Chapter 1

Overview

1.1 Introduction

Magnetic resonance imaging (MRI) is a noninvasive method of mapping the internal structure of the body, which completely avoids the use of ionizing radiation. It employs radio-frequency (RF) radiation in the presence of carefully controlled magnetic fields in producing high quality cross-sectional images of the body in any plane. MRI has the advantage that the contrast between different tissues can be manipulated in order to highlight pathological changes by altering the pattern of applied RF pulses.

The phenomenon of nuclear magnetic resonance (NMR) was first uncovered in 1946 by separate research groups lead by Bloch and Purcell [1] [2], a discovery for which they were jointly awarded the Nobel Prize for physics in 1952. Nearly 20 years later, in 1972, the first MR image was produced by Lauterbur [3], followed by the first human whole body scan by Damadian [4] in 1977. Over the past two decades, there has been a rapid evolution in MRI, from a period when work was confined to laboratory prototypes, to the widespread installation of production machines in clinical use. The rapid growth of MRI is mainly due

to the considerable expertise invested in the design and manufacture of magnet systems as well as the development of signal processing techniques involved in the process from data acquisition to image display.

1.2 Motivation and Aims

The popularity of MRI, over other imaging disciplines such as computed tomography (CT) and ultrasound, depends largely on its high spatial resolution and soft tissue contrast. In order to achieve such high quality, and diagnostically interpretable images, relatively long image acquisition times are required, making it vulnerable to physiological or patient motion-induced artifacts. In two-dimensional (2D) *Fourier imaging* [5], these artifacts appear predominantly as blurring or ghost repetitions of the moving structures. Ghost artifacts can obscure vital anatomical details when superimposed on other structures, hindering the clinical diagnosis.

The techniques that have been proposed for minimizing motion artifacts can be broadly divided into two major categories. Some techniques attempt to *prevent* the motion from corrupting the data, whereas the others use signal processing techniques to *restore* the image, incorporating the information embedded within the corrupted data. Almost all the current techniques in widespread use fall into the former category. These techniques depend basically on restricting or controlling the acquisition times. Although these are effective in limiting artifacts produced by periodic motion, the artifact suppression is inadequate in general, specially against voluntary random movements of the patients. Patient restraints are currently used in preventing such motion. However, this requires active cooperation from the patient. In pediatric studies, sedatives are administered to prevent patient motion [6], which increase both the risk to the patient and the operating cost. It was observed that small angles of rotation can occur in spite

of restraint measures, which are more effective at limiting nodding motions than rotations in the transverse plane [7].

The techniques used in the second category are usually called post-processing methods. These are useful in restoring the images corrupted by motion which are immune to motion prevention techniques. Initial investigations regarding the post-processing techniques for the specific use in motion artifact suppression were instigated by Wood, Henkelman and Haacke simultaneously during 1985-1986 [14] [15]. Since then, a considerable research effort has been invested in developing such methods. Post-processing techniques usually require a fitting motion model which can approximate the actual patient motion. The parameters associated with the chosen model must be measured or calculated using the information extracted from the acquired data. Despite its prominence in the research arena, post-processing methods have not been implemented for widespread use. This is mainly due to the inadequacy of the existing models and the enormous computational effort required for parameter estimation and consequent restoration of data. However, the quest for sophisticated post-processing algorithms has been continued due to the premise that post-processing techniques may be used in conjunction with existing *motion prevention* techniques in order to achieve greater suppression of motion artifacts than either type of technique can provide alone. Therefore, the motion *prevention* and *restoration* techniques should not be viewed as competing, but rather complementing each other.

The primary objective of this thesis is to develop post-processing algorithms for the suppression of artifacts caused by in-plane rotations. In practice, in-plane rotations may occur due to the patient rotating his/her head around the neck, or due to patient limb movements around joints. However, such rotations are usually accompanied by translational motion. The translational motion model has aroused the greatest interest due to the fact that gross translational motion is space invariant and the artifact effects can be modelled using linear sets of

equations [8]. In contrast, the rotational and dilational models have attracted relatively less attention due to their highly non-linear effects. The rotational motion correction techniques already proposed in the literature are limited to the suppression of artifacts caused by *small angle* rotations. Therefore, the techniques developed in this thesis are aimed at suppressing artifacts with no limitations on the extent of the rotation angles. Due to the difficulties encountered in obtaining sufficiently accurate rotation information via sensors, the object rotation function is not assumed to be known *a priori*. Methods of estimating the rotation angles from corrupted data are investigated, in spite of the presence of concurrent translational motion.

A major goal of this work is to develop algorithms that are effective as well as efficient and easy to implement. Therefore, significant time and effort has been invested in simplifying algorithms and developing a computationally efficient motion estimation scheme.

1.3 MR Imaging Techniques

This section is intended to provide a comprehensive background knowledge on MR imaging, from MR signal production to image reconstruction. The 2DFT spin-warp imaging is described in detail, due to its impact on the algorithms presented in this thesis.

1.3.1 MR Signal

MR imaging depends on the properties of the nuclei of atoms. Nuclei suitable for use in MRI are those which have an odd number of protons or neutrons which possess a net magnetic moment, and therefore behave as magnetic dipoles. When

placed in a strong, uniform magnetic field, these nuclei tend to line up with the direction of the magnetic field. However, they do not line up precisely with the external field, instead they stop at an angle, and continue to spin about the direction of the field. This spinning motion is called *precession*, and occurs at a specific resonant frequency for a given nucleus in a magnetic field. This relationship is expressed as the Larmor equation

$$f = \gamma B \tag{1.1}$$

where f is the resonant frequency, γ is the gyromagnetic ratio and B is the applied magnetic field.

MR imaging often involves the element hydrogen, the nucleus of which is a single proton. Hydrogen nuclei are not only the strongest nuclear magnets on a per nucleus basis, but also found in abundance in the human body. This makes hydrogen the strongest MR signal emitting element. Nevertheless, many other nuclei in the body are potential candidates for imaging, such as sodium, potassium and phosphorus. However, the MR signals from non-hydrogen nuclei are very weak, their imaging requires more elaborate equipment, and the resulting anatomic detail is inadequate.

The precessing nuclei can absorb energy if they are exposed to pulses of radio frequency (RF) waves, provided that the RF waves are of the same frequency as the frequency of the nuclear precession. The absorption of energy occurs through the process of *resonance*. When the RF pulse is removed, the excited nuclei relax by releasing the excess absorbed energy in the form of an RF wave, which represents the MR signal, and is known as the *free induction decay* or FID. This MR signal can be received by a sensitive antenna coil. The received signal is then amplified and processed by a computer to produce a cross-sectional image of the scanned object.

The magnitude and length of the FID is determined by the nuclear relaxation

times, which reflect molecular motion. The first of these relaxation times, T_1, or the *spin-lattice* relaxation time, represents the time taken by the system of nuclei to return to thermal equilibrium after the RF pulse, by releasing their excess energy to the general environment or lattice (i.e. arrangement of atoms in a substance). The second relaxation time, T_2, indicates the rate of energy release by excited nuclei through interaction among themselves, and is appropriately named *spin-spin* relaxation time. The signals produced by MR imaging techniques contain a combination of proton density, T_1 and T_2 information. Images weighted toward any one of these three parameters can be obtained by stimulating the nuclei with specific RF pulse sequences. In most imaging schemes, a short T_1 gives a strong MR signal in T_1 weighted images. Conversely, a long T_2 gives a strong signal in T_2 weighted images [9].

The other properties that influence image appearance are flow, magnetic susceptibility, diffusion , chemical shifts, temperature and magnetization transfer.

The RF pulse sequences determine the dependence of the signal intensity of an MR image voxel on the basic image parameters described above. The principal RF pulse sequences used in MRI are,

- *partial saturation (PS)*, which typically utilizes a 90^o RF pulse, but can use a smaller pulse;
- *spin echo (SE)*, which utilizes a 90^o pulse followed at time $TE/2$ by a 180^o pulse, where TE is known as the *echo time*, which represents the time interval from the RF pulse to the middle of the data collection;
- *inversion recovery (IR)*, which utilizes a 180^o pulse followed by a 90^o pulse. The time difference between these pulses is known as the inversion time, denoted by TI.

There are many other pulse sequences such as *steady state free precession (SSFP)* and *echo planar imaging (EPI)*.

At the current levels of magnetic field strength and RF pulse energy transmission, there are no known hazards intrinsic to patient exposure to MR imaging. However, the magnetic field itself is potentially hazardous to certain patients who have artificial metallic or electronic implants inside them. For example, cardiac pacemakers may be adversely affected by the strong magnetic fields. Aneurysm clips on blood vessels within the skull could be twisted by the magnetic field and vessels could be torn. Therefore, caution must be used in allowing patients or visitors to enter the magnet room [9].

1.3.2 Fourier Transform Imaging

The two-dimensional Fourier transform (2DFT) is the most commonly used imaging technique in MRI. The beginning of 2DFT imaging can be traced back to 1974, when Kumar et al. [5] proposed the method of *Fourier Zeugmatography*. The concept behind *Fourier Zeugmatography* is very simple and can be implemented in either the planar or full three-dimensional mode. In this section, a review of 2DFT imaging will be presented, in order to introduce the notation and some of the assumptions that will be used throughout this thesis.

In order to describe the 2DFT imaging technique, the imaged object is placed in an arbitrary coordinate system. The direction of the strong magnetic field conventionally defines the physical Z axis known as the *slice selection axis*, which is generally along the longitudinal axis of the reclined patient in an MRI machine. However, the logical Z axis for slice selection can be along an arbitrary direction in space. The X and Y axes are perpendicular to the Z axis as well as to each other, and represent the transverse plane. It is assumed that the imaged slice is contained entirely within the transverse plane. Most current MRI machines use

the method of selective excitation to isolate the image plane. A gradient field, $G_z(t)$ is applied during excitation by the RF pulse which contains a predetermined narrow band of frequencies. Only those regions, where the local resonant frequency falls within the range of frequencies determined by the gradient field and corresponding to that of the RF pulse, are excited. In this way, imaging can be restricted to a particular slice of the desired thickness.

The spin warp imaging technique is a particular type of 2DFT imaging, which was first proposed by Edelstein et al. [10] in 1980. Figure 1.1 shows a simplified imaging sequence for spin-warp imaging. The selective 90^o RF pulse, in combination with the gradient field $G_z(t)$ excites nuclear spins within the selected transverse plane. The magnetic gradient $G_x(t)$ causes an echo of the signal, and provides frequency encoding of the position in the x-direction, which is also known as the *frequency encoding direction*. The magnetic gradient $G_y(t)$ introduces a phase shift to the signal. This phase shift is a function of the y-direction, which is also known as the *phase encoding direction*.

MR signal data are collected onto a rectangular grid which represents the spatial frequency components of the imaged slice. This MR signal space is also known as the *k-space*. The magnitude of $G_y(t)$ selects the row of the k-space data that is generated by the MR imager. One row of the k-space is also known as a view. Sufficient views are obtained to fill the k-space and to enable the reconstruction of the MR image with sufficient resolution.

