := \frac{1}{\sqrt{2\pi\sigma^2}} \cdot e^{-\left[\frac{(a-\mu)^2}{2\sigma^2}\right]}$$

Gaussian CDF:  $\text{CDF}_{\text{gauss}}(\text{bin}_{\text{min}}, z, \mu, \sigma)$

$$\text{CDF}_{\text{gauss}}(\text{bin}_{\text{min}}, z, \mu, \sigma) := \int_{\text{bin}_{\text{min}}}^z \frac{1}{\sigma \cdot \sqrt{2 \cdot \pi}} \cdot e^{-\frac{1}{2} \cdot \left(\frac{z-\mu}{\sigma}\right)^2} dz$$

Gaussian, white sequence:  $GWN(\Delta t, \text{durSec}, \mu, \sigma, \text{numSwaps})$

$GWN(\Delta t, \text{durSec}, \mu, \sigma, \text{numSwaps})$

Function arguments:

$\Delta t$  - the desired sampling period

$\text{durSec}$  - desired duration of the sequence in seconds

$\mu$  - desired mean

$\sigma$  - desired standard deviation

$\text{numSwaps}$  - number of stochastic swaps.

If  $\text{numSwaps} = 0$ , then default is  $4 \cdot \text{sequence length}$ .

```

GWN( $\Delta t, \text{durSec}, \mu, \sigma, \text{numSwaps}$ ) :=
  numElem  $\leftarrow$  floor( $\frac{\text{durSec}}{\Delta t}$ ) + 1
  for  $j \in 0.. \text{numElem}$ 
     $x_j \leftarrow \frac{j \cdot \Delta t}{\text{durSec} + \Delta t}$ 
  for  $j \in 1.. \text{numElem} - 1$ 
     $y_{\text{old},j} \leftarrow \text{qnorm}(x_j, 0, 1)$ 
  N  $\leftarrow$  last( $y_{\text{old}}$ )
  for  $c \in 0.. \text{if}(\text{numSwaps} = 0, 4N, \text{numSwaps} - 1)$ 
    ACFold  $\leftarrow$  lcorr( $y_{\text{old}}, y_{\text{old}}$ )
    SSold  $\leftarrow$   $\sum_{i=0}^{\text{last}(\text{ACF}_{\text{old}})} (\text{ACF}_{\text{old},i})^2$ 
    j1  $\leftarrow$  floor(rnd(N))
    j2  $\leftarrow$  floor(rnd(N))
     $y_{\text{new}} \leftarrow y_{\text{old}}$ 
     $y_{\text{new},j1} \leftarrow y_{\text{old},j2}$ 
     $y_{\text{new},j2} \leftarrow y_{\text{old},j1}$ 
    ACFnew  $\leftarrow$  lcorr( $y_{\text{new}}, y_{\text{new}}$ )
    SSnew  $\leftarrow$   $\sum_{i=0}^{\text{last}(\text{ACF}_{\text{new}})} (\text{ACF}_{\text{new},i})^2$ 
     $y_{\text{old}} \leftarrow \text{if}(\text{SS}_{\text{new}} < \text{SS}_{\text{old}}, y_{\text{new}}, y_{\text{old}})$ 

```

$$\left| \begin{array}{l} \mu_y \leftarrow \text{mean}(y_{\text{old}}) \\ \sigma_y \leftarrow \text{stdev}(y_{\text{old}}) \\ y_{\text{old}} \leftarrow \sigma \left( \frac{y_{\text{old}} - \mu_y}{\sigma_y} \right) + \mu \\ y_{\text{old}} \end{array} \right.$$

## D. Probability

### PDF and CDF algorithm

Spread 100 (for example) amplitude bins across 8 standard deviations centered at the mean. Count how many outcomes fall into each bin. Divide by total sequence length to find the probability density function.

