

Rapid 3D Phase Contrast Magnetic Resonance Angiography
through High-Moment Velocity Encoding
and 3D Parallel Imaging

by

Nicholas R. Zwart

A Dissertation Presented in Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

Approved July 2011 by the
Graduate Supervisory Committee:

David Frakes, Co-Chair
James Pipe, Co-Chair
Kevin Bennett
Bruce Towe
Josef Debbins

ARIZONA STATE UNIVERSITY

August 2011

ABSTRACT

Phase contrast magnetic resonance angiography (PCMRA) is a non-invasive imaging modality that is capable of producing quantitative vascular flow velocity information. The encoding of velocity information can significantly increase the imaging acquisition and reconstruction durations associated with this technique. The purpose of this work is to provide mechanisms for reducing the scan time of a 3D phase contrast exam, so that hemodynamic velocity data may be acquired robustly and with a high sensitivity. The methods developed in this work focus on the reduction of scan duration and reconstruction computation of a neurovascular PCMRA exam.

The reductions in scan duration are made through a combination of advances in imaging and velocity encoding methods. The imaging improvements are explored using rapid 3D imaging techniques such as spiral projection imaging (SPI), Fermat looped orthogonally encoded trajectories (FLORET), stack of spirals and stack of cones trajectories. Scan durations are also shortened through the use and development of a novel parallel imaging technique called Pretty Easy Parallel Imaging (PEPI). Improvements in the computational efficiency of PEPI and in general MRI reconstruction are made in the area of sample density estimation and correction of 3D trajectories. A new method of velocity encoding is demonstrated to provide more efficient signal to noise ratio (SNR) gains than current state of the art methods. The proposed velocity encoding achieves improved SNR through the use of high gradient moments and by resolving phase aliasing through the use measurement geometry and non-linear constraints.

To my wife, Christine, my mother, Cindy, and my sister, Katrina, for providing me with unequivocal support, and encouragement.

ACKNOWLEDGEMENTS

The undertaking of this degree has taught me that success is rarely achieved by an individual alone. My educational experience has been fostered through the guidance, counsel and encouragement of my committee members and fellow colleagues. For this reason I would like convey my deepest thanks to Dr. James Pipe, Dr. David Frakes, Dr. Josef Debbins, Dr. Bruce Towe and Dr. Kevin Bennett for their mentorship over the years. I'd also like to thank my colleagues at the Keller Center for Imaging Innovation for their teamwork. I would also like to mention that this work was funded in part by GE Healthcare and the American Heart Association.

TABLE OF CONTENTS

	Page
LIST OF TABLES	xi
LIST OF FIGURES	xii
CHAPTER	
1 INTRODUCTION	1
2 BACKGROUND	3
2.1 The MR Signal	3
2.1.1 Spin System	3
2.1.1.1 Spin	4
2.1.1.2 Applied Magnetic Fields	5
2.1.1.3 Spin Dynamics	7
2.1.2 System Excitation and Response	10
2.1.2.1 Transmission	10
2.1.2.2 Reception	11
2.1.3 System Modulation	12
2.1.3.1 Imaging Gradients	13
2.1.3.2 Contrast Preparation	14
2.2 Velocity Encoding	16
2.3 Parallel Imaging	20
3 SPIRAL PROJECTION IMAGING PCMRA	25
3.1 Theory	25
3.1.1 Spiral Projection Imaging	26
3.1.1.1 Data Conditioning	28
3.1.2 Dual VENC Velocity Encoding	28
3.1.2.1 Dual Low VENC	31

CHAPTER	Page
3.1.2.2 Composite Reconstruction	32
3.2 Methods	33
3.2.1 SPI Acquisition & Reconstruction	33
3.2.1.1 Data Conditioning	35
3.2.2 Dual VENC Velocity Encoding	35
3.2.2.1 Dual Low VENC	35
3.3 Results & Discussion	36
3.3.1 SPI Acquisition & Reconstruction	36
3.3.1.1 Data Conditioning	36
3.3.2 Dual VENC Velocity Encoding	40
3.3.2.1 Composite Reconstruction	41
3.3.2.2 Dual Low VENC	41
3.4 Conclusion	43
4 PRETTY EASY PARALLEL IMAGING	45
4.1 Theory	45
4.1.1 Convolution Based PEPI	47
4.1.2 Mask Generation	48
4.1.3 Conjugate Gradient PEPI	50
4.2 Methods	52
4.2.1 FFT Based PEPI	52
4.2.2 CG-PEPI Simulations	53
4.2.3 CG-PEPI and CG-SENSE	54
4.3 Results and Discussion	54
4.3.1 FFT Based PEPI	54
4.3.2 CG-PEPI Simulations	58

CHAPTER	Page
4.3.3 CG-PEPI and CG-SENSE	60
4.4 Conclusion	62
5 SAMPLE DENSITY CORRECTION.....	63
5.1 Theory	64
5.1.1 Gridding Reconstruction	64
5.1.2 Iterative Sample Density Estimation	65
5.1.3 Direct Method Design	66
5.1.3.1 Optimal Convolution Kernel.....	67
5.1.3.2 Computational Optimization	67
5.1.4 Grid Method Design	68
5.1.4.1 Ideal Grid Kernel	68
5.1.4.2 Approximate Grid Kernel	69
5.2 Simulations	71
5.2.1 Data	71
5.2.1.1 Trajectories	71
5.2.1.2 Images	73
5.2.2 Optimal Grid Oversample Factor	74
5.2.3 Direct Method Comparison	74
5.2.4 Analytical DCF Comparison	76
5.2.5 Initial Conditions	76
5.2.6 Parallel Computation	78
5.3 Results	78
5.3.1 Grid Oversample Factor	79
5.3.1.1 Accuracy and Convergence Properties	79
5.3.1.2 Trajectory and Image Variation	79

CHAPTER	Page
5.3.2 Direct Method Comparison	80
5.3.2.1 Accuracy	80
5.3.2.2 Execution Time	81
5.3.3 Analytical DCF Comparison	82
5.3.4 Initial Conditions	84
5.4 Discussion	86
5.4.1 Grid Oversample Factor	86
5.4.2 Direct Method Comparison	87
5.4.3 Analytical DCF Comparison	88
5.4.4 Initial Conditions & Convergence	88
5.5 Conclusion	89
5.6 Code Release	89
6 MULTI-DIRECTIONAL HIGH MOMENT ENCODING	91
6.1 Theory	92
6.1.1 Unaliasing Phase	93
6.1.2 Proposed Method	94
6.1.2.1 Velocity Encoding in 2D	95
6.1.2.2 Velocity Encoding in 3D	97
6.2 Experiments	101
6.2.1 Noise Sensitivity	101
6.2.1.1 Methods	101
6.2.1.2 Results & Discussion	103
6.2.2 Optimal Voxel Shape	106
6.2.2.1 Methods	110
6.2.2.2 Results & Discussion	111

CHAPTER	Page
6.2.3 SNR Optimized Moment Balancing	116
6.2.3.1 Methods	118
6.2.3.2 Results & Discussion	122
6.2.4 Constraints & Reconstruction	127
6.2.4.1 Methods	136
6.2.4.2 Results & Discussion	137
6.2.5 Dual VENC Method Comparison	141
6.2.5.1 Methods	141
6.2.5.2 Results & Discussion	143
6.3 Summary Discussion	148
6.3.1 Noise Sensitivity	148
6.3.2 Optimal Voxel Shape	148
6.3.3 SNR Optimized Moment Balancing	149
6.3.4 Constraints & Reconstruction	149
6.3.5 Dual VENC Method Comparison	149
6.3.6 Conclusion	150
7 DISCUSSION	151
8 FUTURE WORK	154
9 SUMMARY CONCLUSIONS	156
REFERENCES	157
APPENDIX	
A SPATIAL OFF-RESONANCE	168
A.1 Off-Resonance Vs. Head Position	169
B DYNAMIC 2D PEPI RECONSTRUCTION	171
B.1 ISMRM 2010 Data Reconstruction Challenge	172

LIST OF TABLES

Table	Page
6.1 The normalized and effective SNR (time-normalized) using the relative acquisition time for each method.....	104
6.2 Measurement direction matrices (u) and moment sums for each principle direction.....	120

LIST OF FIGURES

Figure	Page
2.1 The effects of spinning.	4
2.2 Spin systems with and without an applied magnetic field at equilibrium.	6
2.3 Spin rotation and precession about an external magnetic field..	7
2.4 A PCMRA pulse sequence diagram for an arbitrary slice select and imaging gradient.	17
2.5 Sagittal PCMRA maximum intensity projections.	19
2.6 Two 2D undersampled trajectories and corresponding aliased images.	21
2.7 Coil combination and SENSE parallel imaging reconstructions.	22
3.1 A spiral projection imaging trajectory diagram.	26
3.2 The effects of MTF modulation for two different trajectory configurations.	37
3.3 SPI radial k-space density.	38
3.4 Properties of 1D spiral arm gridding in SPI.	39
3.5 Signal loss maximum intensity projections for low VENC (a) and high VENC (b) acquisitions.	41
3.6 Sagittal maximum intensity projections of dual-VENC reconstruction (a) compared to composite dual-VENC reconstruction (b).....	42
3.7 1D simulation profiles of the dual-low-VENC method.	43
4.1 PEPI reconstruction flow diagrams.....	46

Figure	Page
4.2 The procedure for generating the MTF mask used in PEPI reconstruction.	49
4.3 Flow diagrams of the SENSE (red) and PEPI (blue) reconstruction pipelines.	51
4.4 Axial slices of reconstructed data.	56
4.5 Sagittal slices of reconstructed data.	57
4.6 Simulated CG-PEPI reconstructed data at various levels of undersampling (R).	59
4.7 Parallel reconstruction of FLORET data at undersample factors of 4 and 9.	61
5.1 An illustration of how Eq. 5.3 is executed on the first iteration.	66
5.2 The spatial and frequency domain profiles of the proposed kernel.	69
5.3 Fully supported sampling of the ideal kernel (a) and aliasing of the truncated kernel (b).	70
5.4 Simulation elements.	72
5.5 Error convergence at various levels of oversampling.	80
5.6 The direct and grid method comparisons.	81
5.7 The grid method is compared to the analytical DCF for the SPI and SoC trajectories.	83
5.8 The error convergence rate of various initial conditions for each trajectory using the grid method.	85

Figure	Page
6.1 Orientations for six velocity sensitive directions that make the vertices of a regular icosahedron.	96
6.2 A 2D example of the solution spaces for two measurement directions.	98
6.3 A 2D example of the solution space for 3 and 4 measurement directions allowing for 2 wraps (ie. $ k_i \leq 2$).	99
6.4 Error sensitivity in terms of consistency error for 6 directions using 1, 2 and 3 wraps (ie. $ k \leq 1$, $ k \leq 2$ and $ k \leq 3$ respectively).	100
6.5 The error sensitivity thresholds of 7 different measurement direction configurations versus input noise.....	104
6.6 Intra-voxel phase induced signal loss.	107
6.7 Intra-voxel phase distribution model and limits.	108
6.8 Additive and subtractive signal biased phase.	109
6.9 A 2D cross-section of a 3D laminar flow model (a) and corresponding MTF filters (b) used in simulation.	110
6.10 The bias effects of voxel shape in terms of resolution and low pass filtering.	112
6.11 3D simulated profile cross sections.	114
6.12 Level curves of maximum bias in the 3D simulated profiles.	115
6.13 An illustration of moment addition in balanced and semi-balanced cases.....	116
6.14 Moment balanced measurement vector configurations.	120

Figure	Page
6.15 Measured noise produced in various moment balancing strategies.	123
6.16 Flow phantom experiments comparing dual-VENC and 6-direction reconstructions.	126
6.17 A 1D cross section of flow for both 6-direction and dual-VENC acquisitions at a 20cm/s VENC.	128
6.18 The sources of phase bias sensitivity in 2D low moment encoding.	129
6.19 Axial PCMRA reconstructions showing signal loss and consistency error specificity.	133
6.20 An approximate relationship between the corrupted PSF in phase biased areas and the average image PSF.	135
6.21 Axial, coronal, and sagittal maximum intensity projections demonstrating each constraint in the proposed reconstruction process.	138
6.22 Measured velocity profiles of the basilar artery for constrained reconstruction compared to an 80cm/s VENC referenced reconstruction.	140
6.23 Maximum intensity projections in the axial, coronal, and sagittal planes.	144
6.24 Maximum intensity projections in the axial, coronal, and sagittal planes of deblurred data sets.	145
6.25 Blurred and deblurred basilar artery cross sections.	146
6.26 A 1D cross section of the basilar artery flow lumen.	146

Figure	Page
A.1 Spatial off resonance at air tissue interfaces for various head positions.....	170
B.1 2D dynamic PEPI reconstruction.....	174

1 INTRODUCTION

Phase contrast magnetic resonance angiography (PCMRA) is an MR modality that is capable of producing quantitative flow velocity information. Multiple images are acquired with velocity sensitivity in each of the principle (x, y and z) directions. A 3D velocity vector can then be constructed using the magnitude and directional information contained in these data sets. Since the construction of a velocity vector requires multiple acquisitions (at least one in each principle axis, three for 3D), the acquisition time of the base imaging sequence more than triples. The long scan durations make PCMRA less viable for clinical use. The goal of this work is to provide mechanisms for reducing the scan time of a 3D PCMRA, so that hemodynamic velocity data may be acquired robustly and with a high sensitivity.

The methods proposed for accomplishing this goal are applied to two areas of the scan technique: the base imaging technique and the contrast encoding technique. The proposed work starts with the combination of a rapid 3D imaging technique (1, 2) with a high gradient moment velocity encoding strategy (3, 4). Parallel imaging acquisition and reconstruction strategies are then explored using a novel time saving reconstruction technique (5–7). In an additional effort to reduce the computational overhead, a 3D sample density correction algorithm (8) is optimized through the combination of two existing techniques (9, 10). The presented work is concluded with the introduction and optimization of a novel 3D PCMRA encoding and reconstruction technique.

The application of these techniques is focused on the construction of a scan and reconstruction technique that can provide full volume 3D neuroangiographic images within clinically viable scan durations. The practical requirement of reducing the scan time provides the technical challenges of main-

taining a minimum signal to noise ratio (SNR) and spatial resolution necessary for providing phase contrast. This work demonstrates the relative quality and computation time in the reconstruction of the uniformly undersampled 3D trajectories such as spiral projection imaging (SPI) (11, 12), and FLORET (13), between the CG-SENSE (14) and proposed CG-PEPI parallel reconstruction methods. The proposed sample density correction method demonstrates substantial time reductions in estimating the densities of 3D center-out trajectories compared to the current state of the art. The proposed method is not only shown to be one of the fastest and most accurate algorithms, it is also completely generic, allowing any arbitrary trajectory to be density compensated extemporaneously. The novel 3D phase contrast method proposed in this work is shown to provide significant gains in signal to noise ratio efficiency and high velocity sensitivity producing more vessel conspicuity as compared to the current state of the art.

2 BACKGROUND

This chapter covers some of the principles of MRI acquisition and reconstruction techniques needed to understand the proposed projects presented in the subsequent chapters. This background information covers the scan trajectory used in most of the preliminary work (spiral projection imaging), velocity encoding, and parallel imaging.

2.1 The MR Signal

This background derives the higher level signal mechanics, used in the discussions about velocity encoding and k-space, starting from the interactions of spins with an external magnetic force. The magnetic resonance (MR) signal is characterized in many different ways depending on the application or level within the imaging system. For this work the focus will be to introduce a basic overview of the MR imaging process used in this work by starting with what is being measured and finishing the discussion with how it is being measured. The information presented here is a summary overview of the MR scanning process. The suggested references for a more complete and detailed description are (10, 15, 16).

2.1.1 Spin System

The MR system that is characterized through experiments, scans and examinations is called a spin system. This system is used as a physical transfer function for modulating input signals. The basic procedure for generating MR signal is to perturb the spin system with various input stimuli and use the measured response to reconstruct an MR image. This section introduces the building blocks of the spin system and the conditions with which stimulating and measuring the system response are possible.

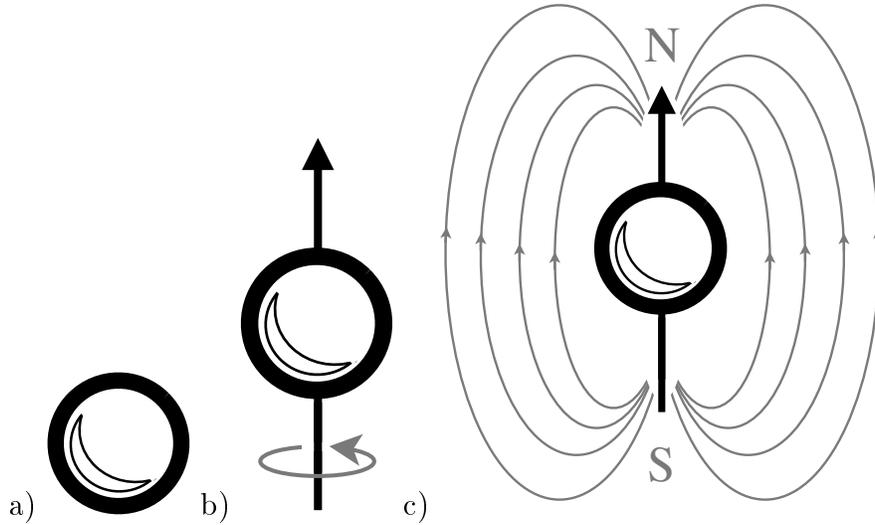


FIG. 2.1: The effects of spinning. (a) a single proton. (b) a proton rotating about a central axis (a 'spin'). (c) magnetic flux lines surrounding a spin, similar to a bar magnet.

2.1.1.1 Spin

The basic element of the spin system is described in terms of the mechanical phenomena that accompanies certain atomic nuclei. Some nuclei possess an angular momentum about a central axis that resists the change in direction of the axis, similar to a spinning top or gyroscope. This angular precession about the central axis is the motivation for the term 'spin'. The implicit charge associated with the subatomic particles contained within the nucleus combine with the spinning property to produce a dipole magnetic field that surrounds the particle.

Nuclei that can be coherently influenced by external magnetic fields are those that possess a spin- $\frac{1}{2}$ quantum number. Spin- $\frac{1}{2}$ nuclei are comprised of an odd number of protons and/or an odd number of neutrons. Examples of spin- $\frac{1}{2}$ particles are hydrogen, sodium, and carbon-13.

Since the hydrogen proton is the most abundant spin- $\frac{1}{2}$ particle in the

human body, it is the most common target spin system used in clinical MRI.

Figure 2.1 shows an illustration of a spin- $\frac{1}{2}$ particle (a), rotating around a central axis (b), and causing a dipole magnetic field (c).

2.1.1.2 Applied Magnetic Fields

The magnetic field surrounding a spin can be influenced by an external magnetic field. The aggregate magnetic moment of a spin- $\frac{1}{2}$ population will align with the magnetic field in either a parallel or antiparallel direction relative to the externally applied magnetic field. These two states are referred to as ‘spin-up’ and ‘spin-down’. Most of the spin system will be balanced, containing equal numbers of spin-up and spin-down particles. Due to the Zeeman effect, a small fraction of the spins will remain spin-up based on the system temperature and strength of the applied magnetic field. These residual spin-up particles constitute the excitable and measurable MR signal. For a proton system (water) at room temperature in a 3 Tesla field this will correspond to approximately 10 parts per million available for MR signal.

