INFERRING CENTRAL MOTOR PLANS FROM ATTRACTOR TRAJECTORY MEASUREMENTS

by

Antony John Hodgson

B.A.Sc. in Mechanical Engineering, University of British Columbia (1984)
M.A.Sc. in Mechanical Engineering, University of British Columbia (1986)

SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY IN MEDICAL ENGINEERING AND MEDICAL PHYSICS

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 1994

Signature of Author ______________________________
Harvard-MIT Division of Health Sciences and Technology
September 1994

Certified by ______________________________
Neville Hogan
Thesis Supervisor

Approved by ______________________________
Roger Mark
Co-Head, Harvard-MIT Division of Health Sciences and Technology

Inferring Central Motor Plans from Attractor Trajectory Measurements

Antony John Hodgson

Submitted to the Division of Health Sciences and Technology on July 6, 1994
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Medical Engineering and Medical Physics

ABSTRACT

A continuing controversy in motor neuroscience is over what aspects of movement are represented and planned centrally (i.e., in the brain). The virtual trajectory hypothesis claims that since the peripheral neuromuscular system appears to translate central commands into stable postures, the brain need only be concerned with the kinematic aspects of a motor task. If this hypothesis were true, we would expect to find that central commands bear a close resemblance to the task description and that the actual limb trajectory would be distorted by dynamic effects. Some have criticized this hypothesis as untestable because a direct test based on interpreting the symbolic content of neural activity is currently beyond our capabilities and indirect tests based on mechanical behaviour have been thought to require models of muscle impedances — a notoriously difficult proposition given their highly nonlinear and time-varying properties.

I have developed an alternative argument which compares the actual trajectory executed by a subject with the trajectory which they would follow if the inertial contributions to the movement could be cancelled (what I call the *attractor trajectory*). I have also devised a novel experimental technique to measure attractor trajectories; in my experiments, the subject interacts with a robotic arm which supplies interaction forces designed to cancel the inertial effects. This technique depends only on having a reasonable estimate of the inertia of the subject's arm and does not require a model of the muscle impedance, so we now have a convincing means to test the virtual trajectory hypothesis.

In point-to-point movement experiments using this technique, I have disproven the virtual trajectory hypothesis by showing that the brain's plans for unconstrained movements are more complex than the actual arm trajectories; the planning process therefore cannot be as simple as the hypothesis suggests, but must incorporate some form of compensation for dynamic effects. However, when I change the speed of the movement or ask the subject to cope with spring-like or coriolis-like forces, I find that the plans are simpler than the actual trajectories and are in fact very similar to the plans for the unconstrained movements. These results indicate that, in the spirit of the virtual trajectory hypothesis, the brain executes a new or unusual task by recycling an already-well-learned plan used in a similar but more common movement.

Thesis Supervisor: Neville Hogan
Title: Professor of Mechanical Engineering & Brain and Cognitive Sciences

To the little nonlinear, time-varying,
chaotic physical system that I'll spend
the rest of my life learning from.

God bless you.

ACKNOWLEDGEMENTS

It feels to me as though I have been working on this thesis the whole time I have been at MIT, though in fact I spent almost half my time there without ever setting foot in the biomechanics lab, inconceivable as that is to me now. It is only because this work has so fully occupied me for the past three or four years that I have such difficulty recalling the earlier times. However, there was a time when working with Neville Hogan would have seemed to me to be a terribly unlikely prospect.

I met Neville in my first term at MIT, when I was preparing for my qualifying exams. I had gone to ask him about the contents of the controls exam — to decide whether or not to take it — and he proceeded to settle the matter, in my mind at least, by giving me a mini-oral on the spot which scared me right into the arms of the heat transfer people. I drifted back after taking Neville's physical systems modeling and control theory course, and eventually persuaded him, aided by claims of self-solvency, to take me on. That return was the wisest choice of my graduate school career. I am most indebted to Neville for what he has taught me about engineering approaches to the study of physical and (the subset of) biological systems, but the way he treats his graduate students as colleagues from whom he expects to learn will always be a model for me of such relationships. Never did I feel that I was working *for* Neville, only *with* him. My only regret is that, with the coming of increasing administrative responsibilities, Neville has lost time for the Friday afternoon sessions in the Thirsty Ear enjoyed by his earlier students, although one should still always block out the 4-7 PM slot for those "one-hour" late-afternoon meetings.

My other committee members have also taught me a great deal. I offer my thanks to Emilio Bizzi for illuminating discussions in my early days in this field as I was taking his seminar course, and for his continuing insights in the following years expressed both in my committee meetings and in his weekly lab group meetings. I have learned a great deal from his prodigious knowledge of neuroscience. Sandro Mussa-Ivaldi has also been invaluable in his close reading and questioning of my arguments. His questions, and those of Jean-Jacques Slotine, forced me to develop the whole notion of an attractor trajectory in order to explain what I really thought, and so brought a measure of theoretical rigour to this otherwise strongly experimental thesis.

Since this is an experimental thesis, I could not have succeeded without the free and willing participation of my friends. Denis Rancourt, Justin Won, Nancy McArdle and Andy Kalt all put up with "the chair" in the early days when my experimental protocol was grindingly tedious (as opposed to *merely* tedious) and the software unreliable. The success of this thesis owes much to their good grace and cheerfulness. I am also grateful for the later participation of Karon MacLean, Cody "Outlying Data Points" Laumeister and Elizabeth Anthony; they helped me refine my experimental technique and added depth to my results. I hope that they all feel the Tosci's ice cream was adequate recompense. Thanks also to my labmates for friendship, helpful critiques of my work, and supporting the lab's infrastructure: Ernie Fasse, Reza Shadmehr, Joe Doeringer and Igo Krebs.

My deepest thanks go to my wife, Elizabeth, for keeping me more sane than most people who do PhD's at MIT are allowed to be. Her encouragement has pulled me back from the edge of despair when my experiments weren't working (again!), and her love has given me joy in the hardest times. Perhaps what I valued most of all when we were both in the midst of our doctoral programs was the commitment we had to having dinner together — those times affirmed that work was not supreme — the ever-present temptation of MIT.

In the same vein, I thank all my friends at St. James's for their unfailing support and for showing me how lives of service and caring are lived. I hope that I have learned these lessons well. A very special thanks to all those who gave me a place to stay when I was working in Boston after Liz started her professorship in South Carolina: Justin Won and Joe Doeringer, Denis Rancourt, Lisa Montouri, Pat Michaels and Laurie Rofinot, Nancy Nies and David Greeley, Nancy McArdle, Cody Laumeister, Lisa Brandes, Dave Woodward, and Elizabeth Anthony. I'm only sorry that I wasn't freer during those visits to enjoy their company more.

As a student in the Medical Engineering program of the Division of Health Sciences and Technology, I have led a split life, shuttling from one end of campus to the other and occasionally leaving it altogether to spend time at Harvard Medical School and various hospitals. Thanks to all the people in HST and mechanical engineering who handled the details of my administratively rather complex life — Pat, Leslie, Patty, Gloria, Elizabeth, Sally, Barbara and Chris. Special thanks to Keiko: I entered HST in the heyday of the Keikodollar, and my life is richer for it.

Finally, I wish to express my great appreciation for the constant love, support, and prayers of my mother and of Liz's parents. Though distant, they have been present for me through their frequent phone calls and letters. I love you all dearly.

This research was performed in the Eric P. and Evelyn E. Newman Laboratory for Biomechanics and Human Rehabilitation in the Department of Mechanical Engineering at the Massachusetts Institute of Technology. I have received financial support from the Natural Sciences and Engineering Research Council of Canada, the Medical Research Council of Canada, the National Institutes of Health (Grant No. AR-40029), the Whitaker Health Sciences Fund, a Poitras Predoctoral Fellowship, and the Johnson & Johnson Research Fund.

TABLE OF CONTENTS

CHAPTER
ONE

INTRODUCTION

The hand is the visible part of the brain - Immanuel Kant

Motivation

Although western society has long espoused a mind-body duality, exalting the activity of the first and despising the demands and constraints of the second, it is clear that from a neurophysiological point of view there is no sharp distinction to be drawn between these two aspects of our being. Indeed, the greater portion of our brain is devoted to processing sensory information and producing motor activity in response to those perceptions. As Kant implies, this motor activity is the only means by which the cogitations of our brain can be made manifest in the external world. Since human beings are particularly adept at interacting with the world around them and manipulating objects in it, one of the most important goals of studies in motor neuroscience is understanding how the brain is able to execute the complex computations necessary to perform these tasks.

Such studies are important for two primary reasons. The first is the potential medical benefits: the more we understand about the processes involved in performing motor tasks, the better our insight will be into the nature of neurological disease and injury and the better able we will be to devise new diagnostic and treatment tools, including improved prosthetic devices intended to replace lost neuromotor function. The second reason for studying an obviously successful manipulator - the human - is to draw inspiration for the design of robotic devices. By understanding the principles by which humans have solved a variety of complex manipulation problems, we will be better able to develop robots with improved dexterity and motor abilities. The primary question addressed in this thesis is that of which aspects of a motor task the brain represents in the planning process. This question is posed in the context of a controversial hypothesis - the virtual trajectory hypothesis - which asserts that the intrinsic behaviour of the peripheral neuromuscular system exhibits strong stability properties which the central nervous system can exploit to simplify the control of movement. One can test such a hypothesis in two complementary ways. First, one can analyze the brain's computational processes directly; by performing detailed neuroanatomical studies and recording the activities of neurons during a variety of tasks, one can attempt to infer the symbolic content of the neural processes. For example, Georgopoulos and his colleagues have determined that the overall behaviour of a large population of neurons in the motor cortex can be viewed as coding for the direction of a desired movement expressed in Cartesian coordinates [Georgopoulos 83, 86, 88, 94, Lukashin 93, Schwartz 93]. The second means of testing motor control hypotheses is through observations of the mechanical behaviour of a subject performing a variety of motor tasks; these studies are known as psychophysical or behavioural experiments. For example, Morasso demonstrated that hand trajectories for tasks in which the subject was asked to move from one point to another were remarkably invariant when they were normalized and expressed in Cartesian coordinates, whereas they showed significant variability and structural differences when expressed in joint coordinates [Morasso 81]; Morasso used this result to argue that, at some stage in the motor planning process, the task itself must be represented in extrinsic or task-relevant coordinates. In other words, the planning process is fundamentally abstract and not tied to the particular set of actuators (muscles) or kinematic linkage (limb) which will ultimately be responsible for implementing the plan. This conclusion is supported by studies of the similarities between handwriting performed using a variety of "execution devices" including hands, feet and mouth [Lacquaniti 89, Morasso 82]. An important component of this thesis will be a detailed analysis of what one can learn about the inner workings of a physical system through such considerations of its mechanical behaviour.

The Nature of the Central Representation

The central claim of the group of ideas which have come to be called the equilibrium point or virtual trajectory hypotheses[1] is that the central nervous system is primarily concerned with generating this abstract, kinematic[2] representation of a motor task, given a task specification, and that the function of the spinal and peripheral neuromuscular system is to translate the kinematic representation into an appropriate set of neural activations [Feldman 66, 76, 86, Bizzi 82, 84]. In other words, there is a partitioning of responsibility for the control of a motor task; the function of the CNS is to fill in the kinematic details of the motor plan given a relatively sparse specification of the task[3] and then to transmit that more complete description to the spinal column.

This virtual trajectory idea is in opposition to the view that the central nervous system is concerned with the dynamics of movement - that is, with the muscle forces needed to produce motion and with the interactions between limb segments which arise due to inertial and velocity-dependent effects such as coriolis and centripetal terms in the equations of motion. This latter view is implicit in any theory of movement which proposes that tasks are optimized according to some dynamic criterion (e.g., minimum

---

[1] I will define the terms equilibrium point and virtual trajectory shortly. For now, I simply use them as names for motor planning hypotheses based on kinematics.

[2] The term *kinematic* refers to the positions and velocities of moving objects, but disregards forces and inertial properties. In contrast, the term *dynamic* includes these effects.

[3] E.g., in a point-to-point movement task, the specification might contain the locations of the points, as well as the desired initiation time and speed of the movement. The filled-in representation would describe either the end-point trajectory or perhaps even the complete limb trajectory, depending on whether the kinematic redundancy issue is handled centrally or not (this question is not addressed in this thesis).

rate of change of joint torque [Kawato 90, Uno 89] or minimum magnitude of muscle force [Dornay 92]). If such a criterion is actually at the heart of the motor planning process, then the central nervous system is in effect required to solve the inverse dynamics problem [Hollerbach 87], which involves creating some internal representation of the relationship between neural activity and the forces ultimately produced by the muscles. This is an extraordinarily difficult computational problem because of the presence of muscle activation dynamics, the complexity of the relationship between the force a muscle generates and its length and velocity, the changing moment arms of the muscles, the existence of more muscles than are needed to generate arbitrary torques at each joint (actuator redundancy), the existence of more degrees of freedom at the joint than are required by the task (kinematic redundancy), and the increasing number of inertia-related interactions between limb segments as the number of involved segments rises. Even in the field of robotics, where it is reasonable to model the actuators as torque sources, thus reducing the complexity of the computational problem, this has proven to be a difficult task as the number of degrees of freedom in the robotic arm approaches that of a human arm.

If the virtual trajectory idea is correct, the subspinal neuromuscular system must exhibit strong enough stability properties to ensure that for routine motor tasks there is a reasonable correspondence between the commanded and actual limb configurations. If such stability properties existed, then the central nervous system could control posture by learning which set of neural activations correspond to which limb configuration, and could generate movement simply by shifting the neural activation pattern from the starting posture to the final posture. Such a control technique would radically simplify the motor computations, although, since dynamic effects are neglected, the actual trajectory will deviate from the planned trajectory because of their influence. Is there then any reason to believe that the virtual trajectory concept is in fact a reasonable description of the way our neuromuscular system is structured? To answer this question, it is helpful to review the history of the development of the idea and to discuss some of the current critiques of it.

History of the Virtual Trajectory Concept

The notion of a kinematic basis for the control of movement was born out of a consideration of the properties of muscle. In contrast to the torque motors commonly used to power robotic devices, muscles exhibit complex length-tension properties; in effect, they behave as nonlinear springs whose rest length (and to a lesser extent, stiffness) depends on the degree of neural activation. Feldman was the first to propose that posture at a joint could be viewed as an equilibrium between the length-tension curves of the agonist and antagonist muscles operating at the joint, and that movement could be generated by shifting the activations of the muscles such that a new equilibrium was established at a new point [Feldman 66]. The actual trajectory followed by the limb would depend on the interplay between the limb's inertia and the muscle forces elicited by the deviation between the current limb position and the equilibrium position specified by the activations of the muscles. In effect, Feldman painted for us a picture of the actual arm being linked by springs (the muscles) to an imaginary arm whose position was directly controlled by the central nervous system. This picture was in striking contrast to the view that the brain computes and controls the forces generated in various muscles - a much more robotic perspective - and so has become a challenge to those who believe that the dynamics of motor tasks are explicitly represented centrally.

Since the early work of Feldman, advocates of the equilibrium point concept, as it was originally called, have divided into two schools of thought. One, headed by Feldman, has pursued the question of the role of reflexes in the generation of muscle forces and has sought to explain the origin of electromyographic (EMG) patterns during movement (perhaps most notably the triphasic burst pattern observed during rapid movements) [Berkinblit 86, Hasan 89, Karst 90, Latash 91, Gottlieb 93]. The second school, headed by Bizzi, has turned its attention to the question of how the CNS utilizes the springlike behaviour of muscle in the control of movement, particularly the movement of multijoint limbs [Bizzi 82, Mussa-Ivaldi 85, Hogan 85, Giszter 91, 93]. Bizzi and colleagues have shown, for example, that reflexes are not absolutely necessary for the stable control of the arm, although their presence certainly enhances performance [Bizzi 84]. In their experiments, they showed that a deafferented monkey was still capable of generating stable pointing movements, even when the monkey's arm was significantly perturbed at arbitrary points in the motion. These experiments also demonstrated that the shift in posture occurred over an interval of several hundred milliseconds, long enough that the movement should be considered controlled throughout its duration; this finding was counter to the claim of the so-called "final position" hypothesis which argued that such transitions were accomplished by an instantaneous switch of the activations from one state to another. For similar reasons, this finding argued against a particular (though not necessary) formulation of Feldman's lambda hypothesis which suggested that the centrally-generated shift occurred over a time scale of approximately 100 ms.

Although Bizzi et al. were not the first to do so [Granit 70, Stein 74], their work has also disproven an earlier hypothesis, known as Merton's servo hypothesis [Merton 53], which claimed that the CNS regulated position by specifying the gamma motor drive and allowing the resulting discrepancy between intra- and extrafusal fiber lengths to elicit alpha motor activity to reduce the discrepancy. In particular, Hogan et al. [Hogan 87] pointed out that the feedback delay times are too long for any significant reflex gain to avoid causing instability; the role of reflexes in generating motion must therefore be limited to low-frequency behaviour.

Regardless of relative contributions of intrinsic muscle properties and feedback to the response of a limb to imposed perturbations, the response itself must exhibit certain characteristics if it is to be considered stable. Most importantly, the mechanical impedance of the limb must be such that it appears to only temporarily store energy, rather than to be tapping an indefinite supply. Mussa-Ivaldi et al. performed a series of two-dimensional stiffness measurements by moving a subject's hand a short distance from a nominal position and measuring the resulting forces when the hand came to rest [Mussa-Ivaldi 85]; when they fitted this data to a stiffness matrix, they found that the off-diagonal terms were essentially identical — a condition necessary for springlike behaviour and compatible, assuming the presence of dissipation, with the requirement of a passive impedance response. If the off-diagonal terms were different, the system could supply energy indefinitely if it were taken along a closed path of the appropriate orientation around the equilibrium point. Hogan pointed out that this passivity property was not guaranteed to be found because it depends on a fine balance between heterogeneous reflexes (i.e., the activation of a muscle at one joint due to the afferent responses from a muscle at a different joint) [Hogan 85]. Similar springlike force fields has been measured independently by Shadmehr, Won, and Dolan [Shadmehr 93, Won 93, Dolan 91].

So, the various experiments of Feldman, Bizzi and Mussa-Ivaldi established that the length-tension behaviour of muscle was such that a constant level and pattern of central neural activity would create an equilibrium configuration for the limb which was both stable and passive. Collectively, these investigators also refined the notion that the CNS could generate movement by shifting activation patterns from the one corresponding to the starting position to the one corresponding to the final position. Hogan introduced the term virtual position to distinguish between Feldman's definition of the equilbrium point as that configuration at which the arm would be at rest when interacting with a constant environmental force (such as lifting a weight, say) and Hogan's own view that the attractor behaviour of a system should be defined independently of interactions with the environment [Hogan 84]. To that end, he defined the virtual position as that position the arm would assume in the absence of inertial or viscous effects and with no interaction forces applied. He then defined the virtual trajectory to be the time history of virtual positions. Given these definitions, the question which came to be called the virtual trajectory hypothesis was posed: does the central nervous system exploit the stability properties of the limb by representing motor tasks primarily (or even solely) in terms of the virtual trajectory?

Mussa-Ivaldi viewed the virtual trajectory hypothesis as drawing together two complementary research threads: studies of muscle elasticity and studies of movement planning [Mussa-Ivaldi 88]. The evidence that movements are planned in kinematic terms comes from several quarters. First, a number of studies provided evidence that tasks were planned in an extrinsic (or task-relevant) coordinate frame [Morasso 81, 82, Abend 82, Atkeson 85, Flash 85, Cruse 87, Uno 89, Flanagan 93, Wolpert 94a,b]. Morasso's work, in particular, showed that two-dimensional hand trajectories in a point-to-point reaching task were extremely regular when viewed in Cartesian coordinates but showed no consistent pattern and were much more complex when viewed in a joint-angle-based coordinate frame; for example, in Cartesian coordinates, most trajectories of unconstrained reaching movements are relatively straight lines, whereas the coordinate changes for the same movements expressed in terms of joint angles are not even necessarily monotonic. Building on Morasso's early work, Flash noted that even when viewed in Cartesian coordinates, the hand paths were not perfectly straight but instead exhibited characteristic deviations or curvatures which depended on the location of the movement in the workspace [Flash 87]. She hypothesized that the curvature in the trajectories was primarily due to the failure of the motor planning process to compensate or account for the dynamic forces which arose during motion and predicted that the virtual trajectories would be straighter than the actual trajectories. Flash measured a broad range of point-to-point reaching movements at a variety of speeds ranging from slow (1.0 s) to fast (0.4 s) at a number of locations across a two-dimensional workspace. She then used the stiffness measurements of Mussa-Ivaldi and certain assumptions about the damping behaviour of the arm[4] to compute the virtual trajectory, $x_o(t)$, as defined by the following model:

$$M(x)\ddot{x}+C(x,\dot{x})+B\dot{x}+K(x-x_0)=0.$$

In general, the computed virtual trajectories were much straighter than the actual trajectories, particularly for movements with durations of greater than 0.5 s, so, using this model, Flash took the virtual trajectory computed for a single trajectory measured at one particular location in the workspace and applied it (after suitable scaling, rotation and translation) to the simulation of the other movements which she had measured. In general, she found extremely good agreement between the predicted and measured trajectories; this simple linear-impedance model was able to reproduce the fine details of the actual trajectories caused by the dynamic effects. Despite the apparent success of her

---

[4] Flash assumed that the stiffness could be scaled in magnitude to provide an optimal fit to the model, but not changed in shape or orientation from its postural values. She also assumed that the damping was relative to ground and that its shape was either the same as the stiffness ellipse or was the geometric mean between the inertia and stiffness ellipses. See Dolan [1991] for a discussion of and evidence in favour of the reasonableness of these assumptions.

simulations, however, Flash's results were widely criticized. Kawato, for example, argued that the stiffness estimates Flash used were unrealistically high, which would necessarily keep the computed virtual trajectory close to the actual trajectory [Kawato 92, 93]. Kawato claimed that recent stiffness measurements in one-dimensional movements by Bennett and Xu showed that not only were stiffnesses during motion comparable to or smaller than those measured at posture, but they in fact dropped mid-movement [Bennett 90, 92, Xu 91]. Since Flash used stiffness values several times larger than postural values, Kawato felt that she was severely underestimating the deviation between the actual and virtual trajectories. Using stiffness estimates derived from Bennett's measurements, Kawato computed the virtual trajectories which would be required to generate straight movements at various locations in the workspace and showed that these trajectories were extremely complex - certainly far more complex than the actual hand trajectories. He felt, therefore, that the virtual trajectory hypothesis could not provide a means for simplifying the control of movement.

A second aspect of Kawato's criticism was that the situation would be even worse as the speed of movement increased, particularly if the stiffness of the arm did not change with speed. Milner, however, has demonstrated recently that elbow stiffness in a one-dimensional task does increase with speed of execution, and Saund showed some time ago that the perturbation response of the elbow is essentially invariant when normalized according to execution speed [Milner 93, Saund 84]. Hollerbach determined that one can preserve the same relationship between the virtual trajectory and the actual trajectory if the damping matrix scales in proportion to speed and the stiffness in proportion to the square of the speed [Hollerbach 82]; Milner's work suggests that such scaling occurs to some extent, while Saund's study indicates that the scaling may be almost precisely what is needed to preserve the trajectory relationships. This result is counterintuitive to some researchers; Hasan, for example, has stated that high muscle stiffness would be energetically inefficient [Hasan 92].

A deeper concern with the virtual trajectory idea was voiced by Smith and Humphreys, who felt that the concept was either an obvious statement of what must be or else it was fundamentally untestable [Smith 91]. The first part of their concern stems from the observation that, if the limb is stable, then one can always define an equilibrium position, so its existence can hardly be considered a hypothesis. As I've discussed above, it is not true that such equilibria must exist, although our observations of intrinsic muscle behaviour and the aggregate response of whole limbs do suggest that they do. However, mere existence of equilibrium points is not the essence of the virtual trajectory hypothesis, but computation based on them. The second part of Smith's concern is that the hypothesis is not physiologically rooted - since the virtual trajectory is not identified with any neurophysiological structures or signals, they claim that we will not be able to devise a test of the hypothesis. As we will see, though, such tests have been created.

Hasan, for example, also felt that the virtual trajectory hypothesis was inadequate as a hypothesis because it did not make any direct predictions about EMG activity [Hasan 89]. In particular, he did not see how the hypothesis would have anything to say about the

EMG activity elicited by movements in different directions from a common starting point (this experiment was performed to test a number of hypotheses relating to the activation of muscles to generate a movement in a given direction). Shadmehr, however, analyzed Hasan's data and showed that all of Hasan's hypotheses were contradicted by that data, whereas the data was in agreement with the hypothesis that the movement was generated by a shift in the equilibrium point from the starting point along a vector oriented in the specified direction [Shadmehr 91].

A number of cortical recording studies can also be interpreted as potential tests of the virtual trajectory hypothesis. For example, Georgopoulos' demonstration that the Cartesian direction of impending motion is represented by an ensemble of cortical neurons suggests that the primary representation of the motor plan is kinematic; if we found similar evidence of central representations of muscle forces or of the dynamics of interactions between different limb segments, I would regard such evidence as disproof of the hypothesis.

Testing the Virtual Trajectory Hypothesis

It is clear then that the virtual trajectory hypothesis is not fundamentally untestable, but simply difficult to devise tests for. Smith and Humphrey's concern that the hypothesis is not physiologically grounded, in the sense that there is no biological structure which we can identify with the virtual trajectory, is important because it points out that we will not be able to actually measure the virtual trajectory. Flash's definition of the virtual trajectory is fundamentally expressed in terms of a model of the impedance of the limb, predominantly the impedance of the muscles which power the limb. Since the behaviour of real muscle is highly nonlinear and time-varying, a simple linear model of its behaviour such as the one Flash used cannot yield produce a quantity which has physiological significance. In other words, in terms of our picture of the control of the arm as the control of an imaginary arm linked to the real arm by (nonlinear) springs, Flash's estimate of the location of the imaginary arm depends critically on her model of the springs; if we get their behaviour wrong, as we certainly must given the complexity of real muscles, we may be able to compute a location for the imaginary arm, but its relationship to the position of the imaginary arm in the actual central representation cannot be determined.

An example of the difficulty in creating an adequate model is discussed by Won, who attempted to find a set of linear impedance parameters and a virtual trajectory which explained both an unconstrained point-to-point reaching movement (in the same vein as Flash) and a corresponding constrain-and-release experiment [Won 93]. Won found that he could accurately reproduce either one of the trajectories considered singly, but that he could not find a single combination of impedance parameters and virtual trajectory which would reproduce both measurements. This indicates that much more complex models of the mechanical behaviour of muscles, tendons, joints, and other arm structures would be needed before any model-based test of the virtual trajectory hypothesis would be generally accepted.

Given that such models are not yet sufficiently sophisticated, would it be possible to devise an alternative, model-independent test of the hypothesis? In the remaining chapters of this thesis, I will do precisely that. In chapter two, I address the question of what we can learn about the internal organization of a physical system based only on measurements made at its interaction ports.[5] I introduce the concept of the attractor trajectory in order to formalize our notion of the attractor behaviour of a system. Until now, we have assumed that the definition of the virtual trajectory adequately expresses this notion, but I show that this assumption is incorrect because the definition relies on a necessarily incorrect or incomplete model of the muscle. In contrast, the definition of the attractor trajectory is based solely on measurable quantities and does not require any model of the muscle. Also until now, statements of the virtual trajectory hypothesis have been relatively vague and critics like Smith have been justified in claiming that the statements do not make testable predictions related to the motor planning process. Flash has come closest to making such statements by arguing that the same virtual trajectory (invariant under rotation, translation and scaling) underlies the generation of a broad range of related movements, namely, point-to-point reaching movements. I consider what a kinematics-based motor-planning strategy would predict about the shapes of various attractor trajectories and present both strong and slightly weaker versions of the hypothesis, along with new tests based on measurements of attractor trajectories. In essence, the new statements claim that the central representations of a motor task are kinematic in nature

---

[5] A port is a connection between a system and its environment across which energy can be exchanged.

and are either simpler (strong claim) or more consistent (slightly weaker claim) than the actual trajectories executed by the arm.

The tests of the new statements of the virtual trajectory hypothesis require measurements of the attractor trajectory of the arm when all inertial effects are eliminated. In the third chapter, I explain how such measurements can be made and outline the particular approach I use for my experiments.

In the fourth chapter, I describe in detail my experimental protocol. I examine a wide variety of point-to-point reaching movements made in different parts of the workspace at speeds ranging from 30-80% of a subject's maximum possible speed. I also consider both unconstrained movements (in which the subject moves an unpowered manipulandum) and those in which the subject is required to interact with a dynamic environment (either a position- or velocity-dependent force field).

In the fifth chapter, I present the results of the experiments and argue that neither version of the virtual trajectory hypothesis is correct. Finally, I discuss in the sixth chapter the implications of these results for the virtual trajectory idea and demonstrate that the results do not support the views of either its proponents or opponents, but perhaps point the way towards a deeper understanding of the processes by which we develop new motor skills and build upon our existing ones.

CHAPTER

TWO

FORMULATION OF HYPOTHESES

Overview

Our goal is to study the means by which the central nervous system represents, plans and executes motor tasks. Our methodology is mechanical interrogation of our human subjects, which means that we place our subjects in a controllable dynamic environment (implemented using a robotic arm) and measure the paths taken and forces generated during execution of the experimental tasks. In particular, we wish to identify and explore those paths which have some sort of "attractor" behaviour. On the basis of the experimental measurements, we intend to draw conclusions about possible strategies which the CNS might be exploiting to successfully perform the tasks. We must therefore ask ourselves two main questions: (1) how can attractor behaviour be measured, and (2) what is the relationship between our measurement and the neural operations underlying the generation of movement?

I will show that for systems which exhibit certain stability properties (there is evidence that the human arm possesses such properties), trajectories with attractor properties can be defined. These *attractor trajectories* will be well-defined at each interaction port in terms of the observed port displacements and will furthermore be independent of assumptions about the internal structure of the system, which implies that no interaction port measurement can give us information about the internal structure unless we make assumptions about or have knowledge of that structure. This statement is related to the claim of Norton's theorem for one-port linear systems that any arbitrary collection of linear elements and time-varying current and voltage sources can, from the point of view of interaction port behaviour, be replaced by an equivalent circuit consisting of a single time-varying current source connected in parallel with a complex impedance. The attractor trajectory which we measure in our experiments is a multiport nonlinear analog to the one-port linear Norton flow source. We need such a definition and measurement of attractor behaviour to avoid the criticisms leveled at those justifications of the virtual trajectory hypothesis (such as Flash's) which rely on oversimplified and physiologically unrealistic muscle models.

Once we have estimates for the attractor trajectories in a variety of tasks, we will want to relate the characteristics of the trajectories to the predictions of various hypotheses of motor control, particularly the so-called *virtual trajectory* hypotheses. Since these hypotheses make statements about the form or meaning of the neural commands to the motor system, it would seem that we would have to have a sufficiently detailed model of that system to allow us to compute from interaction port measurements what the command must have been. Fortunately, this is not the case. It will be sufficient to compare the characteristics of the attractor trajectory of the subsystem which excludes inertial effects with those of the unperturbed movement (which is the attractor trajectory of the system with inertial effects included). Since the inertial effects are not a function of the difference between the interaction port position and the desired (commanded) position, they can only degrade the correspondence between the commanded and the achieved trajectories. The virtual trajectory hypotheses claim that the commanded trajectories are more consistent than the achieved trajectories, so we would predict that a comparison of the achieved trajectories with the attractor trajectories of the subsystem with inertial effects excluded would show that the subsystem attractor trajectories are more consistent since they are closer to the commanded trajectories. If the comparison did not show this, then we have a strong rejection of the virtual trajectory hypotheses because we have considered the attractor trajectory most favourable to the hypotheses — the one linked to the achieved trajectory through the shortest chain of dynamics. This argument does not require us to assume anything about the real motor system except that the inertial effects are a function of absolute rather than relative velocities — a simple consequence of Newton's laws of motion.

In what follows, I will lay out in some detail the terms used in this thesis, elaborate on the question of what we can learn through mechanical interrogation of the human arm, and discuss the issues of how our measurements relate to the various motor control hypotheses. I begin by looking at the terms which already exist in the literature.

Inclusion or Exclusion of Environment

Although the concept of an equilibrium point has become well-established in the motor control literature, there are still a number of conflicting definitions of it, as well as of the associated terms *virtual position* and *virtual trajectory*. In the static domain, Feldman makes a distinction between the λ value, which represents the stretch reflex threshold, and the equilibrium point, which he defines as the position at which the muscle is in equilibrium with the (static) external load. In other words, his definition of the equilibrium point involves the characteristics of the external load. Given a particularly unstable load, Feldman would say that no equilibrium point exists [Feldman 66]. In contrast, Bizzi *et al.* have defined the virtual position to be that position or configuration which a system would settle into in the absence of external loads [Bizzi 92]. According to this view, the dynamic nature of the environment is irrelevant to the definition of the virtual position; if a limb is stable in isolation, coupling that limb to a highly unstable environment may well lead to an unstable interaction, but the limb's virtual position remains well-defined. If the limb is placed in a gravitational field, but no other loads are applied to the limb, the virtual position might well differ from the equilibrium position. I believe that these are useful distinctions to draw, so I will use Feldman's term, equilibrium point, when I consider overall interactive behaviour, but I will use the term virtual position to refer to a property of a system treated independently of its environment. However, the term *virtual trajectory* as defined by Bizzi is not precise enough for my purposes, so I will now propose an alternative definition expressed in terms of measurable quantities.

The Informativeness of Mechanical Interrogation

Defining a System's Attractor Point

My definition of a virtual trajectory is built upon the notion of the attractor point of a system — that point toward which the system is tending at each instant. This attractor point should be definable solely in terms of the system's interaction port behaviour and its definition should, if possible, make no assumptions about the internal structure of the system (we will see later that this is an impossible requirement). For an isolated system, it has been argued that it is reasonable to identify the attractor point with the vector of interaction port positions if that vector is static. If it is moving, it is less clear how the moving attractor point (the *attractor trajectory*) should be defined. The following set of definitions constitute a formal definition for the attractor trajectory.

Definition 1. Let S be a time-varying state-determined system whose interaction port behaviour can be characterized by the vector triple $\{x,v,F\}$ representing position, velocity and force, respectively, at each interaction port.

Definition 2. A system S is said to be in a state of internal equilibrium if all internal effort sources are set to a constant value, all flow sources are set to zero, all internal states are unchanging with time and the interface forces are all zero.

Definition 3. Let S be in internal equilibrium at $t = t_i$. Let G be the set of all interaction port trajectories, $x(t),\ t \geq t_i$, for which F is identically zero for $t \geq t_i$.

Definition 4. A trajectory $x_G(t)$ in G is said to be uniformly stable if for any t_o and for any $R > 0$ independent of t_o, there exists a positive scalar $r(R)$ such that

$$\|x(t_o)-x_G(t_o)\| < r \text{ and } F(t) = 0,\ \forall\, t \geq t_o \implies \|x(t)-x_G(t)\| < R,\ \forall\, t \geq t_o.$$

Definition 5. A trajectory $x_G(t)$ in G is said to be uniformly asymptotically stable if it is uniformly stable, and if there exists a positive scalar R independent of t_o such that

$$\|x(t_o)-x_G(t_o)\| < R \text{ and } F(t) = 0,\ \forall\, t \geq t_o \implies \|x(t)-x_G(t)\| \to 0 \text{ as } t \to \infty$$

and, for any R_1, R_2 satisfying $0 < R_2 < R_1 < R$, there exists $T(R_1,R_2) > 0$ such that

$$\|x(t_o)-x_G(t_o)\| < R_1 \text{ and } F(t) = 0,\ \forall\, t \geq t_o \implies \|x(t)-x_G(t)\| < R_2,\ \forall\, t \geq t_o + T(R_1,R_2).$$

Definition 6. The set A is the set of all uniformly asymptotically stable trajectories in G. If A is the null set, no attractor trajectory is said to exist; if A is not null, any trajectory in A is said to be an *attractor trajectory.*

In effect, the attractor trajectory is defined to be any uniformly asymptotically stable trajectory of the interaction ports of the system which evolves from an equilibrium point in the absence of interaction forces. Note that the definition of stability of a trajectory does not allow one to test the stability of a system during a single excursion from an equilibrium state; rather, it requires a comparison of an unperturbed and a perturbed trajectory and therefore assumes that the commands underlying the execution of the two trajectories are the same in both cases. An alternative test is proposed below, and, though it does not get around this requirement of comparing trajectories with the same underlying commands, it may give somewhat more insight into the strength of the stability of the trajectory.

Conjecture 1: Let $\delta x(t)$ be a deviation from any trajectory, $x_a(t)$, in A. If $\delta x(t)$ is bounded by the maximum r for which the corresponding trajectory, $x_a(t)$, can be shown to be uniformly asymptotically stable, the following inequality holds:

$$E_v = \int_{t_i}^{t_o} F^T \delta \dot{x} \, dt \geq 0.$$

A necessary but not sufficient condition for the equality to hold is that $\delta x(t_o) = 0$. The basis for this conjecture is the notion that causing any deviation from a uniformly stable trajectory requires an input of energy to the system. If the system exhibits no dissipative behaviour, then it is possible to completely recover all the energy, but only if the deviation from the original trajectory is brought to zero. My concern with the theoretical foundation for this conjecture is that the energy required to generate a deviation might not produce a Lyapunov function from which we can deduce the posited stability of the trajectory. Nonetheless, if one could extract energy from the system by taking it through a deviation from a given trajectory, I would claim that this is strong evidence that the trajectory is not stable and therefore not an attractor trajectory.

The Problem of Finite Observation Time

According to Definition 3 above, one way to measure an attractor trajectory is to begin observing the system when it is in an equilibrium state and to ensure that F remains zero ever after. If we begin our observations at some later time, however, at which we cannot verify the equilibrium condition, we may misinterpret a perturbation response as an attractor trajectory.

Consider, for example, a system comprised of a spring and a damper connected in parallel between a reference flow source and the interaction port. The interaction port position of such a system will precisely track the flow source if the interface force is zero. Suppose the attractor trajectory is $x = 0$, $\forall t \in (-\infty,\infty)$, and that a perturbation has been applied to bring the interaction port to a position $x = 1$ at $t = 0$; for all positive times, $F = 0$ and the interaction port will move along an exponentially-decaying trajectory which approaches $x = 0$ as $t \rightarrow \infty$. If we assume (incorrectly) that F has been zero since some assumed time at which the system was in equilibrium, then we would conclude that the current decaying trajectory belongs to G. We might then decide to test whether or not this measured trajectory belongs to A by applying various perturbations beginning at time zero, in accordance with Conjecture 1.[6] It is straightforward to show that for a linear system E_v will be positive for any deviation from the decaying trajectory, so we would conclude that the decaying trajectory was indeed an attractor trajectory. We would be mistaken in this conclusion, however, because in order to get to the given state at $t = 0$, a non-zero force must have been applied to the system at some time prior to that, so the observed trajectory is not a portion of a trajectory in G and so cannot belong to A, which is a subset of G. We make this mistake because we cannot begin our observations in the infinite past and therefore miss the perturbation applied prior to the start of our observations.

Having shown that observing trajectories for finite times *can* lead us to incorrectly classify them as attractor trajectories, is it nevertheless possible to claim that observing particular patterns of interaction port behaviour allows us to conclude that the system is moving along an attractor trajectory? Consider, for example, the following conjecture which suggests that zero port acceleration and velocity at a particular time is an indication of being on a (non-moving) attractor trajectory.

Conjecture 2. If $v(0) = 0$ and the acceleration of the interaction port is also zero, then if $F(t) = 0, \forall t > 0$, the resulting trajectory $\{x(t),v(t)\}$ will be the positive-time portion of an attractor trajectory.

---

[6] We assume either that the system can be restarted and put into the same state it was in when we first began observing it (the ergodicity assumption), or we do a thought experiment and pretend that we have a number of identical systems which have all gone through the same history throughout the infinite past. In either case, we apply various perturbations at time zero.

This conjecture claims that if the initial velocity and acceleration of the interaction port are both zero, the external responses to any earlier perturbations have completely died out and there is no chance that any residual internal effects will evidence themselves at the interaction port. Unfortunately, a simple counterexample will demonstrate the falsity of the conjecture.

Consider a system comprised of a spring and a damper connected between a flow source and a mass; the interaction port is the point of application of a force to the mass and the port's velocity is the velocity of the mass. Suppose the mass is originally resting at $x = 0$ and that the flow source has also been at $x = 0$ for the semi-infinite past. At time $t = -\delta t$, the flow source shifts from $x = 0$ to $x = 1$ over the relatively short interval, δt. If no perturbations where applied, the mass would respond to the spring and drag forces and would eventually settle out at $x = 1$ at some positive time. This unperturbed transition path would be the attractor trajectory for the system. If a perturbation were applied to move the mass to $x = 1$ at $t = 0$ and leave it there with zero velocity, there would be no tendency for the mass to move. If we began our observations at $t = 0$, we would observe zero velocity and acceleration not only at $t = 0$, but at all times, $t > 0$, and so, according to Conjecture 2, would conclude that the attractor trajectory is $x(t) = 1$, for all $t > 0$. By construction of the example, however, this would not be the case; the attractor trajectory does not reach $x = 1$ until some positive time distinctly greater than zero. The conjecture is therefore false.

We can furthermore demonstrate through a consideration of linear systems that no observation of interaction port behaviour will allow us to distinguish between a system in a state of internal equilibrium (which will express itself in equilibrium at the interaction ports) and a system with an internal source varying in such a way as to exhibit apparent equilibrium when viewed through the interaction ports.

Theorem 1. If a system S has multiple interaction ports and the velocities and forces at all of the interaction ports are zero over a finite time interval, $[0,T]$, the port position vector $X(T)$ is not necessarily on an attractor trajectory of the system.

Proof: The following proof constructs two linear systems with identical (zero) port velocities in an interval [0,T] but with different attractor trajectories. Since the port variables in the interval satisify the premise of the theorem, but the two attractor trajectories differ at time T, the theorem is proven.

Let $x_1(t)$ be an attractor trajectory of a linear controllable system S_1 and let $x_1(t)$ be constant over the interval $[0,\infty)$. Let S_2 be a second linear controllable system with an attractor trajectory $x_2(t)$ which differs from $x_1(t)$; in particular, let $x_1(0) \neq x_2(0)$. Now let a perturbation be applied to S_2 prior to $t = 0$ to bring the system to $x_1(0)$ at $t = 0$. For $t > 0$, let $F(t) = 0$. Since S_2 is controllable,[7] we can modulate the sources in such a way as

---

[7] "A system is said to be controllable if and only if it is possible, by means of the input, to transfer the system from any initial state $x(t) = x$, to any other state $x_T = x(t)$ in a *finite* time $T - t \geq 0$." [p 204, Friedland 1986].. Friedland adds that if impulsive inputs are permitted, "it is possible in a controllable system to go from any state to any other state in zero time."

to bring the state of the system to any other state in finite time, and since the system is linear, the actuator efforts are unlimited, so we can accomplish this transition in arbitrarily short times. We can therefore construct the set of source histories which will keep the interaction port of S_2 tracking along $x_1(t), \forall t > 0$; i.e., from the point of view of the observer, the port positions of S_1 and S_2 will remain fixed and equal in the interval $[0,\infty)$ and the interface forces will be zero.

If a system is of at least first order, then the interaction port position cannot instantaneously change. Since the attractor trajectory $x_2(t)$ is a path which the interaction port of S_2 could have followed, $x_2(t)$ cannot change instantaneously either, so having $x_1(0) \neq x_2(0)$ implies that there is a finite interval, $[0,T_m]$, during which $x_1(t) \neq x_2(t)$. If $T < T_m$, then $x(T) = x_1(T)$, but since the system being observed could equally well be S_2 as S_1 (both yield the same observations), the attractor trajectory at time T could be $x_2(t)$, which, being different from $x(T)$, proves the theorem.

Corollary 1. The potential for confusion over the location of the attractor trajectory at time T is independent of knowledge of the system's impedance structure.

Proof: Knowing the impedance structure of a system does not allow one to unambiguously locate the attractor trajectory after a finite observation time. The proof of Theorem 1 makes no assumption about impedance structure. In particular, if S_1 and S_2 share the same (known) linear impedance, control inputs for S_2 are easily constructed to guarantee that $x_1(0)$ differs from $x_2(0)$. The rest of the proof follows as before.

One might accept Theorem 1 and yet argue that we have unnecessarily restricted ourselves to zero interaction forces; by choosing an appropriate (more complex) set of port interactions, we could unambiguously tease out the attractor trajectory of a system. To counter this argument, we propose the following theorem.

Theorem 2. Given a system S with multiple interaction ports, no set of port manipulations can, in general, unambiguously identify any point on an attractor trajectory if observations begin at a time at which the internal equilibrium of S cannot be verified.

Proof: Consider a pair of identical linearly controllable systems with different attractor trajectories as described in the proof of Theorem 1. Since both S_1 and S_2 by construction share the same impedance structure and have the same interaction port path for $t > 0$, they will exhibit the same interaction port behaviour for $t > 0$. However, since S_1 and S_2 have completely different attractor trajectories in the interval $[0,T]$, we have a 1-to-2 mapping from port behaviour to attractor trajectory. No computations based solely on port behaviour in the interval $[0,T]$, therefore, can uniquely identify a point on an attractor trajectory.

The Fundamental Unmeasurability of the Attractor Trajectory

The results of Theorem 2 are somewhat disturbing. We had set out to make precise the notion of "the point to which the system is tending" and we now find that our definition requires that we keep track of the system from a state of equilibrium. How do we determine that our system begins from equilibrium? The following theorem asserts that such a determination is impossible to make.

Theorem 3. Measurements over a finite time interval at the interaction ports of a system provide insufficient information to conclude that the system is at equilibrium at any time within the interval of observation.

Proof: Suppose we begin observing a system, S, at $t = 0$ with the aim of establishing the existence of internal equilibrium of S at some future time, t_i. According to Definition 3, if we can verify that S is in equilibrium at time t_i, any trajectory starting at that point will be in the set G. If, further, we verify that a trajectory including $x(t_i)$ is a uniformly stable trajectory, then $x(t_i)$ is on an attractor trajectory. Theorem 2, however, states that we are unable to unambiguously identify a point on an attractor trajectory if we do not know in advance that the system started from equilibrium. Since that is the case here, and since an equilibrium point is by definition on an attractor trajectory, we cannot identify equilibrium based on our port measurements.

Corollary 3. An attractor trajectory is fundamentally unmeasurable without making assumptions concerning the structure of the system.

If we cannot measure an attractor trajectory without knowing that the system started at equilibrium, and we cannot know that the system is at equilibrium based on measurements of the system, then any estimate of an attractor trajectory is fundamentally based on an assumption as to when the system is in equilibrium, that is, on an assumption about the internal structure of the system. For example, if the internal system dynamics can be characterized by some time constant, we could reasonably claim that the system is in equilibrium if the interaction ports have not moved during an interval several times longer than the characteristic time constant. This is a relatively benign assumption, but we should nonetheless recognize the fact.

We have demonstrated above that the attractor trajectory is fundamentally unmeasurable unless we have some knowledge about the internal structure of the system. In fact, the problem is worse than this. The following theorem asserts that we cannot determine whether or not a given system even has stability properties without making further assumptions. This is so because we cannot tell from interaction port behaviour alone whether the behaviour we observe is due to command changes or to impedance-generated responses.

Theorem 4. Given a system S, no set of interaction port manipulations can determine whether or not the system has stability properties about any trajectory whatsoever.

Proof: The definition of stability about a trajectory (Definition 4 above) requires that the difference between an executed path and the candidate trajectory be bounded by R for any perturbation smaller than r at time t_0. Any conceivable test of stability therefore requires a measurement of this difference. To compute it, however, both the executed path (which results from a perturbation applied to the interaction port) and the unperturbed path must be measured simultaneously, which is impossible because it would require the same interaction port to be at different locations at the same time. No interaction port measurement, therefore, can provide any information relevant to the question of stability.

(Saved By) The Ergodic Assumption

The above discussion about the unmeasurability of attractor trajectories demonstrates that, without making assumptions about the structure of the system under consideration, we are unable to directly or indirectly measure the command input to the system. The typical system identification problem sidesteps this difficulty by considering only those systems for which the inputs (and often the structure) are known - only the parameters remain to be identified. If we can make the assumption, however, that the system is at equilbrium at time t_0, and if we do not apply any forces after that time, then the trajectory which evolves belongs to the set G described in Definition 3. If we can further assume that the system can be brought back to the same equilibrium point at some later time t_1, that the system commands are the same in the interval $[t_0,t_0+T]$ as in $[t_1,t_1+T]$, and that any time-varying properties of the system satisfy $p(t_0+t) = p(t_1+t),\ 0 \leq t \leq T$, then the trajectories which evolve in the two intervals specified will be identical. We could therefore assess the stability of a specified trajectory by comparing an unperturbed trajectory in the first interval, $g_0(t)$, with a perturbed trajectory in the second interval, $g_1(t)$. This set of assumptions is known collectively as the ergodic assumption, by which we mean that observations made on a single system during different time intervals can be compared in the same way in which we would compare observations of several identical systems made during a single time interval.

The Meaning of the Attractor Trajectory

If we feel justified in making the assumptions described in the above sections given a particular system of interest, then we are able to measure attractor trajectories in various circumstances. Since attractor trajectories are fundamentally expressions of interaction port behaviour, what can we say about their relationship to the internal features of the system being analyzed?

Recall that for linear systems, Norton's theorem says that the behaviour at any given interaction port can be reproduced by an equivalent system consisting of a time-varying flow source placed in parallel with a linear impedance element.[8] If the effort at the interaction port is kept at zero, then the path of the interaction port, which is by definition the attractor trajectory, matches that of the Norton-equivalent flow source. For the linear system, then, the attractor trajectory which we measure is precisely the Norton-equivalent flow source. Since the Norton-equivalent circuit is uniquely defined for linear systems, the attractor trajectory is a complete and unambiguous embodiment of the notion of "the point which the system is heading towards"; it is, in essence, the effective command input to the system.

---

[8] Possibly complex, but with no internal effort or flow sources.

Given that the human arm is not a linear system, what can we say about the attractor trajectories measured for the arm? I will show later that certain models for the arm do have generalized Norton equivalent representations and that for these models, the attractor trajectories have the same interpretation as Norton-equivalent sources. However, since any particular Norton-equivalent model can correspond to an infinite number of real systems and configurations, we cannot say in general that the attractor trajectory can be identified with a single source term in the real system. If, however, we have a good model of the real system and the real system satisfies the conditions for having a unique Norton equivalent, we will then be able to establish a unique map from the attractor trajectory to the system's command input. I will return to this issue shortly when we consider the relationship between attractor trajectories and virtual trajectories; in the meantime, I will provide some examples of attractor trajectories to make the definitions more concrete and to respond to some possible objections to them.