Plot: The PDF should be plotted with an index of "bin", with the indep index  $(\text{bin}_{\text{min}} + \Delta_{\text{bin}} \cdot \text{bin})$ . It is Best to keep this as an algorithm rather than a function so can keep values for bins available for plotting.

set temp array:  $\text{ta} := \text{hvel}$        $\sigma_{\text{ta}} := \text{stdev}(\text{ta})$   
 $\mu_{\text{ta}} := \text{mean}(\text{ta})$        $\mu_{\text{ta}} = \blacksquare$

sequence length:  $D := \text{length}(\text{ta})$        $D = \blacksquare$

set number of bins:  $\text{numbins} := 100$        $\text{bin} := 0..(\text{numbins} - 1)$

set number of deviations (+/-) over which to spread bins:  $\text{bindevs} := 4$

calculate bin bounds and increments

$$\text{bin}_{\text{min}} := \mu_{\text{ta}} - \text{bindevs} \cdot \sigma_{\text{ta}} \quad \text{bin}_{\text{min}} = \blacksquare$$

$$\text{bin}_{\text{max}} := \mu_{\text{ta}} + \text{bindevs} \cdot \sigma_{\text{ta}} \quad \text{bin}_{\text{max}} = \blacksquare$$

$$\Delta_{\text{bin}} := \frac{\text{bin}_{\text{max}} - \text{bin}_{\text{min}}}{\text{numbins}} \quad \Delta_{\text{bin}} = \blacksquare$$

$$\text{PDF: } P_{\text{ta}_{\text{bin}}} := \frac{1}{D \cdot \Delta_{\text{bin}}} \cdot \sum_{i=0}^{D-1} \text{if} \left[ (\text{bin}_{\text{min}} + \Delta_{\text{bin}} \cdot \text{bin}) < \text{ta}_i \leq [\text{bin}_{\text{min}} + \Delta_{\text{bin}} \cdot (\text{bin} + 1)], 1, 0 \right]$$

set PDF of orig to temp:  $P_{\text{hvel}} := P_{\text{ta}}$



### Cumulative distribution function

The CDF function evaluated at  $x_0$  represents the probability of an outcome being less than  $x_0$ . The distribution function is the integral of the density function.

$$\text{Cumulative distribution function: } \text{CDF}_{\text{ta}_{\text{bin}}} := \frac{1}{N+1} \cdot \sum_{i=0}^N \text{if}[\text{ta}_i < (\text{bin}_{\text{min}} + \Delta_{\text{bin}} \cdot \text{bin}), 1, 0]$$

set CDF of orig to temp:  $\text{CDF}_{\text{hvel}} := \text{CDF}_{\text{ta}}$

plot on a X-Y plot with axis graph, same arguments as for PDF

## E. Spectrum

Coherence squared function (Mathcad functions): CohSqdMC(x, y, nr, ol, win)

Function uses Mathcad function pspectrum and cspectrum calculate the coherence function between input x and output y. User sets record length, overlap, windowing

$$\begin{array}{l} \text{CohSqdMC}(x, y, nr, ol, win) := \\ \left. \begin{array}{l} S_{xx} \leftarrow \text{pspectrum}(x, nr, ol, win) \\ S_{yy} \leftarrow \text{pspectrum}(y, nr, ol, win) \\ S_{yx} \leftarrow \text{cspectrum}(y, x, nr, ol, win) \\ \text{for } r \in 0.. \text{last}(S_{yx}) \\ \text{cohsqd}_r \leftarrow \frac{S_{yx_r} \cdot \overline{S_{yx_r}}}{S_{yy_r} S_{xx_r}} \\ \text{cohsqd} \end{array} \right\} \end{array}$$

PSD function: PSD(za, zb, RecLen, overlap)