The dynamics of moment alignment with an external field take place on a sub-second time scale. This transient period is where the MR signal is measured. Figure 2.2(a) depicts spins outside of an external magnetic field finding random orientations (due to thermal noise (17)) as they maintain a system equilibrium at the lowest energy state. Figure 2.2(b) depicts system equilibrium in the presence of a magnetic field (B_0), where the net magnetic moment of the system is aligned with the external field B_0 . Thermal energy in the system causes random fluctuations in spin orientation making the simplified example of Zeeman splitting (figure 2.2(b)) look more like the illustration in figure 2.2(c) where the simplified example is a superposition of the randomly

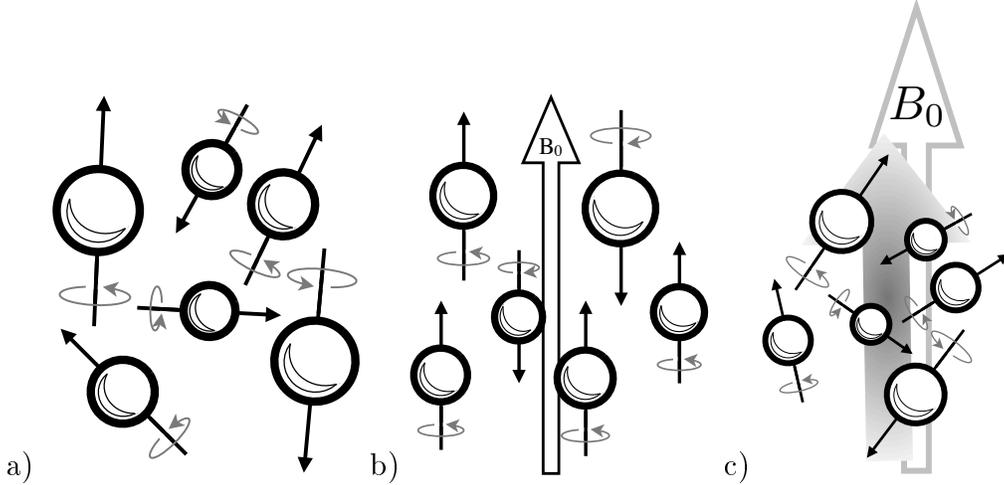


FIG. 2.2: Spin systems with and without an applied magnetic field at equilibrium. (a) a spin system with random individual orientations yielding no net moment. (b) a simplified spin system with net moment aligned to an external magnetic field (B_0). (c) a spin system with net moment (grey arrow) aligned to an external magnetic field (B_0).

oriented moments (18). In the time between the states shown in (a) and (b), the spins precess about the applied magnetic field as illustrated in figure 2.3. A common analogy for spin alignment during this period is the time course of a spinning top. The top resists gravity due to its angular momentum just as the spin resists the pull of the applied magnetic field acting on its own magnetic field. The rotational axis of a spinning top precesses about an axis parallel to gravity as it decays to a position orthogonal to the direction of gravity. The alignment of a spin is opposite, in that its rotational axis becomes parallel to the external magnetic field as shown in figure 2.3. The angle between the rotational axis and the precessional axis is reduced at an exponential decay rate.

In summary, the rotational axis of the spin is also the axis of the dipole magnetic field (figure 2.1(c)). The alignment of the spin field with the applied external field is caused by magnetic attraction, while the spin rotation resists

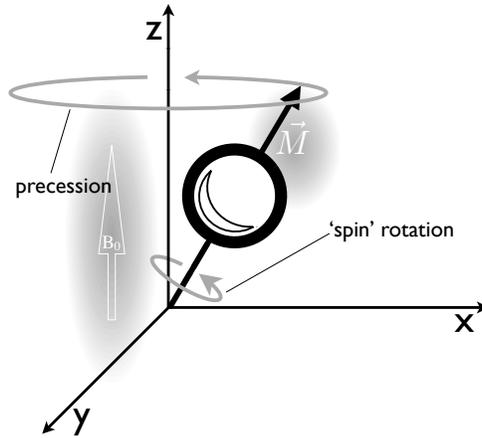


FIG. 2.3: Spin rotation and precession about an external magnetic field. This example is an analogue for a spinning top.

the change in spin orientation causing a precession of the spin field about the applied field.

2.1.1.3 Spin Dynamics

The classical physics description of the spin dynamics in the presence of an applied magnetic field is based on the aggregate magnetic field of a spin population. This allows the system to be described in terms of continuously defined functions. This section recapitulates the qualitative assessment of spin dynamics presented in the previous subsection using the net magnetic moment of a spin population as the basic element of spin motion.

The rate of precession of the rotational axis about the applied magnetic field (B_0) is a function of the field strength and the spin rate. The spin rate is native to each particle. The rate is related by a constant called the gyromagnetic ratio (γ). For a proton, γ is roughly 42.58 MHz/Tesla , and the precessional frequency (ω) is,

$$\omega = \gamma B_0 , \tag{2.1}$$

which is known as the Larmor frequency. The precessional frequency corresponds to the time varying rate of phase change. The rate of precession is therefore the amount of phase change (measured in degrees, cycles, or radians) divided by the frequency. The relative phase relationship between spins can cause their magnetic fields to add coherently or incoherently, producing a larger or smaller aggregate magnetic field.

The precession of a spin in time and space is characterized by the orientation of the spin field in three dimensions \vec{M} and the orientation of the applied field \vec{B} . By holding the applied field as the frame of reference, the change in spin orientation ($\frac{d\vec{M}}{dt}$) is,

$$\frac{d\vec{M}}{dt} = \gamma \vec{M} \times \vec{B} \tag{2.2}$$

where \times represents the cross product operation. The cross product in equation 2.2 relates the precessional rate of change to the physical orientation of \vec{M} about \vec{B} . Spin precession is the key component to the generation of the MR signal. For this reason, MR scans are performed in the presence of a constant high level magnetic field.

The decay (or relaxation) of the angle between \vec{M} and \vec{B} is governed by two independent relationships. From figure 2.3, the component of \vec{M} that lies in the x-y plane is referred to as the ‘transverse component’ and the component in the z direction is called the ‘longitudinal component’. The spin fields orient themselves to minimize the energy in the system, therefore the return of the spins to the equilibrium state is referred to as ‘relaxation’ (19).

The relaxation of the transverse component is due to small changes in temperature that vary within the spin system. These variations randomly change the relative precessional-phase relationship between spins of a sample population. The loss in aggregate magnetic field reduces the transverse component of \vec{M} by a time varying exponential decay,

$$M_{x,y}(t) = M_{x,y}(0)e^{\frac{-t}{T_2}} , \quad (2.3)$$

where $M_{x,y}$ is the time varying transverse magnetization, t is the time, and T_2 is the decay constant that is dependent upon the spin environment.

The relaxation of the longitudinal component of \vec{M} is analogous to a top losing its rotational momentum. The momentum is what opposes the change in \vec{M} orientation. As the analogous rotational momentum is lost, the spin aligns with \vec{B} . The longitudinal component is recovered over time and modeled as,

$$M_z(t) = M_{z,0}(1 - e^{\frac{-t}{T_1}}) , \quad (2.4)$$

where M_z is the time varying longitudinal magnetization, $M_{z,0}$ is the magnetization at equilibrium and T_1 is the recovery constant, which is also environment dependent.

Equations 2.3 and 2.4 are solutions to additional terms of the Bloch equation in 2.2, where the rate of change in spin magnetization is,

$$\frac{d\vec{M}}{dt} = \gamma\vec{M} \times \vec{B} - \frac{(M_x\vec{i} + M_y\vec{j})}{T_2} - \frac{(M_z - M_{z,0})\vec{k}}{T_1} , \quad (2.5)$$

which describes the precession with the effects of relaxation.

2.1.2 System Excitation and Response

As previously mentioned, the spin system is analogous to a transfer function; information about the system is collected by perturbing the system and measuring the response. The system perturbation and response are administered and measured through radio frequency (RF) transmissions. An overview of this procedure begins with a spin system placed in a constant high level magnetic field (B_0). Before an MR scan begins the spin system is allowed enough time to reach an equilibrium state (i.e. $time \gg T_1 \gg T_2$), which is in alignment with the magnetic field B_0 . B_0 remains present over the duration of the scan. The system is then perturbed by an applied RF pulse, which tips the spin magnetization into the longitudinal plane. Since the spin is no longer aligned with B_0 , it continues to precess about B_0 until equilibrium is again reached. The precession of the spin field induces an RF pulse which is measured by the scanner.

This section introduces the fundamentals of transmission and reception of MR-RF signals that comprise an MR scanning system.

2.1.2.1 Transmission

The process of applying RF to the spin system is called ‘excitation’. Spins are excited when they are forced out of B_0 alignment into the precessional state. The RF pulse accomplishes this by creating a momentary magnetic field that is effectively stronger than B_0 .

The RF is applied at a frequency that is on resonance with the precessional frequency ω (equation 2.1). The momentary field B_{RF} is much smaller than the B_0 , however, since the pulse is applied at resonance, the effective strength is higher causing the spins to precess about B_{RF} . B_{RF} is orthogonal to B_0

causing the spins to tip away from the longitudinal axis (z) into the transverse plane (x - y) (figure 2.3).

The duration and amplitude of B_{RF} determines how far the spins are tipped angularly into the transverse plane. This tip angle (or flip angle) is,

$$\theta = \gamma B_{RF} T , \tag{2.6}$$

where T is the RF pulse duration.

The RF wave transmission is a fundamental property of Faraday's law. An alternating current is passed through an electrically conductive coil causing an alternating magnetic flux. The amplitude of the RF wave is proportional to the applied current. The coil surface is oriented parallel to B_0 and as close to the spin system as possible.

2.1.2.2 Reception

The MR signal measured from an excited spin system is received by the same Faraday property used in the excitation of the system. An electrically conductive coil (potentially the same used in RF transmission) is placed near the spin system, the alternating magnetic field generated by the precessing spins induces an alternating current within the coil. This current is measured in time and demodulated at the precessional frequency to produce the relative spin frequency. The relative spin frequency can be spatially varied which yields the information necessary to reconstruct an MR image, as discussed in the next section.

The excitation and signal reception is repeated, generally allowing the spins to reach equilibrium before each repetition. This is known as the repetition time or TR. Hundreds of TRs are measured over the course of an MR scan

using techniques to provide different information about the spin system in each iteration, as will be discussed in the next section.

2.1.3 System Modulation

The spin system can be examined through excitation and measurement of its response. As discussed the previous sections, spins precess at a frequency proportional to the magnetic field strength of an externally applied field. These two principles are combined to gain information about the spin system under various conditions created by changing the spatial distribution of the magnetic field. This is accomplished through magnetic gradients created by electromagnets which are positioned on each of the three principle axes (x, y and z). These gradients can be varied in time and amplitude to produce unique conditions that provide information about the spatial distribution of spins.

Additionally, the materials or biological tissues that comprise the spin system vary in longitudinal and transverse relaxation rates. The RF pulse can also be varied in time and amplitude to produce distinctive signals based on the relationship of these rates.

Finally, spin systems themselves may also vary in time providing additional signals (e.g. from physiologic process) that can be manipulated through the use of RF and magnetic field gradients.

The amalgamation of strategically placed time varying pulses of both RF and magnetic gradient is known as a ‘pulse sequence’ and is the examination protocol for characterizing a spin system. This section covers the imaging, relaxation contrast and physiological contrast mechanisms that are combined to make the pulse sequence that used the body of this work.

2.1.3.1 Imaging Gradients

In order to measure the spatially varying spin distribution of an object the relative intensities of spatial harmonics must be measured. The way samples are interpreted as spatial harmonics depends on an abstraction called k-space.

As briefly discussed in the previous sections, a spin that is placed in a magnetic field will precess about this field. The rate at which it precesses depends on the strength of the field as shown in equation 2.1. If additional linear gradients G are present, the frequency of precession is spatially dependent on position x within the object:

$$\omega = \gamma(B_0 + G_x x) . \quad (2.7)$$

A receive coil placed next to an excited spin system does not selectively measure the magnetic flux of each spin independently, instead, it measures the spin system as a whole as indicated by the following equation.

$$S(t) = \sum_{x \in O} M(x) e^{-i\omega t} \quad (2.8)$$

where S is the time varying MR signal, O is the extent of the object, and M is the spatially varying spin magnetization precessing at a frequency ω . In one dimension, the time varying signal is the sum of all spins in x across the object O .

By substituting the positional dependent frequency ω with equation 2.7, the time dependent signal becomes

$$S(t) = \sum_{x \in O} M(x) e^{-i\gamma(B_0 + G_x x)t} . \quad (2.9)$$

If the system is viewed from the precessional frame of reference, the spatial variation of precession frequency is all that is left. This can also be viewed as demodulation and results in the summation of only the relative change in frequency by dropping the ‘ B_0 ’ term from equation 2.9. Substituting in $k = \gamma G_x t$ yields the identity of the discrete Fourier transform of the magnetization profile M .

$$S(k_x) = \sum_{x \in O} M(x) e^{-ik_x x} \quad (2.10)$$

In terms of the Fourier transform, k is the spatial harmonic and is also the namesake of ‘k’-space. K-space is representative of the Fourier transform frequency domain. Using both equations 2.10 and 2.9 allows mapping of sample points in time to harmonic positions within the Fourier domain. When the Fourier domain has been completely mapped, a Fourier transform is applied to create the final image.

Another result of equation 2.9 is that the Fourier harmonics that are sampled, can be modulated in time by changing the amplitude of the spatial magnetic gradients in time $G_x(t)$. The path, through k-space, created by changing $G_x(t)$ is called a trajectory and can be thought of as tracking a particle through a force vector field. A trajectory velocity and acceleration in k-space are analogous to G_x and $\frac{dG_x}{dt}$ respectively.

2.1.3.2 Contrast Preparation

Spoiled Gradient Recalled Echo (SPGR) The particular class of scan sequence used in this project is called the spoiled gradient echo or SPGR sequence. Spoiling means that the transverse magnetization from a previous TR is actively attenuated before each successive sampling period. A ‘gradient echo’

refers to the way the signal is varied in time by the imaging gradients as previously discussed. The equations presented in this section explore the way MR signal is affected by the timing characteristics of the pulse sequence and the relaxation properties of the spin system.

In order to understand the steady state behavior of an SPGR signal, the effects describing the recovery of longitudinal magnetization (T_1) will be examined. Transverse relaxation (T_2) can be ignored since the signal is attenuated through spoiling. Assuming a perfectly spoiled system, the RF pulse will excite only the recovered longitudinal magnetization. The time course of longitudinal magnetization is, taken from the third term of equation 2.5,

$$\frac{dM_z}{dt} = \frac{M_{z,0} - M_z}{T_1} \quad (2.11)$$

and when solved for M_z becomes equation 2.4 where $M_{z,0}$ represents the equilibrium magnetization. A flip angle (α) will leave a component of longitudinal magnetization equal to $\cos(\alpha)$ times the last available $M_z(0)$ for each TR. The $M_z(0)$ of a subsequent TR is equal to $M_z(TR) \cdot \cos(\alpha)$ of the previous TR. If the TR is sufficiently less than the T_1 , the longitudinal magnetization will not fully recover over contiguous TRs. Eventually the system will reach a steady state where the longitudinal magnetization will be equal across TRs. A relative measure of signal strength at this point is longitudinal magnetization normalize by the equilibrium magnetization at the end of each TR as follows.

$$\frac{M_z}{M_{z,0}} = \frac{1 - e^{-TR/T_1}}{1 - \cos(\alpha) \cdot e^{-TR/T_1}} \quad (2.12)$$

In order to ensure the steady state is at the maximal achievable signal strength, the Ernst angle, calculated by the following equation, is used.

$$\alpha = \arccos(e^{-TR/T_1}) \quad (2.13)$$

The basic signal mechanics presented in this section are applied temporally, within the pulse sequence at the beginning (i.e. the flip angle and RF spoiling), at the end (i.e. the gradient spoiling), and implicitly based on the flip angle to TR relationship.

Other Contrasts The physiological based contrast mechanism used in this work is known as ‘phase contrast’ (20). The physiological signal contrast is encoded in relative spin precessional phase by a preparatory gradient pulse that is applied before the imaging sequence. This phase is induced by virtue of the spin-particle motion through the applied magnetic gradient. Preparatory sequences are a common mechanism used in MRI to encode various types of contrasts (e.g. diffusion weighted imaging (21)). Physiological signals may also be given contrast implicitly through the sequence timing just as T_1 relaxation is in an SPGR sequence (e.g. time of flight MRA (22)).

2.2 Velocity Encoding

Velocity encoding is the mechanism by which flow contrast is encoded into the image phase. The following section covers how velocity encoding is accomplished, and some of the challenges in achieving high quality velocity maps.

In the simplest case, velocity is encoded in the direction of one of the physical gradients. Figure 2.4 is a diagram of a PC-MRA pulse sequence for an arbitrary imaging gradient. The pulses shown in red are bipolar gradient pulses that work together to encode spin velocity as a phase in the reconstructed image. The bipolar pulses are played out after the RF excitation (and possibly a slice selection gradient) and before the imaging gradients, which are shown

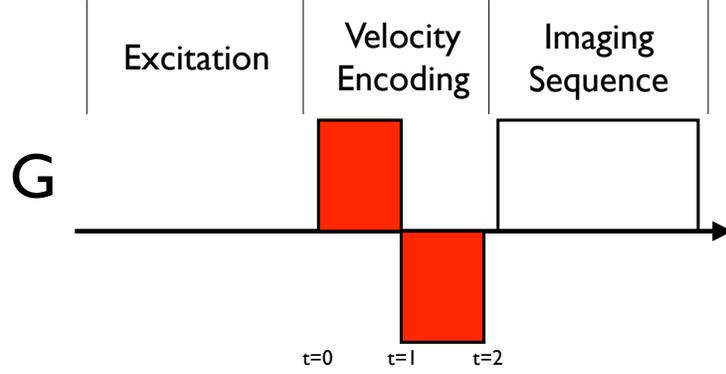


FIG. 2.4: A PCMRA pulse sequence diagram for an arbitrary slice select and imaging gradient.

as a nondescript block. The velocity encoding itself will be described in terms of the time points marked below the bipolar pulses.

The sampling trajectory in k-space is determined by the imaging gradients. The position in k-space is called the ‘zeroth moment’ and is determined by the area of the time varying gradient pulse by,

$$k(t) = \gamma \int_0^t G(u) du \quad (2.14)$$

where k is the position in k-space, $G(t)$ is the gradient amplitude at time t , and γ is the gyromagnetic ratio. Between points $t = 0$ and $t = 2$ the net zeroth moment is zero. The bipolar pulse is essentially a trajectory that has moved out, from the center of k-space, along a k-space axis and then moved back to the center.

Spatially, the spins accrue a phase when an imaging gradient is pulsed. Between points $t = 0$ and $t = 1$ a positive gradient lobe is played out and a phase is accrued in each spin based on its spatial position as,

$$\phi(t) = \gamma \int_{t=0}^{t=1} G(u)x(u) du , \quad (2.15)$$

where ϕ is the phase and $x(t)$ is the location of the spin. Between points $t = 1$ and $t = 2$, all spins that have remained still (at a fixed position x) will have accrued a net phase of zero. All spins that have moved will maintain a residual phase. This residual phase is proportional the velocity of the spin. However, the relative phase angle velocity encoded image alone is arbitrary, and so a reference image with no bipolar lobes is required to estimate the base phase angle.