Another point of view of the action of the phase encoding gradient $G_y(t)$ is that in each imaging sequence N projections are collected onto the x-axis. The projections are different because spins at different locations are given varying amounts of phase shift by the different $G_y(t)$ values, hence the name *spin warp*. This situation differs from X-ray CT scanning, where only density (and not

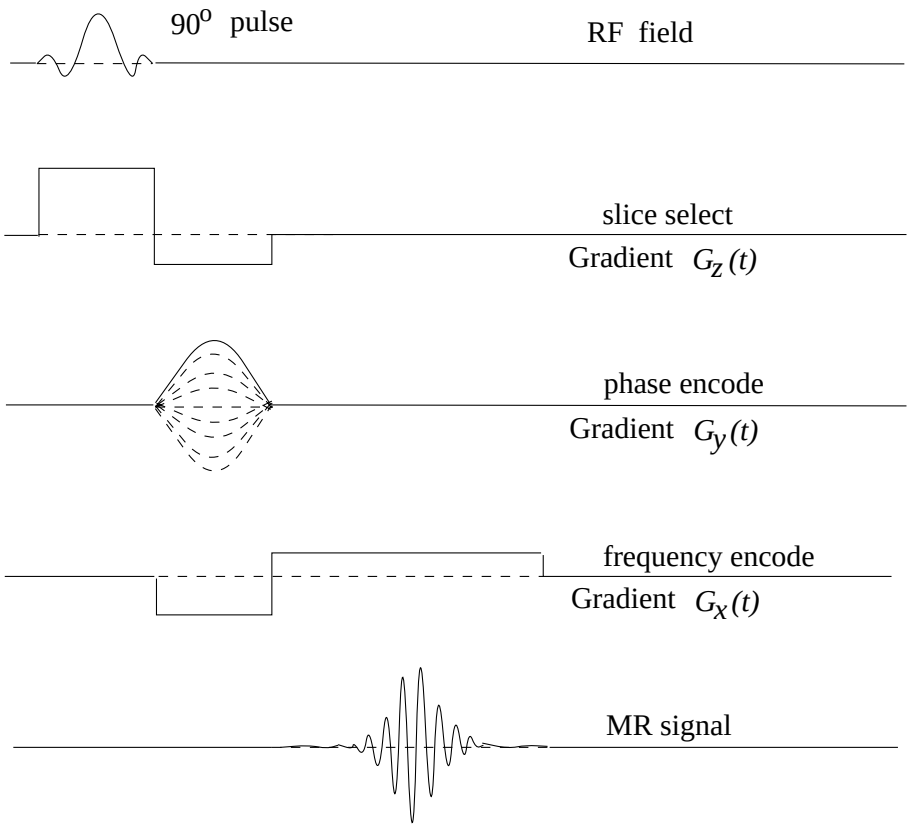

Figure 1.1: Spin warp gradient echo imaging pulse sequence.

phase) information is available, in which case the N projections must be taken from N different angles.

The MR data span the k-space coordinates k_x and k_y in the frequency encoded and phase encoded directions respectively. k_x and k_y can be expressed in terms of the applied magnetic field gradients, as shown below:

$$k_x = \gamma \int_0^{t'} G_x(t)dt \tag{1.2}$$

$$k_y = \gamma \int_0^{\tau} G_y(t)dt \tag{1.3}$$

where γ is the *gyromagnetic ratio*, t' is the scan time of a single view and τ is the duration of the applied phase encoding gradient. Given the transverse magnetization $m(x, y)$ of the object, the detected MR signal is proportional to

$$S(k_x, k_y) = \int\int m(x, y)e^{-i2\pi(xk_x+yk_y)}dxdy \tag{1.4}$$

provided that the relaxation effects are ignored [11]. From Equation 1.4, it is apparent that $S(k_x, k_y)$ is equivalent to the 2D Fourier transform of $m(x, y)$. Since the data are collected onto a uniformly spaced rectangular grid, the MR image is reconstructed by directly applying inverse 2D fast Fourier transform (FFT) on the k-space. The application of linearly varying gradients along x and y directions ensures uniform coverage of the entire frequency space of interest, as shown in Figure 1.2, where the high spatial frequency components are as equally represented as the low spatial frequencies. Hence, data acquisition of one view at a time contributes to high resolution and signal-to-noise ratio (SNR) in the reconstructed images, and avoids the need for data interpolation prior to image reconstruction. However, this technique involves relatively long image acquisition times and is susceptible to motion artifacts.

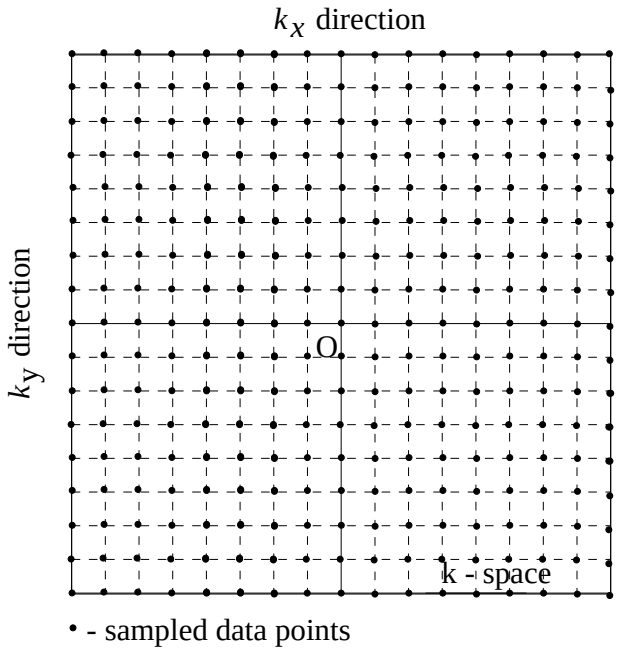

Figure 1.2: Locations of sampled data in the k-space for spin warp imaging.

It is also instructive to compare the spin warp technique with MR imaging by multiple-angle projections [3]. In the latter case, inhomogeneities in the magnetic field cause projected lines to be curved, thereby making reconstruction difficult. Therefore, either the curvature of the projected lines must be known exactly or

the magnetic field has to be very homogeneous. In the spin warp technique, all the projections are along the same direction, and therefore the inhomogeneity manifests itself only as a geometric distortion in the final image, rather than as a smearing of image information.

1.4 Motion Artifacts in MRI

In this section, the k-space concept described in the previous section, is extended to include the imaging of moving objects. The effects of motion on the k-space data can be divided into two categories [12], known as intra-view and inter-view artifacts. Intra-view artifacts are caused by motion during the data sampling period, or between the RF excitation and data acquisition. The intra-view artifacts were studied by Twieg et al.[13] for the specific cases of chemical shift and steady flow. Inter-view artifacts result from sufficiently slow motion that occurs between one excitation and the next. Therefore, movement of the imaged object during the acquisition of a single view is assumed to be negligible, if the object is supposed to undergo inter-view motion. This thesis is mainly concerned with in-plane rigid movements, which are often classified as inter-view motion.

In-plane rigid movements consist of translations and rotations in the transverse plane, and do not include shape changes of the imaged slice due to expansion or contamination resulting from out-of-plane structures. In-plane translations inflict phase errors in the k-space data [8], whereas in-plane rotations relocate the data onto a non-rectilinear grid [7]. The rotational motion affected data can also contain phase errors, if the centre of rotation does not coincide with the centre of the field of view (FOV).

For 2DFT spin warp imaging, the phase-encode gradient (i.e. $G_y(t)$) varies linearly with time (t). Therefore, according to Equation 1.3, k_y is also a linear

function of t. For inter-view motion, since movement of the imaged object during the acquisition of a single view is assumed to be negligible, the degradation of data in the k-space can be modelled as a function of k_y and not dependent on k_x. Therefore, movement of the imaged object slice $m(x, y)$ can be modelled by including the variable k_y as shown below,

$$S'(k_x, k_y) = \int\int \hat{m}(x, y, k_y) e^{-i2\pi(xk_x + yk_y)} dxdy \tag{1.5}$$

where $S'(k_x, k_y)$ is the motion affected signal, and $\hat{m}(x, y, k_y)$ is the moving object slice at the acquisition of the ${k_y}^{th}$ view. For in-plane rigid motion, both $\hat{m}(x, y, k_y)$ and $m(x, y)$ possess the same object shape and details, however the location and the orientation of $\hat{m}(x, y, k_y)$ within the FOV may change at each phase-encode k_y. Therefore, by computing the inverse Fourier transform of $S'(k_x, k_y)$, it is only possible to obtain a blurred and ghost version $m'(x, y)$ of the stationary slice $m(x, y)$.

2DFT MR images posses characteristic artifacts arising from inter-view motion. It has been observed that such motion leads to ghosting and blurring along the phase encoding direction, irrespective of the direction of the motion. In a detailed analysis of the effects of periodic translational motion, Wood and Henkelman [14] determined the point spread function for motion in the frequency encoding direction, phase encoding direction and slice selection direction. It was concluded that the separation between the ghosts is identical, regardless of the direction of motion, and can be more conveniently expressed in terms of image pixels y_G as shown below [15],

$$y_G = f_t \times T_R \times N \tag{1.6}$$

where f_t is the temporal frequency of the motion, T_R is the repetition time and N is the total number of phase encoding steps. The repetition time is the time interval between the acquisition of two consecutive phase encoded MR signals. Ghosts that would appear beyond the boundary of the FOV are wrapped back due to aliasing. An essentially similar result was observed by Twieg et al [16]

using $(k-t)$ *space formalism* to represent the dynamic behaviour of moving objects. This method was later extended by Xiang et al. [17] and Lauzon et al. [20] for motion artifact correction.

(a) *(b)* *(c)*

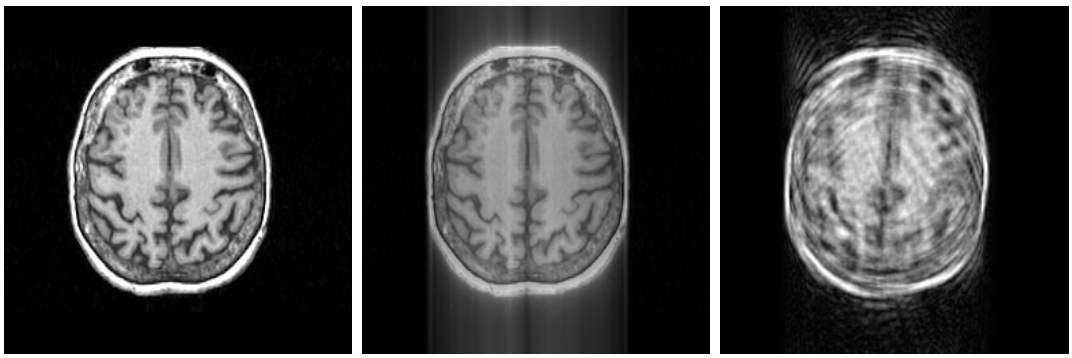

Figure 1.3: *(a)* artifact free image for comparison, *(b)* typical blurring along the phase encoded direction due to pure translational motion, and *(c)* typical ghost edges and blurring inflicted by pure rotational motion.