PSD( ) pulls records of length RecLen from sequences za and zb, calculates the record's DFTs, Za(f) and Zb(f), using Mathcad's complex fft, cfft( ), and estimates the record's cross power spectrum as Za(f)Zb\*(f). The PSDs of each record are averaged to estimate the average PSD of the original sequence. Since the DFTs are calculated with Mathcad's cfft( ) the PSD is the square magnitude of the DFT with no additional scale factor.

```

PSD(za, zb, RecLen, overlap) :=
  las ← min(last(za), last(zb))
  zahold ← submatrix(za, 0, las, 0, 0)
  zbhold ← submatrix(zb, 0, las, 0, 0)
  L ← RecLen
  numRec ← 0
  PSDsumL-1 ← 0
  while L ≤ length(zahold)
    numRec ← numRec + 1
    zat ← submatrix(zahold, 0, L - 1, 0, 0)
    zbt ← submatrix(zbhold, 0, L - 1, 0, 0)
    zaDFT ← cfft(zat)
    zbDFT ← cfft(zbt)
    PSDrec ←  $\overrightarrow{(za_{DFT} \cdot \overline{zb_{DFT}})}$ 
    PSDsum ←  $\overrightarrow{(PSD_{sum} + PSD_{rec})}$ 
    a ← floor[(1 - overlap) · L]
    zahold ← submatrix(zahold, a, last(zahold), 0, 0)
    zbhold ← submatrix(zbhold, a, last(zbhold), 0, 0)
  PSD ←  $\frac{PSD_{sum}}{numRec}$ 
  PSD

```

Coherence squared function: CohSqd(x, y, reclen, overlap)

Function uses user defined PSD( ) to calculate the coherence function between input x and output y. User sets record length and overlap.

```

CohSqd(x, y, reclen, overlap) :=
  Sxx ← PSD(x, x, reclen, overlap)
  Syy ← PSD(y, y, reclen, overlap)
  Syx ← PSD(y, x, reclen, overlap)
  for r ∈ 0..last(Syx)
    cohr ←  $\frac{S_{yx_r} \cdot \overline{S_{yx_r}}}{S_{yy_r} S_{xx_r}}$ 
  coh

```

## F. Correction filters for sensors

Watson sensor 14150 calibration correction: HVcorr14150(hv)


Algorithm adds a correction to the measured voltage so voltage represents velocity

```


HVcorr14150(hv) :=
  N ← length(hv)
  for i ∈ 0..N - 1
    corr ← 0
    corr ← 3.98 · 10-3 · hvi + 0.0398 if -10 ≤ hvi < -7.49
    corr ← 0.01 if -7.49 ≤ hvi < -4.99
    corr ← -4.0 · 10-3 · hvi - 0.0096 if -4.99 ≤ hvi < 0.01
    corr ← -0.01 if 0.01 ≤ hvi < 10.00
    hvcorri ← hvi + corr
  hvcorr

```

$$i := 0..39 \quad hv_i := \frac{(i - 20)}{2}$$



$$hv_{new} := HVcorr14150(hv) \quad corr := hv_{new} - hv$$



## F. Low pass filter algorithms

Noncausal LPF with Kaiser window:  $ncLPFkaiser(\Delta t, f_{pb}, f_{sb}, \delta)$

This algorithm generates the positive elements of a LPF with a Kaiser window. The filter is intended to be used with the convolution function,  $ncSymDtConvolve(h_{alg}, x, \Delta t)$ , which should also be scaled by  $1/Dt$ . This was developed using Oppenheim and Schaffer, Discrete-Time Signal Processing, chaps 5-7.

```

ncLPFkaiser(■) :=
  f_samp ← 1 / Δt
  ω_pb ← (2π Δt) f_pb
  ω_sb ← (2π Δt) f_sb
  ω_co ← (ω_pb + ω_sb) / 2
  ω_trans ← ω_sb - ω_pb
  A ← -20 log(δ)
  M_min ← (A - 8) / (2.285 ω_trans)
  M ← floor(M_min + 1)
  rmdr ← mod(M, 2)
  M ← (M + 1) if |rmdr| < 0.001
  β ← 0.1102(A - 8.7) if A > 50
  β ← 0.5842(A - 21)0.4 + 0.07886(A - 21) if (A ≥ 21) ∧ (A ≤ 50)
  β ← 0.0 if A < 21
  for m ∈ 1 .. (M - 1) / 2
    numArg ← β · [1 - ((2 · m) / (M - 1))2]1/2
    w_num ← I0(numArg)
    w_den ← I0(β)
    h_kaiser_m ← (w_num / w_den) · (sin(m · ω_co) / (π m))
  h_kaiser_0 ← (ω_co / π)
  h_kaiser

