Processing Strategies for
Functional Magnetic Resonance Imaging Data Sets

by

LUIS CARLOS MAAS III

S.M. Electrical Engineering and Computer Science, 1994
S.B. Computer Science and Engineering, 1992
S.B. Electrical Engineering, 1992
Massachusetts Institute of Technology

Submitted to the Harvard University - Massachusetts Institute of Technology
Division of Health Sciences and Technology in Partial Fulfillment
of the Requirements for the Degree of

DOCTOR OF PHILOSOPHY

in Medical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Februrary 1999

/

Signature of Author: _____
Harvard-M.I.T. Division of Health Sciences and Technology
December 7, 1998

Certified by: _____
Perry F. Renshaw, M.D., Ph.D.
Associate Professor of Psychiatry, Harvard Medical School
Thesis Supervisor

Accepted by: _____
Martha L. Gray, Ph.D.
J. W. Kieckhefer Associate Professor of Electrical Engineering
Co-Director, Division of Health Sciences and Technology

Processing Strategies for
Functional Magnetic Resonance Imaging Data Sets

by

LUIS CARLOS MAAS III

Submitted to the Harvard University - Massachusetts Institute of Technology
Division of Health Sciences and Technology on December 7, 1998 in Partial Fulfillment
of the Requirements for the Degree of Doctor of Philosophy in Medical Engineering

ABSTRACT

In recent years, functional magnetic resonance imaging (FMRI) has become a valuable new research tool for the study of the human brain. With this imaging modality, correlates of brain activity are rapidly and noninvasively measured, yielding a combination of temporal and spatial resolution unmatched by earlier functional imaging techniques. To accomplish this, however, FMRI studies generate large time series of images which must be carefully processed in order to reliably detect and quantify changes in brain activity.

A critical step in the processing of functional MRI data is the minimization of brain motion artifacts. While motion artifacts within a single frame have been largely eliminated by high speed imaging techniques, functional MRI experiments remain highly susceptible to artifacts from frame-to-frame motion, as even subpixel shifts may result in false activation artifacts and reduced sensitivity. To address this problem, a novel, high-throughput, two-dimensional registration algorithm is presented. Unlike prior iterative image registration algorithms based primarily in the spatial domain, the method described estimates displacement parameters in the frequency domain, allowing the decoupling of rotations and translations and a fast, non-iterative design. Performance data and an extension of the technique into three dimensions are also presented, and complications of three-dimensional registration are discussed.

New artifacts are introduced into FMRI data sets during image registration, however the nature of these artifacts and techniques to reduce them have not been previously described. A new framework for the discussion of registration-related artifacts is developed, and implications of these artifacts in subsequent data analysis are explored. Novel filtering methods in both the spatial and temporal domains intended to reduce these artifacts are introduced, and data presented demonstrating their effectiveness. Finally, a unified set of tools to analyze functional MRI data integrating the methods developed is described, and experimental data illustrating the utility of these techniques in clinical research presented.

Thesis Supervisor: Perry F. Renshaw, M.D., Ph.D.
Title: Associate Professor of Psychiatry, Harvard Medical School

ACKNOWLEDGMENTS

L.C.M.

BIOGRAPHICAL NOTE

The author, originally from southeastern Michigan, has been living in the Cambridge and Boston area for ten years. He first came to M.I.T. as an undergraduate in 1988, receiving Bachelor's degrees in Computer Science and in Electrical Engineering in 1992. His undergraduate thesis research on computer assisted readout of DNA sequencing gels was supervised by Professor David Gifford. At M.I.T., he was elected to Tau Beta Pi, Eta Kappa Nu, and Sigma Xi.

In the fall of 1992, he joined the Ph.D. program in Medical Engineering and Medical Physics of the Harvard-M.I.T. Division of Health Sciences and Technology, and began research in the Noninvasive Cardiology Laboratory at the Brigham and Women's Hospital under the supervision of Professor Richard Lee. In 1993, he was accepted into the M.D. Program at the Harvard Medical School, beginning his M.D.-Ph.D. career. That year he was also selected to receive the Whitaker Foundation Graduate Fellowship in Biomedical Engineering and the National Science Foundation Graduate Fellowship in Engineering. He served as a Whitaker Graduate Fellow from 1993 to 1998. In 1994, he completed his Master's thesis in Electrical Engineering and Computer Science, entitled *Statistical Methods in Ultrasonic Tissue Characterization.*

In 1994, he began research at the Brain Imaging Center at the McLean Hospital, under the direction of Professor Perry Renshaw, which later evolved into this thesis. He completed his preclinical coursework at Harvard and his doctoral qualifying exams in the M.I.T. Department of Electrical Engineering in 1995. After completing his experimental work, he returned to the clinics in the fall of 1997 to complete his final two years of medical school and the writing of this thesis. He plans to receive both his M.D. and Ph.D. degrees in 1999, after which he intends to complete an internship in internal medicine and a residency in diagnostic radiology.

TABLE OF CONTENTS

TABLE OF FIGURES

7

1 GENERAL BACKGROUND

In the last decade, clinical magnetic resonance imaging (MRI) has become a key component of the medical diagnostic arsenal. MRI provides clinicians with high resolution images of the human body in a noninvasive manner and without the use of ionizing radiation. With resolution approaching that of X-ray techniques and increased sensitivity to differentiate soft tissues, MRI has become the imaging modality of choice for soft tissue structures such as the nervous system. A full understanding of the human body, however, requires not only information about structure, but also about function. Until recently, imaging of human brain function was limited to the domain of radionuclide techniques such as positron emission tomography (PET) and single photon emission computed tomography (SPECT), and surface encephalography techniques such as electroencephalography (EEG) and magnetoencephalography (MEG). Today, these modalities are joined by the emerging field of functional magnetic resonance imaging (FMRI). With spatial and temporal resolution superior to radionuclide techniques and spatial resolution superior to electro- and magnetoencephalography techniques, FMRI is providing investigators with a new window into human physiology.

1.1 MRI

This section briefly describes the principles of magnetic resonance, conventional structural MRI, and the ultra-fast echo planar MRI techniques used in functional MRI, and introduces some of the terminology used in the field of functional magnetic resonance imaging. For a detailed review of these topics, the reader is referred to works of Abragam (1989) and Aine (1995).

8

1.1.1 Basic NMR Physics

Magnetic resonance imaging is based upon the natural phenomenon of nuclear magnetic resonance (NMR). Briefly, nuclei with an odd number of protons or neutrons possess a non-zero nuclear spin, which gives rise to a magnetic moment. In human imaging, relevant elements which have been studied with NMR techniques include ^{1}H, ^{6}Li, ^{11}B, ^{13}C, ^{17}O, ^{19}F, ^{23}Na, and ^{31}P. Of these, ^{1}H possesses both the largest magnetic moment and physiologic abundance. When a hydrogen atom is placed in a large magnetic field $\mathbf{B_0}$ aligned with the z axis, the z component of the nuclear spin occupies one of two quantum energy states. The energy difference depends on both the gyromagnetic ratio (γ) of ^{1}H (an inherent property) and the field magnitude B_0. The resonant frequency corresponding to this energy difference is known as the Larmor frequency.

In the classical physics model, the magnetic moments of the individual nuclei precess about the z axis at the Larmor frequency. At equilibrium, the nuclear moments are randomly distributed around the z axis such that the net magnetic moment $\mathbf{M_0}$ is aligned with $\mathbf{B_0}$. To generate a measurable signal, a second magnetic field $\mathbf{B_1}$ is applied perpendicular to $\mathbf{B_0}$ and rotating around the z axis at the Larmor frequency. As this is equivalent to a constant perpendicular field in the rotating frame, the resultant torque on $\mathbf{M_0}$ brings it off the z axis such that $\mathbf{M_0}$ also precesses about the z axis at the Larmor frequency. As its projection onto the x-y plane thus represents a net time-varying magnetic field, a voltage can be induced in a nearby receiver coil, allowing measurement of the NMR signal. In practice, a $\mathbf{B_1}$ pulse is generated by applying an oscillating current at the Larmor frequency to a transmitting coil. Because the Larmor frequency of hydrogen is in the radio-frequency range for clinical NMR field strengths, the $\mathbf{B_1}$ pulse is often called an RF pulse and its application called an RF excitation.

The NMR phenomenon has been exploited since the 1940s in chemical spectroscopy to determine the molecular makeup of unknown compounds by mapping out the small variations in the Larmor frequency caused by differences in the local molecular environments of the nuclei. Relaxation mechanisms which attenuate the NMR signal can also be measured or exploited as sources of contrast. The two principal relaxation mechanisms of relevance in this thesis are described by the exponential time constants T_1 and T_2. T_1 measures the recovery of the net magnetic moment along the z axis back to its resting state. T_2 describes net signal loss due to the dephasing of individual magnetic moments as they precess around $\mathbf{B_0}$.

1.1.2 Magnetic Resonance Imaging

Because of the natural abundance of hydrogen in living tissue, conventional MRI is based on the hydrogen proton NMR signal. In the early 1970s, researchers successfully applied magnetic field gradients to provide the spatial localization of the proton NMR signal necessary to generate useful images of anatomy. In the clinical context, the phrase "nuclear magnetic resonance" is abbreviated to "magnetic resonance" (MR).

The application of magnetic field gradients alters the precessional frequency of the spins in a spatially dependent manner. By applying a linear "read-out" gradient concurrent with sampling of the MR signal at the receiver, the sampled proton signal is spread out over some bandwidth centered at the Larmor frequency. Thus, the signal amplitude at each frequency corresponds to the proton signal from a unique position along the gradient. After modulating out the Larmor "carrier" frequency, the received signal is equivalent to a Fourier encoding of one-dimensional spatial projection of the proton signal. A second dimension can be added by applying a linear gradient in a direction perpendicular to that of the read-out gradient prior to the onset of sampling. This "phase-encode" gradient imparts a phase on the spins proportional to the position in the second direction. If data are collected following multiple excitations, each with a different phase-encode gradient strength, a Fourier encoding of the MR signal in the second direction results. By combining phase-encode gradients before data sampling with read-out gradients during data sampling, a two-dimensional Fourier encoding is obtained. The final spatial dimension can be added similarly with a second phase-encode gradient perpendicular to both the read-out and first phase-encode gradients or through the use of a "slice-select" gradient. In the latter method, a linear gradient is applied *during* each initial RF excitation. With a spatially dependent Larmor frequency, a narrowband RF pulse is used to selectively excite a "slab" of spins in the desired slice, with the bandwidth a gradient strength determining the thickness of the slice. This direct (non-Fourier) encoding in the third dimension is then combined with the conventional two-dimensional read-out and phase-encode gradients described above.

In conventional imaging, the discrete time samples collected following each RF excitation are referred to as k-space samples, indicating their relation to the spatial frequencies kx, ky, and possibly kz. In the common imaging approach described above, each sampling yields data along a "line" of k-space, reflecting the linear frequency encoding provided by the read-out gradient. With multiple excitations, each followed by appropriate phase-encode and read-out gradients, a two- or three-dimensional matrix of k-space samples can be filled. The completed data matrix is referred to as the "raw" or "k-space" data, and typically, a magnitude image is generated from this raw data by inverse Fourier

transformation. Thus, it is the Fourier encoding of the desired image which is sampled directly by MR imaging techniques and not the image itself, a point which will become important later in this thesis.

In conventional MRI, each phase encode step requires a separate RF excitation as T_2 relaxation limits the useful time available for data sampling. Furthermore, T_1 relaxation limits the pulse repetition time (T_R) since the amplitude of the MR signal is proportional to the magnitude of the net magnetic moment aligned with the field at the time of excitation. If T_R is made too short, too little time is available for signal recovery, and saturation effects reduce the measurable MR signal. As the total time required for image acquisition is equal to T_R times the number of excitations, the acquisition of a complete image, which may have 256 or more phase encode steps, may take several minutes. While a number of variations on conventional techniques allowing shorter imaging times have and continue to be developed, only echo planar imaging, the method most commonly associated with functional magnetic resonance imaging, will be described here. For a recent review of other magnetic resonance methods, the reader is referred to the work of Bushberg (1994).

1.1.3 Echo Planar Imaging

Recent advances in hardware design have allowed the fabrication of faster, more powerful read-out gradients which can be reversed sufficiently rapidly to allow, when paired with modified phase-encoding gradients, the sampling of multiple lines of the Fourier data matrix following each RF pulse. With specialized hardware, resonating read-out gradients can be used together with short phase-encode gradient "blips" to acquire a complete two-dimensional k-space matrix following a single RF pulse. This technique, known as echo planar imaging (EPI; Mansfield 1977), can acquire a "snapshot" of a slice of interest in under 100 ms. By allowing the acquisition of an entire plane in each T_R interval, the temporal resolution of the imaging technique is greatly increased, and "movies" of a brain slice can be collected. Additionally, by interleaving the acquisition of multiple non-overlapping slices, full three-dimensional data sets can be acquired in each T_R interval (Jezzard 1996).

The development of the three-dimensional analog of EPI, known as echo volumar imaging (EVI), continues (Song 1994, Mansfield 1995, Harvey 1996, Yang 1997A), but at present these methods yield very limited spatial resolution or small imaging volumes. However, other three-dimensional methods based on echo planar imaging have also been proposed (Abduljalil 1995) which, while not "snapshot" in nature, achieve temporal resolution on the order of 4 to 5 seconds.

1.2 FUNCTIONAL MAGNETIC RESONANCE IMAGING

With the development of echo planar MRI, researchers were able to more easily study dynamic MR effects. An early application of EPI of the brain was the tracking of an intravascular bolus of injected MR contrast agent, a compound which changes local relaxation properties (Belliveau 1990). With this technique, known as dynamic susceptibility contrast (DSC) MRI, a time course of relative contrast concentration can be derived from the sequence of MR image intensities at each spatial position. Integration of the time course at each pixel yields a map of relative cerebral blood volume (CBV), a physiologic parameter whose measurement was previously possible only with radionuclide modalities. Shortly after the initial DSC MRI experiments, investigators were able to demonstrate CBV changes in the visual cortex of humans during photic stimulation (Belliveau 1991). Thus, echo planar CBV measurements provided the first indirect "functional" MR data, i.e., hemodynamic data related to neuronal activation, and mappings of CBV changes between rest and stimulus states yielded the first MR based "activation maps" of the human brain.

Although MR methods to measure neuronal activation directly have not yet been described, other indirect techniques have since been proposed. These methods have recently been reviewed by Jezzard and Song (1996). Like the initial bolus tracking method, these new techniques link neuronal activation to a physiologic parameter which can be measured with MR methods, such as cerebral blood flow (Edelman 1994, Kim 1995, Chen 1997A, Talagala 1998, Yang 1998, Schwarzbauer 1998) and volume (Sorensen 1997, Lev 1997, Mandeville 1998). All of these methods fall into the general category of functional MRI (FMRI). The most successful and dominant FMRI technique, however, is based on relative blood oxygenation levels, and is described below.

1.2.1 Blood Oxygenation Level Dependent FMRI

The magnetic properties of hemoglobin have long been known to depend on its oxygenation state (Pauling 1936). Unlike oxygenated hemoglobin, deoxyhemoglobin is paramagnetic and thus changes local relaxation properties (Thulborn 1982). Shortly after the development of DSC MRI, Ogawa et al. (Ogawa 1990) theorized that deoxyhemoglobin could be used as an endogenous susceptibility contrast agent for in vivo studies. This was verified in early animal experiments which demonstrated reversible brain signal decreases during hypoxia, presumably due to increased blood carbon dioxide levels and increased deoxyhemoglobin levels (Turner 1991). Earlier work with PET in humans indicated that

during neuronal activation, oxygen extraction by neurons increased only slightly, while blood flow into the activated region increased significantly (Fox 1986). This uncoupling of oxygen utilization and delivery leads to a washout of deoxyhemoglobin by oxygenated hemoglobin in activated brain regions.

With MR pulse sequences sensitive to changes in T_2^*, the effective or observed T_2 relaxation rate, this relative decrease in deoxyhemoglobin levels can be translated into a small signal increase in the MR image in these regions when compared to baseline images. Verifying this basic principle, signal increases were successfully detected following neuronal activation in both visual and motor cortices of human subjects (Bandettini 1992; Kwong 1992; Ogawa 1992). This technique, now known as blood oxygenation level dependent (BOLD) FMRI, is often used in studies of regional brain activation although the precise dynamics relating neuronal activation, deoxyhemoglobin washout, and the MR signal remain active areas of basic research (Ogawa 1993, Ernst 1994, Boxerman 1995, Boxerman 1995A, Frahm 1996, Kim 1997, Dymond 1997, Gati 1997, Zhang 1997, Yang 1997, Hoogenraad 1998, Bock 1998, Buxton 1998).

Other work has centered on applications of BOLD techniques. With the increased availability of echo planar imaging hardware and the continuing development of spiral imaging (Ahn 1986, King 1995, Noll 1995, Yang 1996, Glover 1998) and other fast imaging techniques (Merboldt 1995, Gao 1996, Crémillieux 1997, Noll 1998), BOLD FMRI has established itself as a powerful new tool for the mapping of human brain function. In addition to continuing work with visual (Howard 1995, Ross 1997, Chen 1998A) and motor (Noll 1997) paradigms, BOLD FMRI has also been used to study the other senses (Lin 1996, Yang 1997, Gaschler-Markefski 1997, Yang 1997), word production (McCarthy 1993), memory (McCarthy 1994, Noll 1997), and many other cognitive processes (Kwong 1995, Le Bihan 1995, Le Bihan 1996, Yurgelun-Todd 1997). Studies have been conducted in both healthy volunteers and subjects with neuropsychiatric disorders (Renshaw 1994, Wenz 1994, Connelly 1995, Howard 1996, Yurgelun-Todd 1996, Breiter 1996, Maas 1998), as well as in other species (Chen 1996, Yang 1997, Jezzard 1997).

1.2.2 Analysis Techniques for FMRI Data

As will be described in Chapter 2, this thesis will center on preprocessing methods designed to improve the quality of subsequent FMRI data analyses. Thus, a brief overview of current statistical analysis techniques employed in functional MRI is included here, such that these preprocessing techniques can be understood in the proper context.

At 1.5 Tesla, the typical field strength for medical MR scanners, activation-induced signal changes in the sensory cortices are typically 1-5% of the total signal intensity (Kwong 1995, Jezzard 1996), on the order of the inherent signal-to-noise of the imaging techniques. Thus, the analysis of a large number of images is required to detect the underlying BOLD signal changes. Although a very large number of statistical methods have been applied to FMRI data sets, three pixel-wise methods dominate current BOLD FMRI work: difference t-test methods, correlation coefficient analysis, and variants of Fourier analysis.

In difference t-test methods, an averaged baseline image is subtracted from an averaged stimulus image, to create a pixel-by-pixel "difference" image between the baseline and activated states. Bright regions in the resulting difference image correspond to areas of activation. The Student's t-test is typically used to assign a level of statistical significance to the difference at each pixel, providing a more uniform measure of activation, and often a map of corresponding t-values is used in place of the difference image to visualize activation. In experiments with alternating baseline and test conditions, the partition of the resulting image set into control and stimulus image sets is complicated by the poorly understood hemodynamic transfer function describing the relationship between neuronal activation and the blood flow and volume changes which regulate the activation-induced signal changes (Kim 1997A, Frank 1998, Fransson 1998). Lags of 4 to 8 seconds from the onset or termination of the stimulus to the corresponding change in signal intensity are generally observed in FMRI data (Lee 1995, Jezzard 1996), as well as rise and fall times of up to 10 seconds, an early "dip" (Ernst 1994, Dymond 1997), and a relative signal "undershoot" following the cessation of neuronal activation (Jones 1998). These temporal characteristics tend to obscure stimulus-induced differences in t-test based analysis, reducing the sensitivity of this technique for the reliable detection of activation.

An alternative analysis technique is based on the computation of the coefficient of correlation between a pixel's time signal and a reference time function (Bandettini 1992, 1993). In the simplest case of a "box car" reference function, which models observed activation as a simple on-off phenomenon, correlation methods are equivalent to the t-test methods described above. However, with reference functions based on models of the hemodynamic transfer function, the method can make more complete use of the increased temporal resolution available with high speed imaging. The method can be further extended into a multiple linear regression to compensate for undesired "noise" functions, such as linear trends. Unfortunately, the selection of ideal reference and noise waveforms is complicated by (1) spatial variation of signal characteristics, including lag, rise and fall times (Lee 1995, Mitra 1997, Chen 1998), (2) physiologic noise confounds from respiration and circulation

(Hu 1994, Hu 1995, Biswal 1996, Jezzard 1996), and (3) spatially correlated background brain activity (Biswal 1995).

Less widely used analyses include multiple reference waveforms, including partial and full Fourier analyses (Lee 1995, Bullmore 1996, Howard 1996), to provide some immunity to shifts in the hemodynamic response. Nonparametric analyses, such as pixel-wise computation of the Kolmogorov-Smirnov statistic (Xiong 1996), may provide some immunity from nonnormality of temporal noise in functional MRI data sets, however their use is controversial (Aguirre 1998) and nonnormality is not always observed (Moser 1996).

Although the options for data analysis are varied and new methods continue to be described (Forman 1995, Sekihara 1996, Bullmore 1996, Maas 1997B, Baumgartner 1997, Noll 1997, Gaschler-Markefski 1997, Frank 1998) and compared (Constable 1995, Sorenson 1996, Genovese 1997), it is important to note that most methods remain fundamentally based on the analysis of times series from individual pixels. Thus, in all cases, it is imperative that the data collected represent information from identical anatomic positions through time. As will be described in the next Chapter and addressed throughout this thesis, subject motion during the course of functional experiments is a major factor violating this imperative, but one which can, in most cases, be well remedied with proper processing following acquisition and preceding data analysis.

2 SUBJECT MOTION IN FUNCTIONAL MRI

In a recent joint report by the Committee on the Mathematics and Physics of Emerging Dynamic Biomedical Imaging, the National Research Council, and the Institute of Medicine (1996), four important areas for future research on the processing of FMRI data were identified: (1) the detection and compensation of subject motion, (2) the characterization of the functional response pattern, (3) the characterization of physiological noise confounds, and (4) visualization of results. As will be described in this section, this thesis will address the first of these points: subject motion. Both motion itself and attempts at its correction introduce "errors" which contaminate subsequent analysis and visualization. This thesis will explore data preprocessing strategies to attenuate these errors before statistical analysis and visualization, thereby improving the robustness of the functional MRI technique.

2.1 GROSS SUBJECT MOTION

It is widely accepted that one of the major limiting factors in BOLD functional MRI studies is subject motion (Hajnal 1994; Jezzard 1996; Zeffiro 1996; Maas 1997). With FMRI techniques sensitive to motion on the scale of tenths of pixels (Hajnal 1995, Jezzard 1996), it has been estimated that as much as 30-90% of the variation in pixel time course data may be due to motion-related effects. (Friston 1996) These artifacts predominantly result from frame-to-frame motion (Hajnal 1994), as echo planar acquisition of functional images in less than 100 ms has largely eliminated the intra-image motion artifacts seen in conventional imaging (Duyn 1996).

The sources of subject motion are varied. In addition to the inherent difficulties of remaining motionless inside a scanner for long periods of time, physiologic functions such as respiration, circulation, and swallowing can cause subtle position changes. These factors are exacerbated as functional experiments are lengthened, as with newer protocols lasting longer than twenty minutes (Chen 1998). Slow positional drifts and sudden positional changes

increase the underlying variability in pixel time series, thus reducing the sensitivity of FMRI to detect real activation-induced signal changes (Bandettini 1993). Furthermore, the subject is generally exposed to a series of stimuli or asked to perform one or more tasks. These stimuli or tasks may themselves induce subject motion resulting in stimulus-correlated motion artifacts and associated false activation effects during analysis (Hajnal 1994). In a recent study of healthy volunteers utilizing a simple five-minute photic stimulation paradigm, Moser et al. (1996) observed stimulus-correlated motion in as many as 7% of subjects and motion exceeding one half pixel in 11%. Similarly, Hajnal et al. (1994) observed evidence of drift, sudden position change, or stimulus-correlated motion in four of four healthy subjects for both visual and motor tasks, in spite of subject head immobilization with evacuated pillows. Naturally, larger effects may be expected in less cooperative patient populations (Baudendistel 1995). In studies conducted at the McLean Brain Imaging Center with psychiatric patients, subject motion over the course of an experiment ranging over approximately 1 degree of rotation and 1 millimeter of translation are not uncommon. Thus, with combinations of positional drifts and stimulus-correlated motion, subject motion decreases both the sensitivity and specificity of FMRI to detect activation.

2.2 APPROACHES TO LIMIT MOTION ARTIFACTS

Efforts to directly limit subject motion during imaging sessions have included different mechanisms of subject restraint, including securing the head with medical tape, stereotactic frames (Debus 1996), bite bars (Lee 1995), and custom formed pillows and face masks (Hajnal 1994). In spite of these measures, residual motion effects have remained (Hajnal 1994, Zeffiro 1996). In fact, even if the skull were rigidly fixed, the brain would still be free to move within the cerebrospinal fluid (Poncelet 1992). Additionally, many of these constraints are incompatible with the stimuli or tasks required by certain FMRI paradigms are poorly tolerated by subjects. Another approach under development is to avoid motion by locking the scanner's imaging coordinate system to the subject's frame of reference (Lee 1996, Lee 1998). With these techniques, subject motion is estimated in real-time using navigator echoes and the scanner's imaging parameters are updated on a frame-to-frame basis. Unfortunately, these methods require specialized hardware modifications and corrupt the data with saturation bands associated with collecting navigator echoes in planes orthogonal to the imaging plane.

Thus, most investigators have turned to post-acquisition image registration algorithms to realign FMRI data sets before statistical processing. Early algorithms include landmark

(Hill 1991) and surface matching methods (Pelizzari 1989). As the need for constant operator intervention made their application to increasingly large FMRI data sets impossible, these have been largely superseded by more automated methods. One approach which has achieved popularity in FMRI applications is the method of Woods et al. (1992), adapted from PET imaging, which estimates motion by iterative minimization of the variance of image-to-image ratios. However, other solutions continue to be proposed. Regression based minimizations constitute the majority, including least squares (Hajnal 1995, Eddy 1996, Friston 1996), least absolute (Lange 1994), and least median of squares methods (Alexander 1996). In general these iterative methods are intensity-based, i.e. the error function is determined from the pixel-wise difference between the two images to be registered, but the use of derived edge or contour information has also been described (Alexander 1996, Biswal 1997). Other methods include cross-correlation (Baudendistel 1995, Biswal 1997), Fourier-based methods (Maas 1995, Kassam 1996, Maas 1997, Lee 1998), and center-of-mass tracking (Moser 1996).

2.3 LIMITATIONS OF CURRENT REGISTRATION ALGORITHMS

The ongoing research in the area of image registration stems not only from the crucial need to address the motion-related artifacts described above, but also from the continued shortcomings of available methods. The most important of these limitations are described in this section.

First, all of the methods listed above which estimate translation and rotation simultaneously in image-space are iterative in nature and thus may become computationally intensive when applied to large FMRI series. Furthermore, convergence may be dependent on the choice of initial conditions (Alexander 1996, Eddy 1996). Two of the methods cited above have avoided iteration by omitting estimation of rotation altogether (Baudendistel 1995, Moser 1996). However, as rotations are the predominant form of motion observed in functional data sets (Baudendistel 1995, Lee 1998), this is a sub-optimal solution. Existing Fourier-based (Kassam 1996) and non-linear (Alexander 1996, Alexander 1997) algorithms have also found it necessary to employ iteration. Thus, the identification of a noniterative solution would be desirable.

Second, to generate a corrected image, the original "unregistered" image must be resampled in the correct orientation. Many currently described post-processing algorithms, such as the popular technique of Woods et al (1992), construct this new image by linear interpolation of the original image. While this method was originally employed to register

18

lower resolution PET images where such blurring was largely inconsequential, linear interpolation may lead to unnecessary high-spatial-frequency attenuation in higher resolution functional MRI data (Hajnal 1994, Eddy 1996, Maas 1997). Other algorithms have employed higher order interpolation schemes (Bullmore 1996), but these have also been demonstrated to result in undesired blurring (Unser 1995). Thus, a higher fidelity method is desirable to interpolate the images.

Hajnal et al. (1995) used truncated sinc interpolation to generate images in their iterative registration algorithm, but they found that with the extreme computation demands of larger windows significant truncation was necessary to prevent their algorithm from becoming "impossibly slow." One apparent compromise is to retain linear interpolation during parameter estimation because of the high speed with which it can be performed, and then employ a more accurate interpolation once the iterative technique has converged. Unfortunately, linear interpolation of interim images has been demonstrated to lead to errors in the values of iteratively estimated motion parameters, particularly rotation, in registration methods based on least-squares difference minimization (Hajnal 1995, Eddy 1996). This provides another reason to seek noniterative solutions.

Third, as will be discussed in this thesis, registration techniques themselves may generate new temporal artifacts as a result of incomplete resampling, aliasing, and other side-effects of image reconstruction related to time-varying positional changes. This type of error has not been adequately addressed in the FMRI literature.

In light of the current limitations, it is not surprising that in the report from the National Research Council and the Institute of Medicine mentioned above, the detection and compensation of subject motion was described as "one of the most urgent needs in functional imaging." With this in mind, a novel registration method which addresses many of these issues will be presented in Chapter 4. Furthermore, spatial and temporal filtering methods to reduce new artifacts introduced by all registration methods will be presented in Chapter 5.

2.4 THREE DIMENSIONAL REGISTRATION

Three-dimensional registration has been accomplished with some success in conventional anatomic MRI (Bedell 1996) where the precision required is less than that of functional MRI. The design of improved methods for FMRI is complicated by certain aspects of typical FMRI data sets.

The first complicating factor is that in multi-slice FMRI data sets, there does not exist a true reference point in time, since adjacent slices are acquired at different instants in time.

Thus, "intra-volume" motion occurring during the time required to image the full volume can lead to inaccurate registration and image distortions. This problem is exacerbated in multi-shot methods which pair out-of-plane phase-encoding with echo planar techniques (Abduljalil 1995) since any motion-related signal changes necessarily affect the entire volume (Duyn 1996). These limitations may ultimately be overcome by echo volumar imaging (Song 1994, Mansfield 1995, Harvey 1996, Yang 1997A), the three-dimensional analog of echo planar imaging, but at present these techniques do not yet provide the resolution necessary for FMRI.

A second factor complicating three-dimensional registration is that typical voxels in echo planar imaging are not cubic, with the slice thickness greatly exceeding the in-plane dimensions. This is necessary in multi-slice methods, where a sufficiently thick slice, generally at least 5 millimeters, must be excited to provide an adequate signal-to-noise ratio. Thus, data may be severely undersampled in the out-of-plane direction (Noll 1997A). This problem is further exacerbated when inter-slice spacing is used to prevent saturation effects in neighboring slices from imperfect RF pulses. As a consequence, out-of-plane interpolation can be unreliable or unacceptably sacrifice the overall spatial resolution of the data set (Noll 1997A). Multi-shot methods can provide thinner slices (Abduljalil 1995), at the cost of reduced temporal resolution, but are limited by the intra-volume effects described above.

Another factor is that, while a single image slice typically captures all of the edges of the brain and surrounding background space, the out-of-plane "stack" generally does not extend to the edges of the brain, but rather captures only a section of interest, or, in the extreme, only a single slice. This windowing can disrupt accurate estimation of out-of-plane motion parameters, and consequently, three-dimensional realignment as sections of the brain move in and out of the imaging volume. Additionally, the windowing complicates realignment at the ends of the stack, where no information is available about the adjacent brain content. Naturally, a larger volume can be imaged to reduce this problem, however this introduces other problems, as discussed above. Finally, recent work suggests that differential spin histories within a slice introduced by out-of-plane motion can also lead to signal variation artifacts in single- and multi-slice methods (Friston 1996), further complicating the estimation of displacement parameters.

The difficulties of post-acquisition three-dimensional registration in FMRI has prompted development of a straightforward real-time method to provide prospective correction of linear translation in the out-of-plane direction (Lee 1996). This method can also be extended to the more important out-of-plane rotations (Lee 1998), but both solutions require excitation of slices orthogonal to the imaging plane, producing saturation banding

artifacts in the resultant images. As a result, regions within the saturation bands cannot be analyzed reliably.