The encoded velocities are represented by a phase between $-\pi$ to π , for positive and negative velocities. The proportional relationship between velocity and phase is set by the gradient strength of the bipolar lobes. The velocity component of the spin position is

$$x(t) = x_0 + vt . \quad (2.16)$$

By combining the velocity component with equation 2.15 and finding the net phase of both of the bipolar lobes, the phase can then be calculated as,

$$\phi(t) = \gamma \int_{t=0}^{t=\Delta_t} G(u)x(u)du \quad (2.17)$$

$$- \gamma \int_{t=\Delta_t}^{t=\Delta_t+\Delta_t} G(u)x(u)du \quad (2.18)$$

$$= \gamma \int_{t=0}^{t=\Delta_t} G(u)[x_0 + vu]du \quad (2.19)$$

$$- \gamma \int_{t=\Delta_t}^{t=\Delta_t+\Delta_t} G(u)[x_0 + vu]du \quad (2.20)$$

$$= \gamma[G\Delta_t x_0 + Gv\Delta_t^2 - Gx_0(2\Delta_t) \quad (2.21)$$

$$- Gv(2\Delta_t)^2 + Gx_0\Delta_t + Gv\Delta_t^2] \quad (2.22)$$

$$= \gamma G\Delta_t^2 v , \quad (2.23)$$

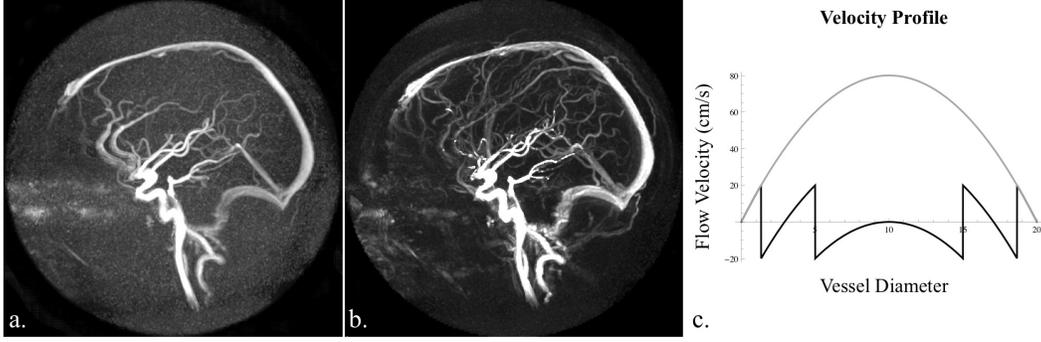


FIG. 2.5: Sagittal PCMRA maximum intensity projections. (a) high $VENC = 100 \frac{cm}{s}$. (b) low $VENC = 20 \frac{cm}{s}$. (c) plot of laminar velocity profile. The grey plot is representative of the actual cross-sectional flow through a vessel. The black plot is representative of a reconstructed phase-aliased version of the flow profile if the max flow velocity is $80 \frac{cm}{s}$ and the VENC is $20 \frac{cm}{s}$.

where Δ_t represents a generalized time difference between points $t = 0, 1$ and 2 . As previously noted, the zeroth moment, relating phase to position x_0 , drops out leaving only velocity sensitivity. The ‘first moment’, m_1 , is represented by $G\Delta_t^2$ in the final product above. m_1 is used to relate the estimated velocity to the encoded phase.

The maximum representable velocity ‘VENC’ is calculated using the first moment and the maximally distinguishable phase, in either direction, as,

$$VENC = \frac{\pi}{\gamma|m_1|} \quad (2.24)$$

Velocities that exceed the VENC cause phase aliasing wraps in the encoded image. Velocities that produce a phase magnitude higher than π can not be distinguished from velocities that produce a phase at any equivalent modulus of 2π . Figure 2.5(c) shows a simulated cross-sectional flow for both sufficient (relatively high) VENC and insufficient (relatively low) VENC. The lower insufficient VENC is phase aliased where the velocity exceeds the VENC.

Typically the maximal expected velocity is lower than $100\text{cm}/s$. Setting the VENC to a high value such as this will allow the high flow to be represented and the slower flow to vanish under the noise floor. The metric that describes this relationship is velocity to noise ratio (VNR) and is calculated as,

$$VNR \propto \frac{V}{VENC} SNR . \quad (2.25)$$

This relationship shows that choosing a VENC that is close the maximum expected flow velocity will be the most effective.

Figure 2.5 show a comparison between the high and low VENC encodings. The images were scanned using a 1mm^3 voxel, 240mm field of view (FOV), and a VENC of $100\frac{\text{cm}}{\text{s}}$ and $20\frac{\text{cm}}{\text{s}}$ for the high and low encodings respectively. The high VENC data set features no phase aliasing in the high flow areas such as the internal carotid arteries. Contrastingly, the low VENC set shows sharp boundaries at the edges of high flow signifying phase aliasing. However, the low VENC image shows a lower noise floor making smaller vessels more visible.

2.3 Parallel Imaging

Parallel imaging is a trajectory and reconstruction method that generally works to shorten the acquisition duration. The time reductions of the scan are accomplished by undersampling (sampling below the Nyquist rate) k-space (5, 6, 23). Parallel reconstruction works to fill in the undersampled areas of k-space by using prior information or by enforcing constraints in an iterative solution process. Other reconstruction strategies utilize the sparsity in a chosen domain of the collected data (24), and some simply enforce Hermite symmetry in the Fourier domain (25–27). The method of parallel imaging used in this work is called sensitivity encoding or SENSE (14, 28).

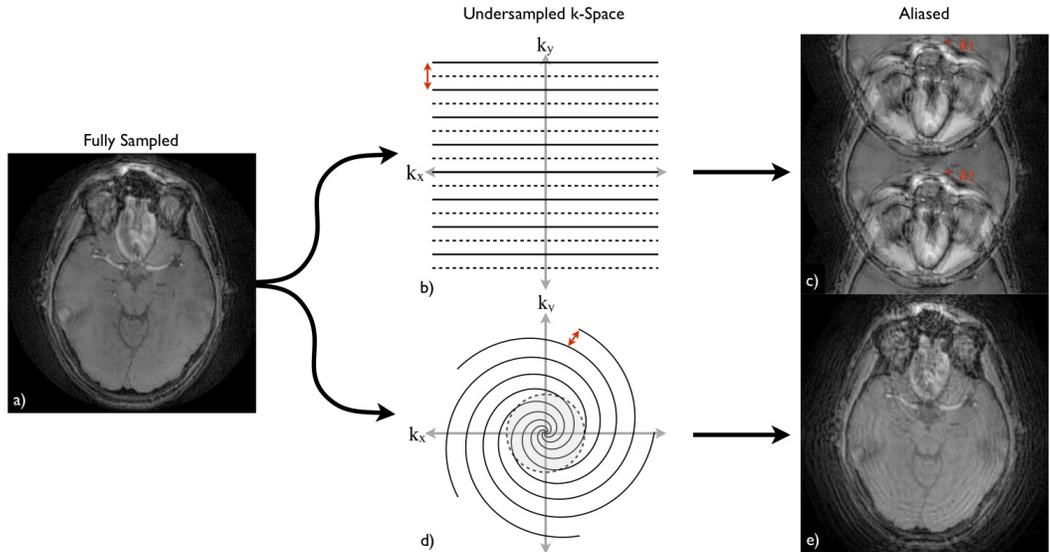


FIG. 2.6: Two 2D undersampled trajectories and corresponding aliased images. A fully sampled image (a), when undersampled in k-space (b) and (d) produces aliased images (c) and (e).

The undersampling of k-space is shaped by the trajectory design. Different undersampling patterns form different aliasing patterns in the image domain. Figure 2.6 is a flow diagram of two different 2D undersampling examples. A fully sampled image is Fourier transformed into k-space, sub-sampled by either a Cartesian trajectory or a variable density spiral trajectory and then transformed back. This is analogous to an MR acquisition that uses one of the example trajectories to sample k-space directly. The missing k-space information manifests in the image as aliasing which is predictable in the Cartesian case and less so in the spiral case. The points marked a_1 and a_2 in figure 2.6(c) represent two places in the image where the nose and cerebral peduncle overlap.

Some additional information inherent in sampled MR data is imposed by the measuring instruments themselves, the coils. This is depicted in figure

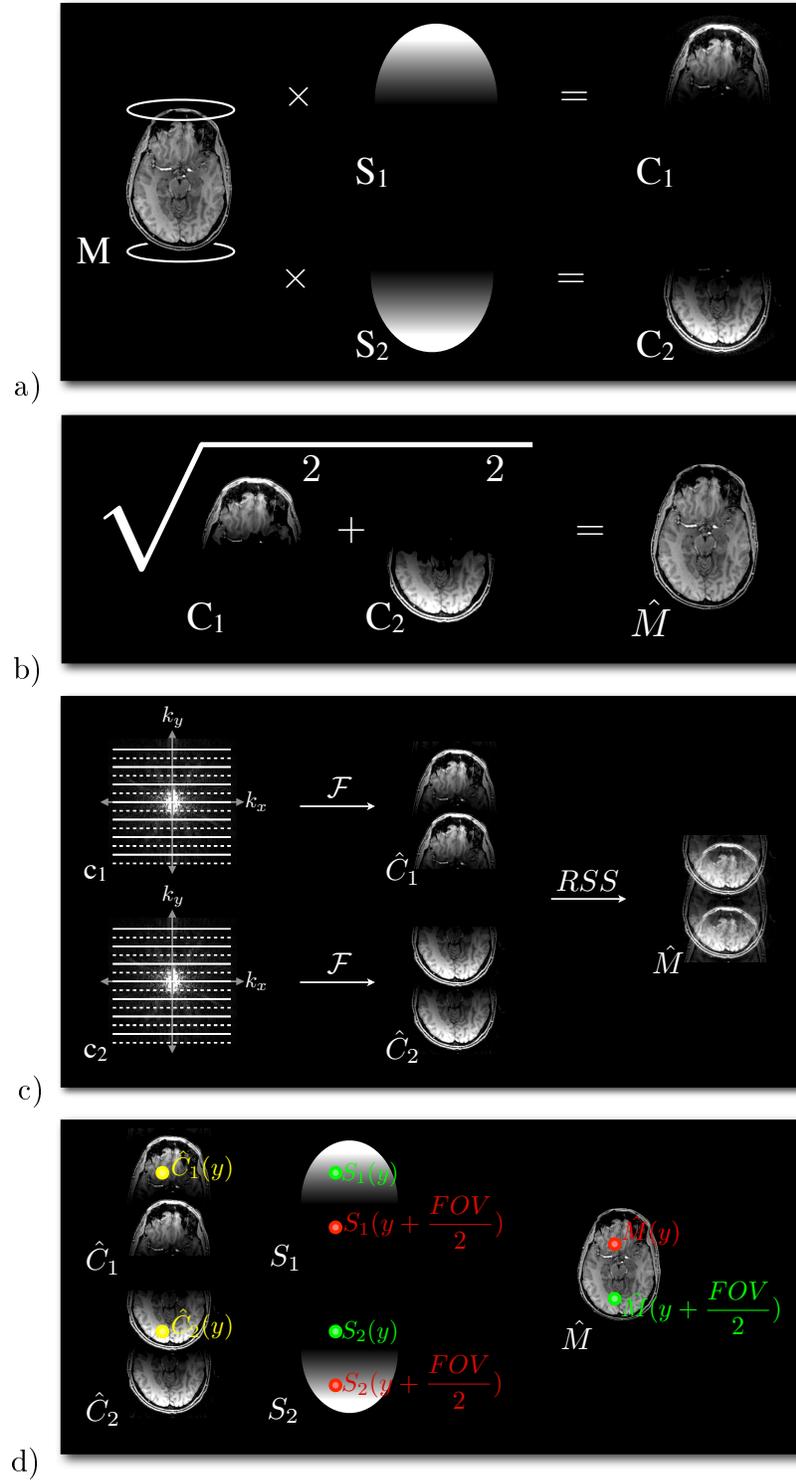


FIG. 2.7: Coil combination and SENSE parallel imaging reconstructions. a) coil weighted images. b) simple RSS coil combination. c) coil weighted and undersampled data. d) Cartesian SENSE parallel imaging reconstruction.

2.7(a). Two coils, placed on either end of a spin system (M) are used to sample k-space information. The coils are most sensitive at close proximity and loose sensitivity to signals originating further away. The coil sensitivity (S) for each coil is represented by the greyscale ellipses S_1 and S_2 . If the k-space data are fully sampled, the reconstructed images made from each set of sampled data results in the coil weighted images C_1 and C_2 .

The reconstruction of fully sampled coil weighted data is performed by combining each reconstructed coil image in a root sum of squares combination. Figure 2.7(b) shows how the two coil images complement each other providing more information where each coil is most sensitive. Since the sensitivity is spatially dependent, the SNR is also spatially dependent. Figure 2.7(c) shows how undersampled data would appear if reconstructed using the basic coil combination of figure 2.7(b).

SENSE (sensitivity encoding) (28) reconstructs the missing k-space information in undersampled acquisitions by using coil sensitivity information which is a fundamental component of the sampled data. This extra information can be thought of as an extra encoding layer produced by the coil sensitivity. In the Cartesian case, the spatial aliasing can be easily determined allowing a system of linear equations to be written for the reconstruction of \hat{M} . The reconstruction of two sample points in \hat{M} can be generated by using the corresponding point from each undersampled coil image and two points from each of the coil sensitivity maps. These points are identified in figure 2.7(d) and are reconstructed using following equations.

$$\hat{M}(y) = S_2(y + \frac{FOV}{2}) \cdot \hat{C}_1(y) - S_1(y + \frac{FOV}{2}) \cdot \hat{C}_2(y) \quad (2.26)$$

$$\hat{M}(y + \frac{FOV}{2}) = S_1(y) \cdot \hat{C}_2(y) - S_2(y) \cdot \hat{C}_1(y) \quad (2.27)$$

Two subtleties of this example are 1) the coil sensitivity maps are not aliased, and 2) the undersampling factor of 2 (in the y direction) relates each pair of points that are half of the field of view apart. Since the spatial sensitivity is slowly varying, the sensitivity maps can be generated using an additional low resolution acquisition or extra k-space samples in the imaging acquisition to fully sample the low spatial frequencies. In Cartesian imaging the undersampling factor corresponds to the distance between points of aliasing in the image.

In the case of a spiral trajectory, the undersampling does not occur uniformly in each direction. Figure 2.6(d) shows how a variable density spiral increases in the amount of undersampling, radially, from the center of k-space. The aliasing pattern in figure 2.6(e) is less coherent than that created by the undersampled Cartesian trajectory because it is distributed angularly through the image. Reconstruction for this trajectory is not as straight forward because a system of linear relationships is not as easily created. The dotted circle shown in the spiral trajectory of figure 2.6(d) indicates an area of critical sampling that can be built into an undersampled spiral scan. Low resolution sensitivity maps can be generated from this k-space data without the need for additional acquisitions. These pieces of information can be used to constrain or condition an iterative reconstruction in order to calculate the undersampled k-space regions (14).

3 SPIRAL PROJECTION IMAGING PCMRA

Phase contrast magnetic resonance angiography (PCMRA) is an MR modality that is capable of producing quantitative flow velocity information. The velocity magnitude is encoded as a phase within a complex valued image. Multiple complex images are acquired with velocity sensitivity in each of the principle (x, y and z) directions (20). A velocity vector can then be constructed using the magnitude and directional information contained in these data sets. Since the construction of a velocity vector requires multiple acquisitions (at least one in each principle axis, three for 3D), the acquisition time of the base imaging sequence more than triples. This increased scan time is a considerable disadvantage especially where large fields of view are needed (e.g. neuroangiography).

The work presented here is focused on the construction of a scan and reconstruction technique that can provide full volume 3D neuroangiographic images within clinically viable scan durations (1, 2, 6, 29). The proposed method is a combination of a rapid 3D imaging technique with a high gradient moment velocity encoding strategy. The practical requirement of reducing the scan time provides the technical challenges of maintaining a minimum signal to noise ratio (SNR) and spatial resolution necessary for providing phase contrast.

3.1 Theory

This section introduces the basic imaging and high moment encoding techniques combined in this work. The first subsection covers the spiral projection imaging technique which is followed by the introduction to a well known three-point moment encoding method.

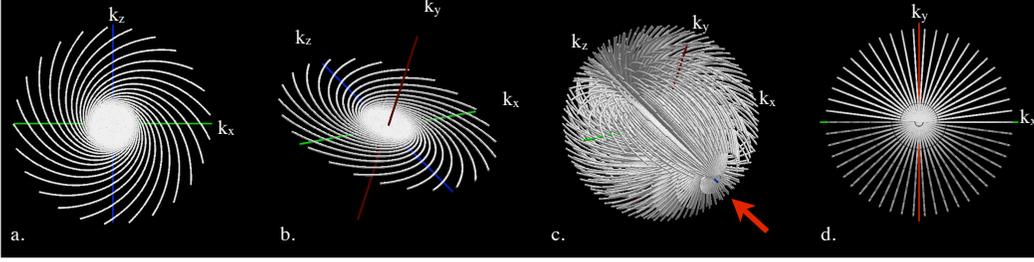


FIG. 3.1: A spiral projection imaging trajectory diagram. A 2D plane of spiral interleaves (a) in k_x, k_z space (b) is successively rotated about the k_z axis (c), until the set of spiral planes fills a sphere in k-space (the red arrow shows the axis of plane rotation). Data may be undersampled angularly by collecting fewer planes (d), or radially by altering the base spiral pattern (a).

3.1.1 Spiral Projection Imaging

Spiral projection imaging (SPI) (11, 12, 30, 31) is a 3D k-space trajectory that is generated using multiple 2D spiral sampled planes. The benefits of SPI include those that are intrinsic to spiral scans such as, reduced sensitivity to bulk motion and highly configurable variable density undersampling schemes. This section introduces the basic construction and configuration of SPI.

A single spiral projection is shown in figure 3.1(a). In an SPI acquisition, successive projections are collected at different orientations as shown in figure 3.1(b). By collecting the projections at even angular distances, rotated about a single axis, 3D k-space is spherically supported by the globe shape shown in figure 3.1(c). This method of plane ordering has dubbed ‘uni-axial’ since the planes are rotated about one axis only.

Three dimensional undersampling of a uni-axial SPI sphere can be accomplished by lowering the number of planes collected, or by undersampling of the spirals within each projection. The latter will be referred to as in-plane undersampling and is accomplished here by changing the ratio of angular to radial progression, which is called variable density spiraling (32) (figure 3.1(a)).

Large strides in scan time reductions can be made through uniform undersampling strategies such as this. This is simply due to the fact that fewer data points are collected meaning less k-space has been traversed. While undersampling leads to aliasing, uniform variable density undersampling incoherently reproduces aliasing allowing the underlying image structure to show through. The level of tolerable aliasing provides a bound on the amount of undersampling allowed, which in turn bounds the level of scan time reduction.

Spiral trajectories, in general, are efficient in terms of echo time and imaging time. The spirals used in this work are center-out trajectories that sample the low k-space frequencies first and gradually moves out to the high spatial frequencies, maximizing the amplitude of the imaging gradients. The short echo time is due to the lack of a prephasing gradient (as needed in Cartesian imaging) which also reduces the moment related phase accrual acquired when sampling central k-space after a prephasing gradient. The redundant low spatial frequency sampling, provided by the multiple spirals collected in-plane and for each plane, reduces the sensitivity to bulk motion because the motion elements are averaged out when all trajectories are combined.