Although the effects of periodic translational motion have been studied in detail in the past, a similar analysis has not been provided for rotational motion. However, it is observed that rotational motion generally inflicts blurring and ghost edges with severe degradation of the original structure, compared to translational motion. Figure 1.3 illustrates typical motion artifacts inflicted by in-plane rigid motion of the imaged object. Figure 1.3(a) shows an artifact free image for comparison, whereas Figure 1.3(b) shows typical blurring along the phase encoded (i.e. vertical) direction due to pure translational motion within the imaged plane. Therefore, the regions to the left and right of the object are unaffected by motion artifacts. In this particular example, the maximum translation distance is limited to 5 pixels in vertical direction, and no horizontal translation is present. Figure 1.3(c) shows typical ghost edges and blurring inflicted by pure rotational motion within the imaged plane. The maximum angular span of the rotation is limited to $\pm 30^o$. It is clearly observable that rotational motion is capable of severe degradation of the original structure, compared to pure translational motion. Therefore, the objective of motion artifact correction algorithms is to

estimate an image which is closest to the artifact free image (i.e. Figure 1.3(a)), using the motion affected k-space data and a chosen motion profile.

1.5 Review of Real-Time Artifact Suppression Techniques

In this section, the methods for suppressing motion artifacts during the data acquisition are briefly reviewed. It is not intended to cover all the related literature, however, the purpose of this section is to describe the main techniques for real-time motion artifact suppression. Most of the techniques discussed in this section are already implemented in commercial MR imagers, and a few are still in the experimental stage.

Methods for suppressing motion artifacts in real-time can reduce either the physical motion of the patient, or the phase encoding errors induced in the captured MR signal. Some methods impose constraints on the selection of pulse sequences and timing parameters or limit the mobility of the patient. The more sophisticated methods do not require monitoring of patient motion and are compatible with the various standard imaging techniques in clinical use. However, many methods require monitoring of motion, which inevitably prolongs the time for patient preparation and data acquisition.

1.5.1 Physical Restraints

During the imaging time, the movement of the patient can be restricted using inflexible bands. These bands need not be wrapped uncomfortably tight for it to be effective. In clinical studies involving abdominal images, it was reported by Wood et al. [21] that the ghost artifact intensity in the image background was

35% greater in the image obtained without the restraining band. In pediatric MR examinations, de Varies et al. [6] used sedatives to restrain patient motion. Children aged from three months to 2 years were administered chloral hydrate via a naso-gastric tube, thus ensuring that a known amount of sedative was received. In rare instances a general anaesthetic was used [6], however, such restraint methods may increase both the risk to the patient and the operating cost. One of the simplest methods for suppressing respiratory motion artifacts requires the patient to refrain from breathing during the scan [22]. Such a restraint is possible in conjunction with rapid MR imaging [23]. In standard spin warp imaging, the T_R must be no longer than 250 msec, for the suspended respiration to be of clinically acceptable length [24].

Physical restraint is one of the few practical methods that can decrease the amount of blurring as well as ghost artifacts. However, restraining the patient only reduces artifacts induced by voluntary patient motion or respiratory motion, and does not suppress artifacts generated by other sources of motion, such as cardiac pulsations, pulsatile flow of blood and peristaltic bowel motion [21]. In addition, physical restraining requires active cooperation from the patient.

1.5.2 Rapid Imaging

Rapid imaging techniques are capable of acquiring the image data in a time short enough to freeze the motion. Fast low angle shot (FLASH) [25] imaging is a method of rapid MR imaging which has been demonstrated to provide abdominal images without artifacts due to respiratory and peristaltic motion. The sequence typically employs 15^o radio-frequency excitation pulses and acquires a free induction decay signal in the form of a gradient echo. Therefore, the measuring times of the FLASH images are reduced to about 2 seconds for a 128×128 pixel resolution. Another rapid data acquisition method is the fast spin echo

(FSE) [26] imaging technique. In FSE, several phase encodes are captured using a single excitation pulse within a single T_R. These are known as "single excitation images". Therefore, the number of T_R periods required to encode an MR image is less than that of conventional spin echo technique. FSE with reduced T_R (40 - 250 msec) and TE (less than 20 msec) allow acquisition of "single excitation images" in 5 - 30 seconds, within a single breath-hold [26].

Some of the other fast imaging sequences are SEPI (spiral-scan echo planar imaging) [27], QUEST (quick echo split) [28] and RARE (rapid acquisition with relaxation enhancement) [29]. Fast MR imaging times can also be achieved by reducing the number of phase encodes required for image reconstruction. One such technique was proposed by Weaver [30] using the Karhunen-Loeve expansion.

A major problem with fast imaging techniques is that the reconstructed images suffer from low contrast and low signal-to-noise ratio (SNR). Since the noise averages incoherently while the MR signal averages coherently, SNR will improve as the square root of the total data acquisition time, T_{acq} as given in Equation 1.7.

$$SNR \propto B \times \frac{\text{voxel volume}}{\text{noise volume}} \times \sqrt{T_{acq}} \tag{1.7}$$

This equation [31] describes the tradeoffs between field strength (B), resolution, acquisition time and SNR. Clearly, the development of imaging sequences that are faster and/or produce images at higher resolution incur an SNR penalty. Efforts to reduce system cost by reducing magnet field strength incur a similar penalty. Therefore, much effort has been invested in methods to reduce noise.

In addition, although respiratory and peristaltic motion artifacts as well as the voluntary patient motion artifacts are greatly reduced, transmitted cardiac and vascular pulsations may still degrade the images.

1.5.3 Signal Averaging

The motivation for signal averaging is not only to reduce random noise, but also to suppress systematic noise arising from ghosts, provided that the position of the imaged object is different during each acquisition. One method is the so called *parallel averaging* [32] where N signal acquisitions are repeated and averaged at each phase encoding step. *Serial averaging* described by Dixon et al. [33] is another method of repeated acquisition where N images are obtained sequentially and averaged later. The effect of averaging data sets is that ghosts interfere destructively and sometimes even cancel each-other completely [34]. Due to such interference, it is shown that the intensity of ghost artifacts is approximately inversely proportional to the number of data sets averaged. Although averaging reduces ghost intensity, it does not reduce blurring [34]. In addition, averaging causes prolonged imaging times. Therefore, it is more suited for short T_R sequences.

A method essentially similar to averaging has been proposed by Xiang et al [18], known as the *three-point ghost phase cancellation*. In this method, three images are acquired such that the phase of the ghost artifact in the second image is halfway between that of the first and the third images. Therefore, pixel-by-pixel averaging of the images produces effective ghost cancellation.

Another method has been developed by Madore et al. [35], in which several degrees of freedom are introduced in the averaging process, providing a choice of different weighting factors for different portions of the k-space. If a weighting factor is associated with each view, then freedom to eliminate motion artifacts is obtained. This process minimizes a quantity known as the gradient energy [19] over a region of interest in the imaging plane, and yields greater motion artifact suppression than that achieved with regular averaging.

1.5.4 Gating Techniques

Gated data acquisition has been used with great success for cardiac imaging [36], where the cardiac motion is considered to be quasi-periodic. In order to execute the imaging sequence in synchrony with the cardiac motion, cardiac gating is implemented, in which the data are acquired only within a limited window of some parameter detected by a motion monitoring device [17]. Three methods of acquiring the gating signal were investigated by Lanzer et al. [36], including an air-filled plethysmograph, a laser-Doppler capillary perfusion flow-meter and an electrocardiograph (ECG). Due to intrinsic delay, variations in the temporal relationship between the R-wave and plethysmographic and laser-Doppler signals were identified, which caused inconsistencies in the timing of imaging sequences. Therefore, the ECG method is preferred over other techniques, since the R-wave of the ECG gating device provides a reliable, intrinsic delay-free gating signal [36].

Gating techniques are also used to suppress respiratory motion artifacts. However, a gating signal analogous to the R-wave of the ECG in cardiac gating has not been found. Current respiratory monitoring systems are based on the movement of the abdomen. Often, the data are acquired during the end expiration, since that is usually the longest constant portion of the respiratory cycle [21]. In the *triggered mode* of respiratory gating, spin echo (SE) sequencing at a predetermined repetition time (T_R) is turned on and off at appropriate times in the respiratory cycle under the control of the gating signal. In this mode, tissues are subjected to long periods of relaxation when the RF pulse is switched off by the gating signal in response to respiratory motion. In the *spin-conditioned mode*, SE sequencing runs continuously at a preset T_R. The enabling gating signal is used to determine which spin echoes from the incoming stream will be kept for image reconstruction [37]. An extension of the same principle is used in retrospective cardiac gating [38].

Despite its effectiveness in motion artifact suppression, gating causes prolonged imaging times due to the necessity of motion detection. However, Ehman et al. [37] reported that the imaging time was prolonged to only $30\% - 50\%$ longer than non-gated sequences, by instituting a prescribed pattern of slow breathing.

A variation of gating that does not require continuous monitoring of the motion is known as *pseudo-gating* proposed by Haacke et al. [39]. In this technique, the total duration of the repeated acquisitions (N) equals the period of the cyclic motion. Pseudo-gating can reduce motion artifacts even when the motion is not exactly periodic, since it is at least an averaging of N acquisitions over most of the motion cycle. Therefore, Pseudo-gating may eliminate ghosting. However, the blurring is not diminished.

Another gating technique is called SLO-MOTION [40], and is capable of eliminating random or sporadic motion artifacts. In this technique, a small globe of a suitable liquid is attached near the head of the patient, using an appropriate band, cap or helmet. The globe is initially placed in a known position, at a far enough distance from the head to ensure that the head will not move into its initial position. A spatial localization technique is used to rapidly sample the signal from the globe before each acquisition of a phase encoded spin echo. If the patient's head moves and causes the signal from the localized volume to drop below a preset threshold, the image acquisition is paused and the localized volume is sampled until the signal returns. This technique is useful for suppressing gross movement artifacts induced by voluntary patient motion.

A gating technique which does not require motion detection hardware is known as the *diminishing variance algorithm* (DVA) proposed by Sachs et al. [41]. In this method, a complete set of preliminary data is obtained along with information about the relative position of the object at each view of the k-space data,

using a navigator echo [42]. When all the preliminary data are acquired, the position information is used to determine which data frames are most corrupted by motion. The phase encoded views are then re-acquired, starting with the most corrupted one. In this technique, the imaging time is prolonged not only by the re-acquisition of data but also by modifying the standard imaging sequence to accommodate the navigator echo.

1.5.5 Phase-encode Re-ordering

In order to eliminate the periodicity of the motion and hence remove motion induced ghost artifacts, the views are ordered to match the amplitude of the respiratory motion. This is accomplished by associating the smallest amplitude with the most negative phase encoding value, and larger amplitudes with larger phase-encoding values. Therefore, the amplitude of motion is forced to be a monotonically increasing function of time. This method is known as respiratory ordered phase encoding (ROPE) [43].

A variation of this approach is known as centrally ordered phase encoding (COPE) [44] where the data are ordered outwards from the central k-space region. It has the advantage of having symmetrically equivalent spatial frequency components and hence can be symmetrically filtered, where all the best data remain intact upon filtering.

The method of phase encode re-ordering only redistributes rather than eliminates the artifacts. Therefore, as the total energy of the harmonic components is always conserved in ROPE and COPE, the ghost artifacts are transformed into local blur and distortion [17].