Continued improvements in imaging techniques and hardware may ultimately overcome the limitations described above. However, the need exists now for an artifact-free method to assess and correct motion in three dimensions, if only to quantitatively identify data sets overly contaminated by out-of-plane effects to exclude from further analysis. To achieve this goal, a subset of the algorithms described above have been extended into three dimensions. As in the two-dimensional case, the majority are image-space methods which simultaneously estimate motion in all six degrees of freedom by iterative least squares (Friston 1996) or related (Woods 1992) methods, and construct corrected images by linear interpolation. Thus, the limitations of these algorithms parallel those of the related two-dimensional algorithms discussed in Section 2.3, as do the motivations for new approaches.

Fourier-based methods, which will be introduced in Chapter 4, allow the separation of parameter estimation into multiple stages thereby reducing the complexity at each stage. A novel three-dimensional Fourier-based registration solution and its advantages over existing Fourier-based methods (Kassam 1996) will be presented in Chapter 6.

3 SPECIFIC AIMS OF THIS THESIS

It is the goal of this thesis to present novel processing and analytic methods for functional magnetic resonance imaging data which address the problems associated with subject motion described in Chapter 2. These include techniques to directly estimate and correct rotational and translational misregistration in image sets in both two and three dimensions, as well as other post-registration techniques to limit additional artifacts introduced by registration. The presentation is divided into three chapters.

In Chapter 4, a novel two-dimensional image registration algorithm is described. Unlike prior spatially based algorithms, this method performs the registration in frequency space, allowing the effective decoupling of rotational and translational elements of the estimation problem. Decoupling allows a fast, non-iterative design representing an improvement over prior iterative methods. The one-pass nature of the algorithm permits the use of a more computationally expensive interpolation method, frequency regridding, in place of the more commonly used bilinear interpolation, resulting in less high-frequency attenuation in the registered data sets. Problems related to implementation are discussed and solutions presented. The algorithm is validated using MR data acquired from imaging phantoms and human subjects and numerical simulations. Limitations of the algorithm are discussed, and a refinement of the original system designed to overcome some of these limitations is presented, implemented, and tested.

In Chapter 5, the issue of new artifacts introduced during image registration is raised and the mechanisms leading to these errors described. It is demonstrated that image registration alone, even with carefully chosen and computationally expensive interpolation routines, necessarily introduces errors which may degrade subsequent analysis. A method to generate simulated data to assess these errors is presented and validated by comparison with experimental data. A spatial filtering method is described which attempts to reduce post-correction variability in the high-frequency content of images. A temporal filtering method which exploits the empirically observed characteristics of post-registration artifacts is also presented.

In Chapter 6, the inherent difficulties of three-dimensional registration are revisited, and an extension of the two-dimensional Fourier-based methods into three-dimensions is proposed. Solutions required to apply the Levenberg-Marquart minimization to estimate the rotational displacements are derived, and implementation issues discussed. Finally, data demonstrating the accuracy of this partially iterative technique in three-dimensions are presented.

Finally, in Chapter 7, analysis of the results of a clinical research study are presented to demonstrate that the combination of these methods into a unified data processing strategy for functional MRI data can yield relevant insights into human physiology even under conditions predisposed to large subject motion effects.

4 IMAGE REGISTRATION IN TWO DIMENSIONS

As described in Chapter 2, the problem of motion artifacts in functional MRI is an important one. In section 4.1, the principle of k-space based decoupling of rotational and translational misregistration is described. In section 4.2, a two-dimensional image registration algorithm which exploits this principle is described. In section 4.3, enhancements to this algorithm are described which permit more accurate estimation of rotational parameters.

4.1 DECOUPLING ROTATION AND TRANSLATION

As described in Section 2.3, most rigid body registration algorithms in use up to now have simultaneously estimated both rotational and translation parameters with iterative techniques. Iterative approaches are computationally intensive since at each iteration an interim image must be computed. Furthermore, purely spatial approaches are complicated by the close interplay of rotation and translation in x,y-space since the center of rotation, or equivalently the linear displacement of the origin, must be estimated concurrently with the rotation angle. The problem of estimating the displacement parameters can be greatly simplified in k-space, where rotation and translation can be decoupled from one another, reducing the number of parameters which must be simultaneously estimated, and allowing efficient independent estimation of rotational and translational parameters. In fact, in the two-dimensional registration problem, a one-pass technique can be implemented, as will be described in Section 4.2.

The decoupling principle is based on the mapping of translations and rotations between x,y-space and k-space. This principle, introduced here in two dimensions, is directly extended to three dimensions in Chapter 6. It consists of two parts:

(A)   if two continuous real images $d_1(x,y)$ and $d_2(x,y)$ are related by a linear shift in x,y-space, then their two-dimensional Fourier transforms $D_1(k_x,k_y)$ and $D_2(k_x,k_y)$ will be related by a phase shift in k-space, leaving the magnitudes of the Fourier transforms unchanged; and

(B)   if two continuous real images $d_1(x,y)$ and $d_2(x,y)$ are related by a rotation about the origin of x,y-space through an angle of ψ, then their two-dimensional Fourier transforms $D_1(k_x,k_y)$ and $D_2(k_x,k_y)$ will also related by a rotation about the origin of k-space through the same angle of ψ.

Since rotation through an angle ψ about any point in x,y-space other than the origin can be further decomposed into a linear translation followed by a rotation about the origin through the same angle ψ, or vice-versa, the magnitudes of the Fourier transforms will always be related by a rotation about the origin of k-space regardless of the center of rotation or superimposed linear shifts in x,y-space. (N.B.: due to the symmetry of the Fourier transform magnitude for real images, a rotation through the angle ψ cannot be differentiated from a rotation through $\psi + 180°$, but this limitation is of no significance in functional MRI applications where all rotations are expected to be much smaller than $180°$.)

By applying this principle and examining only the magnitude of the Fourier transforms, the center of rotation in image-space need not be known in order to determine the rotational angle relating the two images. This can be viewed as a means of "centering" the rotations in k-space, and decoupling them from superimposed translations, which remain encoded in the phase. Thus, in the continuous case, rotation can be *independently* estimated by determining the displacement angle ψ about the origin of k-space which relates the magnitudes of the Fourier transforms. This can be accomplished efficiently in polar k-space coordinates, where rotations about the origin of Cartesian k-space will map to linear translations at a given radii in polar k-space.

As will be described below, these linear shifts are efficiently measured using cross-spectrum techniques. After rotating the Fourier transform of the second image by the computed displacement angle, only a linear shift will remain between the two images. This residual shift, too, can be efficiently measured using cross-spectrum techniques and corrected by reversing the phase shifts. Taken together, these steps constitute the Decoupled Automated Rotational and Translational (DART) registration algorithm (Maas 1997), described in the next section.

4.2 THE DART ALGORITHM

The principles of Section 4.1 have been successfully applied in optical systems for signal detection (Casasent 1976, Jensen 1987), pattern recognition algorithms (Altmann 1984), navigator-echo-based motion detection systems in conventional MRI (Fu 1995), and registration systems for retinal images (Apicella 1989, Cideciyan 1992). Unlike these applications, however, the algorithm described here was developed to exploit the decoupling principle for digital image data without the use of recursive or iterative techniques or navigator echoes, and thereby provide a high throughput, one-pass registration system for functional MRI applications. By avoiding iteration, a higher fidelity resampling method can also be substituted without imposing significant performance costs, thus avoiding the shortcomings of conventional linear interpolation techniques.

4.2.1 Correcting rotational misregistration

As in the continuous case described in Section 4.1, the angle ψ relating two reconstructed magnitude images $d_1[x,y]$ and $d_2[x,y]$ and their discrete Fourier transforms $D_1[k_x,k_y]$ and $D_2[k_x,k_y]$ can be estimated from the transformations of $D[k_x,k_y]$ into polar coordinates. To generate the polar sampling, an approximation of the continuous $D(k_x,k_y)$, labeled $\hat{D}(k_x,k_y)$, must be computed. In the DART algorithm, this is accomplished by interpolation of $D[k_x,k_y]$ with a circularly symmetric Fourier-space Hamming window. From this, we redefine the approximated spectra in polar coordinates, $\hat{P}(k_r,k_\phi)$, as:

$$\hat{P}(k_r,k_\phi) = \hat{D}(k_r\cos k_\phi, k_r\sin k_\phi). \qquad [4.2.1]$$

Applying the principles of section 4.1, the polar spectra $\hat{P}_1(k_r,k_\phi)$ and $\hat{P}_2(k_r,k_\phi)$ of two images related by a rotation of ψ, irrespective of superimposed linear translations in the underlying images, will be ideally related by the expression:

$$\left| \hat{P}_1(k_r,k_\phi) \right| = \left| \hat{P}_2(k_r,k_\phi + \psi) \right|, \qquad [4.2.2]$$

where the vertical bars indicate the absolute magnitude.

In the original description of the DART registration algorithm, $\hat{P}_1(k_r,k_\phi)$ and $\hat{P}_2(k_r,k_\phi)$ were discretely sampled as $\hat{P}_1[k_r,k_\phi]$ and $\hat{P}_2[k_r,k_\phi]$, and the rotational displacement estimate $\hat{\psi}$ between the images was determined by locating the maximum value of the discrete wrapped cross-correlation of the magnitudes of $\hat{P}_1[k_r,k_\phi]$ and $\hat{P}_2[k_r,k_\phi]$ with respect to k_ϕ over the range of k_r of interest. Symbolically:

$$\hat{\psi} = \max_{k_\phi} \sum_{k_r'} \left\{ \left| \hat{P}_1[k_r',k_\phi] \right| \otimes \left| \hat{P}_2[k_r',k_\phi] \right| \right\}, \qquad [4.2.3]$$

where $\max_x f[x]$ is the value of x for which $f[x]$ is maximized and $\otimes$ represents wrapped convolution with respect to k_ϕ, i.e.:

$$A[x] \otimes B[x] = \sum_{x' \in \{0..X-1\}} A[x'] \cdot B\left[(x-x')_X \right] \quad \text{for } x = 0..X-1, \qquad [4.2.4]$$

where $(x)_X$ represents x modulo X. As will be discussed below, this method is limited by the discrete nature of the cross-correlation, and the large number of angular samples of k_ϕ which are required to achieve the necessary accuracy. An improved method for estimating the rotation angle ψ is described in Section 4.3.

After rotational parameter estimation, k-space regridding (Jackson 1991) is employed to generate the rotationally corrected image. To rotate an image using k-space regridding, the k-space sample grid of the original test image is convolved with a continuous filter chosen to have small sidelobes in image space. The specific filter parameters will be described in Section 4.2.3. The resultant continuous k-space data is then resampled along a new grid rotated about the center of k-space by the desired angle, and this new k-space matrix is transformed back to form a magnitude image. Finally, this image is appropriately scaled to compensate for the roll-off characteristics of the filter mainlobe in image-space. Small image-space filter sidelobes reduce the aliasing energy introduced into the new image, and a filter mainlobe similar in size to the image of interest allows accurate compensation of the introduced attenuation. In contrast, inspection of the k-space characteristics of the image-space linear-interpolation family of filters commonly used in image registration algorithms shows how this approach used for rotated resampling may result in uncompensated high-frequency attenuation, which can lead to image degradation.

27

4.2.2 Correcting translational misregistration

After the correction of rotational misregistration, the two images are related by a purely linear translation. A large number of algorithms are available for the estimation and correction of a linear translation between two images. Cross-correlation techniques, such as those implemented by Baudendistel et al. (1995) and Alexander et al. (1996), only estimate pixel displacements to the nearest integer in each direction of interest. Thus, sub-pixel estimates along an axis are not possible without additional image manipulation. By interpolating new values between existing pixels, sub-pixel resolution can be achieved, but at the cost of increased computational time. By considering the problem in the Fourier domain, the shift property of Section 4.1 can be applied. Expressing this principle symbolically:

$$
\begin{aligned}
\text{if} \quad & d_2(x,y) = d_1(x - \Delta x, y - \Delta y), \\
\text{then} \quad & D_2(k_x, k_y) = D_1(k_x, k_y) \cdot \exp\left[-i(k_x \Delta x + k_y \Delta y)\right].
\end{aligned}
\qquad [4.2.5]
$$

In Equation 4.2.5, Δx and Δy are the linear displacements relating the two images, which need not be integers. By computing the cross-power spectrum of the two images, the phase information associated with the images cancels, leaving only the phase associated with the shift. Symbolically:

$$
\begin{aligned}
C_{12} &= D_1 \cdot D_2^* \\
&= D_1 \cdot D_1^* \cdot \exp\left[i \cdot (k_x \Delta x + k_y \Delta y)\right] \\
&= |D_1|^2 \cdot \exp\left[i \cdot (k_x \Delta x + k_y \Delta y)\right] \\
\angle C_{12} &= \Delta x \cdot k_x + \Delta y \cdot k_y,
\end{aligned}
\qquad [4.2.6]
$$

where C_{12} is the cross-power spectrum, the asterisk indicates complex conjugation, the vertical bars indicate absolute magnitude, and $\angle C$ denotes the phase angle of C. As seen in Equation 4.2.6, the phase angle of the cross-power spectrum thus defines a "phase plane" in k_x and k_y with coefficients Δx and Δy equal to the displacements along each image axis. These coefficients are estimated by fitting a plane using the least-squares method, and then used to construct a unity magnitude "conjugate phase plane" image. The k-space samples of the rotationally corrected but translationally misregistered image can then be multiplied by this conjugate phase plane, reversing the linear shift. Inverse Fourier transformation of the result yields the final registered image. Since this plane fitting algorithm does not require integer values for the shift coefficients Δx and Δy, this method permits the rapid identification and correction of sub-pixel shifts.

4.2.3 Implementation and performance

The complete registration algorithm consists of the polar recoding and cross-correlation of the image power spectra to estimate rotation, rotational alignment by k-space regridding, and linear translation estimation and correction by phase-plane conjugation. The algorithm was coded in C and tested on a variety of hardware and operating systems. The choices for implementation parameters and other operational specifications are described in this section.

All k-space data used by the algorithm are back-computed from the magnitude image-space data generated by the echo planar imaging station. All image dimensions are powers of two. The fast Fourier transform, as adapted from Numerical Recipes in C (Press 1994), was used for all necessary transformations. All k-space data was computed at doubled resolution by zero-padding the image data. For small inter-frame rotation angles, correlation of sampling frame dependent aliasing energy in the test and reference images leads to undesired bias towards zero in rotational estimates. This correlation, described in more detail in Section 5.1, is reduced by first rotating the reference image by 45 degrees, such that sensitivity to small rotations is maintained. This rotation of the reference image is performed with the same regridding algorithm used to correct the misregistered images.

Polar resampling of the k-space data is accomplished by convolution with a circularly symmetric 5.1 sample wide continuous Hamming window. Empirically, 512 angular samples were selected, yielding a step size of 0.352 degrees. (The additional factor of two results from the half-plane rotational symmetry of the power spectra.) Values are computed for non-negative range values, with the range step size defined to match to the resolution of the k-space matrix, e.g. 64 range points are needed for a 128x128 matrix. Angular cross-correlations are computed using fast Fourier transform methods. A composite angular cross-correlation is computed by averaging the cross-correlations over the central half of range values, since low frequency values for typical brain images may be relatively symmetrical about the center of k-space, and very high frequencies may contain relatively more noise energy. The maximum offset is determined by linear search of the composite cross-correlation, and corrected for the 45 degree reference image adjustment.

Following estimation of the rotational angle, the misregistered Cartesian k-space matrix is convolved with a circularly symmetric 5.1 sample wide continuous Hamming window to generate the rotated test Cartesian k-space matrix values. The width of the filter was chosen such that the original image space of interest is completely contained within its image-space mainlobe. A magnitude image is computed from the Fourier transform, and cropped to the original image size. To correct for the attenuation introduced by the filter, data values in the resultant image are divided by the corresponding image-space filter mainlobe values. These

filter values are computed by inverse Fourier transformation of the filter sampled with the same k-space matrix. The resulting rotated test image has the same dimensions, resolution, and orientation as the reference image.

After zero-padding and transformation of the rotationally aligned image back to k-space, the cross-power spectrum of this interim image with the reference image is computed by point-by-point conjugate multiplication of the k-space samples. The phase angle of each point of the cross-power spectrum is computed with the arc tangent function as needed. To allow for larger translational displacements, which may result in phase wrapping, an automated first estimate for Δx and Δy displacements is computed using only the first k-space sample along each frequency axis. After the appropriate remodulation of the cross-power spectrum, a second estimate is optionally computed using equal weighting of the first and second samples along each axis. Finally, after a second remodulation of the cross-power spectrum, the phase plane is estimated using the least-squares method. To reduce noise effects during fitting, each frequency is weighted by its cross-power spectrum magnitude. The least-squares fit is performed over the lowest half of frequencies, allowing for a residual motion of at least $\sqrt{2}$ pixels in any direction before phase-wrapping occurs. In practice, the initial estimates were found to be consistently accurate to less than one pixel, allowing a considerable margin of error. After each k-space value of the interim test image is appropriately multiplied by the conjugate phase plane model generated from the estimates, inverse Fourier transformation yields the final registered test image.

The processing time for image registration on a Hewlett-Packard 735/125 with the implementation parameters described above is 7.5 seconds for a 128 x 128 pixel image, and 1.7 seconds for a 64 x 64 pixel image. Similarly, a 64 x 64 pixel image can be processed in 0.7 seconds when the algorithm is run on a Digital Equipment Corporation AlphaStation 500/333. Thus, a typical 100 image 64 x 64 pixel functional MRI time series data set can be aligned in under 3 minutes using the HP system, or just over one minute using the AlphaStation.

4.2.4 Validation: imaging phantoms

The algorithm was tested using data acquired from both imaging phantoms and human subjects. All images were collected using a General Electric 1.5 Tesla Signa MR scanner retrofit with a whole body echo planar gradient set by Advanced NMR Systems, Inc. The imaging phantoms consisted of a 5.5 cm diameter orange and a 7.5 cm diameter tomato. Each phantom was secured to an imaging platform which allowed free adjustment of rotation ± 0.25 degrees about the central axis of the scanner or translation in the axial scanner plane. Neither

platform was detectable by magnetic resonance imaging. Thirty-three images were acquired for the orange over a range of zero to 20 degrees, and thirty-two images were acquired of the tomato over a range of zero to 9.33 mm (axial echo planar spin echo images, quadrature head coil, TE = 100 ms, TR = 2 s, in-plane resolution 1.56 x 1.56 mm, slice thickness 5 mm). Representative images are shown in Figure 4.2.1.
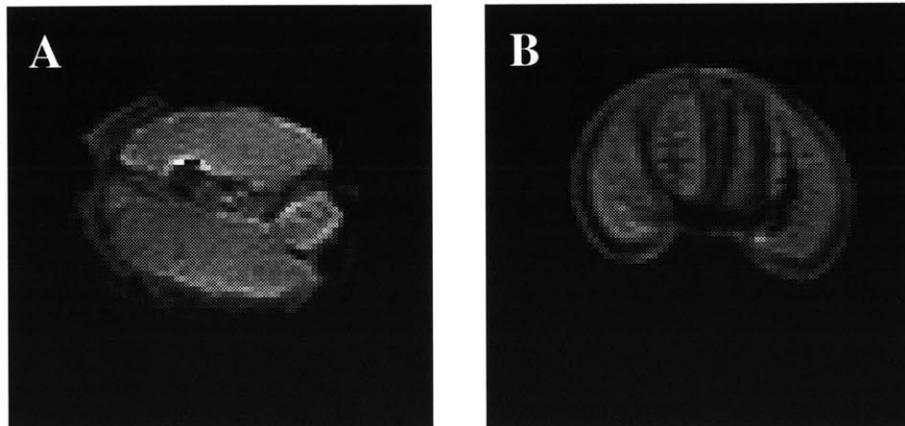
Figure 4.2.1. Imaging phantoms used to measure accuracy of the registration algorithm. (A) Representative image of the orange phantom used to assess rotational estimation. (B) Representative image of the tomato phantom used to assess translational estimation.

For both phantom image sets, the first image in the set was used as the reference image and computed registration parameters were compared with measured values, as presented in Figure 4.2.2. For the rotational data set, the mean and standard deviation of the difference between the measured and computed angles of rotation were -0.09 degrees and 0.17 degrees, respectively. All errors ranged between -0.44 and +0.31 degrees, inclusive. Difference images before and after registration for both small (0.25°) and large (19.42°) rotational displacements are presented in Figure 4.2.3. Evident in Figure 4.2.3D are the effects of slight out-of-plane motion on the difference image. Instead of the typical leading dark edge and trailing light edge that would be expected with residual misregistration, this image illustrates an outer "halo" pattern consistent with out-of-plane motion of the imaging phantom. Additionally, as the orange was injected with water prior to imaging, an air bubble is visible whose position varies with the angle of rotation. Following registration, the shift in the air bubble is more evident.
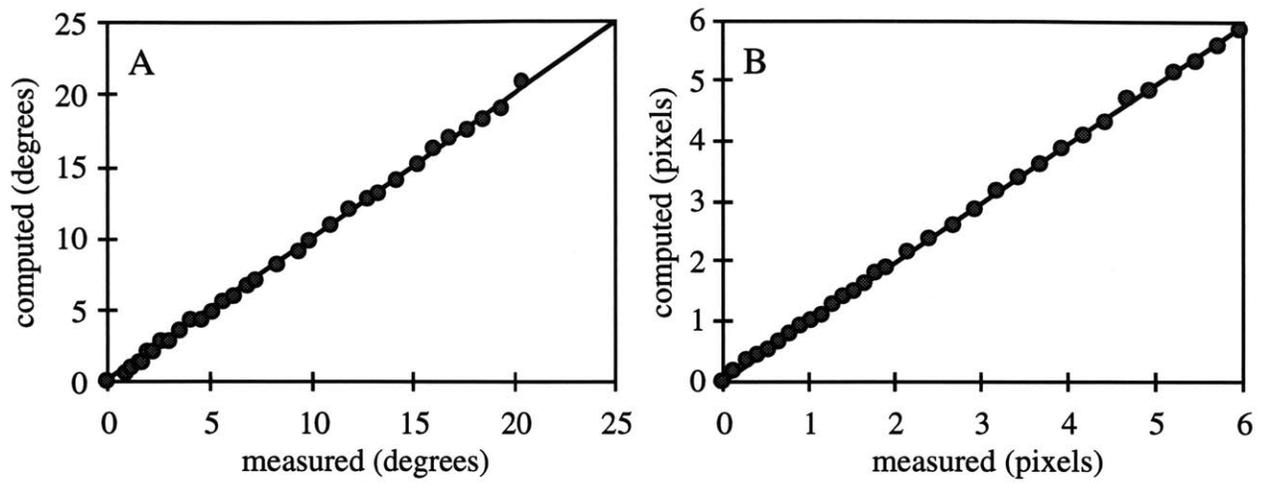
Figure 4.2.2. Computed vs. measured displacements for the two imaging phantoms. (A) Rotation. (B) Translation.
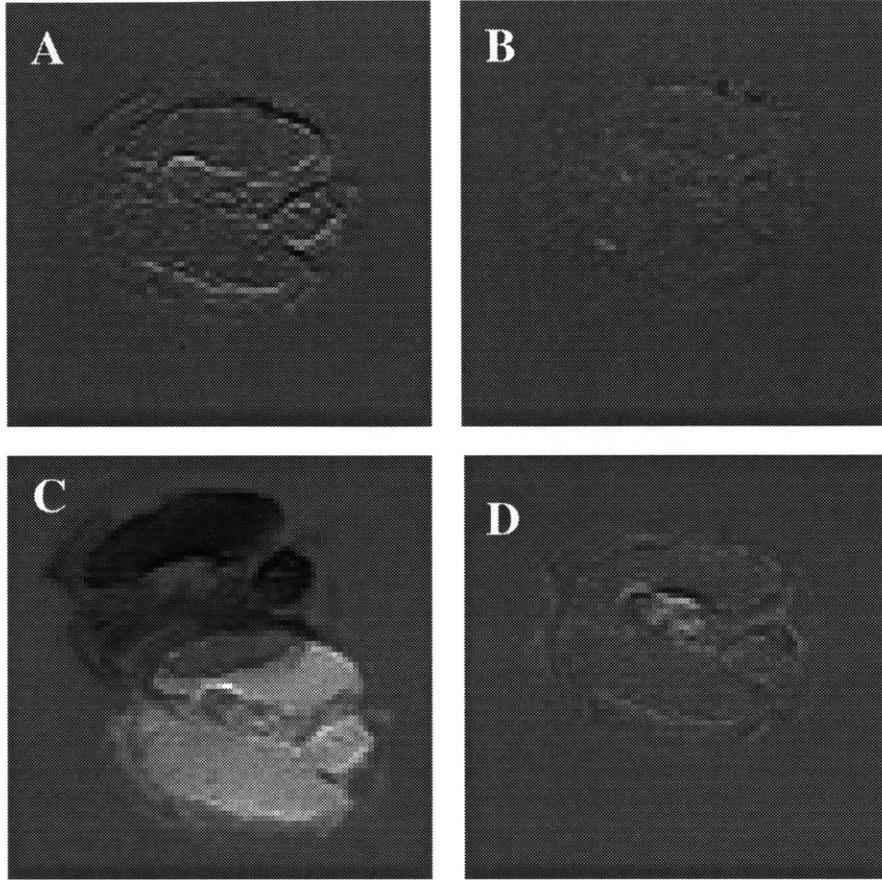
Figure 4.2.3. Correction of small and large rotations. (A) Difference image for rotation of 0.25 degrees. Significant edge-related artifacts are introduced for even small rotations. (B) Residual difference image after registration of images from (A). (C) Difference image for rotation of 19.42 degrees. (D) Residual difference image after registration of images from (C).

For the translational data set, the mean and standard deviation of the difference between the measured and computed displacements were -0.035 pixels (0.055 mm) and 0.054 pixels (0.084 mm), respectively. All errors ranged between -0.16 (-0.25 mm) and +0.06 pixels (+0.094 mm), inclusive. Difference images before and after registration for both small and large translational displacements are presented in Figure 4.2.4.
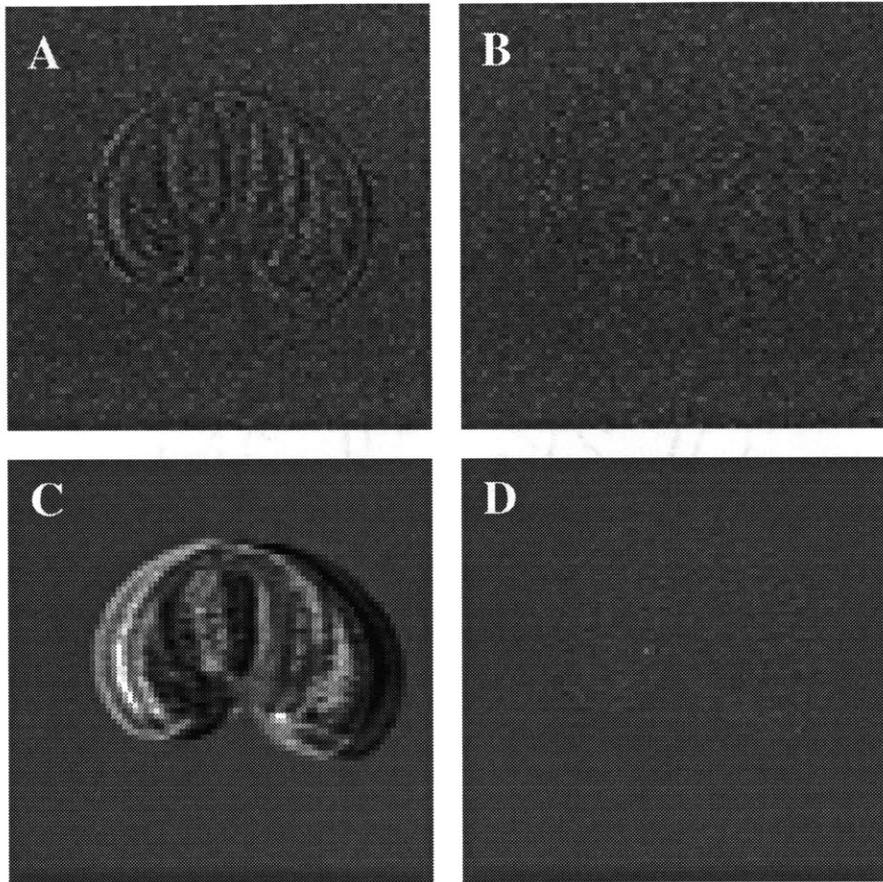
Figure 4.2.4. Correction of small and large translations. (A) Difference image for translation of 0.127 pixels. Significant edge-related artifacts are introduced for even small subpixel shifts. (B) Residual difference image after registration of images from (A). (C) Difference image for translation of 5.97 pixels. (D) Residual difference image after registration of images from (C).

4.2.5 Validation: numerical simulations

Since registration algorithms based upon least-squares or cross-correlation approaches assume that, aside from position, the imaged object is constant from frame-to-frame, these intensity-based methods may incorrectly estimate registration parameters in the presence of signal changes related to cortical activation (Biswal 1997). Numerical simulations were performed to determine whether the DART registration algorithm is sensitive to this type of bias.

A set of 64 x 64 pixel images was constructed to simulate typical 1.5 Tesla gradient echo functional echo planar images. A baseline image of the simulated 1540 pixel brain was registered to several simulated activation images with and without superimposed white noise (SNR = 100) using the algorithm described above. Posterior regions of activation occupying both 99 and 183 pixels at both 3% and 6% signal increase relative to the baseline cortex were analyzed. A smaller posterolateral 26 pixel region of simulated 6% signal increase was also examined. While 3% activation is occasionally reported at 1.5 Tesla, 6% signal increases are generally larger than those expected in typical 1.5 Tesla functional MR experiments (Kwong 1992, Bandettini 1992). Thus, as there were no true rotations or translations in the image sets, any falsely detected misregistration should represent an upper bound on the expected biases introduced by true functional activation.

In all cases, there were no rotational errors introduced by the algorithm. For every simulated activation image, the magnitude of the translational registration error vector was less than 0.004 pixels, corresponding to 0.0125 mm for the simulated field of view of 20 cm, with standard deviations less than 0.002 pixels (n=32 for each case.) This included the analysis of the largest 6% activation region, which filled nearly 20% of the brain pixels.

Additional simulations were conducted to compare the errors introduced by regridding to those introduced by bilinear interpolation. A test image was rotated first clockwise and then counter-clockwise through a test angle with each method. The root mean squared (RMS) difference between the original and rotated images was then computed as a function of rotational angle. The test image was taken from a photic stimulation FMRI data set of a human subject (oblique echo planar gradient echo, 5 inch receive-only surface coil, TE = 40 ms, TR = 1s, flip angle 66°, slice thickness 7 mm, in-plane resolution 3.1 x 3.1 mm). Test angles ranged from 0 to 90 degrees in 0.25 degree steps. The RMS difference error was calculated over all pixels in the central 31 pixel circle. The relative error energies introduced by k-space regridding and bilinear interpolation are presented in Figure 4.2.5A. The reduction in error energy when k-space regridding is used is presented over the range of angles from 0 to 5 degrees in Figure 4.2.5B. Similar results were found with a test image collected using a quadrature head coil. The improvement ranged between 10 dB and 18.4 dB, with the peak improvement occurring at 1.25 degrees of rotation. This peak represents a reduction of 88% in RMS noise, corresponding to a 98.5% reduction in added variance.
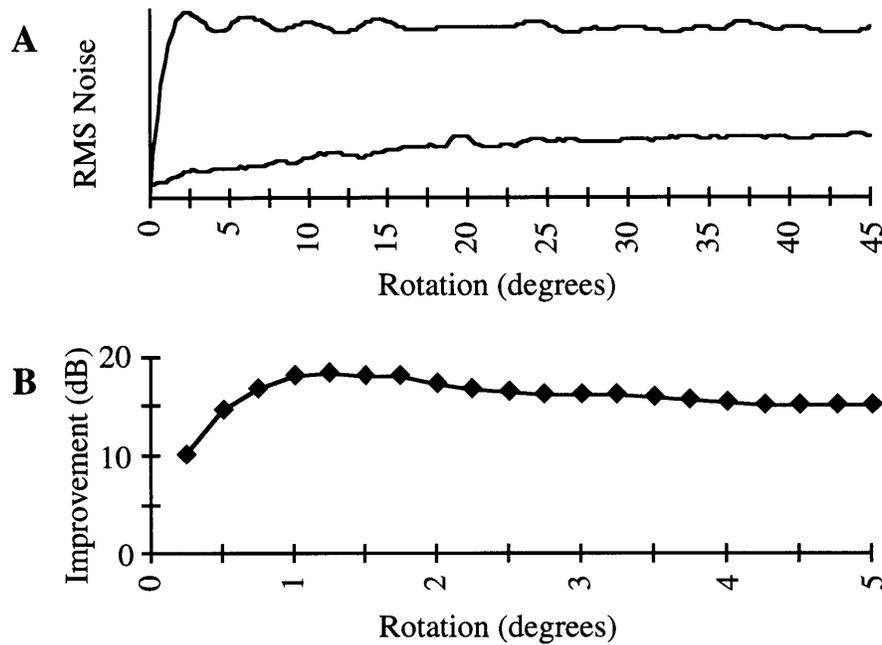
Figure 4.2.5. A comparison of errors introduced by linear interpolation and frequency regridding methods. (A) RMS Noise introduced following consecutive clockwise and counterclockwise rotations through the angle specified for linear interpolation (top) and frequency regridding (bottom). (B) Decibel reduction in noise for small rotations provided by frequency regridding.