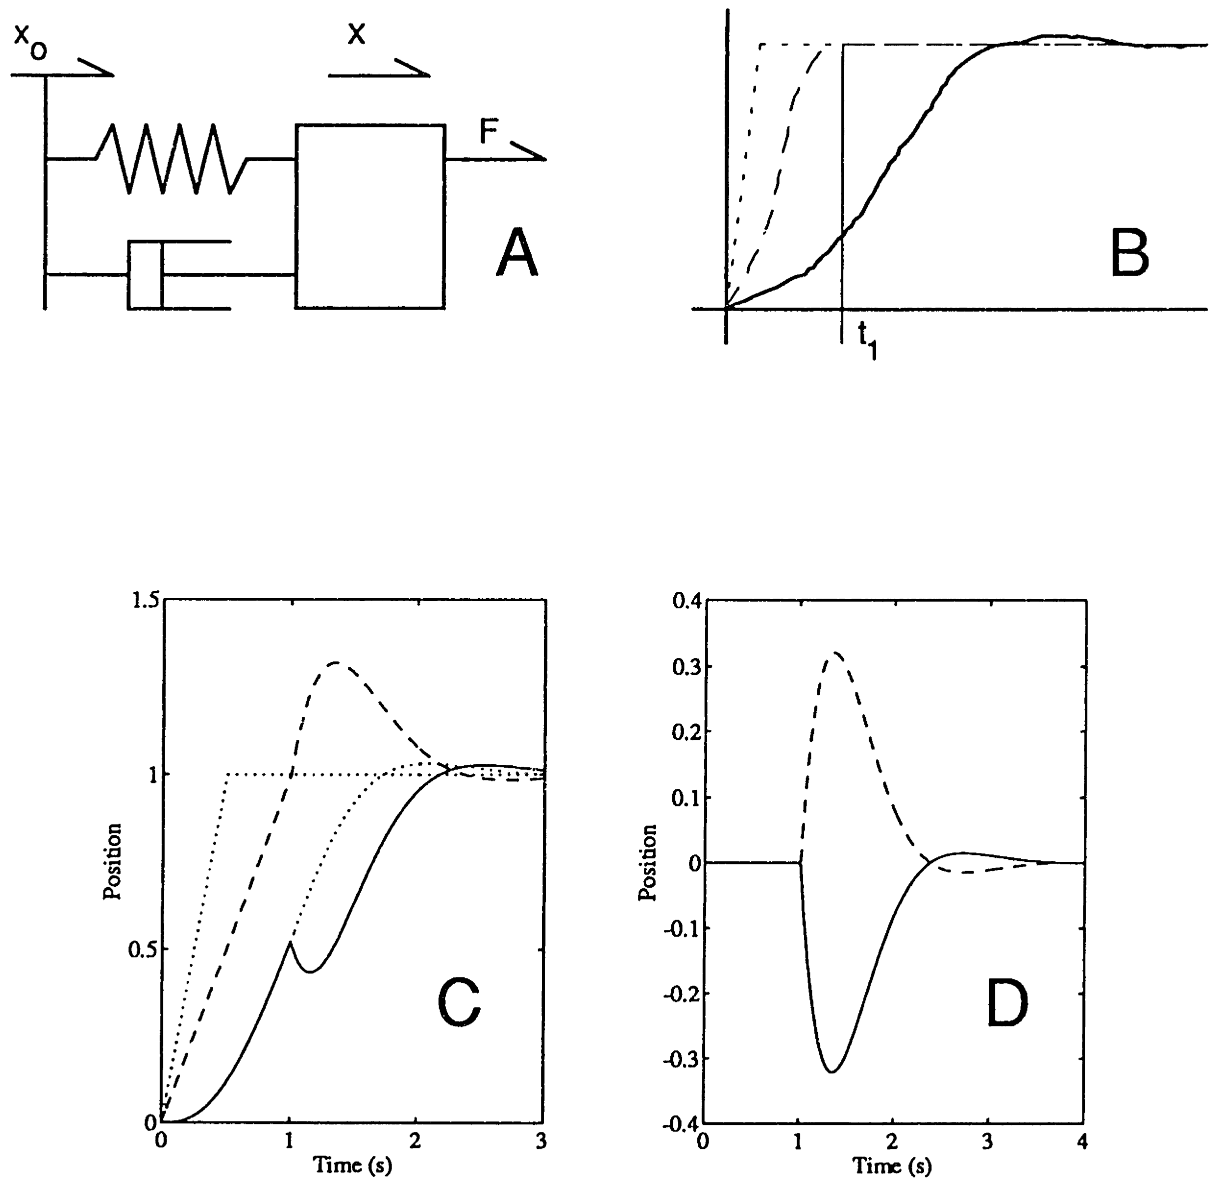

Figure 1. (A) Mass-Spring-Damper system. (B) Trajectories for the system shown in (A). The dotted line represents the transition of the commanded position, the solid line represents the unperturbed path of the mass in response to the transition of the commanded position, and the dashed line represents a perturbed path for which the external force applied to the mass is zero after time t_1. *(C) Perturbation responses to impulsive disturbances of opposite signs but equal magnitude applied at time* $t_1 = 1$ *to the system when it is in the state shown by the dashed and solid lines in (B). The curved dotted line shows the continuation of the unperturbed path of the mass. (D) The impulse responses of the two disturbances shown in (C), relative to the dotted trajectories; note that they are identical except that they are of opposite sign.*

Examples of Attractor Trajectories (Some Counterintuitive Situations)

Mass-Spring-Damper With Rapid Transition of Equilibrium Point

The definitions developed above do imply some results which appear to be counterintuitive at first glance. For example, suppose our system is composed of a mass driven by a spring and damper combined in parallel (see Figure 1A). If the rise time of the reference position (the free end of the spring and damper) undergoes a rapid step transition (Figure 1B), the mass will follow slowly afterwards. However, if the mass is rapidly moved to the final reference position, it will show no tendency to move anywhere else and if it is perturbed further, it will settle again to this new constant reference position. Surely these are grounds for calling the new reference position the trajectory to which the system is attracted. How can one seriously claim that the attractor trajectory is slowly rising towards the new reference position when the mass shows no tendency to return to the so-called attractor trajectory but instead stays put? To see that the claim is plausible, one simply has to change one's point of view. Consider the responses of the system to impulsive disturbances of equal magnitude but opposite sign applied at time t_1 when the system is on the perturbed and unperturbed trajectories, respectively (see Figure 1C). The differences between these responses and the trajectories the system would have followed had the impulses not been applied are shown in Figure 1D. Since both these responses are identical (except for sign), we must conclude that both candidate attractor trajectories exhibit the same dynamic response to a perturbation and it would be mere prejudice to decide that a constant trajectory was the "real" attractor rather than varying

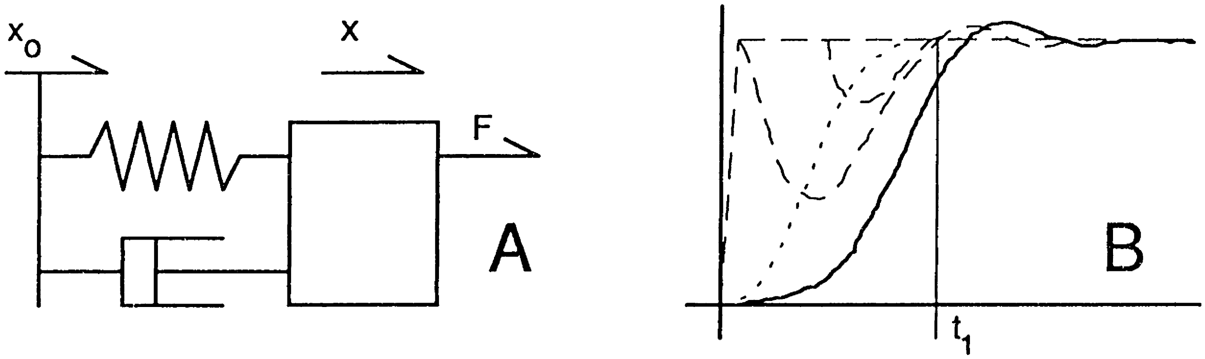

Figure 2. (A) Mass-Spring-Damper system. (B) Trajectories for the system shown in (A). The dotted line represents the transition of the commanded position, the solid line represents the unperturbed path of the mass in response to the transition of the commanded position, and the dashed lines represents perturbed paths for which the external force applied to the mass is zero after time t_1.

one. This paradox is resolved when we consider the implications of Theorem 2: if we did not know the original equilibrium state of the system, we would naturally identify the unforced perturbation response as the attractor trajectory simply because we did not start observing early enough to detect the perturbation. Once we know that the zero-velocity trajectory is the result of an earlier perturbation, then it is apparent that the slowly rising path of the unperturbed mass is the true attractor trajectory, while the seemingly steady (when viewed in absolute coordinates) perturbation response is an illusion, compelling only because of its fortuitous constancy.

Mass-Spring-Damper With Slow Transition of Equilibrium Point

A similar objection to the above definition of an attractor trajectory comes from a perturbation experiment involving deafferented monkeys. In this experiment, the system model is as shown in Figure 2A; the situation is similar to that shown in Figure 1, except that the rise time of the commanded position is extended (see the dotted line). The observation is that if the arm is quickly perturbed to the final reference position and released, then it returns to some intermediate point before rising again to the final position (see the dashed lines). If the arm is held a little longer at the reference position before being released (until time t_1), no return movement is observed; the arm stays in place. Surely this situation is stronger evidence that the true attractor point arrives at the final reference position prior to the mass, whose path we claim to be the attractor trajectory. In fact, if one again compared perturbation responses expressed both in absolute coordinates and in coordinates relative to the trajectory claimed to be the attractor, one would find that the responses would be the same; the system converges equally to both the (assumed) known command trajectory and to the claimed attractor trajectory so there is nothing in the perturbation response to decide between the two. Much of the confusion in this case arises because the selected perturbations are special; they are not impulses, but consist of whatever forces are required to place the arm in a particular state immediately prior to release. Again, we are tempted to attribute significance to the existence of a steady response after release of a perturbation, but such an attribution implicitly assumes that the steady response could only arise from a steady command. I could, however, compute an infinite number of different command histories which would give precisely the same response to that perturbation; for example, any command trajectory which is the sum of the steady response and a solution to the equation $Bv + Kx = 0$ would cause the mass to remain at a constant position upon release of the perturbation.

Relation of Attractor Trajectory to Virtual Trajectory

Identification of Virtual Trajectory With Command

Having discussed the question of what we are in fact measuring when we make measurements at an interaction port, we turn now to the question of the signficance of those observations for the various theories of motor control. In particular, we are interested in assessing the plausibility of a class of hypotheses known as the *virtual trajectory hypotheses*. Historically, the virtual position has been defined as that position which the limb would take up in the absence of external loads. However, definitions of the virtual trajectory have also generally required, at least implicitly, if not explicitly, that inertial loads be treated as external loads. This is clear in the models of Flash, for example, which have linear stiffness and damping elements attached to an inertial linkage [Flash 87]; the springs define an equilibrium position (in the absence of external loads) which can be thought of as the rest position of an equivalent set of springs attached to the hand. Flash identifies the movement of the rest position of the equivalent set of springs with the virtual trajectory, explicitly excluding the inertia of the linkage. She then presents simulations to show how the virtual trajectory differs from the actual hand path (or attractor trajectory for the interaction port at the hand) in a variety of unconstrained movements. Flash in fact goes one step further and treats the damping behaviour of the arm as an external load. In modeling the arm in this manner, she identifies the virtual trajectory with the command input to an assumed neuromuscular apparatus which behaves as a pure spring and argues that with an appropriate choice of impedance values the virtual trajectory exhibits more regular and repeatable behaviour than does the actual hand path.

This then is the central claim of one family of virtual trajectory hypotheses: that the command to the limb is mechanically equivalent to a specification of a reference trajectory with attractor properties and that the time course of this reference trajectory is more closely related to the description of the task than is the actual hand trajectory. For example, we know that it is in principle possible to specify a neural input to a limb which will cause it to track along a straight line between two points. The fact that human subjects deviate from such a path would tend to support the virtual trajectory concept if we could show that some command representation deviated less so. The major difficulty with performing such a test is that there is no single point in the system where a motor command can be unambiguously identified; rather, neural activity generates activity in a large number of motor units whose combined responses implement behaviour which can be partially summarized by locating an attractor point. It is only in this sense, that neural activity has a mechanical manifestation, that we can say (without resorting to direct interpretation of neural activity — a much more difficult undertaking) that the command representation has particular geometric characteristics. In interpreting neural activity in geometric terms, then, we have left physiological reality and are speaking of a model for which the input is a trajectory rather than some symbolic representation; it is this *equivalent* trajectory that we are referring to when we speak of a virtual trajectory.

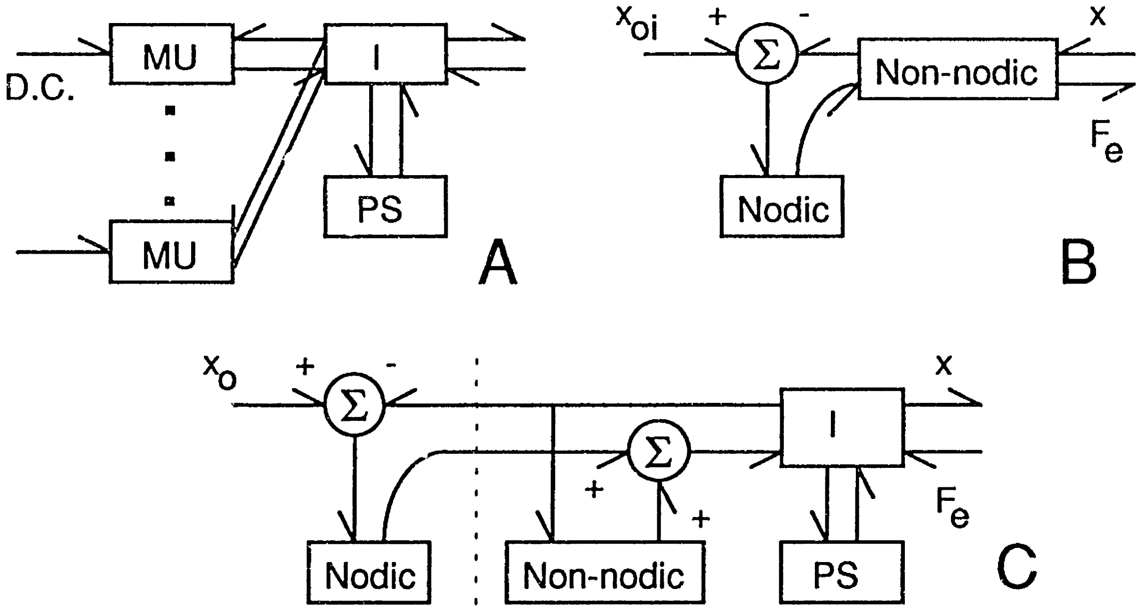

Figure 3. (A) Schematic diagram of subsystems of limb. The inertial linkage, I, is acted upon by forces arising from passive structures, PS, such as joint friction, by forces from the many motor units, MU, which are activated by descending commands, DC, and by external forces applied at the interaction port. Most connections are shown as double bonds to emphasize that they are power-exchanging interaction ports; position and force at each port are not independent. Arrows show input/output relationships. (B) Schematic view of a single motor unit showing that the forces it generates can in general be partitioned into those elements whose behaviour is a function of the difference between the interaction port position and the commanded position (nodic) and those whose behaviour depends only on the interaction port position (non-nodic). (C) Alternative model of entire limb shown with all non-nodic contributions from each individual motor unit lumped together and with the equivalent effect of the descending commands expressed as a virtual trajectory which is the input to an equivalent nodic impedance. The dotted line separates the Norton-equivalent neuromuscular behaviour from that which can be directly modelled.

If we cannot test the virtual trajectory hypotheses by direct interpretation of the symbolic content of neural activity, we must base our test of the hypotheses on mechanically measurable quantities, i.e., an attractor trajectory. Our problem is then reduced to deciding which attractor trajectory to identify as the virtual trajectory. This is a difficult problem because the virtual trajectory as such does not exist — it is simply a partial description of the behaviour of a collection of motor units. If we have, for example, a detailed model of the multitudinous components of a real limb as shown in Figure 3A, what, in the context of the model, is the virtual trajectory? Each motor unit can be described using a more detailed model (Figure 1B) which partitions those force-generating components which depend on the absolute interaction port position (non-nodic[9]) from those which depend on the difference between the interaction port position and some kinematic equivalent of the descending command to that motor unit (nodic). Because of the parallel connection of the motor units to the inertial linkage, the non-nodic components from all motor units can always be gathered together into a single equivalent non-nodic impedance and the nodic components can often, but not always,[10] be consolidated into a single equivalent nodic impedance with a well-defined command trajectory, x_o, the virtual trajectory for the limb (Figure 3C). The attractor trajectory that we identify as the virtual trajectory is the one that remains when all non-nodic phenomena in the limb have been accounted for (measured at the interaction port marked by the dotted line in Figure 3C). In principle, then, we do have a means of directly estimating the virtual trajectory, although in practice many detailed questions about the behaviour of

---

[9] The term *nodic* refers to an operator which is a function only of the differences between its inputs rather than of their absolute values. For example, in heat transfer, the conductive resistance is nodic, $Q = R(T_2 - T_1)$, whereas the radiative resistance is strongly non-nodic, $Q = k_r(T_2^4 - T_1^4)$. In this discussion, we use the term non-nodic somewhat loosely to refer to those impedance contributions which are not referenced to the difference between the endpoint position and the commanded position. Strictly speaking, some of these impedances could be nodic; we will consider them to be non-nodic, however, if they are referenced to something other than the commanded position (as ground-referenced damping would be, for example).

[10] Consider the overlap of two V-shaped impedance functions (energy as a function of position) with different vertex positions but equal slope angles. For each impedance function considered independently, the minimum point is well-defined — it is the vertex — but in the region between the vertices, the energy function is a flat line with no unique minimum at all. If the impedance functions were quadratic rather than linear, however, their summation would yield a unique minimum, and all evidence suggests that the aggregate behaviour of the motor units in an animal limb exhibits unique minima over a wide range of activation patterns [e.g., Giszter 93].

muscle would have to be answered before we could come up with a believable virtual trajectory estimate.

Relationship Between Commanded and Attractor Trajectories

Although our current state of knowledge of neuromuscular dynamics is too limited to enable us to reliably compute the virtual trajectory given a measurement of an attractor trajectory, we might nonetheless wish to speculate as to the possible forms of the virtual trajectory. Since such speculation would require the use of models of the human neuromuscular system, the significance of the results would depend on the plausibility of the models. As already mentioned, Flash used a model in which all non-inertial behaviour was approximated by a linear model and in which the damping was referenced to ground (zero velocity) [Flash 87]. This choice of a reference for the damping is somewhat arbitrary; there is evidence that damping is activation dependent [Feldman 86, Wu 1990], so it is at least possible that the damping reference velocity could be specified by the CNS in the same manner as the position reference. McIntyre points out that damping referenced to the movement of the virtual position would enhance the command-following ability of the arm in single-joint movements, and a similar argument could be made for the two-joint movements we are considering [McIntyre 90, 92]. In effect, McIntyre proposes that the neuromuscular system is structured as a proportional/derivative controller rather than as a pure proportional controller, although he gives no specific physiological argument to support his contention. Since there is as yet no strong argument favouring

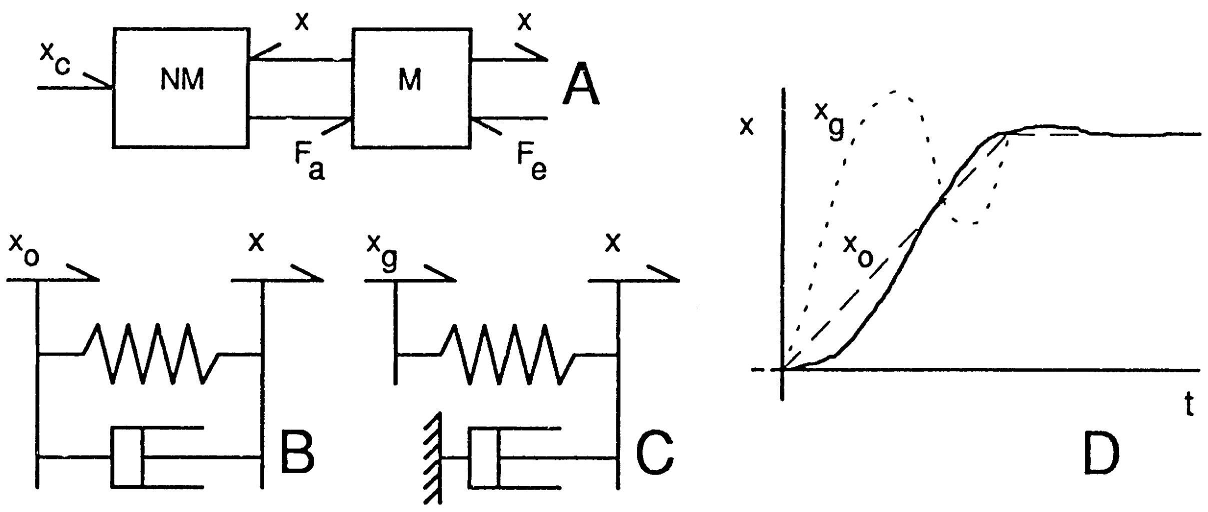

Figure 4. (A) Model of the elbow joint consisting of the mass of the forearm, M, driven by the neuronuscular system, NM, which takes a virtual trajectory as its input. (B) and (C) Spring and damper models of the neuromuscular system: (B) has the damper referenced to the virtual trajectory, whereas (C) has it referenced to ground. (D) Trajectories generated by the models: the solid line represents the attractor trajectory of the entire system, the dashed line represents the attractor trajectory of the subsystem left after accounting for the inertial effects (which is identical to the command input for system B), and the dotted line represents the command input for system C.

either one of these models, we could take them, for speculation purposes, as extremes and ask what effect our choice of model would have on our estimate of the virtual trajectory.

Suppose we observe the one-dimensional system shown in Figure 4A as it executes the trajectory shown as a solid line in Figure 4D when no forces are applied at the rightmost interaction port. After accounting for the inertial effects, we determine that the attractor trajectory of the remaining subsystem is as shown by the dashed line in Figure 4D. If we have no knowledge of the structure of the actual neuromuscular system, we cannot say what the correspondence is between the dashed line and the command input, x_c. Two possible models of the neuromuscular system are shown in Figure 4B and Figure 4C. If the system were actually as shown in B, then the attractor trajectory we measured would in fact be the command, x_o. If, however, the system were actually as shown in C, but we

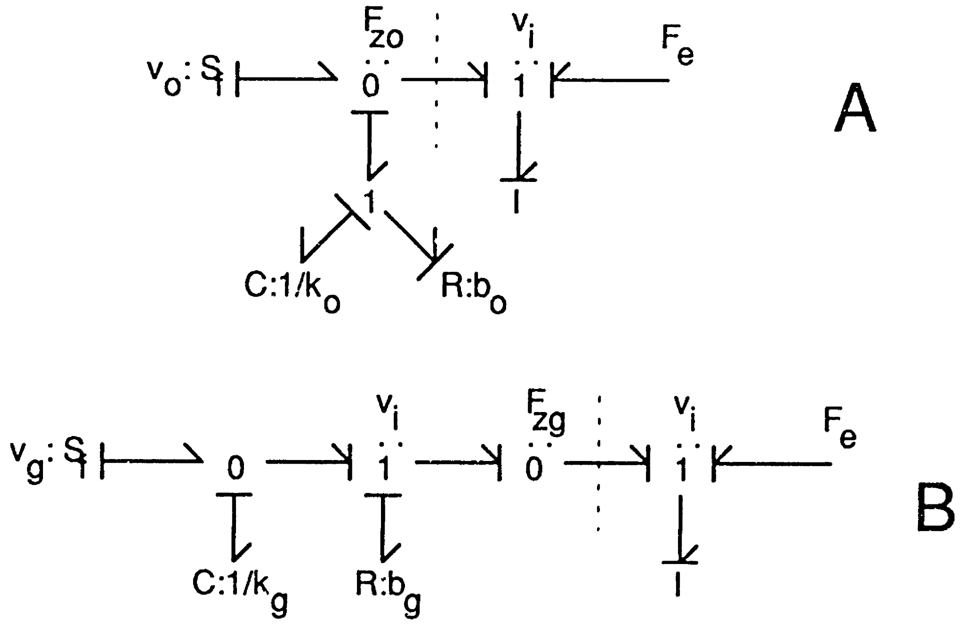

Figure 5. Bond graphs for the two primary neuromuscular models. (A) Model with damping referenced directly to the virtual trajectory. (B) Model with damping referenced to ground. The dotted line indicates the interaction port for the non-inertial components of the limb. In both graphs, the bond on the right-hand edge of the graph is connected to the environment, which is assumed to impose an effort on the neuromuscular system.

had assumed that it was like B, then we would erroneously conclude that the virtual trajectory was the same as the attractor trajectory. Likewise, if we had assumed the system were like C, but it was actually as shown in B, then we would misinterpret x_g as the virtual trajectory.

In the previous section, we noted that one can often, but not always, transform a collection of subsystems, each of which exerts an effort on the same kinematic linkage, into a single equivalent system. Given that we have been discussing two extreme lumped-parameter models for the equivalent system, we might be interested in knowing what conditions must be satisfied to ensure that the two models can be transformed from one to the other. Consider the bond graphs (Figure 5) for the systems which were shown schematically in Figure 4B and Figure 4C. The first bond graph represents the system in which damping is referenced directly to the virtual trajectory, whereas the second shows the damping resistance, R, to be a function solely of the interaction port position and velocity. The VT-referenced damping model has the interpretation, mentioned earlier in the section on the meaning of the attractor trajectory, of being the Norton equivalent representation of the system as viewed from the interaction port indicated by the dotted line in Figure 5A. If the actual system is linear and has the form of Figure 5B, its impedance function will be

$$F_{zg} = k_g x_g - k_g x_i - b_g \dot{x}_i$$

and its Norton representation will be

$$F_{zo} = k_o x_o - k_o x_i + b_o \dot{x}_o - b_o \dot{x}_i.$$

We can determine the trajectory, $x_o(t)$, which renders $F_{zo}(t)$ uniformly zero by observing that $F_{zo}(t)$ and $F_{zg}(t)$ are equivalent, setting $F_{zg}(t) = 0$, and solving for $x_o(t)$ in terms of $x_g(t)$. When we do this, we find that the following differential relationship holds between $x_o(t)$ and $x_g(t)$:

$$\dot{x}_o = b_g^{-1}\left(-k_g x_o + k_g x_g\right).$$

We also note that the two systems can be mapped into one another if we ensure that the stiffness and damping elements have the same values in the two models; in other words, the two systems are formally equivalent in terms of their interaction port behaviour if their source terms are related according to the above equation and if $k_0 = k_g$ and $b_o = b_g$.

If the actual system is not linear, but the impedance elements are nonetheless separable, in the sense that the position-dependent relationship can be separated from the velocity-dependent relationship as shown in Figure 5A and B, then we can write an implicit relation between $x_o(t)$ and $x_g(t)$ by substituting $x_o(t)$ for $x_i(t)$ in the equation for $F_{zg}(t)$ and then setting this equation to zero:

$$\dot{x}_o = b_g^{-1}\left(k_g(x_g - x_o)\right).$$

where b_g and k_g now represent nonlinear functions. If b_g is invertible, then this implicit relationship can be solved, numerically if not analytically. I have not determined whether or not separability of the impedance elements as shown in Figure 5B implies that the Norton equivalent impedance can also be expressed as a separable function, although, as demonstrated above, it is simple to show that the linearized impedance relation is a separable function.

Finally, if the impedance element k_g is not a function merely of relative position, but of velocity also (let us rename this element Z_g to indicate that it represents a generic impedance), and if this function is neither linear nor separable, then we can form an implicit equation relating $x_o(t)$ to $x_g(t)$ as follows:

$$Z_g(\dot{x}_g - \dot{x}_o, x_g - x_o) = 0,$$

but we have in general no guarantee that this relation has a unique solution for $x_o(t)$. Furthermore, it is also impossible to say in general how to derive the equivalent impedance relation, Z_o (the unified replacement element for the C and R elements of Figure 5A) given Z_g. However, if we are able to determine that the system has an attractor trajectory at the pre-inertial interaction port (the dotted lines in Figure 5), then the system either has a structure similar to that of its Norton equivalent or it has a structure which can be mapped into a Norton equivalent; if such a mapping did not exist, we would not be able to find an attractor trajectory.

It is clear that our choice of model strongly affects our interpretation of the attractor trajectories we measure. If the actual system is as shown in Figure 4C, we will not only have to determine the attractor trajectory for the inertialess subsystem — we will also have to invert the ground-referenced damping behaviour and the nodic impedance to deduce the virtual trajectory. Since this process would require a detailed knowledge of the system and would therefore introduce significant measurement errors, we need to have some means of testing the virtual trajectory hypotheses which does not involve such complex (and possibly suspect) modelling assumptions.

Testing Hypotheses Without Measuring the Virtual Trajectory

Since the definition of an attractor trajectory is expressed solely in terms of interaction port behaviour, we can, after determining that a system starts at equilibrium, perform experiments to measure its attractor trajectory. In particular, it is straightforward to estimate the attractor trajectory at the hand; all one needs to do is to ensure that no forces are applied to the hand - the resulting motion is the attractor trajectory. As we try to push deeper into the system, however, the problem becomes more difficult. Our first claim is that subsystem which contains the neuromuscular implementation of the flow source shown in Figure 3C is completely contained within the muscles and tendons; the skeleton is external to this subsystem. If we know the inertia of the system and can accurately measure the velocities and accelerations at the endpoint, we can subtract from the interaction forces measured at the hand the forces required to drive the inertia of the subject's skeleton (as well as any other inertias which are causally dependent on the motion of the skeleton) along the measured path; what remains is the interaction force at the port of the subsystem containing the flow source. If we can find the trajectory for which this interaction force remains zero throughout the movement, we will have found the attractor trajectory for this subsystem.

Limits to Modelling

As we pointed out in the previous section, this subsystem may still not be well-modeled by a single flow source connected through a nodic impedance. For example, there is some rotational friction in all joints. Such friction could be modeled by a resistive element attached to the interaction port at the hand. If we either had a model of this friction or could assume it to be small, we could subtract off the frictional contribution to the measured force and use the remaining force as our estimate of the force at the next interaction port. We could repeat this process, pushing back from one interaction port to the next towards the command input, as long as we had a good model for the dynamics at each step along the way. As soon as we were faced with significant uncertainty about the appropriateness of our model of the real system, however, we would have to stop our attempt to link an attractor trajectory to the virtual trajectory. In our case, that means we are stymied as soon as we have accounted for inertial effects (and perhaps joint friction) simply because we do not have a good enough model of the activation-dependent impedance behaviour of the limb *in toto*.

Hypothesis Definition and Model-Independent Tests

Fortunately, we do not need to have a complete model of the neuromuscular system in order to test some of the virtual trajectory hypotheses; an incomplete model will still allow significant tests to be performed. These tests depend on the fact that if a subsystem with attractor properties is connected through a chain of intervening non-nodic dynamic systems to an interaction port whose movement it is being used to control, the attractor trajectories at each step along the chain will become progressively more distorted when compared with the commanded trajectory. For example, suppose we consider the system described by the sketch in Figure 6A and by the bond graph in Figure 6B, and suppose the command trajectory is a ramp going from 0 to 1 unit in 1 s. If the system is linear with a natural frequency of 0.5 Hz and a damping ratio of 0.7, then the attractor trajectories are as shown in Figure 6C. The actual command trajectory is shown by the dotted line and the path taken by the mass is shown by the solid line. As we expect, the path taken by the mass differs significantly from the command trajectory. More relevant is the relation

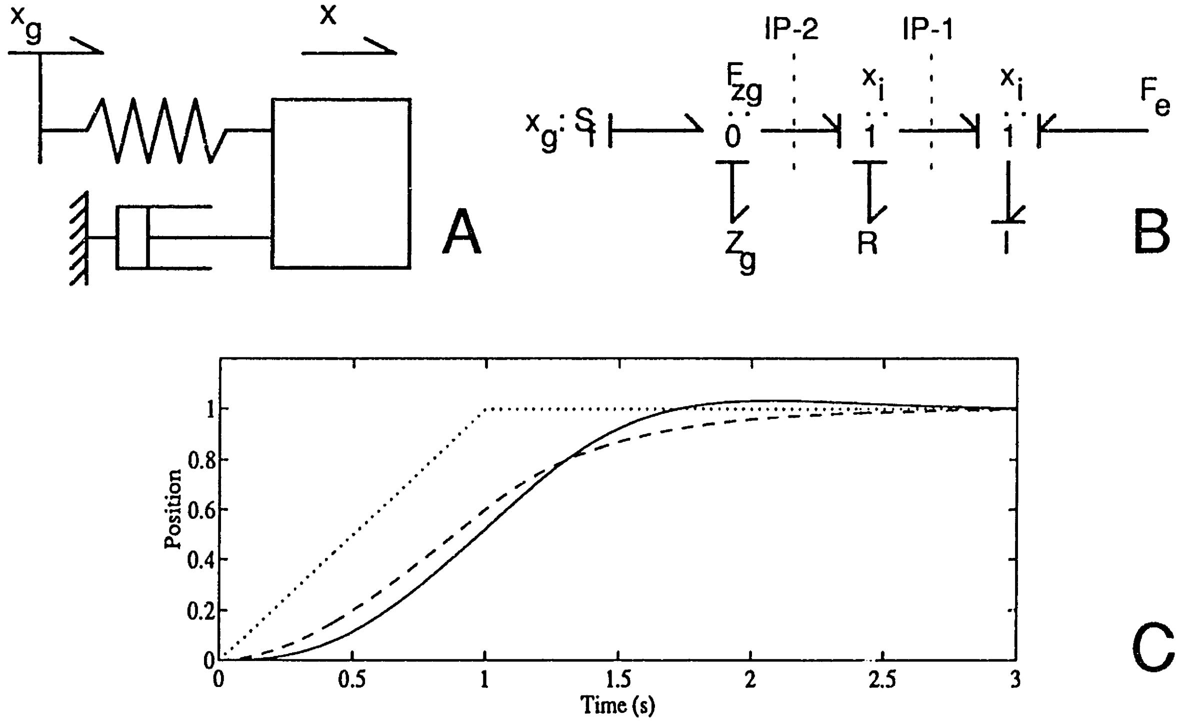

Figure 6. (A) Model of a single-joint system in which the damping is referenced to ground. (B) Bond graph of the model shown in (A). The dotted lines represent interaction ports. (C) Attractor trajectories for the system shown in (B). The solid line is the attractor trajectory at the port where external forces are applied, the dashed line at IP-1, and the dotted line at IP-2.

between the command trajectory and the attractor trajectory at IP-1; the discrepancy between these trajectories is almost as great as that between the command trajectory and the path of the mass, but there is a relatively small discrepancy between the attractor trajectory at IP-1 and the path of the mass. The subsystem to the left of IP-1 therefore exhibits better command-tracking performance than does the one to the left of IP-2.

One might not be inclined to accept without hesitation the above statement that there is progressive distortion along the chain of dynamic subsystems if one recalls that the whole thrust of control theory is to design compensators to enable better command tracking. Surely it is possible to add some compensator outside IP-2 to improve performance. In fact, except in a very limited sense, this is not possible. The techniques of controller design generally involve shaping the response to an error signal generated by the difference between the commanded position and the measured position; any such compensator must therefore be part of the impedance labelled Z_g in Figure 6B. Any compensator added outside of IP-2 will only have access to the measured position and therefore cannot exploit any knowledge of the desired or commanded path. If such a compensator does succeed in enhancing tracking performance, it will only do so somewhat fortuitously and under special conditions. For example, if the commanded trajectory is near resonance for a lightly damped system, large deviations between the commanded and executed trajectories will occur. Adding damping will reduce the size of these resonant oscillations, but will degrade tracking performance when the frequency content of the commanded trajectory is below resonance.

Having established that we expect progressively more distortion of the virtual trajectory as we work our way outward towards the external interaction port (in our case, the hand), we turn now to a discussion of how we might test the virtual trajectory hypotheses without needing to measure the virtual trajectory itself. We consider two major classes of virtual trajectory hypotheses: the *strong correspondence hypotheses* and the *insensitive strategy hypotheses*.

The central claim of the strong correspondence hypotheses is that the virtual trajectory is more closely related to the task description than is the actual trajectory. If we have a measure of closeness to the task description, then we would expect the hand trajectory to be furthest away, and each succeeding attractor trajectory (as we account for more of the limb dynamics) to be closer because we are accounting for more and more of the influences distorting the commanded trajectory. Thus, if we see increasing distance from the task description as we compute deeper and deeper attractor trajectories, we will be able to reject these hypotheses even though we may never reach the true virtual trajectory.

The insensitive strategy hypotheses are less rigid than the strong corresondence hypotheses; they relax the requirement that there be a close connection between the task description and the virtual trajectory. This allows the subject to modify their virtual trajectory to best accomplish a task, particularly in novel or difficult dynamic environments. However, they assert that for a given class of task, the strategy the subject uses should be highly consistent, with most adjustments taking the form of distance or time scaling as necessary to properly execute the task. Large changes in the dynamic environment would be required to cause the subject to adopt a noticeable change in strategy. To test this class of hypotheses, we would choose a measure of self-consistency and argue again that the actual hand paths should exhibit the greatest degree of inconsistency and that as we progress through the various attractor trajectories towards the virtual trajectory, we would expect to see increasing consistency. In particular, if we compare the consistency of the unperturbed hand paths for a given task executed at different speeds in different workspace locations to the consistency of the pre-inertial attractor trajectories for the same task, we would expect the attractor trajectories to show greater consistency.

The Danger of Using Unsubstantiated Models

The tests described above may lead to disproofs of the virtual trajectory hypotheses as stated. If they do not, however, we cannot conclude that the hypotheses are proven; the most we can claim is that they are plausible. We might then decide to perform further tests by accounting for damping behaviour. If we are incorrect in assuming that the damping elements are referenced to ground, however, we might well be able to compute an attractor trajectory, but this computation would be not only uninformative but potentially misleading. Since the computed attractor trajectory would be a transformation of the command to a system which actually exhibits better command tracking performance than our model of the system, the transformed command would have to be more complex than the actual command. If we observed this increase in complexity, we would erroneously reject the hypotheses.

This then is the danger of applying a model for which we have insufficient evidence - that we may observe increasingly complex behaviour and incorrectly conclude that the hypotheses are to be rejected. This is why it is prudent to first assess the plausibility of the hypotheses using the pre-inertial attractor trajectories before engaging in a more detailed modelling effort. It is possible, however, to use unproven models to suggest areas for future research. For example, suppose that the virtual trajectory hypotheses were in fact correct. We will find the most consistent or simple trajectories when our model of the system matches the structure of the real limb; if our model is incorrect, we will observe more complex attractor trajectories, which we could take as a hint that our model was wrong and as a suggestion that we pin down the origin of the phenomenon most recently included in our model. Such a use of unproven models may be helpful even if it does not lead to rigorously supportable conclusions.

Summary

In this chapter, we have considered how mechanical interrogation of a subject can be used to test the various virtual trajectory hypotheses. These hypotheses all make claims about the nature of the motor computations occurring in the central nervous system and of the symbolic content of the instructions transmitted down the spinal column. Although we cannot make direct observations of these neural instructions by mechanical means, we can measure the mechanical consequences of neural activity, which manifests itself in movements of the limb which are stable to perturbation. The definition of an attractor trajectory of a system formalizes our notion of "the time-evolving point the system is tending towards" and prescribes an experiment one can perform to measure the attractor trajectory at an interaction port.

We defined the *virtual trajectory* as the attractor trajectory of the model of the neuromuscular system after all non-nodic effects such as inertia and joint friction were accounted for; it is, in effect, one aspect of the direct mechanical expression of the neural commands. Since attractor trajectories can be measured, so too can the virtual trajectory, at least in principle, although we currently lack sufficient detailed knowledge of the neuromuscular system to make a convincing measurement of the virtual trajectory.

Although it is necessary to measure virtual trajectories to prove the hypotheses (since they make specific claims about the form of the virtual trajectories), it is nonetheless possible to disprove them without making such measurements. Based on the argument that there will be increasing distortion of the attractor trajectories as we consider larger and larger subsystems (starting from the subsystem whose attractor trajectory is the virtual trajectory), we claimed that if the virtual trajectory hypotheses were true, we would observe that the attractor trajectory of the subsystem which excludes inertial effects would be simpler or more consistent across the workspace and across time and magnitude scalings of the task than would be the unperturbed trajectories. If we observe increasing complexity or less consistency of the attractor trajectories after we have accounted for inertial effects, then accounting for other non-nodic phenomena can only result in even more complex or discrepant attractor trajectories, so we can reject the hypotheses without requiring more detailed models of the neuromuscular system. If, therefore, the hypotheses failed this test, the disproof would be strong because it would be independent of assumptions about the limb's neuromuscular dynamics.

CHAPTER

THREE

THEORY OF ATTRACTOR TRAJECTORY MEASUREMENT

Problem Formulation

In the previous chapter, we described what an attractor trajectory is and presented a basic outline of the experiments we will need to perform to test the virtual trajectory hypotheses. Here we describe the means by which we measure attractor trajectories and discuss the various assumptions and parameter computations which are required.

We begin by considering a model of our subjects which allows us to represent interactions between the subjects and their environment through a mechanical interaction port at which we can measure force and velocity (as well as position and acceleration):

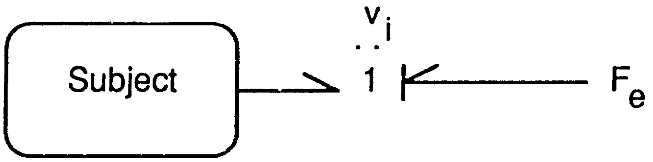

Figure 7. Pseudo-bond graph representing the subjects interacting with their environment.

In this model, all of the internal details are hidden — subjects are capable of initiating and executing motor actions, but all we are able to observe is the final result of those internal processes as they are made manifest at the interaction port. We do know, however, that if the subject is sufficiently constrained (i.e., if their wrist is immobilized, shoulder position fixed, and arm confined to a plane), then the position of their arm can be inferred directly from a measurement of the interaction port position; in other words, if the number of degrees of freedom we allow the subject is equivalent to the number of coordinates needed to express the position of the interaction port, then the subject can be characterized as kinematically non-redundant. This assumption is formalized in Figure 8 where the inertia of the subject's arm is shown to be connected directly to the interaction port (1-junction) and, as discussed in the previous chapter, the mechanical behaviour of the remaining components of the arm (neuromuscular system, joint dynamics, etc.) is represented by a Norton-equivalent subsystem (see dotted box) consisting of a flow source term, $v_o(t)$, connected through a nodic impedance, $Z_o(t)$, to the interaction port. By assumption, this subsystem has attractor properties, so we can in principle determine its attractor trajectory; this will be the path the interaction port will follow if $F_z = 0, \forall t$. Since the impedance, $Z_o(t)$, is nodic, F_z will only be zero if $v_o(t) = v_i(t), \forall t$; this Norton-

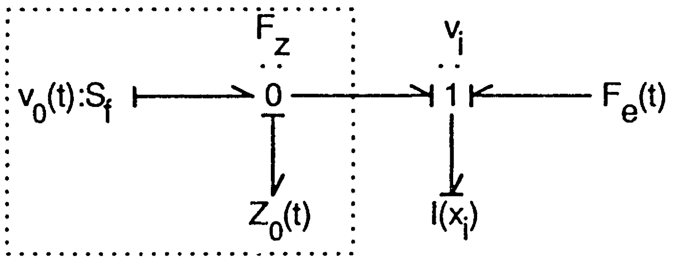

Figure 8. Pseudo-bond graph representing a system in which the inertial effects are separated from the nodic impedance effects and connected directly to the interaction port (the 1-junction where the subject grasps the manipulandum and where velocity is measured).

equivalent flow source, $v_o(t)$, is therefore the preinertial attractor trajectory which we desire to measure.

If we know the inertial parameters for the subject, along with the acceleration and velocity of the interaction port, we can compute the net force applied to the inertial linkage, I; let us designate this force F_i. Knowing F_i, we can compute the impedance force, F_z. from the simple equation $F_i = F_e + F_z$. Note that F_e is to largely under our control; by applying different force histories or supplying different environments for the subject to interact with, we can influence the achieved trajectories, $v_i(t)$. Our goal, therefore, will be to adjust the interaction forces until the achieved trajectory is in fact the attractor trajectory.

Note that the criterion for assessing whether or not a given trajectory is the attractor trajectory is applied after an experimental run — i.e., it is a post-test assessment: if, after applying a sequence of interaction forces, F_e, we found that $F_z = 0, \forall t$, we would conclude that $v_o(t) = v_i(t)$; i.e., that the path we had just measured was the attractor trajectory. To be absolutely sure that we have found a subject's attractor trajectory, we would have to cause their hand to actually track it. However, since the subject's execution of the task is subject to variability, we are not likely to have applied the correct F_e to make $F_z = 0$ for any particular run; the best we can hope for is a series of runs in which $F_z \approx 0$ (relative to the estimated inertial forces, F_i). How then can we derive an estimate for the virtual trajectory from our collection of executed trajectories, v_i, and estimated impedance forces, F_z?

Estimating the Attractor Trajectory from Experiments

If we do indeed have a series of runs in which $F_z \approx 0$, then we could quite reasonably take as our estimate of the attractor trajectory the mean of the collection of the executed trajectories. However, since the attractor trajectory is not known prior to the experiment, it is not likely that our first series of runs will yield $F_z \approx 0$, so we must adopt an iterative procedure in which the data we obtain from a given series of runs allows us to refine our estimate of the attractor trajectory and adjust the forces applied to the subjects so as to take them closer to their attractor trajectories in subsequent experimental runs. This necessarily involves extrapolating from (or, if we have been fortunate enough to have bracketed the attractor trajectory, interpolating between) the measured trajectories.

In principle, we could perform this extrapolation by casting the problem as a function mapping problem and solving it using a neural network. From a function mapping point of view, we are trying to find the trajectory, v_o, which maps to $F_z(t) = 0$. One could envision training a neural network on the collection of (v_i, F_z) and then taking as our estimate of the

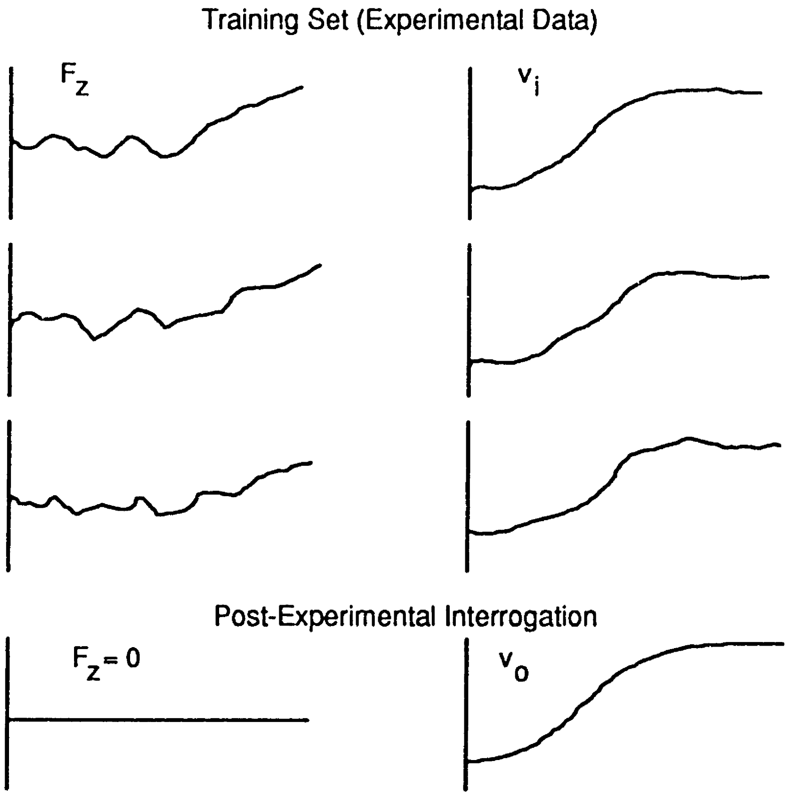
Training Set (Experimental Data)

Post-Experimental Interrogation

Figure 9. Sketch of data pairs which might be used to train a neural network to provide an estimate of the attractor trajectory. After training, a zero impedance force history is supplied as input and the resulting output is taken to be the estimate of the attractor trajectory.

attractor trajectory the output of the network when $F_z = 0$ is supplied as the input (see Figure 9).

If the average F_z were approximately zero throughout the movement, we would expect the output of the network to be a good estimate of the attractor trajectory. However, in those regions of the training set where none of the F_z are near zero (as in the latter half of the histories shown in Figure 9), the neural network is projecting rather than interpolating, and the network's estimate is likely to be incorrect because of the strong nonlinearities in the limb impedance. Nonetheless, if the correction is in the right direction and the step is not too large, the set of impedance forces after the next round of experiments will be closer to zero and we can repeat the procedure until we achieve convergence.

Since the convergence tests are applied directly to the experimental measurements, it is immaterial how we choose to generate the perturbations which produce the measurements, as long as the corrections ultimately result in convergence. In fact, I use a structured model of the impedance based on postural measurements to generate the corrections (discussed in the following section); this generally results in acceptably rapid convergence of the attractor trajectory estimates.

Iterative Convergence Using a Linear Model

Using a Linearized Model to Estimate the Attractor Trajectory

In the experiments we perform, we apply forces to the subject based on our current estimate of their attractor trajectory and then measure both these forces and the position of the manipulandum during the movements. If we find, upon analyzing this data, that we have not found the attractor trajectory (as indicated by non-zero impedance forces), how do we refine our estimate of it? As discussed above, we can assess proximity to the attractor trajectory independently of the means by which we generate the intermediate estimates of the attractor trajectory in the process of convergence. We can therefore use any model of the subject's arm we like to generate a new attractor trajectory estimate without jeopardizing our claim that the final estimate is independent of any modelling

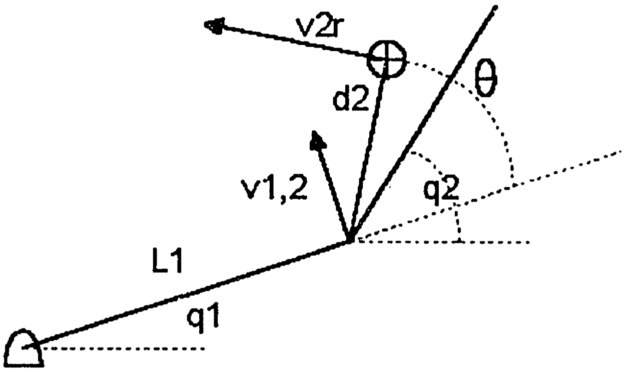

Figure 10. Schematic of two-link arm.

assumptions. Since there is a reasonable collection of data on the stiffness and damping properties of the arm at rest in configurations similar to those used in our experiments, I propose to use a linearized model of the arm's impedance to generate the new estimates of the attractor trajectories. The following section outlines the model I use.