1.5.6 Modified Pulse Sequences

There are generalized motion artifact suppression methods proposed in the literature, which require pulse sequence changes. The method proposed by Lauzon et. al. [20] uses two types of pulse sequences for data acquisition. The first sequence is at moderate speed which produces an image with high spatial and contrast resolution (Mohicon), whereas the second sequence is at high speed and produces an image with high temporal resolution (Falcon). The correction scheme combines the attributes of both the images to suppress motion effects. The major advantage of this method is its ability to handle any type of rigid body motion, including aperiodic motion. However, this method suffers from spatial resolution and signal to noise ratio limitations, and works better for small objects than for large objects. The method also requires longer data acquisition times than usual, due to the fact that it requires one high contrast image and several low contrast images in order to suppress motion effects.

The Short-Tau Inversion Recovery (STIR) [45] method uses an inversion recovery (IR) pulse sequence to make *fat* have little or no image intensity, hence eliminating ghost artifacts induced by the movement of fat. The limitations of this method are reduced SNR, low resolution and low contrast.

1.5.7 Gradient Moment Nulling

Although gating and phase encode reordering techniques are affective against inter-view motion artifacts, they fail to suppress concurrent intra-view motion artifacts that occur during the echo time (TE) or the data sampling period. During intra-view motion, incomplete rephasing of magnetization at the centre of the frequency encode gradient, and at the end of the slice select gradient, results in blurring artifacts [12]. A general solution, that refocuses the magnetization

in the imaging plane and zeroes the intra-view random phase terms, has been derived. The motion is modelled by the Taylor series, where the first three terms are position, velocity and acceleration. De-phasing attributed to motion is minimized by special design of gradient profiles necessary for nulling the moments of these three terms. The number of terms that can be nulled is limited by the shortest echo time required. This technique is formally known as the *gradient moment nulling* [12] [21] [46]. However, it is also known by many other names, including the motion artifact suppression technique (MAST) [12] [47], rephasing gradients technique [48] and compensating gradients technique [49]. Gradient moment nulling can be effectively used not only with spin warp imaging, but also with a host of other imaging methods including fast spin echo technique [50]. With no increase in imaging time, no additional equipment and little appreciable loss in spatial resolution as compared with short TE sequences, gradient moment nulling provides high quality T_2 weighted images [12], and can be combined with most of the other techniques that are capable of suppressing inter-view motion artifacts.

1.6 Review of Post-Processing Techniques

In this section, previously proposed post-processing techniques for motion artifact suppression are reviewed. Such techniques often assume an appropriate motion model to represent the motion. The model must be sufficiently complex to closely match the actual motion, but must also be sufficiently simple so that data degradation can be reversed efficiently. The motion models that have been proposed to date, as well as correction algorithms , have been limited to translation, rotation and dilation, which are over-simplified for widespread use in most practical applications. However, it is envisaged that by amalgamating

the available models, a more useful technique for motion artifact correction can be achieved.

Once a suitable motion model has been found, the next step is estimating the parameters for that model. These parameters may be known in advance, detected during the data acquisition or estimated from the corrupted data. If the motion is due to patient movement, it is unlikely that the motion can be known in advance. However, if the motion is induced by the mechanical movement of the MR imager gantry under operator control, the motion parameters can be predicted as reported by Korin et al. [51]. For rigid patient movements, the motion can be externally measured using sensors or transducers, with a similar procedure adopted in ROPE and COPE. Another method of extracting motion information is from the corrupted data. However, it is considered to be a difficult task, especially if the motion involves irregular movements.

The final step is the data correction. The difficulty level of this step is based on the chosen motion model and the availability of the motion parameters. The post-processing techniques discussed in the remainder of this section are divided into three main motion models, translation, rotation and dilation.

1.6.1 Translational Model

Early investigations on the post-processing algorithms for motion artifact suppression in MRI have mainly concentrated on the translational model for two reasons. Firstly, respiratory motion, which is a major contributory factor to motion artifacts in abdominal images, can be approximately modelled by in-plane translation of the diaphragm [52]. Secondly, and most importantly, gross translational motion is space invariant and the artifact effects can be modelled using linear sets of equations [52][53][54]. Therefore, the correction algorithms

for translational motion artifacts have evolved from in-plane 1D models to out-of-plane models.

Estimation of the translational motion profile has posed the greatest challenge in most artifact suppression techniques. Initial algorithms relied on interleaved navigator echo information acquired at the scan time to determine the motion profile [8] [55] [56] [57]. Navigator echo sequences are usually designed to provide the displacements along the frequency encoded direction (x-axis). The navigator echo is similar to an image echo, except that no phase encoding is applied, which implies that the navigator echo data will vary from view to view only if motion along the x-axis is present [8]. A limitation of such a linear projection echo [8] is that it can only provide information on translational motion along one axis. However, 2D translations can be determined by tracking the location of externally attached markers [58] with more conventional navigator echoes. This method requires patient preparation.

Assuming that only inter-view effects are modelled, the object displacement becomes a function of the phase-encode, for gross in-plane movements [8]. Analysis of the effects of such motion on the data is performed to show that the motion causes a phase error [8] [59], as shown below

$$S'(k_x, k_y) = e^{-i\psi(k_x,k_y)} S(k_x, k_y) \tag{1.8}$$

where $S'(k_x, k_y)$ is the motion corrupted signal and $\psi(k_x, k_y)$ is the phase error introduced by the translational motion. Therefore, if the displacements of the object are known via navigator echoes, the correction algorithm simply shifts the apparent distribution of spins in each successive spin echo to compensate for inter-view displacements by applying a phase shift that is opposite to the observed phase perturbation due to motion [8].

The navigator echo technique usually requires modifications to the standard pulse sequences, specialized hardware, and causes prolonged imaging time. Therefore,

Felmlee et al. [60] developed a method of using the phase encoded data to provide inter-view x-directional displacements. According to the Fourier transform shift property, the effect of x-directional motion is to shift the x-directional inverse Fourier transform of the acquired MR signal in the k-space. This property was first suggested by Felmlee et al. [60], formulated by Tang et al. [52] and later assigned the name *spectrum shift algorithm* by Zoroofi et al. [61]. However, a major limitation of this algorithm is that it could only estimate the x-directional displacements.

A method to estimate and correct for translational motion artifacts along the phase encoded direction (y-axis) is first proposed by Hedley et al. [62], using a modified Gerchberg-Saxton algorithm [63]. The *priori* information used in this algorithm is limited to the magnitude of the acquired data and a finite region of support for the imaged object. The algorithm successively transforms between the spatial and spatial frequency (k-space) domains, where the respective constraints are imposed. For y-directional translations, the phase error is a function of the phase-encode (k_y) only. Therefore, averaging of phase error values for each view has been used to increase the accuracy of the phase error estimations. Another feature of the algorithm is that it is usually possible to calculate the displacement of the object in each view, using the averaged phase error values. The limitations of this algorithm are that, it can only estimate the y-directional displacements, takes much computation time due to the iterative procedure and provides no guarantee of convergence [52].

To avoid an iterative procedure, a new constraint was proposed by Tang et al. [52], with which the motion component and the true image component are separated by a simple algebraic operation. However, the new constraint required was that the proton density distribution be symmetric along the y-direction, hence the phase of the Fourier spectrum along the y-direction is a linear function of y position. Therefore, the spatial frequency components that do not comply

with the linear function, can be separated as representing the motion. This method was later extended [64] to cover translations in both x and y directions by attaching the spectrum shift algorithm [60] to correct for x-directional translations. However, the method deteriorates when the presumption of symmetry of the proton density function is violated.

The modified Gerchberg-Saxton algorithm [63] proposed by Hedley et al.[62] was extended to cover translations in both x and y directions by an iterative algorithm of generalized projections [65]. This is a direct extension of the method of convex projections [66] [67] for non-convex sets. A relaxation parameter [65] has been introduced to increase the rate of convergence of the algorithm. Although this reduces much of the computation time, no guarantee of convergence is provided. This method was further developed by Hedley et al.[69] to have superior convergence properties, by incorporating additional *a priori* information based on a linear phase error relationship. The method involved a phase unwrapping algorithm [70] which required that the phase difference between adjacent data points be less than π radians, an assumption which is often violated if both x and y directional translations are simultaneously considered [61]. Therefore, Zoroofi et al. [61] combined the spectrum shift algorithm [60] and the method of generalized projections [65] to correct for in-plane 2D translations. In this method, the x-directional translations are corrected using the spectrum shift algorithm [60] prior to using phase unwrapping [70] in order to satisfy the required assumption.

Although in-plane 2D translational motion has been studied by several researchers, the effect of rigid out-of-plane translations has been analyzed by only a few. Mitsa et al. [71] proposed a model for the effect of periodic motion of a single slice in the direction of the slice selection axis as amplitude modulation of the raw data, with a motion kernel along the phase encoding direction in the k-space. This study claims that the motion can be detected in 1D projections of the raw data along the frequency encoding direction, which, in combination with

appropriate filtering, leads to the recovery of the motion kernel. It is assumed that the effect of the motion in the raw data domain can be modelled reasonably well with a finite number of terms of a modulated cosine function [71] and the 1D projection is highly peaked around the DC term with other frequencies at least two orders of magnitude below. The removal of the amplitude modulation is accomplished using a band-reject filter on the projection magnitude of the raw data. The filter is designed to preserve the phase information. The limitations of this algorithm are that it is restricted to periodic motion and due to the single slice approach, no compensation is made for the contamination of the signal from structures in adjacent slices. A marginal improvement to this method is proposed by Smith et al. [72], based on the TERA autoregressive moving average modelling algorithm. Smith et al. claim that bandstop filters used by Mitsa et al. [71] are based on removing the motion peaks by replacing them with a local average modified by the phase of the transformed projection data, hence introducing discontinuities into the filtered signal which lead to unnecessary artifacts. Therefore, filters developed from three term Blackman-Harris windows [73] are proposed instead, with a relatively narrow centre lobe and low side-lobes to reduce artifacts.

An alternative technique was proposed by Hedley and Yan [53] [54] using the same motion model proposed by Mitsa et al. [71]. However, the restriction of periodic motion was removed, hence the algorithm is applicable to random translational motion along the slice selection axis. According to this algorithm, the correction process is reduced to a magnitude retrieval problem, which can almost always be solved in theory, provided that the field of view (FOV) is larger than the imaged object. The correction method assumes that the region of support of the imaged object is known *a priori* and hence attempts to set the pixels outside this region to be zero, which forms a set of linear equations that can be solved using Gaussian elimination [75]. However, this method requires long computation

times. Furthermore, the equations are often singular or almost singular, which poses the problem of ill-conditioning in the system of equations. Therefore, singular value decomposition [75] was used to arrive at an approximate solution that adds to the computational demand that was already considered high. In order to avoid the problem of ill-posedness, this method is later improved by Hedley and Yan [74] using the iterative Gerchberg-Saxton algorithm [63].

A method of out-of-plane motion compensation in multislice spin-echo MRI was proposed by Reik et al. [76] [77], based on iterative projection onto convex sets (POCS) [67], provided that the motion kernel and the slice-selection profile are known. However, the exact shape of the slice-selection profile is usually not known in practice, but may be approximated from the RF pulse. The assumption that the motion kernel is known, may or may not be realistic depending on the application.