Spiral trajectories are more efficient as the sampling duration (sampling window) within each repetition becomes longer. Fewer spirals are needed the longer the sampling window is, however, longer windows make the acquisition more susceptible to the dephasing from off-resonant spins. Accrued off-resonant phase in spiral imaging causes spatial blurring in areas of off-resonance (33). The most susceptible areas are identified and discussed in appendix A. This effect imposes an upper bound on the spiral length and therefore the level of scan time reduction.

3.1.1.1 Data Conditioning

Since spiral trajectories are highly oversampled in the central region of k-space. The sample density accumulates near the center of k-space as a result of using a constant sampling rate in time. High sampling density becomes particularly cumbersome for the reconstruction of 3D trajectories such as Spiral Projection Imaging (11, 12) (figure 3.1), where each plane contributes to the density of a sphere about the center of k-space (k_0). The gridding computation time increases linearly with the number of points and the sample density correction (SDC) (9) computation time increases quadratically with the density of the number of points. Computation time also increases steadily for time series reconstructions of dynamic data which requires multiple applications of gridding and SDC. This work presents two conditioning methods for improving reconstruction time. The proposed methods are referred to as ‘separable gridding’ and ‘data thinning’. Each method lowers the sample density (mostly near k_0) resulting in reduced computation time of the reconstruction, particularly in SDC.

The data thinning method reduces the number computations required in the 3D reconstruction by throwing out points that are sub-Nyquist distances apart, in k-space. The separable gridding method reduces the number of points by performing a 1D sampled density correction and gridding to each spiral arm individually. This reduces the number of points required in the 3D reconstruction and maintains the SNR by retaining all of the sampled points.

3.1.2 Dual VENC Velocity Encoding

As mentioned in the introduction, the phase contrast technique requires multiple images each containing different directional velocity sensitivity in order

to reconstruct a complete velocity vector at each voxel. The encoded phase is proportional to the velocity magnitude (parallel to the sensitivity direction), but contains no absolute relativity. For this reason it is common to collect a reference set with no velocity sensitivity, in order to achieve this phase resolution. This method is referred to as referencing (34, 35). At least one reference is required to resolve the phase for multiple velocity encoded images. This makes the total scan time equal to 4 times the base imaging time to acquire 3D velocity data.

In order to maintain a high velocity to noise ratio (VNR) the scan must be designed with equation 2.25 in mind. Considering imaging techniques alone, the VNR in equation 2.25 is improved through increases in the base image SNR, which provides gains proportional to the square root the scan time. In terms of velocity encoding alone, the VNR is increased as the VENC is decreased, which is the same as increasing the gradient strength of the velocity encoding gradients as shown by equations 2.23 and 2.24. This provides VNR increases directly proportional to the VENC. However, as discussed in subsection 2.2, if the flow velocity is higher than the VENC, the phase will be aliased resulting in an erroneous velocity estimate. Dual-VENC encoding is a scan and reconstruction technique that address the phase aliasing concomitant in low VENC scans in order to maximize VNR (4). These VNR gains are shown to be more efficient in terms of scan time, because of the direct proportionality with VENC.

While using multiple VENC is not the only method of phase unaliasing (36–38), it is one of the few that can accomplish unaliasing on a pixel-by-pixel basis. This is preferable since it avoids smoothing and is automatic requiring no user input for seed points or segmentation.

In a neurovasculature flow system, the blood velocity can vary greatly (potentially upwards of 100cm/s). Since the VENC level is driven by the desired VNR, the necessary VENC will usually be lower than the maximum velocity within the flow system. As a consequence of equation 2.23, the encoded phase will exceed 180° making the measured phase indistinguishable from flow in the opposing direction (aliasing). In order to benefit from low VENC acquisitions, this phase aliasing must be resolved.

The dual-VENC (three-point) method proposed in (4), solves the phase aliasing by collecting an additional high-VENC acquisition (which is not aliased) which is used to unalias the low-VENC acquisition. The current method requires the acquisition of seven SPI volumes, one for a baseline measurement, three at a low VENC ($20\frac{\text{cm}}{\text{s}}$) in each principle direction, and three at a high VENC ($100\frac{\text{cm}}{\text{s}}$) in each principle direction. The added acquisitions increase the scan time from 4 times the base imaging technique to 7 times. This initially appears to be less scan time efficient, however, it has been shown to be effective at providing a VNR gain of 4 times.

The dual-VENC unaliasing procedure is as follows. In a typical scan where the max flow velocity is just under $100\frac{\text{cm}}{\text{s}}$ the VENC would have to be set at least this high. In this example, a dual VENC encoded scan, with a low VENC of $20\frac{\text{cm}}{\text{s}}$, would retain a VNR increase by a factor of

$$\frac{VENC_{high}}{VENC_{low}} = \frac{100}{20} = 5 . \quad (3.1)$$

The phase aliasing is dealt with by using the high VENC data to ‘unwrap’ the phase of the aliased low VENC data. Finding the number of wraps depends on the relative VENC. The maximum number of wraps ‘k’ that can be resolved

are

$$k_{max} = NINT\left(\frac{VENC_{high}}{2VENC_{low}}\right) \quad (3.2)$$

where $NINT$ represents the nearest integer function and ‘k’ is always an integer because the VENC represents 180° and phase aliasing occurs first at 180° and then at every successive 360°. The number of wraps for any given phase measurement is

$$k = round\left(\frac{V_{high} - V_{low}}{2VENC_{low}}\right), \quad (3.3)$$

where ‘ $V_{high/low}$ ’ are the measured velocities. Applying the wraps to make the reconstructed velocity ‘V’ is then

$$V = V_{low} + 2VENC_{low}k. \quad (3.4)$$

After the low VENC set is unwrapped, the high VENC set is discarded since its comparative SNR is low provides a negligible benefit if $VENC_{low} \ll VENC_{high}$.

3.1.2.1 Dual Low VENC

A new method to phase unalias is also pursued. The motivation for collecting two low VENC, each within a close VENC level to the other can benefit not only from the increase in gradient moment, but provide better averaging capabilities as well. The theory behind the implementation of this method is covered in this subsection.

A high and low VENC set are still required in order to resolve the level of phase aliasing. The relative VENC are set to the following relation:

$$V_{max} = \frac{1}{\frac{1}{VENC_{low}} - \frac{1}{VENC_{high}}} \quad (3.5)$$

where ‘ V_{max} ’ is the expected maximum velocity to be measured within the flow system.

The unaliased velocity is then the difference between the two measured phases for each VENC, phi_{low} and phi_{high} for the low and high VENC respectively. This difference is then multiplied by the effective VENC (V_{max}) to generate the velocity estimate (\hat{v}) where

$$\hat{v} = (\phi_{low} - \phi_{high}) \cdot \frac{V_{max}}{\pi} \quad (3.6)$$

The phase difference is somewhat counter intuitive, but ϕ_{high} is subtracted from ϕ_{low} since the low VENC set incurs the larger phase wraps of the two. This subtraction is performed using a complex conjugate multiplication.

Since the high and low VENC are very similar in this case, the relative contribution is nearly one to one yielding an added SNR benefit of $\sim \sqrt{2}$ in addition to the low gradient moment.

3.1.2.2 Composite Reconstruction

In order to mitigate the signal biased phase present in the low VENC set, data from the high VENC set can be combined with low VENC data to form a composite image (3). This composite is mostly comprised of the low VENC measurements, and where signal loss is highest, the measurements are replaced by a weighted average based on the level of signal loss. The composite velocity is calculated as,

$$V = W_1 V_{high} + W_2 (V_{low} + 2VENC_{low}k) \quad (3.7)$$

where W_1 and W_2 are the relative matched weights based on the level of signal loss.

3.2 Methods

The following subsections cover the experiments used to analyse each of the proposed methods.

3.2.1 SPI Acquisition & Reconstruction

Undersampling the SPI trajectory can be performed in-plane or through-plane. In a variable density SPI globe shown in figure 3.1(c) the maximal undersampling occurs at the ‘equator’, where the combined through-plane and in-plane sampling is the thinnest. An undersampling factor of 4 in-plane and 4 through-plane yields a maximum undersampling factor approximately 16, at the equator edge. Currently the in-plane and through-plane factors are made equal, although it is unknown which direction is most effective at sustaining a reduction while minimizing aliasing.

The undersample factor is currently dependent upon the sampling time. Although it has been shown that lengthening the sampling time increases the scan efficiency, it also makes the scan more susceptible to field inhomogeneity artifacts. Therefore the undersampling factor is set after the maximum allowable sampling duration is determined. The sampling duration typically used is no longer than *7msec*. This corresponds to a linear undersampled, in-plane, spiral shown in figure 3.1(a). This figure shows that the spiral is critically sampled from the center of k-space to a specific radius (in this case 28% of $1/resolution$). Past the critical sampling radius, the sampling rate linearly decreases (between spirals) at a rate of $\frac{4 \cdot resolution}{FOV}$.

After the sampling duration is determined, the undersampling factor is

then determined by the desired resolution. If the total scan time is fixed, then the resolution is increased at the cost of SNR and aliasing. If the resolution is increased then the scan time is used to spread the sample density thin, reaching further out into k-space. However, to make up for the resulting aliasing, point spread function (PSF) is widened which has a blurring affect on the final image. If the resolution is lowered, then the scan time can be used to fill in more of k-space within a smaller radius. This reduces aliasing, but may fail to capture the signal of fine detail such as smaller vessels.

The reconstruction of the high and low resolutions scans can be characterized by their respective transfer function. The magnitude transfer functions (MTF) shown in figure 3.2(c) and (f) are representations of the filter applied to each spiral plane. This weighting scheme is found using the sample density compensation method found in references (9, 10). The high resolution MTF is shaped like a low pass filter, variably suppressing the contribution of aliased high spatial frequencies in the final image. This variable filter is matched with the density of the trajectory used in 3.2(b) and so the support changes radially from the center of acquired k-space.

The low resolution MTF, shown in figure 3.2(f) is an all-pass filter that directly passes all sampled spatial frequencies because the density is within the Nyquist limits. The collected matrix is fully supported and requires no filtering to lower the energy of spatial aliasing.

The generic MRI reconstruction pipeline has been automated using the Python interpreted language as a backbone to the individual algorithms implemented in C/C++, where algorithmic separability is implemented as parallel computational procedures using the POSIX threading library. The base reconstruction method is gridding and phase array coil combination as proposed

in references (20, 39, 40)

3.2.1.1 Data Conditioning

Separable gridding starts with 1D gridding of the samples along each individual spiral arm in the arc length (or time) direction to maintain critical sampling while averaging points in dense areas. The uniformly sampled arms are then reconstructed with the original pipeline by 3D gridding and SDC. Data thinning is accomplished by dropping samples that exceed the Nyquist sampling rate, in arc length, along the spiral trajectory.

The noise properties of each method were determined through Monte-Carlo simulation of multiple sets of Gaussian k-space noise data. The SPI parameters used in the simulation were: 24cm FOV, 240 diameter matrix, 20 spiral leaves, and 377 planes.

3.2.2 Dual VENC Velocity Encoding

3.2.2.1 Dual Low VENC

The Dual Low VENC simulations were run to ensure the method was robust to changes in the relative phase relationship between the high and low VENC measurements. Specifically, for changes due to phase bias caused by signal averaging.

A one dimensional parabolic phase profile was simulated at a resolution of 4096 points, 2048 across the flow lumen. In order to simulate signal averaging, a low-pass k-space filter was applied reducing the resolution to 409 points, 204 points across the lumen. The k-space matrix was left at the original length, which effectively zero-padded the low resolution simulation by a factor of 10.

The simulated flow was set to a maximum velocity of $5\text{cm}/s$, $10\text{cm}/s$ and $20\text{cm}/s$ for a high and low VENC of $10\text{cm}/s$ and $8.89\text{cm}/s$ respectively. The

high and low VENC combination was chosen to resolve up to an 80cm/s maximum velocity (equation 3.5). The phase profiles were combined with a constant signal magnitude profile over the lumen and wall sections. All points were reconstructed using equation 3.6.

3.3 Results & Discussion

The following subsections cover the experimental results and the viability of each method as a component in the final rapid 3D PCMRA technique.

3.3.1 SPI Acquisition & Reconstruction

Figure 3.2 shows two sagittal PCMRA MIPs of (a) high acquisition resolution (0.8mm) and (d) low acquisition resolution (1.3mm). The scans were both 6min in duration. The sample density correction kernel size was smaller for the high res acquisition set than the low res acquisition set. The smaller kernel size has the effect of blurring the resolution out to 1.3mm , while the larger kernel size is fit to maintain the prescribed resolution of the low res acquisition. Figure 3.2(a) shows more small vessel conspicuity than figure (d) even though the SNR appears is higher and the aliasing energy is lower in (d). The low resolution scan also shows narrow vessels in some regions which may be due to intra-voxel dephasing caused by high velocity gradients or more significant partial volume effect. Insufficient suppression of the undersampled high resolution k-space data is responsible for the diffuse clouded look of the high res image that contributes to a low SNR quality of the image and hides many of the small vessels visible in the low res image.

3.3.1.1 Data Conditioning

Figure 3.3(a) shows the densities for each method along a line orthogonal to the axis of plane rotation for the simulated fully sampled SPI sets. Both

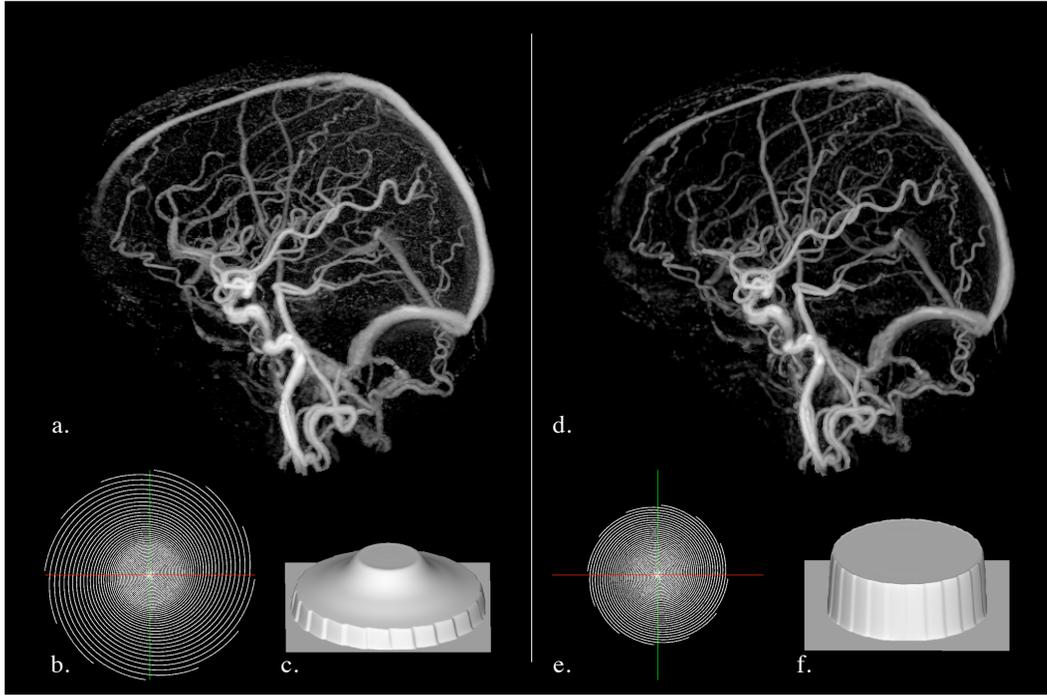


FIG. 3.2: The effects of MTF modulation for two different trajectory configurations. (a) a reconstructed PCMRA from a high resolution scan with a high level of undersampling (scan resolution of $0.8mm^3$ with a reconstructed resolution of $\sim 1.3mm^3$). (b) the relative k-space coverage for the high resolution scan. (c) a cross section of the MTF for high resolution reconstruction (low pass filter). (d) a reconstructed PCMRA from a low resolution data acquisition (scan resolution of $1.3mm^3$). (e) the k-space coverage for the low resolution scan (relative to the high resolution scan in in (b)). (f) the MTF for the low resolution reconstruction.

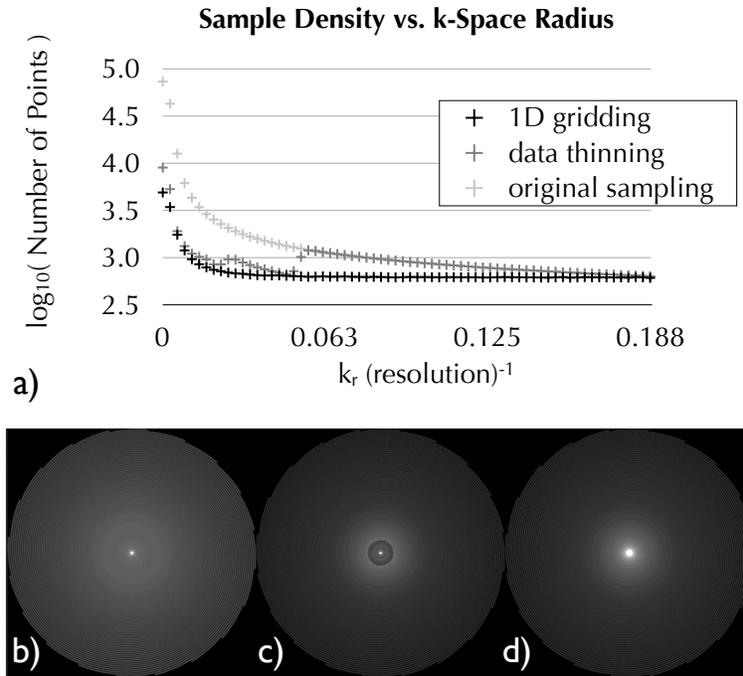


FIG. 3.3: SPI radial k-space density. (a) sample densities versus radial distance from k_0 , on the axis normal to plane rotation, for each presented method using a 3D uni-axial SPI trajectory with 20 Archimedean interleaves per plane, 377 planes, $FOV = 24\text{cm}$ and 1mm^3 resolution. (b-d) sample densities across a 2D spiral plane for 1D gridding, data thinning, and the original sampling density, respectively.

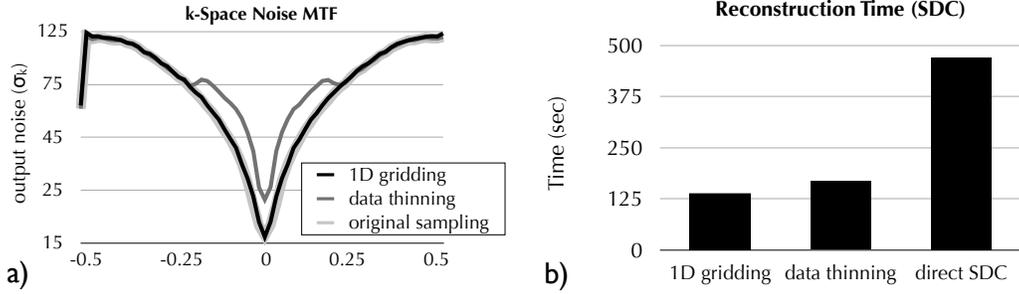


FIG. 3.4: Properties of 1D spiral arm gridding in SPI. (a) a Monte-Carlo simulation (20,000 points, 643 image volume with $FOV = 24cm$, 20 interleaves, 101 planes) showing the standard deviation of noise reproduced across the frequency domain given an input noise of $\sigma_k = 1,000$. In image space, 1D grid noise is 0.06% greater than the original reconstruction and thinning noise is 0.5% greater. (b) the reconstruction times for each method were based on the compute time of the SDC. The compute time scales with the square of the sample density, and number of iterations.

the original sampling and separable gridding provide smooth MTFs which minimize banding in the reconstructed image.