Because of the decoupled nature of the algorithm, a final set of simulations was performed to test the reliability of the algorithm in the presence of simultaneous rotational and translational motion. This was done by considering the effects of residual rotation on the estimation of translation parameters. During registration of test images from functional experiments similar to the one described in the previous paragraph, a rotational bias was introduced ranging over ±5°. The effective center of rotation was then computed from the translational and rotational estimates. For both sagittal and axial imaging geometries, the center of rotation remained registered over the range of residual rotations studied. The point was also found to remain in a stable position, suggesting that the translational estimates are highly reliable in the presence of residual rotations up to 5°, values much larger than the errors demonstrated with the rotational phantom images.

4.2.6 Validation: human studies

Two photic stimulation experiments were conducted to assess the utility of this algorithm in the processing of human functional imaging data. These studies were approved by the McLean Hospital Institutional Review Board. All subjects provided informed consent prior to participating in the studies. The volunteer subjects wore goggles equipped with light emitting diodes (Grass Equipment Co., Quincy, Massachusetts), and were exposed to alternating intervals of darkness and photic stimulation in a checkerboard pattern at 8 Hz. This method has been previously demonstrated to produce robust activation in the visual cortex.

In the first experiment, a 22 year old right-handed male was imaged in an oblique plane running parallel to and through the calcarine fissure using a 5 cm diameter receive-only surface coil positioned at the back of the head (gradient echo, TE = 40 ms, TR = 1 s, flip angle 66°, slice thickness 5 mm, in-plane resolution 3.1 x 3.1 mm, 64 x 64 pixels). The surface coil was used to assess the utility of the registration algorithm in the processing of functional imaging data collected using coils with non-uniform receptivity profiles. The length of each interval of photic stimulation or darkness was 10 seconds. During subject motion, regions of differing contrast shift across pixel boundaries, leading to linear trend artifacts in pixel time series data (Bandettini 1993). Thus, the examination of reductions in linear trends in the pixel time series after registration serves as one marker of the effectiveness of the algorithm in cases where the true displacements are unknown.

The set of 256 echo planar images was registered, with the image at the mid-point of the middle rest interval, i.e. frame 130 out of 256, serving as the reference. A typical image from this data set is presented in Figure 4.2.6A. The standard deviations of the computed rotational and translational displacements were 0.1924 degrees and 0.0243 pixels (0.0759 mm), respectively. Since the mean displacements depend only on the choice of the reference frame, they are not reported here. The correlation coefficient of the time series data was computed on a pixel-wise basis before and after registration using a ramp function as the reference waveform to detect linear trend artifacts (Bandettini 1993). Gray-scale maps of the computed correlation coefficients are presented in Figures 4.2.6B and 4.2.6C. Strong linear trend components are evident along the edges of the cortex in the original, unregistered data set (Figure 4.2.6B). These trends are greatly reduced in the registered data set (Figure 4.2.6C), suggesting that the images have been accurately registered. Complete elimination of linear trends is unlikely as residual out-of-plane motion also contributes to this artifact.
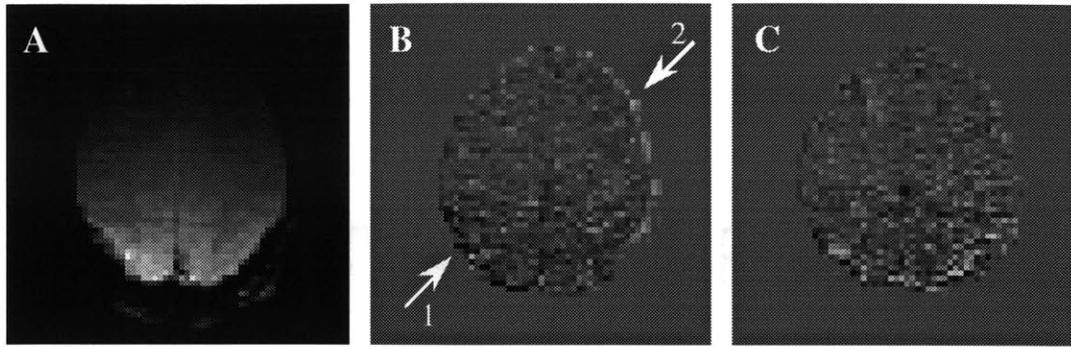
Figure 4.2.6. Reduction in linear trends following image registration. (A) a typical 64 x 64 pixel image from the analyzed functional MRI data set. (B) Map of the correlation coefficients of original pixel times series with a ramp reference function without registration. (C) Correlation coefficients following image registration. The scale for (B) and (C) ranges from -0.75 (black) to +0.75 (white). Only pixels within the brain have been included in the analysis. In the unregistered image set, note the negative linear artifact around the bottom left edge of the image (arrow 1) and associated positive artifact across the top right (arrow 2). Both artifacts are visibly reduced following image registration. N.B.: According to radiological convention, the right side of the image corresponds to the left side of the subject.

In the second experiment, data were acquired from a 23 year old right handed male in a parasagittal plane four millimeters to the left of the interhemispheric fissure using a quadrature head coil (gradient echo, TE = 40 ms, TR = 3 s, flip angle 75°, slice thickness 5 mm, in-plane resolution 3.1 x 3.1 mm, 64 x 64 pixels). The length of each interval of photic stimulation or darkness was 30 seconds. A typical image from this data set is presented in Figure 4.2.7A. The data were examined before and after registration to identify false activation artifacts caused by stimulus-correlated motion (Hajnal 1994). As was discussed in Chapter 2, the reduction of this type of artifact is particularly important in the interpretation of functional imaging data.

The set of 64 echo planar images was registered, using the image corresponding to the mid-point of the second stimulation interval as the reference. The standard deviations of the computed rotational and translational displacements were 0.1626 degrees and 0.0366 pixels (0.1144 mm), respectively. Assuming a 6 second hemodynamic delay (Kwong 1992), correlation coefficeints were computed at each pixel before and after image registration using a boxcar reference function to assess functional activation (Bandettini 1993). Maps of the absolute values of these correlation coefficients before and after registration are presented in

Figures 4.2.7B and 4.2.7C, respectively. A reduction in activation artifact is evident at the inferior posterior edge of the occipital lobe. Because of the high contrast at the edge of the cortex, this region is especially susceptible to motion artifacts. Importantly, the region of activation encompassing the visual cortex has not been affected by the registration algorithm, indicating that preprocessing with the DART algorithm can significantly reduce activation artifacts without disrupting true activation signals.
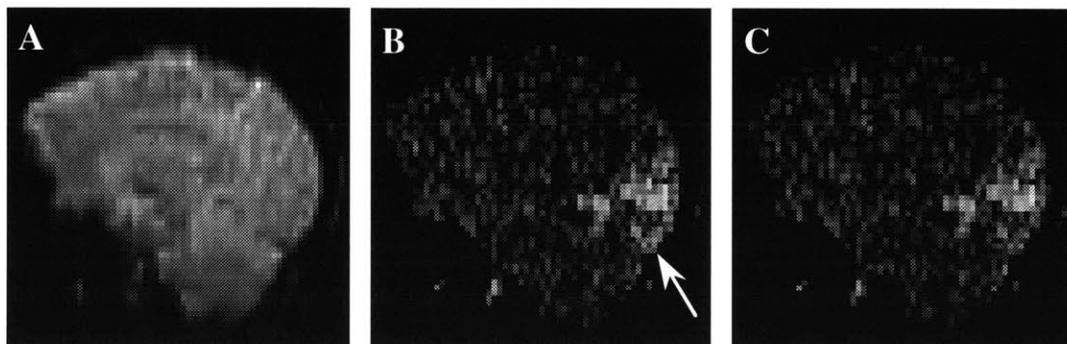
Figure 4.2.7. Reduction in motion-related artifacts in a functional MRI experiment with image registration. (A) Sample sagittal image from the data set of the second experiment. (B) Correlation map of pixel time courses with boxcar reference function before registration. (C) Correlation map following image registration. In the unregistered data, the pixels at the inferior posterior edge of the occipital lobe (arrow) demonstrate high correlation coefficients. These false positives do not correspond to neuroanatomic regions which would be expected to activate with this paradigm, but rather represent stimulus-correlated motion artifacts. Following registration, the apparent activation in this region is greatly reduced and falls below the statistical threshold which identifies the visual cortex. Scale of correlation maps ranges from 0 (black) to 0.80 (white).

4.2.7 Discussion

The sensitivity of functional MRI to motion on the subpixel scale was discussed in Chapter 2. In Figures 4.2.3 and 4.2.4, the clear edge artifacts resulting from image displacements of as small as 0.127 pixels or rotations of as little as 0.25° can be seen, reaffirming the need to address even sub-pixel frame-to-frame motion prior to statistical analysis. Since these displacements may correspond to motion of less than 0.2 millimeters, it is unlikely that typical FMRI experiments are immune from all motion effects. The presence of

motion in even straightforward functional experiments is supported by the human photic stimulation data, where rotational and translational displacements with subpixel standard deviations contributed to visible artifacts.

An algorithm to correct in-plane frame-to-frame motion in large functional MRI data sets was described above. The algorithm was implemented and tested using simulations, echo planar phantom data, and functional imaging data from human subjects. Test results indicate that this algorithm can be used to correct subpixel motion between frames due to both translation and rotation, and reduce motion artifacts which contaminate functional data sets. Processing requires no operator intervention, and the rapid computation is suitable for large FMRI data sets. Although the registration of images collected with surface coils is complicated by non-uniform receptivity profiles, the human data presented suggest that the algorithm may also be beneficial in the processing of these FMRI data sets.

In the implementation described, an angular resolution of 0.352° was selected, such that rotation is ideally corrected to within 0.176°. This bound on accuracy is supported by the results of the rotational imaging phantom experiments. Although in theory arbitrary resolution is attainable by decreasing the angular step size in exchange for increased computation time, this choice has been found empirically to represent a good tradeoff. This correction accuracy leaves a residual motion of less than 0.2 pixels for points at the edges of a 128 x 128 pixel image, or less than 0.1 pixels for 64 x 64 pixel images, with a decreasing residual towards the center. As image information of interest is typically not at the edges of the image, most points in a 64 x 64 pixel time series will be corrected for rotation to within less than 0.1 pixels. While this is a good starting point, Figures 4.2.3A and 4.2.4A demonstrate that even errors on this order may leave undesirable edge artifacts. Thus, an enhancement to the algorithm which overcomes the discrete nature of the cross-correlation will be described in section 4.3.

It has been shown that the reliability of certain intensity based registration algorithms is compromised in the presence of superposed activation (Biswal 1997). The numerical simulations above, however, suggest that this method is not complicated by this source of bias, even in the presence of a 6% localized activation signal occupying nearly 20% of the brain pixels.

This algorithm uses a k-space regridding procedure to correct rotational errors instead of more conventional linear interpolation techniques. The numerical analyses above indicate that this method significantly reduces the additional error introduced in the rotation of an image, with the largest reductions of up to 88% in RMS interpolation error occurring for small angles of rotation. The human data sets presented include a number of images which required small rotational corrections, suggesting that these gains are realized in typical FMRI data sets. Similar improvements may also be obtained with other high-fidelity interpolation methods,

such as two-dimensional sinc interpolation, but the large window sizes required may lead to unacceptable increases in computation time (Hajnal, 1995).

Recently, real-time estimation and compensation for subject motion through the use of navigator echoes has been reported (Fu 1995, Lee 1996, Lee 1998). Based in part on the methods described here, Lee et al. (1998) have developed a functional imaging method in which data is collected from a circle in k-space in a slice parallel to the imaging plane during the time interval between image acquisitions. As this circle provides a direct polar sampling of k-space at a single fixed range, rotation and translation can be estimated in a manner similar to that described above by comparison to a reference navigator echo. Thus, the scanner gradients can be adjusted prior to the acquisition of each functional imaging frame, locking the imaging axes to the object. While promising, further work is necessary to determine the accuracy of estimates computed from a single circle in k-space (Kassam 1996). Furthermore, extending the method to out-of-plane parameters presently requires excitation of a slice orthogonal to the imaging plane, which results in a saturation band and loss of data in the functional images.

This algorithm corrects for rigid-body motion, and is not designed to correct non-rigid deformations, as may occur with eye movements (Chen 1997), swallowing or speaking (Birn 1998), gradient nonlinearities (Bushberg 1995), or as a result of displacement-dependent variation in magnetic susceptibility artifacts (Wu 1997). Although non-rigid motion may complicate the analysis of functional imaging data, the results presented above suggest that gross rigid-body motion of the brain is an important source of artifacts in functional MRI, as these artifacts are greatly reduced by preprocessing with this algorithm. Furthermore, subject movement occurring during the acquisition of a single image has not been considered here. As was described in Chapter 2, the dominant concern in echo planar imaging, where an entire image can be acquired in under 100 ms, is frame-to-frame motion. Finally, the correction of out-of-plane motion has not been considered here, but will be addressed in Chapter 6. The present data suggest, however, that this registration algorithm does accurately accurately correct for in-plane movement, even in the presence of out-of-plane effects.

4.3 EXTENDING THE DART ALGORITHM

As described above, the original DART algorithm estimated rotational displacement by locating the peak of a discrete wrapped cross-correlation to identify the maximum overlap of the polar power spectra. In the implementation described above, this imposed quantized estimation of rotation in $0.37°$ steps, yielding an accuracy limitation of $\pm 0.18°$. Although this choice provided a reasonable trade-off between speed of computation and accuracy of

registration (Maas 1995), it is desirable to eliminate this limitation entirely. Thus, an extension to the algorithm to produce continuous estimates of rotation was designed, implemented, and tested. Importantly, this extension is also applicable in other algorithms which estimate rotation in related manners and are currently limited by the discrete step size imposed by cross-correlation (Kassam 1996, Alexander 1996, Lee 1998).

4.3.1 Continuous estimation of rotational displacement

Continuous estimation is accomplished by employing phase-slope estimation techniques instead of discrete maxima along polar cross-power spectra. As was described in Section 4.1, the polar recoding of the image spectral magnitudes maps rotations to linear shifts. Thus, this extension is conceptually analogous to the estimation method described in Section 4.2.2 for true linear translations. The derivation begins with Equation 4.2.2 for images of two objects related by a rotation through an angle ψ, which is reproduced here as Equation 4.3.1.

$$\left| \hat{P}_1(k_r, k_\phi) \right| = \left| \hat{P}_2(k_r, k_\phi + \psi) \right|. \qquad [4.3.1]$$

Extending to discrete estimates of the polar spectra and taking the one-dimensional Fourier transform of both sides while holding the radial variable k_r constant, the shift properties of the Fourier transform described above yield:

$$F_\phi\left\{ \left| \hat{P}_2[k_r, k_\phi] \right| \right\} \approx F_\phi\left\{ \left| \hat{P}_1[k_r, k_\phi] \right| \right\} \cdot \exp[-i\phi\psi]. \qquad [4.3.2]$$

In Equation 4.3.2, $F_\phi\left\{ P[k_r, k_\phi] \right\}$ is the Fourier transform of P with respect to k_ϕ holding k_r constant, and is a function of ϕ, the Fourier analog of k_ϕ. Assuming equality in Equation 4.3.2, the cross-power spectrum $\hat{C}_{12}[k_r, \phi]$ of $\left| \hat{P}_1[k_r, k_\phi] \right|$ and $\left| \hat{P}_2[k_r, k_\phi] \right|$ can be computed at each value of k_r, yielding:

$$\begin{aligned} \hat{C}_{12}[k_r, \phi] &= F_\phi\left\{ \left| \hat{P}_1[k_r, k_\phi] \right| \right\} \cdot F_\phi\left\{ \left| \hat{P}_2[k_r, k_\phi] \right| \right\}^* \\ &= F_\phi\left\{ \left| \hat{P}_1[k_r, k_\phi] \right| \right\} \cdot F_\phi\left\{ \left| \hat{P}_1[k_r, k_\phi] \right| \right\}^* \cdot \exp[i\phi\psi] \\ &= \left| F_\phi\left\{ \left| \hat{P}_1[k_r, k_\phi] \right| \right\} \right|^2 \cdot \exp[i\phi\psi]. \end{aligned} \qquad [4.3.3]$$

Again, the asterisk denotes complex conjugation. From the individual cross-power spectra, a composite cross-power spectrum $\hat{C}_{12}[\phi]$ is computed over a range K of polar radii k_r of interest:

$$\hat{C}_{12}[\phi] = \sum_{k_r' \in R} \hat{C}_{12}[k_r', \phi]$$

$$= \sum_{k_r' \in R} \left| \hat{C}_{12}[k_r', \phi] \right| \cdot \exp[i\phi\psi]. \qquad [4.3.4]$$

As in the two-dimensional translational case described above, the phase angle of the composite cross power spectrum is a function of ϕ with a coefficient equal to the rotational displacement ψ. As above, this can be solved by a least squares method:

$$\hat{\psi} = \min_m \sum_{\phi} \left(\angle \hat{C}_{12}[\phi] - m\phi \right)^2, \qquad [4.3.5]$$

where $\angle$ is the phase operator, and $\min_x f[x]$ is analogous to $\max_x f[x]$ defined above. The minimization can be performed in a single step by solution of the least squares equations, yielding the estimate of rotational displacement. By allowing continuous estimates, the extension allows a potential reduction in the number of polar points needed to generate equally accurate estimates, such that both higher resolution and speed can be achieved.

4.3.2 Validation

To test the ability of the enhanced estimation method to detect rotations, a set of 16 simulated 64 x 64 pixel MR images was constructed. For each image, the k-space data of a slice through an idealized 32 x 44 pixel rectangular prism was synthesized analytically assuming an incremental rotation through 0.1 degrees in each image, with the first image rotated initially by 20 degrees. Simulated images were then generated by taking the absolute magnitude of the inverse Fourier transform of the simulated k-space matrices, yielding an idealized test set. This data set was registered using the algorithm described in Section 4.2, with the first image in the set serving as the reference. Rotational displacements were first determined by using the cross-correlation rotational estimation method described in Section 4.2.1, and then by using the phase estimation method described in the previous section. In both cases, the reference was rotated by 45 degrees prior to estimation, and the power spectra were generated in polar coordinates using 512 angular samples.

The computed rotational angles are compared to the actual values in Figure 4.3.1 for the two estimation methods. The discrete nature of the estimates computed using the original method, represented by open circles, is clearly evident. Examination of the estimates computed using the enhanced method (solid circles) reveals accurate, continuous, estimates of rotation even for the small displacements tested. The mean and standard deviation of the difference

between the estimated and actual displacements for the original method were -0.025° and 0.102°, respectively. Using the enhanced estimation technique, these values were reduced to 0.005° and 0.005°, respectively, verifying that the new method is largely unbiased, even for small rotations, as well as more accurate than the original method.
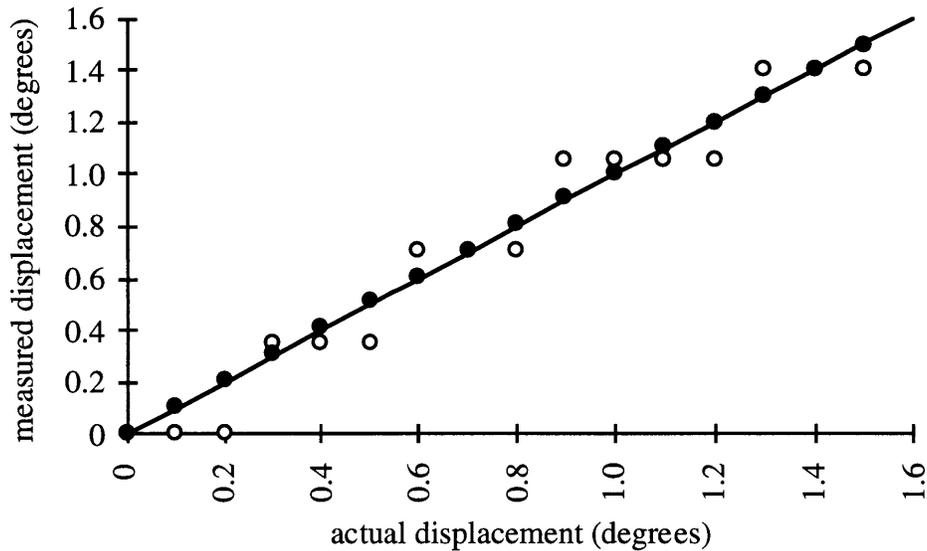
Figure 4.3.1. Measured vs. actual rotational displacements for the original DART algorithm (open circles) and the enhanced DART algorithm (solid circles). The enhanced method yields more accurate, continuous estimation of rotational displacements. Perfect solutions lie along the solid line.

5 RESAMPLING ARTIFACTS IN TWO DIMENSIONS

In Chapter 4, a two-dimensional registration algorithm to align functional magnetic resonance imaging data sets was presented. This and other algorithms have been useful in reducing artifacts caused directly by image misregistration. A critical step in every registration algorithm is the resampling of the misregistered images to bring them into alignment with the reference image, a step which may itself introduce new artifacts into the corrected data sets. Importantly, these artifacts are related to neither the accuracy nor the precision of the registration algorithm used, but rather are intrinsic to the resampling process itself and largely unrelated to the choice of resampling method. The nature of these artifacts and their implications in functional MRI data analysis have not been well characterized.

In Section 5.1, major sources of artifacts introduced by image resampling during the correction of in-plane rotational misregistration in echo planar FMRI are described. In Section 5.2, a spatial filtering approach to the reduction of these artifacts is presented. In Section 5.3, a temporal filtering approach is described. The discussion below will center on rotational correction, the dominant motion in functional MRI (Lee, 1998) and registration utilizing the frequency regridding method described in Chapter 4. The techniques and results, however, can be generalized to other registration methods as well.

5.1 NATURE OF RESAMPLING ARTIFACTS

Given two discretely sampled images $d_1[x,y]$ and $d_2[x,y]$ related by a rotation through an angle ψ, the second image must be resampled in order to bring it into alignment with the first image. A number of resampling methods have been described, including bilinear interpolation (Woods 1992), truncated sinc interpolation (Friston), frequency shearing (Eddy 1996), and frequency regridding (Maas 1997). Although this discussion centers on the frequency regridding method described in Chapter 4, each of these methods gives rise to similar artifacts, the major types of which can be divided into two categories: (1) acquisition-dependent errors, i.e. misalignment of frame-dependent ringing and aliasing in the original

data, and (2) correction-dependent errors, i.e. new ringing, aliasing, and high-frequency loss intrinsic to the resampling method. An overview relevant to both types of errors is presented in Section 5.1.1. Additional discussion of each error type is found in Sections 5.1.2 and 5.1.3. Finally, aliasing effects due to discrete sampling for both error types is discussed in Section 5.1.4.

5.1.1 The object and scanner frames

To better understand resampling artifacts, it is important to briefly review the nature of the acquired data in MRI experiments. As described in Chapter 1, data sampled from the receiver coils is fundamentally encoded in the Fourier domain, i.e. the data represent points in k-space. (N.B.: although multi-slice acquisitions using slice-select gradients combine spatial and spectral encoding, this discussion will center on data from a single anatomic slice only.) As also described in Chapter 1, a common approach to echo planar imaging is rectangular sampling of the k-space plane, generating a matrix of points. Without loss of generality, a square sampling region will be assumed. Furthermore, as the imaged object, e.g. the human head, is spatially limited, it will contain energy outside of the sampled k-space region. Generally, the energy content near the periphery of the sampling region is small, but non-negligible.

If the underlying object is rotated through an angle ψ between two image acquisitions, the principles of the Fourier transform discussed in Section 4.1 mandate that the Fourier representation of the image is also rotated through the same angle. As also discussed in Section 4.1, this is true for the magnitude of the Fourier transform regardless of the center of rotation or superimposed linear translations. Thus, by considering data in the frame of the object and not of the scanner, the boundaries of the frequency sampling grid for the second image are effectively rotated through $-\psi$, as shown in Figure 5.1.1A.
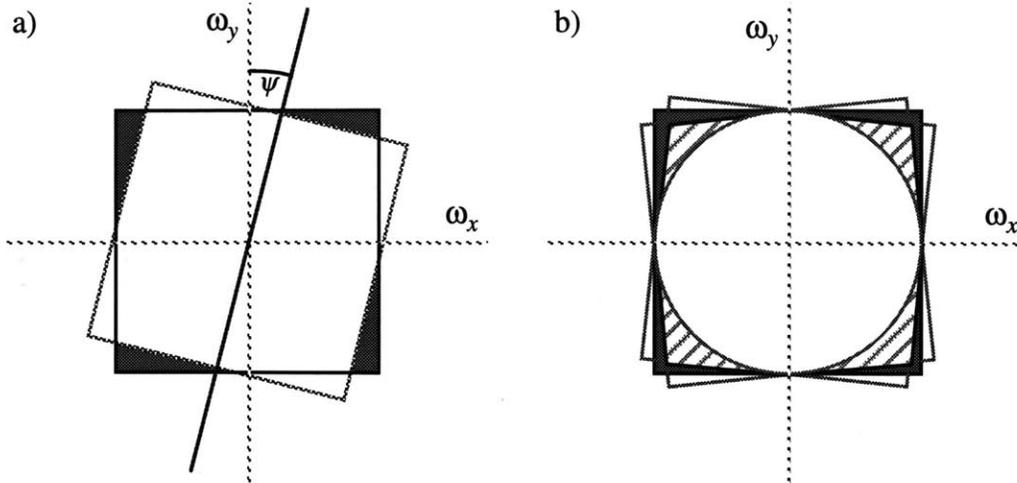
Figure 5.1.1. Effects of rotation on the sampling grid. (A) When considered in the object frame of reference, the clockwise rotation of an object through an angle ψ is equivalent to the counterclockwise rotation of the sampling region (dashed square) through the same angle, leaving several regions undefined (shaded regions) when considered relative to the reference frame's sampling region (solid square). (B) For images rotated through all angles, only the inscribed circle remains well-defined in all images. For a smaller range of rotations, the union of the undefined regions (shaded area) spares additional portions of k-space (hashed regions) extending towards the corners of the sampling region.

As aligning the second image to the reference frame with respect to rotation is analogous to resampling the misaligned frequency data on the original sampling grid, resampling will first be considered in Fourier space. Initially, continuous data within the sampling region will be assumed. While this "infinite sampling" is in fact an impossible case experimentally, it provides a useful starting point for the discussion. The additional aliasing artifacts introduced by discrete sampling are discussed in Section 5.1.4. With continuous data acquired from the sampling regions, acquisition and realignment can be conceptually reconsidered in two steps which relate, respectively, to the two types of error introduced at the beginning of Section 5.1. First, the misaligned data is acquired in a rotated sampling region with respect to the reference frame, leading to high-frequency loss and misalignment of frame-dependent ringing and aliasing. Second, "clipping" the rotated grid to the reference sampling region leads to new ringing, aliasing, and, potentially, further high-frequency loss. These errors are discussed in greater detail below.

5.1.2 Acquisition-dependent artifacts

An examination of Figure 5.1.1A reveals that, when considered in the frame of reference of the object, there are portions of the misaligned Fourier image near the corners of the desired reference sampling region which were *never sampled*, represented by the shaded areas. As this information is required in order to perfectly regenerate the second image in the reference frame, the lack of this information leads to an initial, unavoidable, high-frequency loss for all non-zero values of ψ. Stated differently, following alignment there are always four small triangular sections at the corners of the new sampling grid which are effectively undefined. Although different resampling schemes vary in their treatment of these undefined sections, in all cases there is necessarily a loss of high-frequency information in the resampled image dependent on the angular displacement ψ. Furthermore, in a large image set where the imaged object has undergone a range of rotations, the union of the undefined sections represents a part of k-space which will be variably degraded in the registered image set. As seen in Figure 5.1.1B, all frequencies outside the inscribed circle can be affected in extreme cases. This frame-to-frame variability of high spatial frequency content can thus be considered as a source of temporal variability in subsequent pixel-wise temporal analysis. It is important to stress again that the above description relates to rectangular sampling. Non-rectangular sampling schemes such as spiral imaging sequences may omit the corners of the sampling space altogether in all images, potentially limiting the importance of this type of error. However, even in such methods, the necessary processing step of taking the magnitude image of the inverse Fourier transform of the k-space matrix to generate a real image may reintroduce energy into these high frequency corners.

As windowing of the k-space data to the sampling region is equivalent to convolution of the underlying object image with the inverse Fourier transform of the windowing function, the error can also be considered in the spatial domain. In this case, the data will be considered in the frame of reference of the scanner. Again assuming rectangular windowing, the equivalent convolution kernel is the separable product of two sinc function, one along each of the principal imaging axes. These rectangular frequency windows introduce the expected Gibbs ringing, which can be seen clearly in Figure 5.1.2A. This figure simulates an MR image through a perfect homogeneous rectangular prism. To generate this image, the Fourier representation of the prism was computed analytically at each of the sampling points of a simulated square k-space matrix, which was then reconstructed into an absolute magnitude image. The solid vertical lines in the center are not part of the images, but were drawn in to facilitate identification of the vertical axis. The contrast of the image has been adjusted to highlight the ringing patterns, which can be clearly seen "leaking" from the object along the

two principal spatial axes. The patterns appear as smooth decays instead of high-frequency oscillations because these images represent the absolute magnitude of the inverse Fourier transform, as would be the case with images produced by typical MR scanner software. Naturally, the sinc patterns also ring inward into the object image parallel to the principal axes, although this is not appreciated in Figure 5.1.2 because of the contrast levels.

The frame-dependent, object-independent, nature of the ringing can be appreciated by comparing Figures 5.1.2A and 5.1.2B. In Figure 5.1.2B, the k-space matrix was regenerated analytically assuming an in-plane rotation of the object through twelve degrees. The prism is appropriately rotated in the image, but the ringing patterns remain parallel to the principal axes and not perpendicular to the edges of the object. Thus, when the second image is subsequently brought into alignment with the reference object frame, the ringing patterns, initially parallel to the principal axes in the misaligned images, will become "misaligned." This is appreciated in Figure 5.1.2C, where the image of Figure 5.1.2B has been rotated back through twelve degrees using the frequency regridding method of Chapter 4. Although the prism itself is now in perfect alignment with the prism of Figure 5.1.2A, the original sinc patterns now ring away from the object edges at twelve degrees off axis. This is the spatial domain manifestation of the missing corners of the k-space matrix described above.

While the signal misalignment outside of the image is largely irrelevant in subsequent analysis involving only pixels within the object, it was discussed in Section 4.2.3 that this energy, together with ringing within the object, can bias the estimation of small misalignment angles by leading to erroneous frame-dependent correlation between two misaligned images. More importantly, even if the misalignment angle can be exactly determined, the inward ringing patterns after correction will be different between the reference image and the corrected image, providing another source of rotation-angle-dependent signal variation in the registered data set. In subsequent temporal analysis, this is translated into another source of registration-related variability. Furthermore, these differences will be most important in the vicinity of high-contrast edges where ringing is most pronounced. Unfortunately, the outer cerebral cortex is one of the brain components most often studied with functional MRI. As the cortex constitutes primarily the outer few millimeters of the brain volume, it is particularly susceptible to this form of signal degradation.
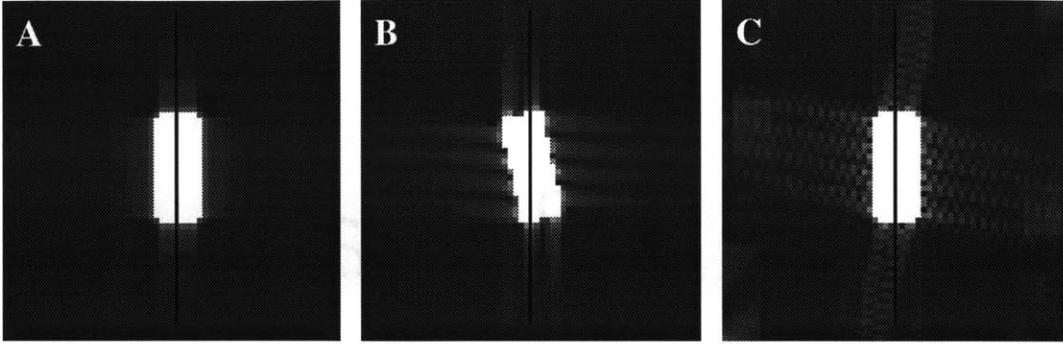
Figure 5.1.2. Simulated magnitude images through a prism windowed to accentuate ringing. (A) Reference image. (B) Rotated image; ringing is not along the axes of the prism, but along the axes of the frame. (C) Image from (B) following registration. Note the original ringing streaks have been correspondingly rotated off axis, a source of error. Also, note new ringing introduced along principal axes. The black vertical lines have been added to demonstrate the y axis direction and are not part of the original images.