1.6.2 Rotational Model

Motion in the transverse plane not only involves translation, but can also involve rotation. Therefore, in order to model the general effects of rigid movements within the plane, rotational motion must be included. The effect of rotation of the imaged object on the k-space data is explained by the Fourier rotation theorem, which states that a rotation of the object by an arbitrary angle α will produce an equal rotation of its Fourier transform. This theorem assumes that the axis of rotation is at the centre of the field of view (FOV). In practice, the axis of rotation is often a point other than the centre. However, it is shown by Korin et al. [7] and later formulated by Zoroofi et al. [78], that the effect of rotation around an arbitrary point can be modelled by a combination of rotation about the centre of the FOV and translation along x and y directions. Therefore, the rotational model is always assisted by the translational model, and the

correction scheme involves appropriate phase correction due to translational motion. However, applying a rotation correction to k-space data, where there is a significant phase-roll across the image due to echo mis-centring, can cause artifacts in the reconstructed image, due to the abrupt phase transitions. Korin et al. [7] suggest that this can be avoided by tuning the pulse sequence beforehand to centre the echo in the data acquisition window.

The correction scheme for the rotational model is more difficult than for the translational model because the data are no longer located on a rectilinear grid. Therefore, the corrected data must be interpolated. Several interpolation techniques are available for this purpose, including bilinear, nearest neighbour and the sinc gridding algorithm [79]. Initially, data interpolation is performed in the k-space [7] [80]. However, Wood et al. [80] reported that interpolation in the spatial domain inflicted comparatively fewer artifacts in the final image. The main contributing factor for this observation is that the transformed spatial domain data exhibit lower variance than the k-space data, hence resulting in lower interpolation errors.

Re-location of data subjected to rotational motion is shown to create sectors of k-space with data voids and data overlaps [7] as shown in Figure 1.4. The shaded regions in Figure 1.4 indicate the available k-space data. Data overlap is illustrated by overlapping shaded regions. The clear regions indicate data voids in the k-space. For small angle rotations, there is only a marginal effect from these regions. Therefore, in many such cases, the data voids are zero filled, and the overlap data are averaged [7] [80]. The algorithms proposed by Zoroofi et al. [78] and Atkinson et al. [81] pay no attention to these k-space sectors, hence producing visible artifacts due to the correction scheme. In this thesis, a new method of managing overlap data via weighted averaging will be presented together with an iterative data estimation scheme for data void regions.

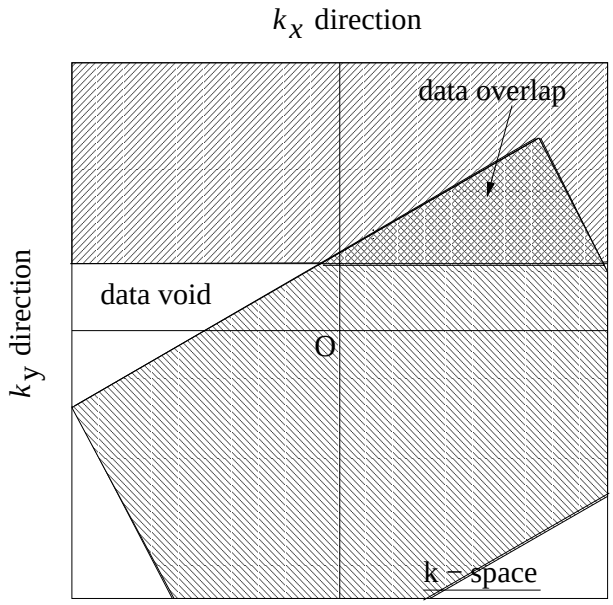

Figure 1.4: Illustration of data overlap and void regions. Data overlap is illustrated by overlapping shaded regions. The clear regions indicate data voids in the k-space.

Rotational motion parameter estimation has always been a cumbersome task, specially due to the concurrent translational motion. Korin et al. [7] used spatial-frequency tuned markers for tracking the rotation angles. The principal spatial frequency of the marker can be varied by altering the spacing of its bright and dark lines and by the appropriate phase encoding of the navigator echo [7]. However, this method requires patient preparation and alterations to standard pulse sequences. Another method of motion detection was proposed by Wood et al. [80], using the paired *t-test* and k-space data. This method identified the phase encode where the subject movement occurred. Movements are expected to make the *t-test* P value small for at least two consecutive phase encode (i.e. k_y) values. The position of the subject was suspected to have changed at a particular k_y if the P values are below a threshold value of 0.01 [80]. The limitation of this method is that it only serves to identify the motion affected views, and does not provide the parameters associated with the motion. The correction scheme is therefore entirely dependent on the uncorrupted data and the conjugate symmetry of the k-space.

Following the work of Wood et al. [80], Zoroofi et al. [78] proposed the *bilinear superposition algorithm* (BSA) for in-plane rigid motion artifact cancellation, where the rotational motion correction and interpolation are performed in the spatial domain. However, Zoroofi et al. [78] also proposed an algorithm for motion parameter estimation using the minimum energy method. This method is similar to the method proposed by Hedley et al. [62], in which the energy outside the region of support is minimized. The energy minimization is performed using an exhaustive search of every possible combination of motion parameters representing in-plane translations and rotations. Therefore, as has been reported by Zoroofi et al. [78] estimation of ten unknown rotation angles with less than one degree error took approximately 4.5 hours, indicating the computational burden associated with the motion parameter estimation scheme. In this thesis, a fast and efficient method of object rotation function estimation will be presented, applicable to in-plane rigid motion involving rotations and concurrent translations.

An alternative correction scheme for in-plane rigid motion was proposed by Atkinson et al. [81], using an *entropy focus criterion* for re-focusing the image. This technique can determine unknown patient motion or use knowledge of the motion from other measures as a starting estimate. The motion estimate is used to compensate the acquired data, and is iteratively refined using the image entropy. The entropy focus criterion favours alterations to data that tend to increase the number of dark pixels. Motion often has the opposite effect, creating ghosts and blurring in the image regions that would otherwise be dark. Therefore, one might expect entropy minimization to aid the search for a motion-free image. However, it is by no means clear that this will always be the case. An exhaustive search on all possible combinations of angles and phase encodes has been performed in which the rotated portion of the k-space is re-gridded, followed by a correction for translational motion. However, in the midst of concurrent

translational motion, this exhaustive search may prove to be impractical due to high computational and time demands.

A method of in-plane rigid motion estimation and correction was also proposed by Guidry et al. [82], using active contours. In this method, the k-space is divided into several sub-bands, such that each sub-band is acquired in a small fraction of the full imaging time. These sub-bands are supposed to create invariant tissue feature maps called sub-band images. Using active contours to find the outer boundary of the sub-band images, the relative motion is analyzed to determine the motion parameters. The estimated motion parameters are used to correct each sub-band image, thus correcting the k-space data. The major limitation of this method is the assumption that each sub-band is motion free. Furthermore, no attempt has been made to detect the changes of orientation and location of the imaged object, using the k-space data. Although it was claimed that each sub-band produces a discernible shape of the object, usually high frequency spatial components provide much less recognizable details for an active contour model to extract the true shape of the object. Therefore, this method can be used only in exceptional circumstances where the assumption of motion-free sub-bands is sufficiently satisfied at low spatial frequency views.

In-plane rigid motion artifact reduction techniques are proposed not only for conventional spin-warp imaging but also for echo planar type projection reconstruction. Van de Walle et al. [83] proposed a method of isolating the motion-distorted projections using interleaved projection-angle order during the acquisition of the MR data. However, the correction scheme only uses the non-distorted projections, hence no motion parameter estimation was performed. For the special case when there is no concurrent translation, the rotation angles are estimated using the root mean squared distance (RMSD) between the images produced by the distorted and non-distorted projections.

1.6.3 Dilational Model

The effect of linear dilational motion on the k-space data have been studied by Atalar and Onural [84]. It has been observed that the linear dilational model leads to both phase error and re-location of the data points on a non-rectilinear grid in the k-space. Therefore, the correction scheme involves compensation of the phase error and interpolation of the data onto a rectangular grid. However, it involves calculation of an inverse of the transformation matrix T, which consists of elements of the form of a sinc gridding function [79]. The matrix T may be almost singular if some non-uniform samples are too close to each other or if there is no sample around a uniform sample position in k-space. In order to overcome this potential problem, singular value decomposition (SVD) has been used and the uniform samples are calculated by multiplying the pseudo-inverse of T with the non-uniform sample vector. The linear and third order Lagrange and cubic spline interpolation methods can also be used, if fast image reconstruction is required. However, the final image quality may be degraded due to interpolation errors. These artifact correction methods impose high computational complexity and long computation times. For example, Atalar and Onural [84] have reported that for a 256×256 image, the average reconstruction time is as high as 56 hours using the SVD algorithm.

A fast iterative scheme for correcting in-plane motion artifacts has been proposed by Steagall et al. [85], using the POCS procedure. The constraints for this scheme have been chosen according to the motion model representing translational, rotational or dilational motion, with only a few degrees of freedom. Further constraints are also applied when it is possible to ensure reasonably rapid convergence. Specifically, the operation is defined via the equation given below:

$$I_k = I_{k-1} + B(L_p I_0 - M L_0 I_{k-1}) \tag{1.9}$$

where I_0 is the ghosted image, L_0 and L_p are constraint operators, M is the motion model operator, B is a relaxation parameter, and I_k is the k^{th} iterated image. The method has been applied to translational and dilational motion, with results indicating a reduction in ghosting to less than 5%. However, the method is sensitive to the choice of the motion model and the parameters, with significant *a priori* knowledge required for POCS formulation.

An alternative method proposed by Tseng et al. [86] uses an iterative correction technique via inverse problem solving and assumes that the corrupted image is close enough to the original image for the convergence of the algorithm. The motion profile of the imaged object is assumed to be known *a priori* via sensors, which limits the applicability of this method in practical clinical scenarios. This technique is also vulnerable to noise and errors embedded in motion measurements.

The post-processing methods discussed in this section are mainly aimed at reducing motion artifacts from spin-warp MR images. However, such techniques are also proposed for diffusion weighted MR images, using navigator echoes [87] and reference phase maps [88]. Methods of artifact reduction in projection reconstruction, applicable to both MRI and CT, have been proposed by Crawford et al. [89] and Gai et al. [90]. A discussion on post-processing techniques mainly involving translational motion model can also be found in the review paper by Hedley and Yan [91].

1.7 Review of Projection onto Convex Sets

In this section, projection onto convex sets (POCS) is reviewed, in order to establish the theoretical background for the algorithms presented in this thesis. POCS is a powerful mathematical tool in reconstructing or restoring a signal or

image from incomplete information. It has been used in medical image processing for over a decade [92]. Popularity of POCS in medical imaging can be attributed to the fact that, often the acquired signal from modalities such as CT and MRI are essentially in the spatial frequency domain, whereas the required image solution is in the spatial domain. Therefore, information required for the reconstruction of the image is found in two different domains. In order to alternate between the domains, Fourier, Radon or Wavelet transformations are used. However, both domains usually provide partial information, making POCS an ideal candidate for solving the reconstruction problem.

The advent of POCS to the medical imaging arena was sparked by the limited-view CT applications [92] and the well known missing cone problem [93]. However, POCS has also found applications in MRI, for reconstructing images from motion affected signals, where partial degradation of the signal in magnitude, or phase, is incurred due to translational motion. Rotational motion artifact correction can also benefit from using POCS to estimate the data missing in the data-void regions. This is the premise which led to the involvement of POCS in the algorithms used in this thesis.