```

Noncausal convolution with IRF symmetric about 0: ncSymDtConvolve(h, x, Δt)

This algorithm separates the causal (indices  $\geq 0$ ) and noncausal portions of the convolution. Expresses the noncausal portion in terms of positive indices. This implementation requires the filter IRF (h) to be symmetric about 0. Scales by Δt to approximate the dt of the continuous convolution equation.

```

ncSymDtConvolve(h, x, Δt) :=
  Mmax ← last(h)
  Nmax ← last(x)
  for i ∈ 0..Nmax
    ycausi ←  $\sum_{j=0}^{\min(Mmax, i)} h_j \cdot x_{i-j}$ 
  for i ∈ 0..(Nmax - 1)
    ynci ←  $\sum_{k=1}^{\min(Mmax, Nmax-i)} h_k \cdot x_{i+k}$ 
  yncNmax ← 0
  for i ∈ 0..Nmax
    yi ← Δt · (ynci + ycausi)
  y

```

## G. IRF algorithms

Correlation (covariance) based IRF:  $\text{IRF}(\text{input}, \text{output}, \Delta t, \text{lengthIRF})$

Calculates IRF between input and output using covariance. Vectors should either have zero mean vectors or user should account for offset when using the IRF to predict output. The algorithm assumes that the input and out vectors have the same number of elements.

The calculation is conducted by setting the first  $M$  elements of the output to zero. This has the effect of only considering output responses that are a result of all elements of the IRF.

```

IRF(input, output, Δt, lengthIRF) :=
  N ← length(output)
  M ← lengthIRF
  y_rec ← output
  x_rec ← input
  for n ∈ 0..M-2
    y_rec_n ← 0
  σ_y ← stdev(y_rec)
  σ_x ← stdev(x_rec)
  C_xy ← σ_x · σ_y · lcorr(x_rec, y_rec)
  C_xx ← σ_x2 · lcorr(x_rec, x_rec)
  for r ∈ 0..M-1
    for c ∈ 0..M-1
      Φ_x_{r,c} ← C_xx_{|r-c|}
  invΦ_x ← Φ_x-1
  C_xy ← submatrix(C_xy, 0, M-1, 0, 0)
  IRF ←  $\frac{1}{\Delta t}$  invΦ_x · C_xy
  IRF
  
```

Multiple record IRF: aveIRF(input, output, Δt, lenIRF, lenRec, lenOL)

The function breaks apart the input and output sequences into records of lenRec length. Each successive record begins lenOL terms before the end of the previous record. The IRF for each record are calculated for lenIRF terms.

The function used the lcorr() function which is a covariance, so the user must account for sequences that do not have zero mean.

```

aveIRF(■) := | N ← length(output)
              | M ← lenIRF
              | R ← lenRec
              | OL ← lenOL
              | for n ∈ 0.. M - 1
              |   IRFn ← 0
              | nr ← floor[ $\frac{(N - 1) - OL}{R - 1 - OL}$ ]
              | for rec ∈ 1.. nr
              |   yrec ← submatrix[output, (rec - 1)(R - 1 - OL), rec·(R - 1) - (rec - 1)·OL,
              |   xrec ← submatrix[input, (rec - 1)(R - 1 - OL), rec·(R - 1) - (rec - 1)·OL, C
              |   for n ∈ 0.. M - 2
              |     yrecn ← 0
              |     σy ← stdev(yrec)
              |     σx ← stdev(xrec)
              |     Cxy ← σx·σy·lcorr(xrec, yrec)
              |     Cxx ← σx2·lcorr(xrec, xrec)
              |     for r ∈ 0.. M - 1
              |       for c ∈ 0.. M - 1
              |         Φxr,c ← Cxx|r-c|
              |     invΦx ← Φx-1
              |     Cxy ← submatrix(Cxy, 0, M - 1, 0, 0)
              |     IRFrec ←  $\frac{1}{\Delta t}$  invΦx·Cxy