5.1.3 Correction-dependent artifacts

While the acquisition-dependent errors described in the previous section are unavoidable even with perfect resampling, the second source of error is intimately linked to the resampling method itself. If perfect resampling were possible, then this error type stems entirely from the "clipping" of the rotated Fourier data set to the reference sampling region as described in Section 5.1.1. As this corresponds to the elimination of data acquired on the rotated frame which does not fall within the desired sampling region in the reference frame, it can be simply regarded as yet another multiplication by the windowing function, this time aligned with the object reference frame. As such it will generate new reference frame-dependent ringing artifacts along the major axes of the corrected image, such as those seen in Figure 5.1.2C. These new artifacts are superposed upon the off-axis, object-independent, rotation-angle-dependent artifacts described above. Since the original outward ringing pattern has been forced to positive values during the magnitude step of the initial reconstruction, the new magnitude reconstruction demonstrates the high frequency oscillations more typical of Gibbs ringing where the original and new ringing patterns overlap. Once again, however, the signal decays gradually without noticeable oscillations in regions where the overlap was minimal, such as upwards from the upper-left corner, a side-effect of generating a new magnitude image.

In practice, however, resampling is not perfect. Thus, correction-dependent artifacts also depend on the specifications of the resampling method itself, specifically the technique used to interpolate data onto the resampling grid. In linear methods, e.g. linear interpolation, frequency regridding, and truncated sinc interpolation, this can be viewed as one additional convolution step just prior to the clipping step. As described in Chapters 2 and 4, the convolution may be in image-space for spatially based methods such as conventional linear and truncated sinc interpolation, or in k-space for frequency based methods such as regridding. In either case, an equivalent multiplicative window in the other domain can be generated, and the characteristics of this window determine additional losses inherent in the particular interpolation method.

As was discussed in Chapter 4, the two-dimensional bilinear interpolation method commonly used in functional magnetic imaging registration algorithms is a particularly poor choice because of its severe high-frequency losses. This effect resulted in significantly greater errors introduced by rotation with bilinear interpolation as compared to frequency regridding techniques, as was seen in Figure 4.2.5. Nonetheless, because of its simplicity and extremely high computation speed, linear interpolation remains the only practical choice in many iterative methods. Extremely truncated sinc interpolation, e.g. four samples wide, has been used in iterative methods as an alternative, with some improvement over linear interpolation (Hajnal 1995). However, as discussed in Section 2.3, interpolation methods with less high-frequency attenuation, such as frequency regridding or less truncated sinc interpolation, should be used when possible to limit avoidable correction-dependent errors.

5.1.4 Aliasing artifacts

For simplicity, the above discussion was generalized to continuous data within the sampling region. In this section, the aliasing introduced into images by discrete sampling of the k-space matrices is discussed. Discrete sampling can be treated as a new step following the continuous windowing of the Fourier data within the sampling region. The sampling grid can be represented as an infinite rectangular grid of impulse functions in Fourier space. Thus k-space sampling over a finite region can be treated as sequential multiplication of the full Fourier representation of the imaged object by (1) the windowing function described above and (2) the sampling grid. This multiplication produces a finite array of points in frequency space corresponding to the sampled k-space matrix. As demonstrated above, the first multiplication introduces frame-dependent ringing which extends beyond the object's boundaries. The second step is equivalent to convolution of this ringing image with an infinite rectangular grid

51

of impulse functions separated by a "field of view" along each axis as determined by the frequency step size.

In practice, the field of view is chosen to be wider in each direction than the imaged object's dimensions, preventing massive aliasing of the image data. However, the field of view is necessarily smaller than the extent of the infinite ringing patterns, which leads to aliasing in the sampled images. This aliasing is evident in Figure 5.1.2B, where the ringing artifact can be seen wrapping from top to bottom and left to right, in spite of an otherwise adequate field of view. Aliasing is also present in Figure 5.1.2A, but cannot be appreciated due to the symmetry of the image. Fortunately, the field of view in two-dimensional and multi-slice methods is generally large enough that most aliasing energy is confined to the regions outside the head. Thus, aliasing is not as important a source of errors as the other sources described above. However, aliasing may become important in echo volumar imaging and other three-dimensional methods which employ out-of-plane phase-encoding, particularly with localized acquisitions where only a section of the brain, extending to one or more edges of the excited volume, is imaged (Yang 1997A). Aliasing effects must also be considered when selecting an interpolation filter or performing other spatial filtering, such as with the method described in the following section.

5.2 SPATIAL FILTERING TO REDUCE RESAMPLING ARTIFACTS

One approach to the reduction of acquisition-dependent registration noise is the introduction of a tailored post-registration spatial filter designed to minimize the high-frequency variability introduced during registration. Specifically, this filter would selectively attenuate the regions of k-space near the corners of the sampling region, i.e. those regions which become undefined following the correction of small rotations. As seen in Figure 5.1.1B, the intersection of the valid resampling regions in all possible orientations is an inscribed circle. Thus, a circularly symmetric low-pass filter represents the extreme solution. However, as this full range of motion does not occur in human FMRI experiments, such a filter would unnecessarily reduce effective image resolution. A better low-pass filter design that approximated the contours of the valid regions over physiologically expected rotations would preserve information otherwise lost by a circularly symmetric filter, yet attenuate regions which may be inconsistently resampled within the image set.

5.2.1 Filter design

Based upon the considerations discussed above, a post-registration spatial filter was constructed to attenuate affected regions over experimentally expected range of rotations. This filter was realized using the two-dimensional frequency transformation method with the McClellan transformation kernel (Lim 1990). In this method, the frequency characteristics of a one-dimensional filter are extended to the two-dimensional contours defined by the transformation. As seen in Figure 5.2.1, the k-space contours specified by the McClellan transformation are suitable for selective attenuation of the undefined regions near the corners of k-space for small rotations. The transformation was based on a one-dimensional Kaiser-Bessel low-pass filter with pass-band $\omega_p = 0.85\pi$, stop-band $\omega_s = 0.97\pi$, pass-band variation of less than 0.02 dB, and stop-band attenuation of at least 60 dB (Oppenheim 1989). The width of the transition band reduces further ringing at a small cost in resolution, while the very flat pass-band limits image degradation. The frequency transformation of this filter was performed using the Clenshaw recurrence algorithm (Press 1994), yielding a finite impulse response filter extending 63 x 63 pixels. This filter size limits the introduction of additional aliasing artifacts as image manipulations on 64 x 64 pixel images are performed after zero padding to 128 x 128. The frequency response of the filter is plotted in Figure 5.2.2. The volume under the squared frequency response curve is 0.725 relative to unit volume for a unit impulse identity filter.

0.8

0.6

0.4

0.2

0

-0.2

-0.4

-0.6

-0.8

-1

-1   -0.5   0   0.5

.97pi

.85pi

.8pi

.6pi

.4pi

.2pi
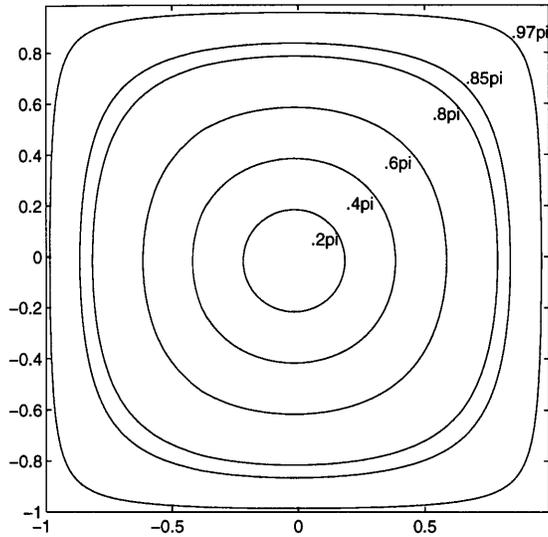
Figure 5.2.1. Frequency contours of the McClellan transformation kernel. The frequency response of the one-dimensional filter is mapped into two-dimensions along these contours using the frequency transformation method. The horizontal and vertical axes represent frequency in x and y directions, respectively, in units of radians divided by pi. The contours corresponding to the edges of the pass-band $\omega_p = 0.85\pi$ and stop-band $\omega_s = 0.97\pi$ of the one-dimensional filter to be transformed are also shown. Note that the stop band includes primarily the undesired corners of k-space.
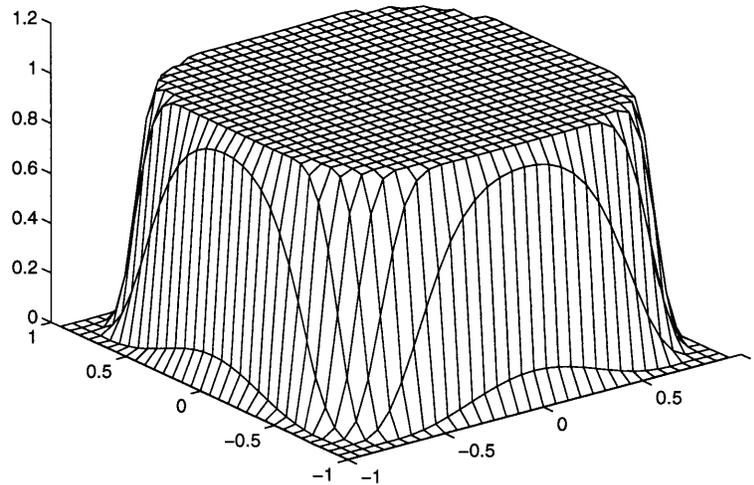
Figure 5.2.2. Frequency response magnitude of the tailored two-dimensional spatial filter. The filter specifications, described in Section 5.2.1, were chosen to maximize attenuation near the corners of k-space while preserving the remainder of k-space. The frequency axes are plotted in units of radians divided by pi.

Software was written in the standard C computing language to generate the filter and to automate its application to FMRI data sets. The code segments related to filter realization are presented in Appendix A.

5.2.2 Experimental Methods

To quantitatively characterize registration-related noise, data was collected with a 1.5 T General Electric Medical Systems Signa scanner equipped with an Advanced NMR echo planar gradient set. A set of 32 echo planar images (gradient echo, TR = 1 s, TE = 40 ms, flip angle = 75 degrees, slice thickness 7 mm) was acquired from a 37 year-old male volunteer with a quadrature head coil in an axial slice through the thalamus as the subject slowly rotated his head at a constant velocity. The images were registered using the DART registration algorithm with the enhancements described in Chapter 4.

A matching set of 32 simulated MR images was synthesized consisting of identical simulated images through a perfectly homogeneous cylinder. The simulated set was

rotationally corrected by a linear function of angles matched to the registration results of the first human data set. As this object would theoretically generate the same k-space matrix data regardless of rotation, any differences in the corrected image set must be due entirely to rotational correction.

For each data set, the registration-related noise in each motion-corrected image was computed as the root mean square (RMS) difference between each registered image and the reference image. The summation was performed only at pixels exceeding a minimum mean threshold selected to exclude the background, such that measured errors would reflect only those pixels which would normally be analyzed with conventional statistical methods. Following this computation, both image sets were filtered using the post-registration spatial filter described above and the resampling noise recomputed.

Finally, to more directly assess the importance of registration-related noise in functional MRI experiments without intentional motion, a set of 640 echo planar images (gradient echo, TR = 1s, TE = 40 ms, flip angle = 75 degrees, slice thickness 7 mm) were collected from the same 37 year-old male volunteer in an axial plane superior to the corpus collosum. The subject, whose head was secured with padding and medical tape prior to imaging, was instructed to remain motionless throughout the 10 minute and 40 second acquisition period. This data set was also registered and the frame-by-frame noise computed as above, with and without post-registration spatial filtering.

5.2.3   Results

The reference image from the human data set, where the subject was instructed to rotate slowly in the axial plane, is presented in Figure 5.2.3A. The angular displacements detected by the registration algorithm for the entire data set are shown in Figure 5.2.4A. The subject performed the task well, rotating nearly linearly through four degrees. A characteristic motion corrected image from the human data set is shown before and after post-registration filtering in Figures 5.2.3B and 5.2.3C, respectively. These images corresponded to a corrected rotation of +1.06° relative to the reference. High-frequency artifacts were visible throughout the corrected image prior to filtering. Following filtering the anatomic features were highly preserved and the high-frequency noise was greatly reduced. The broad extent of the artifacts and their reduction with filtering are more easily seen in difference images between the corrected images and the reference image as in Figures 5.2.3D and 5.2.3E.
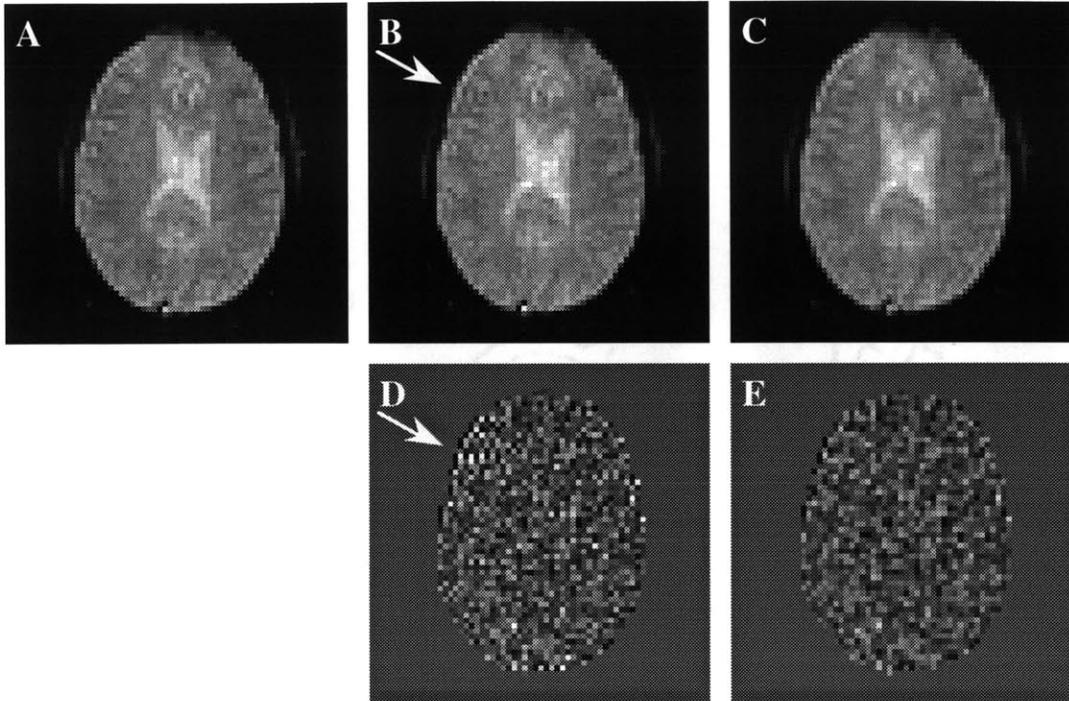
Figure 5.2.3. Reduction of high-frequency artifacts with tailored spatial filtering. (A) Reference image. (B) Image corresponding to approximately +1° of corrected rotation without post-registration spatial filtering. A comparison of (A) and (B) reveals high frequency artifacts throughout the corrected unfiltered image, most prominent at the upper left (arrow). (C) Image (B) after filtering. A comparison of (A) and (C) reveals minimal degradation of anatomic detail with filtering. (D) Difference image between corrected image following motion correction only. (E) Difference image following motion correction and filtering. Images (D) and (E) are masked to include only pixels within the brain, and are scaled from -5% (black) to +5% (white) as a percentage of the mean brain signal. The extent of the high frequency artifacts is more obvious in the difference image (D). As expected, the high frequency noise is greatly reduced in (C) and (E) following spatial filtering.

The RMS difference error between the corrected and reference images with and without post-registration spatial filtering is plotted as a function of rotation angle for the experimental data set in Figure 5.2.4B. Values are plotted as a percent of the mean brain signal. As there was no observable difference for experimental values at equivalent positive and negative rotations, the error is plotted against absolute rotation magnitude. A linear regression was performed to determine noise as a function of displacement angle, also shown in Figure 5.2.4B. The fitted estimated errors were $0.69|\psi|+2.93$ and $0.46|\psi|+2.45$ before and after

post-registration filtering, respectively. The reductions in intercept and slope with filtering were 16.4% and 32.5%, respectively.

The unfiltered intercept can be interpreted as an estimator of the frame-to-frame variation that would occur in the absence of subject motion due to scanner noise and background physiologic fluctuations. If this noise were purely white, one would expect the reduction of the intercept with filtering to be 24.7%, corresponding to the fraction of the unit volume *not* included under the curve of the squared frequency response. The 16.4% reduction suggests predominantly white noise, but with some degree of spatial correlation as well, as would be expected with background physiologic fluctuation (Biswal 1995). If the filtering achieved noise reduction solely by reducing frame-to-frame noise unrelated to motion, one would expect a reduction in intercept alone. However, the 32.5% reduction of the slope demonstrates the effectiveness of the filter in reducing noise related to registration, as it represents a systematic reduction in noise as a function of rotation.
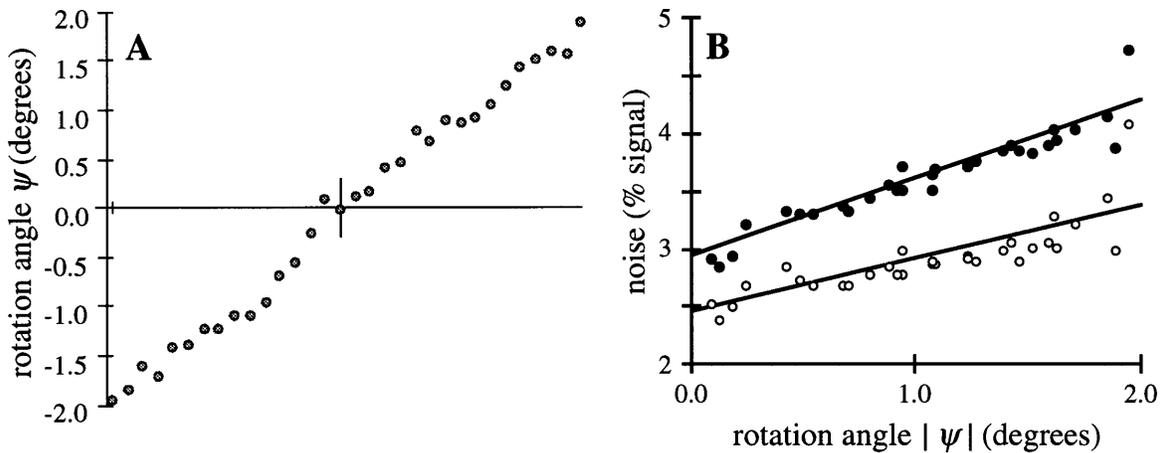
Figure 5.2.4. (A) Experimentally observed rotation angles. Horizontal axis represents sequential frames in the series of 32 images. The reference frame is marked by the vertical hash and defines zero degrees rotation. (B) Root mean square difference error as a percentage of mean brain signal vs. absolute angle of corrected rotation before (solid circles) and after (open circles) post-registration spatial filtering. The 32.5% reduction in the slope of the fitted line following filtering is a measure of the method's performance.

Based on the observed motion in Figure 5.2.4A, the set of 32 identical simulated cylinder images was rotated between -2° and +2° in equal increments to further isolate the

effects of rotational correction. The simulated cylinder image before simulated correction is shown in Figure 5.2.5A. The inward Gibbs ringing pattern caused by the rectangularly windowed sampling of k-space is easily visible in this image because of its regular contour and simulated homogeneity. Also, unlike Figure 5.1.2, the contrast of this image is adjusted to highlight intensity variation *within* the object of interest. Figures 5.2.5B and 5.2.5C show the difference errors introduced following simulated rotational correction of +1° using the regridding method. These images are shown before and after post-registration spatial filtering, respectively. The significant reduction in high-frequency noise following filtering is evident, particularly along the major diagonals of the image.
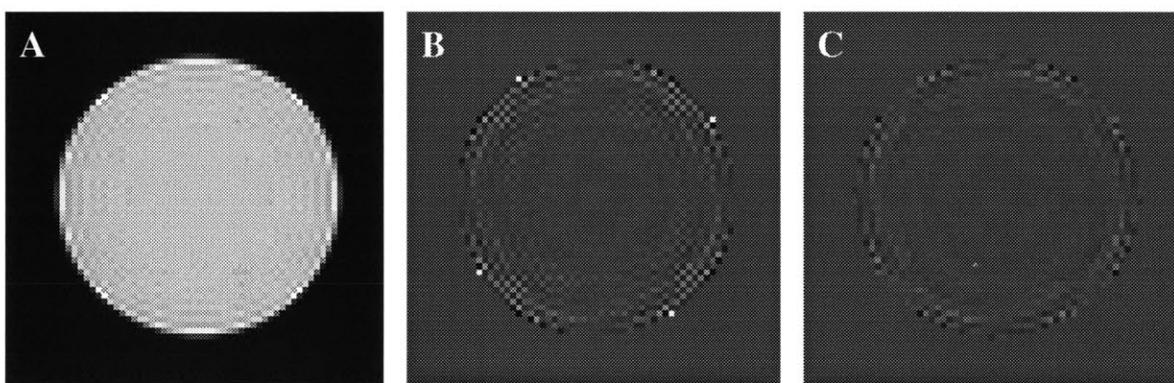
Figure 5.2.5. (A) The simluated image through a homogeneous cylinder. Note the clearly visible ringing within the object. (B) Difference image between cylinder after simulated correction of 1° and the reference frame. Only pixels within the cylinder are shown. Note the high frequency noise pattern near the rim of the object. (C) Difference image following post-registration spatial filtering of the rotated data set. There is a significant reduction in the high-frequency noise throughout. The scale in both images is from -5% (black) to +5% (white) of simulated cylinder intensity (DC).

The root mean square difference noise in the simulated data sets with and without post-registration spatial filtering is plotted as a function of rotation angle in Figure 5.2.6 for frequency regridding and bilinear interpolation. Prior to filtering, bilinear interpolation resulted in the addition of approximately three times the error as frequency regridding, consistent with the lower fidelity of the bilinear interpolation technique. For each method, the reductions in error with filtering were nearly identical at all angles studied, averaging 49.6% ± 0.02% for frequency regridding, and 24.5% ± 0.001% for bilinear interpolation. In these noiseless

examples, the intercept is zero in both cases. Thus, the error reduction of 49.6% in these data can be interpreted similarly to the slope reduction seen in the experimental data above. A larger improvement is seen in this idealized data set as compared to the experimental human data above, which may reflect differences in the frequency content of the simulated data set as well as the introduction of other errors related to angular displacement in the experimental data. The latter type of artifact may include phenomena related to out-of-plane motion.

Although the magnitude of the error reduction with filtering was larger for bilinearly interpolated data, a larger percentage reduction was observed for regridded data. This is not surprising since, as regridding more effectively preserves the high-frequency content near the corners of the sampling matrix, a larger portion of the errors introduced by regridding should be related to the corner-clipping phenomenon and thus be suited to reduction with the described spatial filter. Nonetheless, the significant reductions of errors seen following the filtering of data resampled with bilinear interpolation suggest that filtering is useful in this setting as well.
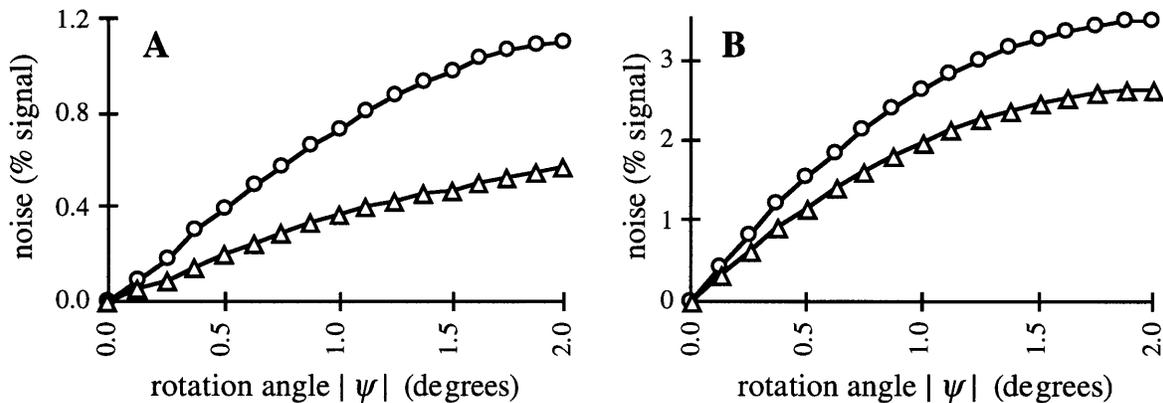
Figure 5.2.6. Root mean square difference error, i.e. registration-related noise, as a percentage of the signal vs. rotation angle in simulated images through a perfect cylinder corrected for rotation. Circles: before spatial filtering. Triangles: after spatial filtering. (A) Frequency regridding, average reduction 49.6%. (B) Bilinear interpolation, average reduction 24.5%.

The observed motion in the 640 image data set is plotted in Figure 5.2.7A. A slow unintentional rotation by the subject through approximately 0.6° is seen during the 10 minute and 40 second acquisition. The root mean square error between registered images and the reference image, computed over brain pixels as determined by a minimum mean threshold, is

shown in Figure 5.2.5B before and after spatial filtering. The fitted estimated errors were $1.31|\psi| + 2.89$ and $1.08|\psi| + 2.48$ before and after filtering, respectively. The reductions in intercept and slope with filtering were 14.0% and 17.4%, respectively. The intercept values and the reduction with filtering are similar to those seen above for a larger range of rotation. This is not surprising given that these values represent an estimate of the intra-frame variability due to all factors excluding motion, which would not be expected to change in a longer acquisition covering a smaller range of rotations. The slopes demonstrated before and after filtering, however, are both somewhat larger than those seen in the previous human experiment with intentional motion over a larger range of rotations, and the reduction is somewhat smaller. A larger slope is also evident in the simulated data of Figure 5.2.6 for the smaller angles of rotation. The smaller reduction may reflect the additional influence of other factors during this prolonged acquisition, such as out-of-plane motion or long-term changes in background physiologic signal variation.
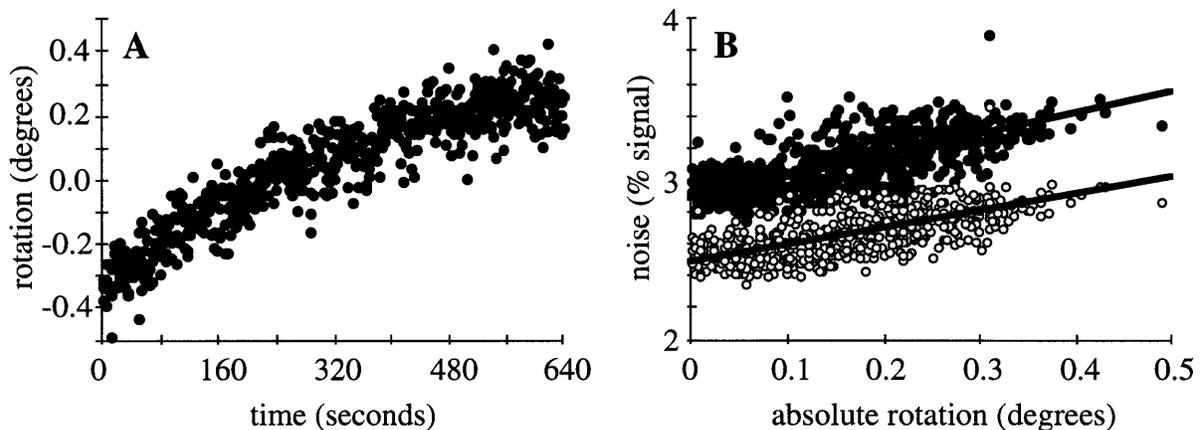
Figure 5.2.7. Spatial filtering in a prolonged FMRI experiment. (A) Observed rotation as a function of time for the 640 second acquisition. (B) Root mean square difference error between registered images and the reference image as a percentage of mean brain signal vs the absolute magnitude of the estimated rotation angle before (solid circles) and after (open circles) spatial filtering. There is a 17.4% reduction in the slope of the fitted line following spatial filtering.

61

5.2.4   Discussion

The data presented suggest that errors introduced into functional magnetic resonance imaging data sets during image registration are not entirely benign. As seen in Figure 5.2.6, for rotations as small as one degree, the variability introduced is 0.74% of the mean signal with regridding methods, or as large as 2.62% with bilinear interpolation. As these noise levels are well on the order of the signal changes sought in BOLD functional experiments at 1.5 Tesla, and the noise increases roughly linearly with rotation angle, attempts to reduce this noise are justified.

The premise behind the filtering technique presented here is that for rectangular k-space sampling matrices, the regions of the matrix near the corners of the sampling region, which are most sensitive to the small rotations seen in functional MRI, are an important source of variability in the registered data sets. Thus, a specialized filter designed to selectively attenuate these frequencies was constructed based on the McClellan transformation of a Kaiser-Bessel low-pass filter. While any number of digital filter design approaches would be suitable, this method was selected because of its finite impulse response (FIR) characteristics, ease of implementation, desirable attenuation contours, and small number of design parameters. A directed frequency-space design, while appealing for its simplicity, would not necessarily guarantee FIR filtering, and might lead to additional unintended aliasing noise, limiting its effectiveness.

The choice of parameters for the underlying Kaiser-Bessel filter represents a limitation of the present design required to keep the filter size small in order to prevent aliasing, and to reduce ringing. The Kaiser-Bessel filter provides a very flat pass-band, limiting distortion in the output, but at the cost of a small reduction in spatial resolution. In the implementation described, the pass-band extended to $\omega_p = 0.85\pi$, corresponding approximately to a 15% reduction in resolution which must be considered in subsequent statistical analysis. Thus, for the detection of very small regions of activation, the reduction in variability may be counteracted by the reduction in effective resolution, complicating the detection of very small regions of activation (Lowe 1997). Further work is necessary to determine whether a threshold range of rotations exists below which filtering is contraindicated.

Importantly, the reductions in variability observed following spatial filtering can be quite large, as seen in Figures 5.2.4 to 5.2.6. For the simulated data set, where all errors are necessarily due to registration, the application of this filter reduces resampling noise by nearly half when used in conjunction with frequency regridding. The improvement is approximately constant over the range of rotations examined. The large reduction is consistent with the high fidelity of regridding in the preservation of high spatial frequencies and the relatively small

amount of aliasing introduced by this method. Similar results would be expected with other high-fidelity interpolation methods such as truncated sinc interpolation. In contrast, spatial filtering yields only a 24.6% reduction when used with bilinear interpolation, which is consistent with the larger degree of high-frequency loss and aliasing inherent to resampling with bilinear interpolation. Both the smaller initial errors and the larger percentage reductions seen with regridding suggest that higher-fidelity resampling methods should be used where possible. Nonetheless, the larger magnitude reduction seen with bilinear interpolation supports the use of filtering in this setting as well.

Reductions in variability were also observed when the post-registration spatial filtering technique was applied to experimental human data. As described above, the baseline noise seen in Figure 5.2.4 represents the combination of scanner noise and background physiologic variation. The effectiveness of spatial filtering in the reduction of registration-related noise is measured by the 32.5% reduction in slope of the noise as a function of absolute rotation. Similarly, the slope of the error in Figure 5.2.7 was reduced by 17.4% following filtering. These improvements are comparable to, although somewhat less than, the 49.6% reduction seen in the simulations. This likely reflects small errors in the estimation of rotation angles in the experimental data, leading to residual misregistration, as well as contributions from additional noise sources which may also be related to in-plane rotation, such as out-of-plane motion. Fortunately, the principles of this method can be easily extended to three-dimensions. However, additional work is necessary to determine whether this problem is equally important with three-dimensional image registration algorithms.