Figure 10 is a schematic of a two-link arm which shows the kinematic variables used in the equations of motion. Note that this schematic allows for the possibility that the center of mass of the forearm does not lie on a direct line between the elbow and the endpoint; although this may be strictly true for arbitrary kinematic linkages, for the human arm, the center of mass is very near the limb's axis, so we will make this approximation (i.e., $\theta = q_2 - q_1$). The meaning of the various symbols in the figure is summarized in Table 1.

The equations of motion for this arm model (inertial effects only) be written as

$$H(q)\ddot{q} + C(q,\dot{q}) = \tau \, .$$

Symbol	Meaning
L_1,L_2	the link lengths (L_1 is the upper arm length, and L_2 the forearm)
d_2	the vector from the elbow joint to the center of mass of link 2; the x-axis of the coordinate frame is oriented along the line connecting the elbow to the endpoint
q_1,q_2	the absolute angles describing the orientation of the links relative to a coordinate frame with its zero at the shoulder joint.
v_{2r}	velocity of the center of mass of link 2 relative to the elbow joint
$v_{1,2}$	velocity of the joint between links 1 and 2

Table 1. Meaning of symbols in Figure 10.

If we model the net effect of the neuromuscular and other non-inertial arm dynamics as shown above in Figure 8, and if, further, we linearize the impedance, Z_o, then we can write that the net applied torque is given by

$$\tau = \tau_e + B(\omega_o - \omega) + K(q_o - q) = \tau_e + B\delta\omega + K\delta q = \tau_e + \tau_z,$$

where the subscript on the position and velocity terms refers to the attractor trajectory, the absence of subscript implies that they are the actual endpoint positions and velocities, and B and K are 2x2 matrix approximations of the subject's impedance. Given a knowledge or model of the inertial parameters in the matrices H and C, and having measured τ_e, we can compute τ_z for any given trial. Assuming that the subject starts at rest and has been at rest for long enough that the attractor trajectory and the actual hand position are identical, then both $\delta q(0) = 0$ and $\delta\omega(0) = 0$. We can therefore write a differential equation for the deviation between the attractor and the actual trajectories as

$$B\delta\omega + K\delta q = \tau_z.$$

This equation is straightforward to integrate, and the result is an estimate of the attractor trajectory based on a linear projection from the trajectory executed during a single movement. The accuracy of the new estimate will clearly depend on several factors, most notably the accuracy and applicability of the linearized impedance matrices, although if the impedance torques are zero throughout the movement, then the deviation between the attractor and actual trajectories is also zero and the estimate is independent of the values of B and K; all we require is that Z_o be zero when the actual trajectory is the attractor trajectory, and this is so by definition. In the following sections, we first consider how we control the subject's interaction with the manipulandum to cause them to track our new estimate of the attractor trajectory, and then we consider in more detail how to assign error bounds to an estimate and how to combine the results of multiple trials to come up with our final estimate of the mean attractor trajectory for a particular task.

Tracking Attractor Trajectories Using Sliding Mode Control

We have described how we develop a refined estimate for the attractor trajectory; how then do we make the subject track along this new estimate? We do so by controlling the manipulandum using a sliding mode control algorithm [Slotine 91]. This class of controller is a trajectory tracker which utilizes a feedforward signal based on the nominal model of the system being controlled (in our case, two coupled inertial linkages since along the attractor trajectory the subject ought not to be developing any non-inertial impedance forces) and a discontinuous feedback control law which compensates for any errors in the nominal model within certain bounds. Sliding mode controllers are robust to modelling errors in the sense that they allow explicit tradeoffs between model uncertainty and control effort and, in the modification incorporated here, between tracking performance and abruptness of the control discontinuity. The design of the controller requires the selection of several parameters including the level of model uncertainty, the time to converge to the control boundary layer, and the control bandwidth (λ). These parameters were selected through experimentation and adjusted so that the typical tracking performance was within 2 cm of the desired position throughout the motion.

Sources of Error

On a single trial, there are three major sources of error when using the linearized model to project an estimate of the attractor trajectory. An analog of the process of finding a multidimensional attractor trajectory, which has both two reference positions and two reference velocities, is finding the zero of a nonlinear spring, as illustrated in Figure 11, which shows the force-displacement relationship for a nonlinear spring; we can understand the sources of error by considering their analogs in this simplified system.

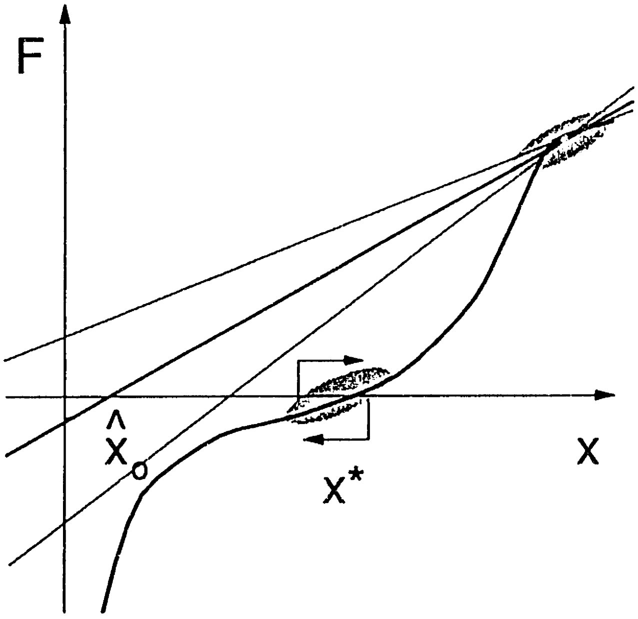

Figure 11. Using a linearized model to estimate the zero of a nonlinear function — in this case, a nonlinear spring. Errors in projecting from a point relatively distant from the crossover point are a function of how nonlinear the function actually is, as well as on how accurately we can determine the local slope. The only way to accurately locate the crossover point is to find an interval (shown by arrows above) across which the function goes from being definitely positive to definitely negative. Noisy measurements will increase the size of this interval.

The first source of error is the degree to which the linear model is a poor approximation to the actual impedance behaviour. Although we know that in general the behaviour of muscles is highly nonlinear, nevertheless a number of studies have had reasonable success in accounting for the response of a single joint (or even the whole arm) to small or moderate perturbations using a linear model; often, the variance in the force records accounted for by the linear model exceeds 70% [Bennett 92]. Since we intend to be concerned about error bounds on our attractor trajectory estimate only at the final stage when we have essentially achieved convergence, we will be projecting over relatively short distances where the linearized approximation has been shown to be reasonable; in reference to Figure 11, we intend to be computing error bounds only when we are near the crossover point (the zero of the function) rather than when we are still a fair distance away. I will therefore not attempt a more detailed analysis of this source of error because in this regime the linear analysis will account for most of the error.

The second major source of error is in the computation of the inertial torques, partly due to errors in estimating accelerations and velocities, but more significantly due to errors in the inertial parameters themselves. In the context of Figure 11, these errors manifest themselves as errors in the force measurements, which would have the effect of shifting all the measured points vertically some distance. In a later section, I will describe how I estimate the inertial parameters and ascribe error bounds to these parameters.

The final major source of error is how accurate the linearized impedance estimates are. An overestimate of stiffness, for example, will clearly lead to an underestimate of the deviation between the attractor and achieved trajectories. Again, in the context of Figure 11, errors in the impedance parameters correspond to errors in the estimate of the slope of the impedance relation in the vicinity of the data points. If we can place bounds on the impedance parameters, however, we will be able to compute bounds on the location of the function's zero (and, by analogy, on the location of the attractor trajectory) — see, for example, the region of the X axis between the two lines projecting from the top right data cluster in Figure 11. We now outline this procedure.

Assigning Error Bounds

Consider again the equation we are integrating:

$$B\delta\omega + K\delta q = \tau_z.$$

Assume for the moment that the problem is one-dimensional (i.e., B, K, τ_z and the position and velocity elements are all scalars). For a given combination of B, K, and τ_z, we can write an expression for $\delta q(t)$ using the convolution operator

$$\delta q(t) = \int_0^t G(\tau)\tau_z(t-\tau)d\tau$$

where $G(t)$ is the impulse response function of the system whose Laplace transform is

$$G(s) = \frac{1}{Bs+K}.$$

If we know that B and K are bounded above and below their nominal values, then $G(t)$ will have a maximum and minimum value for each time, t, which we can compute knowing these bounds on the impedance parameters (see Figure 12). I.e., we define a new impulse response operator, $\bar{G}(t)$, which can take on any value in the range associated with the given bounds on B and K. Similarily, we define a new impedance torque function, $\bar{\tau}_z(t)$, which can take on any value in the range given by the bounds on $\tau_z(t)$. We can then ascribe a minimum and maximum value for $\delta q(t)$ by performing the convolution shown above, choosing at each instant, τ, in the integration interval [0,t] the combination of

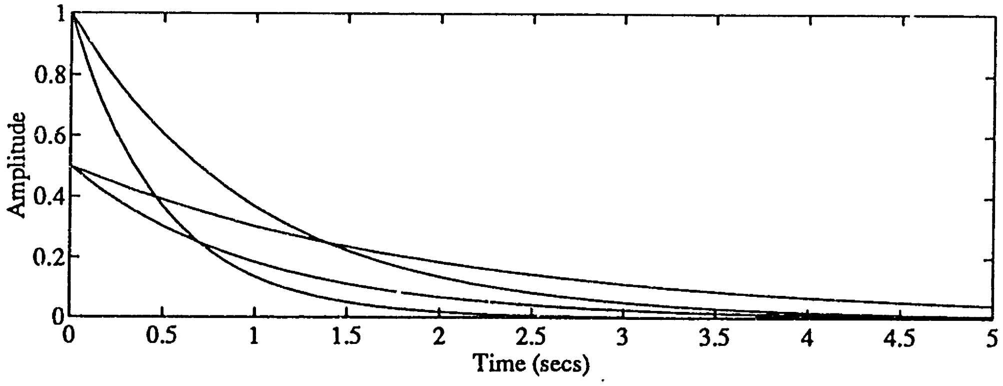

Figure 12. Bounds on the impulse response function of a first-order system as a function of bounds on the stiffness and damping parameters — $1 \leq K \leq 2$, $1 \leq B \leq 2$. The bounds are determined from the topmost and bottommost curves at each instant.

values of $\overline{G}(t)$ and $\overline{\tau}_z(t)$ which minimize or maximize that instant's contribution to $\delta q(t)$. The results of these two convolutions give us our minimum and maximum bounds on $\delta q(t)$.

The above analysis is easily generalized to two dimensions. In 2-D, τ_z and δq are 2-vectors and B, K, and $G(s)$ are 2x2 matrices. To compute the bounds on δq, we simply compute the maximum or minimum dot product of each row vector in $G(s)$ with the impedance torque vector.[11] We compute the bounds in joint coordinates because the impedance matrices are less position-dependent in that frame than in a cartesian frame and then convert the bounds to cartesian space by computing the coordinates of the corners of the rectangle defined in joint space by the various combinations of δq_{min} and δq_{max}.

---

[11] To avoid assigning excessively large (conservative) error bounds, we do make the assumption that each element of $G(s)$ is adjusted in the same direction. For example, if we decide to set $K(1,1)$ to its maximum value, then we also set $K(1,2)$ to its maximum; we do not allow them to vary independently.

Combining Multiple Trials

The procedure outlined above gives us a means of generating an estimate of the attractor trajectory for a given movement, along with associated error bounds. A subject, however, does not perform the task the same way each time. They may start earlier or later, take a longer or shorter time to make the movement, start and stop in slightly different positions, or take a somewhat different path to the target each time. How then do we combine a series of attractor trajectory estimates (with error bounds) into an estimate of the mean attractor trajectory, and how do we combine the error bounds themselves?

The first part of this question is reasonably straightforward to answer. We are, for the most part, concerned primarily with the shape of the trajectories. Before averaging a series of trajectories, then, we would like to normalize them in both position and time so as to preserve as much of the shape as possible.[12] Since each trajectory starts and ends at a well-defined location, we can remap each of them so that they start at (0,0) and end at (1,0). To normalize them in time, we compute the optimal match between each trajectory and some reference (analytical) function whose transition time is easily characterized (such

---

This assumption is based on the observation of Mussa-Ivaldi that the stiffness ellipses are essentially invariant in shape, even under voluntary modulation of overall stiffness levels.

[12] Note that if we do not normalize in time, then the averaging process distorts the shape of a trajectory. Consider averaging two ramps of exactly the same duration, but starting at slightly different times. The average will start rising with half the slope of the ramps when the first ramp rises, will switch to rising at the same slope as both ramps once the second ramp begins rising, and will finish with a third segment, again of half the ramp slope, which ends when the second ramp ends. Non-normalized averaging therefore creates a time-stretched, three-region "average" trajectory out of two perfectly well-defined single-region ramps. Similarily, failing to normalize in position can wash out smaller features of the movement.

as a ramp or minimum-jerk transition, for example); each trajectory is then normalized in time by dividing through by the transition time of the optimally-matching reference function. Once the trajectories have been normalized we compute a weighted-average attractor trajectory by combining all the estimates at each sampling instant in proportion to the inverse of the area enclosed by the error bounds — the smaller the error bounds, the greater the weight given to that estimate.

Combining the error bounds is a more complicated process. The bounds cannot be regarded as standard errors which can be reduced by a greater number of samples. Rather, they reflect our fundamental uncertainty as to the values of the inertial and impedance parameters and so cannot be reduced through analyzing more trials. What is reduced by this process, however, is the variability inherent in the subject's ability to reproduce a given attractor trajectory. We might assume, as a first approximation, that subjects' variations in executing attractor trajectories are comparable to their variations in executing hand trajectories.[13] The standard error of the mean of the attractor trajectory estimates will decrease with the square root of the number of samples, so we can compare the variability of achieved trajectories with the error bounds of the individual attractor trajectory estimates to determine their relative sizes. This will help us decide how many trials we should have the subjects perform before shifting our efforts to reducing the attractor trajectory error bounds directly, since the total error is the sum of the error due

---

[13] This is a reasonable assumption, although since the inertia of the arm functions as a low-pass filter, there may in fact be greater variability in the attractor trajectories without it being apparent in the achieved trajectories.

to inertial and impedance uncertainties and to the variability of the attractor trajectories themselves.

How then do we combine the error bounds from the individual attractor trajectory estimates? I have adopted the pragmatic approach of computing the "average" error bound at a point in time by averaging the error bounds of each attractor trajectory estimate weighted by the inverse of the area enclosed by the error bounds. This has the effect of making the average error bounds no smaller than the smallest bound and no larger than the largest, which is the effect desired since no increase in the number of trials ought to reduce this lower bound. As an example of this averaging scheme, suppose one error bound is characterized by a linear dimension of twice another; the average error bound will have a characteristic dimension of

$$\frac{L_1 w_1 + L_2 w_2}{w_1 + w_2} = \frac{L_1 \frac{1}{L_1^2} + L_2 \frac{1}{L_2^2}}{\frac{1}{L_1^2} + \frac{1}{L_2^2}} = \frac{1 \frac{1}{1^2} + 2 \frac{1}{2^2}}{\frac{1}{1^2} + \frac{1}{2^2}} = 1.2,$$

which is between L_1 and L_2, but decidedly closer to L_1.

Basis for Selecting Linearized Impedance Parameters

In order to implement the scheme outlined above to update the attractor trajectory estimates and to assign error bounds, we have to specify the linearized impedance parameters. On what basis do we do so? There is a very small number of direct studies of the impedance of an arm moving in a plane [Gomi 92, Lacquaniti 93], and relatively few of single-joint motion. Although some of the studies of single-joint motion show a fairly wide range of stiffnesses [Lanman 80, Bennett 82, Saund 84], these studies generally are performed in an experimental setting in which the subject is able to relax almost completely. In constrast, the studies described in this thesis require the subject to support the weight of their own arm and to maintain an active grip on the manipulandum. Under these conditions, the range of achievable stiffnesses is significantly diminished. For example, Hodgson found that in a task which required the subject to resist a force applied perpendicular to an outstretched arm, the total stiffness exhibited by the subject varied by a factor of only two or three from minimum to maximum [Hodgson 90]. Gomi's study of stiffnesses during movement reveals that the minimum stiffness is of the same order as the postural stiffness, while Bennett claims a reduction during motion to a minimum of approximately 50% of the postural levels. Mussa-Ivaldi *et al.* measured the stiffness matrices of a number of subjects at posture in the same configuration we used in this study; they found striking regularities in the shape and orientation of the stiffness ellipses across the workspace [Mussa-Ivaldi 85]. Dolan, in a series of postural measurements, confirmed these results regarding shape and orientation and found that in a set of eleven subjects, the postural levels varied by a factor of ten from the least stiff to the most stiff subject, although it is not clear how much of this range is due to differences in intrinsic muscle properties or tone and how much to differences in active grasp [Dolan 91]; excluding the three most extreme subjects, however, left eight subjects whose stiffnesses differed by less than a factor of two from one another, so for the most part, we can expect our subjects to exhibit relatively similar postural stiffnesses.

For the purposes of updating the attractor trajectory estimates, therefore, I selected as my nominal stiffness estimate a value roughly comparable to that of Mussa-Ivaldi's Subject A; this value is at the high end of the group of subjects in Dolan's study. For the damping matrix, I selected a value equivalent to the stiffness matrix multiplied by a factor of 0.07 *s*, which is a typical ratio determined by Dolan in his postural measurement tests.[14] For the purposes of assessing the error bounds on the attractor trajectory estimates, I selected minimum and maximum impedance matrices equal to the nominal matrices plus and minus 50%; this range covers stiffness values ranging from the midrange of Dolan's core group to values 30-40% larger than the highest stiffness he measured among his eleven subjects

Potential for Model-Independent Estimate of Errors

One possible critique of the procedure described above is that we are relying for our error bound estimates on a model whose structure and parameters are open to question. Is there then any reasonable way to derive a model-independent estimate for our error bounds? If we wish to establish mathematically defensible bounds, then I believe not. Assigning bounds on the possible difference between the attractor and measured trajectories requires that we have some way to convert non-zero force histories to limits

---

[14] See Table 3-9 on page 58 of Dolan's doctoral thesis [1991].

on displacements; by definition, this implies the use of an impedance measure. However, if we are willing to sacrifice some rigour, we may be able to approach the problem indirectly. In the process of converging to an attractor trajectory estimate, we (hopefully) reduce the mean impedance force levels, at least in the early stages. By keeping track of some measure of the changes in the attractor trajectory estimate (say, the RMS distance of the new estimate from the previous estimate) as we change the impedance forces, we may be able to place an upper bound on that measure of change as a function of the magnitude of the impedance force history. Such an estimate would not be rigorously defensible, but it would be based on the argument that, in our experience, we never see changes in the attractor trajectory estimate of more than X units when we reduce the impedance force magnitudes by Y units.

Potential for Estimating Impedance Parameters Experimentally

If part of the reluctance to accept the model-based error estimates is due to concern that the assumed range of the subjects' impedances is too narrow, then it is reasonable to ask whether or not we could estimate their impedance parameters during the experiment. An exploratory analysis of this question, not reported here, concluded that given our experimental paradigm, which requires a series of successful unperturbed movements between the randomly applied perturbations, we would require an unreasonably large number of trials to generate reasonable impedance estimates. The impedance estimates based on twenty perturbations (the upper limit of what we would consider to be an experimentally practical number of perturbations) are too sensitive to variability in the execution of the attractor trajectory and to sensor noise to be useful. One possible alternative, not pursued in this study, would be to dissociate the impedance measurement problem from the attractor trajectory location problem; if, for example, we could develop a technique for measuring the impedance about the nominal trajectory during prototypical movements, we could argue that the error bounds on these impedance estimates should be the error bounds used in the error analysis described earlier. Since the current uncertainties are large, this might be a means of significantly reducing the uncertainties. Lacquaniti recently proposed a method for performing time-varying impedance estimates during multi-joint motor tasks, so if his method proves to be sound, we may be able to adapt it to this problem [Lacquaniti 93].

Estimating Inertial Parameters

We now turn our attention to the problem of identifying the inertia of a human subject. Early on, I considered attempting to measure the inertia matrix of a subject through measurements made at the interaction port while perturbations are applied via the manipulandum. Such measurements would be in the same spirit as Dolan's impedance measurements. Dolan, however, found relatively large errors in the inertia matrix, on the order of 25% or so, particularly when the matrix was expressed in joint coordinates. The largest errors were associated with the estimate of the inertial parameters of the proximal link of the arm, which Dolan attributed to the fact that the excitation was applied directly to the distal link, but only indirectly to the proximal, which resulted in a weaker excitation

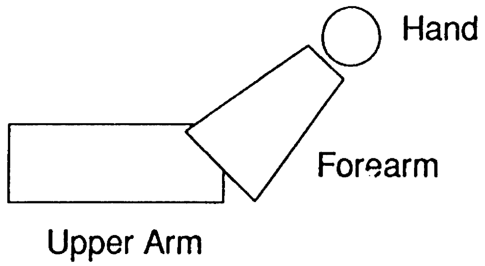

Figure 13. Model of arm.

of the proximal link. Because of these large errors, I felt that we could obtain a better estimate of the inertial parameters of the arm by modelling its shape and computing an estimate based on anthropometric measurements.

Model

The basic model used to compute the inertia of a subject's arm is a cylinder for the upper arm, a truncated cone for the forearm, and a lumped mass at the hand consisting of the mass of the hand (computed from its volume) and the portion of the force transducer handle outboard of the force sensor (see Figure 13). We will attempt to verify the accuracy of this model by comparing the anthropometric volume estimates to direct volume measurements (based on an immersion test), as well as by comparing the computed inertial parameters to those measured by Dolan. We will then assess the sensitivity of the inertial parameter estimates to measurement errors and consider the effect of ignoring structural uncertainties such as residual shoulder and wrist movement.

Accuracy of Anthropometric Estimates

An inertial estimate based on the cone and cylinder model is subject to several sources of uncertainty including errors in measuring the various lengths and circumferences, errors in the assumed tissue density, and deviations of the actual limb shape from the assumed shape. The length measurement errors are quite small - on the order of 0.5-1.0 cm in 20 or 30 cm (about 1.5-3%), and we can compute the sensitivity of the inertia matrix values to such errors. We can get some idea of the deviation of the actual limb shape from the assumed shape by measuring the volume of a limb by immersion in water and then comparing this measured volume to that given by the cone and cylinder model. For example, consider the measurements for one subject shown in Table 2. Comparing the direct volume measurements with the predicted volumes based on the cone and cylinder model, we see that the model underestimates the actual volume by about 10%. This probably happens because the arm "bulges" relative to the simplified geometric forms we assume in the models; we should therefore multiply the models' volume estimates by a factor of about 1.115 (equivalent to multiplying the circumference measurements by about 1.055) to compensate for this.

Measurement Type	Location	Value
Circumferences:	Upper arm	30 ± 0.5 cm
	Forearm (elbow)	27.5 ± 0.5 cm
	Forearm (wrist)	17 ± 0.5 cm
Lengths:	Upper arm	30 ± 1 cm
	Forearm	25 ± 1 cm
Volumes:	Hand	360 ± 10 ml
(from immersion test)	Hand + Forearm	1475 ± 25 ml
	Forearm	1115 ± 35 ml
	Whole Arm	3875 ± 50 ml
	Upper Arm	2400 ± 50 ml
Computed Values:	Forearm Volume	1000 ± 50 ml
	Upper Arm Volume	2150 ± 85 ml

Table 2. Measurements used to compute inertial parameters for one subject.

Estimation of Endpoint Mass

In order to compute the endpoint mass, we need to not only measure the volume of the subject's hand, but also estimate the mass of the force transducer handle. The reason this is slightly problematic is that there is no clear point of division between the "inboard" and "outboard" sides of the transducer because force is not measured at a point but is instead inferred from the response of strain gauges cemented to a set of beams. However, on the basis of an oscillation test in which I mounted various masses to the outboard side of the force transducer and measured the accelerations and forces generated when bandpasslimited random perturbations were applied by the manipulandum, I inferred that the mass of the handle which could be considered to be "outboard" of the sensing gauges was approximately 450 • 25 g.

Sensitivity of the Estimates

Based on the nominal measurements shown in Table 2 and the estimated mass of the handle, the estimated inertia matrix is

$$\begin{bmatrix} 0.2674 & 0.1213\cos(q_2 - q_1) \\ 0.1213\cos(q_2 - q_1) & 0.1095 \end{bmatrix} \text{kg·m}^2.$$

The sensitivity[15] of each of the three constants in the inertia matrix is summarized in Table 3. The inertia estimate is not particularly sensitive to any of the measurement errors; most of the errors enter into the estimate with a sensitivity less than unity. The major exception is the uncertainty in the upper arm length which affects the inertial estimate by a factor of 2.3 because both the volume of the arm and its moment of inertia depend on this measurement.

---

[15] Sensitivity is defined as the ratio of the fractional change in the inertial parameter to the fractional change in the measured variable.

Parameter	h_{11}	h_{12},h_{21}	h_{22}	Expected Error (%)
Handle Mass	0.15	0.38	0.49	5.6%
Hand Volume	0.10	0.21	0.24	2.8%
Upper Arm Length	2.34	1.00	0	3.3%
Forearm Length	0.41	0.90	1.12	4%
Upper Arm Circumference	0.63	0	0	1.7%
Forearm Circumference	0.55	0.31	0.15	1.8%
Wrist Circumference	0.29	0.34	0.24	2.9%
Density	0.85	0.62	0.51	~3%
Forearm Distribution	-0.001	0.17	0.15	~3%

Table 3. Sensitivity of inertial parameter estimates to measurement errors and maximum expected measurement error (%).

The uncertainty in tissue density, which relates mass to volume, is another source of parametric error. A number of cadaveric studies, most notably a detailed Air Force study [Young 83], have generated density estimates for the arm of approximately 1.1X the density of water; this value can vary by a factor of about 3-5%, so our mass estimates will be uncertain to a similar extent, although, since the mass of the handle is measured independently of this density assumption, the sensitivity of the inertia estimates is less than unity (see Table 3).

The final primary source of parametric uncertainty is the effect of the difference in tissue distribution between the cone or cylinder model and the actual arm. This was partially accounted for by the compensation factor determined by the immersion experiment, but without a detailed study of the arm it is difficult to fully account for this factor. One could, however, estimate the effect by assuming a change in the shape of the arm which preserves the arm's volume. For example, suppose we change the description of the cone by adding a couple of centimeters to the wrist measurement and taking away from the elbow circumference an amount sufficient to preserve the arm's volume. This will cause a distal shift in the forearm's center of gravity and will thereby change the inertia. Such a distribution uncertainty enters into the inertia estimate with a sensitivity of less than 0.2, which is quite minimal.

We can compute an estimate of the errors in each component of the inertia matrix by assuming that the square of the overall percentage error is equal to the sum of the squares of the percentage errors in each of the individual measurements, weighted by their sensitivities. By applying this procedure, we find that our parametric inertia matrix estimate is likely to be in error by approximately 7% in each component — a considerably smaller error than we would obtain by applying Dolan's procedure.

Structural Uncertainty

The model described above assumes that the subject's arm has only two degrees of freedom during the experiment. Since we cannot perfectly fix the subject's shoulder, immobilize their wrist, and be certain that their elbow behaves as a simple pin joint, however, this model will be somewhat in error. How much effect will these residual motions have on the subject's behaviour when performing the experiments? I will return to this question in more detail when I consider possible critiques of the experimental methodology in the following chapter, but for now I would argue that since our conclusions will rest primarily on the behaviour of the arm during large scale motion, small movements at the wrist and shoulder ought not to exert much influence on the results; these unmodelled joint movements would be far more problematic if we were trying to interpret small amplitude oscillations as indicative of the inertia of the whole arm.

Comparison With Experimentally-Determined Inertial Estimates

Since we are proposing to use a model to compute the inertial parameters, it would be worthwhile to compare these estimates to some experimentally derived estimates to see if the experimental approach yields plausible values. Although a number of experimental methodologies are possible, Dolan has performed perhaps the most comprehensive set of inertial measurements to date [Dolan 91]. He performed impedance parameter identification experiments at posture using an apparatus quite similar to the one used in these studies. In his experiment, he applied small (< 10 mm) displacements in eight equally spaced directions from the equilibrium position and fitted his data to a model expressed in Cartesian coordinates which did not take account of spatial variations in the impedance matrices. Nonetheless, he was able to obtain reasonably consistent stiffness, damping and inertia estimates, so we will consider his results in more detail.

Figure 14 shows the inertial ellipses measured by Dolan (expressed in joint coordinates), along with the inertial ellipse as computed by the anthropometric procedure outlined

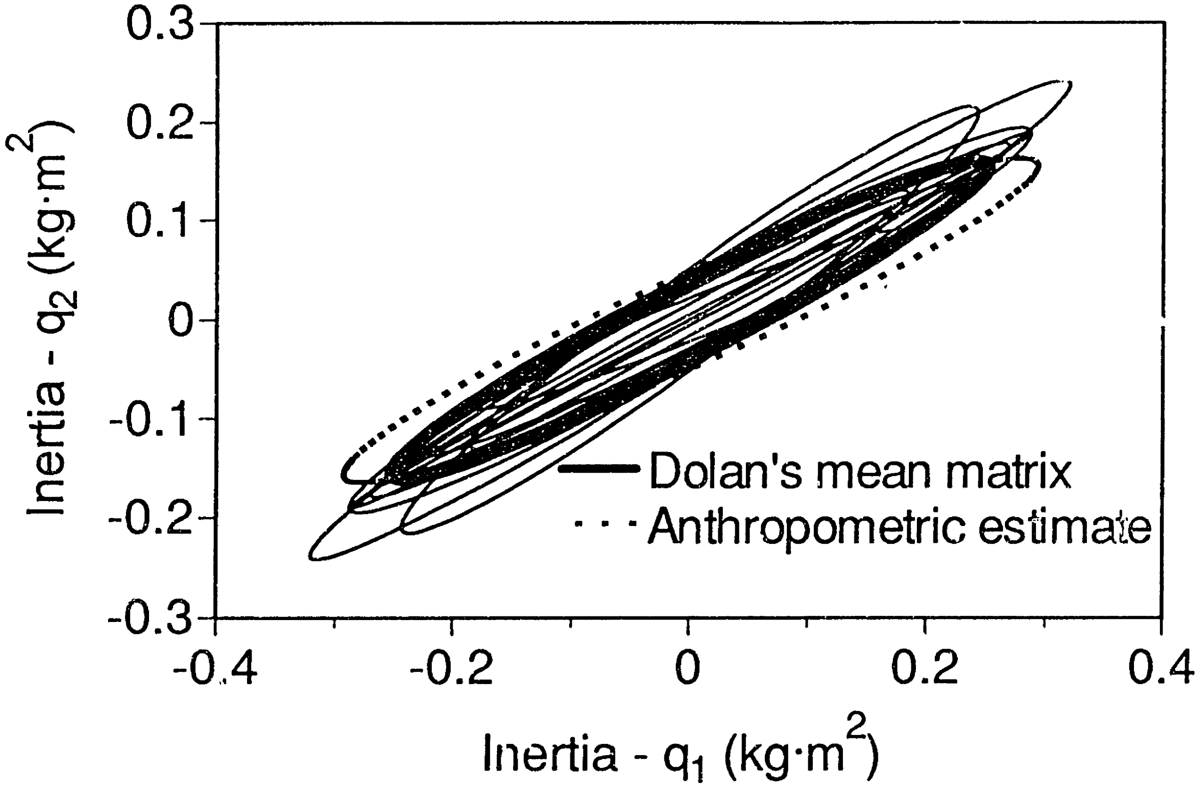

Figure 14. Comparison of Dolan's measured inertial ellipses with the anthropometric estimate.

above (the medium-weight line). The ellipse with the heaviest outline is a "typical" inertial ellipse computed using the mean main diagonal term and the mean ratios of the lower and off-diagonal terms to the main diagonal term. This mean inertial matrix is

$$\begin{bmatrix} 0.210 & 0.145\cos(q_2-q_1) \\ 0.145\cos(q_2-q_1) & 0.065 \end{bmatrix} \mathrm{kg{\cdot}m^2},$$

and it can be scaled up and down by a factor of about 1.3 to cover the range from the smallest to the largest inertia matrix in Dolan's sample of eleven subjects. The parametric estimate is reproduced below to facilitate comparing the two matrices:

$$\begin{bmatrix} 0.2674 & 0.1213\cos(q_2 - q_1) \\ 0.1213\cos(q_2 - q_1) & 0.1095 \end{bmatrix} \text{kg·m}^2.$$

It is clear that the parametric estimate is in the same ballpark as Dolan's estimate, with the biggest discrepancy being in the lower diagonal term, although it is unclear whether or not Dolan's estimate includes the mass of the handle of his apparatus. If we were to recompute the parametric estimate assuming that the mass of the handle were zero, the inertia matrix would be

$$\begin{bmatrix} 0.227 & 0.075\cos(q_2 - q_1) \\ 0.075\cos(q_2 - q_1) & 0.056 \end{bmatrix} \text{kg·m}^2.$$

In this case, both the main diagonal terms are considerably closer to those reported by Dolan, although the discrepancy between the off-diagonal terms has been exacerbated. Dolan's estimate is problematic, however, because the inertia matrix must be positive definite, representing as it does a real kinematic linkage. For this to be true at arbitrary elbow angles, the off-diagonal terms must be lower in magnitude than the geometric mean of the main diagonal terms. For Dolan's mean inertial estimate shown above, this would imply that the maximum off-diagonal term consistent with a real physical system would be 0.117; his value of 0.145 is therefore physically unrealistic. Part of the reason for this is that many of his measurements come from places in the workspace where the position-dependent off-diagonal terms are relatively small by virtue of being multiplied by the cosine term (which is near zero at an elbow angle of 90 degrees), so we cannot expect that his estimate of the position-independent terms will be particularly accurate because they are derived by dividing a small number by a cosine value close to zero; this will inevitably lead to large errors. I am therefore inclined to believe the parametric estimates, especially for the off-diagonal terms, because they are derived from a physically realistic model and they agree with the experimental measurements for the main diagonal terms.

Summary

The tests of the virtual trajectory hypotheses proposed in the previous chapter require comparisons between the trajectory executed by the subject interacting with the nominal dynamic environment and the attractor trajectory of the subsystem which excludes the inertial effects. We showed that this "preinertial" attractor trajectory is the source term in a nonlinear Norton-equivalent model of the non-inertial components of the arm and that it can be recognized by the fact that along this trajectory the applied forces precisely match the inferred inertial forces — i.e., that the impedance forces are zero. However, since the assessment of the trajectory is performed after it has been completed, we are unlikely to have succeeded in making the impedance force history zero. How then do we use our knowledge of the subject's response during the just-completed trial to improve our estimate of the attractor trajectory?

In principle, if we had a series of trials in which the impedance force histories were all near zero ("near" relative to the inertial forces), we could interpolate between these data records in some manner. For example, we could cast the problem of determing the relationship between an impedance force history and a measured trajectory as a function-mapping problem and use a neural network to solve it. If the set of impedance forces do not in some multidimensional sense bracket zero, however, then we would have to extrapolate from our data set. We proposed using a linearized impedance model with parameter values derived from postural measurements. Although projections using such a model will certainly be in error, as long as they point in the right direction, we should converge after several iterations. We demonstrated how to use the linear model to generate the new attractor trajectory estimate and described the use of a sliding mode controller to compel the subject to move along a path near to our current attractor trajectory estimate.

Given the intrinsic variability in the repeated performance of motor tasks, however, no single trial will give us a perfectly zero impedance force history. We therefore considered how we might characterize the error associated with our measurements and how we might combine attractor trajectory estimates from different trials to generate our estimate of the mean attractor trajectory. Each trajectory measured with the sliding mode controller engaged is an estimate of the attractor trajectory for that task; we showed how the error in each estimate is bounded by a convolution of the impedance forces and the impulse response function of the linearized impedance model. Given these error bounds on the individual attractor trajectory estimates, we described how to combine these estimates and their error bounds using an averaging scheme in which the weights are inversely proportional to the area enclosed by the error bounds. Although the current scheme still relies on a model to estimate the error bounds, we proposed an model-independent approach whose utility will be further assessed in the results chapter. We also assessed the prospects for determining the impedance parameters experimentally, but concluded that such experiments were not currently feasible because of the large number of trials necessary to obtain reliable estimates.

Finally, we considered how we might estimate the inertial parameters used to compute the impedance forces. Experimental approaches have led to large errors (> 25%) in the inertia matrix, so we decided that we could do better with a parametric model of the arm which represented the upper arm as a cylinder, the forearm as a truncated cone, and the hand as a lumped mass. The volumes predicted using this model were corrected for the shape of actual human arms by means of an immersion test. We assessed the sensitivity of the inertial estimates to errors in the anthropometric measures — virtually all sources of error enter into the estimate with sensitivities less than unity. We estimated the overall error to be approximately 7%, which is considerably lower than the errors associated with Dolan's experimental approach, even though the estimates based on the two different approaches are reasonably consistent with one another (except for the off-diagonal terms from the experimental measurements, which are physically impossible).

CHAPTER

FOUR

EXPERIMENTAL APPARATUS AND METHODS

In this chapter, I describe the experimental paradigm used to test the virtual trajectory hypotheses, the apparatus used and the tests performed. The primary experiments are designed to estimate the pre-inertial attractor trajectories of the subjects; secondary experiments are also needed to estimate the inertial parameters of the apparatus and of the subjects. Possible critiques of the experimental paradigm used to identify attractor trajectories are addressed at the end of the chapter.

Basic Experimental Paradigm

Selection of Prototypical Task

The primary experimental goal of this thesis is to make measurements of the pre-inertial attractor trajectories subjects exhibit during "normal" two-joint motor tasks executed at various locations about the workspace and at various speeds. A well-studied "normal" task is that of unconstrained point-to-point reaching movements. This class of motor task has been observed in both planar and fully three-dimensional movements and under conditions in which the subject's arm is well-modelled as either redundant[16] or non-redundant (the latter is typically through the use of some immobilizing device) [Atkeson 85, Cruse 87, Flash 87, Wolpert 94a,b, Lackner 94]. For the purpose of testing the virtual trajectory hypotheses, it is sufficient to consider a planar task in which the subject is restricted to two degrees of freedom.

We want an easily understood, easily learned task which can be executed repeatedly and which elicits significant dynamic forces. A point-to-point movement satisfies these criteria, especially at higher speeds. As I pointed out in the earlier chapter on measuring attractor trajectories, it is important to be able to make the assumption that the subject starts at equilibrium. We therefore cannot use the sorts of cyclic movements used by Bennett to measure impedance during motion [Bennett 92]; instead, we must begin with a rest period sufficiently long to make it plausible that the subject begins at equilibrium. Although it is not necessary from the point of view of making attractor trajectory measurements to require that the subject also come to rest at the end of a movement, such a requirement is easily communicated to the subject. Starting and ending at rest is also a

---

[16] The term redundant implies that the configuration of the arm is not determined by the task specification. If the task specification does determine the arm configuration, then the arm is said to be non-redundant with respect to that task. Note that these definitions are task-specific; if the number of constraints which a task imposes changes during different phases of the task, then the arm may switch from being redundant and non-redundant from one phase to another.

fundamentally symmetric task specification, so it will be easier to recognize and define "simple" reference trajectories against which we compare the measured trajectories . A point-to-point task is also a useful reference task because the same task can be performed in a variety of dynamic environments (for example, position- or velocity-dependent force fields) which will prove useful for more general tests of the virtual trajectory hypotheses. The task is performed without impairing the view of the subject in any way because this is characteristic of the conditions under which most normal motor tasks are performed. Nonetheless, I did perform tests to assess the effect of removing visual feedback during execution of the task and found that at least the first 75% of moderate speed movements is performed without showing evidence of visually-based corrections.

Repeated Execution of Discrete Tasks with Occasional Perturbations

In the earlier discussion of the iterative approach to identifying an attractor trajectory, I stated that forces would have to be applied to the subject's arm in order to drive the arm along the estimated attractor trajectory. In these experiments, the forces are applied by a robotic arm programmed to track along a specified trajectory — typically the current estimate of the attractor trajectory. These forces are not applied every time the subject executes the task, however, because it is fundamental to the method that we be sure that the subject is generating the same command history of neural activations every time the task is performed, even when the perturbation is being applied. To maximize the chance that this occurs, we train the subjects thoroughly on all the tasks they are asked to perform and require that they execute the unperturbed task successfully a number of times in succession prior to applying a perturbation. We also ask the subjects not to intervene when they do feel a perturbation, but simply to play out the command they were in the process of executing.[17] A discussion of the justifications for using this "do not intervene" paradigm can be found at the end of this chapter. The number of successful completions required prior to a perturbation is randomly determined to minimize the possibility of the subject anticipating a perturbation; the subject must successfully execute at least four consecutive unperturbed movements before qualifying for a perturbation. Once the subject is qualified, a perturbation is applied on a randomly selected trial; the cumulative probability of having received a perturbation approaches 1.0 after ten successful movements with a mean of approximately six movements. If the subject has qualified and then fails to make a successful movement, their "success" count drops to three. Their next movement, if successful, will requalify them for perturbations. If they fail on subsequent attempts, their success count is decremented each time (to a minimum of zero). This process ensures that at the time the perturbation is applied, the subject has demonstrated that they are "in the groove" — routinely performing the unperturbed task successfully. It does not seriously penalize a single error, but becomes less tolerant of repeated errors.

---

[17] "Not intervening" is defined for the subjects as not consciously changing what they are doing as a result of the perturbation. They are instructed to neither consciously resist nor give in to the perturbation. In particular, they are told not to be concerned about reaching and staying at the target if a perturbation prevents them from doing so.

Prototypical Point-to-Point Task

In a typical experimental session, the subject is seated in front of a two-link robotic arm (or *manipulandum* — see following section for details about the hardware) and strapped in to the chair with a four-point seat-belt system to minimize movement of the shoulder (see Figure 15). At the end of the manipulandum's "forearm", a handle is mounted vertically and the entire robot is adjusted vertically so that the midpoint of the handle is level with the subject's shoulder. The subject's forearm is placed in supination and the wrist is immobilized in a lightweight cuff which essentially prevents flexion/extension and abduction/adduction and strongly inhibits pronation/supination. Subjects grasp the manipulandum's handle and place their arm in the horizontal plane. With the limb in this position, the subject's arm segment lengths are measured; the length of the upper arm is taken to be the distance from the tip of the acromion to the anterior edge of the lateral epicondyle, and the length of the forearm is taken to be the distance from the anterior edge of the lateral epicondyle to the centre of the manipulandum's handle. Two markers are placed just above the level of the handle at X-Y positions of (-0.2m,-0.5m) and (+0.2m,-0.5m) to indicate the starting and ending points for the movement (see Figure 15).[18]

---

[18] For one subject, the marker at (-0.2m,-0.5m) was moved to (-0.1m,-0.5m) when the subject's shoulder was placed in line with the two markers. This was done because in the original position the marker was beyond the subject's reach.

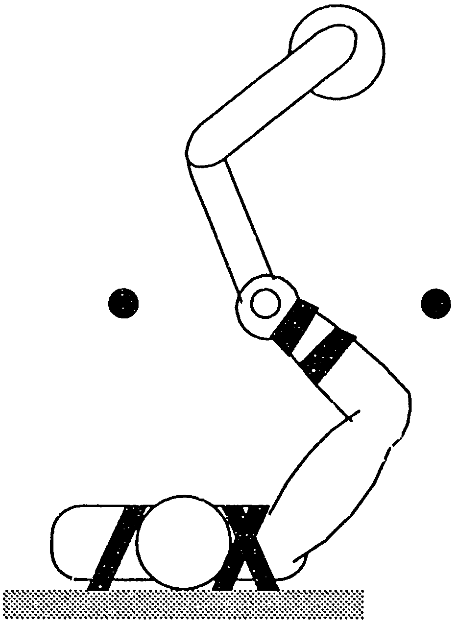

Figure 15. Experimental set-up.

Subjects were tested in four different positions relative to the manipulandum, as shown in Figure 16. In all cases, the movements remained the same in the manipulandum's frame of reference; the subjects were repositioned to produce movements at different locations in

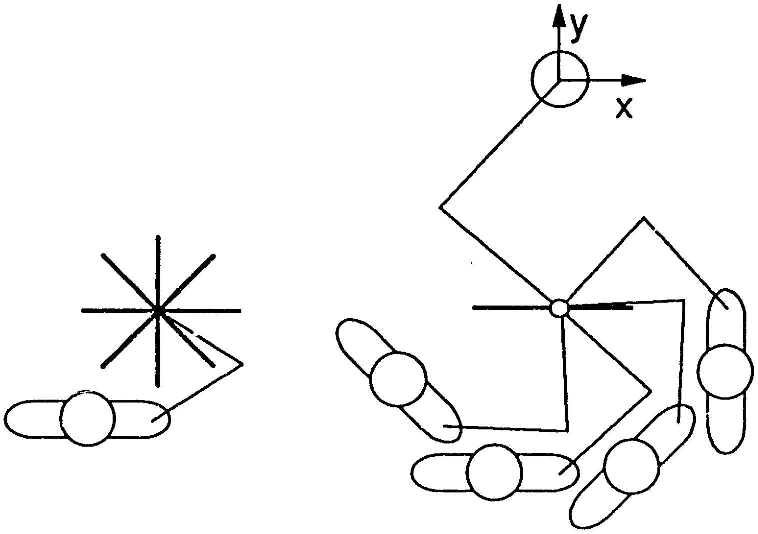

Figure 16. Task orientations. Subjects make movements in four different orientations relative to their trunk, but they are moved relative to the manipulandum so that the inertial behaviour of the manipulandum remains constant in the task-based coordinate frame.

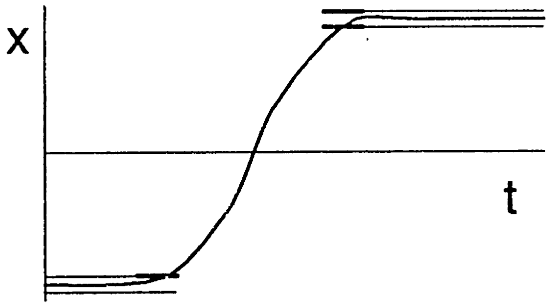

Figure 17. Timing requirements. Subjects must hold at the starting point for a specified time, leave within a short window, arrive within another short window, and hold at the target for approximately one second.

their workspaces. Since the manipulandum does not move relative to the targets, its inertial behaviour remains invariant across the different tasks, so any difference in the executed trajectory will be due solely to the reorientation of the subject.

The distance, duration and onset of the movements are all controlled. As mentioned above, the distance is indicated by the position of two markers placed above the starting and ending points; the computer requires that the movements begin and end within a target zone of radius 1.5 cm about these points (see Figure 17). Audio cues indicate the onset and duration of the movements — a pair of pacing beeps counts down to the start of the movement and a long sustained tone indicates the duration of the movement. The pacing beeps are triggered by the subject moving the manipulandum handle into the starting zone and remaining there for at least one second; by the time of departure, the subject has been at rest for approximately two seconds to satisfy the requirement that the subject start each movement from a state of equilibrium. The interval between the pacing signals is half the nominal duration of the movement. At the end of the movement, the computer issues one or two "report" tones to indicate either success (a single high-pitched tone) or failure (a low tone followed by a high tone if the movement was too fast or a low tone if it was too slow). If the movement was successful, but the subject left the destination zone too early (either by bouncing back or by overshooting slightly), the high tone was followed by a low tone as a warning to pay attention to the task requirement of coming to rest at the target, although the movement was still deemed acceptable (except in difficult dynamic environments, this was a relatively rare occurrence). If the subject left the starting point too early, the long sustained tone was not sounded and they were drawn back to the starting point. Similarly, if they left too late, the sustained tone was prematurely terminated and the subject drawn back to the starting point. The windows about the nominal departure and arrival times were both set to ±14% of the duration of the movement; this value was found to be near the limit of what subjects could consistently achieve, particularly for the faster movements. The centre of the arrival window was adjusted based on the moment of departure — if the subject left at the early end of the departure window, the arrival window was shifted earlier by a similar amount. This adjustment was made to encourage consistent movement durations.

Since the attractor trajectory identification process is iterative, subjects participate in several experimental sessions separated by intervals of days to weeks, depending primarily on the subjects' schedules; most subjects require five or six sessions to reach convergence of the attractor trajectory estimates and each session lasts for approximately ninety minutes. The first session is devoted primarily to training; the limb parameters used to estimate inertias are measured, as are the distances from the manipulandum's motor axes to the subject's shoulder, and the subject is introduced to the requirements of the task. It generally takes a naive subject approximately fifty attempts at a moderate speed unconstrained movement before they are able to perform five successful movements in succession and over one hundred attempts before they feel confident that they can perform the task comfortably. Subjects are encouraged to take as much time as they wish in the training mode (no perturbations applied) before moving on to the experiments themselves. During this training period, I coach them on the various aspects of the task: interpreting the meaning of the tones, figuring out how to leave on time and remain in transit for the specified duration, avoiding overshoot, and remaining at rest at the target for a sufficient period of time. When the experiments begin, subjects are often surprised by the perturbation and react strongly to it. This voluntary reaction is reduced, however, both through coaching and through exercises in which they are asked to consciously respond to the perturbation by returning to the starting point when they feel it; as they become familiar with the perturbation and conscious of what a voluntary response feels like, they are able to suppress such responses and learn to substantially ignore the perturbation. To complete the first session, subjects usually attempt several experimental tasks. Each task (defined as a particular combination of workspace location, movement duration, movement direction, and dynamic environment) is performed until a specified number of perturbations have been applied (generally two to four perturbations per session), at which time another task is performed. In the first session, subjects are usually able to complete only three or four tasks and typically express a great deal of frustration; "I'll never be able to do this," is a common sentiment. However, when they return for the second session, especially if the second session occurs within three or four days of the first, their performance is remarkably improved; they can often complete as many as fifteen tasks in an hour. By the third and fourth sessions, they have generally reached their steady-state ability to perform twenty to twenty-five tasks per hour. Other than the timing constraints within a particular movement, I impose no other timing requirements. Within a particular task, the subject controls the rhythm of the movements by deciding when to enter the starting zone; if they wish to rest, they simply hold the handle outside this zone until they are ready. They are free to request a break at any time and most do so every twenty to thirty minutes or during a particularly frustrating trial. After each session, the data is analyzed and a new estimate of the attractor trajectory for each task attempted is computed; this new estimate forms the basis for the perturbations applied during the execution of that same task in the next experimental session.