Figure 3.4(a) shows the results of a Monte-Carlo simulation. Separable gridding reproduces noise to a level similar to the original reconstruction process. Image space noise is increased by 0.06% with respect to the original scheme (Figure 3.4(a)). Thinning increases the noise level by 0.5% and changes the spectral noise profile. Figure 3.4(b) shows that for 1 iteration, 3D SDC has shortest compute time with separable gridding, followed by data thinning, which are both over 3 times faster than direct 3D SDC. The added computation of 1D gridding is $\approx 30sec$ on a single core of an Intel Quad Core Duo 3GHz, however, the method is scalable on a multi-CPU system. The sample density generated by thinning causes discontinuities in sample density (Figure 3.3(a)) wherever the sample distance is not evenly divisible by the Nyquist distance. This appears as concentric rings in the spiral plane shown in figure

3.3(c).

The artifact imparted by each of the proposed methods was determined to be unreasonably conspicuous to be used in the PCMRA reconstruction. While the tradeoff in reconstruction time is high, the relative SNR difference is negligible. The data thinning method produces banding which appears as radially distributed interference. The separable gridding method produces an apodization in the final image that has yet to be determined. For these reasons the direct SDC method was used for the final reconstruction.

3.3.2 Dual VENC Velocity Encoding

The signal loss has an added affect on the velocity encoded phase. In areas where the spatial change in velocity is significant, relative to the resolution, the resulting measured phase will be biased by the signal lost to this high velocity gradient. This gradient moment induced dephasing lowers the signal in the immediate vicinity of of the high flow gradient, which corrupts the phase measurement by unevenly averaging velocities within the intra-voxel distribution. The corresponding signal loss associate with these areas can be mapped, as shown in figure 3.5, by taking the difference, in magnitude, between the encoded image and the reference image (non-encoded).

As shown in figure 3.5, the losses due to dephasing are mitigated in the high VENC image. This is due to the low moments used for higher VENC imaging. Figure 3.5 also shows how these corrupted areas mostly occur in vessels of high flow which are adequately characterized in the high VENC data. The unbiased information in these regions are used to correct the areas that cannot be unaliased using equation 3.3. The new unaliased velocity is a composite of the high and low VENC data using the method proposed in (3).

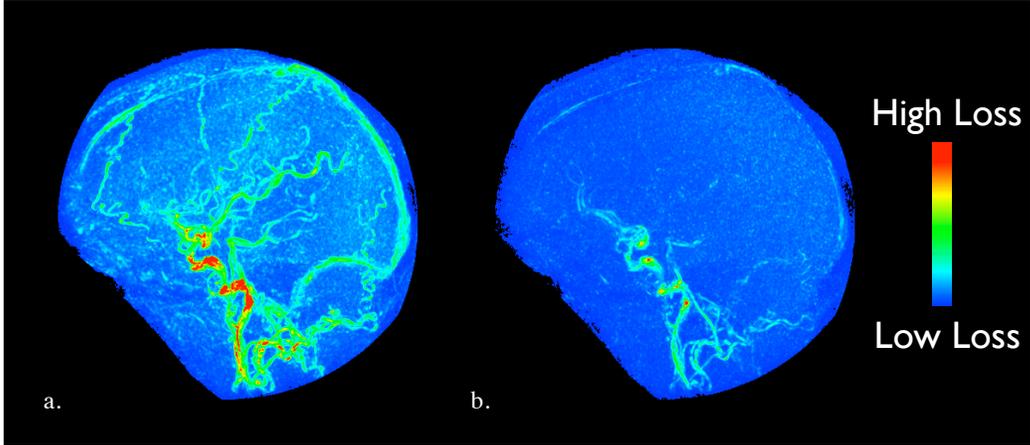


FIG. 3.5: Signal loss maximum intensity projections for low VENC (a) and high VENC (b) acquisitions. The high losses are colorized as red and low losses are colorized as blue as shown in the spectrum on the right. The VENC levels used are $20 \frac{cm}{s}$ and $100 \frac{cm}{s}$ for (a) and (b) respectively.

3.3.2.1 Composite Reconstruction

Figure 3.6 shows the added affect of using the signal loss information to make a composite image. The high velocity gradients usually occur at the vessel boundaries causing a discretized vessel edge due to erroneous unaliasing or signal bias, shown in figure 3.6(a). These edges are corrected at the cost of VNR by preferentially weighting the high VENC data in these areas (figure 3.6(b)).

3.3.2.2 Dual Low VENC

The drawback of this method is the sensitivity to signal biased phase present in each of the measurements. The simulation shows that the method is highly sensitive, even in high resolution cases that are well beyond practical bounds. It can be seen from equation 3.5 that a small change in either of the VENC result in a relatively large change in V_{max} . This sensitivity is translated to the reconstruction of ϕ_{low} and ϕ_{high} . Additionally it is hypothesised that the

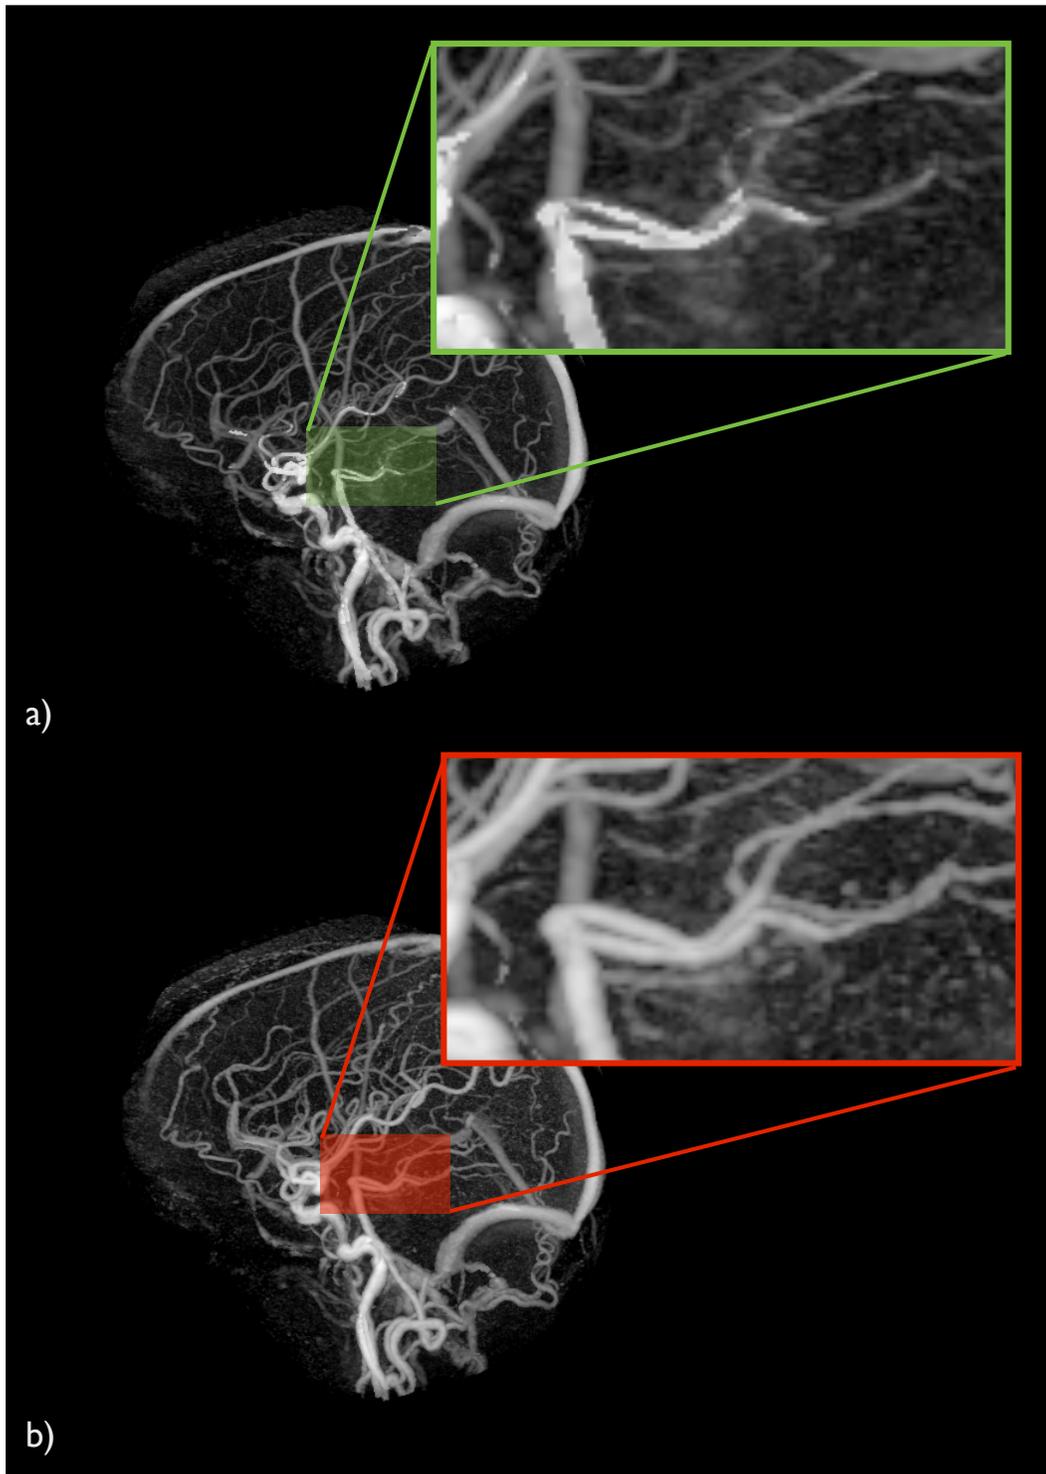


FIG. 3.6: Sagittal maximum intensity projections of dual-VENC reconstruction (a) compared to composite dual-VENC reconstruction (b).

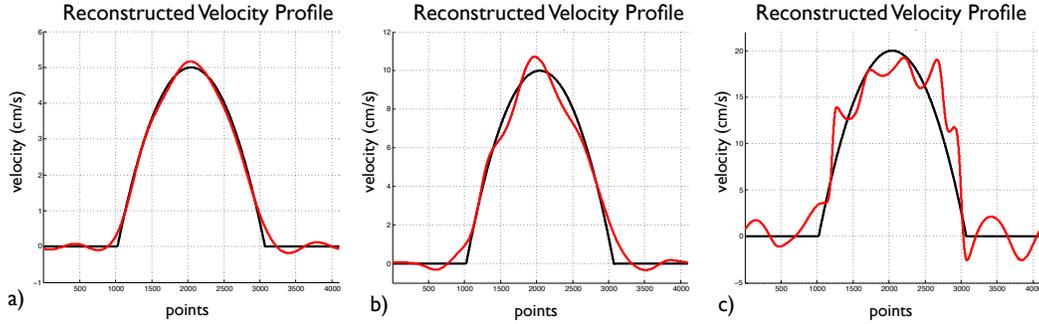


FIG. 3.7: 1D simulation profiles of the dual-low-VENC method. The parabolic profile apex is 5cm/s , 10cm/s , and 20cm/s for (a), (b) and (c) respectively.

method is more sensitive to bias than noise, because of the high bias correlation between the two measurements. Unlike the composite method presented in subsection 3.1.2.2, this dual low VENC method cannot be as easily corrected using signal loss information. These deficiencies were the motivation for continuing on with the dual-VENC method introduced in (3, 4).

3.4 Conclusion

Spiral projection imaging is demonstrated to be a suitable basis for rapid 3D PCMRA. The variable density and high uniformity of undersampling minimize coherent aliasing and allow large strides in scan time reduction.

MTF filtering through modulation of the sample density calculation allows for more efficient data acquisition. It has been demonstrated that high resolution k-space data can be acquired and filtered to produce low spatial aliasing and higher small vessel conspicuity. The SDC method is also highly configurable allowing full advantage to be taken of lower resolution, less undersampled data.

The dual-VENC method was shown to be robust to incoherent spatial aliasing. The addition of signal based composite imaging proved to be an

important component in the efficacy of phase unaliasing in the presence of high signal bias. The phase unaliasing errors increase as the VENC level is lowered. The increased signal loss at low VENC is mitigated as the resolution increases, despite the associated increase in incoherent spatial aliasing. This relationship becomes trivial with compositing since the high VENC data set produces negligible losses from either of these factors.

4 PRETTY EASY PARALLEL IMAGING

Parallel image reconstruction methods synthesize data to replace undersampled or non-sampled gaps in k-space. The SENSE parallel imaging algorithm presented in (14, 28) is generalized for the reconstruction of non-cartesian k-space trajectories through the use of a gridding/de-gridding step within the iteration loop. In an effort to reduce the number of computations in each iteration, a method of masking k-space, called PEPI (5, 7, 41), was introduced as a replacement for the gridding/de-gridding step. This work demonstrates the relative quality and computation time in the reconstruction of a uniformly undersampled 3D trajectory, FLORET (13), between the CG-SENSE and CG-PEPI methods at two different levels of undersampling.

4.1 Theory

Pretty easy parallel imaging (PEPI) was introduced as a 2D parallel imaging reconstruction method in (23), as a simple alternative to the complicated parallel imaging techniques available. As introduced, the focus of the PEPI technique is to simplify the gridding and de-gridding operations involved in the data consistency segment of the iterative SENSE reconstruction (14). Previous methods (7) have been proposed to replace these operations with a multiplication through the introduction of a masking technique. PEPI introduces additional theory to the mask generation technique which considers the sampling density of newly generated points along with points sampled along the imaging trajectory by making use of the extensible sample density correction algorithm proposed in (8, 9, 42).

Computationally, the gridding and de-gridding require $O(2 \cdot \frac{4}{3}\pi(L_k/2)^3 \cdot R_m^3 \cdot N)$ number of operations, where L_k is the diameter of the convolution kernel, R_m is the oversampling factor, and N is the number of trajectory points. The

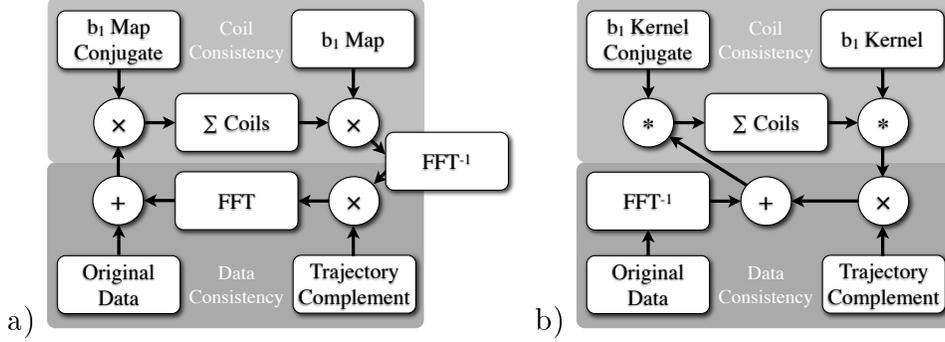


FIG. 4.1: PEPI reconstruction flow diagrams. (a) the Fast Fourier Transform based PEPI. (b) the convolution based PEPI.

factor of 2 corresponds to the grid and degrid pair. In contrast a multiplication requires $O(R_m^3 \cdot \frac{4}{3}\pi(M/2)^3)$ where M is the diameter of a cubic matrix. For 3D SPI data set undersampled by a factor of 3, the number of grid points within a sphere contained within a cubic matrix defined by M is approximately equal to the number of sampled trajectory points N . This leaves reduction factors in reconstruction time (for the data consistency segment) dependent on the size of the grid kernel L_k .

PEPI is an iterative reconstruction algorithm that focusses on two principles: 1) data consistency and 2) coil consistency. The reconstruction pipeline in figure 4.1(a) shows a coil consistency segment that is multiplication based using the speed of the FFT to effectively compute the convolution. Coil consistency is enforced by the process highlighted on the top half of the flow diagram. Starting with the ‘Original Data’ gridded and transformed (for each coil) into image space, the data are then multiplied by the conjugate coil sensitivity maps (‘ b_1 Map Conjugate’) to remove coil phase from the individual images. The coil images are then summed to produce a single image volume that is free of the coil sensitivity profile. The coil maps (‘ b_1 Map’) are then multiplied by the single coil combined image to reproduce individual coil im-

age volumes. This process has the effect of convolving, in k-space, the coil support with the sampled k-space data, which extends data from the sampled areas into non-sampled areas of k-space.

Data consistency is enforced by the process highlighted in the bottom half of figure 4.1(a). The new reproduced coils for the current iteration are Fourier transformed in to k-space. The new extended data from the non-sampled regions are kept by masking out the convolved data from the original sample locations. Figure 4.2 shows the sampled spiral trajectory (a) and the trajectory complement (b). The trajectory complement is used as the mask. After the original sampled locations are cleared, the originally sampled data are reinserted into those locations. This process continues iteratively until the process ceases to extend new, significant, data into non-sampled regions (usually in about ten iterations).

4.1.1 Convolution Based PEPI

The reconstruction pipeline in figure 4.1(b) shows a convolution based coil consistency segment that computes both data consistency and coil consistency in the k-space domain directly. The coil sensitivity profiles are converted into k-space convolution kernels that are small enough to maintain computational efficiency without the need for FFTs. The advantage of staying in k-space, and not transforming back and forth between domains, is the removal of Gibbs ringing from the iterative system. Ringing is mitigated in the FFT-PEPI method by increasing the grid size by a sampling factor R_m . Convolution-PEPI requires no such increase effectively removing R_m and adding a Cartesian convolution kernel to the computational complexity.

The coil sensitivity based convolution kernel starts with the same procedure

as the b_1 map generation outlined in (28). Image modulus is removed from the coil sensitivity and the sensitivity profile is smoothed to minimize the spectral extent required to characterize the profile in the k-space domain. The smoothing of the profile is key to minimizing the kernel size, and thus, the number of computations required in the convolution.

Additionally, the convolution is speed up by only choosing the principle kernel components. This has a unique advantage in 3D, since the level of coil support can be different in each direction, the kernels can be made more compact in directions with lower support by applying this prior knowledge (43).

The current methods of generating the coil sensitivity based convolution kernels fail to produce sizes that are computationally competitive with the FFT-PEPI method variant. However, the convolution based PEPI method is capable of achieving the same reduction factors as the FFT-PEPI counterpart with a relatively large kernel extent. Kernel generation methods such as those presented in (44) provide compact spectral support and may potentially be adaptable to this method in the future. Since the current state of this method variant does not provide gains in reduction factors over the FFT based method, despite its potential advantages, it was not further explored in this work.