Finally, it is important to note that the ultimate manifestation of registration-related noise in functional MRI analysis is in the temporal domain. Application of post-registration spatial filtering should reduce the temporal variability of time sequences which have been corrected for rotational errors, potentially increasing the power to detect activation, and complements temporal approaches to correct for motion-related errors. Importantly, in cases of stimulus-correlated motion (Hajnal 1994), the resampling errors at any given pixel may also be stimulus-correlated, such that post-registration spatial filtering will reduce false activation related to this problem. As described above, however, spatial filtering does reduce the effective resolution of the images, which may potentially counteract the reduction of variability in the detection of very small regions of activation. Additionally, the increased correlation between time series at neighboring pixels must be considered in subsequent statistical analysis.

5.3 TEMPORAL FILTERING TO REDUCE RESAMPLING ARTIFACTS

As described above, errors introduced during registration are inherently spatial artifacts, and thus are suited to direct reduction by spatial techniques, such as the post-registration spatial filtering technique of the previous section. However, it has also been stated that the artifacts are ultimately manifested as temporal noise during later analysis. Given the nature of these artifacts, it is reasonable to suspect the noise at any given pixel may be related, at least in part, to the degree of misalignment. In this section, methods are presented to model noise following registration as functions of rotation angle in the time domain and to remove these registration-related temporal artifacts efficiently from image sets by regression techniques.

5.3.1 An orthogonal model

To characterize registration-related noise in the time domain, the following simplified additive model for the pixel time series $p_{x,y}[n]$ at position x,y of image n will be used:

$$p_{x,y}[n] = \bar{p}_{x,y} + a_{x,y}[n] + r_{x,y}(\psi[n]) + e_{x,y}[n] \qquad [5.3.1]$$

where $\bar{p}_{x,y}$ is the true "baseline" value at the pixel, $a_{x,y}[n]$ is an overlying, possibly null, "true" activation time series related to a stimulus or task, $r_{x,y}(\psi[n])$ is the registration-related noise term at each pixel dependent on the rotational displacements $\psi[n]$, and $e_{x,y}[n]$ is an additional noise term related to scanner and background physiologic noise..

The registration-related noise term $r_{x,y}(\psi[n])$ is further modeled as a linear combination of M orthonormal basis functions $\phi_{1...M}[n]$ which are functions of $\psi[n]$. In this section, principal component analysis will be used to extract the temporal basis functions explaining the largest variance in the image data set, and analytic functions of $\psi[n]$ sought to model the observed eigenvectors.

5.3.2 Experimental Methods

Data was collected with a 1.5 T General Electric Medical Systems Signa scanner equipped with an Advanced NMR echo planar gradient set. A set of 32 echo planar images (gradient echo, TR = 1 s, TE = 40 ms, flip angle = 75 degrees, slice thickness 7 mm) was acquired from a 37 year-old male volunteer with a quadrature head coil in an axial slice through the thalamus as the subject slowly rotated his head at a constant velocity. This data set was described in Section 5.2.

As in section 5.2, a matching set of 32 simulated MR images was synthesized consisting of identical simulated images through a perfectly homogeneous cylinder. The simulated set was rotationally corrected by a linear function of angles matched to the registration results of the first human data set. As this object would generate the same k-space matrix data regardless of rotation, all data should be identical before and after rotational correction. Thus, any temporal noise observed in the corrected synthetic image set must be due entirely to the rotational correction process.

In both sets of images, the resampling noise was modeled as a linear combination of orthonormal basis functions as described above. Principal component analysis (PCA) was performed to extract the temporal basis functions explaining the largest variance in each data sets. Data were entered into an R-mode analysis (Basilevsky 1994, Maas 1997A) with pixels as observations and time-points as variables. To perform the PCA, the eigenvalues and eigenvectors of the sample temporal covariance matrix were computed using the TRIRED and TRIQL functions as implemented in the Interactive Data Language software package (Research Systems, Inc., Boulder, CO). The TRIRED routine reduces the covariance matrix to tridiagonal form using Householder's method and the TRIQL routine uses the QL algorithm to determine the eigenvalues and eigenvectors. For a detailed description of these algorithms, the reader is referred to Press et al. (1994). For each component, the eigenvalue is equal to the total variance captured by component, and the eigenvector indicates the component's temporal pattern.

Following principal component analysis, an analytic model was selected to approximate the most significant observed eigenvectors as functions of the rotational angle. The model consisted of mutually orthogonal terms based on the normalized Legendre polynomials (Press 1994). The terms are given by:

$$P'_k[n] = P_k\left(\frac{2(n-1)}{N-1} - 1\right)$$

[5.3.2]

where $P_k(x)$ is the kth Legendre polynomial.

A final data set of 640 images was collected as in Section 5.2. During this acquisition, the subject was instructed to remain motionless. The images were registered using the algorithm of Chapter 4. The post-registration spatial filter was not applied. Following registration, the first six terms of the polynomial noise model were tested in all image pixels using linear regression techniques, with significance defined at the $p = 0.05$ level.

5.3.3 Results

The angular displacements detected by the registration algorithm for the first human data set, where the subject was instructed to rotate at a constant rate, were presented in Figure 5.2.4A. The subject succeeded in estimating linear rotation through a range of approximately 4 degrees. The simulated data set through a perfect cylinder was synthesized and corrected for rotation with a linear function through 4 degrees using the regridding technique described in Chapter 4 yielding a matched data set.

The results of the principal component analysis are presented in Figure 5.3.1. The six experimental eigenvectors computed from the registered experimental human data set (dots) accounted for 26.2%, 11.4%, 7.7%, 5.0%, 3.3% and 2.6% of the image set variance, for a total of 56.1%. PCA of the simulated data set yielded similar results for the largest eigenvalues (dashed lines). The six eigenvectors accounted for 61.1%, 27.0%, 6.0%, 2.0%, 1.2% and 0.6% of the simulated image set variance, for a total of 97.7%. The similarity of the six eigenvectors shown for both the experimental and simulated data sets supports the hypothesis that the components are related to resampling noise and not to other noise processes. The emergence of a small number of dominant eigenvectors supports the choice of an incomplete set of orthogonal basis functions to model registration noise. Furthermore, the data support the empirically selected model based on the normalized Legendre polynomials.
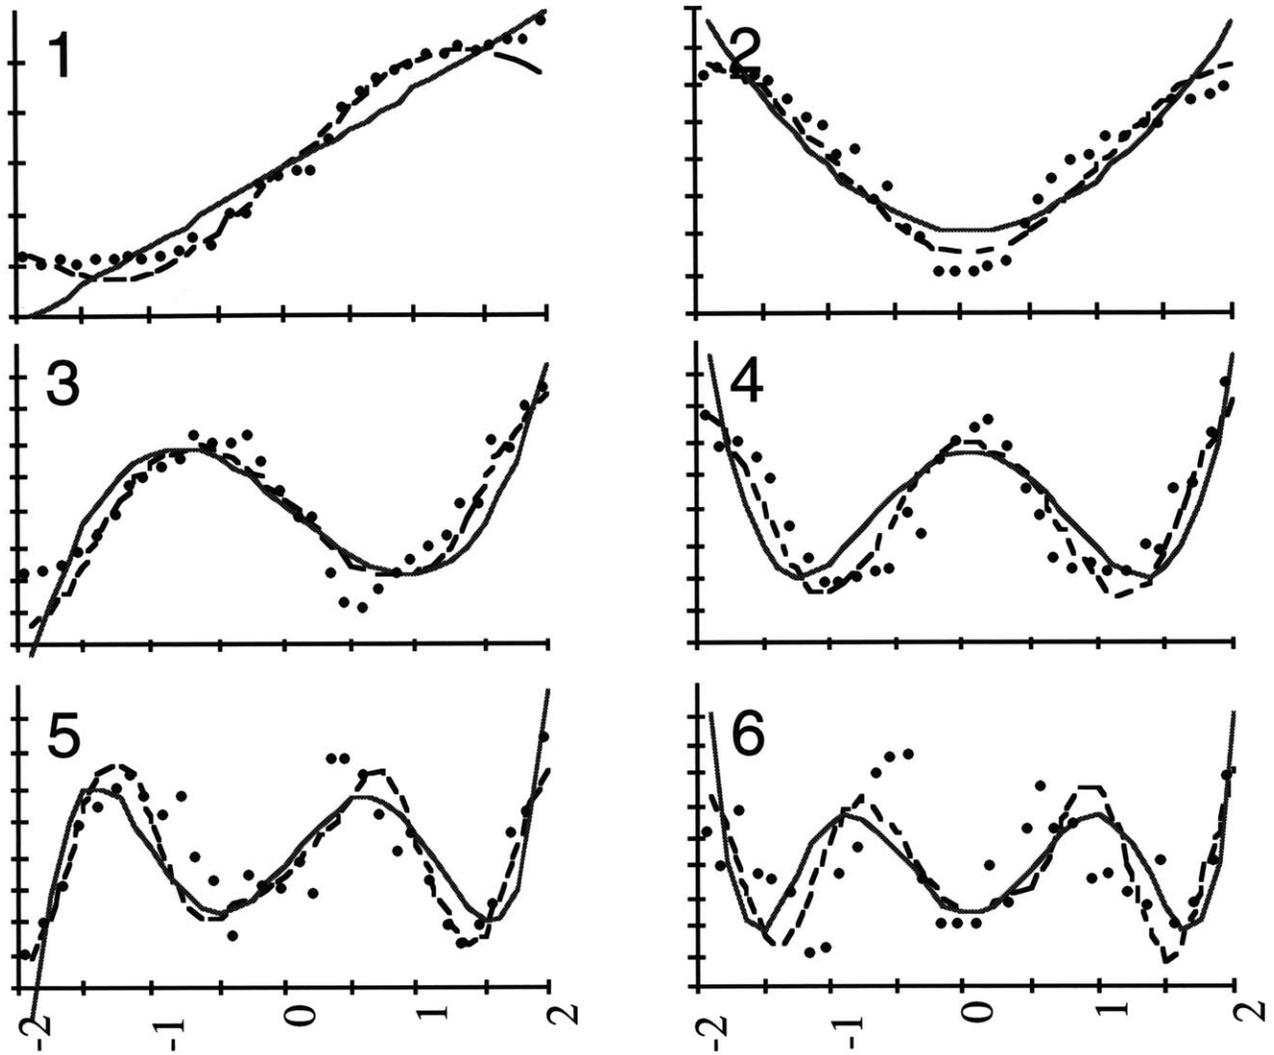
Figure 5.3.1. Temporal data representing the first six eigenvectors of the simulated data set (dashed line) and the experimentally observed human data (dots), and the first six terms of the Legendre polynomial model (solid line). X-axis: degrees of rotation.

Using the polynomial model, the significance of registration noise was assessed in the final functional MRI experiment where the subject was instructed to lie motionless during data collection. As previously shown, image registration revealed that the subject had in fact moved with a slow rotational drift covering 0.6° during data acquisition. Analysis of the registered data with multiple linear regression using the first six terms of the Legendre polynomial model revealed that the registration noise model was significant, capturing 18.8% of the total variance after registration, in the *absence* of any deliberate motion. Additionally, the first four terms

were significant when taken individually. Finally, the data set was divided into two 320 image data sets each covering approximately 0.3° of rotation, and four 160 image sets each covering approximately 0.15° of rotation, which were analyzed in the same fashion. The model was statistically significant in all cases. Average results for each image size are presented in Table 5.1.

Table 5.1. Average energy in the first six normalized Legendre polynomial terms for brain pixels from experimental data sets. † arbitrary units (DC = 361)

Data segment	Short	Medium	Long
points/segment	160	320	640
number of segments	4	2	1
approx rot/seg	0.15°	0.30°	0.60°
%var P'1	4.74 %	8.52 %	11.0 %
%var P'2	1.36	2.12	3.84
%var P'3	1.06	0.97	1.79
%var P'4	0.80	0.79	0.98
%var P'5	0.76	0.54	0.65
%var P'6	0.69	0.46	0.55
full model	9.42	13.4	18.8
regression var †	5.22	7.97	12.2
residual var †	50.2	51.6	52.6
total var †	55.4	59.6	64.8
F	2.65	8.06	24.5
df v1, v2	6, 153	6, 313	6, 633
p	<0.05	<<0.01	<<0.01

Not surprisingly, the total variance increases with the range of rotation. This is due, in part, to the greater number of images rotated through larger angles during resampling. Other factors, such as the increased probability of out-of-plane motion over the longer sampling times, may also contribute to the increased variance. With increasing rotational range, the variance captured by the normalized Legendre polynomial model also increases accordingly, while the residual variance remains relatively constant. Although a number of sources may contribute to the linear term, this finding supports the hypothesis that as the rotational range increases, the relative importance of the registration noise also increases.

Additionally, as the sampled normalized Legendre polynomials are nearly orthogonal for the values of N examined, the significance of each term can be considered separately, and tested statistically in the presence of the other variables using an F test with $v_1 = 1$ and v_2 taken from Table 5.1. Defining significance at the $p = 0.05$ level, the linear term is significant in the 0.15° case, the first and second order terms are significant for the 0.3° case, and all of the first four terms are significant in the 0.6° case. Both the data length and the rotational range influence the importance of each of the terms of the registration noise model.

While the above results characterize the significance of the registration noise model over the entire brain, another consideration is the importance of the registration noise at individual pixels. Considering that a typical activation in a functional experiment is on the order of 1% peak-to-peak, corresponding to a variance contribution of $1/2\sqrt{2}$ % of the mean signal, it is meaningful to count the number of brain pixels where the model exceeds this threshold for one or more terms. This data is presented in Table 5.2. The large number of pixels affected indicate that registration-related noise is likely to be important in typical functional MRI experiments. Again, the increasing importance of higher order polynomial terms is evident as rotational range increases. Furthermore, as the total variance was shown in Table 5.1 to increase slowly with increasing rotational, this observation is strengthened by the reduced likelihood of spurious energy in each term as N increases, reducing the number of pixels expected by chance.

Table 5.2. Percentage of brain pixels where the RMS energy explained by each model term exceeded $1/2\sqrt{2}$ % of the average brain signal, the equivalent energy of a sinusoid with peak-to-peak amplitude of 1% DC.

Data segment	Short	Medium	Long
points/segment	160	320	640
number of segments	4	2	1
approx rot/seg	0.15°	0.30°	0.60°
%pixels P'1	38.0 %	52.4 %	56.7 %
%pixels P'2	11.9	25.5	37.1
%pixels P'3	8.66	10.2	21.2
%pixels P'4	5.73	6.11	9.68
%pixels P'5	4.59	3.25	6.11
%pixels P'6	3.76	1.53	4.39
% by chance	< 5 %	< 1 %	<< 1 %

5.3.4 Discussion

In Section 5, it has been shown that the correction of slow in-plane rotational drifts in functional MRI data sets may be associated with a characteristic residual temporal noise pattern, which can be modeled as an additive linear combination of orthonormal basis functions. This model was supported by principal component analysis of both experimental and simulated data sets, which yielded small numbers of significant basis functions. The similarity of the eigenvectors extracted from both the experimental and simulated data sets indicates that they are in fact related to registration noise and not other noise processes, and that the simulated model can be used to gain additional insight into the nature of the noise process. Examination of the principal component results further supports the modeling of basis terms as normalized Legendre polynomials. While this choice is convenient due to the mutually orthogonal and easily computable nature of the terms, it is important to note that any set of polynomial basis functions is equally suitable for this method if regression is properly performed.

The importance of noise resulting from image registration was demonstrated in a functional MRI experiment where the subject was instructed to lie motionless without time-varying stimuli during data collection. As discussed in Chapter 2, subject motion is nearly ubiquitous in functional MRI, even in the absence of stimuli, and subsequent image registration in this case revealed that the subject had nonetheless moved with a slow rotational drift covering $0.6°$ during data acquisition. Analysis of the registered data with multiple linear regression using the first six terms of the Legendre polynomial model revealed that the registration noise model was significant for data corrected for rotation spanning ranges of $0.15°$, $0.3°$ and $0.6°$. In these cases, the six-term model captured 9.42%, 13.4%, and 18.8% of the total variance, respectively. Additionally, the first four terms were found to be significant when the 640 image data set was analyzed in its entirety. Furthermore, when pixels within the brain were considered individually, a significant number of pixels demonstrated expression of each of the first six polynomial terms at a physiologically significant threshold. Taken together, these results support the conclusion that linear regression using the registration noise model presented may be useful in the analysis of functional MRI data sets demonstrating slow rotational drifts prior to image registration.

It has been previously recognized that the removal of linear trends in data, even after image registration, greatly improves the results of subsequent analysis (Bandettini 1993). The strong presence of the linear error term in the data sets described above supports the removal of such trends and provides some explanation of the nature of this improvement. It is important to note, however, that the correction of slow rotational drifts is unlikely to be the only source of linear trends. Another major source is out-of-plane drift during the experiment, supported by the correlation among the six degrees of rigid body freedom observed by others (Hajnal 1993). Additionally, regression of these trends may compensate for local deviations from the rigid-body assumptions.

A second order term, on the other hand, is not commonly considered in the analysis of functional data sets. Friston et al. (Friston 1995) have proposed a model of differential spin histories due to small shifts in the planes of excitation which supports the inclusion of second order functions of measured motion parameters, together with time-lagged first and second order terms, but the relation between second order terms and image registration has not been previously proposed. As the data above indicate that such terms may account for a significant portion of the variance in registered functional data sets, one may conclude that inclusion of a second order term, together with the linear term, may further improve the results of subsequent analysis. The sources of error proposed here, however, do not necessitate the inclusion of time-lagged terms.

The data above also suggest that the higher order polynomials may be useful with larger rotations or larger data sets. In practice, data with rotational ranges greater than a degree may be of marginal utility and are often discarded, limiting the utility of higher order terms in short data sets. However, for longer data sets, such as the 640 image data set presented here, the first four polynomial terms independently influence the data set at a statistically significant level. The choice of the optimal number of terms to use in a particular analysis may be best determined by a preliminary analysis of variance similar to that employed above.

In this work, the polynomials are computed assuming perfect linearity to the rotational drift. While this is generally a good approximation of rotations measured in actual experiments, as supported by the results of the 640 point data set presented above and the empirically established utility of removing linear trends in general use, a useful alternative would be to generate the polynomial based directly upon the estimates of rotation $\hat{\psi}$. While such methods should become more useful as the measured rotations deviate further from a linear model, it remains to be seen whether the current model can be extended to data sets with large discontinuities in ψ caused by abrupt head motion or by non-linear or piece-wise linear drift, other forms of motion commonly observed in functional MRI.

In summary, these results suggest that linear regression using the registration noise model presented may be useful in the analysis of functional MRI data sets demonstrating slow rotational drifts prior to image registration and may represent an improvement over the practice of removing linear trends from functional data before analysis. The number of terms required for optimal improvement vs. degrees of freedom lost and the applicability of this system to non-uniform motion remains to be investigated.

6 IMAGE REGISTRATION IN THREE DIMENSIONS

In Chapter 2, the problem of motion artifacts in functional MRI was introduced and several of the major factors limiting accurate three-dimensional image registration were discussed. For example, out-of-plane information is typically undersampled in multi-slice functional echo planar studies, where slice thickness of 5 or more millimeters are generally required. Also, since echo planar slices are not collected simultaneously, there exists no complete reference image set at any instant in time, making imaging volumes susceptible to intra-volume motion effects. One approach has been to avoid these difficulties altogether by limiting registration to two-dimensions, even when three-dimensional data is available. The two-dimensional registration algorithm described in Chapter 4, for instance, does register images accurately within the imaging plane even in the presence of out-of-plane motion.

In spite of the multitude difficulties, an extension of registration methods to three dimensions is still a worthwhile endeavor. While it may ultimately prove unsound to attempt direct correction for out-of-plane motion under all circumstances, it may still be possible to generate reasonably robust measures of out-of-plane motion in functional experiments. At a minimum, such measures can be used to exclude overly contaminated data. Furthermore, as fast imaging technologies continue to improve, so too will the potential quality of data registered in three-dimensions. Thus, a three-dimensional extension of the methods of Chapter 4 is presented in this chapter for the estimation and correction of rotations and translations in three dimensions.

6.1 A THREE-DIMENSIONAL IMAGE REGISTRATION ALGORITHM

The motivations for a decoupled three-dimensional algorithm are the same as those in two dimensions. As was described in Chapter 2, image-space based methods, which are by their nature fully iterative, require the simultaneous estimation of all motion parameters. As in two dimensions, the application of the decoupling principle in three dimensions allows the separation of the registration problem into two parts, rotations and translations. Thus, simpler

solutions may be sought for each type of motion. The potential gain in performance is increased in three dimensions, where the number of degrees of freedom increases to six.

One approach to decoupled three-dimensional image registration was recently described by Kassam and Woods (1996). Their algorithm estimated the three rotational parameters in two stages: (1) an overall axis of rotation, accounting for two of the three rotational degrees of freedom, was identified first, and (2) using the estimated axis, the final rotational parameter was computed by two-dimensional techniques related to those described in Chapter 4. The axis of rotation was identified by a brute-force search of k-space for a null line in the spectral magnitude difference between the reference image and the test image. This null corresponds to the line in k-space on which the two spectral magnitudes correspond exactly, i.e. where no rotation has occurred, and is guaranteed to cross through the origin. While the simplicity of this approach is appealing, Kassam and Woods found that several weighting and masking refinements were necessary to improve the stability of the algorithm, which nonetheless broke down for small displacements. Additionally, the accuracy of their brute-force search strategy is limited by the granularity with which the difference in spectral magnitudes is calculated. Furthermore, as their approach examined only single lines through the large k-space array, it would be reasonable to expect that an approach which uses more of the data would yield more stable and accurate estimates of the rotational parameters.

With these considerations in mind, a new, decoupled three-dimensional registration algorithm was designed and implemented based on the methods presented in Chapter 4. Unfortunately, in three dimensions, the problem of simultaneously estimating three rotational degrees of freedom is not transformed into a linear problem as easily as the single parameter problem of two dimensions. Thus, a partially iterative solution is required. It is expected, however, that the iterative estimation of three rotational parameters will still lead to a large-scale savings in time and computation when compared to the six-parameter image-space alternatives. Once the rotational parameters are estimated, a three-dimensional extension of the regridding algorithm described above is used to align the images, after which a noniterative three-dimensional analog of the phase-plane conjugation technique of Chapter 4 can be applied to estimate and correct residual linear translations.

6.2 THEORY

The methods presented here are largely direct extensions of those presented in Chapter 4. However, there a number of new theoretical issues which are unique to the three-dimensional estimation of rotation. These are discussed below in Sections 6.2.1 and 6.2.2.

The extension of the phase-plane method of Section 4.2.2 for translations is described in Section 6.2.3.

6.2.1 An error function for iterative estimation of rotation

As in the two-dimensional case, rotational displacements are estimated first. These estimates are computed by iterative minimization of an error criterion. The criterion selected is the sum of squared differences between the spectral magnitudes of the image to be registered and of the reference image. For convenience, the matrix of spectral magnitudes of the unregistered image will be written as the function $D[\mathbf{k}]$, where $\mathbf{k} = \begin{bmatrix} k_x & k_y & k_z \end{bmatrix}^T$ is a vector argument specifying a discrete k-space position. Similarly, the reference image spectral magnitude will be written as $Y[\mathbf{k}]$.

Rotational displacements will be represented by the vector $\mathbf{a} = \begin{bmatrix} \theta & \phi & \psi \end{bmatrix}^T$ specifying rotations θ about the k_x-axis, ϕ about the k_y-axis, and ψ about the k_z-axis. The function $\hat{Y}[\mathbf{k};\mathbf{a}]$ resamples $D[\mathbf{k}]$ at discrete values of $\mathbf{k}$ following the rotations specified by $\mathbf{a}$. Thus, for a given reference image $Y[\mathbf{k}]$ and unregistered image $D[\mathbf{k}]$, the error criterion can be written as a function of the rotation vector $\mathbf{a}$:

$$E^2(\mathbf{a}) = \sum_{\mathbf{k}' \in \mathbf{K}} \left[Y[\mathbf{k}'] - \hat{Y}[\mathbf{k}';\mathbf{a}] \right]^2 . \qquad [6.2.1]$$

In Eq. 6.2.1, $\mathbf{K}$ represents the set of values of $\mathbf{k}$ over which the error $E^2(\mathbf{a})$ is calculated. It may include all points at which the original spectral matrices are defined, i.e. the sampling points in the reference frame, or only a subset of interest, as will be described later. It is the rotation vector $\mathbf{a}$ which will be estimated iteratively to minimize the error criterion.

Although a three-dimensional extension to the regridding algorithm will be used to generate the final rotationally aligned image after convergence of the parameters, this algorithm is too computationally expensive for the iterative portion of the algorithm. Instead, the interim image $\hat{Y}[\mathbf{k};\mathbf{a}]$ will be generated using a spherically symmetric Hamming window $h(\mathbf{k}_c)$ of half-width M to interpolate the spectral data of $D[\mathbf{k}]$ at the positions determined by $\mathbf{a}$. This window is analogous to the circularly symmetric filter used in the two-dimensional algorithm, and is defined by:

$$h(\mathbf{k}) = \begin{cases} \gamma + \gamma \cdot \cos \dfrac{\pi |\mathbf{k}|}{M} & |\mathbf{k}| < M, \\ 0 & \text{otherwise,} \end{cases} \qquad [6.2.2]$$

where $\gamma = \left(\dfrac{4}{3} \pi - \dfrac{8}{\pi} \right) M^3$ is a constant chosen to give $h(\mathbf{k})$ unit mass. The round brackets and subscripted argument identify that $h(\mathbf{k}_c)$ is a continuous function also defined at points off the matrix. Using this interpolation filter, the interim image can be written as the convolution sum:

$$\hat{Y}[\mathbf{k};\mathbf{a}] = \sum_{\mathbf{k}' \in \mathbf{K}} D[\mathbf{k}'] \cdot h(\mathbf{k}_0[\mathbf{k};\mathbf{a}] - \mathbf{k}').$$

[6.2.3]

In Eq. 6.2.3, $\mathbf{K}$ may represent the full set of discrete matrix points at which $D[\mathbf{k}]$ is defined, or, with equivalent summation results, only those points where the filter term is not zero. The function $\mathbf{k}_0[\mathbf{k};\mathbf{a}]$ returns the position in the frame of reference of $D[\mathbf{k}]$ which, when rotated according to $\mathbf{a}$, maps to the matrix point $\mathbf{k}$. The value returned by $\mathbf{k}_0[\mathbf{k};\mathbf{a}]$ need not, and typically will not, lie on the matrix, but its argument $\mathbf{k}$ always will.

For the purpose of computing $\mathbf{k}_0[\mathbf{k};\mathbf{a}]$, rotations on the uncorrected image are performed in the following order: about the k_x-axis, then about the k_y-axis, and finally about the k_z-axis. Positive angles are assigned the counter-clockwise direction with respect to the axis. These rotations can be written as sequential matrix multiplication by the transformation matrices $\mathbf{T}_\theta$, $\mathbf{T}_\phi$, and $\mathbf{T}_\psi$, with $\mathbf{k}$ and $\mathbf{k}_0[\mathbf{k};\mathbf{a}]$ related by:

$$\mathbf{k} = \mathbf{T}_\psi(\psi)\mathbf{T}_\phi(\phi)\mathbf{T}_\theta(\theta)\mathbf{k}_0[\mathbf{k};\mathbf{a}],.$$

[6.2.4]

where:

$$\mathbf{T}_\theta^{-1} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\theta & -\sin\theta \\ 0 & \sin\theta & \cos\theta \end{bmatrix} \quad \mathbf{T}_\phi^{-1} = \begin{bmatrix} \cos\phi & 0 & \sin\phi \\ 0 & 1 & 0 \\ -\sin\phi & 0 & \cos\phi \end{bmatrix} \quad \mathbf{T}_\psi^{-1} = \begin{bmatrix} \cos\psi & -\sin\psi & 0 \\ \sin\psi & \cos\psi & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

[6.2.5]

Thus, $\mathbf{k}_0[\mathbf{k};\mathbf{a}]$ has the inverse solution:

$$\mathbf{k}_0[\mathbf{k};\mathbf{a}] = \mathbf{T}_\theta^{-1}(\theta)\mathbf{T}_\phi^{-1}(\phi)\mathbf{T}_\psi^{-1}(\psi)\mathbf{k},$$

[6.2.6]

6.2.2 The gradient and Hessian matrix of the error function

From Eq. 6.2.1, the gradient of $E^2(\mathbf{a})$ and a first order estimate of its Hessian matrix $\mathbf{D}$ can be derived:

$$\nabla E^2(\mathbf{a}) = -2\sum_{\mathbf{k} \in \mathbf{K}} \left[Y(\mathbf{k}) - \hat{Y}[\mathbf{k};\mathbf{a}] \right] \frac{\partial \hat{Y}[\mathbf{k};\mathbf{a}]}{\partial \mathbf{a}},$$

[6.2.7]

$$D_{ij}(\mathbf{a}) \approx 2 \sum_{k' \in K} \frac{\partial \hat{Y}[\mathbf{k};\mathbf{a}]}{\partial a_i} \frac{\partial \hat{Y}[\mathbf{k};\mathbf{a}]}{\partial a_j}. \qquad [6.2.8]$$

Since the generation of the interim image in Eq. 6.2.3 can be described analytically using the equations of the previous section, the partial derivatives of the interim image in Eqs. 6.2.7 and 6.2.8 can also be derived analytically. Thus, this problem is well-suited for the Levenberg-Marquardt minimization routine (Press 1994). In this method, the gradient and Hessian matrix of the error function at an interim value of $\mathbf{a}$ are used to compute the gradient of a truncated Taylor series estimate of $E^2(\mathbf{a})$ in the vicinity of $\mathbf{a}$, i.e.:

$$E^2(\mathbf{a} + \Delta\mathbf{a}) \approx E^2(\mathbf{a}) + \nabla E^2(\mathbf{a}) \cdot \Delta\mathbf{a},$$
$$\nabla E^2(\mathbf{a} + \Delta\mathbf{a}) \approx \nabla E^2(\mathbf{a}) + \mathbf{D}(\mathbf{a}) \cdot \Delta\mathbf{a}. \qquad [6.2.9]$$

As the gradient of $E^2(\mathbf{a})$ must be zero at a minimum, the left side of Eq. 6.2.9 can be nulled to estimate the step size $\Delta\mathbf{a}$ required to reach the minimum from $\mathbf{a}$:

$$\Delta\mathbf{a} = -\mathbf{D}^{-1}(\mathbf{a}) \, \nabla E^2(\mathbf{a}). \qquad [6.2.10]$$

By repeating these steps at each iteration, the Levenberg-Marquardt algorithm provides improving estimates of the value of $\mathbf{a}$ which minimizes the least squared error $E^2(\mathbf{a})$. To increase the stability of the algorithm, the Hessian matrix $\mathbf{D}$ is actually replaced by $\mathbf{D} + \mathbf{I}\lambda$, where λ is a weighting which is reduced towards zero when the correction $\Delta\mathbf{a}$ reduces the error, and away from zero when the error worsens, shifting smoothly between the solution of Eq. 6.2.10 for small λ and a more conventional gradient descent for large λ (Press 1994).