Apparatus and Preliminary Data Processing

Manipulandum and Sensors

The robotic arm used in these experiments is a two-link manipulandum driven through a parallel four-bar linkage by two axially-opposed, low inertia, direct drive DC motors. The design of the arm is substantially similar to the device designed and built by Faye [Faye 86]. The manipulandum is equipped with precision potentiometers to measure absolute joint angles and a six-axis force transducer (ATI Model 15/50) mounted at the end of the outer link to measure interaction forces. The force transducer has a range of ±65 N and a precision and accuracy of approximately 0.05 N. The position sensors are precise to approximately 0.001 radians and accurate to 0.0035 radians; the equivalent endpoint precision in Cartesian coordinates is approximately 0.5 mm across the workspace and the accuracy is approximately 2.0 mm. The position signals are low-pass filtered for antialiasing purposes by a fourth-order Butterworth filter with a cutoff at 50 Hz. The signal from the force transducer during a typical task was measured at a frequency of 500 Hz and it was found that >99.9% of the power of the signal was contained in the frequency band below 50 Hz; because of the low noise content of the signal, I chose not to apply an antialiasing filter. Both position and force signals were sampled at 200 Hz during all experiments, and the commands to the motors were updated at this rate also. The motors (PMI JR16M4CH "pancake" motors) were driven by voltage-controlled servo amplifiers. We had a problem with oscillations in the current output from the amplifiers unless the onboard command filters were engaged; the filter characteristics could not be measured directly, but since they eliminated 20 Hz oscillations, we infer that their characteristic response time is on the order of 20-50 ms.

Data Processing

Data processing occurs in three stages. In the first stage, during data acquisition, the angular position signals are converted to Cartesian coordinates. In the second stage (post-processing), the Cartesian position measurements are time-shifted by an amount equal to the delay induced by the antialiasing filters to bring them into temporal alignment with the unfiltered force signals. These signals are then passed through a dynamic programming filter [Dohrmann 88, Hodgson 94] to compute the velocities and accelerations and converted to angular coordinates in a frame centred on the subject's shoulder. The third stage of processing prior to analysis of the data involves estimating the inertial contributions to the force signals, subtracting this contribution from the measured forces to estimate the impedance forces, and using these forces to compute estimates of the pre-inertial attractor trajectories, along with the associated error terms, as described in the earlier analysis section.

Computing Discrepancy and Complexity Measures

Definitions

Once we have estimates for the attractor trajectories, we can test the virtual trajectory hypotheses by computing and analyzing relevant function comparison measures. The purpose of the measures described in the following paragraphs is to quantify the degree of discrepancy between two different trajectories. In general, we are only concerned with the shape of the trajectories, so, because there is intrinsic variability in the timing, duration and distance of individual movements, we normalize all trajectories in both position and time; this has the effect of making all trajectories leave the point (0,0) at time 0 and arrive at the point (1,0) at time 1. We record for further analysis the timing and distance values used to normalize the different realizations of the trajectories. I will describe the normalization process in more detail after defining the discrepancy measure used to compare two trajectories.

Definition: Given a pair of normalized trajectories, $x(t)$ and $y(t)$, we define a discrepancy measure as follows:

$$D(x,y) = 100\sqrt{\int_0^1 \|x(t) - y(t)\|^2 \, dt}.$$

The measure, D, is the normalized RMS deviation between the two trajectories and the factor of 100 is applied to scale typical values into the range of 0 to 10 rather than 0 to 0.1. This definition can be used to compare the shapes of any two trajectories. One statement of the virtual trajectory hypothesis, however, claims that the attractor trajectories are simpler than the control trajectories. How then do we define the concept of "simple"?

If we select a particular trajectory shape, $x_{ref}(t)$, that embodies our notion of simplicity, then we could define the complexity of a given trajectory, $x(t)$, as the discrepancy between it and the reference shape:

$$C(x)\big|_{x_{ref}} = D(x, x_{ref}).$$

What would constitute a simple trajectory? Simplicity implies that the trajectory can be described by the fewest possible number of parameters. Since our task is to start at rest, move to a second point, and remain at rest there, the minimum number of parameters necessary to describe that movement is five: the two coordinates for each of the two target points and the duration of the transition between them. I would argue that, given a Cartesian coordinate frame, the simplest trajectory between two points would be straight and symmetric in its velocity profile; if the trajectory were not straight, then an additional parameter would be needed to specify the degree of curvature; similarly, if it were not symmetric, an additional parameter would be needed to specify the asymmetry. Examples of straight, symmetric trajectories are ramps (which minimize mean squared velocity) or minimum jerk transitions (which minimize mean squared jerk — the third derivative of position). Since the minimum jerk (min-jerk) velocity profile is a good match to those of stereotypical unconstrained planar reaching movements, I will use a min-jerk transition as my reference "simple" trajectory.

Trajectory Normalization

Since subjects begin and end their movements in target zones whose dimensions are on the order of certain features of the movement (such as the lateral deviation, for example), we wish to avoid interpreting slight shifts in the starting position, movement direction or movement distance from one trial to another as changes in strategy. We minimize this possibility by mapping subjects' actual motions onto a reference coordinate system which is (0,0) at their actual starting position and (1,0) at their actual ending position. The starting position is fairly straightforward to compute because they are required to be at rest in the target zone for approximately two seconds prior to beginning their movement. We take the starting position to be the median position occupied in the 100 ms window beginning 300 ms before the nominal departure time. Similarly, since subjects are also required to come to rest at the end of the movement, we take the ending position to be the median position occupied in the 200 ms window beginning 200 ms after the latest permitted arrival time at the target.

Subjects also execute movements at different rates. Again, since we are primarily interested in trajectory shapes, we normalize these trajectories in time. We do this by finding the time shift and scaling parameters which yield the smallest discrepancy measure when compared against a minimum-jerk trajectory which goes from (0,0) to (1,0) in unit time. More precisely, if we have a position-normalized trajectory, $x(t)$, we define a set of variants of $x(t)$ parameterized by a two-vector, p, consisting of an non-dimensional offset time, t_o, and a time scale value, τ. We define the normalized discrepancy of $x(t)$ from a reference trajectory as

$$\tilde{C}(x) = \min_{p \in R^2} C(x(t_0 + \frac{t}{\tau})); \quad p = [t_0 \quad \tau]$$

From these shift and scaling parameters, we compute the nominal duration and midpoint of the trajectory's transition between the target points; these values are used to compare the behaviour of the control and attractor trajectories (see Results chapter).

Examples of Discrepancy Computations

Having defined a measure which we will use to compare trajectories, it is helpful to see what the value of the measure is for easily-characterized test cases so that we can later assess the significance of differences in the discrepancy measure for experimentally-derived trajectories. It is also useful to understand how random changes in the parameters which characterize a trajectory affect the probability distribution of the discrepancy measure. In this section, we will compute the discrepancy measure between two simple trajectories — a ramp and a minimum-jerk transition, show how lateral deviation affects the measure of a nearly perfect minimum-jerk transition, and show the effect of time asymmetry (i.e., non-symmetric velocity profiles).

Ramp and Minimum-Jerk

Consider the class of all functions which describe a straight-line trajectory from one point to another. Two of the simplest members of this class are the ramp and the minimum-jerk transition; both have the same paths, differing only in their time histories. As described above, to compute the discrepancy measure, we normalize the trajectories in position (i.e., so that they start at [0,0] and end at [1,0]) and then determine the time shift and scaling factor which brings them into closest alignment. The result of this alignment process is shown in Figure 18: the top figure shows the time histories of the aligned trajectories and the bottom figure shows the distance between them as a function of time. For these two trajectory shapes, which are quite similar to one another, the discrepancy measure is 1.98 and the duration of the minimum-jerk transition is 1.61 times the duration of the ramp.

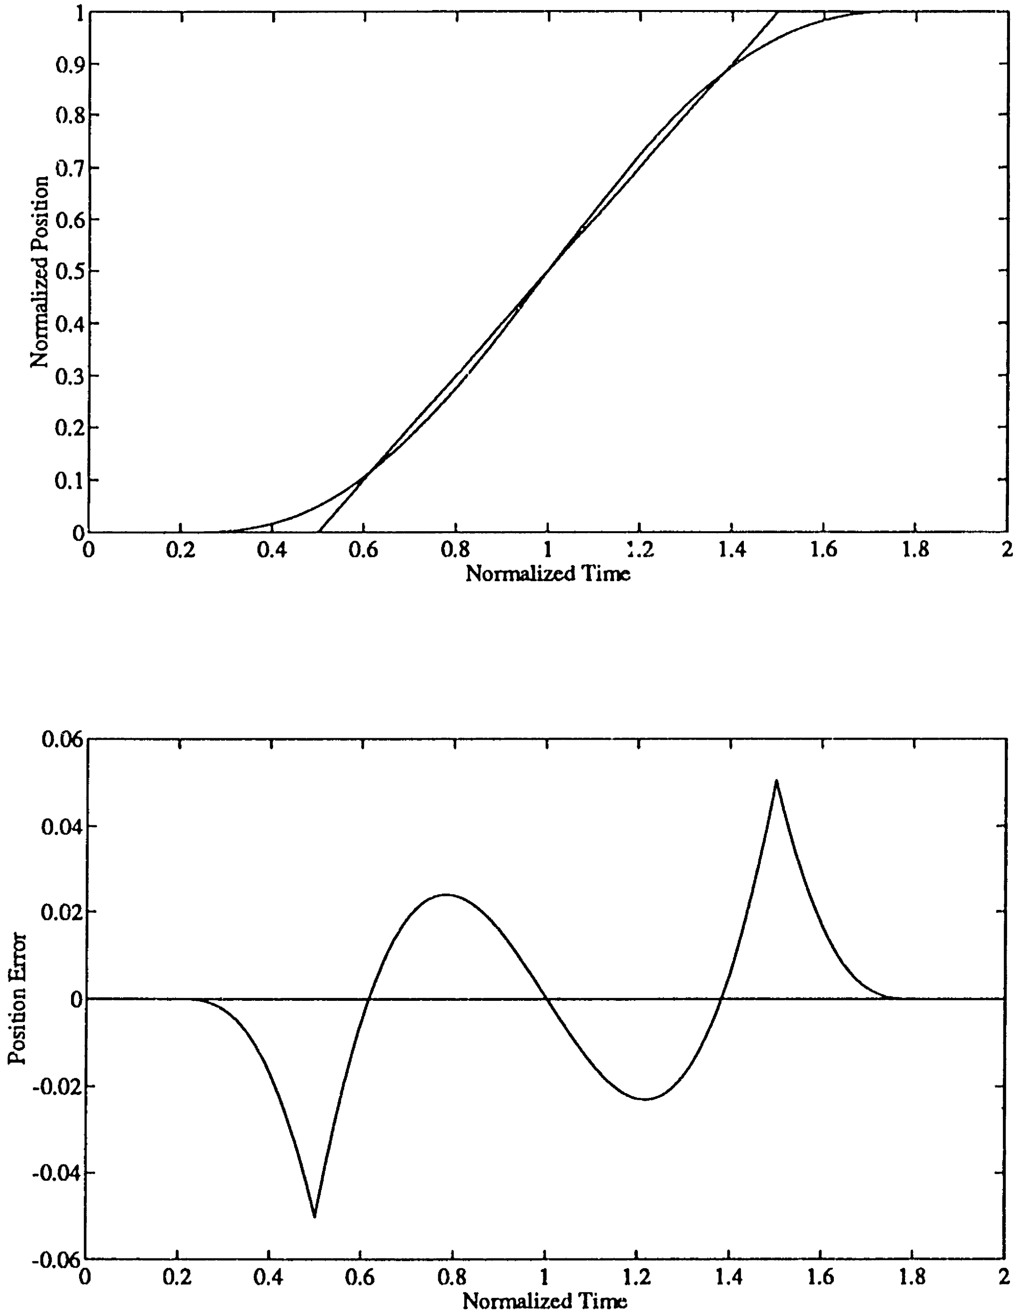

Figure 18. Optimal match (top) and discrepancy (bottom) between ramp and min-jerk trajectories.

Minimum-Jerk with Lateral Deviation

Many of the experimentally-observed trajectories, both control and attractor trajectories, are not perfectly straight but exhibit some degree of curvature. How does this curvature affect the discrepancy measure? Figure 19 shows a set of curved trajectories; all have circular paths and therefore all can be parameterized by a number such as the maximum lateral deviation or the radius of the circle on which they lie. Unlike the case of the comparison between the ramp and the min-jerk transition, here the paths are all different, but the tangential velocity profiles remain identical. However, since the paths of the circular trajectories are different from the path of the min-jerk reference trajectory, no amount of time-scaling or shifting can completely eliminate the discrepancy between them.

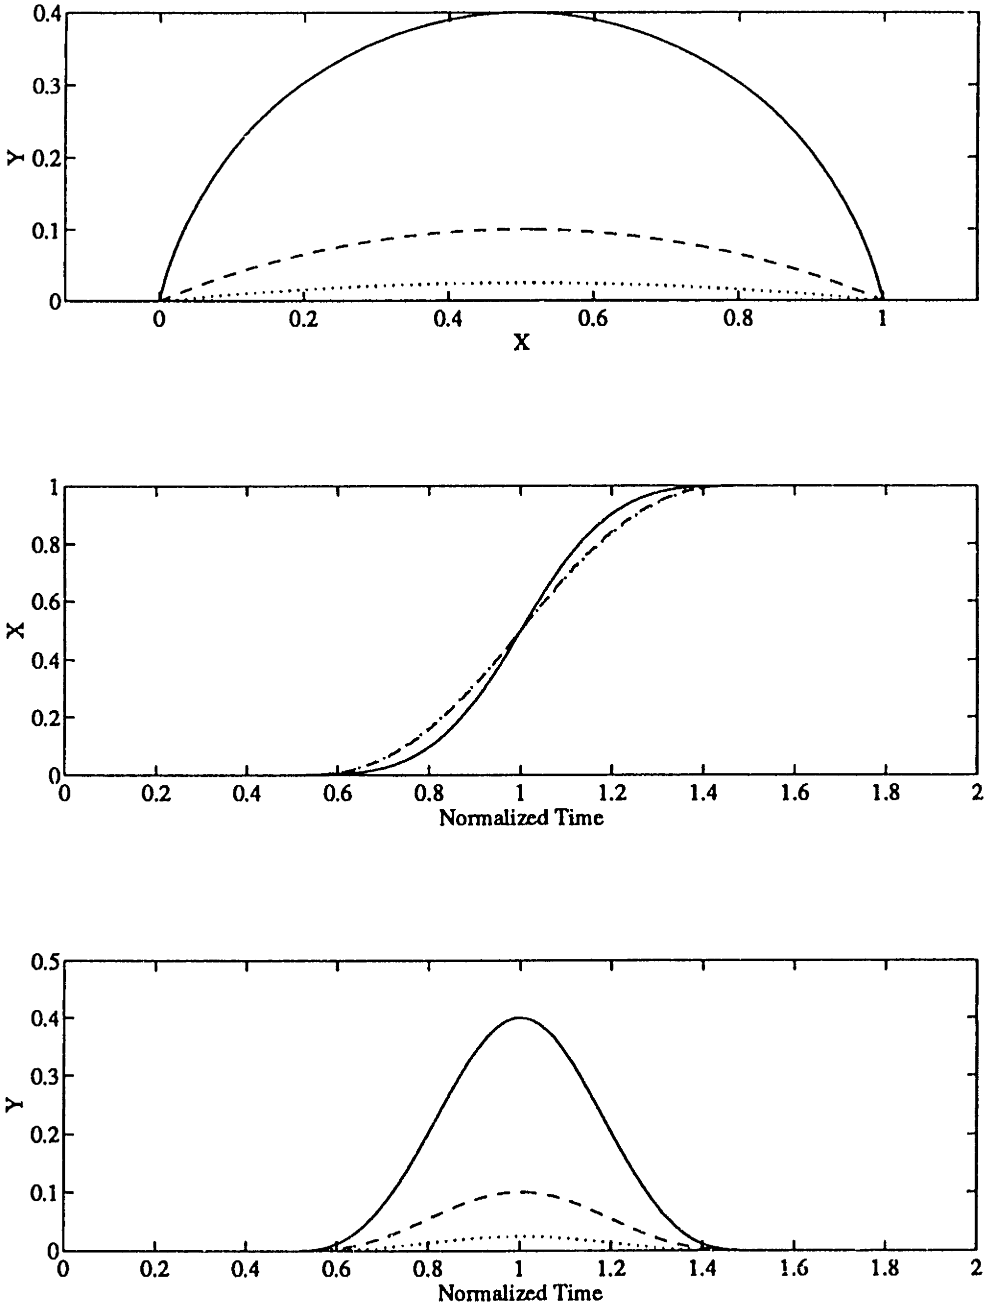

Figure 19. Min-jerk trajectories along a circular path. Top: paths. Middle and bottom: time course.

The plots in Figure 19 show the path and time histories of the circular trajectories at three different lateral deviations — 2.5%, 10% and 40% of the distance between the target points — while Figure 20 shows how features of the fitting trajectories depend on the deviation. The key point is that the complexity measure is nearly linear with lateral deviation out to quite large deviations (on the order of 0.4). For these circular trajectories, the constant is approximately 0.565 units/% deviation; different path shapes which are characterized by a single parameter with units of normalized length will also exhibit a linear relationship between the parameter and the complexity measure for low parameter values, although the constant will differ with the path shape.

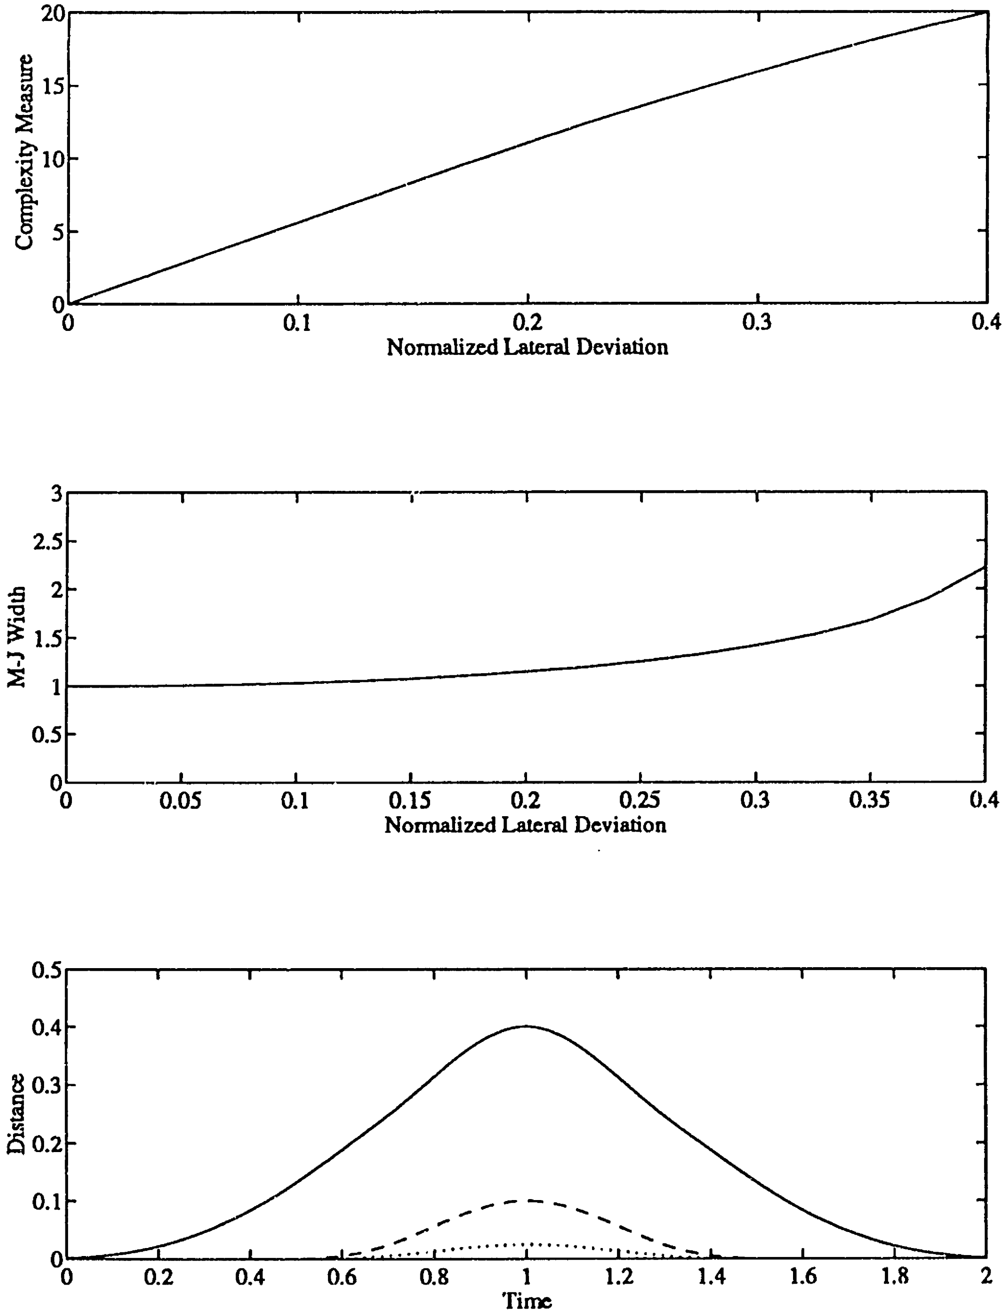

Figure 20. Top: Complexity measure of circular min-jerk trajectory as a function of lateral deviation. Middle: Width of the optimally-matched straight min-jerk trajectory as a function of lateral deviation. Bottom: Distance between circular and best-fit straight min-jerk trajectory.

This relation between the discrepancy measure and the lateral deviation will prove helpful when we assess whether differences in discrepancy measures obtained under different task conditions are significant or not. For example, for each task and subject, we will generally have a dozen or so samples of the unperturbed trajectory from which, given some reference trajectory, we can derive an estimate of the mean discrepancy measure and can estimate its statistical properties. We should be aware, however, that the statistics of the discrepancy measure can depend significantly on the proximity of the reference trajectory to the mean trajectory in the set of trajectories being analyzed. For example, consider a set of trajectories belonging to the family of circular trajectories described above. Let this set of trajectories be characterized by some mean lateral deviation and let each trajectory in the set have some randomly determined lateral deviation drawn from a normally-distributed population centred on the mean for the set as a whole. The probability distribution function of the complexity measure, which is computed by comparing the each trajectory in the set to a straight minimum-jerk reference trajectory, will depend on how close the mean lateral deviation of the set is to zero. The top plot of Figure 21 shows the probability distribution functions for the complexity measure as a function of the mean lateral deviation (normalized by the standard deviation of the lateral deviations of the elements of the set), while the bottom plot shows the mean of these PDF's. Note that because the complexity measure treats deviations in both the "rightward" and "leftward" senses identically, assigning positive values in either case, the mean complexity measure for a set of trajectories which lie precisely along the reference trajectory is not zero; only as the difference between the mean lateral deviation gets large does the mean complexity measure approach the (nearly) linear relation between discrepancy measure and lateral deviation described earlier. In other words, the mean complexity measure of a set of trajectories is higher than the complexity measure of the mean trajectory of the set, but the two approach each other with increasing mean lateral deviation. The correct statistical test, therefore, to test the hypothesis that the mean trajectory of a set of trajectories is the same as some reference trajectory is not to see if the discrepancy measure is zero, but if the probability distribution function is what we would expect if the mean trajectory were the same as the reference.

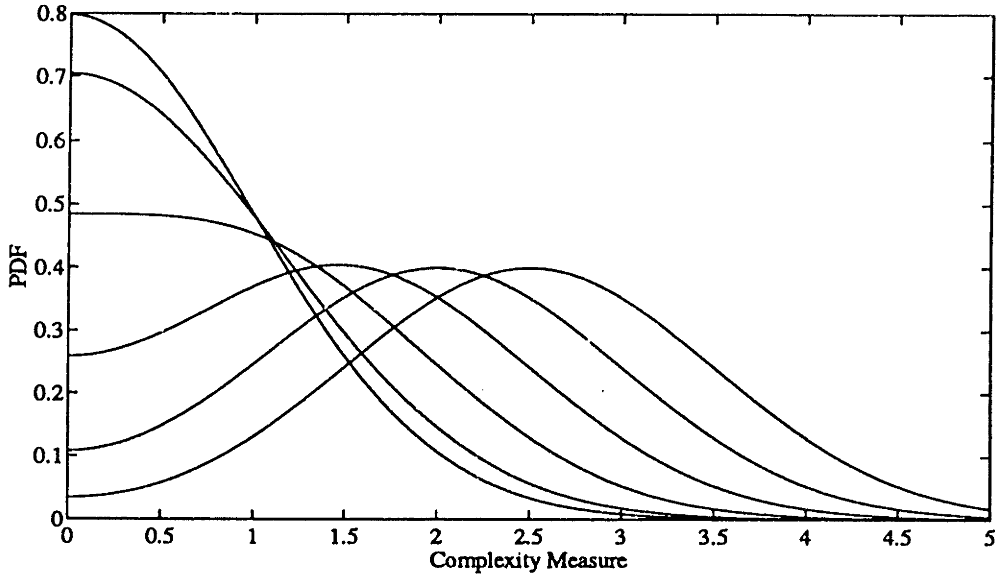

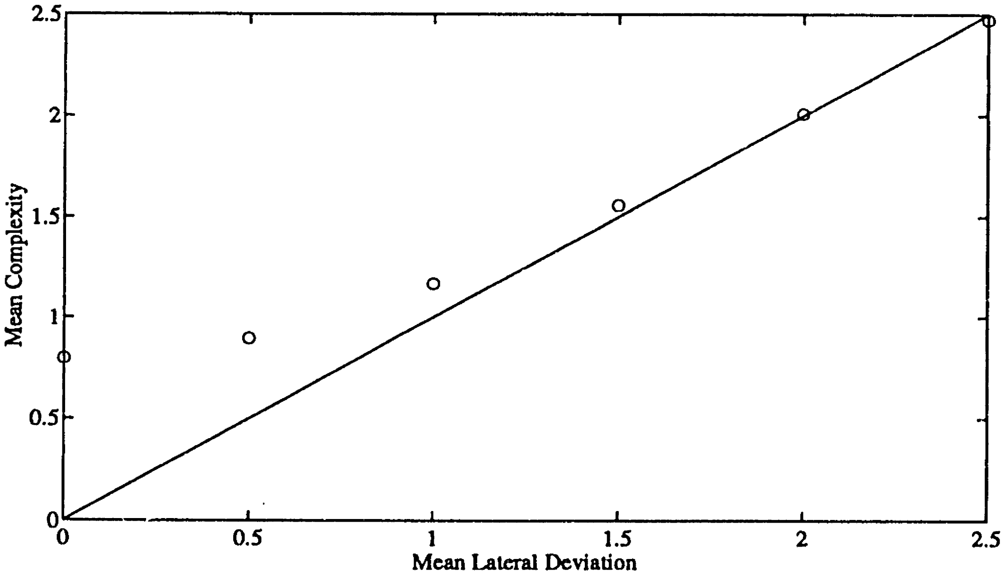

Figure 21. Top: Probability distribution functions for the complexity measure as a function of the mean lateral deviation. Increasing deviations shift the PDF to the right. Bottom: Mean value of the PDF as a function of the mean lateral deviation. Note that because the complexity measure is non-negative, the mean complexity measure for a group of trajectories with zero mean lateral deviation is greater than zero. As the mean deviation increases, the mean complexity approaches the complexity of the mean trajectory.

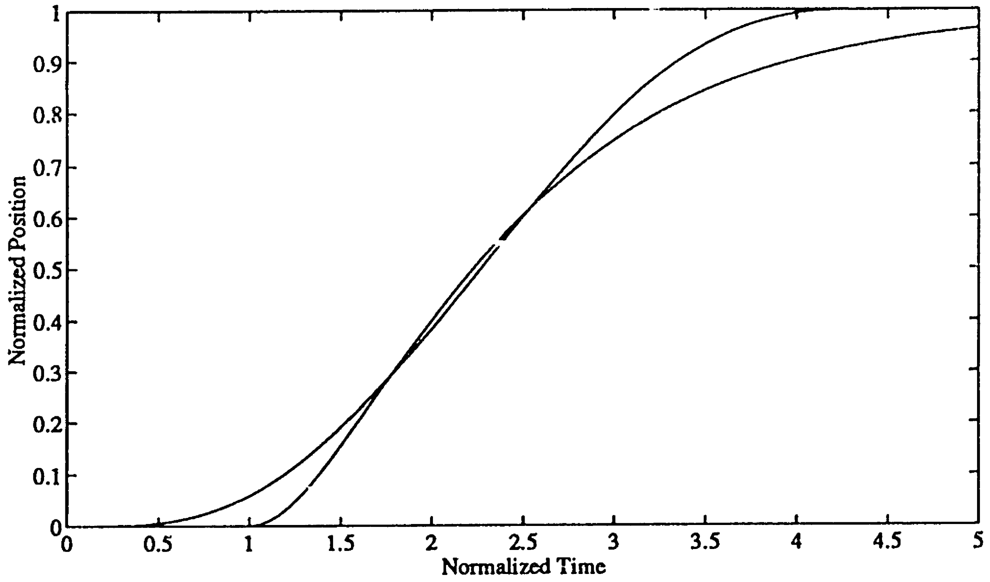

Figure 22. Asymmetrical and optimally-matched minimum-jerk transitions.

Minimum-Jerk vs. Asymmetrical Transition

The experimentally-determined attractor trajectories typically do not exhibit the velocity profile symmetry of the minimum-jerk trajectory we have been using as our reference trajectory for assessing complexity; instead, they generally rise quite rapidly at the beginning and approach the target asymptotically. An analytical approximation which has similar behaviour is shown in Figure 22, along with the best-fitting minimum-jerk transition. This function has the form, $f(t) = (1-exp(-t))^2$, and its complexity measure (relative to the minimum-jerk transition) is 3.53.

Error Bounds

I have outlined above the way in which we compute a complexity measure for a particular trajectory or a discrepancy measure between two trajectories. Experimentally, however, we will never be in the situation of having perfectly known trajectories; there will always be some error associated with their measurements. For example, every time a subject performs the same task, they do it slightly differently. For the most part these trial-to-trial variations are small, but they are not zero. How does that uncertainty get reflected in our complexity measure? As described above in the section on the complexity measure associated with lateral deviations, the mean complexity of a set of trajectories is not the same as the complexity of the mean trajectory of the set. Given a set of trajectories, however, one could compute the complexities of each trajectory considered individually as well as the complexity of the mean trajectory and then approximate the error in the complexity of the mean trajectory by the standard error of the mean of the complexities of all trajectories considered individually. I will present an example of this computation in the results section.

The attractor trajectory complexity measure is more difficult to assess. We have a single trajectory estimate derived from a number of trials, but we also have uncertainty bounds associated with each point on the trajectory. We can therefore compute the minimum and maximum distances of the attractor trajectory from the reference trajectory at each sampling instant and so can place minimum and maximum bounds on the complexity measure. Similar calculations are performed when computing discrepancy measures.

Experiment Design

Subjects

A total of nine subjects participated in the study. There were six men and three women between the ages of 21 and 32 with a median age of 29. Their heights ranged from 5'6" to 6'2". One of the subjects was left-handed, but all used their right arm for the experiment. None of the subjects had any known neurological disorders or disabilities of their arms. Three of the subjects (two men and one woman) were used solely to develop the experimental protocols. The remaining six (four men and two women) performed different sets of experiments to explore a variety of dynamic environments, but all of this group of six performed a core set of trials of unconstrained movements to test in detail two of the virtual trajectory hypotheses and to permit tests of intersubject variability. All but one subject could perform the task in arbitrary orientations; this subject could not perform the 40 cm radially-directed reaching movement and so was asked to perform the same task over a distance of 30 cm.

Tests of Hypotheses

The central claim of the strong correspondence hypothesis is that the virtual trajectories will be of similar complexity to the task description. One of the strongest counterclaims is that this idea cannot hold for medium and high-speed movements where the discrepancy between the virtual position and actual position necessarily becomes large. Kawato, for example, has claimed that for fast movements, the VT bears little resemblance to the relatively smooth unperturbed movement [Kawato 92]. However, if impedance scales in the appropriate manner with increasing speed, the correspondence could be maintained. To assess this claim, we will measure and compare control and attractor trajectories for unconstrained point-to-point reaching movements executed at various speeds ranging from about 30-80% of a maximum (ballistic) effort. To ensure that our results are valid throughout the workspace, we will repeat the measurements in each of four different orientations (as shown earlier in Figure 16) and for movements in both directions between each pair of targets. We can use the same experiments to test the insensitive strategy hypothesis, which claims that comparisons between tasks executed at different speeds or in different locations in the workspace will show that the attractor trajectories are more consistent that the unperturbed trajectories. As we will see in the next chapter, the experiments show that the trajectories executed at different speeds at the same workspace location are relatively similar to one another, be they attractor or actual trajectories. However, the attractor trajectories are significantly more complex and different from one another at different workspace locations than are the actual (control) trajectories.

In addition to the unconstrained movements, we test subjects' behaviour when the movements are made in the presence of two dynamic environments — one position-dependent and the other velocity-dependent. One reason for these studies is that observations of unconstrained movements have revealed that the hand trajectories for such movements are extraordinarily insensitive to the location in the workspace where they are performed and exhibit remarkable symmetry in their velocity profiles. It is therefore difficult to imagine that the attractor trajectories could be demonstrably simpler (based on our min-jerk criterion) than these unconstrained trajectories. A finding that the attractor trajectories are only slightly more complex than the unperturbed trajectories could therefore be criticized on the grounds that the selected task left no room for a clear measurement of a simpler attractor trajectory and that the higher complexity values are simply an artifact of the inherently greater variability of attractor trajectories (see the section on critiques of the methodology below). By presenting a task in which the actual hand path is encouraged to be complex, we are able to check whether or not the attractor trajectory in such situations is actually simpler.

A second and fundamentally more important reason for studies in dynamic environments is that unconstrained tasks represent only part of our motor repertoire. A great deal of our typical motor behaviour involves manipulating our environment, but relatively few studies have addressed the control of such tasks. It is therefore helpful to have available experimental data from a broad range of activities when developing and testing models of the neuromuscular system as the models may represent certain tasks well but fail dramatically when applied to other tasks. Won, for example, found in simulations using a linear model that he could find two sets of impedance parameters which would reproduce the features of either an unconstrained or a constrained response when considered independently, but that he could not find a single set of parameters which would reproduce the response of both tasks [Won 93].

The position-dependent environment can be envisioned as an energy landscape superimposed on the point-to-point reaching movement (see Figure 23); I call these environments "hills" because one experiences either a retarding force ("uphill") or an assistive force ("downhill") depending on the direction in which one is moving. All hills are the same height, in the sense that they require the same total transfer of energy from the subject to the manipulandum,[19] but their position along the line from start to target, as well as their steepness, can be varied. By placing a steep hill near the beginning of a movement, I hope to induce an asymmetry in subjects' hand trajectories such that their movements are retarded as they encounter the hill and accelerated as they "pop over" the top. Similarly, by placing the hill at the end of the movement, I hope to induce a more rapid deceleration phase at the end of the movement. These asymmetries open up the possibility that the attractor trajectory could be simpler than the achieved trajectory. The shorter, steeper hills undergo a transition over 40% of the range of the movement, whereas the longer, gentler hill rises over the full range of the movement.

---

[19] Approximately 2.5 Nm, which corresponds to a maximum force of 10 N for a hill which extends over the entire distance between the target points or 25 N for a short, steep hill extending over 40% of the distance.

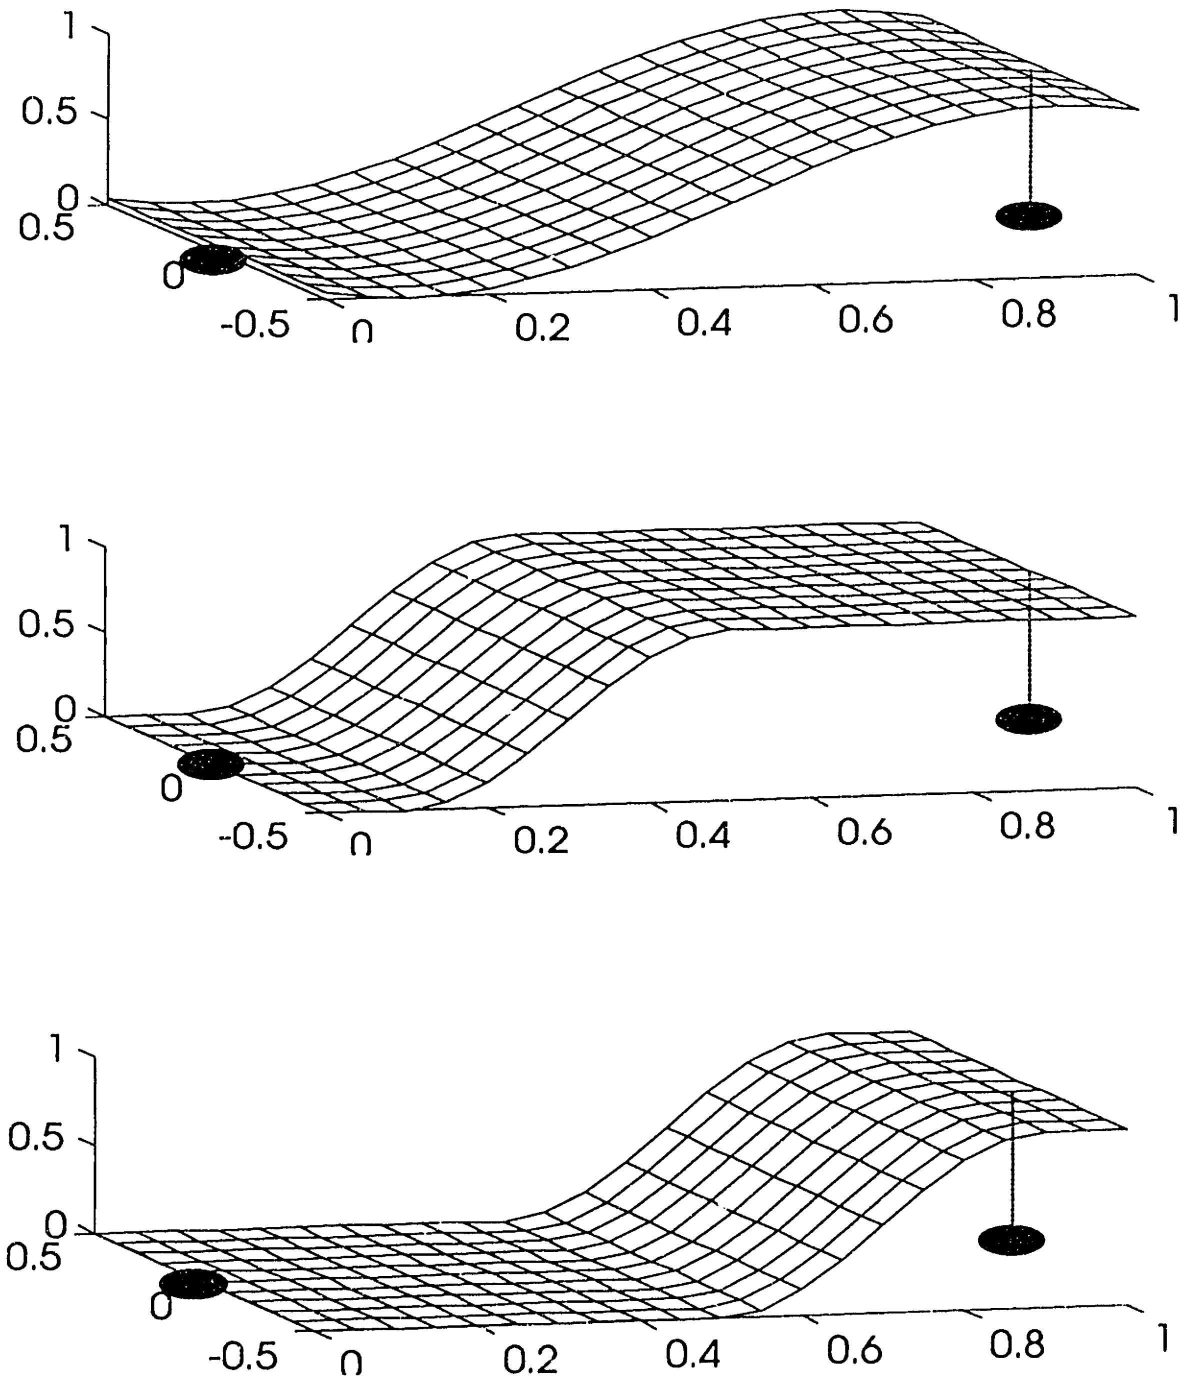

Figure 23. Energy landscape for the hill tasks. Top: A hill spanning the full distance from the starting point to the target point. Middle and bottom: Hills with short transition regions near one or the other of the points.

The velocity-dependent environment was inspired by Lackner's observation that subjects asked to make reaching movements in a spinning room made errors when pointing to targets in the dark [Lackner 94]. This result, if confirmed, would be a strong counterargument to the claim that the central nervous system directly represents the equilibrium location of the limb because the Coriolis forces existing in the spinning room disappear when the limb comes to rest.[20] By programming the manipulandum to exert a force perpendicular to the line connecting the target points in proportion to the speed of movement parallel to that line (see Figure 24), we can emulate the Coriolis effects without adding to them any steady-state centrifugal effects. Such a force field will tend to exaggerate the curvature of an already curved movement if the forces are directed outwards and to diminish it if the forces are directed inwards. If we do succeed in encouraging more complex hand trajectories than in the unconstrained case, we may again open up the possibility of seeing demonstrably simpler attractor trajectories. In the experiments, we use up to three different constants of proportionality ranging from levels which are nearly imperceptible to the subject to those which are clearly apparent, and we apply the forces in both a righthanded and lefthanded sense.

---

[20] The centrifugal forces which do remain when the subject's arm is at rest in the spinning reference frame are insufficient to cause errors of the reported size.

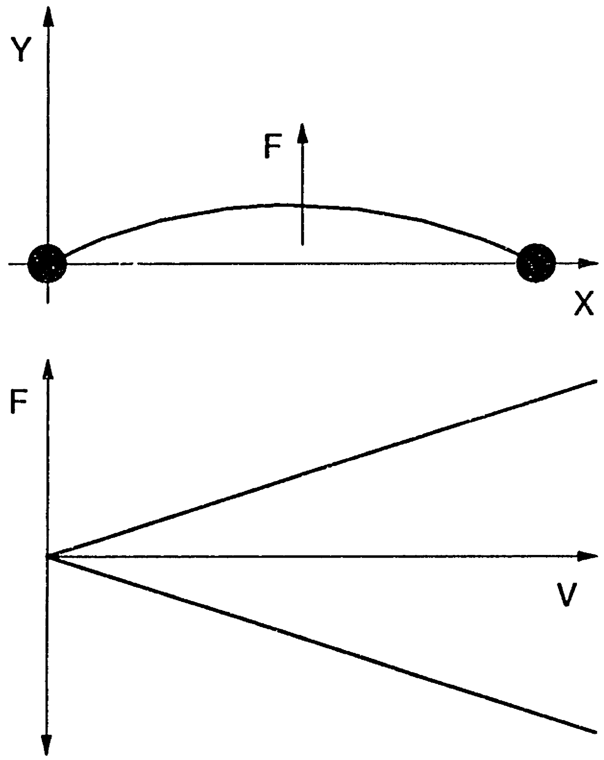

Figure 24. Lateral deviation forces as a function of target-directed velocity. Top: A typical path resulting from application of a velocity-dependent force oriented as shown. Bottom: Sample force-velocity relationships.

Note that neither the hills nor the lateral deviations require the subjects to change their hand trajectories; they are given as much time as they want to learn the task and are free to adopt whichever strategy feels to them most comfortable, and they are not told that straight symmetric trajectories are valued, nor do they receive any feedback on whether or not they achieved such a trajectory; their only feedback is based on their own proprioceptive and kinesthetic cues. This is in contrast to Shadmehr's study in which subjects were required to learn to make essentially straight movements in complex force fields [Shadmehr]. The two dynamic environments here are chosen simply to encourage asymmetries or lateral deviations in the hand trajectories. Three subjects participated in these dynamic environment studies, all of which were carried out in the transverse orientation only (with the exception of some exploratory studies not reported here). All of these subjects did in fact exhibit actual trajectories which deviated considerably from the corresponding trajectories in the unconstrained tasks, although, most interestingly, their attractor trajectories exhibited surprising consistency, in accordance with the hypotheses.

Critiques of Experimental Methodology

Use of Learned Tasks

The methodology described above assumes that the subject executes the same motor program each time a perturbation is applied and requires measurement of a number of perturbations to identify the attractor trajectory. This approach therefore fundamentally requires the subject to have learned a motor task well, so it can be argued that the results do not necessarily apply to motor acts more typical of those generated in daily activities; such acts are generally not performed under the relatively strict timing and positioning requirements imposed on the subjects in these experiments and so may be executed using completely different planning strategies.

While I do not disagree that the experimental task is somewhat contrived, I believe that it has strong connections to more "natural" motor acts. Even though the great majority of our daily motor acts such as reaching for objects are performed in the absence of explicit timing and positioning requirements, these acts must be considered learned actions. Studies of skill acquisition show that in the early stages of developing a new skill, subjects tend to make highly segmented movements and to exhibit high joint stiffnesses [Bullock 91, Vereijken 92]. As the skill is acquired, the movement segments flow together more gracefully and limb stiffness drops. These are the characteristics of our typical daily motor actions. Similarly, although the subjects in these experiments initially experience a great deal of difficulty in performing the tasks, they all reach a point where their execution of the tasks becomes almost subconscious. We are clearly not looking at first-time execution of a task by naive subjects; by the time we actually use data to test our hypotheses, it has come from subjects who have engaged in the activity on half a dozen different occasions and who have developed their own style of execution (as indicated by comments such as, "I've found the sweet spot" or "I'm thinking of making an '*S*'"). There is no indication that they are changing their strategies at any significant rate at the end of several experimental sessions, so the results are not open to the charge that they derive from only partially learned behaviours. Although I have not explicitly measured limb compliance at the beginning and end of the learning process, subjects do report that after the first couple of sessions they are much more comfortable and relaxed than when they began, so it seems plausible that the experimental tasks are executed in a manner similar to those of typical daily motor tasks.

Even if one accepts that both "natural" and experimental tasks are learned, one might still claim that the imposition of timing and positioning requirements makes the two classes of tasks fundamentally different. The experimental tasks admittedly take a great deal of concentration to synchronize one's movements with the externally generated pacing cues, but this is precisely the sort of timing constraint that one encounters in any number of sporting situations where one's opponent initiates the event (returning a serve in tennis, for example). The need for concentration does not imply that a task is neither learned nor natural. Of course, the relevance of a study of such learned movements to an understanding of less precise, possibly more common, human movements executed in daily life without such severe timing and distance constraints is open to debate. However, the discussion in the preceding paragraph concerning the acquisition of skilled movements strongly suggests that skilled actions develop out of and are a refinement of generic motor abilities. Certainly the fact that subjects are able to master a variety of experimental tasks over a time span of a few hours argues that these tasks are in some manner similar to those which they routinely perform.

Implied Constraints

In interpreting the experiments, we propose to compare unperturbed (or control) trajectories with pre-inertial attractor trajectories to decide which are simpler. However, since humans are extraordinarily capable at performing motor tasks, it is possible that the subjects can modify their "natural" behaviour so as to satisfy certain implied constraints. For example, the experimental tasks are point-to-point movements with relatively tight tolerances on the timing and distance of the movements. Although the path connecting the target points is not explicitly specified, subjects may well interpret the task specification as implying a certain trajectory shape such as a path as straight as possible and a symmetric velocity profile. If they do so, then they will presumably be able to change their motor commands so as to ensure that the achieved path matches their interpretation of the task specification. If the unperturbed path then turns out to be simpler than the attractor trajectory, one could charge that that result does not reflect natural behaviour, but instead simply reflects the subject's compliance with the implied constraints. To test this possibility that the subjects are adapting their natural behaviour to produce the implied trajectory, we imposed various dynamic environments which would tend to induce deviations from the unconstrained paths. For example, subjects interact with the simulated hills described earlier; these hills are placed either at the beginning or the end of the movement in order to introduce some longitudinal asymmetry into the task. If the subjects are attempting to generate symmetric achieved paths, they would adapt to both of these environments.[21] Similarly, when they interact with the lateral deviation forces (proportional to longitudinal velocity), these forces tend to displace them to the side. Again, if the subjects experience a strong compulsion to generate paths of a particular shape, they will be able to adapt to this environment. So then, we have two tests of the claim that subjects are imposing more constraints on their motions than the experimental protocol requires; these tests are based on their behaviour in position- and velocity-dependent force fields. As we will see, subjects are generally content to allow these interaction forces to displace them from the nominal trajectory measured in

---

[21] Shadmehr had subjects interact with an unstable velocity-dependent force field. His subjects were asked to make movements between points 10 cm apart in randomly selected directions. Although he did not explicitly require his subjects to make straight movements, after several hundred trials, his subjects typically adapted to the complex dynamic environment he had imposed and generated straight line paths [Shadmehr 1994].

unconstrained movements. Furthermore, the general shape of the subjects' trajectories during unconstrained movements are strikingly similar to those measured in other studies which did not impose the tight timing and accuracy constraints of this study (e.g., *cf* Flash's study [1987]), so there is no evidence that these constraints render the movement particularly unnatural.

Voluntary Response

The experimental paradigm described above relies on the assumption that the subjects maintain (in a mean sense) the same temporal sequence of motor commands during perturbations as they use during the unperturbed movements. Since the perturbations are applied starting at a time 100 ms or more before the visible onset of movement in the unperturbed case, one might reasonably expect that the subject would be aware of the perturbation and able to initiate a voluntary response prior to the completion of the movement. In such a case, what guarantee do we have that the measured behaviour corresponds with the commands the subject would have issued in the absence of a perturbation? There are several arguments one might make.

First, although the voluntary response time to a visual stimulus has often been cited as being on the order of 200 to 300 ms [Jeannerod 88], Won found response times of 500 ms or more for a mechanical perturbation [Won 93]. In his experiment, subjects were asked to make repeated self-paced point-to-point reaching movements in the horizontal plane. On randomly determined trials, one of the joints of the manipulandum held by his subjects would be locked, thereby constraining the endpoint to follow a gentle arc from the starting point to the target point; this arc would generate a maximum lateral deviation from the unconstrained path of 1.8 cm.[22] Won tested the voluntary response time by asking his subjects to return to the starting position as soon as they detected a perturbation. He consistently found that subjects were unable to generate a deviation in the force history in less than 500 ms from the time that the joint was locked. Although my experiment is somewhat different, I can perform a similar test; the results are shown in Figure 25. I find that typical response times are on the order of 300 ms, with one subject able to respond in less than 250 ms. These response times are long enough that, in medium speed movements, the subject is typically 75% of the way to the target before any deviation is observed.

---

[22] The arc had a radius of 23.2 cm and the distance between the target points was 18 cm. The maximum lateral deviation is 1.8 cm.