4.1.2 Mask Generation

The novel contribution of PEPI to this type of reconstruction technique is in the generation of the trajectory complement. The complement is created through the use of an iterative convolution style sampling density compensation technique (9, 10). The points of interest are not only the sampled locations but the non-sampled locations that fall on the grid (the data to be synthesized

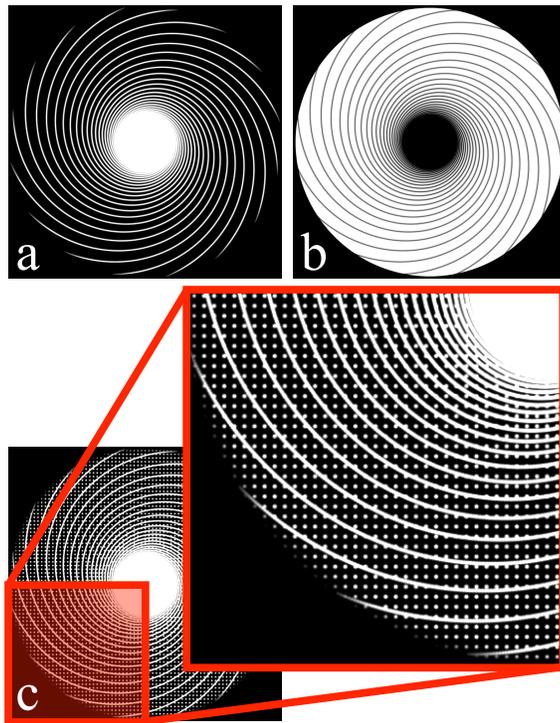


FIG. 4.2: The procedure for generating the MTF mask used in PEPI reconstruction. (a) a sample density corrected SPI cross section. (b) the cross section complement. (c) a close up of the density corrected SPI sample trajectory and cartesian spaced PEPI synthesized data. The complement is density corrected by considering sampled points on the trajectory (a) as well as points that will be synthesized during iterative PEPI reconstruction, which lie off of the trajectory (c).

in each PEPI iteration). The sampling density correction is calculated with both sets since the final reconstructed set will contain both.

Figure 4.2(a) shows an example of a non-Cartesian 2D cross section of a gridded 3D trajectory. The data consistency mask is essentially the complement of this gridded trajectory. However, the complement is generated by first evaluating the sample density of both the trajectory and, prospectively, the new data points to be synthesized by the parallel reconstruction process, together. The synthesized points are those that fall on the Cartesian grid between the trajectories covered by the acquisition sequence as depicted in figure 4.2(c). After the relative density of the combined set have been determined, the Cartesian points are gridded separately, forming the mask in 4.2(b).

4.1.3 Conjugate Gradient PEPI

Some of the difficulties in tuning the iterative PEPI method is in finding the relative scale between each set. The sampling density correction used to generate the initial condition, gridded originally sampled data, and the mask. This scales each data volume by an amount based on its relative level of undersampling. For this reason, and for the reductions in reconstruction time, the conjugate gradient (CG) minimization technique implemented in (7, 14), for non-Cartesian SENSE, was extended to PEPI.

The CG technique sends image residuals through the system over each iteration. For this reason the mask generation method is slightly altered to produce a complement of the non-CG mask. This is done after the combined sample density is determined. The weighted trajectory points are gridded as opposed to the Cartesian points, providing a mask that looks similar to 4.2(a).

An overview of the CG-PEPI method is compared to the CG-SENSE

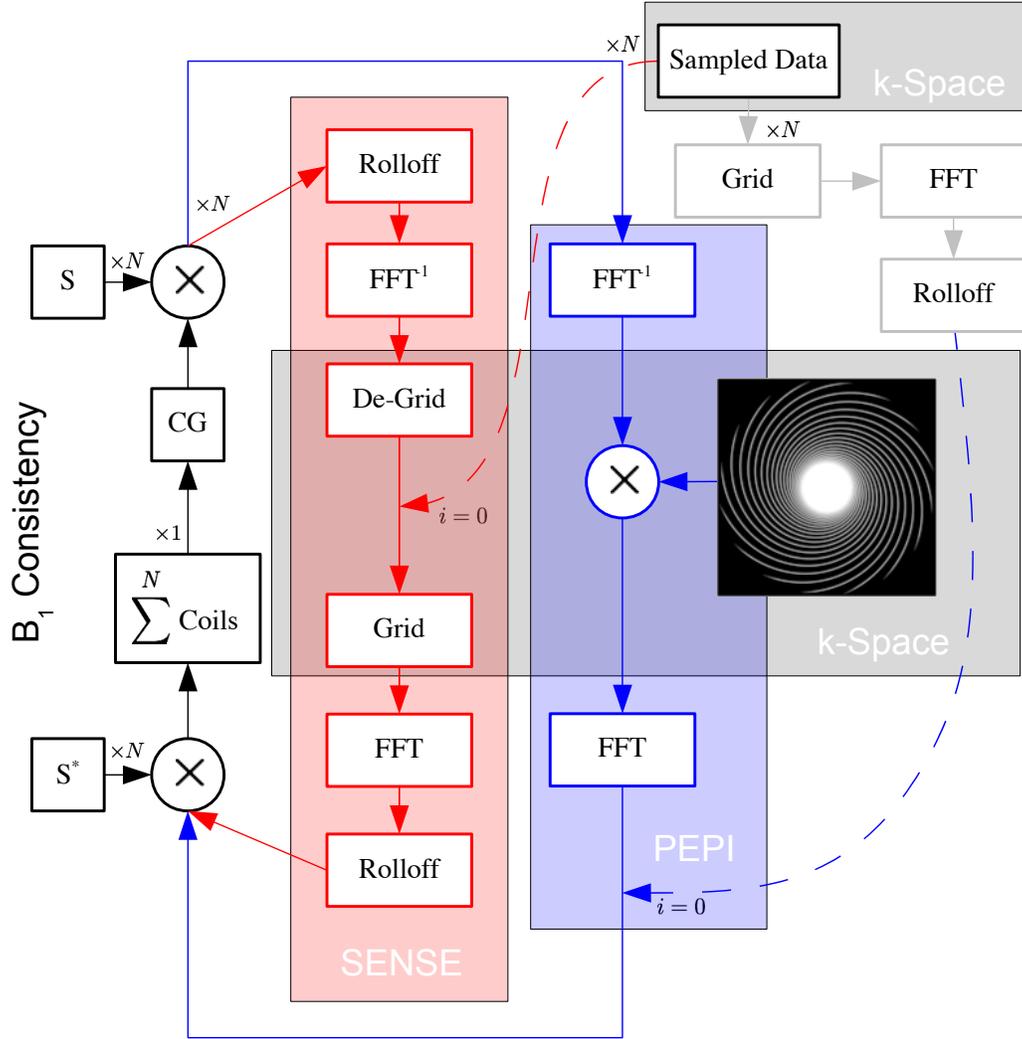


FIG. 4.3: Flow diagrams of the SENSE (red) and PEPI (blue) reconstruction pipelines. Each algorithm processes the k-space data differently as shown in the k-space box (grey). The k-space mask (W_c) is that of figure 4.2(b). The initial preparation ($i = 0$) for PEPI is a gridded reconstruction (upper right). B_1 consistency (left) for N coils is reinforced using coil sensitivity maps S and complex conjugate S^* .

method in figure 4.3. The boxes shaded in red and blue correspond to the portion of the algorithm that is different between the two SENSE and PEPI methods respectively. The diagram shows that the PEPI method requires only a multiplication in the data consistency portion of the loop where the SENSE method employs a relatively expensive grid and degrid step. The diagram also shows the initial preparation required for each method. For SENSE the sampled data can be used directly. For PEPI the sampled data must first be gridded, transformed, and deapodized. This loop may also be applied to the convolution based PEPI, however, additional FFTs must be applied before and after the CG logic.

4.2 Methods

This section covers the experiments performed for the FFT based PEPI method, the CG-PEPI method and the comparison between CG-PEPI and CG-SENSE.

4.2.1 FFT Based PEPI

The FFT based PEPI method was performed using an SPI based imaging technique (5, 11, 41). The central k-space was critically sampled in order to generate the coil sensitivity map. Undersampling was performed angularly by omitting planes. This produces undersampling in the axial plane where most of the coil variation exists for the 8-channel head coil used. The scan parameters used were as follows: FOV 24cm, 240 diameter matrix, 384 planes (fully sampled), 75 spiral arms, $TE/TR = 0.6/9.9ms$, total scan time of 5 minutes (fully sampled). The fully sampled set was scanned on a GE 3 Tesla Signa Excite System with a standard 8-channel head coil. Sub-sampling of the full set was performed by removing planes creating reduction factors of 4, 8 and 12 corresponding to 96, 48 and 32 planes respectively.

The reconstruction was performed using C coded algorithms with the POSIX threading library, on an 8-CPU 3GHz Linux platform.

4.2.2 CG-PEPI Simulations

The reduction factor limits of the FFT-based CG-PEPI method were determined through simulation of 3D coil sensitivity weighed data.

The base simulation 3D image data consisted of a 3D Cartesian SPGR scan of a GE resolution phantom. The scan produced a $1mm$ resolution image volume, that was thresholded, providing sharp, high resolution spatial transitions with little noise. A quadrature head coil was used to collect the data providing a single channel reconstructed data. The thresholding was also used to remove any residual coil weighting from the synthesized image.

A ball phantom was scanned using a combined head and neck coil array to acquire coil sensitivities without underlying image contrast. These coil sensitivities were applied to the synthesized base image before data simulation. A spherical crop, the size of the FOV, was applied to the base image before simulation to ensure the image was fully contained within the FOV.

Each coil-image combination was used to synthesize data for various levels of Cartesian undersampling using the direct Fourier transform. The three dimensional Cartesian trajectories uniformly undersampled in the following patterns: 1) in x by a factor of 2, 2) in x by a factor of 2 and in z by a factor of 1.5, and 3) in x and z by a factor of 2. Each set contained a supplemental critically sampled region in the center of k-space to simulate the central region of center-out trajectories. The radius of the centrally critically sampled region is on tenth of the extent of k-space in each principle direction. The simulated trajectories were created for a reconstruction matrix of 128 points in diameter.

4.2.3 CG-PEPI and CG-SENSE

The implemented 3D CG-PEPI and CG-SENSE methods were compared using a 3D spiral sampling technique called FLORET (45). A fully sampled, 4 times and 9 times undersampled neuro-imaging acquisitions were collected in-vivo using a standard 8-channel head coil. Each trajectory was created for an FOV of 24cm , 14.2msec sampling window and 240 diameter matrix. The $R = 4$ set has a little over twice the number of sample points as the $R = 9$ set. Scans were performed on a GE Signa Excite 3Tesla scanner.

Reconstruction benchmarks were performed using algorithms coded in C and the POSIX threading library on a 12-CPU 2.6GHz Linux platform.

4.3 Results and Discussion

This section discusses the efficacy of the PEPI parallel imaging method through in-vivo and simulated experimental results. The final subsection compares the PEPI method to a 3D implementation of the widely used SENSE method proposed in (14).

4.3.1 FFT Based PEPI

Figure 4.4 shows an array of axial cross sections of 3D PEPI reconstructions. The SPI imaging method was reconstructed using a basic root sum of squares coil combination (top row) which is compared to the PEPI reconstruction (bottom row). The RMS error shows the relative difference between each cross section and a cross section of the fully sampled set labeled as truth. The PEPI reconstruction significantly reduces the level of aliasing with negligible losses in resolution.

The reduction in aliasing is most noticeable in the sagittal cross sections shown in figure 4.5. The sagittal images show a significant improvement in the

cerebellar regions where the aliasing appears to be the most disruptive. Since the images fall outside of the FOV, signal pileup occurs at the edges of the FOV towards the base of the head. This pileup makes it difficult to compare the relative full volume aliasing reductions as it skews the error between the fully sampled reconstruction.

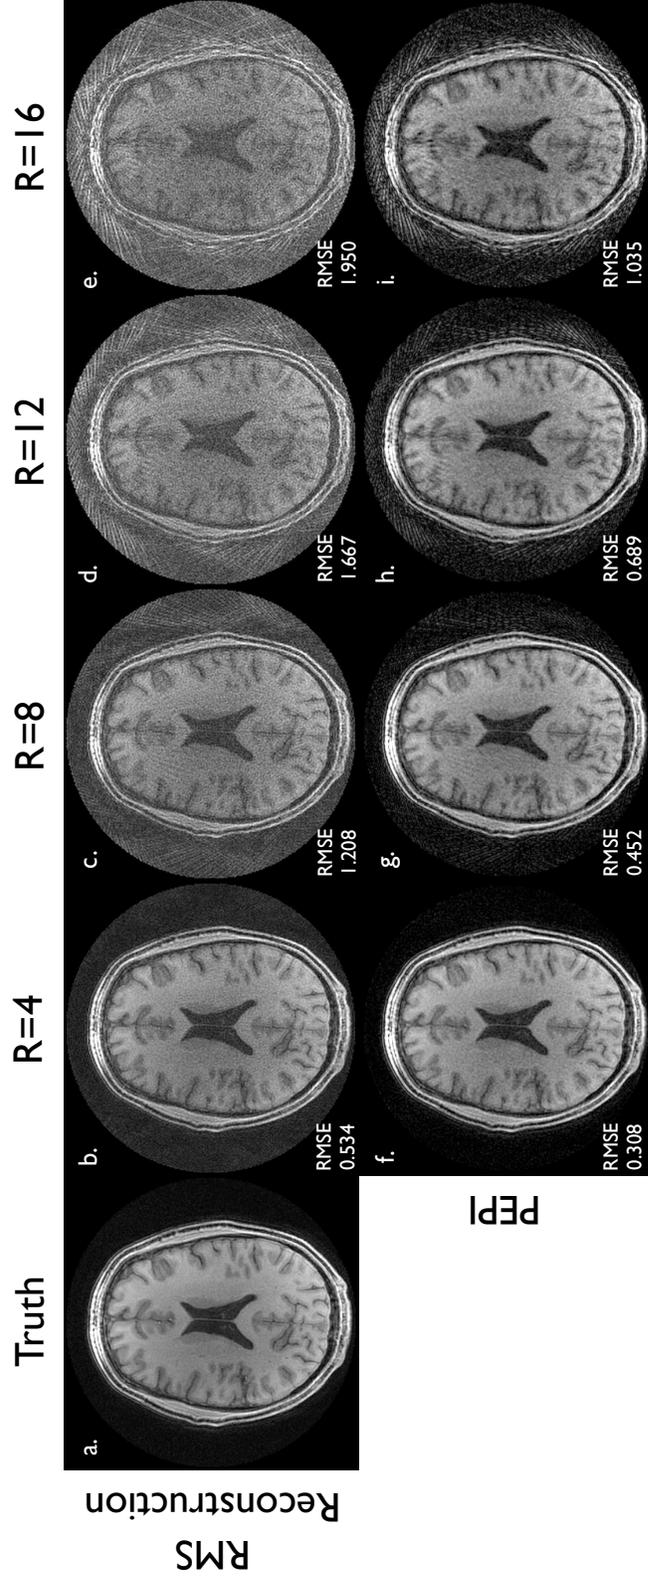


FIG. 4.4: Axial slices of reconstructed data. (a) a fully sampled set. (b-d) under-sampled sets. (e-g) PEPi reconstructed undersampled sets. Undersampling factors (R) pertain to the number of planes used: 384 (full), 96 (R=4), 48 (R=8), 32 (R=12). In SPI, the under-sample factor is the maximum Nyquist distance between two points.

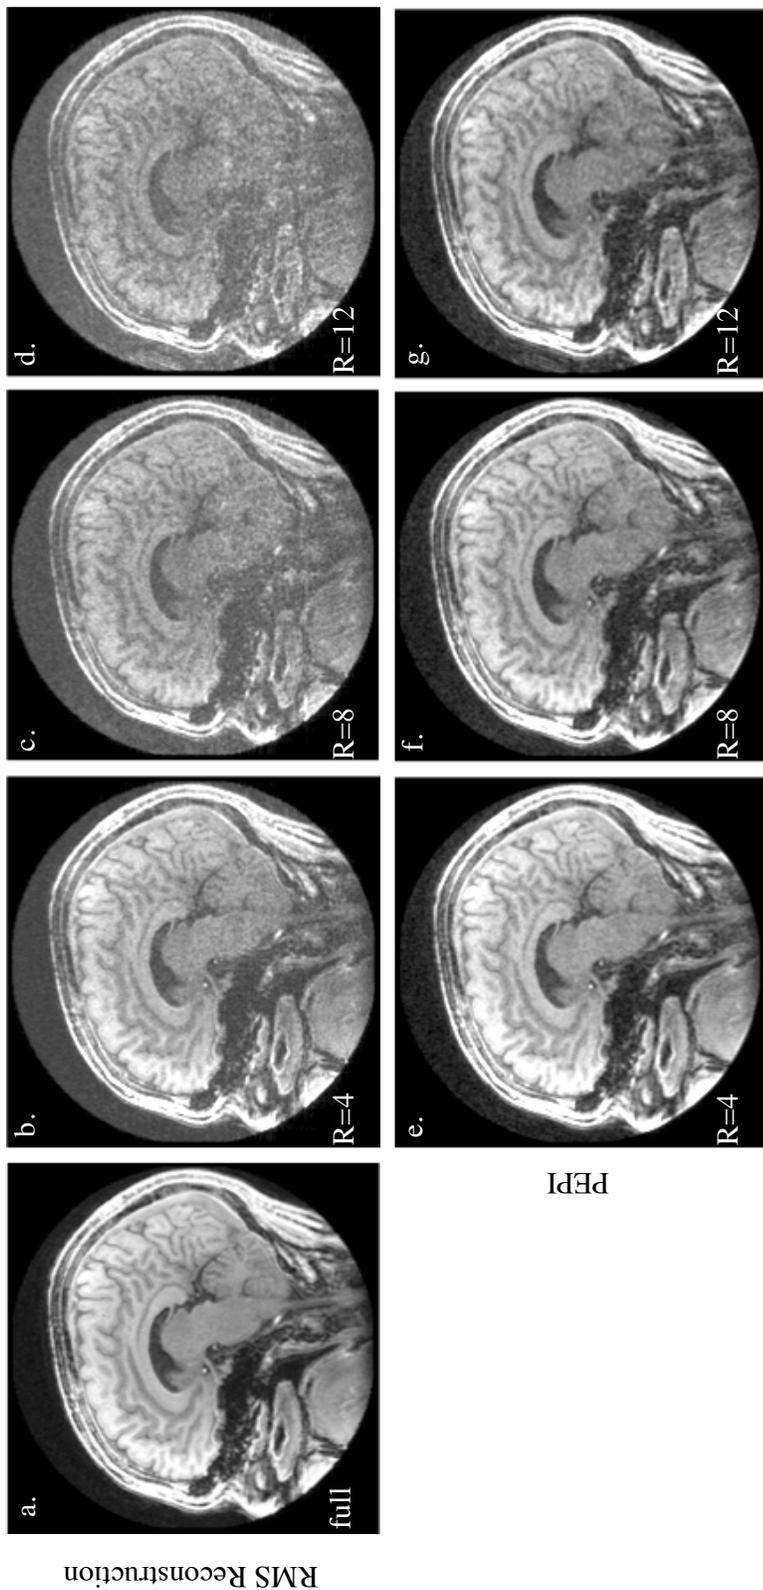


FIG. 4.5: Sagittal slices of reconstructed data. (a) a fully sampled set. (b-d) under-sampled sets. (e-g) PEPI reconstructed undersampled sets. Undersampling factors (R) pertain to the number of planes used: 384 (full), 96 ($R=4$), 48 ($R=8$), 32 ($R=12$). In SPI, the under-sample factor is the maximum Nyquist distance between two points.