To implement the algorithm, the partial derivatives of Eq. 6.2.3 must be derived with respect to each of the components of $\mathbf{a}$. Including only the non-zero range of Eq. 6.2.6 in the summation, and defining the convenience function $s(\mathbf{k}_0, \mathbf{k}') = \mathbf{k}_0[\mathbf{k};\mathbf{a}] - \mathbf{k}'$, application of the chain rule yields:

$$\frac{\partial \hat{Y}[\mathbf{k};\mathbf{a}]}{\partial a_i} = \sum_{\substack{k' \in K \\ |s| < M}} D[\mathbf{k}'] \sum_m \frac{\partial h(s)}{\partial s_m} \sum_n \frac{\partial s_m(\mathbf{k}_0, \mathbf{k}')}{\partial k_{0_n}} \frac{\partial k_{0_n}[\mathbf{k};\mathbf{a}]}{\partial a_i}, \qquad [6.2.11]$$

where the second and third summations are performed over the three components of the position vector $\mathbf{k}$. As $\frac{\partial s_m}{\partial k_{0_n}}$ equals unity for $m = n$ and zero elsewhere, Eq. 6.2.11 can be simplified to:

$$\frac{\partial \hat{Y}[\mathbf{k};\mathbf{a}]}{\partial a_i} = \sum_{\substack{\mathbf{k}' \in \mathbf{K} \\ |\mathbf{s}| < M}} D[\mathbf{k}'] \sum_j \frac{\partial h(\mathbf{s})}{\partial s_j} \frac{\partial k_{0_j}[\mathbf{k};\mathbf{a}]}{\partial a_i}.$$ [6.2.12]

The first remaining partial derivative of Eq. 6.2.12 is solved using Eq. 6.2.6 and the equality $|\mathbf{s}|^2 = \mathbf{s}^T \mathbf{s} = \sum_i s_i^2$:

$$\begin{aligned}
\frac{\partial h(\mathbf{s})}{\partial s_j} &= \frac{\partial}{\partial s_j} \left\{ \gamma + \gamma \cos \frac{\pi}{M} \left[\sum_i s_i^2 \right]^{1/2} \right\} \\
&= \gamma \left\{ -\sin \frac{\pi |\mathbf{s}|}{M} \right\} \cdot \frac{\pi}{M} \frac{1}{2} \left[\sum_i s_i^2 \right]^{-1/2} \cdot 2 s_j \\
&= -\frac{\gamma \pi}{M} \frac{s_j}{|\mathbf{s}|} \left\{ \sin \frac{\pi |\mathbf{s}|}{M} \right\}.
\end{aligned}$$ [6.2.13]

This partial derivative is also zero outside the range of $|\mathbf{k}| < M$. The second remaining partial derivative of Eq. 6.2.12 can be found from Eqs. 6.2.5 and 6.2.6, and by noting that each transformation matrix depends on exactly one independent component of $\mathbf{a}$, such that:

$$\frac{\partial k_0[\mathbf{k};\mathbf{a}]}{\partial \mathbf{a}} = \left[\frac{\partial \mathbf{T}_\theta^{-1}}{\partial \theta} \mathbf{T}_\phi^{-1} \mathbf{T}_\psi^{-1} \mathbf{k} \quad \mathbf{T}_\theta^{-1} \frac{\partial \mathbf{T}_\phi^{-1}}{\partial \phi} \mathbf{T}_\psi^{-1} \mathbf{k} \quad \mathbf{T}_\theta^{-1} \mathbf{T}_\phi^{-1} \frac{\partial \mathbf{T}_\psi^{-1}}{\partial \psi} \mathbf{k} \right] \equiv \mathbf{U}.$$ [6.2.14]

In Eq. 6.2.14, $\dfrac{\partial k_{0_j}}{\partial a_i} = u_{ji}$, i.e. the element of $\mathbf{U}$ at row j and column i.

Substituting the results of Eqs. 6.2.13 and 6.2.14 back into the Eq. 6.2.12, the final equation for the partial derivative of the interim image with respect to the rotational parameters is derived:

$$\frac{\partial \hat{Y}[\mathbf{k};\mathbf{a}]}{\partial a_i} = \sum_{\substack{\mathbf{k}' \in \mathbf{K} \\ |\mathbf{s}| < M}} -D[\mathbf{k}'] \frac{\gamma \pi}{M} \frac{1}{|\mathbf{s}|} \left\{ \sin \frac{\pi |\mathbf{s}|}{M} \right\} \sum_j s_j u_{ji}.$$ [6.2.15]

Thus, with Eqs. 6.2.3 and 6.2.15, the gradient and Hessian matrix of Eqs. 6.2.7 and 6.2.8 may be computed as required by the minimization algorithm to yield the improved motion estimates at each iteration given by Eq. 6.2.10.

6.2.3 Estimating translation in three dimensions

As in the two-dimensional case, following the estimation and correction of the rotational displacements the partially corrected image will be related to the reference by a linear translation. The shift relation of Eq. 4.2.5 is readily extended to three dimensions, yielding:

$$\text{if} \quad d_2(x,y,z) \quad = \quad d_1(x - \Delta x, y - \Delta y, z - \Delta z),$$
$$\text{then} \quad D_2(k_x,k_y,k_z) \quad = \quad D_1(k_x,k_y,k_z) \cdot \exp\left[-i\left(k_x\Delta x + k_y\Delta y + k_z\Delta z\right)\right]. \qquad [6.2.16]$$

In Eq. 6.2.16, $d(x,y,z)$ and $D(k_x,k_y,k_z)$ are the continuous image and its three-dimensional Fourier transform, respectively, and Δx, Δy, and Δz are the linear displacements relating the two images. As in Chapter 4, the displacements need not, and typically will not, be integers. Paralleling the two-dimensional approach, computation of the cross-power spectrum of the two images causes the phase information associated with the images themselves to cancel, leaving only the phase associated with the shift. Extending Eq. 4.2.6 yields:

$$C_{12} = D_1 \cdot D_2^*$$
$$= D_1 \cdot D_1^* \cdot \exp\left[i \cdot \left(k_x\Delta x + k_y\Delta y + k_z\Delta z\right)\right]$$
$$= |D_1|^2 \cdot \exp\left[i \cdot \left(k_x\Delta x + k_y\Delta y + k_z\Delta z\right)\right]$$
$$\angle C_{12} = \Delta x \cdot k_x + \Delta y \cdot k_y + \Delta z \cdot k_z, \qquad [6.2.17]$$

where C_{12} is the cross-power spectrum, the asterisk indicates complex conjugation, the vertical bars indicate absolute magnitude, and $\angle C$ denotes the phase angle of C. As in the two dimensional case the phase angle of the cross-power spectrum in Eq. 6.2.17 defines a linear function in k_x, k_y, and k_z with coefficients Δx, Δy, and Δz equal to the displacements along each image axis. Thus, the coefficients in three dimensions are easily estimated by least-squares method, allowing the efficient estimation of sub-voxel translation. The rotationally aligned images can then be brought into full register by appropriate modulation in k-space and inverse Fourier transformation.

6.3 IMPLEMENTATION

All code to implement this algorithm was written in the standard C programming language. Portions of the code related to the minimization algorithm are presented in Appendix B. Other specific information regarding the implementation is presented below.

6.3.1 Estimation and correction of rotation

As described in Section 6.1, the rotational displacements θ, ϕ, and ψ between the "test" image to be registered and the reference image, are estimated by the minimization of the sum of squared differences in the spectral magnitudes between the test and reference images. This minimization is performed using the Levenberg-Marquardt algorithm, as implemented by Press et al. (1994), over a frequency range of interest. This range of interest, denoted I_{rot}, consists of frequency points lying between two hemispherical surfaces centered at the origin of k-space, and constitutes the set of points $\mathbf{K}$ in Eq. 6.2.1.. The default range has been empirically set between radii of 0.2 and 0.4 times the sampling frequency in the y direction, denoted $I_{rot} = (0.2, \ 0.4)$, although smaller intervals may be safely used with larger data sets. The reasons for selecting a subset of frequencies within the half-plane are the same as in two-dimensions, i.e. avoidance of relatively noisier higher frequencies and relatively uniform lower frequencies. Inclusion of only one half-plane is sufficient due to the symmetry of k-space. If necessary, scaling along the k_z direction compensates for an unequal field of view in the out-of-plane direction. In the present implementation, it is assumed that the field of view is equal in x and y, but additional scaling can be included if this is not the case.

The spectral magnitudes used to estimate rotation are computed by taking the magnitude of the three-dimensional fast Fourier transformation of the original image data output from the MR scanner. Iteration of the minimization algorithm is stopped after the changes in tested parameter values are less than 10^{-5} degrees or after forty iterations. During registration of a sequence of images, the final rotational parameters from the prior image, if available, are used as the starting estimates for rotation, otherwise null displacements are used. Values for the test image spectrum for a given θ, ϕ, and ψ are computed by interpolation of the original spectrum with a spherically symmetric Hamming window of Eq 6.2.2, which is rewritten as:

$$h(k_x, k_y, k_z) = \lambda + \lambda \cos\left[\frac{\pi}{M}\left(k_x^2 + k_y^2 + k_z^2\right)^{-1/2}\right] \text{ for } k_x^2 + k_y^2 + k_z^2 < M^2, \qquad [6.3.1]$$

and zero elsewhere, where λ is equivalent to γ in Eq. 6.2.2 corrected for out-of-plane scaling, if necessary, and M is the half-radius of the filter. Prior to minimization, the reference image spectra is pre-filtered with this same Hamming function to guarantee that all images have been identically treated prior to computing the error function and its derivatives. The default filter size has been empirically set at 2.5 times the step size along k_y, denoted $M_{rot} = 2.5$, although a larger radius may be required for images with unequal fields of view. As in the two-dimensional case, the symmetry of the k-space magnitudes prevents the

discrimination of rotation through an angle α from rotation through $\alpha + 180°$ about any of the axes. In practice, however, this does not represent an important limitation and the smaller rotation can be safely assumed.

After the estimation of the rotational displacement parameters, three-dimensional k-space regridding was used to generate the rotationally aligned test image. This method is a direct extension of the two-dimensional method described in Section 4.2.1. The Fourier transform of the original test image is computed after zero-padding to double the field of view in each direction and reduce subsequent aliasing artifacts. An aligned frequency space matrix is generated using the final rotational estimates by interpolation with a second spherically symmetric Hamming window. The default half-radius for the second window has also been set empirically at 2.5 times the step size along k_y of the doubled transform, denoted $M_{grid} = 2.5$, although a larger radius may be required for images with unequal fields of view. Following interpolation, the frequency image is inverse transformed and each pixel in the original field of view divided by the corresponding pixel in the inverse transform of the frequency interpolation filter. The image is cropped to the original field of view yielding a rotationally aligned test image.

6.3.2   Estimation and correction of translation

The final translation required to bring the rotationally aligned test image into full register with the reference image is estimated and corrected by the three-dimensional extension of the phase-plane method described in Section 6.2.3. Specifically, the cross-power spectrum between the transform of the rotationally aligned test image and the complex conjugate transform of the reference image is first computed by point-wise multiplication. The phase angle of the cross-power spectrum, computed with the arc tangent function, is used in a three-parameter least squares regression to determine the displacements coefficients Δx, Δy, and Δz fitting Eq. 6.2.17. As in the two-dimensional case, the regression is weighted by the magnitude of the cross-power spectrum at each point in k-space. Errors due to phase wrap are again avoided by incorporating a first estimate based on the first non-D.C. frequencies along each frequency axis, and limiting the regression to a frequency band of interest. The default interval was empirically set to lie between hemispheres of radii 0.1 and 0.3 times the sampling frequency in y, denoted $I_{trans} = (0.1, \ 0.3)$. Again, only a half plane is analyzed due to the symmetry of k-space.

After estimation of the parameters, the Fourier transform of the rotationally aligned test image is appropriately modulated to reverse the translational displacement in three dimensions, and the final fully registered image is constructed by inverse transformation.

6.4 VALIDATION: NUMERICAL METHODS

The accuracy of the method was quantitatively analyzed using numerical simulations in which idealized data sets with known displacements were constructed and registered using the methods described above. The sensitivity of parameter estimation to image noise was also assessed through numerical simulations.

6.4.1 Methods

Two 64 x 64 x 32 pixel simulated image sets were constructed by analytic computation of the frequency space of an ideally imaged object. The first set was used to assess algorithm accuracy; the second to assess algorithm stability in the presence of noise. In both cases the field of view in the third (z) dimension was assumed to be 50% of the x and y fields of view, simulating cubic voxels.

In the first case, a set of 64 images was generated of a simulated 38 x 30 x 24 pixel ideal homogeneous rectangular prism with an inset 34 x 26 x 20 pixel prism of 90% relative intensity of the larger prism. Images were constructed over a range of rotational angles and translational displacements. A slice through the center of the reference image in the x-y plane is shown in Figure 6.4.1A. To generate this reference image from a centered position aligned with the spatial axes, the prism was translated by -0.7 pixels along the x axis, -1.1 pixels along the y axis, and +1.2 pixels along the z axis, and then rotated sequentially through -0.9° about the x axis, -1.2° about the y axis, and +1.5° about the z axis. These values were arbitrarily selected to create a reference image which was slightly off axis. The reference k-space samples corresponding to this positioning were then synthesized analytically.
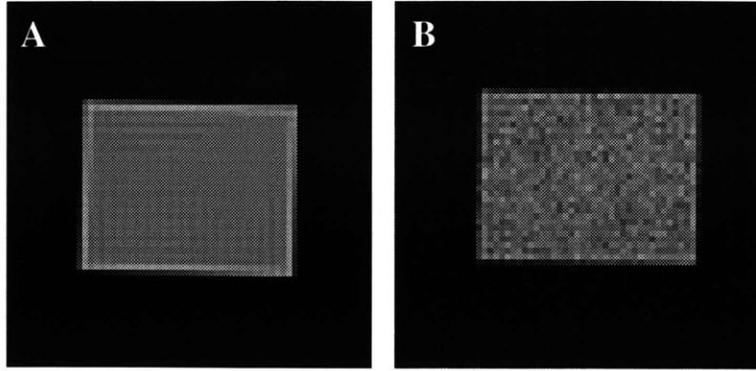
Figure 6.4.1. Simulated objects used to assess algorithm accuracy and stability in the presence of noise. (A) A slice through the z=0 plane of the noiseless reference image of a 38 x 30 x 24 pixel prism with an inset 34 x 26 x 20 pixel prism of 90% relative intensity. (B) A slice through the z=0 plane of the 38 x 30 x 24 pixel prism with simulated signal-to-noise of 20:1. The orientations and positions of these objects are described in the text.

The remainder of the images in the first data set were generated by assuming that the reference image was further translated by one of 0, 0.3, 0.6, or 0.9 pixels along each axis, and then sequentially rotated by one of 0, 0.3, 0.6, or 0.9 degrees about each of the x, y, and z axes. Again, the k-space matrices for each image were generated analytically, i.e. no regridding or other manipulations were performed. The angular and translational displacements about or along each axis were equally distributed such that all images were subjected to both translation and rotation, and all pairs of images had differences in at least one rotational and one translational parameter. A set of inverse Fourier transformed magnitude images was generated from the simulate k-space matrices and registered as described above. The mean and standard deviation of the estimate error were computed for each value of each estimated parameter.

In the second case, the inset prism was omitted from the simulated object. The reference image of this data set was generated assuming translation of -0.9 pixels along each axis and sequential rotation through -9° about the x, y, and z axes, beginning from the centered position as above. Prior to inverse Fourier transformation, white noise was added to the simulated sampling of k-space. Four groups of 16 images were generated from the reference k-space matrix, corresponding to signal-to-noise ratios of 50:1, 20:1, 10:1, and 5:1. An example image from the 20:1 signal-to-noise group is shown in Figure 6.4.1B. The images within each group were registered to the first noisy image of that group, and the mean and standard deviation of the parameter estimates were computed. As there was no actual

difference in position between the underlying prisms, the computed values provide a measure of performance degradation in the presence of noise.

In both cases, the registration parameters were set to $I_{rot} = (0.1,\ 0.2)$, $M_{rot} = 5.0$, $M_{grid} = 2.2$, and $I_{trans} = (0.1,\ 0.3)$. The larger initial filter size was required due to the sparser sampling in frequency space along the third dimension corresponding to the 50% reduction in field of view. To assess further improvement in accuracy with larger window sizes, the analysis was also repeated with the initial window size set to $M_{rot} = 10.0$.

6.4.2 Results

Processing time for both simulated data sets on a Digital AlphaStation 500/333 averaged 50.1 seconds per frame, of which 2.22 seconds (4%) was used to estimate rotation, 30.7 seconds (60%) for regridding, 17.2 seconds (33%) to estimate and correct translation, and the remainder spent accessing data from a network server. Of total processing time, 55% was spent executing calls to fast Fourier transform routines.

The errors in estimation of rotational angles and translational displacements are listed in Table 6.4.1. Taking all values together, the mean and standard deviation of rotation and translation were -0.026 ± 0.019 degrees and +0.001 ± 0.043 pixels, respectively. There is a small negative bias in the estimated rotation angles which increases with the displacement, likely related to the frame-related ringing noise described in Chapters 4 and 5, which may benefit from reference image pre-rotation as was described in the two-dimensional case. Translations appear largely unbiased, and the deviation of the errors does not appear to increase with increasing displacements for either rotations or translations.

When the analysis was repeated with initial filter size set to $M_{rot} = 10.0$, the mean and standard deviation of the estimation errors of rotation and translation were -0.010 ± 0.013 degrees and 0.001 ± 0.043 pixels, respectively. Increasing the window size provided only a small improvement in the rotational accuracy of the algorithm, but significantly increased the processing time to estimate rotation from 2.22 seconds to 12.9 seconds per frame.

Errors introduced by registration in the presence of noise are summarized in Table 6.4.2. Not surprisingly the standard deviations of the estimate errors increase with increasing noise levels. However, even in the presence of a signal-to-noise ratio of 5:1, the additional error remained remarkably small, demonstrating the reliability of the algorithm in the presence of noise.

Table 6.4.1. Accuracy of estimated registration parameters from simulated images for concurrent small rotations and translations. As the results were similar for all three angles estimated, the errors are pooled in the table. As this was also true for translations, these are also pooled. s.d. = standard deviation.

	Rotation (degrees)				
	0.0°	0.3°	0.6°	0.9°	Combined
mean error	-0.009	-0.013	-0.028	-0.053	-0.026
s.d. error	0.019	0.017	0.016	0.022	0.019

	Translation (pixels)				
	0.0	0.3	0.6	0.9	Combined
mean error	-0.000	0.001	0.003	0.000	0.001
s.d. error	0.037	0.040	0.048	0.049	0.043

Table 6.4.2. Degradation of estimated registration parameters from simulated images in the presence of noise. As in Table 6.4.1, the results are pooled for all three rotation estimates and for all three translation estimates. The results are presented as mean ± standard deviation.

	Signal-To-Noise Ratio			
	50:1	20:1	10:1	5:1
Rotation error (degrees)	+0.002±0.007	-0.000±0.013	+0.021±0.025	+0.002±0.058
Translation error (pixels)	-0.000±0.002	+0.001±0.003	+0.001±0.005	-0.002±0.007

6.5 VALIDATION: HUMAN SUBJECTS

To further assess the utility of the registration algorithm, data were collected from human subjects. While the actual displacements of the subjects brain cannot be known under these conditions, comparison of images before and after registration can be used to qualitatively evaluate the accuracy of the method under experimental conditions which might be encountered in practice, and difference images between pairs of volumes before and after registration used to quantify reductions in image misregistration. Addtionally, such an analysis provides some insight into the types of artifacts which might be expected to limit the accuracy of three-dimensional registration.

6.5.1 Methods

A set of 25 echo planar images were collected from a 20 year old right handed male volunteer with a 1.5 T General Electric Signa scanner retrofit with whole-body resonating gradients (gradient echo, TR = 6 s, TE = 40 ms, 32 slices, 5 mm slice thickness, field of view 20 x 20 x 16 cm, quadrature head coil). An interleaved acquisition was used to prevent cross-slice saturation effects caused by imperfect slice select gradients, allowing contiguous volumes to be acquired without inter-slice spacing. Thus, for each imaging volume, all odd slices were acquired first, then all even slices. During the acquisition, the subject was asked to move his head to a different position every 30 seconds, and to hold that position until the next movement. The registration algorithm was applied using the default filter widths and intervals specified in Section 6.3. Image registration was visually assessed for the large intentional subject movements by direct image comparison before and after registration. For unintentional movements, difference images were examined between volumes acquired sequentially in the intervals between subject movements.

6.5.2 Results

The total time to register the full set of 25 image volumes was 33 minutes and 11 seconds (33'11"). Of this, 12'27" (37%) was spent performing Fourier transformations. Excluding 18.63" of total initialization time, the average time spent registering each image volume was 1'18.9", of which 12.45" (15%) was spent iteratively estimating rotation, 47.53" (60%) was spent correcting rotation with regridding, 18.34" (23%) was spent estimating and correcting translation, and the remainder was spent on disk read-write and other overhead operations. As in the numerical simulations, the bulk of the processing time is dedicated to the high-fidelity resampling operation, while the iterative estimation steps remain relatively fast.

86
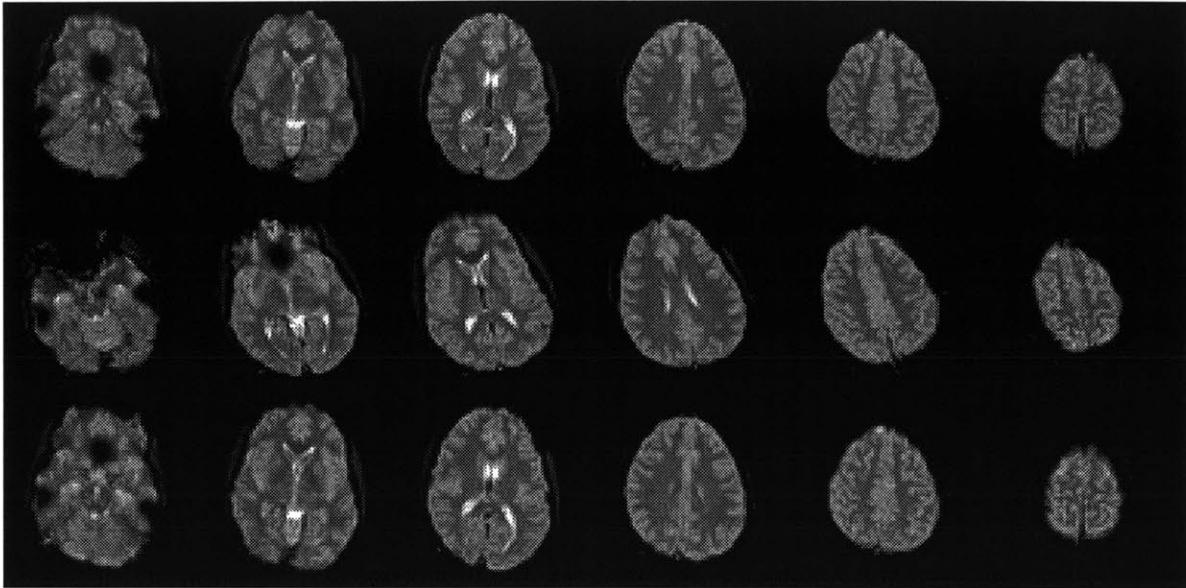
Figure 6.5.1. Three dimensional registration of echo planar imaging volumes. Top row: Selected axial slices at 1.5 cm intervals through the reference volume. Middle row: The same slices after the subject was instructed to shift in the scanner. Bottom row: The volume from the second row following three-dimensional registration . The algorithm detected and corrected rotations of -14.6° about x, +3.25° about y, and -19.1° about z and translations of +1.98 pixels along x, +0.136 pixels along y, and -0.780 pixels along z.

A typical set of images from the volume before and after registration is presented in Figure 6.5.1. Every third slice through brain is shown. As clearly evident, the reference volume (top row) and volume to be registered (middle row) were collected with the subject's head in different positions. The algorithm detected and corrected rotations of -14.6° about x, +3.25° about y, and -19.1° about z, and translations of +1.98 pixels along x, +0.136 pixels along y, and -0.780 pixels along z, yielding the images in the bottom row of the figure. These results indicate algorithm stability in the presence of large rotations. Importantly, these rotations are orders of magnitude larger than would be expected in a typical functional MRI experiment.

The effects of image registration on small unintentional movements are illustrated in Figure 6.5.2. In this figure, difference images between the reference volume and the unregistered volume are compared with difference images between the reference volume and

the registered volume. Every slice through significant amounts of cortex is shown. The reference and test volume were collected consecutively 6 seconds apart during a segment of the experiment where the subject was instructed to remain motionless. Nonetheless, the algorithm detected and corrected rotations of -0.489° about x, -0.016° about y, -0.012° about z, and translations of +0.031 pixels along x, -0.029 pixels along y, and -0.036 pixels along z. This pattern of displacements is consistent with a nodding motion by the subject, i.e. sagittal rotation, which continues to be one of the most common patterns of rotation observed in functional experiments (Lee 1998).

Figure 6.5.2. (following page) Reduction in artifacts following three dimensional registration for small rotation angle. The figure is arranged in three sets of three rows. The first row of each set includes selected axial images from the reference volume at 5 millimeter increments from inferior (towards the neck) to superior (towards the crown). The second row of each set shows difference images between a second volume acquired six seconds after the reference volume. The images are scaled such that black pixels represent a 5% or greater decrease in signal image (as a percentage of the maximum intensity in the volume) and white pixels represent a 5% or greater increase. Clear artifacts are visible throughout, with pronounced artifacts in the most superior images. These are uniformly reduced in the third row of each set, demonstrating the second volume following image registration. The algorithm detected and corrected rotations of -0.489° about x, -0.016° about y, -0.012° about z, and translations of +0.031 pixels along x, -0.029 pixels along y, and -0.036 pixels along z. This pattern of displacements is consistent with a nodding motion by the subject.
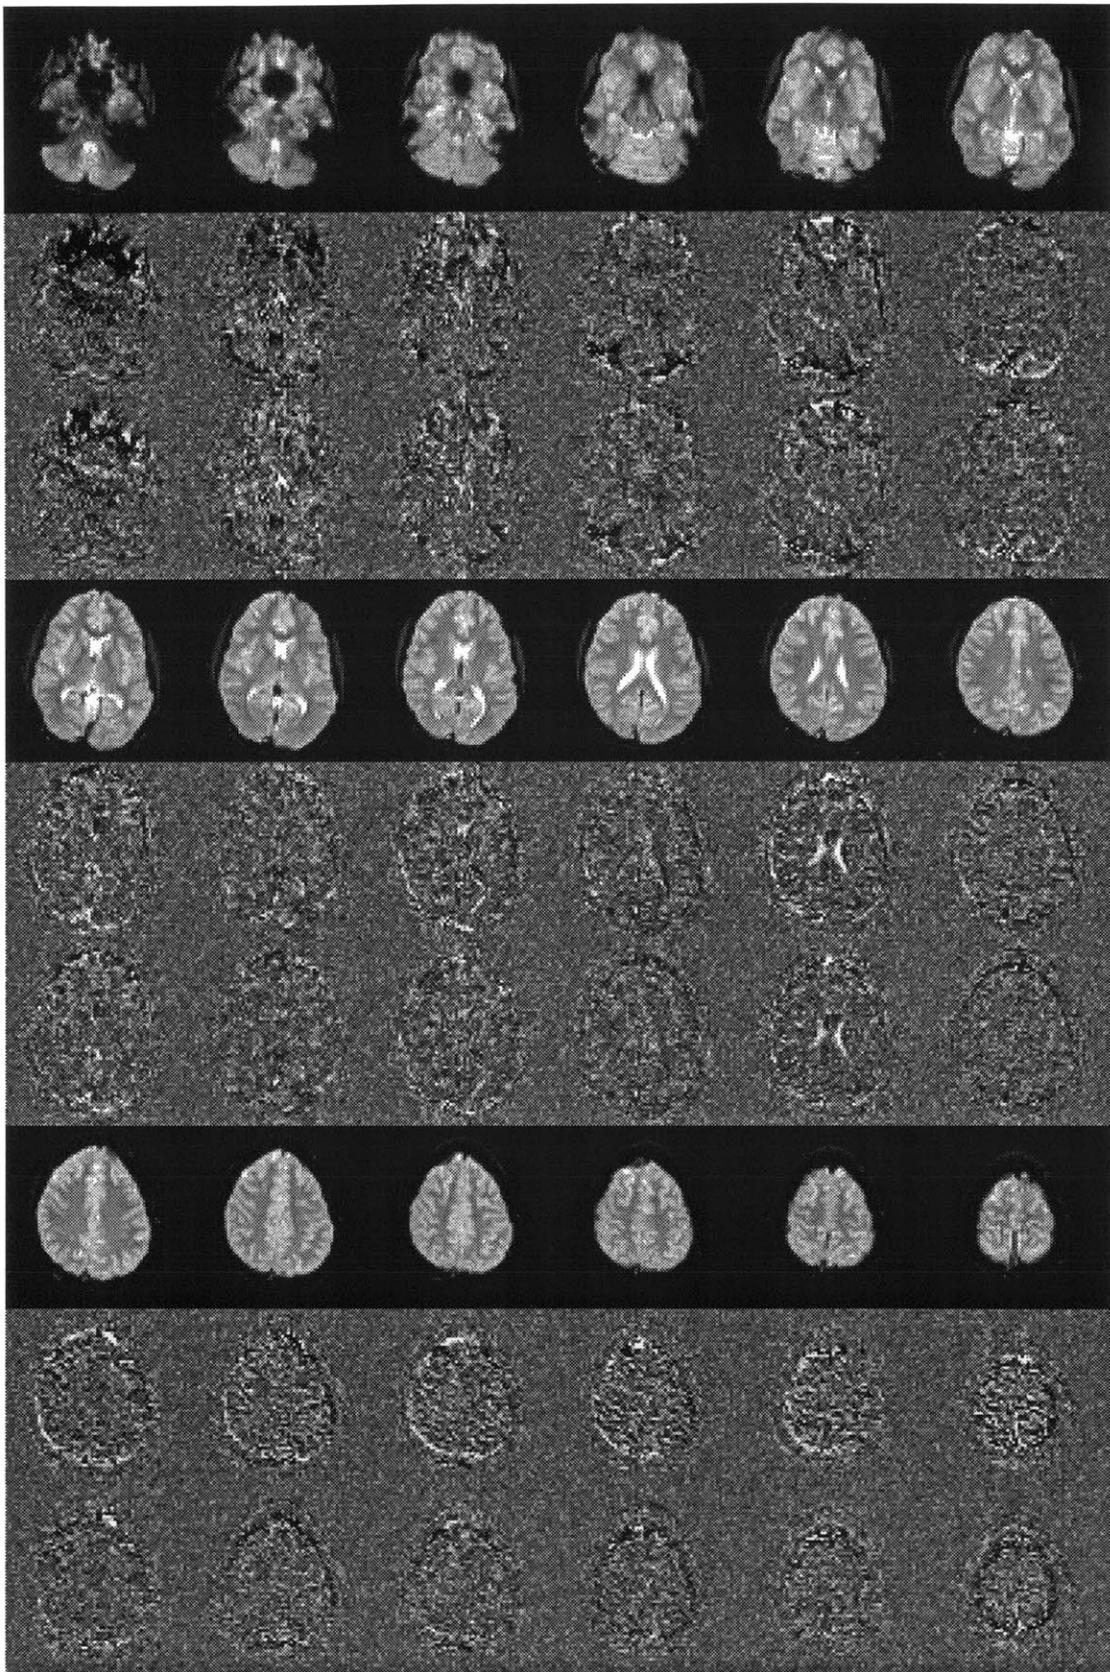
Figure 6.5.2. The caption for this figure appears on the preceding page.

6.6 DISCUSSION

A three dimensional registration algorithm has been presented which extends the decoupled methods of Chapter 4. The new algorithm is partially iterative, using the Levenberg-Marquardt algorithm to estimate rotational displacement by minimizing the squared difference error between the magnitude of the spectra of the images to be aligned. Once rotation has been estimated, a three-dimensional frequency regridding algorithm is used to generate rotationally aligned volumes. Residual translation is efficiently estimated with cross-spectrum techniques and corrected by frequency modulation. The algorithm is fully automated and rapidly processes image sets, allowing for the use of a computationally intensive high-fidelity resampling method.

The analysis of simulated data demonstrated the accuracy of the algorithm and robustness in the presence of significant noise. The values presented in Tables 6.4.1 and 6.4.2 indicate that the method provides excellent sub-voxel registration over the range of rotations and translations which might be expected in functional MRI experiments. These data, however, represent an upper bound on the accuracy which might be expected in practice, as they do not reflect the additional inaccuracies which might be introduced by other experimental factors not related to the algorithm. For instance, scanner gradient non-linearities, field inhomogeneities, amplifier instabilities, and other technical limitations would be expected to lessen performance. While these factors might be reflected in data collected from an appropriate three-dimensional, six degree of freedom imaging phantom, the construction of such a phantom is not trivial.

The utility of the algorithm, however, can still be assessed by examination of data collected from human subjects, where the true displacements are unknown. The data presented in Figure 6.5.1 demonstrate visually the registration of images with large displacement parameters. More important, however, are the findings in Figure 6.5.2, where two image volumes collected sequentially only six seconds apart were analyzed before and after registration.