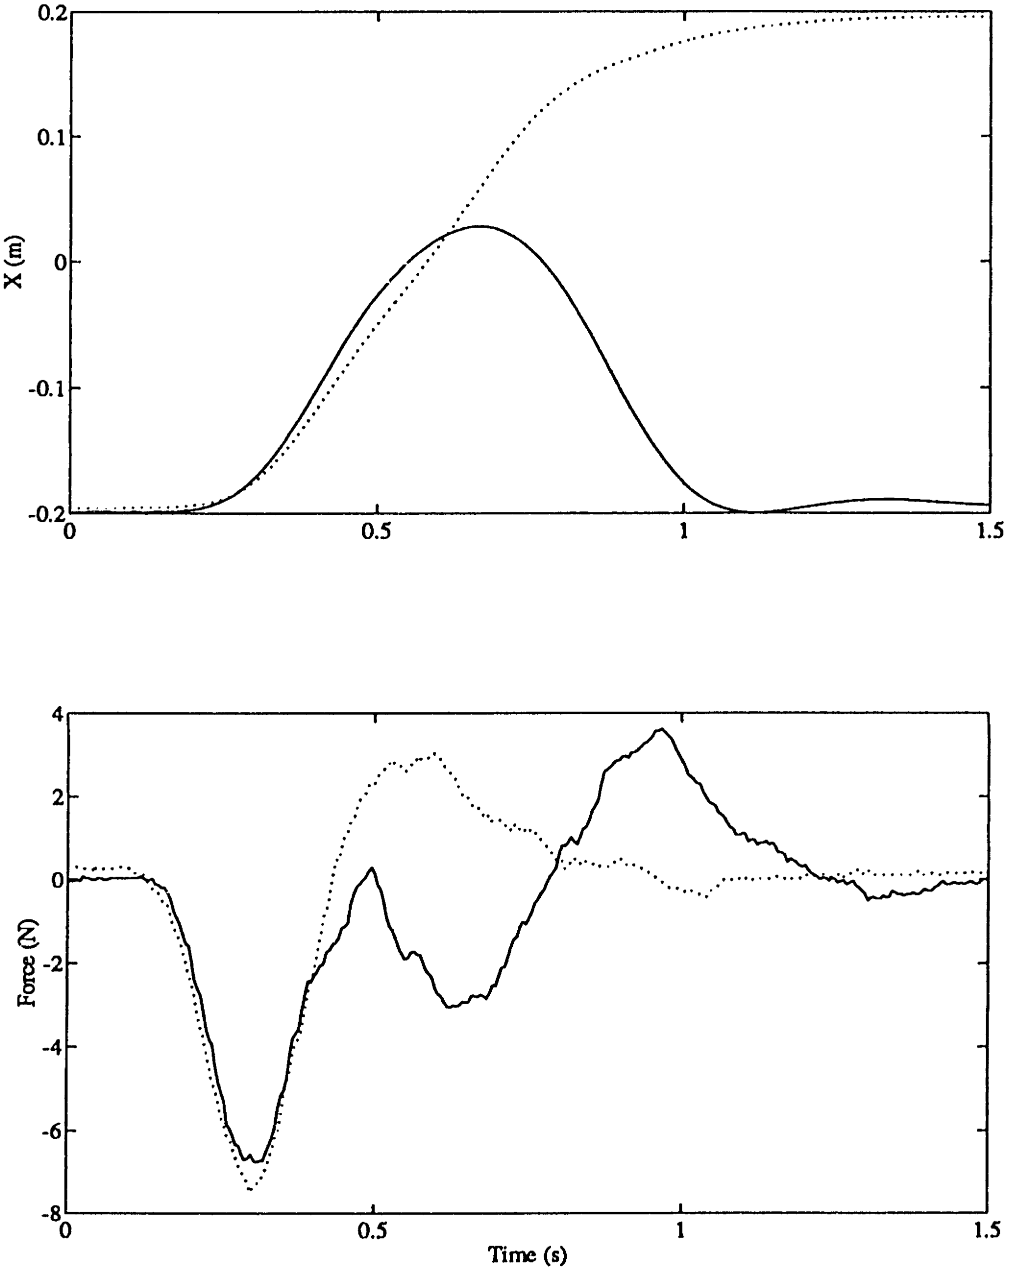

Figure 25. Minimum voluntary response time test. Dotted lines indicate results from normal perturbation test; solid lines indicate subject's response when asked to return to the starting point as soon as the perturbation is felt. Top: Time course of movement along line connecting the starting and target points. Bottom; Forces in the direction of the line connecting the two points. Note evidence of abrupt negatively-directed deviation in voluntary force beginning approximately 350 ms after onset of perturbation.

Even though subjects are able to respond in 400 ms or so, there is evidence that they can suppress such a voluntary response if asked to. This is the so-called "do not intervene" paradigm which has been used successfully by several investigators [Latash 93, Shadmehr 93, Mussa-Ivaldi 85, Won 93]. Latash performed an explicit test of the "do not intervene" paradigm in horizontal elbow movements and found that correlation coefficients between angle and torque remained at or above 0.9 for periods of up to 1000 ms, whereas the correlation coefficients dropped dramatically after 250-300 ms if the subject was instructed to compensate for the perturbations and maintain the desired trajectory or to return to the starting position when they felt a perturbation (analogous to Won's test described above). This is strong evidence that the subject is able to effectively ignore the application of a perturbation during movement, even when the perturbation on a given trial is unexpected.

Finally, since voluntary responses introduce an element of unpredictability into the net response, we would expect to see divergence in the responses to identical perturbations. Indeed, this is what we find in the early cycles of training a subject. However, by the time they are well-trained at the experimental tasks, their responses to the same nominal perturbations are very similar, indicating that they are executing the underlying motor commands without attending to the perturbations.

Visual Feedback

In order to make the experimental tasks as normal as possible, I allow subjects to view both their arms and the targets at all times. To what extent does the availability of visual feedback compromise the assumptions underlying the experiments? One might argue that since the duration of the task exceeds that of the visual response time the unperturbed task is not performed as a single uninterrupted movement, but instead exhibits corrective responses at discrete intervals, as Milner found in a peg insertion task [Milner 90].

A second argument would be that during perturbations, subjects observe the deviation from the intended path and compensate for it. Drawing on the discussion in the previous section concerning voluntary response times, we generally find that the subject's response is extremely consistent, which suggests that they are ignoring the perturbations.

To test our presumption that visual feedback makes little difference to the execution of the movement, we studied a subject performing an unconstrained task with their eyes closed and compared the results with the same subject's "open-eyes" trials. We found some evidence that visual feedback did affect the way the subject performed the task, but that there was no observable effect in the first 400 ms of the movement, which is long enough to allow the subject to cover approximately 75% of the distance between the two reference points. These results are presented in more detail in the next chapter.

Use of Large Perturbations

Since impedance is defined as the relationship between imposed displacements and the resulting forces, most of the existing studies of the impedance of a limb during motion involve the application of perturbations throughout the movement. Such studies have been criticized because the parameters describing the impedance response of muscle depend critically on how the measurement is made; real muscle exhibits complex viscoelastic time-varying behaviour and experiments which attempt to capture this behaviour using a few linearized components are almost always problematic.[23] Especially problematic are those studies which attempt to use linearized impedance estimates to project the behaviour of the limb at some position other than that at which the measurement was made. However, this is precisely what has been done in previous attempts to estimate the virtual trajectory of a moving limb since the location of the equilibrium point is, by definition, not coincident with the actual limb position during motion except for those instants when the actual trajectory crosses the virtual trajectory. Often, the perturbations have been continuously applied [Bennett 92, Xu 91], which tends to excite the reflex gains in a systematic manner and so makes the experiment subject to the critique that it is eliciting non-natural behaviour. Such perturbations have been justified, however, on the grounds that they do not substantially interfere with the subject's ability to perform the task. The perturbations applied in my experiments, in

---

[23] Linearized impedance estimates, for example, are unable to reproduce the so-called "short-range" stiffness of muscle — the markedly higher stiffness of muscle when stretched from rest before the cross-bridges are broken.

contrast, are much larger and by design prevent the subject from following their unperturbed path. The reason for this is, as discussed earlier, that because of the nonlinearities in the subject's impedance behaviour, the only way to be sure that we have found the attractor trajectory is to apply a perturbation large enough to place the subject's limb there. It would therefore be plausible to argue that since large perturbations are required, the subject does not experience a subtle and easily ignored change in the task and so is much more likely to generate a conscious response to the perturbation. My response to this argument is essentially the argument outlined above that the subject is able, with training, to suppress voluntary responses to the large perturbations.

Incorrect Coordinate Frame

The definition of complexity implicitly assumes a particular coordinate frame for the representation of the trajectories. If we are viewing the trajectories in the wrong coordinate frame, our definition of complexity will be incorrect and the conclusions will not be strong. I have selected a Cartesian representation for the trajectories because many earlier studies of reaching movements have demonstrated strong kinematic invariants when trajectories are viewed in this frame [Morasso 81, Abend 82, Flash 87] and wide disparities when viewed in other representations such as joint angles.

Even if we are using an incorrect coordinate frame ("incorrect" in the sense that the chosen frame is not explicitly being used in the representation of the motor plan) and are thus prevented from making strong conclusions about the relative complexity of different trajectories, we will nonetheless be able to make arguments about the relative consistency of the attractor and control trajectories elicited by different dynamic environments applied during a task executed at the same workspace location. This is true because discrepancies between two trajectories connecting the same two points will be apparent regardless of the coordinate frame used to express the trajectories; the only effect of deriving a discrepancy measure in different frames will be that the weighting of the differential contributions of trajectory segments to the overall discrepancy measure will differ, so that using one coordinate frame will emphasize discrepancies in one portion of the trajectory over another.

Definition of Complexity

As noted above, the definition of complexity relies on the selection of a reference trajectory. If one does not agree with the reasoning behind my choice of a straight minimum-jerk trajectory as the reference, one might plausibly claim that my definitions of complexity are flawed. However, particular trajectories are implied by a choice of the smoothness criterion [Poggio 90]; for any given smoothness criterion, a corresponding reference trajectory can be computed. The two trajectories I have evaluated are derived from minimization of time-integrated jerk or velocity (the ramp) and the results in both cases are qualitatively similar. Although one could propose that the optimization is that of acceleration or of snap, there is little reason to suspect that the results would be significantly different — both of these criteria would also result in symmetric straight reference trajectories. I believe the onus is on the critic to propose an alternative definition of smoothness or complexity which results in significantly different reference trajectories.

Attractor Trajectory Necessarily More Variable at Given Location

One might argue that since the preinertial attractor subsystem drives an inertial linkage, the inertia will always act as a low-pass filter on the attractor trajectory thus ensuring that the achieved (control) trajectory will be less variable than the attractor trajectory. We would therefore be bound to find that the insensitive strategy hypothesis is false.

Although this argument about variability may well be valid at a given location, the conclusion is false in general because we are comparing the changes in the mean attractor and control trajectories at different locations in the workspace rather than their variabilities. Regardless of the variability of the attractor trajectory at a particular location, it is still meaningful to ask whether or not the mean trajectory changes significantly with workspace orientation. Furthermore, one should note that low-pass filtering in the time domain does not necessarily smooth the path of the arm; as Flash showed, a perfectly straight and smooth virtual trajectory will turn into a relatively complex path due to the effects of inertial filtering.

Validity of Inertial Parameter Modelling

The subject's arm is modelled as a two degree of freedom system — only the elbow and shoulder are assumed to be free to move. This cannot be strictly true, however, because the arm has at least seven degrees of freedom: three at the shoulder, one at the elbow, rotation of the forearm (supination/pronation), and two at the wrist — not to mention two or three degrees of freedom in positioning the shoulder. To what extent, then, is our assumption of limited degrees of freedom justified?

The wrist cast effectively prevents flexion/extension, abduction/adduction and pronation/supination. Since the manipulandum's handle fits firmly against the cast itself, if there is any wrist flexion despite the snug fit of the cast, the pivot point would be only 2 or 3 inches from the elbow and its effect would probably be indistinguishable from elbow flexion. Such limited wrist flexion may be problematic for the interpretation of small amplitude oscillations since the effective inertia would be primarily that of the segment distal to the last joint (i.e., the wrist in this case), but because the excursions of the hand are a large fraction of the workspace dimension in our experiments, the proximal segments are involved in virtually the entire motion, so small flexion effects would make almost no contribution to the overall force histories.

The shoulder is slightly more problematic. Experimentally, I have restricted movement by strapping the subject's back to the chair and have observed that the shoulder excursion is typically on the order of 1 cm. Dolan demonstrated that the endpoint inertia matrix (in Cartesian coordinates) was fairly easy to characterize but that the inertial parameters of the upper arm were quite uncertain when inferred from the endpoint inertia matrix [Dolan 91]. He concluded that the endpoint inertial behaviour was comparatively insensitive to uncertainties in the parameters of the upper arm, although I have not verified his conclusion through modelling. However, given that the endpoint excursions are so large, the small movements of the shoulder ought not to have much effect on the endpoint behaviour.

Summary

The tests of the virtual trajectory hypotheses require comparing the attractor and actual trajectories which subjects generate during a relatively normal motor task. We selected planar point-to-point reaching movements as our prototypical task because they have been extensively studied in a variety of related investigations. Since perturbations must be applied from time to time and the perturbed tasks compared with similar unperturbed tasks, we had to enforce consistent performance of the tasks. We did so by having subjects synchronize their movements with pacing tones and imposing stringent timing requirements. We also required subjects to consistently perform the task successfully before we applied the perturbations. Subjects typically learn to satisfy these requirements fairly quickly and we were generally able to reach convergence of the attractor trajectory estimates after five or six sessions.

Once we have estimates for the attractor and actual trajectories, we compute complexity measures by comparing the normalized trajectories to a reference simple trajectory - a minimum-jerk transition between the two points - and use these complexity measures to compare the ensemble of data obtained from the large number of tasks executed by multiple subjects in a variety of workspace locations and at different speeds.

We test the two statements of the virtual trajectory hypothesis by determining whether the attractor trajectories are simpler or more consistent than the actual trajectories for both unconstrained movements and movements in more complex dynamic environments, represented here by the hill tasks and lateral deviation tasks. As we will see in the following chapter, the results are mixed; for unconstrained movements, the attractor trajectories are generally more complex, particularly at different workspace locations, but in the dynamic environments, they are generally simpler.

Finally, we considered a number of critiques of the experimental methodology ranging from implied constraints on the task to validity of the "do not intervene" paradigm, as well as questions concerning the accuracy of the inertial estimates and our definition of complexity, and outlined the steps we have taken to minimize the complications from each of these sources. We therefore proceed now to a detailed consideration of the experimental results.

CHAPTER

FIVE

TESTS OF THE VIRTUAL TRAJECTORY HYPOTHESES

The tests of the virtual trajectory hypotheses are based on arguments of the similarity (or lack thereof) of the unperturbed and attractor trajectories to each other, to reference trajectories, and to those of other subjects and of the same subject under different experimental conditions. This section demonstrates how successful the iterative technique is in achieving convergence, responds to some of the major critiques of the technique, and shows the results of the tests carried out to attempt to disprove the virtual trajectory hypotheses.

Convergence of Attractor Trajectory Estimates

The attractor trajectory estimation technique described in Chapter 3 implicitly assumes the existence of an attractor trajectory. If an attractor trajectory does not exist, then the iteration algorithm will not converge. In this section, we consider how to determine whether or not convergence is obtained and present evidence that in virtually all tasks studied, such convergence was obtained. This evidence comes from four related observations: (1) the convergence of attractor trajectory estimates, even when starting from different initial estimates, (2) the decrease in the estimated contribution of impedance forces to the perturbed movement, (3) the progressive tightening of the error bounds, and (4) the convergence of the attractor trajectory estimate onto the perturbed trajectory. After presenting this evidence in detail, we will consider why the error bounds might be considered conservative and what measurements we might make to further tighten these bounds. Finally, we point out conditions under which the iteration scheme can fail to converge due to the choice of impedance estimates (rather than because the attractor trajectory itself does not exist) and suggest ways to avoid this situation.

Progression of Attractor Trajectory Estimates

In this section, we will look at changes in the attractor trajectory estimate with repeated iteration of the estimate/perturb/analyze procedure for one particular task and one particular subject (see Figure 26). While this task and subject has been chosen to clearly demonstrate the process, the results are not at all atypical; in cases which did not exhibit as complete a convergence as is shown here, any uncertainty as to the location of the attractor trajectory is reflected in (generally slightly) increased error bounds. The subject shown in the following figures had an average-sized inertia matrix and was performing an unconstrained movement from left to right in the coronal plane approximately 40 cm in front of their shoulder.

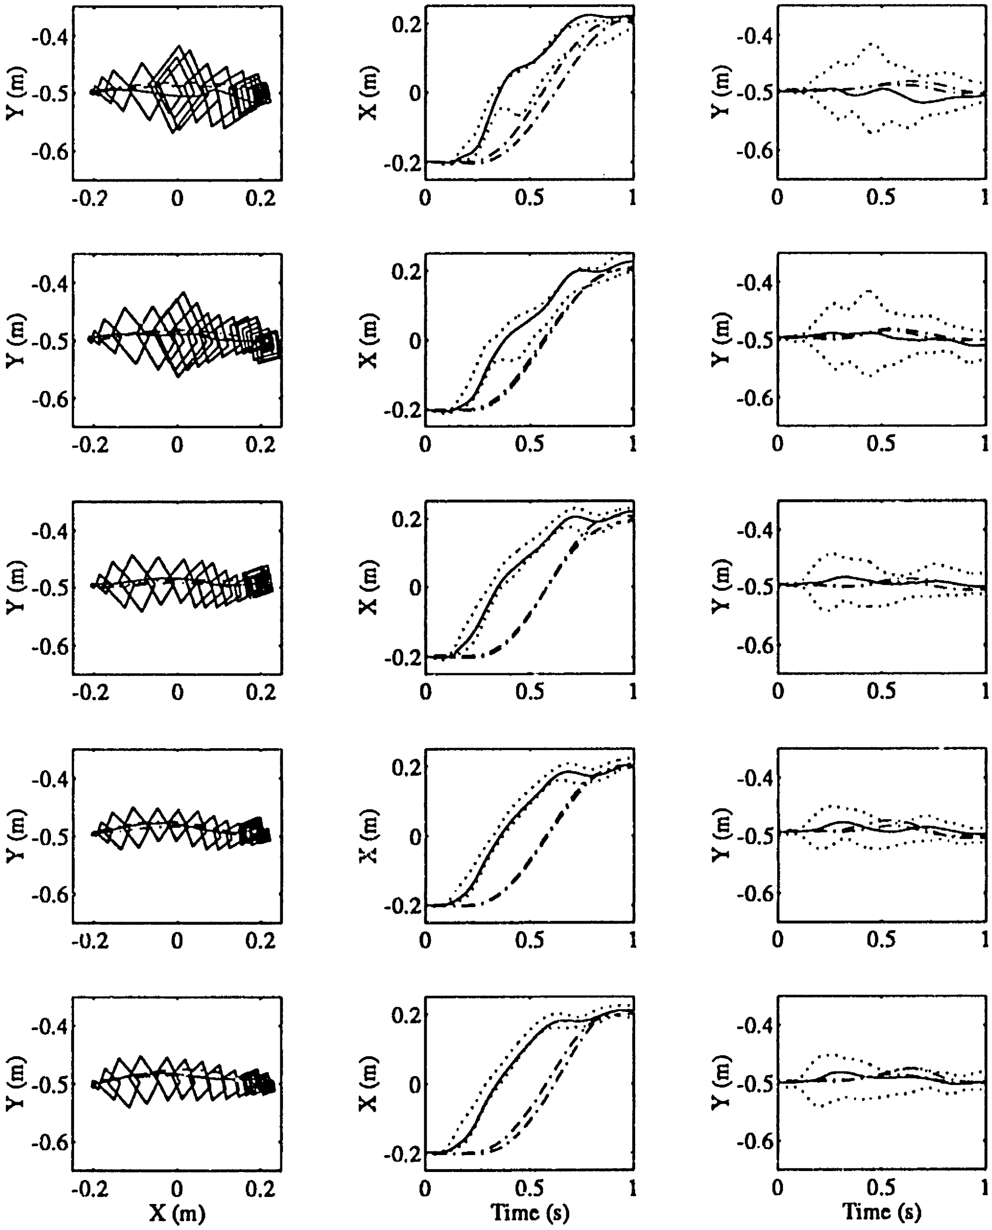

Figure 26. Attractor trajectory estimates at several stages during convergence. The first column shows the paths (solid lines), along with the error bound estimates (squares and dotted lines) and the unperturbed paths (dashed). The second and third columns show the time courses of the X and Y components of the trajectory, respectively. Lower rows correspond to later iterations.

In the first iteration stage, the subject is still learning the task; this is evident from the relatively wide spread in the timing of the transition (compare the time shifts in the *X vs t* plots here with the later stages), and also from the differences in the interaction force histories (not shown here). The initial attractor trajectory estimate is fairly complex, indicating a rapid transition to approximately 60% of the distance between the points, followed by a second slower transition to the target. The *Y* time history also shows evidence of this two-stage behaviour; it has two minima, each corresponding to the midpoint of one of the *X* transitions. These features are in some sense illusory, however, because they are not found in later attractor trajectory estimates. It is unclear whether they represented a real feature of the attractor trajectory which the subject used while learning the task or whether they were simply an artifact of projecting too far from the perturbed trajectory at this stage where the location of the attractor trajectory is as yet undetermined.

After the second iteration stage, the mid-transition "bump" has been significantly reduced and the degree of *Y* movement is similarily reduced. By the third iteration, the bump is only barely in evidence and the path is essentially identical to what it will be after the fifth and final iteration. The results after the fourth and fifth stages are virtually indistinguishable from one another; the bump has disappeared and been replaced by a gentle rolling off of the rate of transition, and the attractor trajectory appears to hold just shy of the target point for approximately 150 ms to allow the hand path to catch up before making a small final movement to the target.[24] We see, then, that even though we started with a complex attractor trajectory estimate, the process of iteration led incrementally to a stable estimate.

To demonstrate that the final attractor trajectory estimate is independent of the starting estimate, a different subject was asked to repeat the same task twice during each experimental session with the perturbations used in each repetition being determined by the attractor trajectory estimate derived from the corresponding perturbation in the previous session. The order of presentation of the perturbations in different sessions was varied randomly to prevent bias due to systematic changes in performance of the task during a single session. The two initial attractor trajectory estimates used for this second subject were (1) a minimum-jerk transition and (2) the final attractor trajectory estimate for another subject performing the same task; the results after one and four rounds of iteration are shown in Figure 27. Both iteration processes have produced nearly identical attractor trajectory estimates, the two differing only in some minor details concerning the final longitudinal shift near the target. The paths are nearly indistinguishable from one another and the features of the lateral movements, though small, are similar.

---

[24] This small final shift in the attractor trajectory ought not to be viewed as a correction based on perceiving an error in execution because it appears in most of the final attractor trajectory estimates, which indicates that it is part of the fundamental strategy of performing the motor task.

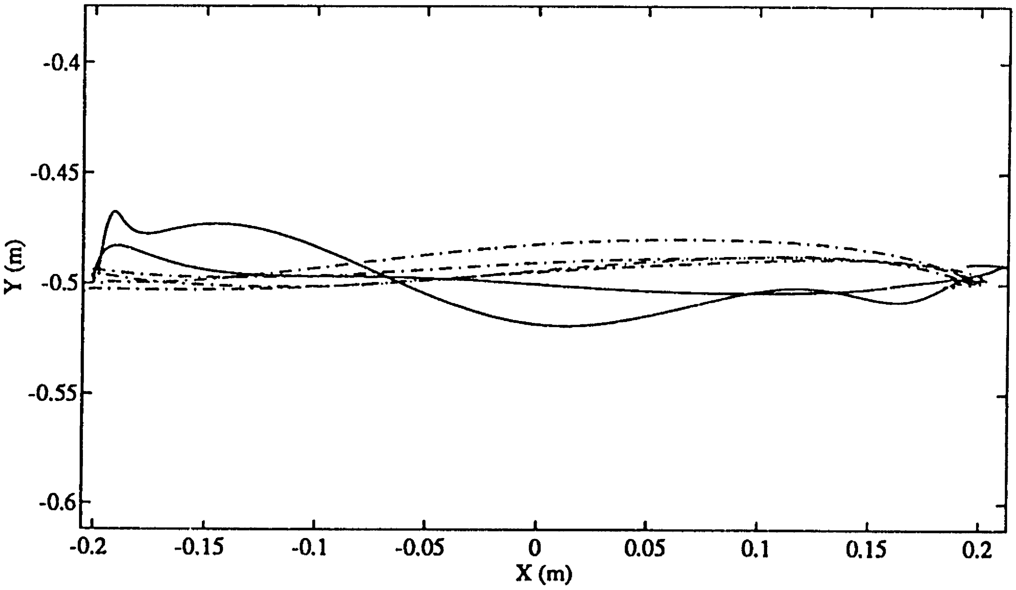

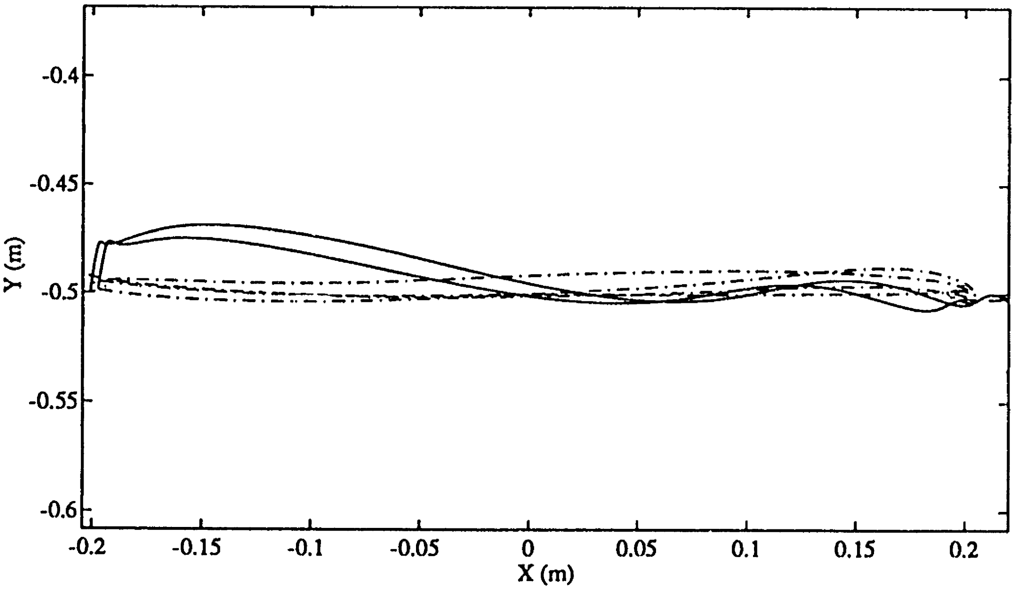

Figure 27. Attractor trajectory estimates (solid lines) after one (top) and four (bottom) rounds of iteration starting from different initial guesses. The corresponding unperturbed trajectories are shown as dashed lines.

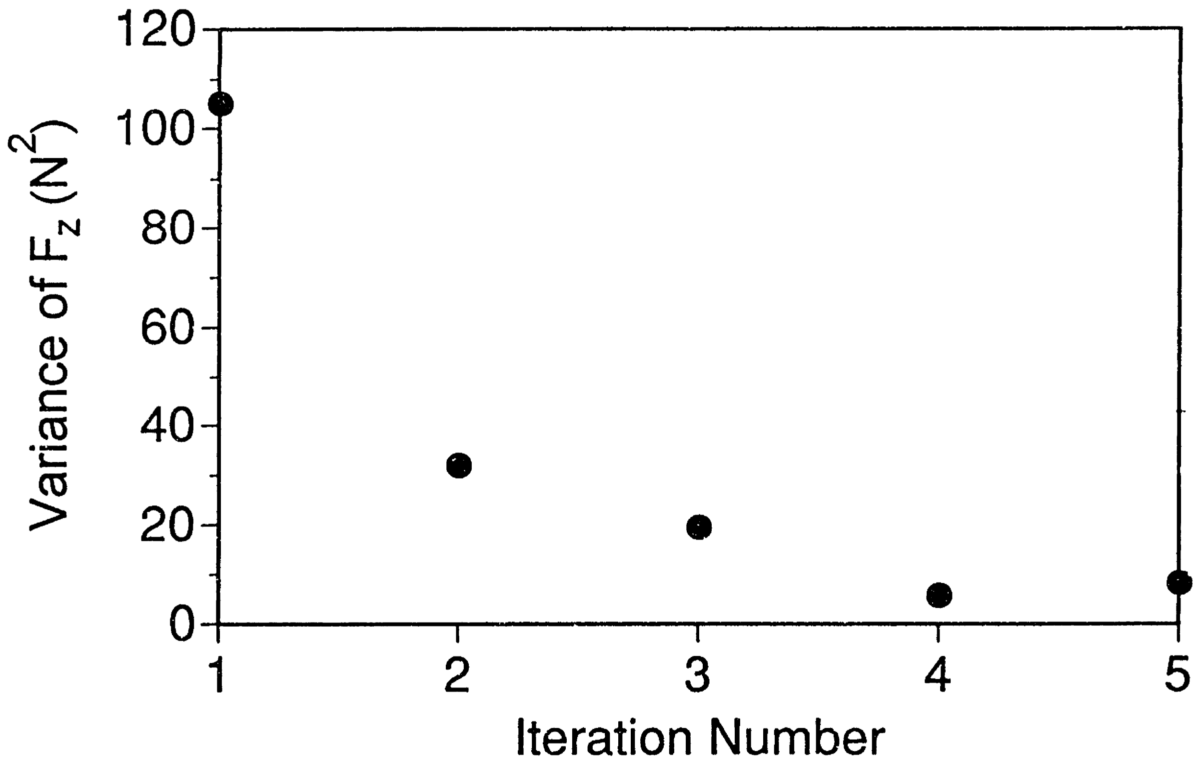

Figure 28. Variance of imp⁻dance force histories as a function of iteration number.

Diminishing of Impedance Force

We return now to the first subject described in the previous section and consider other evidence that the trajectory estimate has converged. Recall that the idea at the heart of the iterative method is that we seek to supply via the manipulandum the forces required to drive the inertia of the arm along the attractor trajectory; if our approximation is improving, then the computed impedance forces (those non-inertial forces supplied by the subject) should be decreasing with each iteration. In the earliest iteration, the peak impedance forces are on the order of 20 N - the magnitude of the inertial forces themselves. After the second and third iterations, these peak forces have been reduced to approximately 10 N, and after the final two iterations, they have been further reduced to

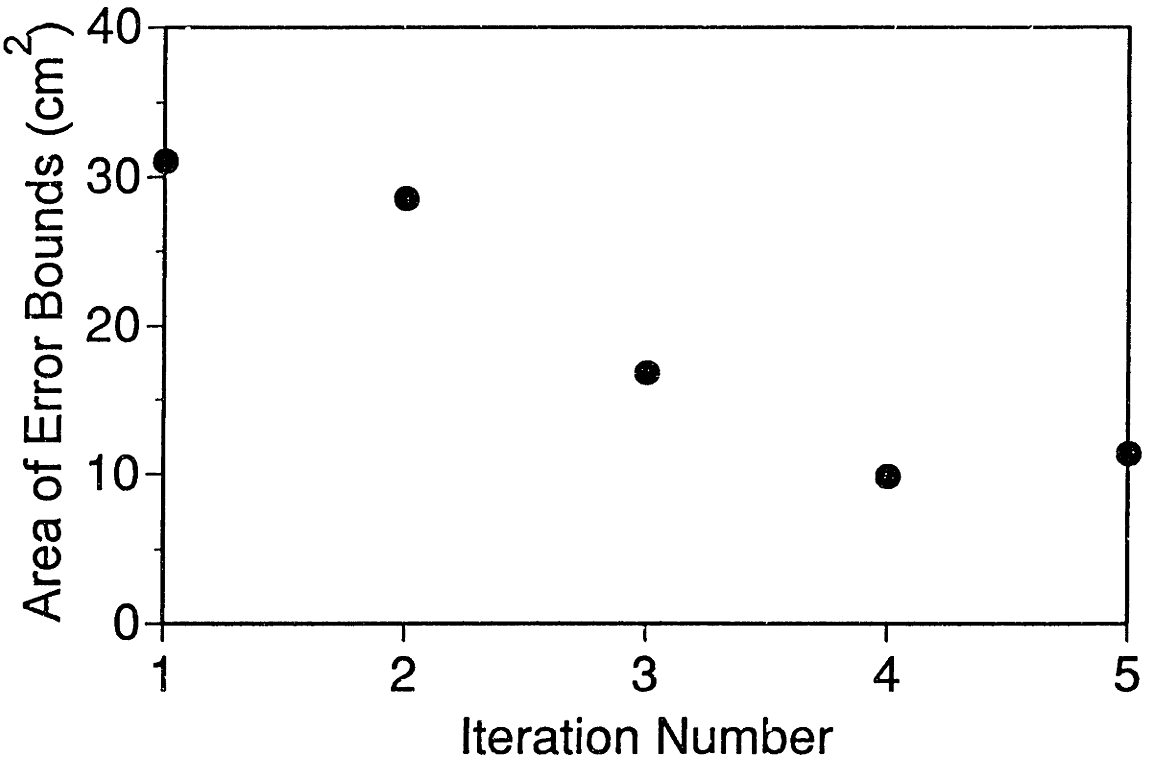

Figure 29. Mean area of uncertainty quadrilaterals as a function of iteration number.

approximately 5 N, or 25% of the peak inertial forces. Figure 28 shows the relative variance of the impedance force history during the transition period as a function of the iteration number. This variance decreases quite dramatically over the five iterations, with final values on the order of <10% of the inertial force variance, which indicates that we have substantially succeeded in supplying the inertial forces via the manipulandum.

Progressive Shrinking of Error Bounds

Since the impedance forces arise because of a discrepancy between the current position and the attractor position, decreased impedance forces ought to correspond to decreased uncertainty as to the location of the attractor trajectory. As described in the earlier section on methods, we use the time history of estimated impedance forces, along with estimates of the subject's stiffness and damping, to place bounds on the possible location of the attractor trajectory. Figure 29 shows the mean area of the uncertainty quadrilaterals as a function of the iteration number and demonstrates that reducing the impedance forces also reduces the error bounds since there is a steady decrease in uncertainty with each iteration until, by the final two iterations, a steady state is reached corresponding to a mean uncertainty of approximately 1.5 cm. To decrease the uncertainty further would require improved estimates of the subject's inertial parameters and impedance behaviour, as well perhaps as improvements in the strength of the manipulandum and of its tracking controller.

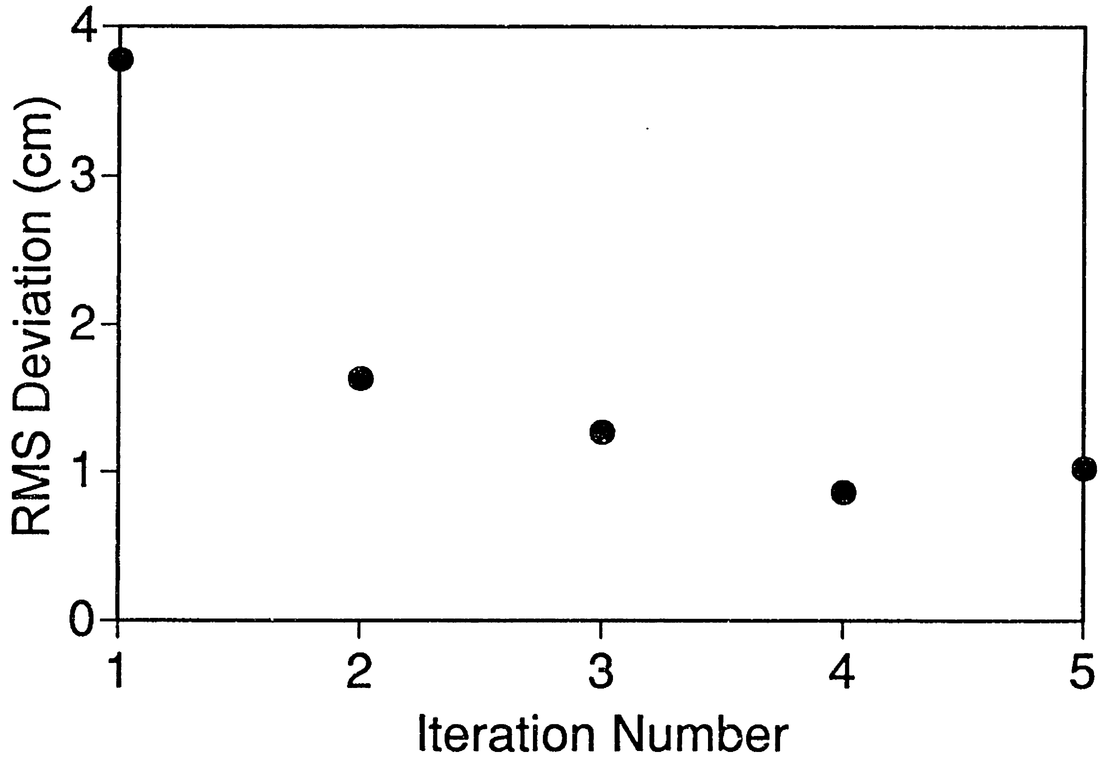

Figure 30. RMS deviation between attractor trajectory and perturbed trajectory as a function of iteration number.

Proximity of Estimated Attractor Trajectory to Perturbed Trajectory

The final measure of convergence which we can invoke is the proximity of the estimated attractor trajectory to the perturbed trajectory. Since the impedance forces are not zero, we use a linearized model of the subject's impedance to improve our estimate of the attractor trajectory by projecting from the (slightly) non-zero impedance forces to the trajectory along which the model predicts they would be zero. If the projection distance is small and our model of the subject's impedance is good, then we can reasonably conclude that we have found the attractor trajectory. We therefore compute the mean deviation between the estimated attractor trajectory and the perturbed trajectory for each stage of the iteration; these results are shown in Figure 30. As with the two previous plots, we again see that the attractor trajectory estimate gets progressively closer to the perturbed trajectory with subsequent iterations. Although this conclusion could be challenged by arguing that we are using an impedance estimate so high that the attractor trajectory estimate could not be anywhere other than immediately adjacent to the perturbed trajectory, this assertion can be dismissed in turn by noting that in the first stage of the iteration, the projection distance is significant - on average, approximately 10% of the distance of the movement. Since the same impedance estimates are used at each stage, they are clearly realistic enough to permit significant projection distances for impedance forces in the range encountered under experimental conditions.

Failure to Converge

Although the examples shown above are representative of many of the tasks studied, some of them (particularly the faster movements of larger subjects) failed to give acceptable attractor trajectory estimates. In this section, we consider the two primary reasons for this failure: manipulandum performance limitations and significant underestimates of the subject's impedance.

Manipulandum Performance Limitations

As the speed of the task is increased, the accelerations required to drive the subject's arm along a given trajectory increase. Since the impetus for these accelerations comes primarily from the manipulandum when the perturbed path is close to the attractor trajectory, the torques required of the manipulandum also increase with task speed. For sufficiently fast motions or for subjects with large inertia matrices, the required torques can exceed the limitations of the torque motors. The only remedy for this problem is more powerful motors.

A second problem is delays on both the sensing and actuation sides of the manipulandum. The main sensor delay is due to the anti-aliasing filters on the position sensors (cutoff at approximately 50 Hz, so characteristic response time is approximately 3 ms), while the main actuation delay is due to the precompensation circuitry of the servo amplifiers. This latter delay is difficult to characterize as we do not have direct access to the internal circuitry of the amplifiers. However, we do know that in order to eliminate oscillations in the output torques, we did have to adjust the precompensation potentiometer. This action introduced an actuation delay of approximately 10-20 ms, judging from the lag in response on the force plots of some of the faster movements. Again, the best remedy for this problem is more modern equipment which does not suffer from this performance problem.

Excessively Low or High Impedance Estimates

In virtually all cases, the iteration process appeared to converge after four to seven rounds. For a small number of subjects and tasks, however, the attractor trajectory estimate appeared to oscillate from one round to the next. For example, consider three sequential attractor trajectory estimates for the largest subject performing a fast movement (Figure 31).

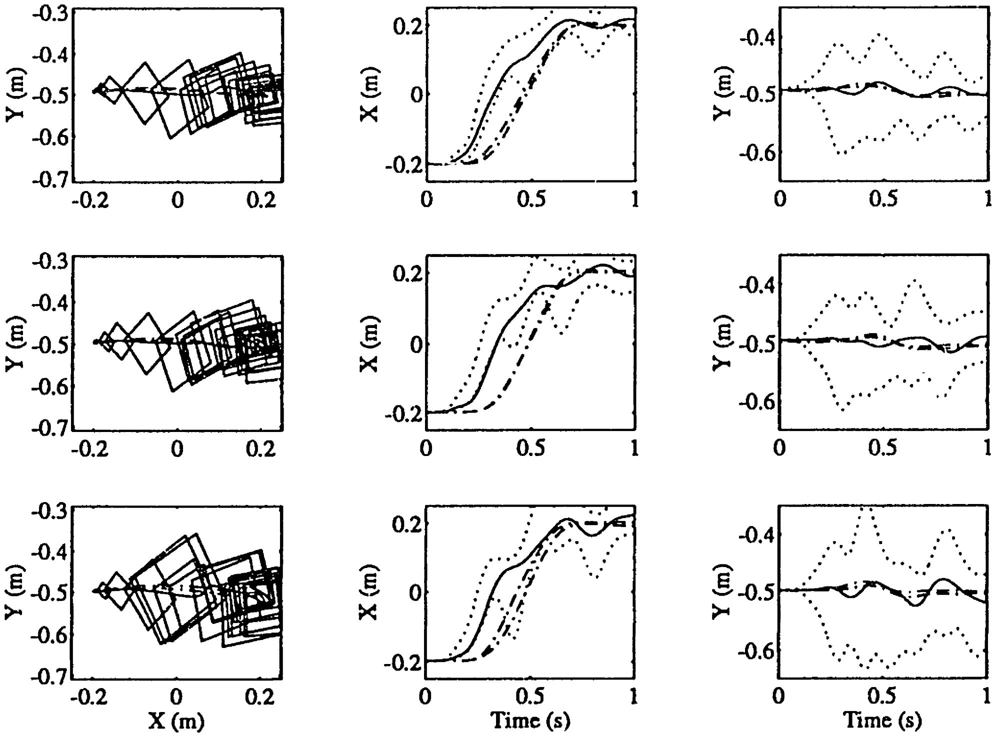

Figure 31. Attractor trajectory estimates for large subject performing fast movement.

Note that at the first stage, there is a noticeable vertical bump at approximately $X = 0.1$ m. At the second stage, this bump has disappeared, but it reappears again at the third stage. Also note the increase in vertical oscillation with successive iterations. Although part of this effect may be due to errors induced by insufficient manipulandum strength, this is much what we would expect if our estimate of the subject's impedance was too low — large impedance forces are attributed to large discrepancies between the current and attractor positions, and if the impedance estimate is sufficiently low, the estimated distance to the attractor may well exceed twice the actual distance. If this happens, then the next iteration will give a result which is even more in error and the cycle will continue.

Fortunately, the remedy is straightforward: choose an impedance value which is appropriate to the subject. There is little difficulty if the impedance estimate is too large; this will simply slow down the convergence process.

Error Bounds

Conservativeness of Error Bounds

The results shown above for two different subjects are typical of all subjects on most tasks — the attractor trajectory estimate converges to a region substantially smaller than that described by the uncertainty quadrilaterals. This suggests that the error bounds themselves are overly conservative. In this section, we consider the reasons why this might be so and what measurements we would need to make to reduce the conservativeness of the error bounds.

The first reason why the error bounds are conservative is that they are computed in differential joint space and, since it was more complex to incorporate the transformation from cartesian space to joint space into the computation of distance, I elected to compute the bounds on differential joint angles independently of one another and then translate these bounds into quadrilaterals in cartesian space.[25] Since it is unlikely that both joint angle extremes would be achieved simultaneously, it would be reasonable to inscribe an

---

[25] Strictly speaking, the joint angle bounds translate into curves in cartesian space, but for sufficiently small excursions, the curves are well approximated by straight lines.

ellipse inside each of the uncertainty quadrilaterals to tighten up the error bounds somewhat.

A more compelling reason why the error bounds are conservative is that their computation assumes that at each instant in the past the impedances and the impedance forces take on extreme values which maximize the possible distance from the estimated attractor trajectory. Furthermore, when the forces switch sign, the impedance values must potentially also make an abrupt transition from one extreme to another. Since it is unlikely that the impedances and impedance forces will take on extreme values throughout the movement and that the impedances will undergo extremely rapid transitions from one extreme to another,[26] none of the extreme deviations is at all likely to occur.

Finally, since the attractor trajectory estimate is independent of the subject's impedance once we have achieved convergence (i.e., there is no projection error), the error bounds are dependent primarily on the interplay between the uncertainties in the inertial forces and the impedances. The inertial parameter uncertainties are likely to have been overestimated, but even if they are uncertain to the extent assumed (+/-25%), the actual inertial parameters remain constant throughout the motion. The uncertainty in the impedance forces (which, at convergence, is identical with the uncertainty in the inertial forces) is therefore more structured than is suggested by the absolute error bounds, which were computed using the assumption that at any instant the impedance forces could take

---

[26] See, for example, the work of Bennett [1992] who shows that for cyclic elbow movements, the impedance changes take place over time intervals of the order of 20% of the cycle's half-period.

on any value within the bounds. Since the inertial forces in any particular trial are not subject to the rapid swings allowed for in the error analysis, the bounds are necessarily conservative. A less conservative error analysis would search over the space of all possible trajectories for which the inertial parameters are constant throughout the movement, but still bounded by the given bounds. Simulations have shown that even the extreme choices of the inertial parameters (+/-25%) tend to have little effect on the estimated attractor trajectory.

Potential for Reducing Bounds

Given the reasons above for conservativeness in the error bounds, can we say what the most effective means would be to tighten the bounds? Consider the size of the inertial forces relative to the total impedance forces for one particular subject and task whose attractor trajectory estimate can be considered to be essentially converged (see Figure 32). The ratio of the RMS impedance force to the RMS inertial force is approximately 30% (lower than this in the peak force regions). Only part of this non-zero impedance force is due to intrinsic variability in the subject's execution of the movement; the bulk of it is due to incomplete convergence, which in turn is primarily due to the manipulandum performance limitations described above. Fortunately, however, indicators such as the proximity to the perturbed trajectory, for example, suggest that this small amount of incomplete convergence will make little difference to our attractor trajectory estimate since it is already likely to be within 1.5 cm of the correct value throughout the movement.

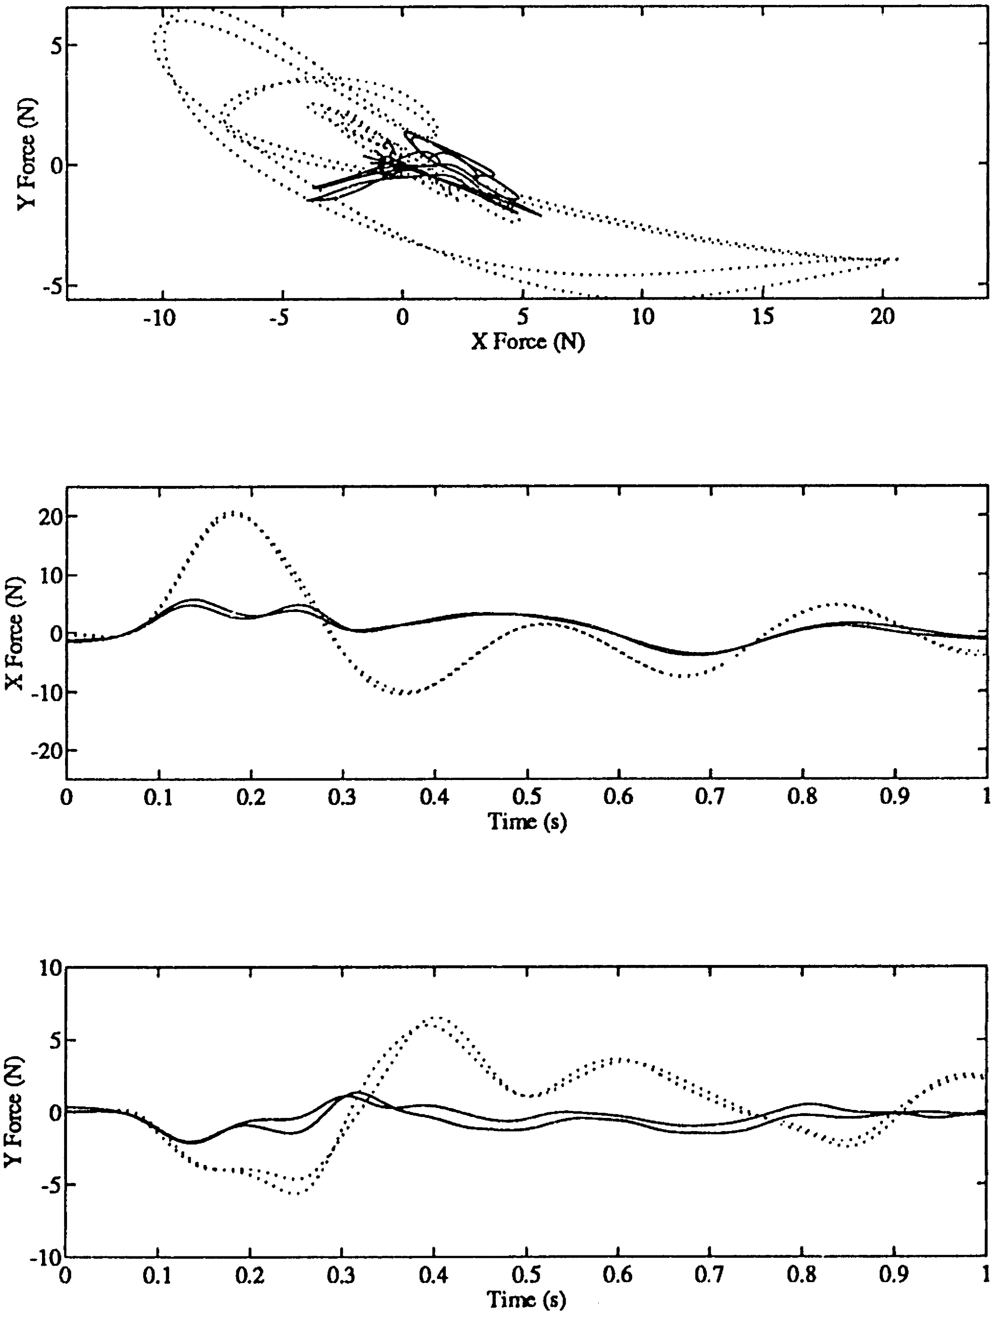

Figure 32. Relative sizes of impedance forces (solid) and inertial forces (dotted) for an attractor trajectory estimate near convergence. Top: Y vs. X forces. Middle and bottom: Time history of forces.