4.3.2 CG-PEPI Simulations

Figure 4.6 shows axial and sagittal cross sections of the simulated sets using CG-PEPI reconstructions. Reconstructions at 1 iteration are shown for images (b), (c) and (d). Reconstructions that provided the minimum error where chosen at the corresponding iteration in images (e), (f) and (g). The high spatial frequency content generated by the sharp edges of the synthesized base image pushed the reconstruction to the limits. Here the weakness of the method are exploited for observation. The 1st iteration provides a sense of how much aliasing is initially present. In each simulation, undersampling by a factor of two in x provides a replicate image circularly shifted by half the FOV in the axial cross sections shown to the right of each pair. Undersampling in z provides more replications as shown in the sagittal cross section on the left of each pair. While the reconstruction method is mostly successfully at synthesizing k-space data for gaps in k-space that are spaced by integer pixel amounts, fractional pixel amounts are reconstructed less effectively. The middle set, where $R_z = 1.5$, shows reclaimed edge detail from the smoothed appearance of the initial iteration. However, the residual aliasing energy is still high. This is thought to be a consequence of the resolution of the data consistency mask. While increasing the grid sampling (R_m) can potentially alleviate this issue, it significantly contributes to the number of reconstruction computations. It is expected that for non-Cartesian trajectories, such as 3D spirals, the reconstruction is more efficacious in areas of k-space where the trajectory spacing is closer to integer distances. For variable density spiral, this is hypothesized to result in aliasing reductions grouped in bands of equal radial distances.

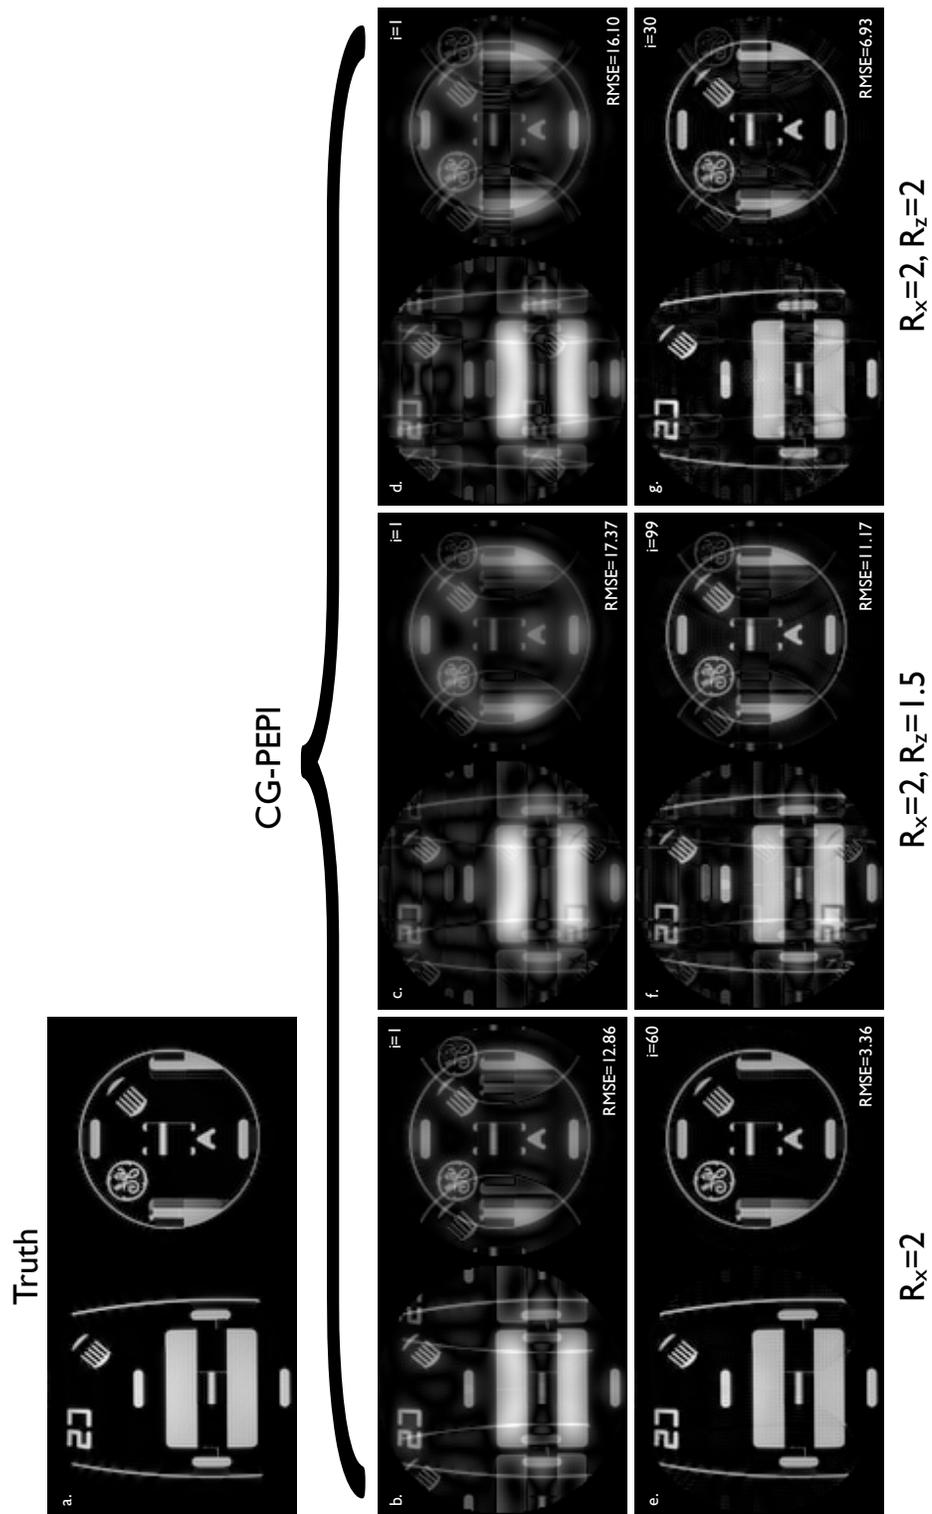


FIG. 4.6: Simulated CG-PEPI reconstructed data at various levels of undersampling (R).

4.3.3 CG-PEPI and CG-SENSE

The columns in figure 4.7 represent (from left to right) fully sampled coil combined, SENSE, PEPI and RSS coil combined reconstructions for undersample factors of 4 (top row) and 9 (bottom row). CG-PEPI shows a moderate reduction in the level of reconstructed aliasing while CG-SENSE yields the largest reductions. Figure 4.7(e) shows a few variable density spiral-cones of the FLORET trajectory around the k_x axis. The FLORET trajectory is comprised of spirals such as these coiled around k_x , k_y , and k_z with varying densities to supporting a sphere in k-space. The uniform undersampling of the FLORET trajectory is thought to impede the some of the possible aliasing reductions provided by PEPI because there are fewer opportunities for integer spacing between trajectories.

The time per iteration for PEPI is 10 seconds for both $R = 4$ and $R = 9$ (each requiring 5 iterations). The time per SENSE iteration is 50 seconds for $R = 4$ and 10 seconds for $R = 9$ (each requiring 15 iterations). The prep time for PEPI is about 10 minutes for $R = 4$ versus 2 minutes for SENSE.

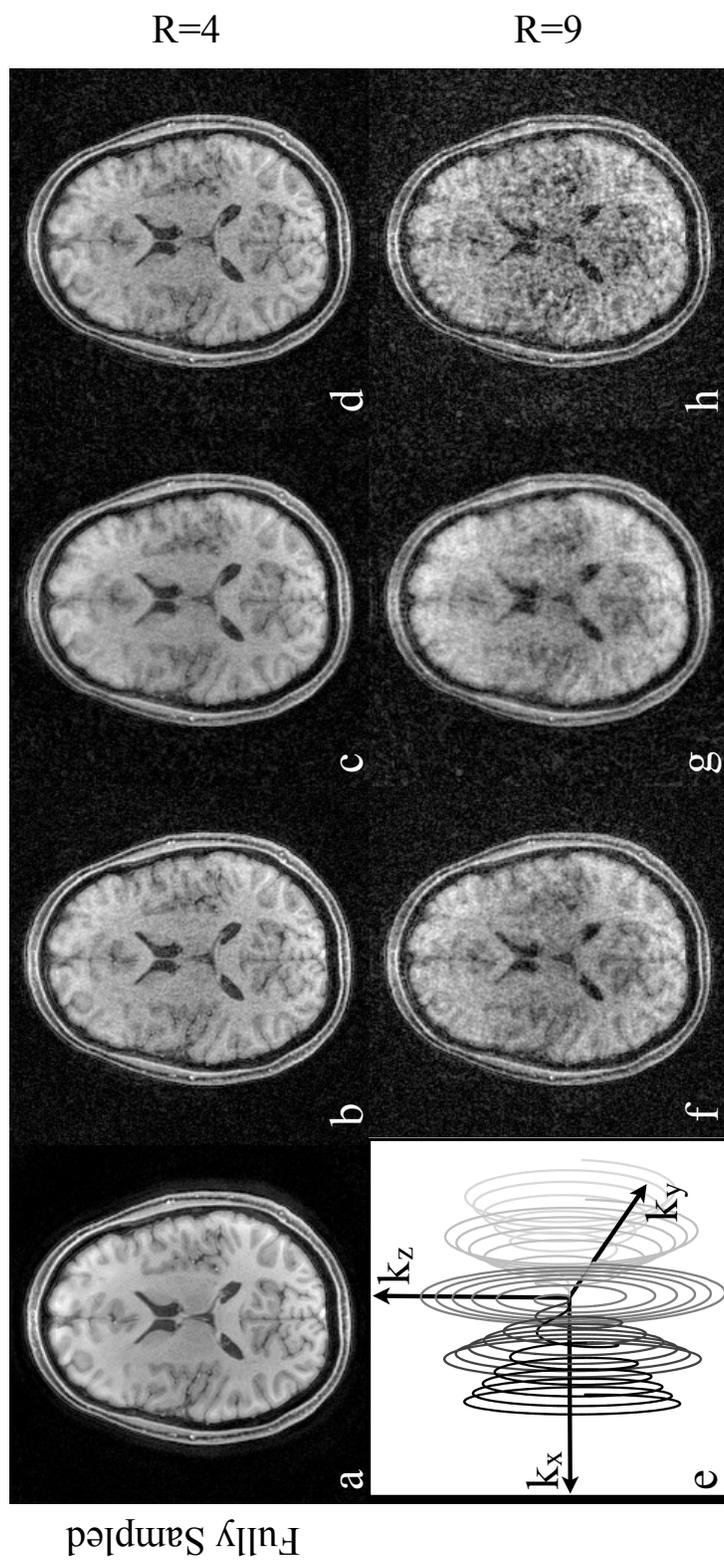


FIG. 4.7: Parallel reconstruction of FLORET data at undersample factors of 4 and 9. (a) a fully sampled FLORET data set reconstructed using RMS coil combination. The remaining top row is FLORET undersampled with $R = 4$; the bottom row is undersampled with $R = 9$. Images (b) and (f) were reconstructed using CG-SENSE, (c) and (g) with CG-PEPI, and (d) and (h) without parallel imaging. A sample FLORET is shown in (e).

4.4 Conclusion

CG-SENSE is practical, on multi-CPU platforms, and is the chosen method at this junction for the continued exploration of rapid 3D PCMRA. PEPI has trajectory independent, short, iteration time but may require extended prep time, which is mostly spent on the calculation of sampling density. The increased prep time and reduced iteration time potentially makes PEPI well suited for dynamic reconstruction applications (2D dynamic PEPI, APPENDIX B).

5 SAMPLE DENSITY CORRECTION

The reconstruction of non-Cartesian MRI trajectories requires estimation of the non-uniform densities in sampled k-space prior to gridding or direct Fourier transform reconstruction. Accurate density estimations are essential to faithful image reconstruction. Problems that may hinder this estimation are those that cause trajectories to cross non-uniformly or inconsistently. For some trajectories, such as rosette (46) and lissajous (47), this is by design. For retrospective correction techniques, such as motion compensation (11, 48), gradient delay correction (49), and dynamic imaging (50), the trajectories may be repositioned or temporally windowed such that the samples overlap in a non-predetermined way. Density estimations that fail to take this into account inherently incur error due to these effects.

Sample density calculations are often specific to a trajectory, exploiting prior knowledge in order to minimize computation time. Previous work done by Johnson et al. (9) presented a method that was shown to be faster and more accurate, without any loss of generality, than several other methods, some of which were trajectory specific (51–57). However, densely sampled areas of k-space may still require a considerable computation time using Johnson’s method. Highly localized density is common in center-out trajectories such as spiral or radial acquisitions. High central density can be further compounded in 3D methods, where trajectories are not restricted to a single plane. Center-out 3D trajectories emanate in all directions, appreciably oversampling the center of k-space (12).

This chapter presents a method which combines the optimal kernel design specified in (9) with the iterative method presented in (58), yielding optimal error suppression with a high computational efficiency (8, 59). The method

proposed in this work is compared to Johnson’s implementation for both accuracy and execution time, as well as two analytical methods. The methods presented here also demonstrate the flexibility of the algorithm due to its trajectory non-specificity, the significance of which is discussed for trajectory design and compatibility with trajectory corrective techniques.

5.1 Theory

The proposed sample density estimation method is introduced in terms of its application in gridding reconstruction, followed by descriptions of the constituent methods previously proposed in (58) and (9). The design considerations for the proposed method are discussed in the final section.

5.1.1 Gridding Reconstruction

Sample density estimates are commonly employed in the gridding reconstruction process as detailed in references (58) and (55). An Eq. representing the gridding process can be defined as

$$M_{\text{III}} = ((M \cdot S \cdot W) \otimes K) \cdot \text{III}_r \otimes^{-1} K . \quad (5.1)$$

K-space data M are sampled at trajectory points S and multiplied by a weighting function (or a sample density compensation function (DCF)) W . Sample points are convolved (\otimes) onto the reconstruction grid III_r by the gridding convolution kernel K . Deapodization is performed in the spatial domain and is equivalent to the deconvolution (\otimes^{-1}) of the gridded points with the grid kernel K . The resulting data are denoted by M_{III} .

The weighting function W averages data that are oversampled to various degrees throughout sampled k-space. Samples within areas of high density are multiplied by lower weighting values while samples from areas of low density

receive higher weight. The balanced weighting function ideally provides a modulation transfer function (MTF) that is unity across sampled k-space.

5.1.2 Iterative Sample Density Estimation

As succinctly described in (32) the ideal weighting function is the solution to $S \cdot (W \otimes C) = S$. The iterative method for conditioning W , detailed in (58), is

$$W_{i+1} = \frac{W_i}{W_i \otimes C} . \quad (5.2)$$

C is a convolution kernel that attenuates sharp transitions across the MTF. Analogously, the spatial profile of C (i.e. the Fourier transform of C) modulates the point spread function (PSF). For each iteration i , the current weighting estimates W_i are conditioned through division by the modulated weights (i.e. $W_i \otimes C$).

The convolution in the denominator of Eq. 5.2 is a non-uniform convolution. As Pipe has shown for 2D trajectories (58), this convolution can be evaluated directly or through a two stage gridding process as in,

$$W \otimes C_{direct} \approx (((W \otimes C_{grid}) \cdot \text{III}) \otimes C_{grid}) \cdot S . \quad (5.3)$$

The two stage method, shown on the right hand side of Eq. 5.3, involves an intermediate grid step using a temporary grid III (not to be confused with the reconstruction grid III_r). As depicted in Fig. 5.1, the weighting function is convolved twice, once onto the intermediate grid, and from there, back onto the sample locations S .

The convolution kernels used in each method are labeled as C_{direct} for the ‘direct method’ (i.e. the left hand side of Eq. 5.3) and C_{grid} for the ‘grid

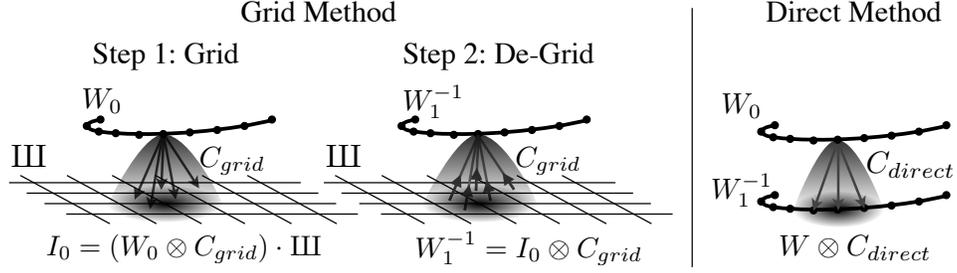


FIG. 5.1: An illustration of how Eq. 5.3 is executed on the first iteration. I represents an intermediate step, III the grid, W the weighting function, and C_{grid} the convolution kernel used in the grid method. For each iteration, the weighting function is gridded using the convolution kernel and de-gridded using the same kernel. In the first iteration $W_0 = 1$. The de-gridded points of W_1^{-1} are a measure of the density. The relative weights are determined by $W_1 = 1/W_1^{-1}$.

method' (i.e. the right hand side of Eq. 5.3). Since the grid method uses two convolutions, the grid kernel is designed to have the same net effect as the direct kernel (58); in the spatial domain this design requires

$$c_{grid} = \sqrt{c_{direct}}. \quad (5.4)$$

Eq. 5.4 shows that the grid kernel can be designed in the spatial domain using the Fourier convolution theorem, where c_{grid} and c_{direct} are the spatial domain transforms of each respective kernel.

5.1.3 Direct Method Design

The direct method, as implemented by Johnson (9), employs both an optimal convolution kernel design for conditioning the PSF and a method for reducing the computational load of the non-uniform convolution.

5.1.3.1 Optimal Convolution Kernel

The optimal conditioning profile is constructed by convolving the region of the signal source with the region over which error is to be minimized. Both regions can be represented by the field of view (FOV), defined in this work as

$$\psi(\vec{r}) = \begin{cases} 1, & |\vec{r}| \leq \zeta/2 \\ 0, & |\vec{r}| > \zeta/2. \end{cases} \quad (5.5)$$

In this case, ψ represents a parameterized spherical FOV bounded by a diameter ζ in terms of the spatial radius \vec{r} . ψ is both the region of signal source and the region over which error should be minimized, therefore

$$C_{direct} = \mathcal{F}\{\psi \otimes \psi\} \quad (5.6)$$

where $\mathcal{F}\{\cdot\}$ represents the Fourier transform.

5.1.3.2 Computational Optimization

As detailed in (9), the number of operations required to evaluate the non-uniform convolution of Eq. 5.2 was significantly reduced by presorting the trajectory coordinates into compartments of equal size and space. The boundaries of each compartment are spaced by the diameter of the convolution kernel. This ensures that the search for neighboring trajectory points during the convolution is confined to the immediate and adjacent compartments only. Consequently, the number of points that fall within a compartment is trajectory dependent. The number of operations within a given region of k-space is $O(N^2)$, where N is the number of points within the immediate and adjacent compartments.

5.1.4 Grid Method Design

This section covers the design considerations used in implementing the grid method. The theory for an ideal grid kernel is discussed as an introduction to the design parameters, then the limitations are covered in terms of a practical implementation.

5.1.4.1 Ideal Grid Kernel

An optimal grid method kernel (C_{grid}), of full spectral support, is designed for a spherical FOV by generating C_{direct} , outlined in the previous section, and then applying Eq. 5.4. Figure 5.2(a) shows a plotted radius of the 3D kernel, C_{grid} . A radius of the corresponding spatial profile (c_{grid}) is plotted in Fig. 5.2(b). The width of the spatial profile is twice the diameter of the FOV as a consequence of Eq. 5.6 (i.e. 2ζ).

The minimum grid resolution necessary to support an FOV of ζ is $1/\zeta$. Resolutions higher than the minimum are denoted by a grid oversample factor R . For example, in order to adequately support the spatial profile of the ideal C_{grid} kernel, the intermediate grid (III) must have a minimum resolution equal to $1/(R \cdot \zeta)$ where $R = 2$.