Two things are particularly evident in Figure 6.5.2. First, the very small amount of motion by the subject over this short time period has resulted in significant edge artifacts throughout the brain. These bright and dark rings appear at the boundaries of the cortex, the ventricles, and even at the boundaries of white and gray matter, and are comparable with those seen in the two-dimensional imaging phantom experiments previously shown in Figures 4.2.3 and 4.2.4. The artifacts are particularly severe in the more superior slices, consistent with the

90

mechanics of nodding. Importantly, the artifacts are significantly reduced by registration, illustrating the potential practical value of the algorithm.

Second, the artifacts in the unregistered images of Figure 6.5.2 are not equal in the even and odd images of the acquisition, i.e. they are more severe in the first, third, and fifth images of each row, especially in the bottom two rows. Recalling that this was an interleaved acquisition, in which all the odd slices were acquired before the even slices, this finding is a good example of the intra-volume motion artifacts described in Chapter 2. Such a pattern would be consistent with unintentional subject movement beginning near the middle of the six second interval required to collect a single volume or other related patterns of movement. As previously stated, this type of intra-volume movement complicates the estimation of displacement parameters, where equal motion affecting every slice is assumed. As a result, the estimated parameters likely reflect an averaging of the motion across all slices, and thus, none of the slices is perfectly corrected. This is suggested in Figure 6.5.2, where the positive intensity artifacts in the odd slices are largely, but not incompletely reduced, and the smaller positive artifacts in the even slices are replaced by small negative intensity artifacts. Unfortunately, given the present limitations of three dimensional acquisition techniques, such artifacts are unavoidable. However, as techniques are developed to allow faster acquistion, such artifacts can be expected to be reduced in importance.

In summary, a registration algorithm has been presented which extends the decoupling principle previously applied in two dimensions to three dimensions. The algorithm is rapid, accurate, and robust in the presence of noise and large rotations, and provides a high-fidelity resampling method.

7 CLINICAL RESEARCH APPLICATION

The methods described in this thesis were applied in the analysis of data from an ongoing study of the activation of brain regions during cue-induced cocaine craving in normal controls and in individuals addicted to crack cocaine (Maas 1998). The nature of the study made motion compensation and noise reduction especially important, as will be described below. In fact, the differences which will be described were only detectable following appropriate processing.

7.1 STUDY BACKGROUND

The craving response is an important measurable component of drug abuse whose neurobiological basis remains an enigma (Koszlowski 1987). Recently, functional neuroimaging methods have been used to identify brain structures that may mediate cue-induced craving. Using fluorodeoxyglucose positron emission tomography (FDG PET), Grant et al. (Grant 1996) demonstrated metabolic increases in a number of brain regions, including the dorsolateral prefrontal cortex and medial temporal lobes, in cocaine users after exposure to cocaine-related cues. In a preliminary study, Childress et al. (Childress 1996) used PET with oxygen-15 labeled water to demonstrate significant increases in blood flow in the amygdala, anterior cingulate, and temporal poles in the brains of cocaine users after they had viewed videotapes containing cocaine-related cues.

As described in Chapter 2, echo planar functional magnetic resonance imaging (FMRI) based upon blood oxygenation level dependent (BOLD) contrast is an alternative method for the measurement of brain activation in response to stimuli, providing high spatial and temporal resolution without radioactive tracers. Unlike the kinetic parameters estimated with PET, the BOLD signal, as described in Chapter 1, is related to changes in local blood flow and deoxyhemoglobin levels subsequent to neuronal activation. This study was performed to

evaluate whether cocaine cue-induced regional cerebral activation previously detected with PET is also detectable with BOLD FMRI.

7.2 METHODS

Six male subjects (mean age 37 ± 7 years) with a history of crack cocaine use at least once every two weeks during the past six months and six comparison male subjects (mean age 32 ± 4 years) were studied. Subjects in the cocaine group reported that crack cocaine was their preferred drug of abuse, but a history of other drug use was not grounds for exclusion. Comparison subjects reported no history of cocaine use or drug abuse of any kind. All subjects were screened using the SCID-I/P structured interview (First 1995). Individuals with a history of psychotic illness or a current episode of an Axis I mood disorder were excluded. All subjects were in good general health and had no known history of neurological disorder. The protocol was approved by the McLean Hospital Institutional Review Board. All subjects provided written informed consent.

Using segments of the videotape described by Childress and colleagues (Childress 1996), a ten-minute videotape consisting of four contiguous 150-second segments of alternating neutral and cocaine-related scenes and sounds was constructed. This tape was presented to each subject by means of a non-ferrous audio-visual system during a single FMRI scanning session (1.5 Tesla, 7 mm thickness, 3 mm skip, gradient echo, TE = 40 ms, TR = 5 s, flip angle = 90°, in-plane resolution = 3.125 x 3.125 mm), wherein 120 echo-planar images were collected at 5 second intervals in each of 16 coronal slices. To quantify each subject's level of craving, a standardized visual analog scale questionnaire on the subject's strength and frequency of desire for cocaine (Weiss 1995) was administered twice, immediately before and after scanning.

To reduce subject motion artifacts, all FMRI data were spatially registered prior to analysis using the two-dimensional registration algorithm described in Section 4.2 with the extensions described in Section 4.3. As will be discussed below, three-dimensional registration was not used because of the large distance between slice centers required to provide adequate temporal resolution for this study. Following registration, the data were processed with the spatial and temporal filtering techniques described in Chapter 5.

Data were analyzed with a custom software package written in the Interactive Data Language (IDL, Research Systems Inc.). This package allows the definition of regions of interest within each image set and automated statistical analysis. Ten neuroanatomic regions (Table 1) were selected based upon the results of the prior PET studies. The regions were

identified from high resolution MR images and mapped to the functional image set by two raters (the author and Marc Kaufman, Ph.D.).

In each region, two scalar measures were used to assess BOLD activation: (a) the magnitude of activation was defined as the percent change in the mean regional signal intensities measured between the neutral and cocaine-cue segments of the acquisition, and (b) the spatial extent of activation was defined as the fraction of pixels in each region that exceeded an arbitrary activation threshold of $r = 0.3$ using the correlation coefficient detection technique of Bandettini et al. (1993) with a box-car reference function incorporating a presumed six second hemodynamic delay. Reliability of the region selection technique was verified by computation of the intraclass correlation from four subjects for both activation magnitude ($r1 = 0.87$) and spatial extent ($r1 = 0.84$).

7.3 RESULTS

Activation data are summarized in Table 7.3.1. Significant differences in activation magnitude between cocaine and comparison groups were detected in the anterior cingulate and left dorsolateral prefrontal cortex, but not in the other neuroanatomic regions of interest. A representative example showing a single coronal slice in one of the subjects with the anterior cingulate and left and right prefrontal dorsolateral cortical regions of interest outlined is shown in Figure 7.3.1. The average time course signals from the cocaine group in the two significantly activated regions are demonstrated in Figure 7.3.2. For the cocaine group, the regional signal intensity time courses were comparable for the first and second presentation of each type of stimulus, indicating that there were minimal activation carry-over effects.

Although significant group differences in spatial extent were not found, spatial extent, but not activation magnitude, correlated significantly with the change in self-reported desire for cocaine in the cocaine group (anterior cingulate: $r = 0.88$, $t = 3.67$, $p = 0.022$; left dorsolateral prefrontal cortex: $r = 0.83$, $t = 2.97$, $p = 0.041$; linear regression, 4 df.). As the craving questionnaire was administered only twice, segment-to-segment and carry-over effects on craving were not assessed in this study. Significant pre- to post-scan craving score changes, however, were observed only in the cocaine group (mean change = 2.83, sd = 3.31, $t = 2.96$, $p < 0.025$ by two-tailed paired t test, 11 df.).

94

Table 7.3.1. Mean ± standard deviation signal intensity increase in each region as a percentage of baseline intensity. All p-values were computed by two-tailed Student's t-test with 10 degrees of freedom. ant. cing. = anterior cingulate. * p < 0.05.

| | ant. cing. | Dorsolateral Prefrontal Cortex | | Temporal Poles | |
		right	left	right	left
Cocaine group	0.29 ± 0.15	0.15 ± 0.14	0.25 ± 0.10	0.07 ± 0.13	0.02 ± 0.23
Control group	0.00 ± 0.12	0.06 ± 0.10	0.05 ± 0.16	0.09 ± 0.15	0.01 ± 0.12
t	3.70	1.32	2.59	–0.20	0.14
p	0.0041 *	0.22	0.027 *	0.85	0.89

| | Basal Ganglia | | Parietal Lobes | | |
	right	left	right	left	Cerebellum
Cocaine group	–0.01 ± 0.10	0.03 ± 0.09	0.11 ± 0.20	0.03 ± 0.13	0.09 ± 0.14
Control group	0.10 ± 0.17	0.09 ± 0.14	0.06 ± 0.12	0.05 ± 0.15	0.04 ± 0.16
t	–1.42	–0.92	0.52	–0.16	0.58
p	0.19	0.38	0.61	0.87	0.57
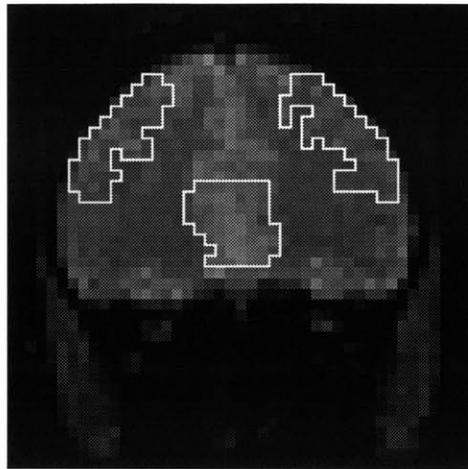
FIgure 7.3.1. Representative image from data set showing three regions of interest. As this image is presented according to radiological convention, the upper left region is the right dorsolateral prefrontal cortex, and the upper right region is the left dorsolateral prefrontal cortex. The central region is the anterior cingulate. The anterior cingulate and left dorsolateral prefrontal cortex were significantly activated in this study.
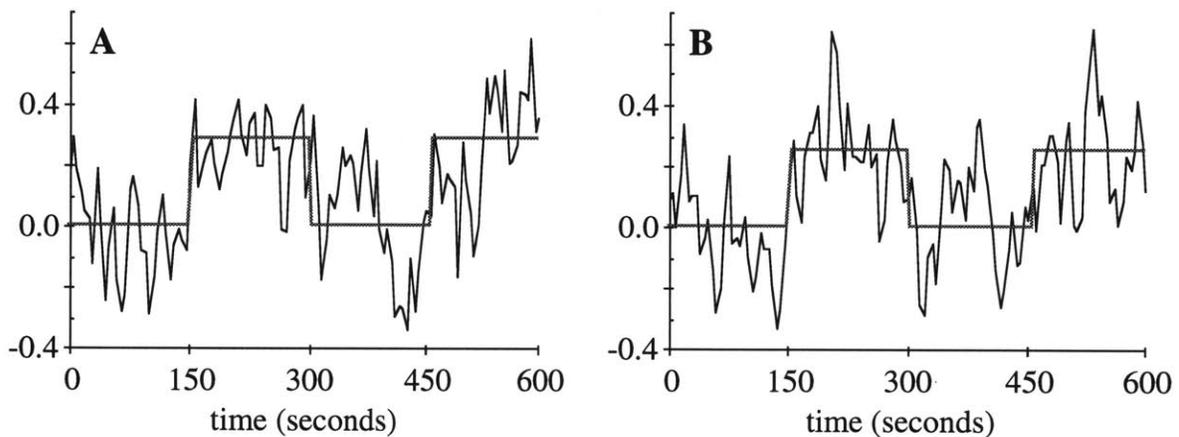
Figure 7.3.2. Activation of 0.29% was observed in the anterior cingulate gyrus (A) and 0.25% in the left dorsolateral prefrontal cortex (B) of subjects in the cocaine group. These activations were significantly different from the comparison subjects. The black line in each figure indicates the average time course for all subjects in the cocaine group. The grey line indicates the model box-car function scaled to the observed activation. The cocaine cues were presented between 150 and 300 seconds and between 450 and 600 seconds. Data from comparison subjects are not shown.

7.4 DISCUSSION

Regional cerebral activation in response to cocaine-related cues was evaluated using echo planar FMRI. Significant activation was detected in the anterior cingulate and left dorsolateral prefrontal cortex of crack cocaine smokers. These findings are consistent with the metabolic (Grant 1996) and blood flow (Childress 1996) changes found in prior PET studies. As in the FDG PET study (Grant 1996), change in the self-reported level of cocaine craving was found to correlate with a measure of activation in these regions. Susceptibility artifacts near air-tissue interfaces limited consistent imaging of the medial aspect of the temporal lobe in the present investigation, and may also have contributed to the negative findings in the temporal poles.

There are a number of procedural differences which may also contribute to the slightly different activation profile observed in the present study. First, in the PET studies, subjects were exposed to cocaine-related cues for 30 minutes. In contrast, the subjects of this study were exposed to the stimuli for only 150 seconds at a time. Second, prior studies have only imaged one condition (i.e., drug-related or neutral-related) per session, allowing subjects at least one hour (Childress 1996) or one week (Grant 1996) between conditions. In contrast, this study sought to determine if activation could be observed within a single 10 minute session by alternating relatively brief segments of neutral and cocaine-related cues. The results suggest that FMRI can detect activation during this short period without substantial carry-over effects. Third, prior study designs have also included tactile stimuli or the anticipation of cocaine use to increase the likelihood of eliciting craving (Grant 1996). In the present study, significant activation and craving were detected using only visual and auditory cues. Collectively, these procedural differences may be viewed as a strength of this study in that significant patterns of activation were observed during brief periods of stimulation encompassing only two sensory modalities.

Several factors in this study highlight the importance of the image registration and post-registration correction techniques used in data analysis. First, the activation detected is relatively small, on the order of 0.25%, and would be easily obscured by motion and registration related noise effects. Additionally, the individual scanning sessions are relatively long for functional magnetic resonance imaging, lasting ten minutes, increasing the likelihood that subject motion artifacts would be present. Furthermore, the subjects were exposed to stimuli which, in many cases, elicited a strong drug-craving reaction. These reactions can be expected to impair the ability of the subjects to remain motionless. As the raw, uncorrected

data were largely uninterpretable, the application of the methods of this thesis were critical in the analysis of this data.

Scanner hardware limitations permitted the acquisition of only sixteen slices within the desired five second repetition interval. Thus, in order to cover the entire brain, including all neuroanatomic regions of interest, a relatively wide distance between slice centers was necessary, limiting out-of-plane resolution. As discussed in Chapter 2, the lower resolution in the out-of-plane direction made registration in three dimensions impractical, and two-dimensional registration was used to prevent image degradation.

Using FMRI techniques and the image registration and post-registration correction methods described in this thesis, this clinical research study demonstrated that the presentation of cocaine cues activated anterior cingulate and dorsolateral prefrontal cortex, consistent with previous PET results. These results suggest that functional MR imaging methods may be a valuable tool in the identification of the neurobiological basis of craving and in the evaluation of new agents to modify or reduce craving (Robbins 1992, O'Brien 1992). Additionally, they demonstrate the utility of proper processing strategies to compensate for subject motion in clinical applications.

8 CONCLUSIONS

8.1 SUMMARY

In this thesis, several processing methods for functional magnetic resonance imaging were presented. These methods were designed to address sources of noise related to subject motion. The key reasons to compensate for subject motion effects were introduced in Chapter 2, and the limitations of currently available techniques were examined. With these limitations in mind, new techniques were developed and explored.

In Chapter 4, motion artifacts in two dimensions were considered. With the goal of attaining a noniterative solution to the two-dimensional rigid body registration problem suitable for use with FMRI data sets, a Fourier domain decoupling principle was described which allowed separate solutions for the rotational and translational components of misregistration. Based upon this principle, a new registration algorithm was presented for rapid, automated in-plane registration of large high-resolution functional MRI time series data sets. Efficient independent parameter estimation using cross-power spectrum techniques provided accurate estimation of registration parameters. Additionally, corrected images were generated using a high-fidelity alternative to linear interpolation to reduce image degradation. Results of numerical simulations and testing with data collected from imaging phantoms and human subjects demonstrated the improvement provided by this registration algorithm when applied in functional MRI applications. Further refinements allowing improved estimation of rotational displacements were also described.

In Chapter 5, the problem of registration-related artifacts in post-registration data sets was explored. These artifacts were demonstrated to occur even with ideal registration of simulated images. The sources of these artifacts were discussed and two approaches to reduce added noise were presented. First, a spatial filtering method which selectively attenuated regions near the corners of the k-space matrix *after* registration was described. Testing with simulated data and experimental data from human subjects revealed a significant reduction in the noise introduced by registration following application of the filter, approaching 50% in

methods based upon frequency regridding techniques. A complementary temporal filtering approach was also described based upon an additive model relating the estimated registration parameters and the noise introduced into the pixel time series. The noise was found to be well approximated by polynomial functions of measured rotations, and thus well suited to reduction with multiple regression techniques.

In Chapter 6, a three-dimensional extension of the decoupled registration algorithm was presented. A partially iterative approach was used to estimate rotation, while a one-pass method for translation was retained. The algorithm was demonstrated to provide accurate estimates of rotational and translational parameters for simulated data sets and robustness in the presence of noise. Application to human data provided a reduction in motion-related artifacts. A number of the factors complicating three-dimensional registration were discussed, and examples of three-dimensional intra-volume artifacts presented.

In Chapter 7, the methods of this thesis were applied in a clinical FMRI study of cue-induced cocaine craving. Factors in this study which made the paradigm particularly sensitive to motion effects were discussed. Following appropriate compensation for motion in the data sets, functional activation was detected in anterior cingulate and dorsolateral prefrontal cortex, consistent with previous PET results. These positive results, which had not been obtained without motion compensation, demonstrated the utility of the described methods in real applications.

8.2 DIRECTIONS FOR FUTURE WORK

The quality of image registration can be improved further by incorporating additional information relating to motion effects into registration algorithms. As respiration and circulation contribute to subject motion, real-time measurements of physiologic data related to these processes may permit direct estimation of their effects on motion. This will be facilitated by continued improvments in MR scanner-compatible physiologic monitoring devices. Other processes which are difficult to avoid, such as swallowing and eye, tongue, and jaw motion complicate the estimation of motion parameters. While masking out these pixels in images may aid in registration, masking introduces its own biases. Thus a careful analysis of the effects of each, and, where possible, real-time monitoring, may ultimately lead to more accurate image registration.

The continued development of real-time methods to compensate for subject motion prior to acquisition is particularly exciting. This has recently been demonstrated by Lee et al. (1998) for five of the six degrees of freedom in three dimensions, and extensions to a full three dimensions seem possible. Specifically, three-dimensional navigator echoes to uniformly

sample a sphere in k-space during volume acquisitions may ultimately provide sufficient information to accurately adjust scanner parameters in three dimensions without the saturation effects limiting current methods. Importantly, the rapid techniques described in this thesis for parameter estimation would be well suited for incorporation into such a system, and direct application to the k-space data sampled by the scanner would further increase speed by reducing the number of Fourier transforms required.

While the importance of motion in three dimensions is clear, a number of factors complicating accurate three-dimensional registration have been demonstrated and discussed in this thesis. As the utility of three-dimensional algorithms will continue to be limited until these technical issues, particularly the suppression of intra-volume motion effects, are resolved, this represents an important area for further study. Fortunately, the continued development of faster three-dimensional volume acquisition methods, such as echo volumar imaging, may ultimately provide a solution to this problem.

8.3 CLOSING COMMENTS

The recent development of functional magnetic resonance imaging has provided a dramatic new opportunity to study human brain function. Everyday, FMRI techniques are being used successfully in an increasing number of ways to gain insight into the processes that make us human. As the technical issues related to image acquisition, paradigm design, image processing, and data analysis continue to be studied and solutions refined, the quality of data provided by FMRI studies can only continue to improve. It is the genuine hope of this author that the work presented in this thesis will contribute to this process.

APPENDICES

A SOFTWARE CODE - SPATIAL FILTER REALIZATION

This section lists software code related to the realization of the post-registration spatial filter described in Section 5.2. The core code is dependent upon the routine klh(), adapted from Press (1994) which generates the impulse response of a one-dimensional Kaiser-Bessel lowpass filter.

```
/* K-B low/high pass filter, adapted from Numerical Recipes.

    lh    = 1 for low-pass, -1 for high-pass
    fs    = sampling frequency
    fpass = pass-band cut-off frequency
    fstop = stop-band cut-off frequency
    Apass = maximum amplitude variation allowed in pass-band (dB)
    Astop = minimum attenuation required in stop-band (dB)
    N     = filter size required
*/

double * klh (short lh,double fs,double fpass,double fstop,
          double Apass, double Astop, int *N);


/* Create spatial post-registration filter.  This procedure returns
   the two-dimensional frequency response of the 2-D frequency
   transformation of the 1-D Kaiser-Bessel low-pass filter
   (KBLPF) specified by the filt_xx parameters.  The McClellan
   transformation kernel is used in Clenshaw's recurrence.

   sizeX,sizeY: maximimum spatial dimensions
   filt_ws: stop-band cutoff frequency divided by pi (0 to 1)
   filt_wp: pass-band cutoff frequency divided by pi (0 to 1)
   filt_pv: maximum amplitude variation allowed in pass-band (dB)
   filt_sa: minimim attenuation required in stop-band (dB)
*/

DoubleImage GenPostRegFilt (short sizeX,short sizeY,
      double filt_ws, double filt_wp, filt_pv, filt_sa);
{

   double     *a;              /* pointer to 1-D KBLPF */
   int        N;               /* size of 1-D KBLPF */
```

```
  DoubleImage   Hwxwy;              /* pointer to frequency response record*/
  int           loopwx,loopwy,lc;   /* loop variables*/
  double        *data;              /* pointer to frequency response data */
  double        *y;                 /* array for recurrence algorithm */
  double        cwx,cwy,T;          /* variables for recurrence algorithm */

  /* generate 1-D low-pass filter */
  a=klh(1,2,filt_wp,filt_ws,filt_pv,filt_sa,&N);

  /* check required 2-D filter size bounds */
  if (N*2+1 >= sizeX || N*2+1 >= sizeY) {
    Error("Filter specified is larger than available space");
    return(NULL);
  }

  /* set up recurrence array */
  y=(double*)malloc(sizeof(double)*(N+3));
  y[N+2]=0.0;
  y[N+1]=0.0;

  /* create new complex data record for frequency response */
  Hwxwy = CreateDoubleImage(sizeX,sizeY,COMPLEX);

  data = Hwxwy->data;
  /* loop through all frequencies */
  for (loopwy=0;loopwy<sizeY;loopwy++) {
    cwy=cos((double)(loopwy-sizeY/2)/sizeY*2*3.14159265);
    for (loopwx=0;loopwx<sizeX;loopwx++) {
      cwx=cos((double)(loopwx-sizeX/2)/sizeX*2*3.14159265);

      /* compute value of McClellan transformation kernel */
      T=0.5*(-1+cwx+cwy+cwx*cwy);

      /* perform Clenshaw's recurrence */
      for (lc=N;lc>=1;lc--)
        y[lc] = 2*T*y[lc+1] - y[lc+2] + a[lc];

      *(data++) = -1*y[2] + T*y[1] + a[0];
      *(data++) = 0;
    }
  }

  free(a);
  free(y);

  return(Hwxwy);

}
```

B SOFTWARE CODE - LEVENBERG-MARQUARDT MINIMIZATION

This section lists software code which performs the Levenberg-Marquardt minimization for the estimation of the rotational parameters in the three-dimensional DART algorithm. The routines mrqmin() and mrqcof() are adapted from Press (1994).

```
#define HSIZE 2.5          /* the Hamming window size */
#define HSUM  ((4.0/3.0*3.14159265-8.0/3.14159265)*HSIZE*HSIZE*HSIZE)

const double M2 = HSIZE*HSIZE;
const double pi_over_M = 3.14159265/HSIZE;
const double Hamming_gamma = 1.0/HSUM;
const double Hamming_gamma_prime = -3.14159265/HSIZE/HSUM;

extern int data[];         /* the unchanged test image data   k-space matrix*/
extern int ref[] ;         /* the unchanged ref image data k-space matrix */
extern int ndata[3];       /* the size of the data matrices */

double Tinv[9];            /* the transformation matrices */
double dTinv_da1[9];
double dTinv_da2[9];
double dTinv_da3[9];

void mrqmin (float x[], int y[], float sig[], int ndata[], float a[],
             int ia[], int ma, float **covar, float **alpha, float *chisq,
             void (*funcs)(int[], float[], float *, float [], int),
             float *alamda);

void fitfunc(int r0[], float a[], float *y, float dyda[], int ma);

/* Multiply two 3x3 matrices and put result into a third : P = AB*/
void mulMat(double *A,double *B,double *P)
{
  int i,j;
  double *PP=P;

  for (i=0;i<9;i+=3)
    for (j=0;j<3;j++)
      P[i+j] = A[i+0]*B[0+j] + A[i+1]*B[3+j] + A[i+2]*B[6+j];
}

/* Compute the transformation matrices and their derivatives for  a
 * given set of rotational parameters a[] */
void UpdateRotationMatrices (float a[])
{
  double a1=a[1],a2=a[2],a3=a[3];
  double c1=cos(a1),c2=cos(a2),c3=cos(a3);
  double s1=sin(a1),s2=sin(a2),s3=sin(a3);

  /* some temporary matrices */
  double  Tp1[9],   Tp2[9],   Tp3[9];
  double dTp1[9],  dTp2[9],  dTp3[9];
  double temp[9], temp2[9];
```

```
   Tp1[0]=1;   Tp1[1]=0;   Tp1[2]=0;
   Tp1[3]=0;   Tp1[4]=c1;  Tp1[5]=-s1;
   Tp1[6]=0;   Tp1[7]=s1;  Tp1[8]=c1;

   dTp1[0]=0;  dTp1[1]=0;   dTp1[2]=0;
   dTp1[3]=0;  dTp1[4]=-s1; dTp1[5]=-c1;
   dTp1[6]=0;  dTp1[7]=c1;  dTp1[8]=-s1;

   Tp2[0]=c2;  Tp2[1]=0;   Tp2[2]=s2;
   Tp2[3]=0;   Tp2[4]=1;   Tp2[5]=0;
   Tp2[6]=-s2; Tp2[7]=0;   Tp2[8]=c2;

   dTp2[0]=-s2; dTp2[1]=0;  dTp2[2]=c2;
   dTp2[3]=0;   dTp2[4]=0;  dTp2[5]=0;
   dTp2[6]=-c2; dTp2[7]=0;  dTp2[8]=-s2;

   Tp3[0]=c3;  Tp3[1]=-s3; Tp3[2]=0;
   Tp3[3]=s3;  Tp3[4]=c3;  Tp3[5]=0;
   Tp3[6]=0;   Tp3[7]=0;   Tp3[8]=1;

   dTp3[0]=-s3; dTp3[1]=-c3; dTp3[2]=0;
   dTp3[3]=c3;  dTp3[4]=-s3; dTp3[5]=0;
   dTp3[6]=0;   dTp3[7]=0;   dTp3[8]=0;

   mulMat(Tp1,Tp2,temp);
   mulMat(temp,Tp3,Tinv);              /* Tinv = Tp1 Tp2 Tp3 */
   mulMat(temp,dTp3,dTinv_da3);        /* dTinv_da3 = Tp1 Tp2 dTp3_da3 */

   mulMat(dTp1,Tp2,temp);
   mulMat(temp,Tp3,dTinv_da1);         /* dTinv_da1 = dTp1_da1 Tp2 Tp3 */

   mulMat(Tp1,dTp2,temp);
   mulMat(temp,Tp3,dTinv_da2);         /* dTinv_da2 = Tp1 dTp2_da2 Tp3 */
}


void mrqmin (float x[], int y[], float sig[], int ndata[], float a[],
            int ia[], int ma, float **covar, float **alpha, float *chisq,
            void (*funcs)(int[], float[], float *, float [], int),
            float *alamda)
{
```
This code for this routine is based upon the `mrqcof()` routine in "Numerical Recipes in C" by Press et al. (1994) and is not reproduced here. It has been modified to accept as the evaluation function argument `funcs()` a routine which takes a vector position of type `int[]` as its first argument instead of a single argument of type `float` to specify the position at which to evaluate the function.
```
}

void mrqcof(float x[], int y[], float sig[], int ndata[], float a[],
           int ia[], int ma, float **alpha, float beta[], float *chisq,
           void (*funcs)(int[], float [], float *, float [], int))
{
```
The code for this routine is based upon the `mrqcof()` routine in "Numerical Recipes in C" by Press et al. (1994) and is not reproduced here. As with `mrqmin()`, it has been modified to accept as the evaluation function argument

funcs() a routine which takes a vector position of type int[] as its first argument instead of a single argument of type float to specify the position at which to evaluate the function, and to call funcs() with the appropriate first argument. Additionally, a single call to UpdateRotationMatrices(a) is made at entrance, and the summation loop is performed over only a half-plane shell as described in Chapter 4.

}

```
void fitfunc(int r0[], float a[], float *y, float dyda[], int ma)
{
```

This routine is intended to be passed as the argument funcs() to mrqmin(). It evaluates the interim image and its partial derivatives with respect to each of the rotational parameters at the specified position r0[] and returns these values in *y and dyda[].

```
/* r0[]:    the desired position vector: x,y,z
 * a[]:     the rotational parameters
 * *y:      the brain pixel value at position x
 * dyda[]:  the partial derivative with respect to each parameter
 * ma:      the number of parameters
 */

long   r[3];
double x0t[3],x0[3]; /*x0t in -/+ coords, x0 shifted to >=0 */
double x0_r[3];
double 12_x0_r;
double 1_x0_r;
double fac;
long   rlow[3],rhigh[3];

long index;

double dx0_da1[3],dx0_da2[3],dx0_da3[3];

*y=0.0;              /* initialize accumulators */
dyda[1]=dyda[2]=dyda[3]=0.0;

/* NB: x0 is r0 in new notation, r0 is r */
x0t[0]=Tinv[0]*r0[0]+Tinv[1]*r0[1]+Tinv[2]*r0[2];
x0t[1]=Tinv[3]*r0[0]+Tinv[4]*r0[1]+Tinv[5]*r0[2];
x0t[2]=Tinv[6]*r0[0]+Tinv[7]*r0[1]+Tinv[8]*r0[2];

x0[0]=x0t[0]+ndata[0]/2;
x0[1]=x0t[1]+ndata[1]/2;
x0[2]=x0t[2]+ndata[2]/2;

dx0_da1[0]=dTinv_da1[0]*r0[0]+dTinv_da1[1]*r0[1]+dTinv_da1[2]*r0[2];
dx0_da1[1]=dTinv_da1[3]*r0[0]+dTinv_da1[4]*r0[1]+dTinv_da1[5]*r0[2];
dx0_da1[2]=dTinv_da1[6]*r0[0]+dTinv_da1[7]*r0[1]+dTinv_da1[8]*r0[2];

dx0_da2[0]=dTinv_da2[0]*r0[0]+dTinv_da2[1]*r0[1]+dTinv_da2[2]*r0[2];
dx0_da2[1]=dTinv_da2[3]*r0[0]+dTinv_da2[4]*r0[1]+dTinv_da2[5]*r0[2];
dx0_da2[2]=dTinv_da2[6]*r0[0]+dTinv_da2[7]*r0[1]+dTinv_da2[8]*r0[2];

dx0_da3[0]=dTinv_da3[0]*r0[0]+dTinv_da3[1]*r0[1]+dTinv_da3[2]*r0[2];
dx0_da3[1]=dTinv_da3[3]*r0[0]+dTinv_da3[4]*r0[1]+dTinv_da3[5]*r0[2];
dx0_da3[2]=dTinv_da3[6]*r0[0]+dTinv_da3[7]*r0[1]+dTinv_da3[8]*r0[2];
```

```
/* convolve with a spherically symmetric Hamming window */

rlow[0]=x0[0]-HSIZE+1; if (rlow[0]<0) rlow[0]=0;     /* valid bounds */
rlow[1]=x0[1]-HSIZE+1; if (rlow[1]<0) rlow[1]=0;     /* for summation */
rlow[2]=x0[2]-HSIZE+1; if (rlow[2]<0) rlow[2]=0;

rhigh[0]=x0[0]+HSIZE+1; if (rhigh[0]>ndata[0]) rhigh[0]=ndata[0];
rhigh[1]=x0[1]+HSIZE+1; if (rhigh[1]>ndata[1]) rhigh[1]=ndata[1];
rhigh[2]=x0[2]+HSIZE+1; if (rhigh[2]>ndata[2]) rhigh[2]=ndata[2];

for (r[2]=rlow[2];r[2]<rhigh[2];r[2]++) {
  x0_r[2] = x0[2] - r[2];
  for (r[1]=rlow[1];r[1]<rhigh[1];r[1]++) {
    x0_r[1] = x0[1] - r[1];
    index = ((r[2]*ndata[1]+r[1])*ndata[0])+rlow[0];
    for (r[0]=rlow[0];r[0]<rhigh[0];r[0]++,index++) {
    x0_r[0] = x0[0] - r[0];
    if ((l2_x0_r=x0_r[0]*x0_r[0]+x0_r[1]*x0_r[1]+x0_r[2]*x0_r[2])<M2) {

        fac = 1.0 + 1.0*cos(pi_over_M*(l_x0_r=sqrt(l2_x0_r))); /* h(s) */
        *y += (data[index] * fac);
        if (l_x0_r > 0) {
          /* avoid divide by zero, as final product would be zero */
          fac = data[index]*sin(pi_over_M*l_x0_r)/l_x0_r;
          dyda[1] += fac *
            (x0_r[0]*dx0_da1[0] + x0_r[1]*dx0_da1[1] + x0_r[2]*dx0_da1[2]);
          dyda[2] += fac *
            (x0_r[0]*dx0_da2[0] + x0_r[1]*dx0_da2[1] + x0_r[2]*dx0_da2[2]);
          dyda[3] += fac *
            (x0_r[0]*dx0_da3[0] + x0_r[1]*dx0_da3[1] + x0_r[2]*dx0_da3[2]);
        }
      }
    }
  }
}

*y *= Hamming_gamma;   /* complete scalings */
dyda[1] *= Hamming_gamma_prime;
dyda[2] *= Hamming_gamma_prime;
dyda[3] *= Hamming_gamma_prime;
}
```

REFERENCES

Abduljalil AM, Aletra AH, and Robitaille P-ML (1995). 3D echo planar imaging: application to the human head. Magn Reson Med 34: 144-148.