1.7.1 History of POCS

Prior to the advent of CT, material scientists concerned with x-ray diffraction, astronomers concerned with star-size measurements, and optical physicists concerned with superresolution, showed interest in restoration algorithms that can effectively use available partial information. As a result, POCS was first reported by Russian scentists Bregman [94] and Gubin et al. [95] in the 1960's. However, for the specific case of intersecting linear varieties (i.e. hyperplanes), POCS is identified to be similar to Von Neumann's alternating projection theorem [96], which dates back to 1937.

A class of iterative algorithms, that is essentially a subset of POCS, were independently developed by Gerchberg and Saxton in 1971 [63]. The original algorithm [63] was used in electron microscopy for determining the complete wave function from intensity measurements in the imaging and diffraction planes. This algorithm was also used by Gerchberg to increase the resolution of diffraction limited images of a finite object (super-resolution) [97]. In 1975, Papoulis published a variation of the algorithm for increasing the spectral resolution determined from a segment of a band-limited signal [98]. This algorithm was consequently known as the Gerchberg-Papoulis algorithm, or error-reduction algorithm (see Figure 1.5).

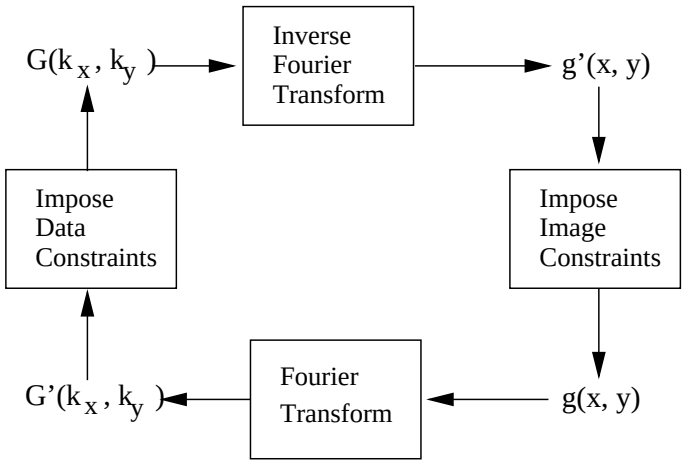

Figure 1.5: Gerchberg-Papoulis algorithm.

The Gerchberg-Papoulis algorithm alternates between the image and signal domains, imposing the respective constraints. This is illustrated in Figure 1.5. Image reconstruction problems are always started with the acquired data $S'(k_x, k_y)$ in the signal domain, by setting

$$G_1(k_x, k_y) = S'(k_x, k_y) \tag{1.10}$$

The algorithm starts each iteration n, by transforming the data to the image

domain,

$$g'_n(x,y) = \mathcal{F}^{-1}\{G_n(k_x,k_y)\} \tag{1.11}$$

where $\mathcal{F}^{-1}$ represents the inverse Fourier transform. The image domain constraint for most reconstruction problems is that there is a known, finite region of interest (ROI) for the imaged object. The constraint is simply imposed by setting pixels outside the region of support to zero.

$$g_n(x,y) = \begin{cases} g'_n(x,y) & \text{if } (x,y) \in \text{ROI} \\ 0 & \text{otherwise} \end{cases} \tag{1.12}$$

The image is then transformed back to the data domain,

$$G'_{n+1}(k_x,k_y) = \mathcal{F}\{g_n(x,y)\} \tag{1.13}$$

where the data constraint is imposed. These steps are iterated until a satisfactory image, $g_n(x,y)$, has been found.

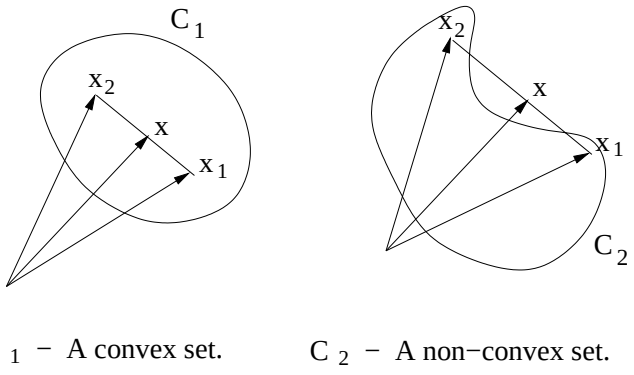

Figure 1.6: *(a)* An example of a convex set C_1, and *(b)* an example of a non-convex set C_2, where a line connecting any two points x_1 and x_2 in the set does not reside wholly within the set C_2.

The major limitation of the Gerchberg-Papoulis algorithm was that it used only two constraints, hence requiring a large number of iterations for convergence to an acceptable solution. A generalization of Gerchberg-Papoulis type linear iterative algorithms was performed by Youla and Webb [67], in which the important questions regarding the ill-posedness or noise-sensitivity was also addressed for

image reconstruction procedures. The reconstruction theory discussed by Youla and Webb [67] has a significant advantage over the Gerchberg-Papoulis and related algorithms, in that it enables a large number of *a priori* known constraints to be incorporated in the algorithm through the mechanism of projection on to convex sets. However, it was observed that, not all constraints can be formulated as convex sets [99] [100]. In a convex set, a line connecting any two points in the set resides wholly within the set. This fundamental property of convexity is illustrated in the set C_1 of Figure 1.6(a), with x denoting any point on the line segment between x_1 and x_2, where $x = \mu x_1 + (1 - \mu)x_2$ for $0 \leq \mu \leq 1$. Non-convex sets (i.e. C_2) do not obey this property for every pair of points in the set as shown in Figure 1.6(b). There are many sets of non-convex nature encountered in physical and engineering problems. A few examples of constraints that result in non-convex sets are listed below [68]:

- Reconstruction of a signal $y(t)$, using the measure of the instantaneous power $|y(t)|^2$ of the signal.
- Reconstruction of a signal $y(t)$, given the magnitude of the Fourier transform $|Y(\omega)|$.
- Reconstruction of a signal $y(t)$, given the upper and lower bounds of the magnitude of the Fourier transform $Y(\omega)$.
- Reconstruction of a signal $y(t)$, knowing that the point-spread function $h(t)$ being linear and space invariant, when performing blind convolution with an unknown original source $x(t)$.
- Reconstruction of a signal $y(t)$, given that $y(t)$ take on binary values ± 1.

Therefore, generalized projections algorithm was developed by Levi [99] as a generalization of the POCS algorithm to non-convex sets. However, unlike POCS, this does not guarantee that the algorithm will converge, not even weakly.

Projection algorithms (including POCS and generalized projections) form one of the more general image reconstruction and restoration methods available at present. Due to the general formulation of the projection methods, inclusion or disposal of a particular constraint can be performed with ease, avoiding reformulation of the problem in the form of a new optimizing functional. POCS is still being developed into a more versatile algorithm, with the use of fuzzy convex sets [102] [103], approximate projections [104] and adaptive constraints [105].

Comparison studies of POCS with the Gerchberg-Papoulis algorithm were conducted by Sezan and Stark, in order to establish the robustness of POCS for noiseless [100] and noisy circumstances [101].

1.7.2 Theory of POCS

The conceptual basis of POCS is that the available constraints can be represented as convex sets in a Hilbert space, $\mathcal{H}$ [67]. For image reconstruction, the Hilbert space is the set of all 2D square integrable functions with the usual inner product

$$(f, g) = \int\int f(x, y)g^*(x, y)dxdy \tag{1.14}$$

where g^* represents the complex conjugate. Therefore, by definition,

$$\|f\| = \sqrt[+]{(f, f)} \geq 0 \tag{1.15}$$

is the "length" or norm of f, and the sequence $\{f_k\}$ is said to converge to f if

$$\lim_{k\to\infty} \|f_k - f\| = 0. \tag{1.16}$$

However, the sequence $\{f_k\}$ is said to converge *weakly* to f if

$$\lim_{k\to\infty} (f_k, g) = (f, g) \tag{1.17}$$

for every $g \in \mathcal{H}$. It should be noted that each element of the Hilbert space is an image.

A set $\mathcal{C}$ is known to be convex if and only if $f \in \mathcal{C}$ and $g \in \mathcal{C}$ results in

$$[\alpha f + (1-\alpha)g] \in \mathcal{C} \quad \text{for all} \quad 0 \leq \alpha \leq 1$$

In the class of image reconstruction problems, every known property of the original f restricts it to lie in a corresponding closed convex subset of $\mathcal{H}$. In general, therefore, m such properties will generate m well-defined closed convex sets $\mathcal{C}_i$, where $i = 1, 2, ..., m$, and therefore,

$$f \in \mathcal{C}_0 = \bigcap_{i=1}^{m} \mathcal{C}_i. \tag{1.18}$$

The intersection, $\mathcal{C}_0$, is also closed convex and contains f. Consequently, irrespective of whether $\mathcal{C}_0$ contains elements other than f, the problem of reconstructing f from its m properties is included in that of finding *at least one* point belonging to $\mathcal{C}_0$.

If the projection operator P_i, projecting onto its respective convex set $\mathcal{C}_i$ is effectively realizable for $i = 1, 2, ..., m$, the problem is recursively soluble. Clearly, if $f \in \mathcal{C}_i$, then $P_i f = f$, therefore every element of $\mathcal{C}_i$ is a *fixed* point for P_i. This rule, which assigns to every $x \in \mathcal{H}$ its nearest neighbour $f \in \mathcal{C}_i$, defines the nonlinear projection operator $P_i : \mathcal{H} \to \mathcal{C}_i$, unambiguously, by means of the following minimality criterion.

$$\|x - P_i x\| = \min_{f \in \mathcal{C}_i} \|x - f\| \tag{1.19}$$

Consider the composition operator

$$T = P_m P_{m-1} ... P_1. \tag{1.20}$$

Although, in general, T is not the projection operator onto $\mathcal{C}_0$, every point of $\mathcal{C}_0$ is clearly a fixed point of every P_i and therefore of T. Conversely, if $\mathcal{C}_0$ is non-empty, every fixed point of T is a point of $\mathcal{C}_0$. Hence, an iterative scheme for generation of the fixed points of T is provided by the standard recursion

$$x_n = T^n x, \qquad n = 0 \to \infty \tag{1.21}$$

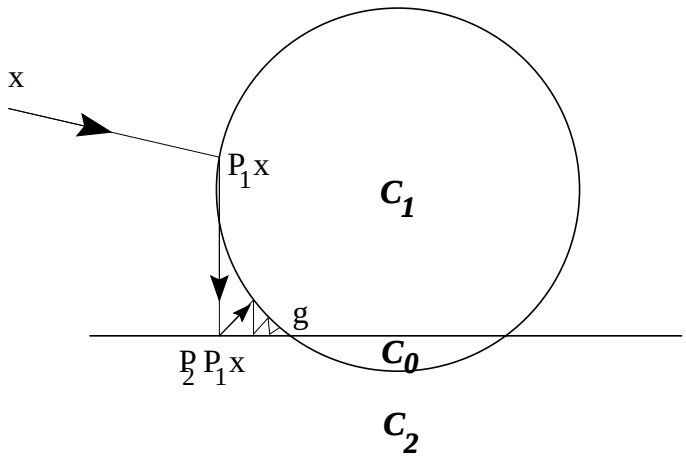

Figure 1.7: An illustration of the POCS iterations for two convex sets $\mathcal{C}_1$ and $\mathcal{C}_2$ with non-empty intersection $\mathcal{C}_0$. The iterations start at the initial point x. The first projection is onto $\mathcal{C}_1$. The projected point P_1x on $\mathcal{C}_1$ is the closest point to x. The second projection is onto $\mathcal{C}_2$. The projected point P_2P_1x on $\mathcal{C}_2$ is the closest point to P_1x. Similarly, repeated projections onto $\mathcal{C}_1$ and $\mathcal{C}_2$ ensure that the solution converges to point g on $\mathcal{C}_0$.

where x is an arbitrary starting point. This process is graphically illustrated in Figure 1.7. Evidently, $T^n x$ weakly converges to $g \in \mathcal{C}_0$. In order to expedite the convergence, over-relaxation can be used for extending the projections beyond the boundaries of the convex sets. The relaxed projection operators are defined as

$$T_i = I + \lambda_i(P_i - I) \tag{1.22}$$

where I is the identity operator and λ_i are appropriately pre-selected relaxation parameters. The fixed points of T_i are the same as P_i, provided that λ_i is non-zero.