Given that a more realistic estimate of the inertial uncertainty would be on the order of 10% (roughly half of the RMS impedance forces), it would appear that the error bounds might be reduced by perhaps another factor of two through improved manipulandum performance, with any further reduction to come primarily from improved characterization of the subject's (nonlinear) impedance properties during movement.

The final consideration briefly mentioned above, that of the intrinsic variability in the execution of the task, is also worth treating properly, although its contribution to the uncertainty in the attractor trajectory estimate is much smaller than those of the other factors already described. Our current approach to computing the average attractor trajectory involves estimating the attractor trajectory underlying each perturbation, normalizing these estimates in both position and time, and finally averaging the normalized estimates. We use this approach because it limits the shape distortion introduced by the averaging process (see comments in previous chapter regarding averaging ramps with different onset times), but its disadvantage is that since the uncertainties in one measurement are not independent of those from another trial (the subject's impedances are likely to vary in similar ways from one trial to another), we do not reduce the error bounds by averaging; all we reduce is the standard error associated with the intrinsic variability in the attractor trajectory shape. The alternative is to average the impedance forces prior to computing the attractor trajectory. If the sample set is properly distributed around the mean attractor trajectory, then these forces will average out to nearly zero, but the averaging process will distort the shape of the attractor trajectory estimate because we have no basis upon which to perform position normalization or time scaling or shifting. Since the performance limitations of the manipulandum generally prevented us from bracketing the mean attractor trajectory, especially for the faster movements, we would have realized little benefit from averaging the impedance forces and so elected to average the attractor trajectory estimates themselves to minimize shape distortion.

If the inertial uncertainties are small compared with the impedance forces which arise due to incomplete convergence and the intrinsic variability of execution of the task, then we need to concentrate on minimizing the effect of non-zero impedance forces on our attractor trajectory estimate. The main strategy here would be to redesign the apparatus so that it was more capable of tracking a specified trajectory during interaction with a human subject. It would also be helpful to more accurately characterize the impedance histories which subjects use in executing tasks and to use these improved estimates in a more sophisticated error analysis which explicitly recognizes limitations on the rates of change of the impedance components.

Conclusions

In this section, I have demonstrated that the iterative convergence procedure developed in the previous chapter does indeed converge for the great majority of subjects and tasks tested. The attractor trajectory estimates so derived are believable (in the sense that they start at rest, undergo a transition to the target, and end at rest, in contrast to Kawato's "virtual" trajectories which begin and end at locations other than the target points) and, of crucial importance for our tests of the virtual trajectory hypotheses, are independent of knowledge of or a model of the subjects' impedance properties. I have also provided a means for estimating the error bounds on the attractor trajectory estimate, although the actual estimate is likely to be more accurate than is suggested by those bounds because the error estimation procedure does not fully exploit our knowledge of the structure inherent in the physical system we are analyzing. We could improve these estimates both by using better hardware and by making and incorporating explicit impedance measurements; the purpose of this study, however, was to assess the feasibility of the iterative convergence approach given our current knowledge and capabilities and to determine what changes would be needed to make improved measurements in the future.

Intra- and Intersubject Consistency

The purpose of this section is to assess how consistently a single subject performs a task and how great the differences are between the performance of the same task by different subjects. These measures will help us judge the significance of changes in the attractor and control trajectories due to changes in the conditions under which a task is performed (changes in speed or workspace location, for example).

Intrasubject Consistency

To determine whether or not subjects are performing consistently from one experimental session to another, we compare their control trajectories from different sessions and from a number of different tasks to their mean control trajectories for each of the tasks. Since we are primarily interested in the subject's steady-state behaviour (i.e., after they have learned the task), we do not analyze those control trajectories from the earliest iteration cycles; instead, we consider only the trajectories from the last three or four cycles of each task. Using a total of approximately eighty realizations (six subjects, two or three tasks each, four to six control trajectories from each task), we compute the cumulative probability distribution function of the discrepancy measures; this function is plotted in Figure 33, along with an analytical approximation to the function from which we can derive an estimate of the probability distribution function (see Figure 34). Roughly 80% of the control trajectories are within 1.0 units of the mean control trajectory for the corresponding task, and almost 100% are within 2.0 units. The peak of the probability distribution function is at approximately 0.65 units, although its shape is somewhat different from what we would have expected based on the analysis presented in Chapter 4 of the distribution of discrepancy measures; there, we had argued that the PDF would be non-zero at zero discrepancy. However, the model in the methods chapter was that the discrepancy was controlled by a single random variable, the lateral deviation. Here, we probably have at least two means by which a trajectory can differ from one realization to the next (lateral deviation and tangential asymmetry, for example); if this is so, then we would expect just what we have found — a zero PDF at zero discrepancy.

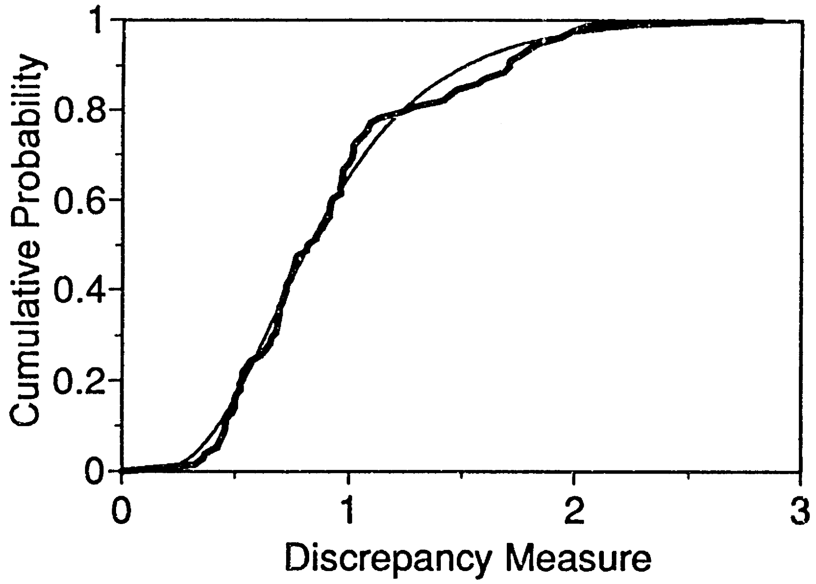

Figure 33. Cumulative probability distribution function for intrasubject control trajectory discrepancy measures.

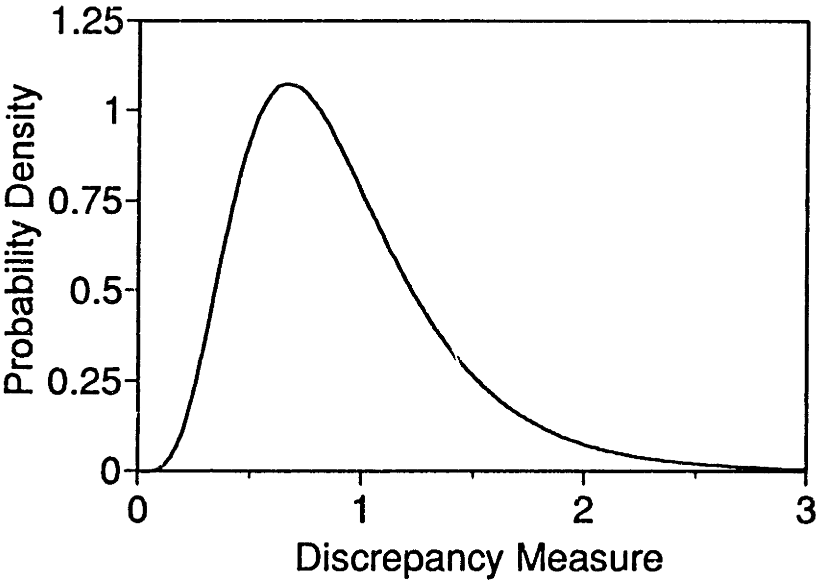

Figure 34. Probability distribution function corresponding to the analytical approximation to the CPDF shown in Figure 33.

In any case, regardless of the details of the PDF, the above analysis gives us a scale for what we might consider significant variations in complexity or discrepancy values — anything less than 2.0 units is relatively insignificant. To put this value in perspective, recall from Chapter 4 that the ratio of the complexity measure for a circular trajectory was approximately 0.56/% lateral deviation; a complexity of 2.0 therefore corresponds to a 3.5% lateral deviation, or ~1.4 cm in a 40 cm movement, or somewhat less than 1.4 cm if coupled with changes in the symmetry of the movement.

Analyzing the control trajectories alone only gives us part of the picture; perhaps the subjects are able to generate an approximately constant control trajectory while at the same time making coordinated changes in their attractor trajectories and impedances as they learn the task. Unfortunately, we are unable to check this possibility because the attractor trajectory identification process is itself iterative, so we cannot be sure what the attractor trajectories are early in the learning cycle; we can only make a defensible estimate after the subject has learned the task. In future, however, we may be able to monitor the subject's progress by assessing the self-discrepancy of the control trials from each cycle of iterations by recording all control trials at each stage, rather than only the one prior to the perturbation as was done in this study. We would expect that early on the subject would be less consistent and would improve with practice; the rate of improvement would be an estimate of the current rate of learning and, if sufficiently small, would allow us to claim that learning had essentially ceased. Furthermore, we could assess changes in impedance by applying a standard perturbation and observing how the response changes from session to session.

The final intrasubject consistency check is the consistency of performance of the task when performed blind and when visual feedback is available. We performed one study with a well-trained subject who was able to perform the task with their eyes closed, although the subject was permitted to open their eyes to recenter the manipulandum handle on the starting point provided they closed their eyes prior to the onset of movement. In all other respects, the two tasks were identical. A comparison of the paths and attractor trajectory estimates measured under the two conditions is shown in Figure 35. In the absence of visual targeting, there appears to be evidence of more scatter in the control trajectories, although it is difficult to confirm this because we have only recorded two control trajectories. The paths of the attractor trajectories look quite similar, with the "blind" trajectory apparently rotated slightly clockwise. The most significant difference, however, is revealed in the *X vs t* plot. Here we see that the normal, eyes-open attractor trajectory reaches a plateau just shy of the final target point and holds there for 100 ms or so before making a final transition to the target point. In the eyes-closed condition, there is no evidence of such a plateau, although there is some discrepancy between the perturbed trajectories and the attractor trajectory estimate because we did not iterate until convergence in this condition. Nonetheless, because we used the same attractor trajectory estimate for the perturbations applied in both conditions, the different responses do signify some contribution of visually-derived feedback to the latter stages of execution of the movement. This divergence is not apparent, however, until at least 450 ms after the onset of the attractor trajectory transition, at which point the attractor trajectory has nearly completed its transition and the hand is approximately halfway between the starting and target points. It is therefore likely that vision is used in normal execution of these point-to-point reaching movements, although the question remains open as to whether or not

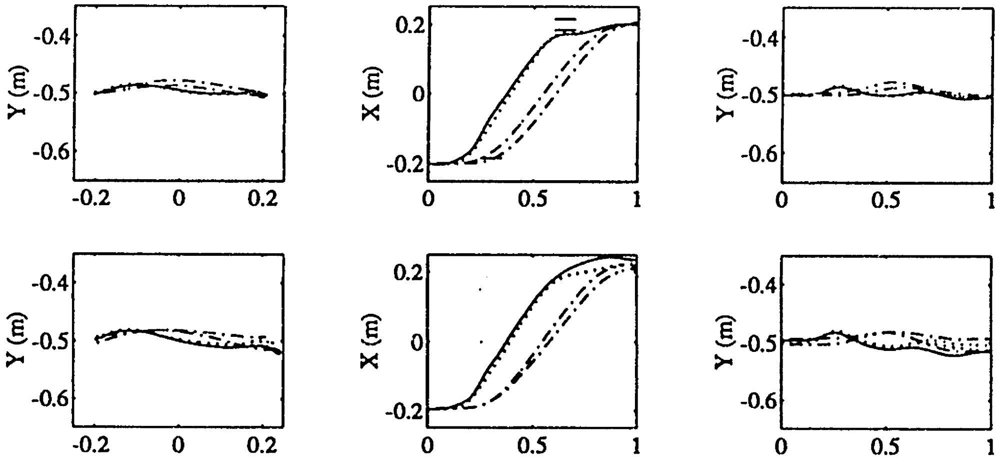

Figure 35. Comparison of unconstrained movements performed in the transverse orientation by the same subject with eyes open (top row) and eyes closed (bottom row). The first column shows the paths, while the second and third columns show the position histories in the X and Y directions, respectively.

sufficient information is available to actually generate corrections to the movement; it may be that visual perception of the early phases of the movement simply acts as a trigger for the timing of the plateau phase of the movement. The consistency of the perturbation response in the eyes-open condition certainly argues that the plateau phase is a planned part of the movement; were it a corrective movement, we would expect to see more divergence in the later portions of the perturbation response.

Intersubject Consistency

In interpreting the results of these experiments, we would like to know whether or not the individual subjects are performing the tasks in much the same manner as one another; this will help to give us a sense of how applicable the conclusions we draw from our moderatesized subject pool would be to a broader population. We do this by computing discrepancy measures for both the control and attractor trajectories for pairs of subjects performing the same task in the same orientation and at the same speed. The results of forty such comparisons for unconstrained movements are shown in Figure 36. One of the subjects tended to make control movements which were considerably more curved than those of the other five subjects; this is confirmed by a two-dimensional Kolmogorov-Smirnov[27] test (see Table 4) which shows that the data from subject B is different from the others with a significance level of 0.001.

Subject	D Statistic	Probability D > Observed	Effective # of Data Points
A	0.40	4.6%	12.4
B	0.56	0.1%	13.0
C	0.31	19%	13.0
D	0.25	47%	12.4
E	0.25	42%	13.0
F	0.35	32%	7.3

Table 4. 2-D Kolmogorov-Smirnov statistics for intersubject comparisons of the complexity measures.

---

[27] The Kolmogorov-Smirnov test is a non-parametric test to determine whether or not two distributions of points in a 2-D space are the same or not. It is based on the expectation of the relative fractions of points from both distributions lying in each of the four natural quadrants about each data point; if there is a large fraction from one set of points and a small fraction from another in one of these quadrants, then the two distributions are deemed to be different. The test is accurate when the effective number of points is ~ 20 or greater and the probability is less than 20% [Press 1992].

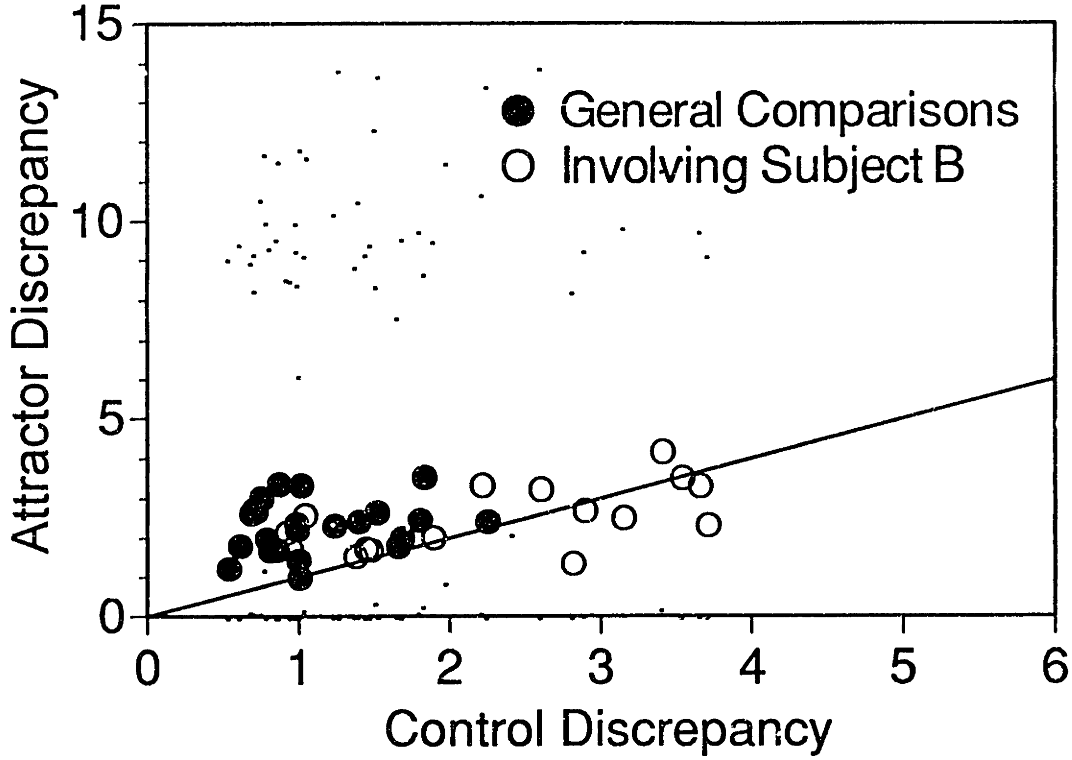

Figure 36. Intersubject discrepancy measures for different subjects performing the same unconstrained tasks. The small dots above and below each large marker indicate the maximum and minimum attractor discrepancies given the error bounds associated with each particular task. The solid line indicates where the attractor discrepancy would equal the control discrepancy.

If we exclude all comparisons with subject B (whose data is indicated by open circles in Figure 36), then we find that the control discrepancies differ by less than 2.3 units between subjects. This value is comparable to the intrasubject discrepancies — Figure 37 shows the cumulative distribution function for these intersubject comparisons, along with the two analytical functions which give the best matches to the experimentally-determined CDF's for the intra- and intersubject comparisons. The analytical function which best matches the intersubject CDF is simply scaled in the X direction (along the discrepancy measure axis) by a factor of 1.2;[28] that is, the distribution of discrepancies between the mean control trajectories of different subjects on the same range of tasks is only (on average) 1.2X larger than the discrepancies between the control trajectories for one subject repeating the same task. We can reasonably conclude, then, that for the most part different subjects are executing trajectories which look very similar to one another (with the obvious exception of the one subject mentioned).

The attractor trajectory measurements are subject to somewhat greater measurement errors, so we are not surprised to find that the attractor discrepancies are generally slightly larger than the control discrepancies, though in almost all cases they are smaller than 3.5 units. We cannot say for certain how much of this value is due to measurement error and how much to actual differences between subjects, but it seems reasonable to conclude that there is little evidence of dramatic differences between the attractor trajectories of different subjects performing the same task.

We check this conclusion by analyzing the intersubject consistency for the three subjects who performed tasks in the two dynamic environments studied in these experiments; these results are shown in Figure 38. From this figure, it is apparent that different subjects respond to the environments in different ways, so we will have to be somewhat careful in drawing conclusions from our relatively small pool of subjects. However, by analyzing

---

[28] With this value of scaling along the X axis, the 1-D Kolmogorov-Smirnov test for similarity of distributions of control complexities gives a D statistic of 0.14, a probability of 76% that D should be greater than this if the distributions are in fact identical, and an effective number of points of 21.8. I.e., this analytical distribution is a good match to the data.

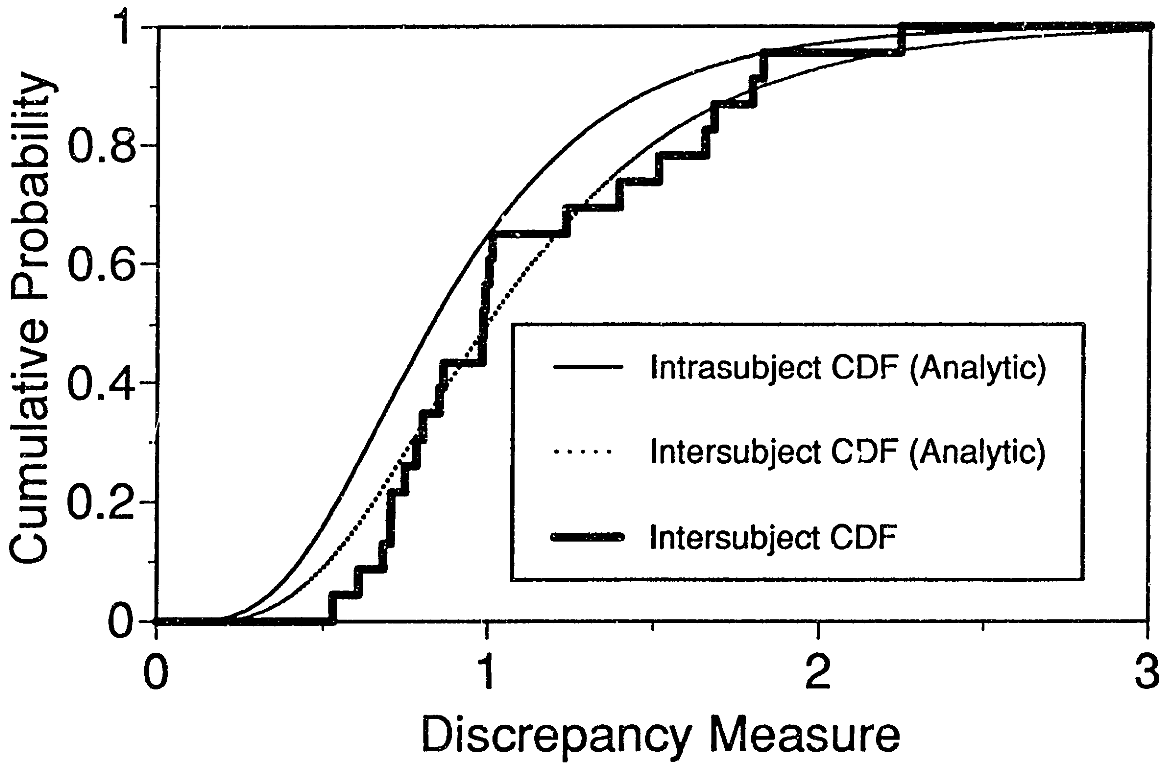

Figure 37. Cumulative probability distribution function for intersubject discrepancy measures.

these cases in detail, we may be able to understand why these intersubject differences occur.

For example, the reason that two control discrepancies are greater than 5.5 for the lateral deviation task is that one of the three subjects allowed the applied lateral forces to deflect their hand, while the other two subjects reported making conscious choices to stiffen up to resist the deflection. As a result, comparisons of the control trajectories between either of these two subjects and the one who chose not to resist the perturbation show significant discrepancies, even though their attractor trajectories are all fairly similar. This finding is comprehensible once we understand what choices the subjects made in performing the tasks.

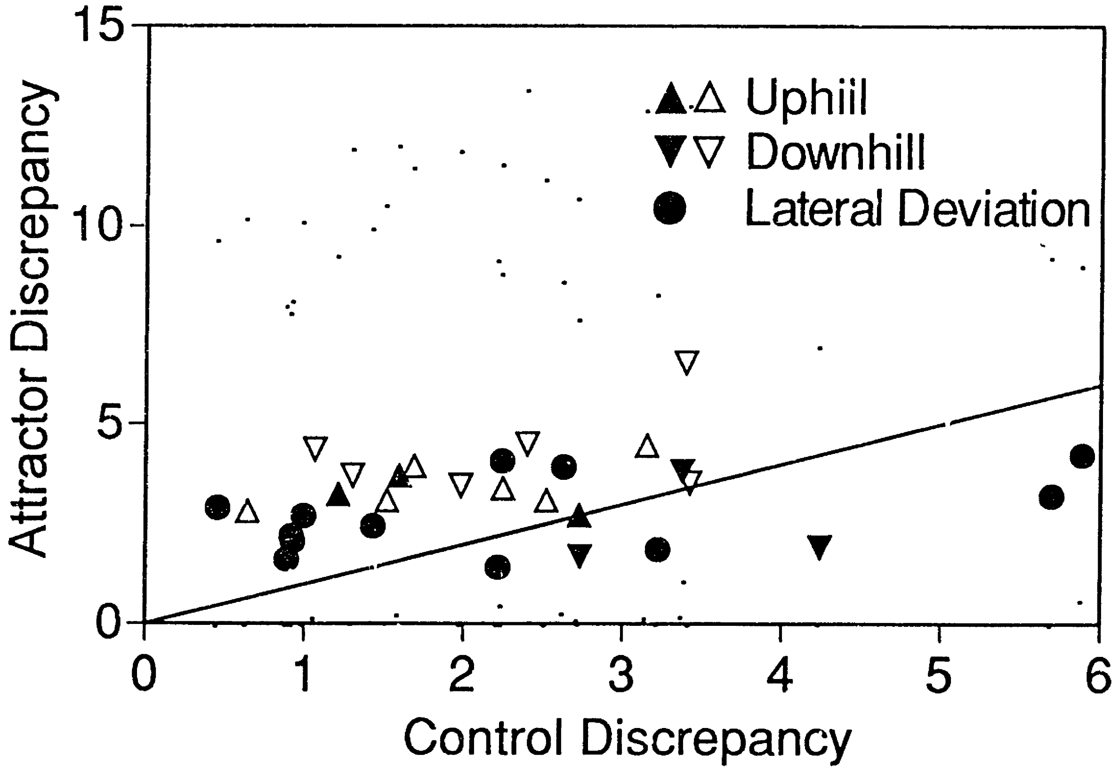

Figure 38. Intersubject discrepancy comparisons for dynamic tasks. The open symbols refer to comparisons involving one subject whose attractor trajectory estimates were not as reliable as those of the other two subjects.

On the whole, the intrasubject discrepancy during dynamic tasks is somewhat greater than that found during unconstrained tasks; typical control discrepancies are less than 3.5, while attractor discrepancies are less than 5.0. These larger discrepancy values will make it difficult to argue that the behaviour of any one subject is representative of the entire group, so we will restrict our arguments in the following sections to intrasubject comparisons.

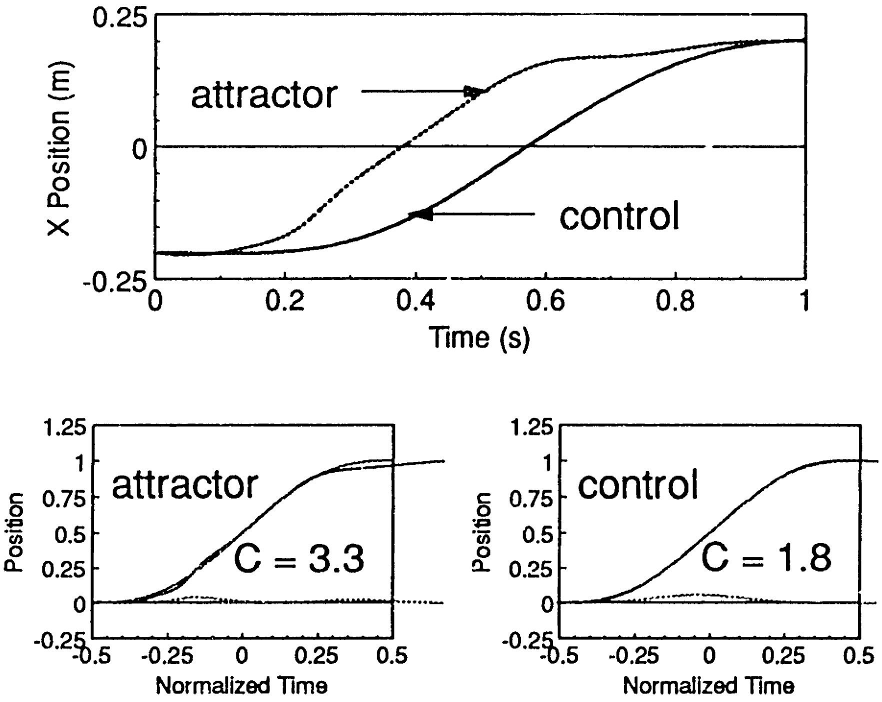

$$C = 100\sqrt{\int_0^1 \left(\bar{x} - \bar{x}_r\right)^2 dt}$$

Figure 39. Sample computation of complexity measure for an unconstrained task.

Strong Correspondence Hypothesis: Complexity

We now turn our attention to the experiments designed to test the strong correspondence version of the virtual trajectory hypothesis. In these experiments, we compare the unperturbed (control) trajectories and the estimated attractor trajectories to a minimum-jerk reference trajectory chosen to embody our notion of simplicity. This reference trajectory has the properties of straightness in task space and symmetry in its velocity profile. The computation of the complexity measure was described earlier in the methods chapter and is demonstrated for real data from an unconstrained movement task in Figure 39.

Ensemble Results

We first consider the results for the unconstrained movements *en ensemble*; Figure 40 shows a plot of the complexity values for the attractor trajectories of a wide variety of unconstrained movement tasks[29] plotted against the complexity values for the corresponding unperturbed trajectories. The line superimposed on the plot is a line of equicomplexity — points above the line indicate more complex attractor trajectories, while points below it indicate more complex unperturbed trajectories. If the strong correspondence hypothesis were correct, we would expect to find most of the points below the equicomplexity line. Clearly, this is not the case. Only three out of a total of approximately 90 tasks (defined as a combination of subject, orientation, speed and direction) show any indication at all that the attractor trajectory is simpler than the unperturbed trajectory, and all of these points are essentially on the line. One might argue that the uncertainties in the complexity measure (indicated by dots above and below each primary data point) are so large as to preclude us from drawing this conclusion, but at least 15 of the 90 points have lower limits which are above the line; we therefore have

---

[29] The figure is based on data from six subjects performing the unconstrained movement task in four different orientations (one subject only performed the task in two orientations), at two to four different speeds, and in both leftward and rightward directions, for a total of approximately 90 different tasks; each point in the plot is based on a set of eight to twelve measurements of the unperturbed and perturbed trajectories.

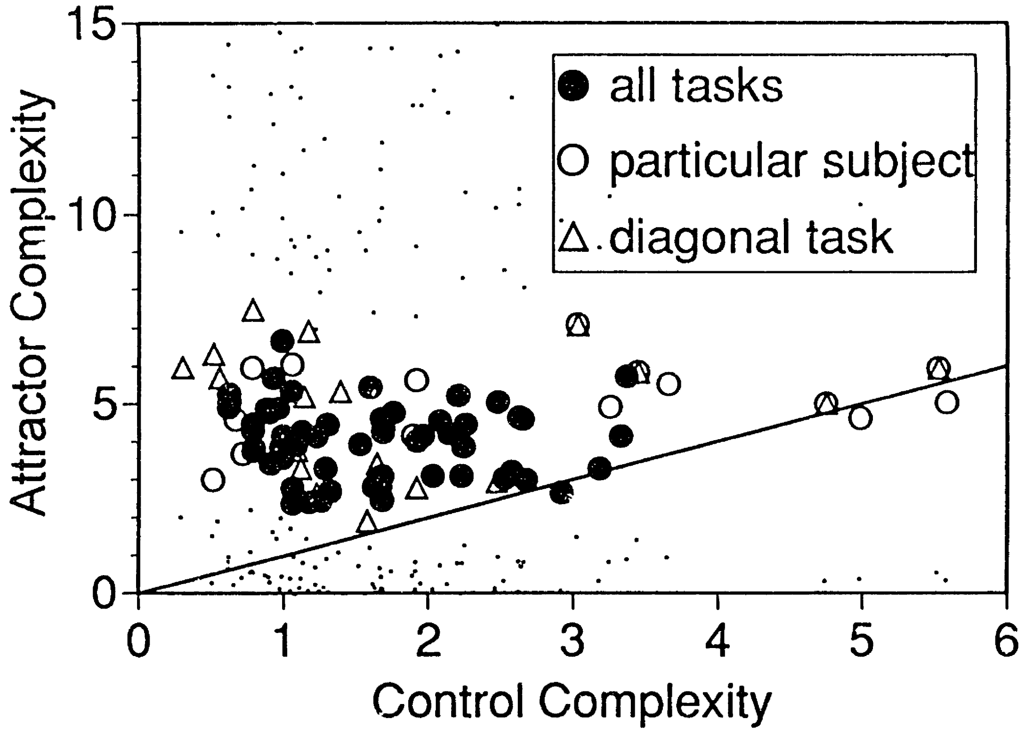

Figure 40. Attractor vs control complexities for the ensemble of unconstrained tasks.

sufficient evidence to dismiss the claim that attractor trajectories are always simpler than the unperturbed trajectories, and the weight of the evidence is that they are almost never simpler.

Figure 40 also hints that there is some structure to the distribution of the data points. For example, there are eleven points with complexity measures greater than 3.0 for the control trajectories. What is the reason for this? It turns out that eight of these points come from a single subject (indicated by the open circles in the plot) who tended to make significantly more curved control movements than the other subjects.[30] Similarily, fully half of the

---

[30] Recall from the discussion in the methods chapter of the complexity measures of circular trajectories that there is a nearly linear relationship between maximum lateral deviation and complexity measure for

points with attractor complexities greater than 5.0 and control complexities less than 1.5 were from movements along one of the diagonal orientations (indicated by open triangles). We will return to a more detailed assessment of these findings shortly, but will first discuss what can be inferred from this complexity plot.

Although the preponderance of data points lie above the equicomplexity line, we cannot simply compute the means and standard deviations of the control and attractor complexities and argue that the attractor complexities are, on average, significantly greater than the control complexities. The reason for this is that we have chosen a particular mix of tasks; since they have not been selected randomly from a universe of possible tasks, we cannot reasonably claim that their properties are randomly distributed. However, we have selected tasks which are broadly representative of the possible set of unconstrained planar reaching movements, so if there were some subset of tasks in which the control complexity were consistently greater than the attractor complexity, it is likely that we would have found an example of such a task.

Unconstrained Movements

In the following sections, we independently consider the effects of workspace location, execution speed, and direction of motion on the complexity measures.

---

moderate lateral deviations; if the trajectories were circular, then a complexity value of 3 would correspond to a lateral deviation of approximately 2 cm (5%), whereas a value of 6 would correspond to a lateral deviation of 4 cm (10%).

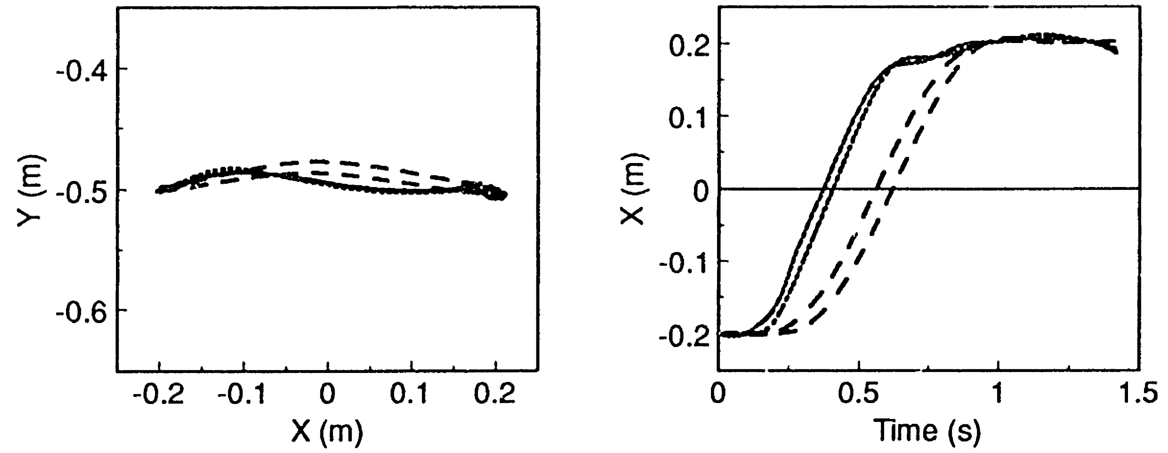

Figure 41. Unconstrained movement in transverse direction (coronal plane). Solid line - attractor trajectory estimate; dotted lines - perturbed trajectories; dashed lines - unperturbed trajectories.

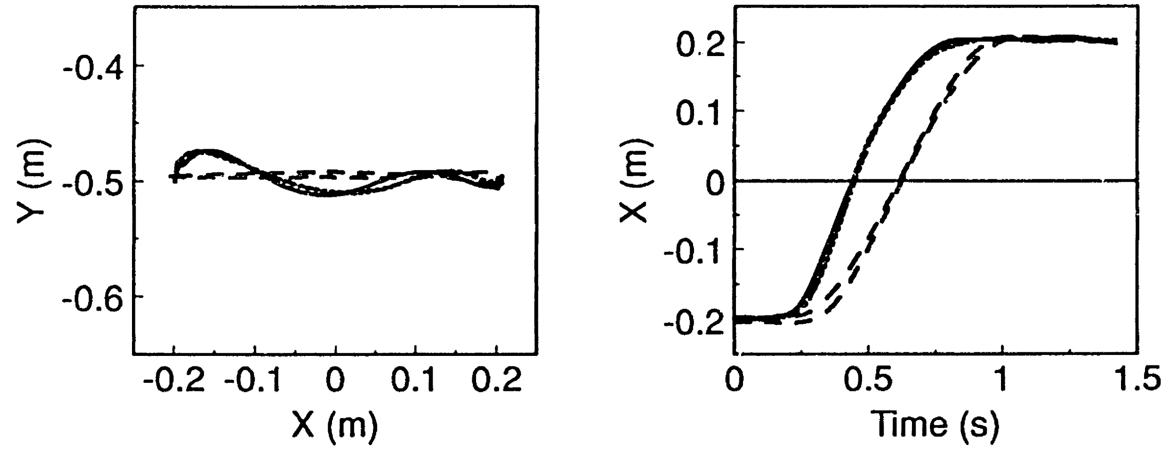

Figure 42. Unconstrained movement in sagittal plane. Solid line - attractor trajectory estimate; dotted lines - perturbed trajectories; dashed lines - unperturbed trajectories.

Workspace Location

Figure 41 and Figure 42 show two typical moderate speed, unconstrained movements executed at right angles to one another. Although many of the features of the two movements are similar, there are some obvious differences. The lateral deviations in the unperturbed transverse movement, for example, are considerably larger than those in the

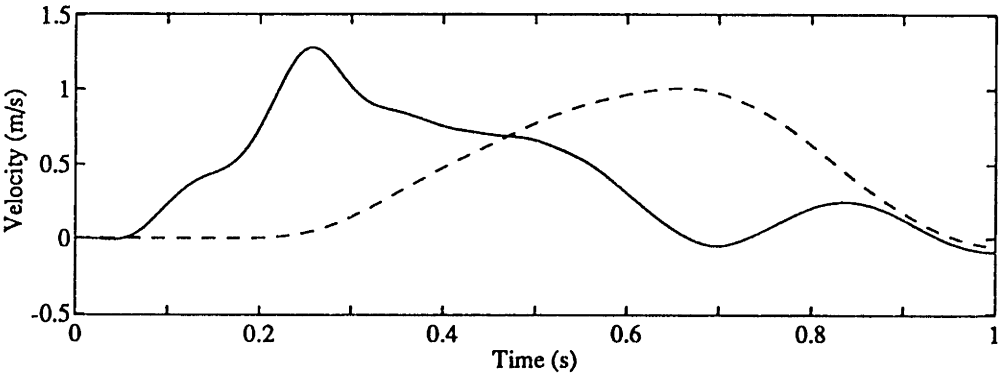

Figure 43. Tangential velocity profile for movement shown in Figure 41.

sagittal movement, while in contrast the lateral deviations in the sagittal attractor trajectory are more pronounced than those in the transverse movement. The computations of the complexity measures for the transverse movement were shown earlier in Figure 39; it would appear that the tangential component of the unperturbed trajectory is well-approximated by the min-jerk trajectory, so that most of its complexity is due to its lateral deviations. In contrast, the maximum lateral deviation of the attractor trajectory is smaller, but the attractor nonetheless has a larger complexity measure associated with it, primarily because of asymmetries in the tangential velocity profile (see Figure 43).

Figure 44 shows the complexity measures for the unconstrained movements broken down by position in the workspace (refer to methods chapter for sketch of locations of movement relative to subject's shoulder). As mentioned above, the four points with control complexities of approximately 5.0 all come from a single subject. All the other data, which includes tasks from all four workspace locations, are extensively intermingled. There may be some difference in the distribution of points from one workspace location to

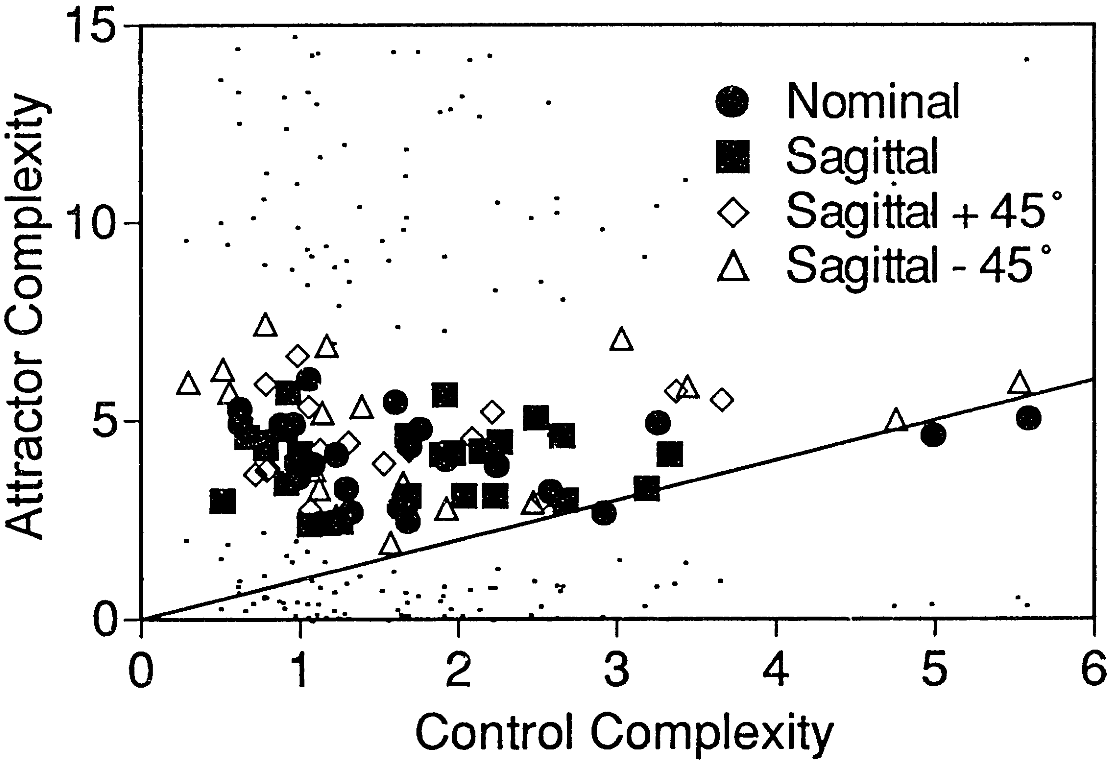

Figure 44. Complexity measures for unconstrained movements, broken down by location of the task in the workspace.

another, since there is a cluster of points in the upper lefthand corner of the plot from one of the diagonal tasks,[31] although the level of significance for drawing this conclusion is only slightly over 90% based on a two-dimensional Kolmogorov-Smirnov test (see Table 5), and, in any case, as is indicated in the footnote, we are on the borderline of having a sufficient number of points on which to base this conclusion.

We conclude, therefore, that the attractor complexities bear approximately the same relationship to the control complexities independent of the location of the task in the

---

[31] These tasks generally have roughly straight hand trajectories and relatively complex attractor trajectories.

workspace, even if the details of how the attractor and control trajectories differ at each location vary.

Location	D Statistic	Probability D > Observed	Effective # of Data Points
Nominal	0.19	64%	17.3
Sagittal	0.25	28%	17.7
Sagittal + 45°	0.27	31%	14.2
Sagittal - 45°	0.35	8%	14.8

Table 5. 2-D Kolmogorov-Smirnov statistics for comparison of the complexity data taken at different workspace locations.

Location	% (#) ATC > CTC	% (#) ATC >> CTC	# Points at Location
Transverse	88% (21)	17% (4)	24
Sagittal	100% (25)	8% (2)	25
Sagittal + 45°	100% (18)	22% (4)	18
Sagittal - 45°	100% (19)	26% (5)	19

Table 6. Proportions of tasks with attractor trajectory complexities (ATC) greater (2nd column) and significantly greater (3rd column — lower bound on ATC > CTC) than control trajectory complexities (CTC) at the four different workspace locations.

Table 6 summarizes the data in Figure 44 and points out that virtually all of the points lie above the equicomplexity line (with the exception of four points corresponding to 12% of the tasks performed in the transverse position). Since there are points from tasks in each position whose lower bounds on the attractor complexity are greater than the control

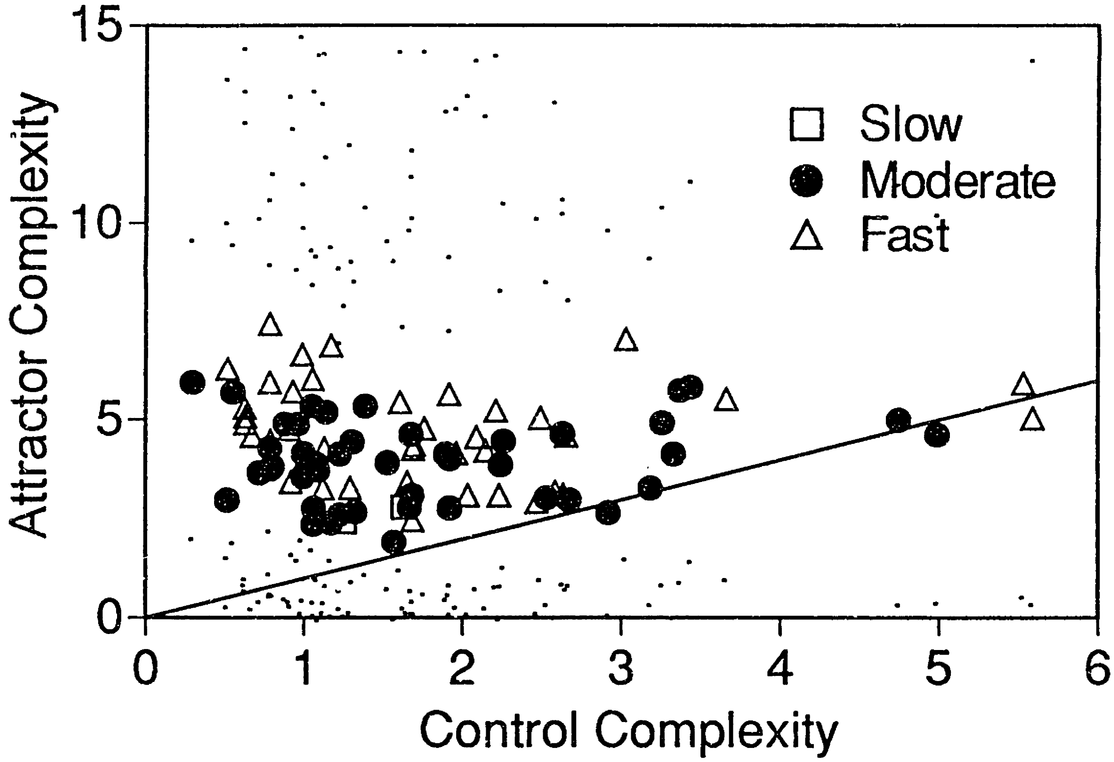

Figure 45. Complexity measures for unconstrained movements executed at different speeds.

complexity, the strong correspondence hypothesis would not appear to hold at any workspace location for the mix of unconstrained tasks studied.

Speed Dependence

Figure 45 shows a comparison of control and attractor trajectory complexity measures for movements in all workspace locations executed at both moderate (~50% V_{max}) and high (~75% V_{max}) speed movements. There is a slight difference in the distribution of the points as determined by a 2-D Kolmogorov-Smirnov test (*D* statistic: 0.33; probability of *D* being greater than this if the two distributions are identical: 4.2%; effective number of points: 21.0), although the two distributions are clearly extensively intermingled, even though the maximum speeds plotted approach 80% of a maximum-effort movement.

Kawato's critique of the virtual trajectory hypothesis, however, would suggest that at higher speeds, there would necessarily be an increasing discrepancy between the virtual and actual trajectories if the same hand path is followed at all speeds. Why then do we not see this here?

The answer lies in the fact that the command trajectory is only half of the story — the other half is the impedance around the command trajectory. Even though we are not explicitly measuring the subjects' impedances, we can develop some insight into their properties by considering the relative durations and onsets of the attractor and control trajectories. Recall that the durations and onsets of these trajectories are computed by finding the best fit between a min-jerk reference trajectory and the given trajectory; the midpoint and duration of the min-jerk trajectory are used to characterize the measured trajectory.

Consider first the relationship between the lead time of the attractor trajectory (the interval between the time the attractor trajectory reaches the midpoint of the movement and the time the control trajectory does) and the duration of the control trajectory, shown in Figure 46. Kawato implicitly assumes that the impedance displayed by the subject will not change with increasing execution speed. If this is so, then the attractor trajectory will lead the control trajectory by a constant amount. However, it is clear that despite the spread in the data points in Figure 46, faster movements (shorter durations) have attractor trajectories with shorter lead times; that is, increasing movement speed *decreases* the

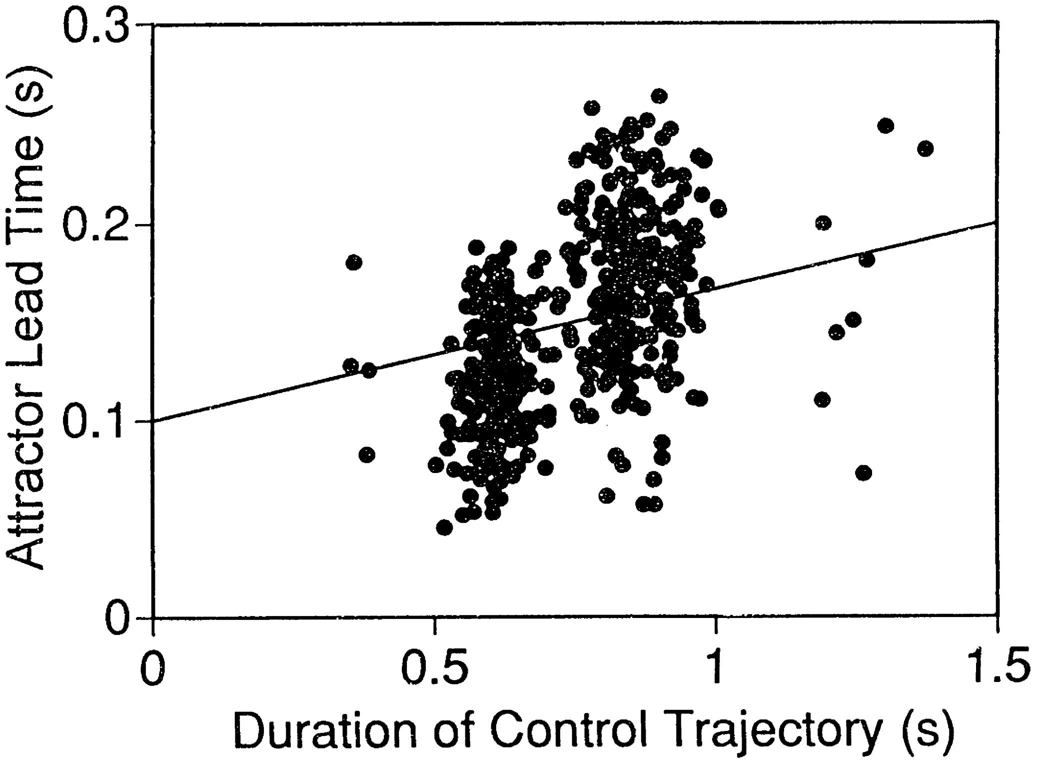

Figure 46. Lead time of the attractor trajectory as a function of the duration of the control movements.

temporal discrepancy between the attractor and control trajectories — a clear contradiction of Kawato's criticism.