Abragam A (1989). Principles of Nuclear Magnetism. Oxford University Press, New York.

Ahn C, Kim J, and Cho Z (1986). High speed spiral scan echo planar NMR imaging. IEEE Trans Med Imaging MI-2: 2-7.

Aine CJ (1995). A conceptual overview and critique of functional neuroimaging techniques in humans: I. MRI/fMRI and PET. Critical Reviews in Neurobiology 9: 229-309.

Alexander ME and Somorjai RL (1996). The registration of MR images using multiscale robust methods. Magn Reson Imaging 14: 453-468.

Alexander ME, Scarth G, and Somorjai RL (1997). An improved robust hierarchical registration algorithm. Magn Reson Imaging 15: 505.

Altmann J and Reitböck (1984). A fast correlation method for scale- and translation-invariant pattern recognition. IEEE Trans Pat Anal Mach Intel PAMI-6: 46-57.

Apicella A, Keppenhan J, and Nagel J (1989). "Fast multi-modality image matching." Proc. SPIE Medical Imaging III: Image Processing.

Aguirre GK, Zarahn E, and D'Esposito M (1998). A critique of the use of the Kolmogorov-Smirnov (KS) statistic for the analysis of BOLD fMRI data. Magn Reson Med 39: 500-505.

Bandettini P, Jesmanowicz A, Wong E, and Hyde J (1993). Processing strategies for time-course data sets in functional MRI of the human brain. Magn Reson Med 30: 161-173.

Bandettini PA, Wong EC, Hinks RS, Tikofsky RS, and Hyde JS (1992). Time course EPI of human brain function during task activation. Magn Reson Med 25: 390-397.

Basilevsky A (1994). Statistical Factor Analysis and Related Methods: Theory and Applications. John Wiley and Sons, New York, pp. 97-346.

Baudendistel K, Schad LR, Friedlinger M, Wenz F, Schröder J, and Lorenz WJ (1995). Postprocessing of functional MRI data of motor cortex stimulation measured with a standard 1.5 T imager. Magn Reson Imaging 13: 701-707

Baumgartner R, Scarh G, Teichtmeister C, Somorjai R, and Moser E (1997). Fuzzy clustering of gradient-echo functional MRI in the human visual cortex. Part I: reproducibility. J Magn Reson Imaging 7: 1094-1102.

Bedell BJ, Narayana PA, and Johnston DA (1996). Three-dimensional MR image registration of the human brain. Magn Reson Med 35: 384-390.

Belliveau JW, Rosen BR, Kantor HL, Rzedzian RR, Kennedy DN, McKinstry RC, Vevea JM, Cohen MS, Pykett IL, and Brady TJ (1990). Functional cerebral imaging by susceptibility-contrast NMR. Magn Reson Med 14: 538-546.

Belliveau JW, Kennedy DN, McKinstry RC, Buchbinder BR, Weisskoff RM, Cohen MS, Vevea JM, Brady TJ, and Rosen BR (1991). Functional mapping of the human visual cortex by magnetic resonance imaging. Science 254: 716-719.

Birn RM, Bandettini PA, Cox RW, Jesmanowica A, and Shaker R (1998). Magnetic field changes in the human brain due to swallowing or speaking. Magn Reson Med 40: 55-60.

Biswal B, Yetkin Z, Haughton VM, and Hyde JS (1995). Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. Magn Reson Med 34: 537-541.

Biswal B, DeYoe EA, and Hyde JS (1996). Reduction of physiological fluctuations in fMRI using digital filters. Magn Reson Med 35: 107-113.

Biswal BB and Hyde JS (1997). Contour-based registration technique to differentiate between task-activated and head motion-induced signal variations in fMRI. Magn Reson Med 38: 470-476.

Blamire A, Ogawa S, Ugurbil K, Rothman D, McCarthy G, Ellermann J, Hyder F, Rattner Z, and Shulman R (1992). Dynamic mapping of the human visual cortex by high-speed magnetic resonance imaging. Proc Natl Acad Sci USA 90: 11069-11073.

Bock C, Schmitz B, Kerskens CM, Gyngell ML, Hossmann K-A, and Hoehn-Berlage M (1998). Functional MRI of somatosensory activation in rat: effect of hypercapnic up-regulation on perfusion- and BOLD-imaging. Magn Reson Med 39: 457-461.

Boxerman JL, Bandettini PA, Kwong KK, Baker JR, Davis TL, Rosen BR, and Weisskoff RM (1995). The intravascular contribution to fMRI signal change: Monte Carlo modeling and diffusion-weighted studies in vivo. Magn Reson Med 34: 4-10.

Boxerman JL, Hamberg LM, Rosen BR, and Weisskoff (1995A). MR contrast due to intravascular magnetic susceptibility perturbations. Magn Reson Med 34: 555-566.

Breiter HC, Rauch SL, Kwong KK, Baker JR, Weisskoff RM, Kennedy DN, Kendrick AD, David TL, Jiang A, Cohen MS, Stern CE, Belliveau JW, Baer L, O'Sullivan RL, Savage CR, Jenike MA, and Rosen BR (1996). Functional magnetic resonance imaging of symptom provocation in obsessive-compulsive disorder. Arch Gen Psychiatry 53: 595-606.

Bullmore E, Brammer M, Williams SCR, Rabe-Hesketh S, Janot N, David A, Mellers J, Howard R, and Sham P (1996). Statistical methods of estimation and inference for functional MR image analysis. Magn Reson Med 35: 261-277.

Bushberg JT, Seibert JA, Leidholdt EM Jr., Boone JM (1994). The Essential Physics of Medical Imaging. Williams and Wilkins, Baltimore.

Buxton RB, Wong EC, and Frank LR (1998). Dynamics of blood flow and oxygenation changes during brain activation: the balloon model. Magn Reson Med 39: 855-864.

Casasent D and Psaltis D (1976). Position, rotation, and scale invariant optical correlation. Appl Opt 15: 1795-1799.

Chen Q, Andersen AH, Zhang Z, Ovadia A, Gash DM, and Avison MJ (1996). Mapping drug-induced changes in cerebral R2* by multiple gradient recalled echo functional MRI. Magn Reson Imaging 14: 469-476.

Chen Q, Siewert B, Bly BM, Warach S, and Edelman RR (1997A) STAR-HASTE: perfusion imaging without magnetic susceptibility artifact. Magn Reson Med 38: 404-408.

Chen W and Zhu X-H (1997). Suppression of physiological eye movement artifacts in functional MRI using slab presaturation. Magn Reson Med 38: 546-550.

Chen W, Kato T, Zhu X-H, Strupp J, Ogawa S, and Ugurbil K (1998A). Mapping of lateral geniculate nucleus activation during visual stimulation in human brain using fMRI. Magn Reson Med 39: 89-96.

Chen W, Zhu X-H, Kato T, Andersen P, and Ugurbil K (1998). Spatial and temporal differentiation of fMRI BOLD response in primary visual cortex of human brain during sustained visual stimulation. Magn Reson Med 39: 520-527.

Childress A, McElgin W, Mozley D, Reivich M, and O'Brien C (1996). Brain correlates of cue-induced cocaine and opiate craving [abstract]. Soc Neurosci Abstr 22(2): 933.

Cideciyan AV, Jacobson SG, Kemp CM, et al. (1992). Registration of high resolution images on the retina. Medical Imaging VI: Image Processing, SPIE Vol 1652, 310-322.

Cohen MS and Rohan ML (1989). "3D volume imaging with instant scan." Proc. Society for Magnetic Resonance in Medicine, 8th annual meeting, p. 189.

The Committee on the Mathematics and Physics of Emerging Dynamic Biomedical Imaging, the National Research Council and the Institute of Medicine (1996). Mathematics and Physics of Emerging Biomedical Imaging. National Academy Press, Washington, D.C.

Connelly A (1995). Ictal imaging using functional magnetic resonance. Magn Reson Imaging 13: 1233-1237.

Constable RT, Skudlarski P, and Gore JC (1995). An ROC approach for evaluating functional brain MR imaging and postprocessing protocols. Magn Reson Med 34: 57-64.

Cox RW, Jesmanowicz A, and Hyde JS (1995). Real-time functional magnetic resonance imaging. Magn Reson Med 33: 230-236.

Crémillieux Y, Wheeler-Kingshott CA, Briguet A, and Doran SJ (1997). STEAM-Burst: a single-shot, multi-slice imaging sequence without rapid gradient switching. Magn Reson Med 38: 645-652.

Debus J, Essig M, Schad LR, Wenz F, Baudendistel K, Knopp MV, Engenhart R, and Lorenz WJ (1996). Functional magnetic resonance imaging in a stereotactic setup. Magn Reson Imaging 14: 1007-1012.

Duyn JH, Yang Y, Frank JA, Mattay VS, and Hou L (1996). Functional magnetic resonance neuroimaging data acquisition techniques. Neuroimage 4: S76-S83.

Dymond RC and Norris DG (1997). Mechanism and echo time dependence of the fast response in fMR. Magn Reson Med 38: 1-7.

Eddy W, Fitzgerald M, and Noll D (1996). Improved image registration by using Fourier interpolation. Magn Reson Med 36: 923-931.

Edelman RR, Siewer B, Darby DG, Thangaraj V, Nobre AC, Mesulam MM, and Warach S (1994). Qualitative mapping of cerebral blood flow and functional localization with echo-planar MR imaging and signal targeting with alternating radio frequency. Radiology 192: 513-520.

Ernst T and Hennig J (1994). Observation of a fast response in fMR. Magn Reson Med 32: 146-149.

First MB, Spitzer RL, Gibbon M, and Williams JBW (1995). Structured Clinical Interview for DSM-IV Axis I Disorders - Patient Edition (SCID-I/P, Version 2.0). Biometrics Research Department, New York State Psychiatric Institute.

Forman SD, Cohen JD, Fitzgerald M, Eddy WF, Mintun MA, and Noll DC (1995). Improved assessment of significant activation in functional magnetic resonance imaging (fMRI): use of a cluster-size threshold. Magn Reson Med 33: 636-647.

Fox PT and Raichle ME (1986). Focal physiological uncoupling of cerebral blood flow and oxidative metabolism during somatosensory stimulation in human subjects. Proc Natl Acad Sci USA 83: 1140-1144.

Frahm J, Krüger G, Merboldt K-D, and Kleinschmidt A (1996). Dynamic uncoupling and recoupling of perfusion and oxidative metabolism during focal brain activation in man. Magn Reson Med 35: 143-148.

Frank LR, Buxton RB, and Wong EC (1998). Probabilistic analysis of functional magnetic resonance imaging data. Magn Reson Med 39: 132-148.

Fransson P, Krüger G, Merboldt K-D, Frahm J (1998). Temporal characteristics of oxygenation-sensitive MRI responses to visual activation in humans. Magn Reson Med 39: 912-919.

Friston KJ, Jezzard P, and Turner R (1994). Analysis of functional MRI time-series. Human Brain Mapping 1: 153-171.

Friston KJ, Williams S, Howard R, Frackowiak RS, and Turner R (1996). Movement-related effects in fMRI time-series. Magn Reson Med 35: 346-355.

Fu ZW, Wang Y, Grimm RC, Rossman PJ, Felmlee JP, Riederer SJ, and Ehman RL (1995). Orbital navigator echoes for motion measurements in magnetic resonance imaging. Magn Reson Med 34: 746-753.

Gao J-H, Xiong J, Lai S, Haacke EM, Woldorff MG, Li J, and Fox PT (1996) Improving the temporal resolution of functional MR imaging using keyhole techniques. Magn Reson Med 35: 854-860.

Gaschler-Markefski B, Baumgart F, Tempelmann C, Schindler F, Stiller D, Heinze H-J, and Scheich H (1997). Statistical methods in functional magnetic resonance imaging with respect to nonstationary time-series: auditory cortex activity. Magn Reson Med 38: 811-820.

Gati JS, Menon RS, Ugurbil K, and Rutt BK (1997). Experimental determination of the BOLD field strength dependence in vessels and tissue. Magn Reson Med 38: 296-302.

Genovese CR, Noll DC, and Eddy WF (1997). Estimating test-retest reliability in functional MR imaging I: statistical methodology. Magn Reson Med 38: 497-597.

Glover GH and Lai Song (1998). Self-navigated spiral fMRI: interleaved versus single-shot. Magn Reson Med 39: 361-368.

Grant S, London ED, Newlin DB, Villemagne VL, Liu X, Contoreggi C, Phillips RL, Kimes AS, and Margolin A (1996). Activation of memory circuits during cue-elicited cocaine craving. Proc Natl Acad Sci USA 93: 12040-12045.

Hajnal J, Myers R, Oatridge A, Schwieso J, Young I, and Bydder G (1994). Artifacts due to stimulus correlated motion in functional imaging of the brain. Magn Reson Med 31: 283-291.

Hajnal JV, Saeed N, Soar EJ, Oatridge A, Young IR, and Bydder GM (1995). A registration and interpolation procedure for subvoxel matching of serially acquired MR images. J Comput Assist Tomogr 19: 289-296.

Harvey PR and Mansfield P (1996). Echo-volumar imaging (EVI) at 0.5T: first whole-body volunteer studies. Magn Reson Med 35: 80-88.

Hill DLG, Hawkes DJ, Crossman JE, Gleeson MH, Cox TC, Bracey EE, Strong AJ, and Graves P (1991). Registration of MR and CT images for skull base surgery using point-like anatomical features. Br J Radiol 64: 1030-1035.

Hoogenraad FGC, Reichenbach JR, Haacke EM, Lai S, Kuppusamy K, and Sprenger M (1998). In vivo measurement of changes in venous blood-oxygenation with high resolution functional MRI at 0.95 Tesla by measuring changes in susceptibility and velocity.

Howard RJ, Bullmore E, Brammer M, Williams SCR, Mellers J, Woodruff P, David A, Andrew C, Allin M, Simmons A, and Cox T (1995). Activation of area V5 by viusal perception of motion demonstrated with echoplanar imaging. Magn Reson Med 13: 907-909.

Howard RJ, Ellis C, Bullmore ET, Brammer M, Mellers JD, Woodruff PW, David AS, Simmons A, Williams SC, and Parkes JD (1996). Functional echoplanar brain imaging correlates of amphetamine administration to normal subjects and subjects with the narcoleptic syndrome. Magn Reson Imaging 14: 1013-1016.

Hu X and Kim S-G (1994). Reduction of signal fluctuation in functional MRI using navigator echoes. Magn Reson Med 31: 495-503.

Hu X, Le TH, Parrish T, and Erhard P (1995). Retrospective estimation and correction of physiological fluctuation in functional MRI. Magn Reson Med 34: 201-212.

Jackson J, Meye C, Nishimura G, and Macovski A (1991). Selection of a convolution function for Fourier inversion using gridding. IEEE Trans Med Imaging 10: 473-478.

Jensen AS, Lindvold L, and Rasmussen E (1987). Transformation of image positions, rotations, and sizes into shift parameters. Appl Opt 26: 1775-1781.

Jezzard P, Le Bihan D, Cuenod C, Pannier L, Prinster A, and Turner R (1993). "An investigation of the contribution of physiological noise in human functional MRI studies at 1.5 Tesla and 4 Tesla." Proc. SMRM 12th Meeting.

Jezzard P and Balaban RS (1995). Correction for geometric distortion in echo planar images from B0 field variations. Magn Reson Med 34: 65-73.

Jezzard P and Song AW (1996). Technical foundations and pitfalls of clinical fMRI. Neuroimage 4: S63-S75.

Jezzard P, Rauschecker JP, and Malonek D (1997). An in vivo model for functional MRI in cat visual cortex. Magn Reson Med 38: 699-705.

Jones RA, Schirmer T, Lipinski B, Elbel GK, and Auer DP (1998). Signal undershoots following visual stimulation: a comparison of gradient and spin-echo BOLD sequences. Magn Reson Med 40: 112-118.

Kassam A and Woods M (1996). Fourier registration of three-dimensional brain MR images: exploiting the axis of rotation. J Magn Reson Med 6: 894-902.

Kim S-G, Ashe J, Hendrich K, Ellermann J, Merkle H, Ugurbil K, and Georgopoulos A (1993). Functional magnetic resonance imaging of motor cortex: hemispheric asymmetry and handedness. Science 261: 615-617.

Kim S-G (1995). Quantification of relative cerebral blood flow change by flow-sensitive alternating inversion recovery (FAIR) technique: application to functional mapping. Magn Reson Med 34: 293-301.

Kim S-G, Richter W, and Ugurbil K (1997A). Limitations of temporal resolution in functional MRI. Magn Reson Med 37: 631-636.

Kim S-G and Ugurbil K (1997). Comparison of blood oxygenation and cerebral blood flow effects in fMRI: estimation of relative oxygen consumption change. Magn Reson Med 38: 59-65.

King KF, Foo TK, and Crawford CR (1995). Optimized gradient waveforms for spiral scanning. Magn Reson Med 34: 156-160.

Koszlowski L and Wilkinson D (1987). Use and misuse of the concept of craving by alcohol, tobacco and drug researchers. Br J Addict 82: 31-36.

Kucharczyk J, Moseley M, and Barkovich AJ (1994). Magnetic Resonance Neuroimaging. CRC Press, Boca Raton.

Kwong KK, Belliveau JW, Chesler DA, Goldberg IE, Weisskoff RM, Poncelet BP, Kennedy DN, Hoppel BE, Cohen MS, Turner R, Cheng H-M, Brady TJ, and Rosen BR (1992). Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. Proc Natl Acad Sci USA 89: 5675-5679.

Kwong K (1995). Functional magnetic resonance imaging with echo planar imaging. Magn Reson Q 11: 1-20.

Lange N (1994). Some computational and statistical tools for paired comparisons of digital images. Stat Meth Med Res 3: 23-41.

Le Bihan D, moderator (1995). Functional magnetic resonance imaging of the brain. Ann Intern Med 122: 296-303.

Le Bihan D (1996). Functional MRI of the brain: principles, applications and limitations. J Neuroradiol 23: 1-5.

Lee AT, Glover GH, and Meyer CH (1995). Discrimination of large venous vessels in time-course spiral blood-oxygen-level-dependent magnetic-resonance functional neuroimaging. Magn Reson Med 33: 745-754.

Lee CC, Jack CCR Jr., Grimm RC, Rossman PJ, Felmlee JP, Ehman RL, and Riederer SJ (1996). Real-time adaptive motion correction in functional MRI. Magn Reson Med 36: 436-444.

Lee CC, Grimm RC, Manduca A, Felmlee JP, Ehman RL, Riederer SJ, and Jack CR Jr. (1998) A prospective approach to correct for inter-image rotation in FMRI. Magn Reson Med 39: 234-243.

Lev MH, Kulke SF, Sorensen AG, Boxerman JL, Brady TJ, Rosen BR, Buchbinder BR, and Weisskoff RM (1997). Contrast-to-noise ratio in functional MRI of relative cerebral blood volume with Sprodiamide injection. J Magn Reson Imaging 7: 523-527.

Lim J (1990). Two-Dimensional Signal and Image Processing. Prentice Hall, New Jersey.

Lin W, Kuppusamy K, Haacke EM, and Burton H (1996). Functional MRI in human somatosensory cortex activated by touching textured surfaces. J Magn Reson Imaging 6: 565-572.

Lowe MJ and Sorenson JA (1997). Spatially filtering functional magnetic resonance imaging data. Magn Reson Med 37: 723-729.

Maas LC, Frederick BB and Renshaw PF (1995). "Fast decoupled rotational and translational registration of fMRI data sets." Proc Society of Magnetic Resonance, Third Scientific Meeting, Nice, France, August, 1995, p. 815

Maas LC, Frederick BB and Renshaw PF (1997). Decoupled automated rotational and translational registration for functional MRI time series data: The DART registration algorithm. Magn Reson Med 37: 131-139.

Maas LC, Harris GJ, Satlin A, English CD, Lewis RF, and Renshaw PF (1997A). Regional cerebral blood volume measured by dynamic susceptibility contrast MRI in Alzheimer's disease: a principal components analysis. J Magn Reson Imaging 7:215-219.

Maas LC, Anderson CM, Teicher MH, and Renshaw PF (1997B). "Hurst analysis of physiological fluctuations in functional MRI." Proc Third International Conference on Functional Mapping of the Human Brain, Copenhagen, Denmark, May, 1997.

Maas LC, Lukas SE, Kaufman MJ, Weiss RD, Daniels SL, Rogers VW, Kukes TJ, and Renshaw PF (1998). Functional MRI of human brain activation during cue-induced cocaine craving. Am J Psychiatry 155:124-146.

Maas LC and Renshaw PF (1998B). "Resampling noise in motion corrected functional MRI data sets." Proc International Society for Magnetic Resonance in Medicine, Sixth annual scientific meeting, Sydney, Australia, April, 1998, p. 717.

Mandeville JB, Marota JJA, Kosofsk BE, Keltner JR, Weissleder R, Rosen BR, and Weisskoff RM (1998). Dynamic functional imaging of relative cerebral blood volume during rat forepaw stimulation. Magn Reson Med 39: 615-624.

Mansfield P (1977). Multi-planar image formation using NMR spin echoes. J Phys C 10: L55-L58.

Mansfield P, Harvey PR, and Stehling MK (1995). Echo volumar imaging. MAGMA 2: 291-294.

McCarthy G, Blamire AM, Rothman DL, Gruetter R, and Shulman RG (1993). Echo-planar magnetic resonance imaging studies of frontal cortex activation during word generation in humans. Proc Natl Acad Sci USA 90: 4952-4956.

McCarthy G, Blamire AM, Puce A, Nobre AC, Bloch G, Hyder F, Goldman-Rakic P, and Shulman RG (1994). Functional magnetic resonance imaging of human prefrontal cortex activation during a spatial working memory task. Proc Natl Acad Sci USA 91: 8690-8694.

Merboldt K-D, Krüger G, Hänicke W, Kleinschmidt A, and Frahm J (1996). Functional MRI of human brain activation combining high spatial and temporal resolution by a CINE FLASH technique. Magn Reson Med 34: 639-644.

Mitra PP, Ogawa S, Hu X, and Ugurbil K (1997). The nature of spatiotemporal changes in cerebral hemodynamics as manifested in functional magnetic resonance imaging. Magn Reson Med 37: 511-518.

Moser E, Teichtmeister C, and Dimling M (1996). Reproducibility and postprocessing of gradient-echo functional MRI to improve localization of brain activity in the human visual cortex. Magn Reson Imaging 14: 567-569.

Noll DC and Schneider W (1994). "Respiration artifacts in functional brain imaging: sources of signal variation and compensation strategies." Proc. Society of Magnetic Resonance, Second Scientfic Meeting, San Francisco.

Noll D, Cohen J, Meyer C, and Schneider W (1995). Spiral k-space MR imaging of cortical activation. J Magn Reson Imaging 5: 49-57.

Noll DC, Boada FE, and Eddy WF (1997A). A spectral approach to analyzing slice selection in planar imaging: optimization for through-plane interpolation. Magn Reson Med 38: 151-160.

Noll DC, Genovese CR, Nystrom LE, Vazquez AL, Forman SD, Eddy WF, and Cohen JD (1997). Estimating test-retest reliability in functional MR imaging II: application to motor and cognitive activation studies. Magn Reson Med 38: 508-517,

Noll DC, Peltier SJ, and Boada FE (1998). Simultaneous multislice acquisition using rosette trajectories (SMART): a new imaging method for functional MRI. Magn Reson Med 39: 709-716.

O'Brien C, Alterman A, Childress A, and McLellan A (1992). Developing and evaluating new treatments for alcoholism and cocaine dependence. Recent Dev Alcohol 10: 303-325.

Ogawa S, Lee T, Kay A, and Tank D (1990). Brain magnetic resonance imaging with contrast dependent on blood oxygenation. Proc Natl Acad Sci USA 87: 9868-9872.

Ogawa S, Tank D, Menon R, Ellermann J, Kim S-G, Merkle H, and Ugurbil K (1992). Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. Proc Natl Acad Sci USA 89: 5951-5955.

Ogawa S, Menon RS, Tank DW, Kim S-G, Merkle H, Ellermann JM, and Ugurbil K (1993). Functional brain mapping by blood oxygenation level-dependent contrast magnetic resonance imaging, a comparison of signal characteristics with a biophysical model. Biophys J 64: 803.

Oppenheim AV and Schafer RW (1989). Discrete Time Signal Processing. Prentice Hall, New Jersey.

Pauling L and Coryell CD (1936). The magnetic properties and structure of hemoglobin, oxyhemoglobin and carbonmonoxyhemoglobin. Proc Natl Acad Sci USA 22: 210.

Pelizzari CA, Chen GTY, Spelbring DR, Weichslebaum RR, and Chen C-T (1989). Accurate three-dimensional registration of CT, PET and/or MR images of the brain. J Comput Assist Tomog 13: 20-26.

Poncelet B, Wedeen V, Weiskoff R and Cohen M (1992). Brain parenchyma motion: measurement with cine echo-planar MR imaging. Radiology 185: 645-51.

Press WH, Teukolsky SA, Vetterling WT, and Flannery BP (1994). Numerical Recipes in C: The Art of Scientific Computing, 2nd edition. New York, Cambridge University Press.

Renshaw PF, Yurgelun-Todd DA, and Cohen BM (1994). Greater hemodynamic response to photic stimulation in schizophrenic patients: an echo planar MRI study. Am J Psychiatry 151:1493-1495.

Robbins S, Ehrman R, Childress A, and O'Brien CP (1992). Using cue reactivity to screen medications for cocaine abuse: a test of amantadine hydrochloride. Addictive Behaviors 17: 491-499.

Ross MH, Yurgelun-Todd DA, Renshaw PF, Maas LC, Mendelson JH, Mello NK, Cohen BM, and Levin JM (1997). Age-related reduction in functional MRI response to photic stimulation. Neurology 48: 173-176.

Schwarzbauer C and Heinke W (1998). BASE imaging: a new spin labeling technique for measuring absolute perfusion changes. Magn Reson Med 39: 717-722.

Sekihara K and Koizumi H (1996). Detecting cortical activities from fMRI time-course data using the MUSIC algorithm with forward and backward covariance averaging. Magn Reson Med 35: 807-813.

Song AW, Wong EC, and Hyde JS (1994). Echo-volume imaging. Magn Reson Med 32: 668-671.

Sorensen AG, Tievsky AL, Ostergaard L, Weisskoff RM, and Rosen BM (1997). Contrast agents in functional MR imaging. J Magn Reson Imaging 7: 47-55.

Sorenson JA and Wang X (1996). ROC methods for evaluation of fMRI techniques. Magn Reson Med 36: 737-744.

Talagala SL and Noll DC (1998). Functional MRI using steady-state arterial water labeling. Magn Reson Med 39: 179-183.

Thulborn KR, Waterton JC, Matthews PM, and Radda GK (1982). Oxygenation dependence of the transverse relaxation time of water protons in whole blood at high fields. Biochim Biophys Acta 714: 265-270.

Turner R, Le Bihan D, Moonen C, Despres D, and Frank J (1991). Echo-planar time course MRI of cat brain oxygenation changes. Magn Reson Med 22: 159-166.

Turner R, Jezzard P, Wen H, Kwong K, Le Bihan D, Zeffiro T and Balaban R (1993). Functional mapping of the human visual cortex at 4 and 1.5 Tesla using deoxygenation contrast EPI. Magn Reson Med 29: 277-279.

Unser M, Thevenaz P, and Yaroslavsky (1995). Convolution-based interpolation for fast, high-quality rotation of images. IEEE Trans Image Proc 4: 1371-1381.

Weiss R, Griffin M, and Hufford C (1995). Craving in hospitalized cocaine abusers as a predictor of outcome. Am J Drug Alcohol Abuse 21: 289-301.

Weisskoff R, Baker J, Belliveau J, Davis T, Kwong K, Cohen M, and Rosen B (1993). "Power spectrum analysis of functionally-weighted MR data: what's in the noise?" Proc. SMRM 12th Meeting,

Wenz F, Schad LR, Knopp MV, Baudendistel KT, Floömer F, Schröder J, and van Kaick G (1994). Functional magnetic resonance imaging at 1.5T: activation pattern in schizophrenic patients receiving neuroleptic medication. Magn Reson Imaging 12: 975-982.

Woods RP, Cherry SR, and Mazziotta JC (1992). Rapid automated algorithm for aligning and reslicing PET images. J Comput Assist Tomog 16(4): 620-633.

Wu DH, Lewin JS, and Duerk JL (1997). Inadequacy of motion correction algorithms in functional MRI: role of susceptibility-induced artifacts. J Magn Reson Med 7: 365-370.

Xiong J, Gao J, Lancaster J, and Fox P (1996). Assessment and optimization of functinoal MRI analyses. Human Brain Mapping 4: 153-167.

Yang QX, Dardzinski BJ, Li Shizhe, Eslinger PJ, and Smith MB (1997). Multi-gradient echo with susceptibility inhomogeneity compensation (MGESIC): demonstration of fMRI in the olfactory cortex at 3.0 T. Magn Reson Med 37: 331-335.

Yang X, Hyder F, and Shulman RG (1998). Functional MRI BOLD signal coincides with electrical activity in the rat whisker barrels. Magn Reson Med 38: 874-877.

Yang Y, Glover GH, van Gelderen P, Mattay VS, Santha AKS, Sexton RH, Ramsey NF, Moonen CTW, Weinberger DR, Frank JA, and Duyn JH (1996). Fast 3D functional magnetic resonance imaging at 1.5 T with spiral acquisition. Magn Reson Med 36:620-626.

Yang Y, Mattay VS, Weinberger DR, Frank JA, and Duyn JH (1997A). Localized echo-volume imaging methods for functional MRI. J Magn Reson Imaging 7: 371-375.

Yang Y, Frank JA, Hou L, Ye FQ, McLaughlin AC, and Duyn JH (1998). Multislice imaging of quantitative cerebral perfusion with pulsed arterial spin labeling. Magn Reson Med 39: 825-832.

Yurgelun-Todd DA, Waternaux CM, Cohen BM, Gruber SA, English CD, and Renshaw PF (1996). Functional magnetic resonance imaging of schizophrenic pateints and comparison subjects during word production. Am J Psychiatry. 153:200-205.

Yurgelun-Todd DA, Baird A, Gruber S, Steingard R, Maas LC, Cohen BM and Renshaw PF (1997). "FMRI studies of normal adolescents: recognition of emotion in facial expression." Proc International Society of Magetic Resonance in Medicine, Fifth annual scientific meeting, Vancouver, B.C., April, 1997, p. 719.

Zeffiro T (1996). Clinical functional image analysis: artifact detection and reduction. Neuroimage 4: S95-S100.

Zhang R, Cox RW, and Hyde JS (1997). The effect of magnetization transfer on functional MRI signals. Magn Reson Med 38: 187-192.