The composite operator T is now given by

$$T = T_m T_{m-1} ... T_1. \tag{1.23}$$

According to the analysis by Youla [67], the sequence $\{T^n x\}$ converges weakly to a point of $\mathcal{C}_0$, for every $x \in \mathcal{H}$, provided that the relaxation parameters are selected within the interval $0 < \lambda_i < 2$. However, if $\mathcal{C}_i$ is a linear subspace, then

the optimal value is $\lambda_i = 1$, which effectively transforms T_i back to the unrelaxed operator P_i.

The slow convergence behaviour of the algorithm is postulated to be due to *traps* and *tunnels* in the Hilbert space [106]. A trap is an image which is a fixed point of T_1T_2, so that it satisfies

$$g_{n+1} = T_1T_2g_n, \tag{1.24}$$

but is not a fixed point of both T_1 and T_2 individually. This can only occur when one of the constraints is a non-convex set, as is the case for generalized projections, however, it cannot occur in POCS. An image, g_n is in a tunnel if Equation 1.24 is almost satisfied, resulting in very slow convergence. It is possible for tunnels to occur in POCS, and this is illustrated in Figure 1.8 for two convex constraints. Although it can be inferred that the algorithm is in a tunnel by noting a very small change in the image error between successive iterations, it is not analytically known how to use the available information to move the iterations to another region of the Hilbert space where convergence will be more rapid.

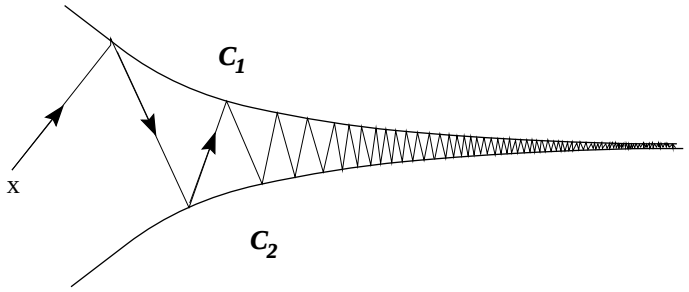

Figure 1.8: An illustration of a *tunnel* in Hilbert space, formed by two convex sets $\mathcal{C}_1$ and $\mathcal{C}_2$.

Another potential problem with POCS is the inconsistency of the available *a priori* constraints. Decades of research has established the fact that the quality of the solution largely depends on *a priori* knowledge of the original form of

the image and the signal degradation model. A major problem is utilizing all available knowledge optimally, since not all *a priori* knowledge is equally reliable. A POCS algorithm which bases its projections on unreliable information usually generates inferior or unacceptable results. In POCS, unreliable or inaccurate *a priori* information leads to inconsistent convex sets where the intersection $\mathcal{C}_0$ is essentially empty. If there are only two sets, the convergence is to the cycle between the closest points of the sets in the mean square sense. If there are more than two sets, POCS converge to greedy limit cycles that are dependent on the ordering of the projections [107].

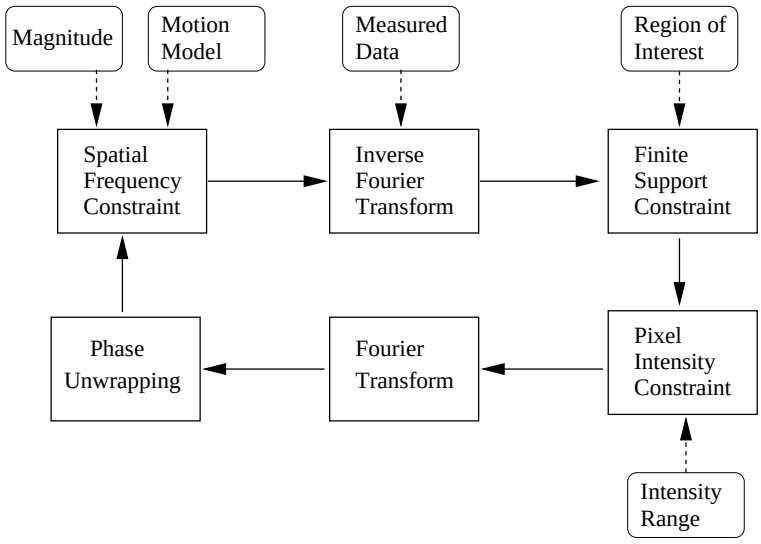

Figure 1.9: Block diagram on POCS iterations for translational motion artifact correction in MRI.

Reconstruction of an image from partially defined spatial frequency information is frequently encountered in CT and MRI. In general, the available spatial frequency information is derived from the data generated by the MR imager. For example, translational motion artifact correction in MRI involve restoration of corrupted phase information using the available uncorrupted magnitude information in the spatial frequency domain, and prior knowledge on the motion model. Restoration of motion affected phase also involve phase unwrapping [142], for the estimation of true phase error. Useful spatial domain constraints include the finite support

of the imaged object as well as pixel intensity constraints. A block diagram of the POCS iterations for translational motion artifact correction is shown in Figure 1.9.

The measured data are usually subjected to measurement errors, interpolation or computational errors and electronic noise. Each of the available spatial frequency values can lead to a separate convex set of the form of a linear variety in the Hilbert space. A linear variety is defined as a translation of a subspace by a fixed vector [108]. This fixed vector is computed from the known spatial frequency value, and defines the location of a particular linear variety in the Hilbert space.

If the reliability of each computed spatial frequency value is known, each resultant linear variety can be fuzzified with a membership function determined by the reliability operator. Depending on its membership, each vector can be given freedom to move , so that a set of fuzzy linear varieties can be formed which do not contradict the spatial constraints defined in *a priori* information. Therefore, proper fuzzification of linear varieties can transform an inconsistent set of constraints to a group of convex sets that are *close* enough to produce a useful reconstructed image. A detailed discussion on the fuzzification of convex sets and the formulation of fuzzy POCS is included in Chapter 6.

1.8 Review of Fuzzy Sets

Crisp sets are the sets most readers are familiar with. In a crisp set, an element is either a member of the set or not. For example, a car belongs in the class of vehicles whereas a house does not.

Fuzzy sets, on the other hand, allow elements to be partially in a set. Each element is given a degree of membership in the set. This membership value can range from 0 (i.e. not an element of the set) to 1 (i.e. a full member of

the set). It is clear that if one only allowed the extreme membership values of 0 and 1, the resultant set would actually be equivalent to a crisp set. A membership function is the relationship between the values of an element and its degree of membership in a set. An example of a membership function is shown in Figure 1.10. In this example, the set of numbers defined as "about 10" are depicted using fuzzy membership. The formulation of a fuzzy membership function is known as "fuzzification". There can be many different formulations of fuzzy membership functions for the representation of set of numbers "about 10", as given in Equations 1.25 to 1.28. Therefore, it should be clearly noted that any parameters associated with such formulations are subjective rather than precise.

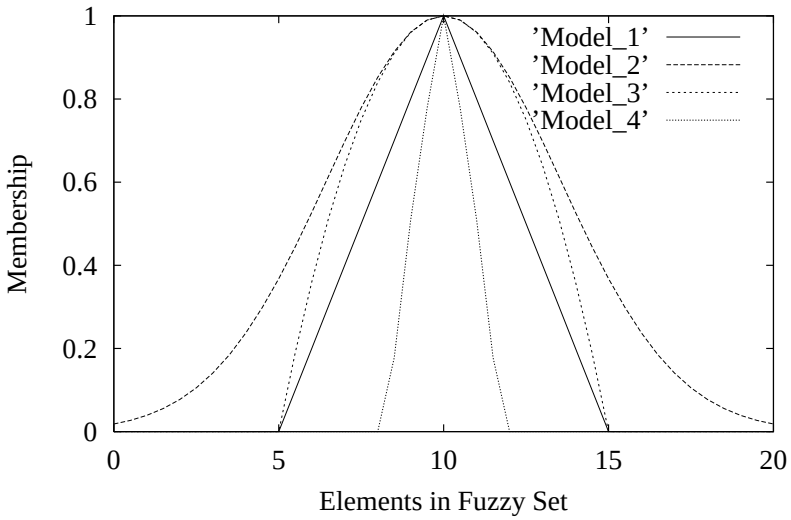

Figure 1.10: Fuzzy membership functions (i.e. fuzzy models) representing the set of numbers defined as "about 10".

Fuzzy model 1:

$$\mu(x) = 1 - \left| \frac{x - 10}{5} \right| \tag{1.25}$$

Fuzzy model 2:

$$\mu(x) = \exp\left(-\left|\frac{x-10}{5}\right|^2\right) \tag{1.26}$$

Fuzzy model 3:

$$\mu(x) = 1 - \left|\frac{x-10}{5}\right|^2 \tag{1.27}$$

Fuzzy model 4:

$$\mu(x) = 1 - \left|1 - \exp\left(\left|\frac{x-10}{5}\right|^2\right)\right| \tag{1.28}$$

1.8.1 Fuzzy combinations

In making a fuzzy rule, the concepts of "intersection" and "union" are often used. The sections below describe the most common definitions of these "fuzzy combination" operators. Fuzzy combinations are also referred to as "t-norms" (intersection) and "t-co-norms" (union) [131].

The fuzzy intersection is written as

$$T[\mu_A(x), \mu_B(x)] = A \bigwedge B \tag{1.29}$$

where $\mu_A(x)$ is read as "the membership in class A" and $\mu_B(x)$ is read as "the membership in class B". There are many methods to compute the fuzzy intersection. The two most common methods are:

1. Zadeh technique [131] - named after the inventor of fuzzy set theory, simply computes the intersection by taking minimum (i.e. $\min(\mu_A(x), \mu_B(x))$) of the two (or more) membership values. This is the most common definition of the fuzzy "intersection".

2. Product - (i.e. $\mu_A(x) \times \mu_A(x)$). This techniques computes the fuzzy "intersection" by multiplying the two membership values.

Both techniques have the following two properties:

$$T[0,0] = T[\mu_A(x), 0] = T[0, \mu_B(x)] = 0 \tag{1.30}$$

$$T[\mu_A(x), 1] = T[1, \mu_A(x)] = \mu_A(x) \tag{1.31}$$

The fuzzy union is written as

$$S[\mu_A(x), \mu_B(x)] = A \bigvee B \tag{1.32}$$

Similar to the fuzzy intersection, there are two techniques for computing the fuzzy union:

1. Zadeh technique [131] computes the fuzzy union by taking the maximum (i.e. $\max(\mu_A(x), \mu_B(x))$) of the two (or more) membership values. This is the most common method of computing the fuzzy "union".