A second piece of evidence for the claim that the shape of the attractor trajectory is not changed significantly as movement speed increases is that the ratio of the rise-times of the attractor and control trajectories show no systematic change with movement speed (see Figure 47). Note that the high ratios at shorter duration movements are predominantly artifactual; the attractor trajectories for the larger subjects at these high speeds have significant uncertainties associated with them, so the estimates are not necessarily well approximated by a min-jerk transition; this implies that the timing values do not necessarily well characterize the attractor trajectories (see Figure 31). In any case, we

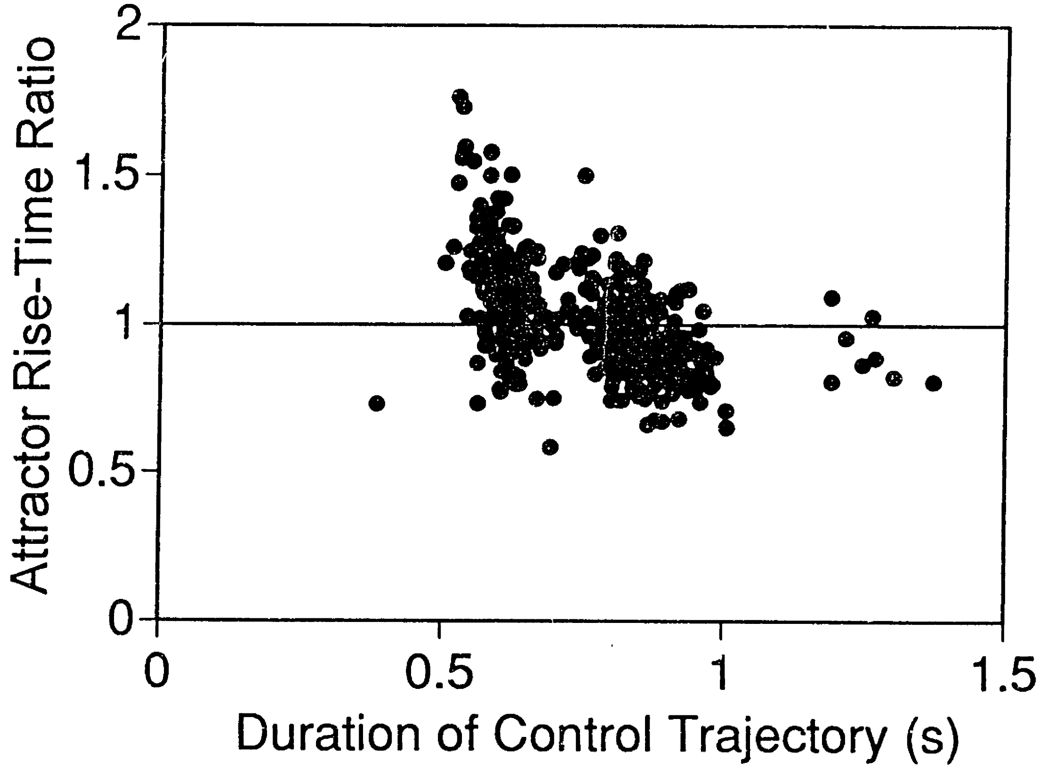

Figure 47. Ratio of attractor trajectory rise-time to control trajectory rise-time as a function of duration of control movements.

would have expected to see the attractor trajectory rising extremely rapidly relative to the control trajectory to generate the forces needed to move the hand more quickly along the nominal path. The fact that the rise-time of the attractor trajectory is, on average, the same as the rise-time of the control trajectory suggests that changes in execution speed are achieved by time-scaling the attractor trajectory and appropriately adjusting the impedances so as to maintain the correspondence in shape between the attractor and control trajectories.

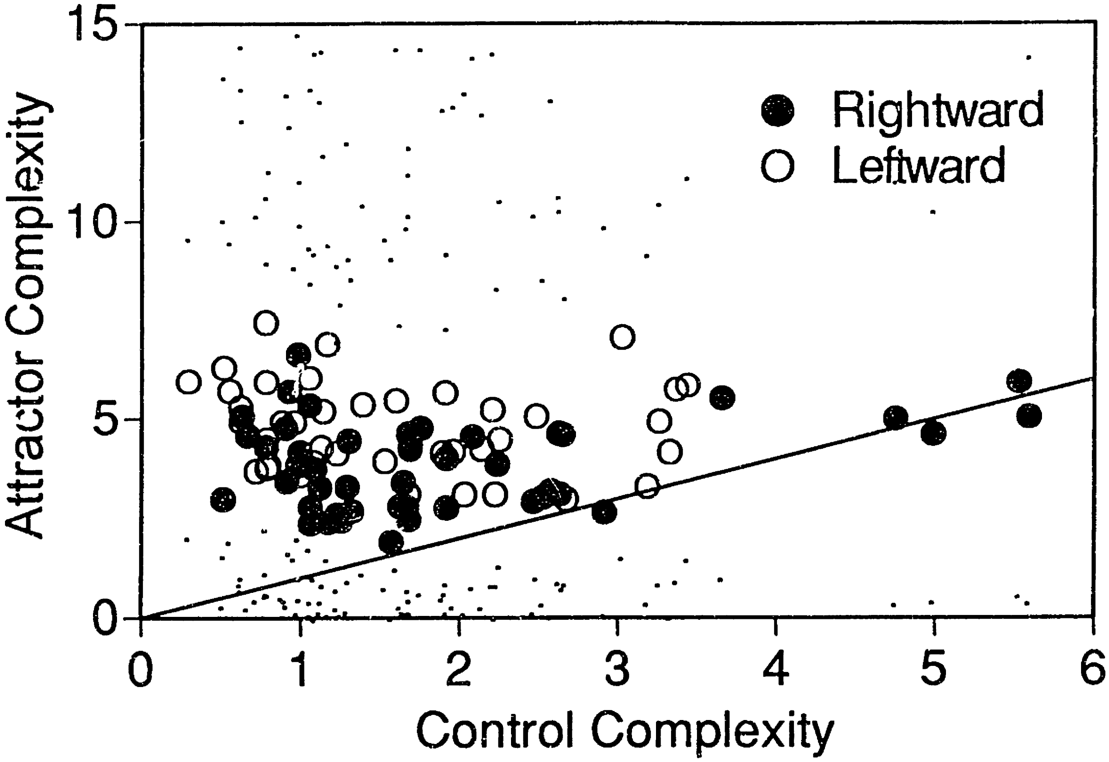

Figure 48. Complexities for unconstrained movements in opposite directions.

Direction

As discussed in the methods chapter, we cannot directly compare the leftward and rightward motions because the inertial behaviour of the manipulandum is different in the two cases when expressed in a task-based coordinate system (with the X-axis running from the starting position to the target). We cannot therefore be sure whether any observed difference in behaviour is due to changes on the subject's side of the interaction port or on the side of the environment. We can, however, ask whether or not changing the movement direction for a well-characterized and easily-learned task has a significant effect on the planning strategy. Figure 48 shows that, in fact, the leftward movements do seem to be associated with slightly greater attractor complexities, even though, with the

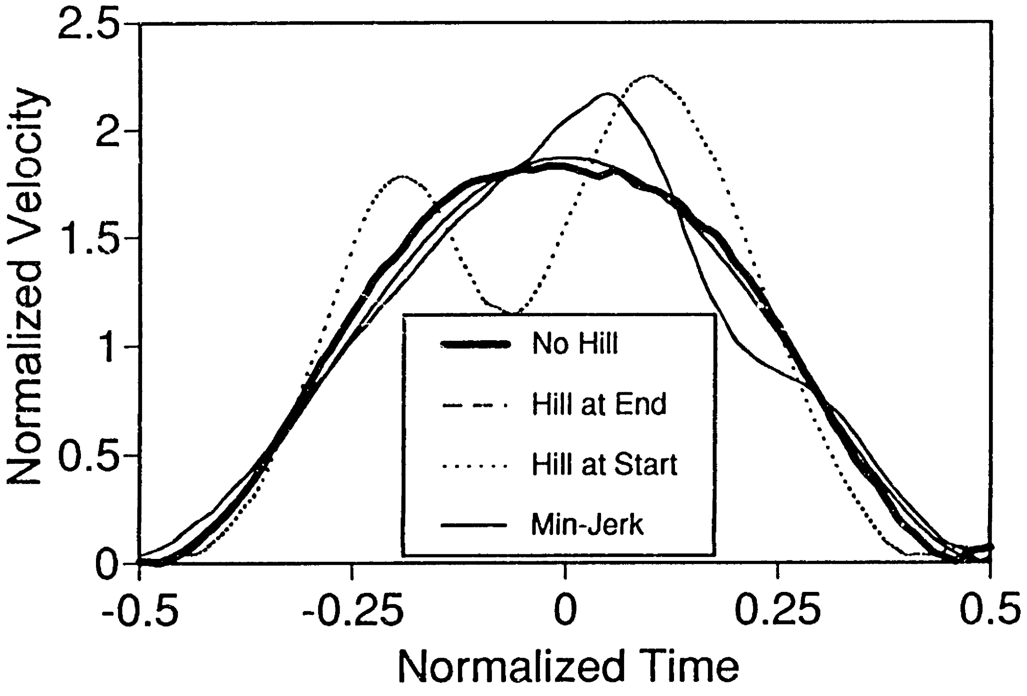

Figure 49. Tangential velocity profiles for hill movements.

exception of the one subject discussed earlier, the control trajectories are of comparable complexity in the two directions. This finding of a slight difference is confirmed by a 2-D Kolmogorov-Smirnov test (*D* statistic: 0.37; probability of *D* being greater than this if the two distributions are identical: 1.5%; effective number of points: 21.4), although the degree of intermingling of the distributions is nonetheless considerable.

Hill Tasks

Distortions in Unperturbed Task

The purpose of the "uphill" and "downhill" tasks was to give subjects the opportunity to execute more asymmetric trajectories than are typically observed in the more unconstrained movements. In fact, subjects did execute control trajectories which were

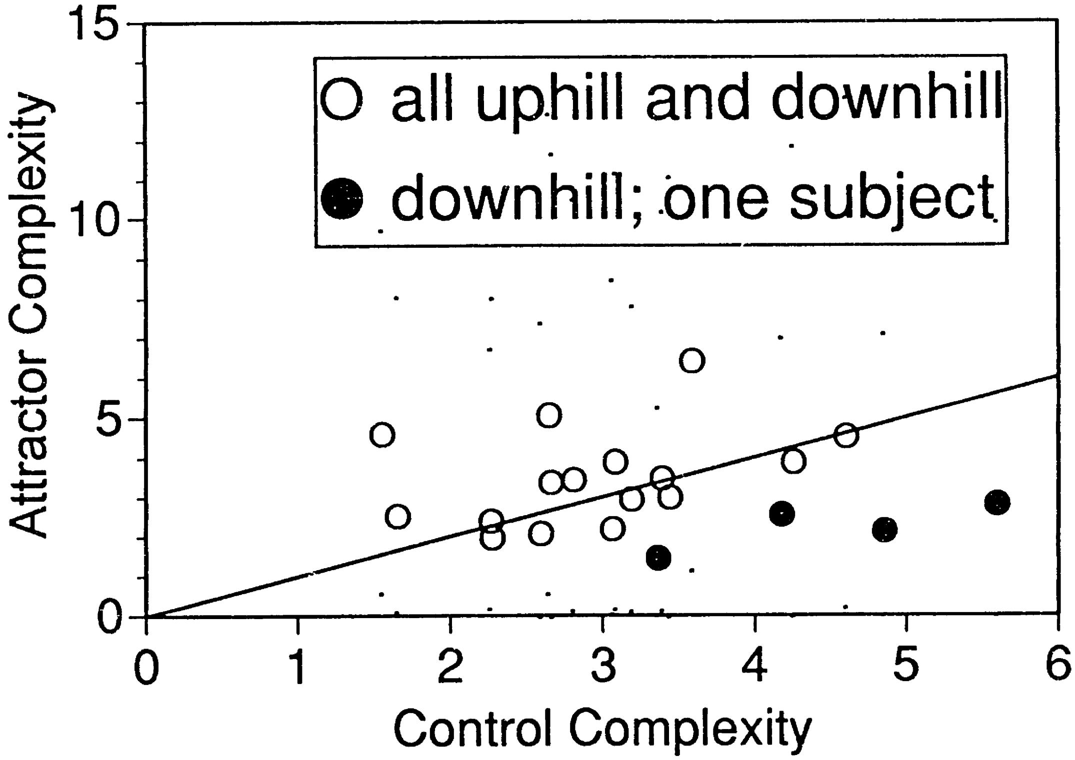

Figure 50. Complexity plot for hill movements.

measurably less symmetric, as is apparent in a plot of the tangential velocity profile of an unconstrained movement and two uphill movements, the two latter profiles resulting from hills placed at opposite ends of the movement (see Figure 49).

This induced asymmetry is also reflected in the relatively high mean complexity values for the control trajectories, as shown in Figure 50 and as confirmed by a 1-D Kolmogorov-Smirnov test comparing the distribution of control complexities in the hill tasks with the control complexities in the full set of unconstrained tasks; the K-S test had a *D* statistic of 0.63, a probability of $1.5\text{x}10^{-4}\%$ that *D* would be this large if the two distributions were identical, and was based on the equivalent of 16.4 data points. The complexity distribution for the hill tasks was a factor of 2.16 greater than that for the unconstrained task, based on the scaling factor that would bring the distributions into maximum agreement, although even so the shapes were not a good match — with scaling, the maximum probability that the two sets came from the same distribution was less than 23%.

Figure 50 also shows that, in striking contrast to the results for the unconstrained movements, here fully half of the attractor trajectories are less complex than the corresponding control trajectories. This is most true for one subject in particular whose downhill movements exhibit considerable asymmetry; it is this subject's trajectory data which is plotted above in Figure 49 and whose corresponding complexity measures are shown as filled circles in Figure 50. The existence of relatively simple attractor trajectories in the hill tasks is also suggested by a comparison of the unconstrained attractor complexities with the hill task attractor complexities. Using a 1-D Kolmogorov-Smirnov test, we found a *D* statistic of 0.43 when all unconstrained tasks were considered (0.31 when only the moderate speed unconstrained movements in the transverse position were considered); these *D* statistics yielded a probability of 0.3% based on 16.4 equivalent data points that the hill task attractors were as complex as those of the unconstrained tasks in general and a probability of 18% based on 11.7 equivalent data points that they were as complex as those of the matched unconstrained tasks. There is therefore no evidence that interaction with the dynamic environments resulted in more complex attractor trajectories for the matched tasks; in fact, there is a slight probability

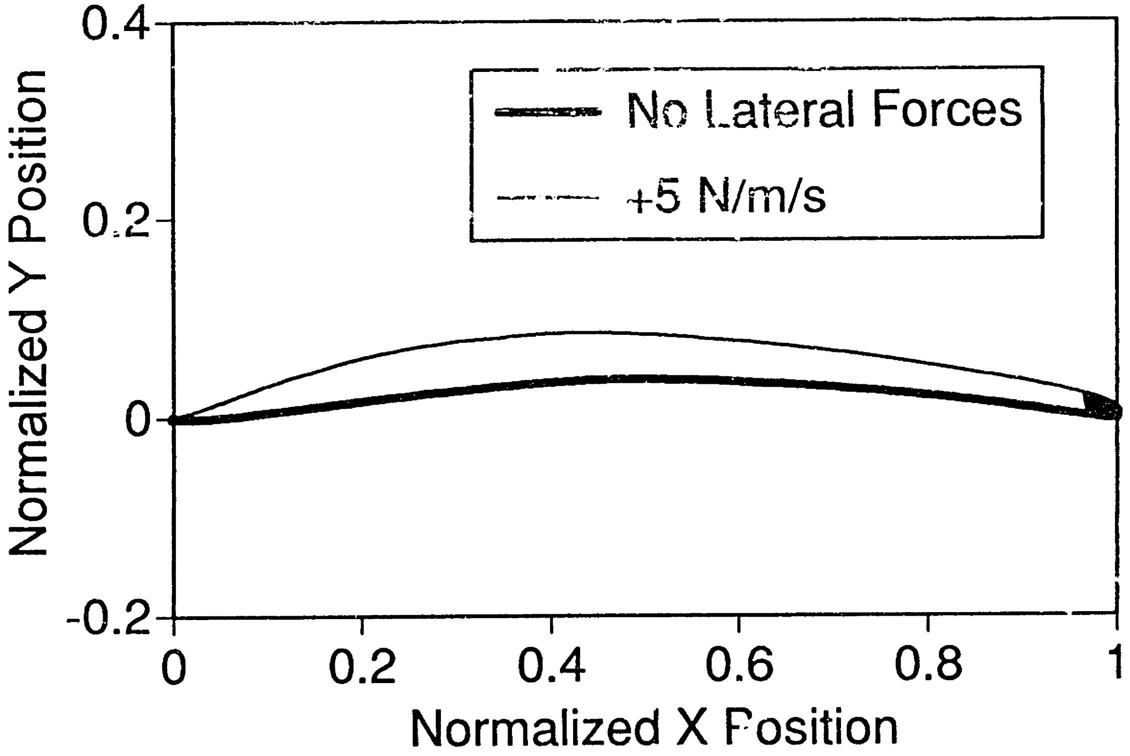

Figure 51. Sample trajectories during lateral deviation task.

(82%) that interaction with the dynamic environments actually produces simpler attractor trajectories than are found in the unconstrained movements.

The existence of these relatively simple attractor trajectories during a dynamic task suggests that the subjects are not consciously controlling their hand paths, but are instead content to let the interaction with the dynamic environment determine the path between the target points, so long as they are able to control their departures and arrivals. In the upcoming section on discrepancy measures, we will consider the evidence that the attractor trajectories are actually more consistent than the control trajectories in different dynamic environments.

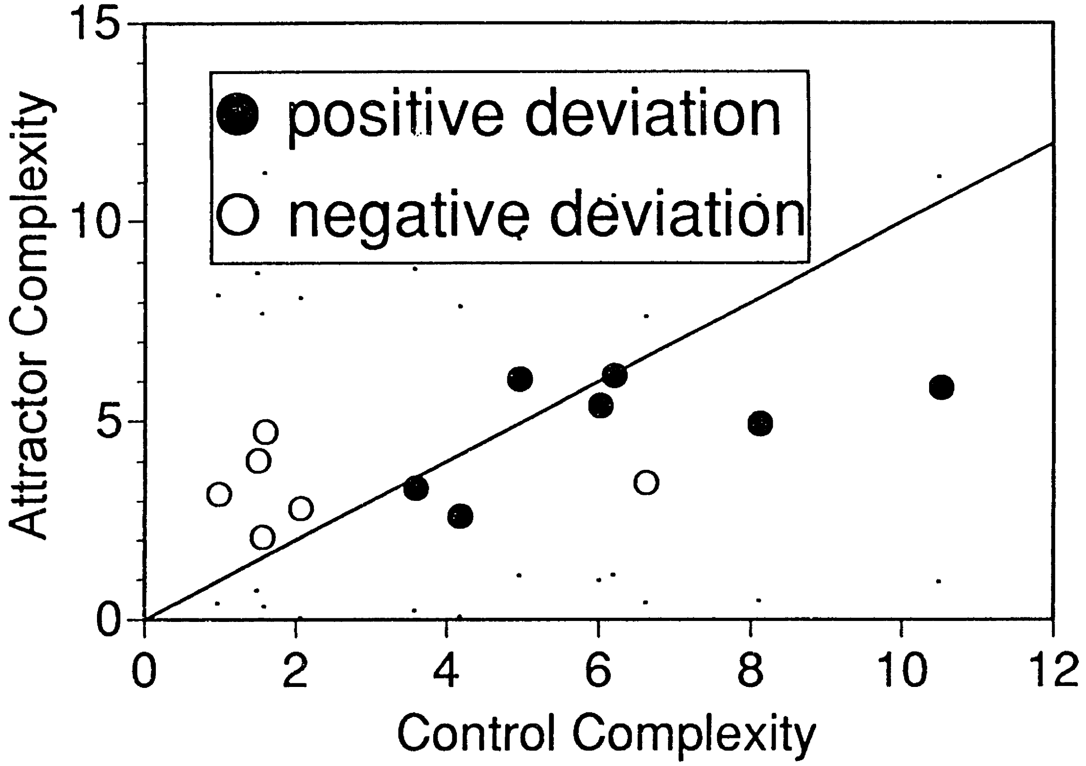

Figure 52. Complexity measures during lateral deviation tasks.

Lateral Deviation Tasks

Distortions in Unperturbed Task

As with the hill tasks, the purpose of the lateral deviation tasks was to provide an opportunity for subjects to execute comparatively more complex trajectories than they would in an unconstrained task. As shown in Figure 51, velocity-dependent lateral forces applied in the appropriate direction generally caused subjects to trace a more curved path than they would in the absence of the forces. These lateral deviations are reflected in control complexity measures as much as two or three times greater than those of unconstrained movements (see Figure 52), and this difference is confirmed by a 1-D

Kolmogorov-Smirnov test (*D* = 0.54; *P* = 0.15%; equivalent points = 11.4) which shows that the control complexities for the lateral deviations are more complex than those for the attractor trajectories. Virtually all of the cases in which the applied forces were directed so as to enhance the existing curvature in the corresponding unconstrained movement had attractor trajectories which were less complex than the control trajectories. Furthermore, a 1-D K-S test (*D* = 0.18; *P* = 80%; eq. # = 11.4) showed no discernible difference between the attractors in the lateral deviation task and the unconstrained tasks, although there was some minor evidence that the attractors of the matched unconstrained tasks (i.e., those at the same speed and workspace location) were slightly less complex than those in the lateral deviation tasks; this is based on a 1-D K-S test (*D* = 0.43; *P* = 5.6%; eq. # = 8.9). Again, as with the hill tasks, we infer that subjects are not intrinsically concerned with generating a particular endpoint trajectory unless it is required by the experimental protocol; their underlying attractor trajectories remain relatively unaffected by their interactions with the lateral perturbing forces.

Conclusions

The complexity analysis outlined above reveals no evidence that the attractor trajectories during unconstrained tasks executed at more than 30% of subjects' maximum possible speed are less complex than the achieved hand trajectories; in fact, in virtually all cases examined, the attractor trajectories were either demonstrably or probably more complex than the control trajectories. Changing the speed of movement from 50% to 75-80% of maximum had little appreciable effect on the distribution of the attractor and control complexities, which suggests that the subjects scale up their impedances with execution speed.

When we attempted to change the dynamic forces experienced during the motion by supplying a position- or velocity-dependent field for the subject to interact with, the results were more in accord with the strong correspondence hypothesis. The hand trajectories became more distorted than they were during the unconstrained tasks, and the subjects showed no inclination to "correct" these trajectories. The attractor trajectories, in contrast, did not exhibit these distortions and instead remained approximately as complex as they were during the unconstrained tasks.

We conclude that the results of the unconstrained movement complexity analysis disprove the strong correspondence hypothesis, but that the results in the two dynamic environments tested do not contradict the hypothesis.

Insensitive Strategy Hypothesis: Discrepancy

Since the complexity analysis presented above has disproven the strong correspondence hypothesis, we now consider whether or not the insensitive strategy hypothesis is plausible. Recall that this hypothesis claims that although the attractor trajectories may not be simpler than the control trajectories they may nonetheless be more consistent as the speed of a task or its orientation in the workspace is changed. To assess these claims, we compute discrepancy measures between the control trajectories in different tasks and between the corresponding attractor trajectories in the same tasks and compare these discrepancy measures. As with the complexity measures described earlier, we will plot attractor discrepancies against control discrepancies; if the insensitive strategy hypothesis is correct, most of the points will lie below the equidiscrepancy line.

Unconstrained Movements

As with the complexity analysis, we will begin the discrepancy analysis by considering the unconstrained movements first, followed by those in the position- and velocity-dependent environments. A typical situation we will be analyzing is shown in Figure 53. Here we have the paths of two unconstrained movements made by the same subject at the same speed in two different workspace locations — a sagittal movement and one at approximately 45° to the sagittal movement. The data from two subjects is shown to demonstrate the consistency of the measurements across subjects. This example is one of a number of situations in which the control trajectories look relatively similar to one another while the corresponding attractors are much less similar; such examples disprove the insensitive strategy hypothesis.

Subject 1: Sagittal and Diagonal Movements

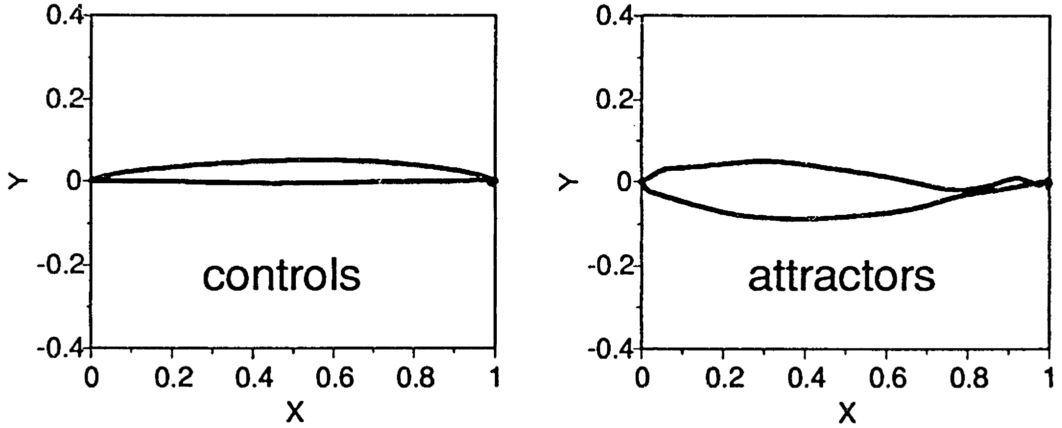

Subject 2: Sagittal and Diagonal Movements

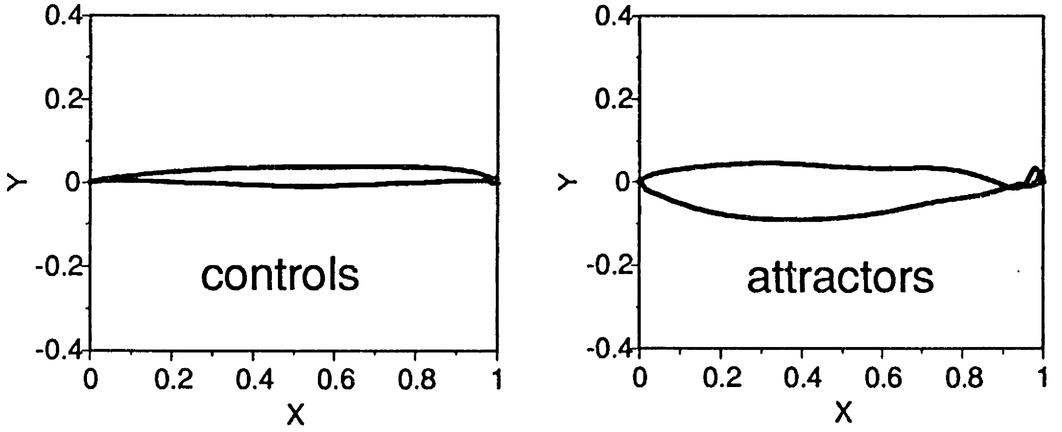

Figure 53. Comparison of unconstrained movements, along with their attractor trajectories, in two different workspace locations.

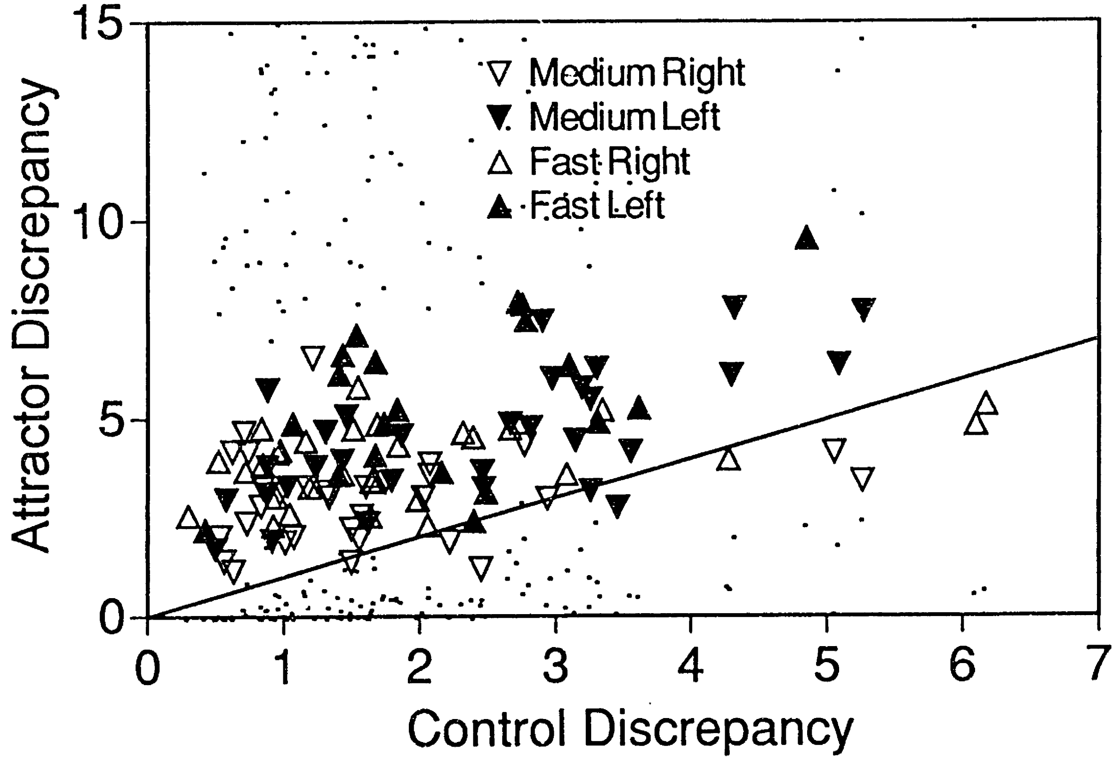

Figure 54. Discrepancy plot for unconstrained movements performed at different workspace locations.

Workspace Dependence

The above example suggests that if we inspect the same movement executed at different locations throughout the subject's workspace, we will find a number of examples in which the attractors exhibit greater discrepancy than do the control trajectories. Figure 54 shows a plot of the attractor and control discrepancies for over 100 comparisons of unconstrained tasks executed at various workspace locations, with each type of task indicated by appropriate symbols.[32] Although the results are not quite as unambiguous as those for the complexity measures, there being only about five points whose lower discrepancy bound is above the equidiscrepancy line, the great majority of the points are nonetheless above the line. Note that all five of the points which lie below the line and which have control discrepancies greater than 4.0 are from a single subject — the same subject who generated the largest control complexities (discussed above in the section on complexity analysis) — so if we exclude the results from this subject, virtually all of the data points would be above the equidiscrepancy line. Even so, we do not have to demonstrate that all of the comparisons are above this line to disprove the hypothesis; we merely have to show that *some* variations in workspace location result in greater attractor discrepancies. To that end, we consider some of the most convincing examples of attractor discrepancies being greater than control discrepancies.

Figure 55 shows comparisons of the paths of two different subjects performing different pairs of moderate speed unconstrained movements in different workspace locations; in both cases, the minimum attractor trajectory discrepancy is either comparable to or somewhat greater than the control discrepancy. Amongst the 100+ comparisons in Figure 54, there are six to ten such cases, so it is clear that the insensitive strategy hypothesis does not hold when the position of the task in the workspace is changed.

---

[32] In each comparison, the speed, direction and subject are held constant — only the orientation is changed.

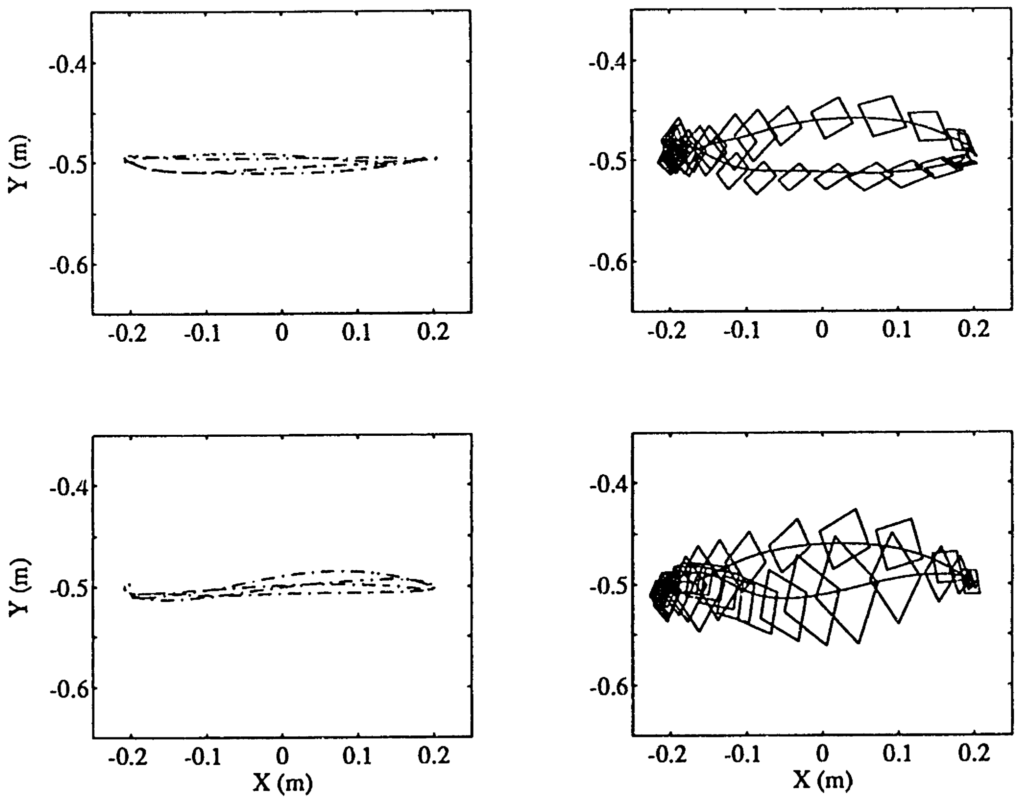

Figure 55. Comparison of paths of two subjects in different workspace locations.

Speed Dependence

By changing the position of the task in the workspace, we have significantly changed the coriolis and centrifugal loads the arm experiences during the movement. We can also generate changes in these dynamic forces by changing the speed of execution of the movement, in which case we would like to know whether or not the attractor trajectories change shape significantly. The results of forty-four such comparisons are shown in Figure 56, most of which (40 of 44 cases) are comparing 50% V_{max} movements with 75%

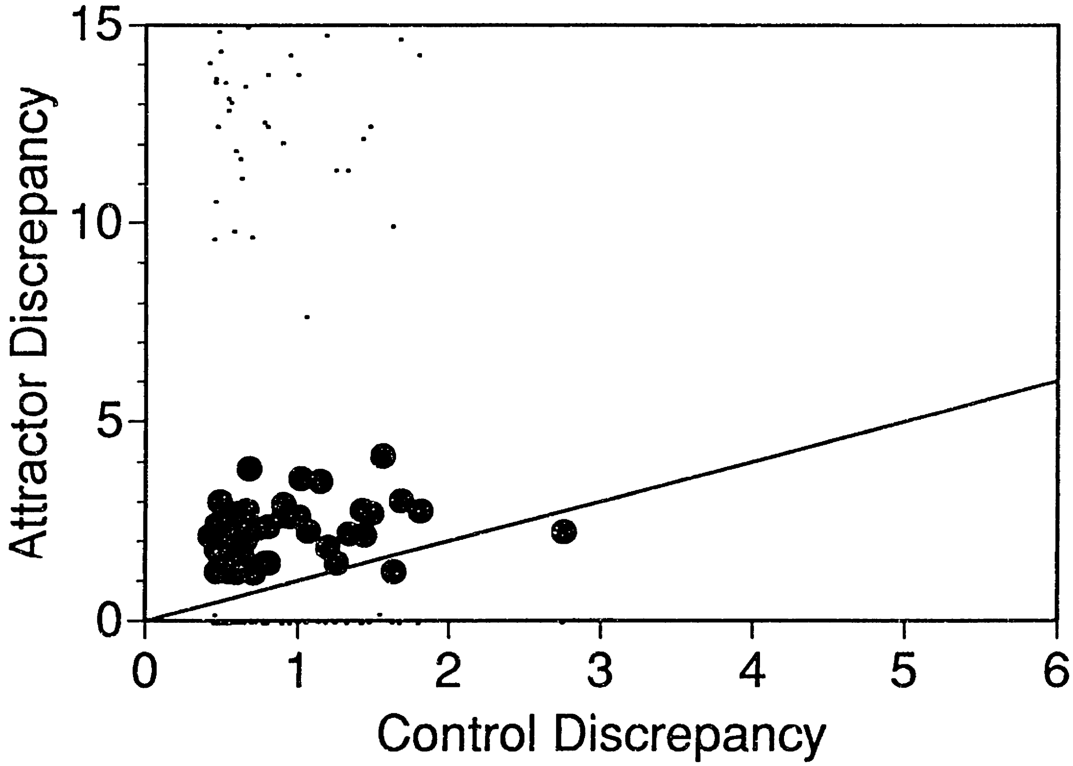

Figure 56. Discrepancy measures for the same task performed at different speeds.

V_{max} movements (the remaining four cases compare 30%-max-speed movements with 75% V_{max} movements).

Although we again see that virtually all (only two possible exceptions) of the data points are above the equidiscrepancy line, this time the results are not as compelling as the comparisons between different workspace locations since none of the points has its minimum attractor discrepancy value above the line. Furthermore, almost all of the discrepancy measures are relatively small compared with those found with changes in workspace location. In the latter case, control discrepancies ranged up to 5 or 6 and attractor discrepancies up to 8 or more; here, we find control discrepancies almost always lower than 2 (comparable to the intrasubject discrepancies on repeated performance of the

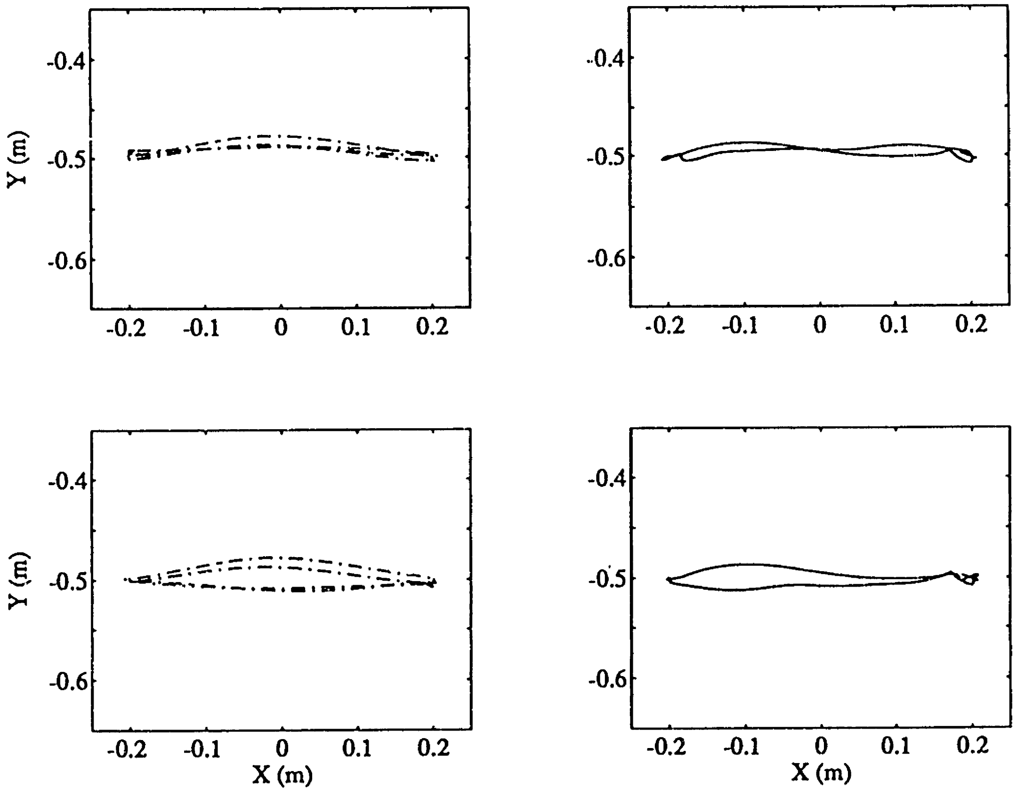

Figure 57. Comparison of rightward and leftward movements in (A) a fixed external coordinate frame, and (B) a task-based reference frame.

same unconstrained task — see Figure 34) and attractor discrepancies generally below 3 or 4 (the two largest attractor discrepancies are for 30% V_{max}/75% V_{max} comparisons).

Such comparatively small discrepancy measures make it difficult to claim that the attractors are truly more variable than the control trajectories. Instead, we would do better to claim that there is little evidence that changing the speed over the range of 30% to 75% of a maximum effort movement has any dramatic effect on the shape of either the attractor trajectories or the control trajectories.

Direction Dependence

When the subject executes a point-to-point task in both directions, the paths often look similar when viewed in a fixed external coordinate frame (see Figure 57). In a task-based coordinate frame, however, the movement in one direction exhibits a rightward curvature, whereas the reverse movement exhibits a leftward curvature. Since the virtual trajectory hypotheses claim that only the task-based coordinate frame is relevant, we would expect to see comparatively similar attractor trajectories when the direction of movement is changed. Figure 58 shows forty comparisons of unconstrained movements at both moderate and fast speeds executed in both directions. The four extreme points with control discrepancies greater than 4.0 are for the same subject noted earlier whose paths were significantly more curved than the other five subjects. From this plot, it appears that there are two or three examples in which the lower bound on the attractor discrepancy is comparable to the control discrepancy, and only three out of thirty-six cases (excluding the four extreme points) in which the attractor discrepancy is lower than the control discrepancy; i.e., the great majority of the data points are above the equidiscrepancy line. These results would suggest that the subjects change strategy (indicated by changes in the attractor trajectory) when reversing their direction of movement, although we must use caution in drawing this conclusion because we cannot be certain that the change in behaviour is solely attributable to the change in direction; the observed responses may be due to the changes in the inertial behaviour of the manipulandum as viewed in the task-based coordinate frame.

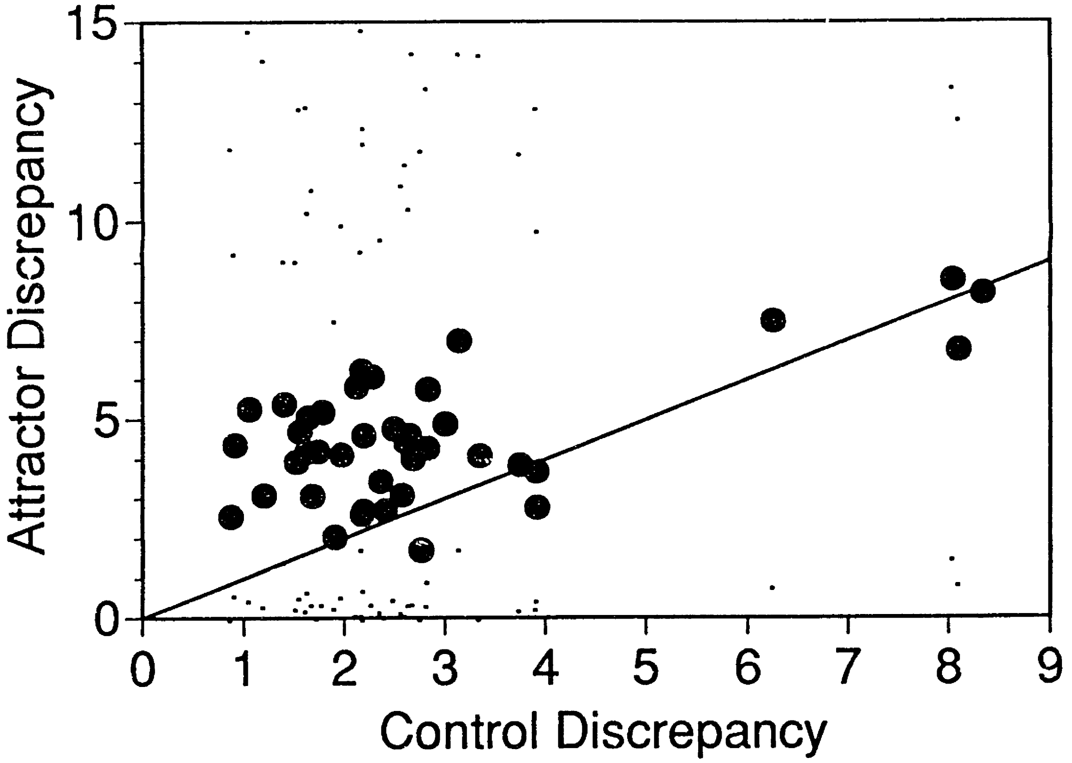

Figure 58. Discrepancies between rightward and leftward movements.

Hill Tasks

The unconstrained movement tasks have left us in a somewhat ambiguous position: the subjects appear to change their attractor trajectories substantially in response to changes in the orientation or direction of the task, but may simply time-scale them to change the speed of execution. Why is there this difference in response? The hill and lateral deviation tasks are a means by which we can change the forces during movement without changing the task orientation and thereby offer us an alternative test of the emerging hypothesis that the pattern of central commands are primarily determined by the orientation of the task in the workspace.

The attractor and control discrepancy measures for twenty-three pairs of hill tasks are shown in Figure 59; each of these points represents a comparison either between hills sloping in the same direction but located at different positions between the target points or between hills and unconstrained movements; uphill and downhill movements are not directly compared to one another. The results are quite surprising. In contrast to the cases in which the task orientation was changed, here we find that sixteen of the twenty-three cases are below the equidiscrepancy line, and none of the points above this line are significantly above it (in the sense that the minimum attractor discrepancy is above the line). Perhaps the most surprising result is that the majority of the points with control discrepancies greater than 2.5 correspond to comparisons between hills placed at the beginning and end of the movement; this implies that even though the interactions with the two different hills induced significant differences in the actual hand trajectory, the underlying attractor trajectory changed almost not at all (on par with the distribution of individual unconstrained trajectories). These results, then, support the notion that at a given workspace location, a change in the forces of interaction is accomodated without changing the attractor trajectory, at least at this first stage of learning. It is, of course, always possible that with extensive practice the subjects will learn a different strategy which causes the actual trajectory to more closely resemble the actual trajectory in the unconstrained case.

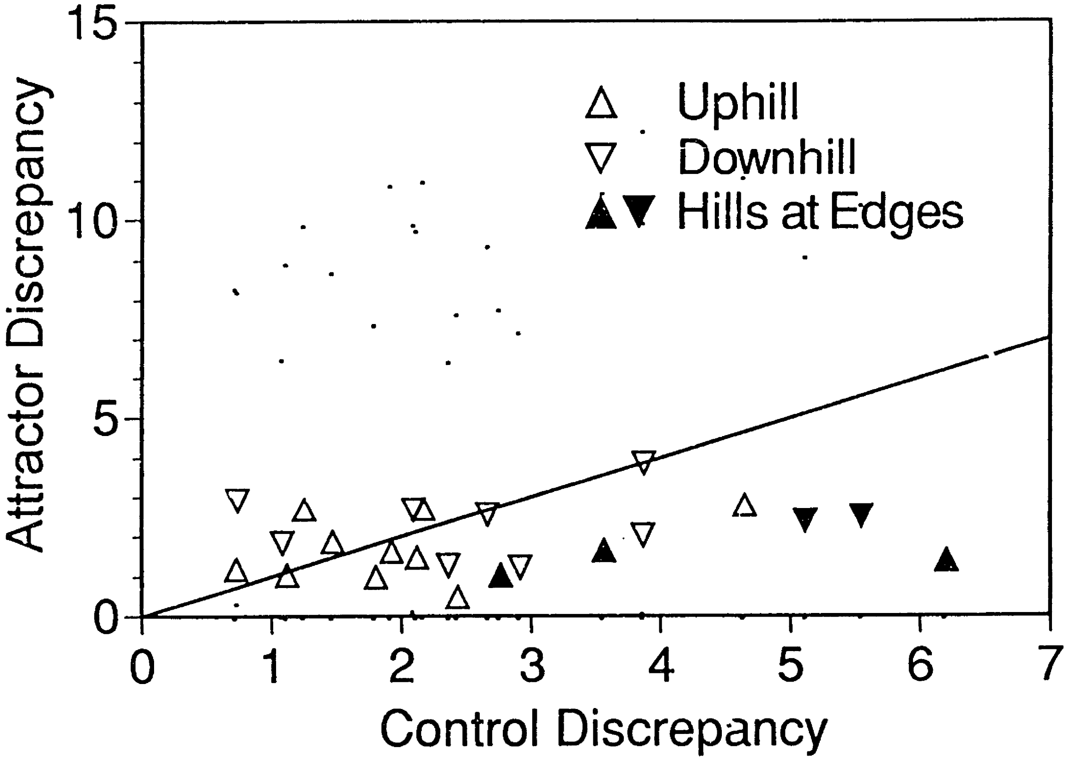

Figure 59. Discrepancies for hill tasks.

If the attractor trajectory does not change when hills are imposed, does that imply that no change at all occurs in subjects' mechanical responses? I believe that there are other changes, although they cannot be detected directly from the plots already shown but are instead inferred from observations of subjects' experiences in learning how to perform the dynamic task. I should first point out that even though we found little change in the shape of the attractor trajectories in the presence of the different hills, they certainly were not all of the same duration. As might be expected, the attractor trajectories of downhill movements had much longer durations than those of uphill movements, with the durations of the unconstrained movements lying in between. When subjects first encounter a hill after performing a series of unconstrained movements, they not surprisingly fail to perform the task succesfully. After a few more tries, they usually succeed. This learning process reflects the subjects' need to learn to time-scale their command trajectories to successfully meet the timing requirements of the task. What is more interesting is their response when the position of the hill changes. On the first trial after the change, the subjects again typically fail in the task. Again, they adjust after a few more tries. What we find, however, when we analyze their attractor trajectories is that not only the shapes but the durations also are extremely similar. That is, if the impedances that the subjects implemented were the same for both hill positions and their attractor trajectories were the same, then they should have succeeded when the hill position was changed. The fact that they did not indicates one of two things: either the change in the hill position caused a dramatic enough change in the feel of the task to cause them to abort their unfolding motor plans, or switching from one hill position to another required a change in their arm impedance which in turn required a few trials for recalibration. The comments of subjects when performing the task tend to support the latter conclusion; most subjects found that having the hill at the beginning of the movement was the easiest condition to learn — for uphill movements, they typically spoke of trying to "pop over" the hill by moving more quickly towards the target — whereas they often reported difficulty in trying to "stiffen up" at the target to prevent overshoot when the hill was near the end of the movement.