```



$$\left| \begin{array}{l} \text{IRF} \leftarrow \text{IRF} + \text{IRF}_{\text{rec}} \\ \text{IRF} \leftarrow \frac{1}{nr} \cdot \text{IRF} \\ \text{IRF} \end{array} \right.$$

Input Toeplitz matrix (N-M+1 x M): InputToeplitzM(M, x)

InputToeplitzM( ) puts an input signal into a Toeplitz matrix that represents input when the convolution is broken into a series of equations for elements M-1 to N-1.

$$\text{InputToeplitzM}(M, x) := \left| \begin{array}{l} N \leftarrow \text{length}(x) \\ \text{for } r \in 0..N-M \\ \quad \text{for } c \in 0..M-1 \\ \quad \quad X_{r,c} \leftarrow x_{M-1-c+r} \\ X \end{array} \right.$$

IRF based on series of correlation equations: IRFlinalg(input, output, Δt, lenIRF)

The algorithm calculates the IRF using the pseudoinverse to solve the matrix series of equations that relate outputs M-1 through N-1 to the convolution sum. The function requires the user defined function InputToeplitzM( ) which arranges the input elements into a matrix.

$$\text{IRFlinalg}(\text{input}, \text{output}, \Delta t, \text{lenIRF}) := \left| \begin{array}{l} x \leftarrow \text{input} - \text{mean}(\text{input}) \\ y \leftarrow \text{output} - \text{mean}(\text{output}) \\ N \leftarrow \text{length}(y) \\ M \leftarrow \text{lenIRF} \\ \Delta T \leftarrow \Delta t \cdot \text{identity}(M) \\ \text{inv}\Delta T \leftarrow \Delta T^{-1} \\ Y \leftarrow \text{submatrix}(y, M-1, N-1, 0, 0) \\ X \leftarrow \text{InputToeplitzM}(M, x) \\ \text{inv}X^T X \leftarrow (X^T X)^{-1} \\ \text{IRF} \leftarrow \text{inv}\Delta T \cdot \text{inv}X^T X \cdot X^T Y \\ \text{IRF} \end{array} \right.$$

(MxM) cross-covariance matrix: CrossCovarMatrixM(x, y, M)

Calculates a cross covariance matrix with 0 lag on the diagonal, negative lags in upper triangle, and positive lags in lower triangle. Sequence x is lagged relative sequence y, e.g. Rxy

```

CrossCovarMatrixM(x, y, M) :=
|
|  $\sigma_x \leftarrow \text{stdev}(x)$ 
|  $\sigma_y \leftarrow \text{stdev}(y)$ 
|  $R_{xy} \leftarrow \sigma_x \cdot \sigma_y \cdot \text{lcorr}(x, y)$ 
|  $R_{yx} \leftarrow \sigma_y \cdot \sigma_x \cdot \text{lcorr}(y, x)$ 
| for  $r \in 0..M-1$ 
|   for  $c \in 0..M-1$ 
|      $XY_{r,c} \leftarrow \text{if}[(r-c) \geq 0, R_{xy_{r-c}}, R_{yx_{c-r}}]$ 
|
| XY