2. Difference - (i.e. $\mu_A(x) + \mu_B(x) - \mu_A(x) \times \mu_B(x)$). This technique uses the difference between the sum of the two (or more) membership values and the product of the membership values.

Both techniques have the following two properties:

$$S[1,1] = S[\mu_A(x), 1] = S[1, \mu_B(x)] = 1 \tag{1.33}$$

$$S[\mu_A(x), 0] = S[0, \mu_A(x)] = \mu_A(x) \tag{1.34}$$

1.8.2 Convexity and α-cut of a fuzzy distribution

An α-cut of the membership function $\mu_A(x)$ is the set of all x such that $\mu_A(x)$ is greater than or equal to α_A. A *strong α-cut* is the set of all x such that $\mu_A(x)$ is *strictly greater* than α_A. Mathematically, an α-cut is defined by Equation 1.35 and a strong α-cut is defined by Equation 1.36

$$A^\alpha = \{x | \mu_A(x) \geq \alpha_A\} \tag{1.35}$$

$$A^{\alpha-strong} = \{x|\mu_A(x) > \alpha_A\} \tag{1.36}$$

The α-cut (or the strong α-cut) of a fuzzy set A is the crips set A^{α} (or the crisp set $A^{\alpha-strong}$) that contains all the elements of the universal set X, whose membership grades in A are greater than or equal to (or only greater than) the specified value of α_A. A^{α} and $A^{\alpha-strong}$ are crisp sets because a particular value x either is or is not in the set. There is no partial membership.

The "support" of a fuzzy set A is all $x \in X$ that have non-zero membership in A, which is identical to the crisp set $A^{\alpha-strong=0}$. $A^{\alpha=1}$ represents all $x \in X$ that are completely in A (i.e. no other sets), and is called the "core" of A. The "height" of a fuzzy set A, $h(A)$, is simply $\max(\mu_A(x))$ for all $x \in X$, or all x in the support, $A^{\alpha-strong=0}$. A fuzzy set is "normal" when $h(A) = 1$ and "subnormal" when $h(A) < 1$.

A fuzzy set A on R (i.e. a membership function that depends upon one variable R) is convex if and only if

$$\mu_A(\lambda x_1 + (1-\lambda)x_2) >= \min\{\mu_A(x1), \mu_A(x2)\} \tag{1.37}$$

for all x_1 and x_2 in R and all λ in the range [0,1].

As λ varies from 0 to 1, $x = \lambda x_1 + (1-\lambda)x_2$ varies from x_2 to x_1. This definition states that a fuzzy set is convex if any point in between two other points resides in the alpha-cut that is largest. This would be true, for example, if the membership function $\mu_A(x)$ were a Gaussian function. Note that this function is not convex in itself since points on the line joining two interior points can be outside the curve (see Section 1.7.1), but the fuzzy set defined by this membership function would be convex by the preceding definition. If this membership function had a dip at the maximum (to form a 2-hilled function), then the fuzzy set would not be convex since x_1 and x_2 could be on either side of the central minimum.

1.9 Test Images

The test images used in the simulation experiments throughout this thesis are introduced in this section. These images are used to verify and quantify the performance of the various stages of the proposed algorithms, and to compare with the results available in literature.

(a) *(b)* *(c)*

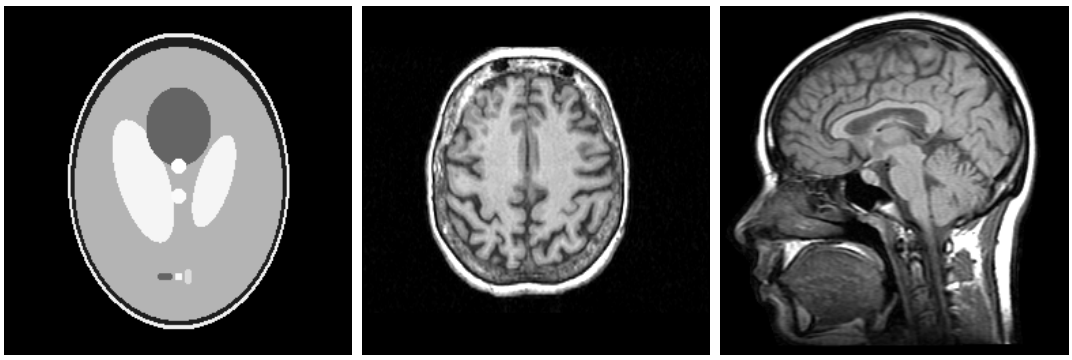

Figure 1.11: Images used for the experiments: *(a)* Shepp & Logan phantom, *(b)* an axial head slice, and *(c)* a sagittal head slice.

x_c	0	0	2	0	0	-40	-11	11	70	70	70
y_c	0	0	0	26	-26	0	0	0	-10	0	7
d_x	218	212	201	71	94	58	11	11	6	6	11
d_y	163	159	153	25	37	48	11	11	11	6	6
θ	0	0	0	-18	18	0	0	0	0	0	0
Intensity	240	30	180	245	245	100	255	250	100	250	220

Table 1.1: Ellipse specifications for the Shepp and Logan phantom

A single synthetic phantom image, known as the Shepp & Logan phantom, is used primarily due to the ease of comparison of the results to that of previous work [78]. Table 1.1 lists the ellipse specifications for the Shepp & Logan phantom [121]. Two real MR spin echo images, consisting of axial and sagittal head images, are also used to verify the performance of the proposed algorithms. These test images are shown in Figure 1.11.

1.10 Contributions of the Thesis

The aim of this thesis is to provide an extensive study into the methods for correcting in-plane rigid rotational motion artifacts in MRI. The general techniques that were previously proposed in literature, for motion artifact correction in MRI, are reviewed and categorized into real-time and post-processing techniques. For this thesis, the classification of post-processing techniques is reserved for those methods that utilize some motion model. Although motion artifact suppression in MRI has been a mature research topic for some time, only a limited number of research papers have been published on post-processing techniques, due to the difficulties in accurately modelling and estimating complicated patient motion. The correction algorithms for in-plane rigid rotational motion artifacts, proposed in the past are lacking in the following aspects:

- There is no theoretically derived method, supported by a comparative study, for management of overlapping data.
- There are no proposed data estimation schemes for filling in the data-void regions.
- There are no proposed rotational motion estimation schemes, capable of extracting the rotational parameters efficiently, in the midst of concurrent translations.
- Many motion estimation algorithms are based on minimizing the image energy outside the region of interest. However, there are no convincing methods in the past literature for extracting the object boundary from the motion corrupted image.

This thesis reports on a detailed analysis of these problems, in order to find effective and efficient solutions.

Although the creation of data overlaps and voids in the k-space, due to data correction via interpolation and re-gridding has been noted in past literature [7] [80], only simplistic solutions have been proposed. In this thesis an in-depth analysis of this problem is performed, and a new weighted averaging method of overlap data is suggested. A comparative study of different weighted averaging methods is conducted, to put the new method into perspective. An iterative estimation technique for filling the data void regions is also developed using POCS. The constraint sets for POCS are derived from acquired data and available *a priori* knowledge. The iterative algorithm is shown to diverge at times from the required image, after a finite number of iterations, due to inconsistency among the constraint sets. In order to avoid the divergence and computation of a complex cost function representing the image quality after each iteration, soft constraint sets and fuzzy projections are proposed and developed in this thesis.

One of the constraints applied in the iterative algorithm is the finite support of the imaged object, marked by the outer boundary of the region of interest (ROI). However, the object boundary extraction directly from the motion affected MR image can be difficult, especially if the motion function of the object is unknown. This thesis presents a new ROI extraction scheme based on entropy minimization in the image background.

The object rotation function is usually unknown or unable to be measured with sufficient accuracy. The motion estimation algorithm proposed in this thesis is based on maximizing the *similarity* among the k-space data subjected to overlap. This method is different to the typically applied parameter estimation technique based on the minimization of pixel energy outside the ROI. It has higher efficiency, and ability to estimate the rotational motion parameters in the midst of concurrent translational motion.

1.11 Scope and Contents

This section outlines the scope and contents of the remainder of the thesis. Each chapter describes a particular problem encountered in forming the overall rotational motion artifact correction algorithm. The algorithms aimed at solving these problems are presented with simulation experiment results.

The algorithm that was developed for ROI boundary extraction from the motion affected data is presented in Chapter 2. The views, least affected by motion, are selected using an entropy minimization criterion. Image background is further cleared using a pixel intensity based fuzzy model. The ROI contour extraction method using the active contour model (snake) is described in detail together with a brief review of the development of snake algorithms.

An analysis of the affects of in-plane rigid rotational motion on the acquired data is presented in Chapter 3. The possibility of modelling the rotations about any arbitrary axis, by a combination of translations and rotations about the centre of the FOV is explored and formulated. Due to the abundance of algorithms for translational motion artifact correction in past literature, the remaining algorithms in this thesis are directed towards suppressing artifacts arising from the rotations about the centre of the FOV. In order to limit the number of motion parameters, initial analytical studies are conducted on constrained models such as rotation at constant angular velocity and single step gross rotation.

General motion parameter estimation schemes for in-plane rigid rotations are discussed in Chapter 4. Real-time navigator echo technique is reviewed, in order to establish the advantages of motion estimation schemes in the post-processing stage. A new motion estimation algorithm is also proposed, which is based on maximizing the *similarity* among the k-space data subjected to overlap. This

method has higher efficiency and ability to estimate rotational motion parameters in the presence of translational motion.

The k-space data correction algorithm is presented in Chapter 5, with an analytical description of the creation of data overlap and void regions. A new weighted averaging scheme is introduced for managing data overlap regions, and compared with other possible alternative techniques. Iterative estimation of missing data in data void regions, is performed using POCS. The algorithm is shown to diverge at times after few iterations, due to inconsistency of the applied constraints. In order to find the point of termination of iterations, a new regulatory error metric is computed. Although this error metric provides the quality of the image at each iteration, it is computationally intensive and demands long computation times.

To avoid the use of regulatory error metric and the divergence of the POCS algorithm, a new method of regularization using soft constraints and fuzzy POCS is presented in Chapter 6. A study of the affect of fuzzy POCS parameters on the regularization process is also conducted, in order to establish guidelines for selecting appropriate parameters. The concluding chapter summarizes the results of research carried out for this thesis, and presents possible directions for future research work.

It should be noted that this research has been conducted based on the in-plane rigid movements of the imaged object. Therefore, the proposed algorithm for motion correction is limited by the following assumptions:

- No out-of-plane motion
- No shape distortions of the object (i.e. rigid motion)
- No magnetic field inhomogeneities.
- No changes to the main magnetic field.

- No multiple objects
- No relative motion of different parts of the same object.

According to the above assumptions, the proposed algorithm can be applied to 2DFT gradient-echo scans. In a gradient-echo scan, a stationary object can produce ghosts due to moving materials next to it. However, this situation will not occur under the above assumptions. In order to include such effects caused by the change in the main magnetic field, the motion model proposed in this thesis has to be supplanted with a model for magnetic field variations.