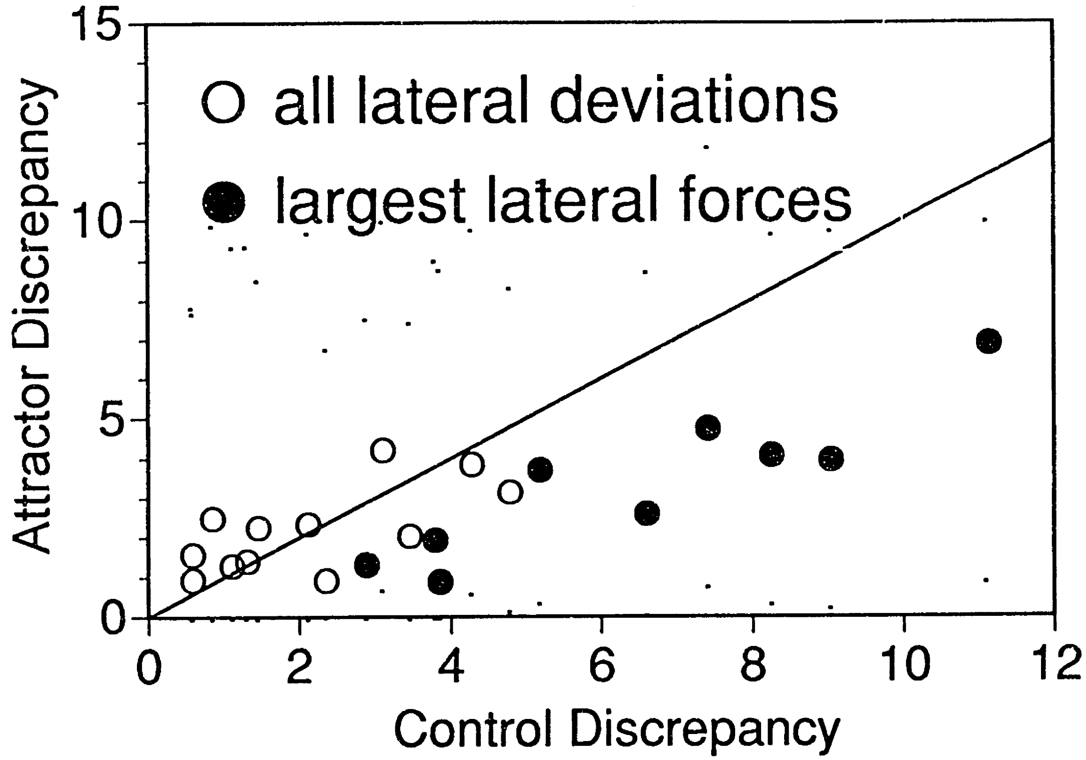

Figure 60. Discrepancies for lateral deviation tasks.

Lateral Deviation Tasks

Whereas the hill tasks encourage changes in the symmetry of the task along the line connecting the target points, the lateral deviation tasks encourage discrepancies between tasks by applying forces perpendicular to the direction of motion. Figure 60 compares the control and attractor discrepancies for twenty-one comparisons of these lateral deviation tasks. In thirteen of the cases, the attractor discrepancy is lower than the control discrepancy, and none of the eight remaining cases have attractor discrepancies significantly greater than the corresponding control discrepancies; in fact, all of these eight cases have attractor discrepancies less than 2 units greater than the control discrepancies. The solid dots in the figure indicate that for the largest changes in the applied lateral forces (> 5 N/m/s for one task; < -5 N/m/s for the second task in the comparison), the attractor trajectory changes relatively little. Although there may be some compensation for the induced increases in curvature of the movements, this compensation is only partial; in effect, the subjects almost ignore the presence of the lateral forces once they realize that they can perform the task without explicitly compensating for them.

Conclusions

The results of the discrepancy analysis are somewhat similar to those of the complexity analysis presented prior to it, although perhaps not quite as convincing. The control trajectories do appear to be more consistent than the attractor trajectories during unconstrained movements, particularly with changes in workspace location and possibly with direction of movement, although this conclusion must be tempered by the recognition that the inertial environment relative to the task coordinates changes when the movement direction is changed. There is, however, no evidence that changing the speed of the movement from 50% to 75-80% of the maximum possible speed produces any significant change in the shape of the attractor trajectory.

When subjects interacted with the hill and lateral deviation environments, their attractor trajectories were often more consistent than were their control trajectories. So, while the insensitive strategy hypothesis is disproven on the basis of the inconsistencies between attractor trajectories at different workspace locations during unconstrained tasks, it appears that for a given workspace location, the insensitive strategy hypothesis may well describe how the subject copes with changes in the dynamic environment.

CHAPTER
SIX

DISCUSSION AND CONCLUSIONS

The results described in the previous chapter present an ambiguous picture. On the one hand, the data from unconstrained movements in different parts of the workspace clearly shows that both the strong correspondence and insensitive strategy hypotheses should be rejected. On the other hand, when movements in different dynamic environments at the same workspace location are considered, the results do seem to agree with the hypotheses. How then are we to understand these results? In this chapter, we consider the strength of the various results and attempt to reconcile the seemingly contradictory data by asking questions about how we learn motor tasks. As we consider the data, we will pose many other questions which can inspire future work, but before we analyze the data in more detail, I will begin by briefly restating the problem I had set out to explore.

The Hypothesized Kinematic Basis of Motor Control

The major question in motor neuroscience which inspired the study reported in this thesis is how the central nervous system represents and plans the execution of a motor task. In particular, we are interested in knowing whether the brain represents all aspects of a task or if it is able to rely on the intrinsic mechanical behaviour of the spinal and subspinal components of the neuromuscular system to provide stability and robustness, thus freeing it to plan the higher level aspects such as the desired trajectory.

We noted earlier that these questions are important in helping us understand how manipulation tasks are performed because, despite our best robotic efforts, the most facile manipulators in existence are humans and our simian friends. These biological manipulators are clearly "constructed" using very different control principles than those we have used to construct robots. The typical approach in robotics is to use high-precision sensors, relatively ideal actuators (e.g., torque motors, which deliver a desired torque almost independently of the applied load), and full state feedback with short feedback delays. In contrast, our muscles are extremely nonlinear and compliant actuators, our reflex delay times are long, ranging from 30-400 ms, and our sensors do not produce outputs which are straightforwardly related to the length, velocity and force measurements typically used in robotic designs. Moreover, recent psychophysical studies suggest that humans do not even use an internally consistent representation of the world in planning our motor activity; Fasse, for example, has demonstrated that our length judgments do not agree with our angle judgments (i.e., given the task of deciding whether a given triangle was isoceles or not, our answer would depend on whether we based our judgment on a comparison of the two approximately equal sides or of the two approximately equal angles) [Fasse 92]. All current robotic control approaches are based on a consistent representation of the world. If we seek then to understand the human system in the same terms as we use to design robots, the question we must face is how the central nervous system is able to compensate for an extremely sloppy set of sensors and actuators. Since this approach is likely to be fruitless, a major challenge for motor neuroscientists is to explain how the organization of the neuromuscular system facilitates human adeptness at manipulation. It is interesting to note that the structure of biological systems is already recognized as being advantageous for manipulator design; some of the emerging robotic designs actually emulate the compliant behaviour of muscle [Mansfield 92, Oshima 94, Fasse 94].

The central idea underlying the various permutations of the equilbrium point or virtual trajectory hypotheses is that there is a partitioning of control responsibility between the proximal and distal portions of the neuromuscular system. Although it is difficult to make a clear distinction between the high-level and low-level systems because of the extensive neural interconnections and continuum of information processing times (ranging from 30 ms for monosynaptic reflex loops to several hundred milliseconds for central cognitive processing), we are speaking roughly of a division at the top of the spinal column. Giszter's work on frogs, for example, demonstrates the existence of functional modules which exist in the spinal column and which produce behaviourally significant responses to non-specific stimuli [Giszter 91, 93]. According to the views of Bizzi et al., the spinal circuitry, coupled with the intrinsic muscle properties, enables the limb to behave essentially passively during interactions with the environment [Bizzi 92, 93]. This guarantee of coupled stability opens the possibility that the higher levels of the central nervous system could carry out motor planning in purely kinematic terms and transmit a kinematically-based task specification to the spinal circuits for implementation. In particular, the virtual trajectory notion posits a strong relationship between posture and movement; the spinal circuits are viewed as translating neural codes representing a posture into the set of muscle activations required to achieve it, and movement is seen as arising as a shift in the postural specification. One implication of this view is that the limb must be stable during movement since it is stable at posture. This prediction is in contrast to dynamic theories of movement which claim that the central nervous system computes the forces that muscles need to generate to achieve the desired motion. Since such theories are not fundamentally predicated on the existence of stability properties of the peripheral neuromuscular system, they predict nothing about the limb's response to perturbations.

The goal of my research, then, was to design and perform a test of this proposal that the central planning process is carried out in primarily kinematic terms. Such a test must rely on or involve a definition of the previously loosely defined notion of "attractor" behaviour, so the first question I addressed in this thesis was how we define attractor behaviour in terms of the measurements we can make. The second question was how we interpret the attractor behaviour we do measure in physiological terms so that we can relate these interpretations to the hypotheses we are testing. The final question was what the characteristics of a motor task would be if it were planned using a kinematics-based strategy.

A Definition of Attractor Behaviour

The answer to the first question is a rigorous definition of an attractor trajectory expressed solely in terms of interaction port behaviour; this definition is extremely general and can be applied to any system or subsystem with interaction ports. In essence, an attractor trajectory is defined to be a stable trajectory of a system which starts at equilibrium and which evolves under the condition that the interaction port forces are zero at all times. This definition makes specific our concept of the virtual trajectory as being that time-evolving point towards which the system is tending at each instant; I use the term attractor trajectory to refer to a measurable quantity and reserve the term virtual trajectory to refer to the kinematic equivalent of the neural command.

The interpretation of the attractor trajectory in physiological terms is somewhat more ambiguous because of the complex chain of dynamic events linking neural activity to mechanical consequences. If we accept the view that descending neural activity effectively determines muscle rest lengths, then there exists a mapping from each neural command to a corresponding equilibrium limb configuration (this can be a many-to-one mapping). The dynamics of the musculoskeletal system can then be regarded as determining the response

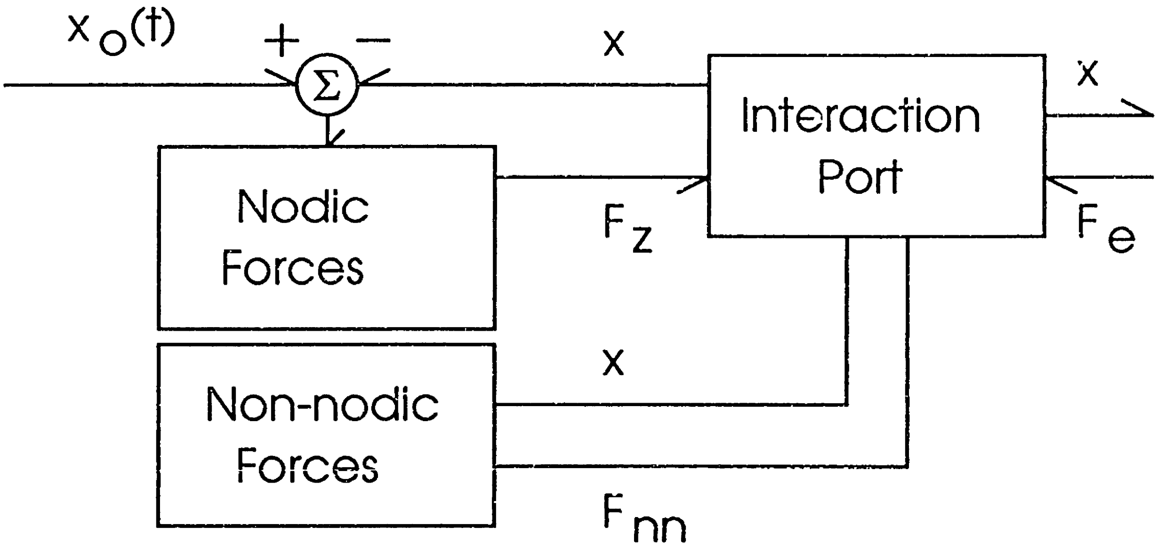

Figure 61. Limb model partitioned into those elements whose behaviour is a function of the difference between the interaction port position and the virtual position (nodic) and those whose behaviour depends only on the interaction port position (non-nodic).

of the limb to a kinematic input: the time sequence of equilibrium configurations. It is this input which we refer to as the virtual trajectory.

Note that the virtual trajectory is an equivalent representation of the command — we may not be able to identify symbolic features in the neural activity itself, but the activity will have mechanical consequences and hence an equivalent mechanical description. The existence of this equivalence allows us to develop a model of the musculoskeletal system (see Figure 1) in which all internal forces on the limb can be partitioned into those (which we will call *nodic*) which depend on the difference between the current limb position (and velocity) and the virtual trajectory and those (which we will call *non-nodic)* which depend on the absolute limb position and velocity (e.g., inertial forces and joint friction). If we consider the subsystem comprised only of the nodic elements, then the virtual trajectory can be shown to be the attractor trajectory of that subsystem.

Identifying a measured attractor trajectory with the mechanically equivalent command to the muscles clearly requires an adequate model of these complex tissues. In principle, any model of the system which is controllable could exhibit the same interaction port behaviour, but only a small subset of these models would competently describe the actual processes in a real limb. Since the virtual trajectory hypotheses make claims about the actual commands, it is imperative that the model we use be strongly grounded in physiological reality. Unfortunately, there are still many unanswered questions about the dynamic behaviour of the limb. For example, only recently have the nonlinear static stiffness fields of the limb been measured over a significant range [Shadmehr 93]; we still have little idea how these fields vary during movement or during interactive tasks. Our knowledge of the damping properties of muscle is likewise limited to individual muscle models or to linear approximations at posture [Dolan 91]; we have little knowledge of how to characterize damping during motion.[33] In particular, we do not know whether the damping forces belong to the set of those internal forces which depend on the absolute endpoint trajectory or to the set of those which depend on the difference between it and the virtual trajectory.

This lack of knowledge of the limb's impedance characteristics prevents us from making a strong claim that any attractor trajectory we might compute corresponds directly to an actual neural command. However, by accounting for as many of the non-nodic forces as

---

[33] There are more complex models for damping in isolated muscle (see, for example, Wu [1990], but these models have not yet been integrated into a model of the whole arm.

possible and by treating all other forces as though they were nodic, we can design tests which have the potential to reject certain classes of the virtual trajectory hypotheses regardless of the details of the actual neuromuscular apparatus. Such tests are possible because each layer of dynamics added to the output of a subsystem which exhibits an attractor trajectory degrades the correspondence between the command to the subsystem and the output of the aggregate system, while the inverse process of removing layers of dynamics will result in progressively closer correspondence between the virtual trajectory and the attractor trajectory of the remaining subsystem. Since the virtual trajectory hypotheses claim that there is a consonance between the task description and the virtual trajectory, or at least a consistency between virtual trajectories executed under similar dynamic conditions but at different points in the workspace, we would expect to find that each time we account for a significant non-nodic phenomenon, such as inertia, the attractor trajectories of the remaining subsystem will show greater consonance or consistency across tasks than did the attractor trajectories at the previous stage. If we do not observe this progression, then we can be assured that any further accounting for non-nodic phenomena will not reverse the trend and so we are justified in rejecting any hypothesis which claims that the virtual trajectory of *any* subsystem connected through the non-nodic dynamics which we have already accounted for will exhibit greater consonance or consistency than the measured paths.

We thus have a means for testing several variations of the so-called virtual trajectory hypotheses; we have both a practical definition of attractor trajectories which allows them to be measured and a prediction about the relationship between the unperturbed motion and the pre-inertial[34] attractor trajectory. Failing to satisfy this prediction would constitute a strong disproof of certain versions of the virtual trajectory hypotheses because the prediction makes no appeal to assumptions about the actual neuromuscular dynamics; if the hypotheses fail, the test conclusions would be compelling.

Advantages of and Concerns With Approach

The primary advantage of my approach to testing the virtual trajectory hypothesis is that, in contrast to earlier model-based tests, my test does not require a model of the muscles' mechanical impedance. The attractor trajectory is defined solely in measurable terms and the measurement of the preinertial attractor trajectory requires only an estimate of the inertia of the limb of the subject. Furthermore, our experience with the technique indicates that even substantial errors in the inertia estimate (larger than would reasonably be expected using the estimation techniques we have employed) cause little change in the attractor trajectory estimate because the stiffness of the arm is sufficiently high. Because we use an iterative procedure to converge to our attractor trajectory estimate, we are not required to project from a minimally perturbed trajectory, as Flash did; this fact strengthens our claim that the procedure reliably locates the attractor trajectory, which is, in essence, the zero of a multidimensional impedance relation.

---

[34] The pre-inertial attractor trajectory is the attractor trajectory of those components of the limb model which remain after accounting for inertial effects.

Although I have dealt with some of the technical objections to the iterative approach to locating attractor trajectories in earlier chapters, I would like to briefly reiterate my response to the four primary objections: the accuracy of the inertial model, the validity of the "Do not intervene" instruction, the question of the completeness of learning, and the adequacy of the error bounds on the attractor trajectory estimates.

For the purposes of this study, we have modelled the arm as being comprised of two segments — an upper arm and a forearm — which lie in a horizontal plane and which are connected to each other, to the torso, and to the environment by one degree-of-freedom pin joints at the shoulder, elbow and wrist. In reality, of course, the shoulder is a complex joint with three degrees-of-freedom at the ball-and-socket joint, compliance in the socket itself, and some limited movement possible relative to the torso. Similarily, the elbow and wrist joints are actually more complex than our model of them, although we have taken steps to ensure that they conform as closely as possible to the model. For example, the subject's shoulder is strapped to a chair and they wear a wrist brace which strongly inhibits motion there. The subject's elbow is not constrained to lie in the horizontal plane, but observation of their elbows during movements shows that they generally stay within 2-3 cm of the plane. Dolan demonstrated that such small excursions have virtually no effect on the inertial behaviour of the arm in the plane because the effect is proportional to the square of the cosine of the angle of elevation of the elbow; for small angles, the square of the cosine is small compared with unity [Dolan 91]. The freedom of movement at the shoulder joint is also likely to have a relatively small effect on the inertial behaviour of the arm both because the excursions themselves are small (on the order of 1-2 cm) and because the excursions are not free but elicit reaction forces at the joint which are likely to dominate the inertial forces required for the small excursions. Finally, the small amount of flexion at the wrist joint is unlikely to be significant because the perturbed hand trajectories cover a large fraction of the workspace (large in comparison with the flexional freedom) so that the bulk inertial properties are likely to dominate the inertial response. Flexion would only be a problem for high frequency, low amplitude perturbations such as those used by Bennett for identification of inertial parameters (although in his case the issue of wrist flexion was less problematic because his apparatus allowed him to apply perturbations at the elbow). So, because the perturbations are low frequency and of large amplitude, and because I have taken measures to limit joint movement, I conclude that the inertial model is valid for the analysis of this experiment, although I have allowed parametric uncertainties of greater than three times what would be expected from my inertial estimation technique to partially cover the possible inaccuracies in my inertial model.

Using the "Do not intervene" experimental paradigm is required to obtain the preinertial attractor trajectory measurements. It can be criticized on the grounds that the perturbation is large enough that the subject must necessarily respond to it, although the successful experience with the technique of a number of researchers including Latash and Shadmehr argues that subjects can learn to ignore the response. Even if one doesn't accept this argument, at least the first 250-400 ms (based on reaction time experiments) ought to be considered free of voluntary interventions. In principle, one could apply perturbations beginning at different times throughout the movement to check whether the early responses to later perturbations match the later responses to perturbations applied closer to the onset of the movement. We did not perform such tests because the manipulandum we used was insufficiently strong to generate the required torques. However, the faster movements generally lasted less than 400 ms, so they can be considered essentially complete before any voluntary intervention could manifest itself, and since the data from the faster movements is in substantial agreement with that from slower movements, it is reasonable to conclude that subjects followed the "Do not intervene" sufficiently well to allow us to accept the data from slower movements.

The third primary objection relates to the seemingly contradictory result that the attractor trajectories were simpler and more consistent than the actual trajectories in the presence of the position- and velocity-dependent force fields (the hill and lateral deviation tasks). Is it possible that we obtained this result because the subjects had insufficient practice with the task and so had not learned it as well as the unconstrained movement tasks? Based on the practice opportunities in the specific context of these experiments, the answer would be "No" - even though subjects were introduced first to the unconstrained tasks in order to learn the basic experimental paradigm (i.e., synchronized point-to-point movements), they were introduced to the dynamic environments during their first or second experimental session and performed roughly as many movements for each dynamic environment task as for each unconstrained movement task. It is true that the total number of unconstrained movement tasks was greater than the total number of dynamic environment tasks, but for all tasks the changes in the mean achieved trajectory from session to session typically approached zero after two (or maybe three) sessions, which indicates that even for the more complex dynamic tasks, learning of the task had reached a plateau.

The final primary objection is that the error bounds on the attractor trajectory estimates are improperly formulated and so are insufficiently tight to support our claims of having disproved the virtual trajectory hypothesis. The basis for this objection is the use of a linearized model of the muscle impedance to bound the error estimates. Since we have already argued that these impedances are not known precisely enough to allow us to compute a physiologically-justified virtual trajectory, why are we justified in using them to produce error bounds on our attractor trajectory estimates? We are justified in so using them because of the evidence of a number of studies which show that the stiffness of the arm during movement is a substantial fraction of its stiffness during posture [Won 93, Milner 93, Gomi 92, Bennett 92]; even Bennett shows that the minimum stiffness is on the order of 50% of the postural value. Since we are unlikely to find the arm exhibiting impedances extremely different from the postural values, it is reasonable to use ranges based on postural values to bound our attractor trajectory estimate. Furthermore, we use these impedance-based estimates primarily out of convenience because the attractor trajectory estimate can in fact be bounded by other means, as described in an earlier chapter, although doing so would require much more experimentation. The error is due to two factors. First the average impedance force history may be non-zero because we have not actually tracked along the attractor trajectory. This error can be reduced to the intrinsic variability in the execution of the task through the iterative convergence procedure. If this variability is sufficiently small, we could regard the associated error as being practically nil. The second source of error is due to errors in the inertial estimate. We could repeat the convergence procedure for a range of possible inertial values to get a sense of the error bounds on the attractor trajectory estimate. In simulations, this has proven to yield much tighter bounds than the impedance-based estimates we have actually used. The reason for the extreme conservativeness of the impedance-based error estimates is that they ignore the structure inherent in variations in inertial parameters; the model-based estimates assume that the inertial parameters can vary arbitrarily rapidly throughout the movement, whereas in fact they must remain constant (except for the dependence on the relative angle between the forearm and the upper arm). For these reasons, as well as those outlined in an earlier chapter, the attractor trajectory error bounds are likely to be quite conservative.

So, the major results of the study are relatively well-established. In order to contradict the finding that the attractor trajectories are in general more complex and variable than the actual trajectories, one would have to make at least one of the following claims: (1) that the inertial model is so flawed that even the large assumed error bounds do not adequately bound the possible range of behaviours of the actual limb, (2) that the "Do not intervene" paradigm does not yield believable estimates even over time scales shorter than those obtained from a voluntary response test, (3) that learning in the dynamic environments is incomplete, or (4) that the attractor trajectory error estimates are excessively tight. Since any of these claims would be difficult to support, we turn now to a consideration of the implications of the results of the experiments.

Conclusions

One of the primary results of this study is in agreement with the recent results of Won and of Milner, namely, that the arm exhibits significant stability throughout its movement [Won 93, Milner 93]. This is apparent from three pieces of evidence. First, for the unconstrained movements, the ratio of the rise time of the attractor trajectory to that of the actual trajectory is approximately constant across speeds ranging from 30-80% of the maximum speed subjects can achieve. If stiffness remained essentially constant as speed varied, we would expect to see the attractor trajectory rise comparatively more quickly than the actual trajectory with increasing speed. Since we do not observe this, and since the general shape of the attractor trajectory changes little with speed, we conclude that the limb's stiffness must increase with speed. This conclusion is supported by the observation that the temporal lead of the attractor trajectory over the actual trajectory actually decreases at higher speeds; again, if stiffness did not rise with speed, we would expect the opposite result. The third and final piece of evidence of stability during movement is the fact that the iterative algorithm for identifying attractor trajectories converges for all the tasks studied. This algorithm depends on the existence of a stable limb impedance for convergence, so failure to converge over certain portions of the trajectory would indicate regions of instability and conversely success in converging over the entire trajectory indicates stability throughout the trajectory.

Although it is apparent that the arm is stable throughout the entire movement, the basis for this stability is still unclear. Crago argues that the static muscle properties are not sufficient to guarantee stability of the elbow joint throughout the entire range of motion and so suggests that any stability that exists is due to some combination of reflex activity and the so-called short-range stiffness (attributable to forces set up in the crossbridges during stretches short enough that the crossbridges are not broken [Crago 94]; this effect disappears over a time course of seconds or tens of seconds as the crossbridges cycle through their detachment-reattachment process and so it is not apparent on the static length-tension curve). However, Mussa-Ivaldi's, Shadmehr's, and Dolan's measurements are all based on perturbations large enough to get beyond the range of the short-range stiffness effect and they all indicate stable static stiffness behaviour [Mussa-Ivaldi 85, Shadmehr 93, Dolan 91]. Co-contraction is a likely source of enhanced stiffness at higher speeds; when muscle is activated, it takes a relatively long time to relax, so at higher speeds when the agonists and antagonists are activated in close succession, both may in fact be active simultaneously. Since the stiffnesses of these muscles add, the overall stiffness is enhanced during fast movements.

A second important conclusion is that the actual trajectories at different workspace locations are in fact more similar to one another than would be expected if the brain ignored dynamic effects. Given that some sort of compensation is clearly possible, it is interesting to enquire why the compensation is not complete or perfect. In other words, if the general tendency of the subjects is to produce relatively straight paths in cartesian space, why do they not produce nearly perfectly straight lines at all locations in the workspace? Is it that the dynamic compensation ignores certain terms in the equations of motion, or is there some other pattern to the deviations? Cruse suggests that the reason could be quite complex and perhaps unrelated to dynamics at all [Cruse 87]. He found that subjects who were asked to make point-to-point movements using a long stick (approximately 1 m long) did not move the tip of the stick in a straight line, even when the movement duration was on the order of 2 s so that the dynamic effects would be negligibly small. Subjects are, however, capable of generating much straighter paths when requested to do so, although in our experiments straightness was neither required nor encouraged, nor were subjects given any explicit feedback at all about the trajectory they followed between the starting and ending points. The choice of non-straight paths therefore is likely to contain useful information concerning the nature of the motor plan.

What then are we to make of the results of the experiments? They do not unambiguously support either the proponents or the opponents of the virtual trajectory hypothesis, nor do they vindicate those who feel that the entire discussion of the hypothesis is irrelevant to a proper understanding of motor control. The proponents of the virtual trajectory idea will have to come to terms with the fact that the attractor trajectories in unconstrained motions are more complex and variable than are the actual trajectories, particularly as the location of the task in the workspace is varied (which indicates that there is at least some workspace-dependent compensation for dynamic effects). On the other hand, the opponents of the idea may be correct, but not for the reasons they have argued. The limb does exhibit stable behaviour with impedances high enough to significantly affect the movement, particularly in response to perturbations. Those who have argued that the virtual trajectory must necessarily be far from the actual trajectory or must execute a considerably more complicated path will have to deal with the observation that there is a close correspondence in shape between the attractor and actual trajectories. Finally, the relative similarity of attractor trajectories used to control movements performed under different dynamic conditions demonstrates that the endpoint trajectory is not always the object of the optimization underlying the generation of the motor plan. It appears that the central nervous system has the ability to roughly compensate for routine variations in inertial loads as a function of the location of the task, but that at a given location the basic motor plan is simply scaled in time (with perhaps some impedance modifications such as cocontraction) to handle a wide range of execution speeds and dynamic environments. Since motor tasks are often learned in stages, it may be that this scaled plan is only the first means the subjects have found to accomplish the task and that with considerably more practice they will develop new strategies for performing these more complex tasks. Even so, the connectionists will have to explain why this is the particular form that the learning algorithm takes.

The Importance of Stability for Learning

It is sometimes not appreciated by the motor control community how important both isolated and coupled stability are, not only for the execution of movements and performance of motor tasks, but also for the learning of such activities. Indeed, even among roboticists, this recognition is only now developing. Historically, robots have been used as high precision placing or tracking machines and so have been designed to reject or overcome any disturbances. In practice, this philosophy requires that robots be designed to be stiff so that interactions with the environment do not displace them from their commanded paths. Such approaches work well when there are no unexpected obstacles in the path, but when there are these types of robots often exhibit unstable and self-destructive behaviour (e.g., the well-known contact instability problem). Only fairly recently has this behaviour been understood theoretically [Colgate 88]; the primary result is that the requirement that the robot remain stable when coupled to a second system forces us to explicitly consider the mechanical impedance which the robot exhibits at its interaction ports and requires that the impedance be shaped so that the overall interaction can remain stable. In particular, to guarantee stable interaction with arbitrary passive environments, the manipulator must itself exhibit a passive impedance. The fact that the planar stiffness matrices measured by Mussa-Ivaldi (and independently by Shadmehr and Dolan) have essentially zero curl shows that at least the static component of our impedance is passive and suggests that we are intrinsically capable of stable interaction with a broad array of physical systems, including the whole class of passive systems.

What is perhaps not so well known is that stable impedances are required for motor learning. For example, suppose muscle activity were directly controlled by the brain and it was trying to learn what muscle activities would produce a desired trajectory. If the limb were intrinsically unstable, then any slight computational error or interaction force would generate a growing divergence between the intended and actual trajectories that would be fundamentally unrelated to the computation of the nominal muscle activities. In such a context, how could the brain interpret the error signals? It would have no way of knowing whether to attribute the errors to noise or computational error, so learning would be fundamentally impossible. As our experience with neural networks has developed, we have learned that such networks must remain stable throughout the learning process, and so too must the system that we are trying to control. It is, of course, possible that the system be unstable mid-movement if the regions of stability at either end of the movement are large enough to catch any divergence which occurs mid-movement, although in such a situation one cannot be said to be truly controlling the trajectory. An example of this strategy in natural movements is the snatch-and-jerk maneuver in weight-lifting in which the lifter initially imparts upward momentum to the weights and uses the deceleration time to drop below the weights prior to making the final phase of the lift. This maneuver can only be performed at high speeds since the middle portion of the trajectory is unstable; at low speeds, the time spent in the unstable region of the trajectory is so great that the weight drops to the floor and the maneuver cannot be completed. In our experiments, the finding that the attractor trajectory is essentially invariant when the dynamic environment changes suggests that the entire trajectory is stable and that the trajectory learned for the unconstrained movements serves as the basis or template for learning the new movements.

Future Work

What we have learned from the studies reported in this thesis is that the brain is not performing its computations on solely kinematic quantities, but is somehow taking account of workspace-dependent dynamic effects. However, the motor plans for kinematically similar but dynamically distinct tasks appear to be strongly related. These observations raise interesting questions concerning the purposes of various structures in the brain. For example, it is well known that cerebellar injury has a strong influence on the smoothness or grace of movements. Is it possible that the cerebellum is responsible for computing the workspace-dependent adjustments to the virtual trajectory? If this were so, then we would expect that the higher level representations of the movement would not exhibit these characteristic adjustments. We could test this idea using the methodology of Georgopoulos, who found that the ensemble behaviour of a group of cortical neurons fairly accurately predicted the upcoming direction of movement of the limb [Georgopoulos 83, 86, 88]. One could test if the dynamic adjustments were made at the cortical level by observing the ensemble behaviour of the neurons to see if the predicted movement was most like the preinertial attractor trajectory or the actual trajectory.

An issue not directly addressed in this thesis is how kinematic redundancy is handled by the central nervous system and how the actions of the whole body are coordinated in more natural motor tasks. Solving the kinematic redundancy problem (an excess number of degrees of freedom relative to those required to perform a motor task) is not straightforward. For any desired hand position and orientation, for example, there is an infinite number of limb configurations which can be adopted. Which of these does the brain select? Furthermore, how does the brain choose one of the infinite number of configurational changes which will move the hand along a desired trajectory? The usual engineering approach when dealing with an underdetermined problem is to select an optimization criterion which regularizes the problem and leads to a unique solution. However, Mussa-Ivaldi and Hogan have pointed out that not all local regularizations have desirable global properties [Mussa-Ivaldi 91]. In particular, some conventional optimization criteria, such as the minimization of joint angle changes (which generates the Moore-Penrose pseudoinverse), are not integrable; that is, if one moves the hand in a closed path, updating the limb configuration at each instant according to the pseudoinverse, the limb configurations at the start and the end of the motion will not be the same. Since we observe that humans do in fact return to the same configuration, any regularization criterion which does not predict this is not an adequate description of the human solution to the redundancy problem. Mussa-Ivaldi shows that basing the pseudoinverse on the springlike properties of the muscle does in fact yield an integrable pseudoinverse. Much work remains to be done to both characterize the impedances humans use when performing motor tasks (some work in this area is currently being performed by Rancourt [1994], as well as by Lacquaniti [1993]) and determine to what extent the impedances chosen are under central control.

The experimental results reported here also leave open the question of what is being optimized when movements are planned. Although the actual trajectories are relatively straight, they are not perfectly so, and the attractor trajectories are even more complex. Cruse and Bruwer have also presented results which suggest that the endpoint trajectory as such is not the quantity which drives the optimization (except insofar as the trajectory relates to the goal of the movement) because the trajectory varies significantly with the location of the task in the workspace [Cruse 87]. Perhaps the grail of a global optimization criterion which explains the organization of all motor tasks will forever elude us; the criterion may be task-specific and may not be kinematic or dynamic at all but may instead be related to the difficulty of computation. We must also recognize that ongoing motor plans can be modified by sensory inputs; Cruse's results might well be due to adjustments for discomfort experienced during slow execution of a path which was originally intended to be straight.

Perhaps the most fundamental challenge to the basic picture of motor control I have presented here comes from recent work by Lackner, who studied reaching movements performed by subjects in a rotating room [Lackner 94]. His subjects experienced velocity-dependent perturbations directed perpendicular to their direction of motion (much like the lateral deviations described here) which disappeared when their velocity dropped to zero. Lackner found that subjects who had practiced making the reaching movements in a stationary room produced systematic errors when making the same movements in the rotating room. With practice, his subjects tended to make movements as straight as those in the stationary room, but the final position error persisted. This result challenges our picture that the central nervous system bases its control of movement on the ability to specify an equilibrium configuration for the limb because such a configuration ought to have been reached in Lackner's experiments. Since it was not, his results imply either that our mode of control is not fundamentally based on a set of postural maps or that the particular form of perturbation causes an unnoticed (reflexive?) change in the equilibrium position. There is some support for Lackner's results from both Soechting and Georgopoulos: Soechting found that perturbations applied to the limb during a motion caused a temporal displacement in the ongoing motor plan [Soechting 88], while Georgopoulos found that the ensemble behaviour of the cortical neurons was most closely related to the direction and speed of the upcoming movement rather than to its amplitude or position [Georgopoulos 94]. These results suggest that position as such is perhaps not as fundamental to movement control as is claimed by the virtual trajectory hypotheses. It would be useful to recreate Lackner's test using our apparatus in order to determine whether or not the source of the perturbation has a significant effect on the outcome. In any case, much work needs to be done to reconcile these studies with those showing a strong stability property of the arm.

In closing, I would like to note that the results of Won and of Russell for constrained movement tasks demonstrate that we cannot restrict our analysis to unconstrained movements [Won 93, Russell 90]; we must also explore the full range of interactive behaviour in order to expose our hypotheses to situations which will truly challenge them and point out their failings. Fortunately, in the notion of the attractor trajectory and the approach I have developed to measure it, we now have the tools to explore a representation of more complex tasks which is closer to the central representation than is the observed unperturbed trajectory. We might draw an analogy between visual and motor studies: progress in vision research has proceeded rapidly because we have had the tools to control the presentation of stimuli; with our programmable manipulandum, we now have a similar capability in the realm of motor control and we ought to look forward to what we can learn about how we perform multiphasic motor tasks in complex dynamic environments.

REFERENCES

Abend W, Bizzi E and Morasso P (1982) Human arm trajectory formation. *Brain* 105:331-48

Atkeson CG & Hollerbach JM (1985) Kinematic features of unrestrained vertical arm movements. *Journal of Neuroscience* 5:2318-30

Bennett DJ (1990) The control of human arm movement: models and mechanical constraints. Doctoral dissertation, Department of Brain and Cognitive Sciences, MIT

Bennett DJ, Hollerbach JM, Xu Y and Hunter IW (1992) Time-varying stiffness of human elbow joint during cyclic voluntary movement. *Experimental Brain Research*

Berkinblit MB, Feldman AG & Fukson OI (1986) Adaptability in innate motor patterns and motor control mechanisms. *Behavioral and Brain Sciences* 9:585-638

Bizzi E (1993) Intermediate representations in the formation of arm trajectories. *Current Opinion in Neurobiology* 3:925:931

Bizzi E, Accorneuro N, Chapple W, & Hogan N (1982) Arm trajectory formation in monkeys. *Experimental Brain Research* 46:139-43

Bizzi E, Accorneuro N, Chapple W, & Hogan N (1984) Posture control and trajectory formation during arm movement. *Journal of Neuroscience* 4:2738-44

Bizzi E, Hogan N, Mussa-Ivaldi FA & Giszter S (1992) Does the nervous system use equilbirium-point control to guide single and multiple joint movements? *Behavioral and Brain Sciences* 15(4):603-13

Bullock D and Grossberg S (1991) Adaptive neural networks for control of movement trajectories invariant under speed and force rescaling. *Human Movement Science* 10:3-53

Colgate JE and Hogan N (1988) Robust control of dynamically interacting systems. *International Journal of Control* 48:65-88

Crago PE (1994) Muscle stiffness and neuroprosthetic control of the upper extremity. *2nd World Congress of Biomechanics*, Amsterdam, the Netherlands

Cruse H & Bruwer M (1987) The human arm as a redundant manipulator: the control of path and joint angles. *Biological Cybernetics* 57:137-44

Dohrmann CR, Busby HR & Trujillo DM (1988) Smoothing noisy data using dynamic programming and generalized cross-validation. *ASME Journal of Biomechanical Engineering* 110:37-41

Dolan JM (1991) An investigation of postural and voluntary human arm impedance control. Doctoral dissertation, Department of Mechanical Engineering, Carnegie Mellon University

Dornay M, Uno Y, Kawato M & Suzuki R (1992) Simulation of optimal movements using the minimum-muscle-tension-change model. In: Advances in Neural Information Processing Systems 4, RP Lippmann, JE Moody & DS Touretzky (Eds.), 627-634, Morgan Kaufmann

Fasse ED (1992) On the use and representation of sensory information of the arm by robots and humans. Doctoral dissertation, Department of Mechanical Engineering, MIT

Fasse ED, Hogan N, Gomez SR & Mehta NR (1994) A novel variable mechanical-impedance electromechanical actuator. *ASME Winter Annual Meeting*

Faye IC (1986) An impedance controlled manipulandum for human movement studies. Master's thesis, Department of Mechanical Engineering, MIT

Feldman AG (1966) Functional tuning of the nervous system during control of movement or maintenance of a steady posture. II. Controllable parameters of the muscle. *Biophysics* 11:565-78

Feldman AG (1976) Control of postural length and force of a muscle: Advantages of the central co-activation of alpha and gamma static motoneurons. *Biophysics* 21:187-89

Feldman AG (1986) Once more on the equilibrium-point hypothesis (lambda model) for motor control. *Journal of Motor Behavior* 18(1):17-54

Flanagan JR, Ostry DJ and Feldman AG (1993) Control of trajectory modifications in target-directed reaching. *Journal of Motor Behavior* 25(3):140-52

Flash T (1987) The control of hand equilibrium trajectories in multi-joint arm movements. *Biological Cybernetics* 57:257-74

Flash T and Hogan N (1985) The coordination of arm movements: an experimentally confirmed mathematical model. *Journal of Neuroscience* 5:1688-1703

Friedland B (1986) Control System Design: An Introduction to State-Space Methods. McGraw-Hill

Georgopoulos AP & Ashe J (1994) New concepts in generation of movement. Publication status unknown

Georgopoulos AP, Caminiti R, Kalaska JF & Massey JT (1983) Spatial coding of movement: a hypothesis concerning the coding of movement direction by motor cortical populations. *Experimental Brain Research (suppl.)* 7:327-36

Georgopoulos AP, Kettner RE & Schwartz AB (1988) Primate motor cortex and free arm movement to visual targets in three-dimensional space. II. Coding of the direction of movement by a neuronal population. *Journal of Neuroscience* 8:2928:2937

Georgopoulos AP, Schwartz AB & Kettner RE (1986) Neuronal population coding of movement direction. *Science* 233:1357-40

Giszter SF, Mussa-Ivaldi, FA and Bizzi E (1991) Equilibrium point mechanisms in the spinal frog. In: Visual Structures and Integrated Functions, MA Arbib and JP Ewert (Eds), Plenum Press

Giszter SF, Mussa-Ivaldi, FA and Bizzi E (1993) Convergent force fields organized in the frog's spinal cord. *Journal of Neuroscience* 13:467-91

Gomi H, Koike Y & Kawato M (1992) Human hand stiffness during discrete point-to-point multi-joint movement. *Proceedings of the IEEE Engineering in Medicine and Biology Conference*, Paris, France

Gottlieb GL (1993) A computational model of the simplest motor program. *Journal of Motor Behavior* 25(3):153-61

Granit R (1970) The Basis of Motor Control, Academic Press

Hasan Z (1992) Is stiffness the mainspring of posture and movement? *Behavioral and Brain Sciences 15* (4):756-8

Hasan Z & Karst GM (1989) Muscle activity for initiation of planar, two-joint arm movements in different directions. *Experimental Brain Research* 76:651-5

Hodgson AJ (1990) Interaction strategies for positioning an object. *IEEE Engineering in Medicine and Biology Conference*, Philadelphia, PA, 1919-20

Hodgson AJ (1994) Considerations in applying dynamic programming filters to the smoothing of noisy data. *ASME Journal of Biomechanical Engineering*, appearing in November

Hogan N (1984) An organizing principle for a class of voluntary movements. *Journal of Neuroscience* 4:2745-54

Hogan N (1985) The mechanics of multi-joint posture and movement. *Biological Cybernetics* 52:315-31

Hogan N, Bizzi E, Mussa-Ivaldi FA and Flash T (1987) Controlling mult-joint motor behavior. *Exercise and Sports Sciences Reviews* 15:153-90

Hollerbach JM, and Atkeson CG (1987) Deducing planning variables from experimental arm trajectories: pitfalls and possibilities. *Biological Cybernetics* 56:279-92

Hollerbach JM & Flash T (1982) Dynamic interactions between limb segments during planar arm movement. *Biological Cybernetics* 44:67-77

Jeannerod M (1988) The Neural and Behavioural Organization of Goal-Directed Movements. Oxford Psychology Series No. 15, Clarendon Press

Karst GM & Hasan Z (1990) Direction-dependent strategy for control of multi-joint arm movements. In: Multiple Muscle Systems: Biomechanics and Movement Organization, JM Winters & SL-Y Woo (Eds), Springer-Verlag

Katayama M & Kawato M (1993) Virtual trajectory and stiffness ellipse during multijoint arm movement predicted by neural inverse models. *Biological Cybernetics* 69:353-62

Kawato M, Katayama M, Gomi H & Kioke Y (1992) Coordinated arm movements: Virtual trajectory control hypothesis and learning inverse models. *Proceedings of the International Symposia on Information Sciences*, Iizuka Kyusyuu, Japan

Kawato M (1990) Computational schemes and neural network models for formation and control of multijoint arm trajectory. In: Neural Networks for Control, WT Miller, RS Sutton and PJ Werbos (Eds), MIT Press

Lackner JR & Dizio P (1994) Rapid adaptation to coriolis force perturbations of arm trajectory. *Journal of Neurophysiology* 72, in press

Lacquaniti F (1989) Central representation of human limb movement as revealed by studies of drawing and handwriting. *Trends in Neuroscience* 12(8):287-91

Lacquaniti F, Carrozzo M & Borghese NA (1993) Time-varying mechanical behavior of multijointed arm in man. *Journal of Neurophysiology* 69(5):1443-1464

Lanman JM (1980) Movement and the mechanical properties of the intact human elbow joint. Doctoral dissertation, Department of Psychology, MIT

Latash ML, and Gottlieb GL (1991) An equilibrium-point model for fast single-joint movement. I. Emergence of strategy-dependent EMG patterns. *Journal of Motor Behavior* 23:163-77

Latash ML (1993) Virtual trajectories and joint stiffness reconstructed with smooth perturbation technique under different instructions. *International Society of Biomechanics*, Paris, France, 758-9

Lukashin AV & Georgopoulos AP (1993) A dynamical neural network model for motor cortical activity during movement: population coding of movement trajectories. *Biological Cybernetics* 69:517-24

Mansfield JM (1992) Functional assessment of amputee performance. Doctoral dissertation, Department of Mechanical Engineering, MIT

McIntyre J and Bizzi E (1992) Servo models for the biological control of movement. *Journal of Motor Behavior* (in press)

McIntyre J (1990) Utilizing elastic system properties for the control of posture and movement. Doctoral dissertation, Department of Brain and Cognitive Sciences, MIT

Merton PA (1953) Speculations on the cervocontrol of movement. In: The Spinal Cord, GEW Wolstenholme, (Ed), Churchill

Milner TE & Ijaz MM (1990) The effect of accuracy constraints on three-dimensional movement kinematics. *Neuroscience* 35(2):365

Milner TE (1993) Dependence of elbow viscoelastic behavior on speed and loading in voluntary movements. *Experimental Brain Research* 93:177-80

Morasso P (1981) Spatial control of arm movements. *Experimental Brain Research* 42:223-7

Morasso P & Mussa-Ivaldi FA (1982) Trajectory formation and handwriting: A computational model. *Biological Cybernetics* 45:131-42

Mussa-Ivaldi FA & Hogan N (1991) Integrable solutions of kinematic redundancy via impedance control. *International Journal of Robotics Research* 10:481-91

Mussa-Ivaldi FA, Hogan N and Bizzi E (1985) Neural, mechanical and geometric factors subserving arm posture in humans. *Journal of Neuroscience* 5:2732-43

Mussa-Ivaldi FA, Morasso P and Zaccaria R (1988) Kinematic networks. A distributed model for representing and regularizing motor redundancy. *Biological Cybernetics* 60:1-16

Oshima T & Kumamoto M (1994) Mechanism and control of robot arm with bi-articular muscles. *2nd World Congress of Biomechanics*, Amsterdam, the Netherlands

Poggio T & Girosi F (1990) A theory of networks for learning. *Science* 247:978-82

Press WH, Teukolsky SA, Vetterling WT & Flannery BP (1992) Numerical Recipes in C: The Art of Scientific Computing, 2nd Edition. University of Cambridge Press

Rancourt D (1994) The role of arm posture in power tool use. *2nd World Congress of Biomechanics*, Amsterdam, the Netherlands

Russell DL (1990) An analysis of constrained motions in manipulation. Doctoral dissertation, Department of Mechanical Engineering, MIT

Saund E (1984) Estimation of equilbirium position, stiffness and viscosity during limb movement. Research paper, Department of Psychology, MIT

Schwartz AB (1993) Motor cortical activity during drawing movements: population representation during sinusoid tracing. *Journal of Neurophysiology* 70:28-36

Shadmehr R (1991) A computational theory for control of posture and movement in a multi-joint limb. Doctoral dissertation, Center for Neural Engineering, University of Southern California

Shadmehr R & Mussa-Ivaldi FA (1994) Adaptive representation of dynamics during learning of a motor task. *Journal of Neuroscience* 14:3208-3224.

Shadmehr R, Mussa-Ivaldi FA and Bizzi E (1993) Postural force fields of the human arm and their role in generating mult-joint movements. *Journal of Neuroscience* 13(1):45-62

Slotine J-JE & Li W (1991) Applied Nonlinear Control. Prentice-Hall

Smith AM & Humphrey DR (1991) What do studies of specific motor acts such as reaching and grasping tell us about the general principles of goal-directed motor behavior? In: Motor Control: Concepts and Issues, DR Humphreys & HJ Freund (Eds), John Wiley and Sons, 357-81

Soechting JF (1988) Effect of load perturbation on EMG activity and trajectories of pointing movements. *Brain Research* 451:390-6

Stein RB (1974) The peripheral control of movement. *Physiological Reviews* 54:215-43

Uno Y, Kawato M & Suzuki R (1989) Formation and control of optimal trajectory in human multijoint arm movement — Minimum torque-change model. *Biological Cybernetics* 61:89-101

Vereijken B, van Emmerik REA, Whiting HTA & Newell KM (1992) Free(z)ing degrees of freedom in skill acquisition. *Journal of Motor Behavior* 24(1):133-42

Wolpert DM, Ghahramani Z & Jordan MI (1994a) Perceptual distortion contributes to the curvature of human reaching movements. *Experimental Brain Research* 98:153-156

Wolpert DM, Ghahramani Z & Jordan MI (1994b) Are arm trajectories planned in kinematic or dynamic coordinates? An adaptation study. submitted to *Experimental Brain Research*

Won J (1993) The control of constrained and partially constrained arm movements. Master's thesis, Department of Mechanical Engineering, MIT

Wu C-H, Houk JC, Young K-Y & Miller LE (1990) Nonlinear damping of limb motion. In: Multiple Muscle Systems: Biomechanics and Movement Organization, JM Winters & SL-Y Woo (Eds), Springer-Verlag

Xu Y, Bennett DJ, Hollerbach JM and Hunter IW (1991) An airject actuator system for identification of the human arm joint mechanical properties. *IEEE Transactions on Biomedical Engineering* 38:1111-22

Young JW, Chandler RF, Snow CC, Robinette KM, Zehner GF & Lofberg MS (1983) Anthropometric and mass distribution characteristics of the adult female. FAA Technical Report FAA-AM-83-16