```

Percent variance accounted for: VAF%(x<sub>est</sub>, x, M)

```

VAF%(xest, x, M) :=
|
|  $N \leftarrow \text{length}(x)$ 
|  $x_{\text{est}} \leftarrow \text{submatrix}(x_{\text{est}}, M, N-1, 0, 0)$ 
|  $x \leftarrow \text{submatrix}(x, M, N-1, 0, 0)$ 
|  $x_{\text{err}} \leftarrow x_{\text{est}} - x$ 
|  $\lambda_x \leftarrow \text{var}(x)$ 
|  $\lambda_{\text{err}} \leftarrow \text{var}(x_{\text{err}})$ 
|  $\text{VAF} \leftarrow 100 \left( 1 - \frac{\lambda_{\text{err}}}{\lambda_x} \right)$ 
|
| VAF

```

g3 IRF - Single record, covariance, output elements set to zero:  $g3_{\text{corZ}}(M, \Delta t, pt, hv, ev)$

The "input" is the pt-hv ccf, and the "output" is the pt-ev ccf. The first M-1 elements of the ev sequence are set to 0.

```

g3corZ(M, Δt, pt, hv, ev) :=
  N ← length(ev)
  for n ∈ 0..M - 2
    evn ← 0
  σpt ← stdev(pt)
  σhv ← stdev(hv)
  σev ← stdev(ev)
  Cpthv ← σpt · σhv · lcorr(pt, hv)
  Chvpt ← σhv · σpt · lcorr(hv, pt)
  for r ∈ 0..M - 1
    for c ∈ 0..M - 1
      Φpthvr,c ← if[(r - c) ≥ 0, Cpthvr-c, Chvptc-r]
  invΦpthv ← Φpthv-1
  Cptev ← σpt · σev · lcorr(pt, ev)
  Cptev ← submatrix(Cptev, 0, M - 1, 0, 0)
  g3 ←  $\frac{1}{\Delta t}$  · invΦpthv · Cptev
g3

```

g3 IRF - multi-rec, covar, M-1 output set to zero: g3avcorZ(pt, hv, ev, Δt, lenIRF, lenRec, lenOL)

The function breaks apart the input and output sequences into records of length lenRec. This length must be chosen so the record is representative of the system, and large relative to the IRF. Each successive record overlaps the previous record by lenOL terms. Zeros are set in the first M-1 elements of each record. lenOL should be greater than M to ensure all data is considered. The IRF for each record is calculated for lenIRF terms.

The function used the lcorr() function which is a covariance, so the user must account for sequences that do not have zero mean.

```

g3avcorZ(■) :=
  N ← length(ev)
  M ← lenIRF
  R ← lenRec
  OL ← lenOL
  for n ∈ 0..M-1
    IRFn ← 0
  nr ← floor  $\left[ \frac{(N-1) - OL}{R-1-OL} \right]$ 
  for rec ∈ 1..nr
    ptrec ← submatrix[ pt, (rec-1)(R-1-OL), rec·(R-1) - (rec-1)·OL
    hvrec ← submatrix[ hv, (rec-1)(R-1-OL), rec·(R-1) - (rec-1)·OL
    evrec ← submatrix[ ev, (rec-1)(R-1-OL), rec·(R-1) - (rec-1)·OL
    for n ∈ 0..M-2
      evrecn ← 0
    σpt ← stdev(ptrec)
    σhv ← stdev(hvrec)
    σev ← stdev(evrec)
    Cpthv ← σpt·σhv·lcorr(ptrec, hvrec)
    Chvp ← σhv·σpt·lcorr(hvrec, ptrec)
    for r ∈ 0..M-1
      for c ∈ 0..M-1
        Φpthvr,c ← if  $\left[ (r-c) \geq 0, C_{pthv_{r-c}}, C_{hvp_{c-r}} \right]$ 

```

```

invPhi_pthv ← Phi_pthv-1
C_ptev ← sigma_pt · sigma_ev · lcorr(pt_rec, ev_rec)
C_ptev ← submatrix(C_ptev, 0, M - 1, 0, 0)
IRF_rec ← (1 / Delta t) · invPhi_pthv · C_ptev
IRF ← IRF + IRF_rec
IRF ← (1 / nr) · IRF
IRF

```

3053-85