Figure 5.3 illustrates that when the grid kernel is sampled by the grid in the frequency domain, replicates are generated in the spatial domain. If the grid is oversampled by a factor of R , then the center to center spacing between the spatial profile and the replicated profiles is $R \cdot \zeta$. In the case shown in Fig. 5.3(a), the kernel is fully supported and is therefore spatially bound between $|\vec{r}| = \zeta$. This example also shows that an $R = 2$ causes no aliasing overlap between the conditioning profile and the replicate profile.

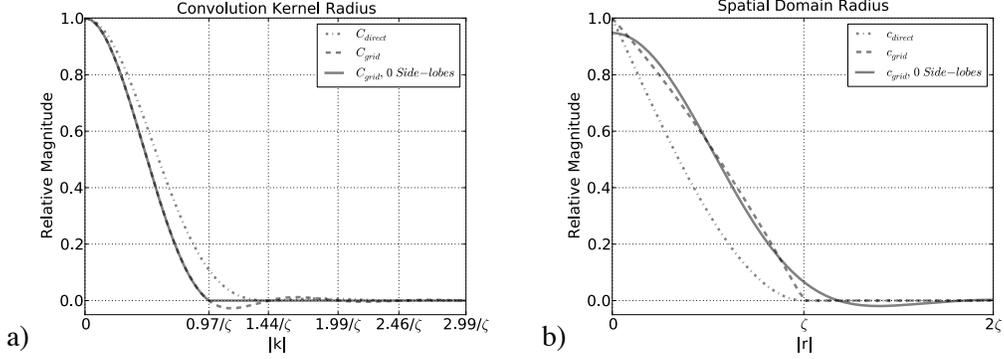


FIG. 5.2: The spatial and frequency domain profiles of the proposed kernel. The radius of the spherically symmetric kernel is plotted in the frequency domain (a) and the spatial domain (b). The full spectral support (solid) is compared to a truncated spectral support (dashed).

5.1.4.2 Approximate Grid Kernel

Realistically, the C_{grid} kernel cannot be fully supported. The number of operations for a gridding convolution is proportional to $O(L_C^3 \cdot R^3 \cdot N)$, where L_C is the kernel diameter and N is the total number of trajectory points. It is, therefore, advantageous to minimize the kernel size in order to reduce the computational time of each convolution.

Truncating C_{grid} to leave only the main-lobe (Fig. 5.2(a), $C_{grid, 0 Side-lobes}$) causes ringing in the spatial domain that extends past $|\vec{r}| = \zeta$ (Fig. 5.2(b), $C_{grid, 0 Side-lobes}$). Figure 5.2(b) shows how most of the energy of the spatial profile is still contained within a 1ζ radius. The corresponding example illustrated in Fig. 5.3(b) shows that ringing which extends outside of the boundary will alias back into the conditioning profile. Figure 5.3(c) and (d) show the effect that these aliasing ripples have on the reconstructed image at two different levels of convergence for $R = 1$. Since the ringing diminishes with $|\vec{r}|$, the amount of aliasing error is decreased with increasing R .

The convolution kernel may be designed analytically, as shown by Johnson,

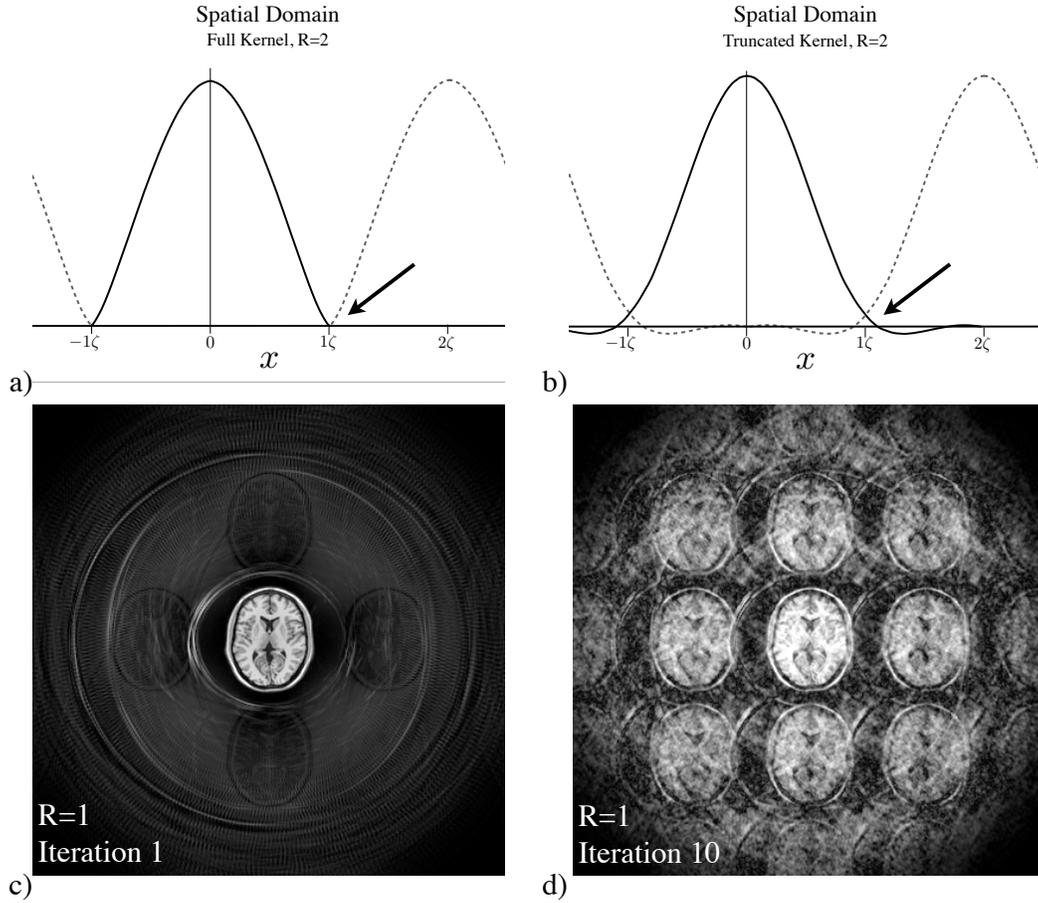


FIG. 5.3: Fully supported sampling of the ideal kernel (a) and aliasing of the truncated kernel (b). The fully supported kernel provides a spatial profile that is not aliased for $R \geq 2$ (a). The truncated kernel profile produces a spatial profile with Gibbs ringing (b). The ringing is not spatially bounded and will therefore alias into the conditioning profile, as indicated by the arrow. A cross section of a 3D reconstruction for insufficient oversampling ($R = 1$, C_{grid} , 0 Side-lobes) after the first iteration (c), and tenth iteration (d). The reconstruction grid of Eq. 5.1 was oversampled by a factor of 4.

using continuous functions for Eq. 5.5 and 5.6, and then evaluated for each grid point during the convolution of Eq. 5.2. The kernel implemented in this work was generated numerically, starting with a discrete ϕ and using the FFT to evaluate the Fourier transform. A kernel table containing a presampled radius of 10,000 points was used to reduce the computation time of the convolution operations in both the grid and direct method implementations used in this work.

5.2 Simulations

Simulated data were used to characterize the grid method for the purposes of determining the optimal grid oversample factor and for benchmark comparisons. The grid method is compared to the direct method in convergence rate, computation time, and accuracy. The grid method is also compared to two analytical methods in terms of accuracy. Options for various initial conditions are presented followed by the parallelization strategy used in the implementation of the grid method which is covered in the final section.

5.2.1 Data

5.2.1.1 Trajectories

Three 3D k-space trajectories were used to characterize any trajectory dependence of the proposed sample density estimation method. Since the proposed method is generic and makes no assumptions about the underlying trajectory, it was applied with no specific or additional parameter adjustments for each trajectory. Each sample point of the whole trajectory is evaluated for each iteration of Eq. 5.2 without regard to symmetry or redundancy within the trajectory. Initial conditions that take advantage of symmetry (58) within the trajectory will be covered in a later section to show how they may be calcu-

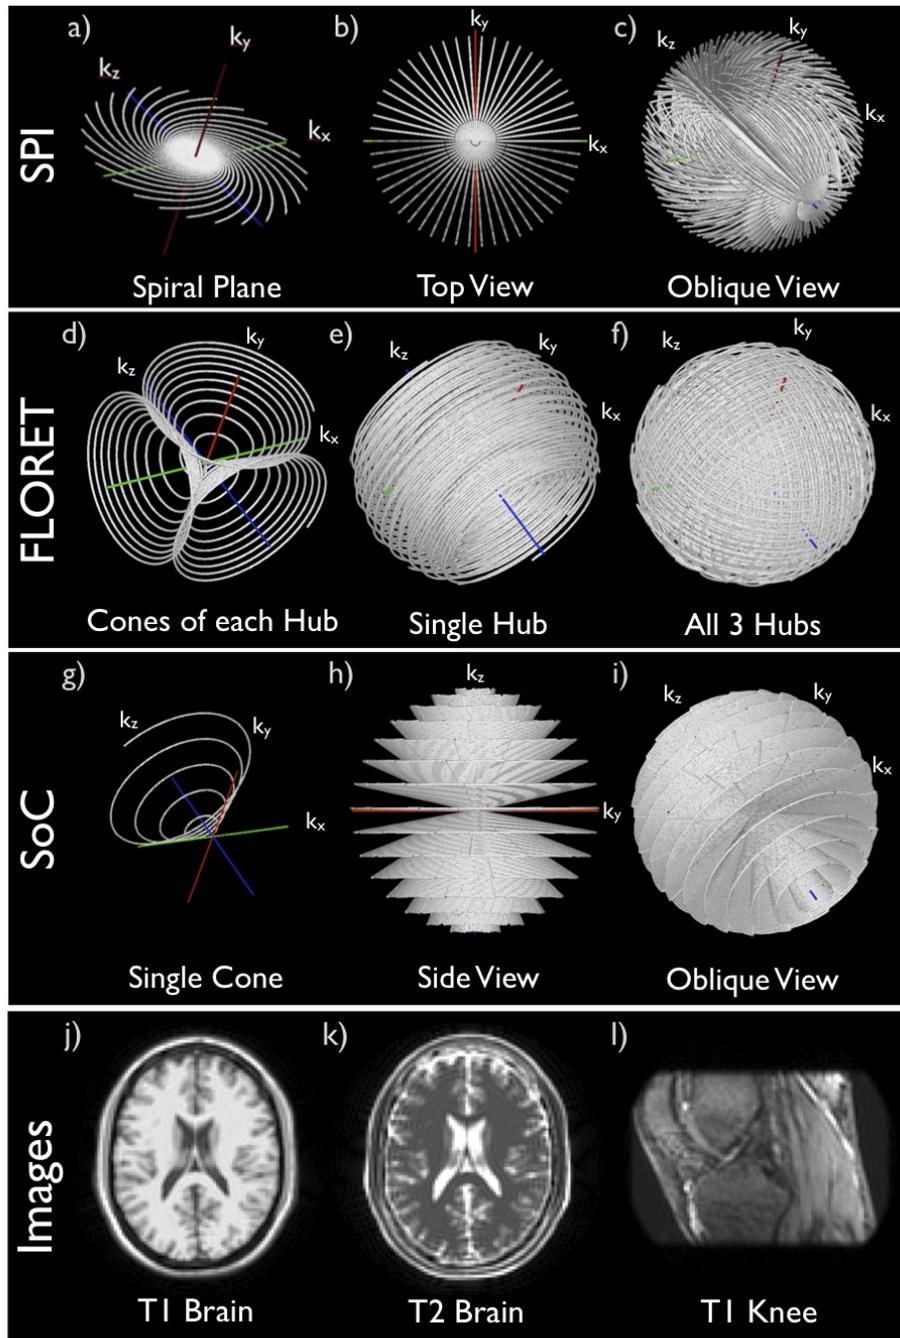


FIG. 5.4: Simulation elements. (a-i) component and fully assembled views of the SPI, FLORET, and SoC trajectories. (j-l) mid level slices of the images volumes used for each simulation.

lated and used to increase the rate of convergence. The following trajectories used in each simulation were designed to support a 100^3 grid matrix.

The spiral projection imaging (SPI) method is a hybrid of spiral and projection trajectories (11). Fully sampled Archimedean spirals are spaced at even angular distances on a 2D plane (or projection) as shown in Fig. 5.4(a). Multiple spiral planes are collected at even angular distances rotated about the k_z axis to sample 3D k-space (Fig. 5.4(b) and (c)). For this simulation, each 2D projection is comprised of 11 spiral interleaves, each containing 1,370 sample points. 157 projections were used to fully sample in the radial direction.

The FLORET trajectory consists of 3 hubs, oriented along the primary axis, each containing multiple variable density spiral cones (Fig. 5.4(d)), rotated by the golden angle (45). 3D k-space is fully sampled when 3 orthogonal hubs are combined. Each of the 3 hubs were designed with 575 spiral cones with each spiral containing 1,262 sample points.

The 3D stack of cones (SoC) trajectory used here was generated with the code supplied by Gurney et al. from reference (60). Figure 5.4(h) shows a stack of fully sampled spiral cones. Each cone contains a variable number of spirals to evenly support the area of each cone. The total number of cones used was 159 which supported a total of 1,498 spiral trajectories. Each trajectory contained about 1,242 sample points.

5.2.1.2 Images

Three image volumes (Fig. 5.4(j-l)) were used to characterize the impact of sampled data on the reconstruction error. Two of the image volumes used were brain simulations, of T_1 and T_2 contrast, from the McGill University, Brain Imaging Centre, <http://www.bic.mni.mcgill.ca>. The third is a T_1

weighted knee image from the ISMRM 2010 Reconstruction Challenge, ‘Piece of the Puzzle’ contest, http://www.ismrm.org/mri_unbound. In order to ensure the error measured in each simulation was on the same relative scale, the image energy of each volume was normalized to a value of 100^3 .

5.2.2 Optimal Grid Oversample Factor

The optimal oversampling factor was determined empirically by minimizing the error of the density estimates and the computational time. Error in the density estimates was calculated indirectly through the RMS difference between truth and the reconstructed images. The full volume RMSE was calculated for each reconstruction. Gridding reconstruction (61) was used instead of the direct Fourier transform to minimize the simulation time.

The error convergence was characterized for R ranging between 1 to 4 in increments of 0.1, over 100 iterations, for the 9 image-trajectory combinations. The resulting data were analyzed for R dependent convergence rate and final accuracy. The results were also compared for inter-image and trajectory variation.

5.2.3 Direct Method Comparison

The proposed grid method was compared to Johnson’s implementation of the direct method for both computation time and accuracy. This section covers the design parameters specific to Johnson’s implementation, the convergence comparisons between the two methods, and the parameters used for the timing benchmarks.

The number of compartments used in the direct method are based off of the level of kernel truncation (i.e. kernel diameter). The kernel diameter, for a 0 side-lobed kernel used in the direct method, is approximately $1.43/\zeta$.

Therefore, the number of compartments used for a 100^3 supported matrix is $\lceil 100/1.43 \rceil^3 = 328,509$. The 1 side-lobe kernel has a diameter of about $2.5/\zeta$ yielding 64,000 compartments.

The convergence and final accuracy were characterized by acquiring the reconstruction error for each of 100 iterations for 4 parameter variations (2 each) using gridding reconstruction. Both the grid and direct method were run for each of the 9 image-trajectory combinations. The trade-off between time and accuracy was evaluated for each method by performing a more computationally demanding setting for higher accuracy, and a less demanding setting, yielding lower accuracy. The grid method was performed with oversample factors of $R = 2.1$ and 3, using the truncated kernel (C_{grid}) of 0 side-lobes (Fig. 5.2(a)). The direct method was performed using a kernel (C_{direct}) with 0 side-lobes and with 1 side-lobe.

The total computation time depends on the rate of convergence (the number of iterations required) and the time per iteration. As previously discussed, the computational time of the direct method depends on the variation of sample density across k-space, where compartments may contain different numbers of samples. The grid method is linearly dependent on the number of samples used in the whole trajectory, independent of local variations in density. The computational dependence of the two methods was differentiated by measuring the time per iteration on a series of SPI trajectories of variable density with the same total number of trajectory points. The grid method was timed using $R = 2.1$, for single and parallel execution, and $R = 3$. The direct method was timed using kernels of 0 and 1 side-lobes.

The variable density SPI trajectories were generated for a 100 matrix, with 157 projections, using 11, 21, 31, 41, and 51 spiral interleaves per projection.

The maximum slew rate was varied to maintain a relatively equal number of sample points between each set. Benchmarks were made on an 8-core, 3.2 GHz Intel CPU. The average durations over 25 successive executions of each method were measured. To get a sense of the relative central density of each trajectory, the sample density (as determined using the grid method) was reported for a trajectory point positioned at k_0 .

5.2.4 Analytical DCF Comparison

The proposed method was compared to analytical density compensation functions for the SPI and SoC trajectories. The T_1 weighted brain simulation was used to compare gridding reconstructions for each method. The analytical density estimation presented by Gurney et al. (60) was used to calculate the weighting for the SoC trajectory.

For SPI, the density of the spiral and projection aspects were respectively calculated by employing the 2D spiral compensation presented in (51) and the density compensation of a 2D projection imaging trajectory for the combined expression,

$$W_{SPI} = (\vec{k}_s \cdot \vec{g}_s) \cdot \sqrt{k_x^2 + k_y^2} \quad (5.7)$$

where \vec{k}_s and \vec{g}_s are the corresponding k-space and gradient waveforms for a 2D spiral as shown in Fig. 5.4(a). The magnitude of the k_x, k_y vector represents the 2D radial density as shown in Fig. 5.4(b).

5.2.5 Initial Conditions

The proposed method takes an initial condition which is by default $W_0 = S$. This default initial assumption is that all sample locations have equal density. However, the convergence of the algorithm can be shortened if some

prior knowledge of the density is supplied. This information can be obtained analytically or by taking advantage of symmetry within the trajectory. The following simulations explore several options for initial conditions to the grid method, using the T_1 weighted brain images as the basis.

The initial condition can be generated by the proposed method itself using an oversample factor that shortens its execution time. This will be referred to as ‘cascaded operation’. In the first stage, the grid method was run with $R = 1.5$ for 15 iterations, second with $R = 2$ for 15 iterations and $R = 2.1$ for the remaining iterations. The average iteration time was measured for each stage of the cascade.

The SPI trajectory is an illustrative example of symmetry as it contains the same 2D spiral trajectory in each projection. The grid method was run on a single 2D spiral for 10 iterations with $R = 1.5$. The densities determined for the 2D spirals were then multiplied by the radial density in the k_x, k_y plane. The result was used as an initial condition for the grid method with $R = 2.1$ for the remaining iterations. This method will be referred to as single projection preconditioning. The average iteration time was measured for the 2D spiral stage.

An analytical approximation was also used as an initial condition for each trajectory. Since each of the trajectories tested is a center-out trajectory, a rough approximation of the density is that of a 3D projection trajectory, where the weighting is proportional to the radius squared (50). This approximation, defined as $W_{PR} = |k|^2$, was used with subsequent iterations of the grid method performed with $R = 2.1$. Additionally, the analytical density compensation functions presented in the previous section were used as initial conditions for the SPI and SoC trajectories with subsequent iterations performed using $R =$

2.1. The average iteration time for each configuration was measured. The convergence error was compared against the calculated average reconstruction time at each time point.

5.2.6 Parallel Computation

The proposed sample density estimation algorithm was split into parallel procedures for both the gridding and degriding operations separately. This implementation of the gridding process was designed for an 8-CPU platform. The 3D grid is split into equal sized octants which overlap by an amount equal to the radius of the convolution kernel, similar to the concepts proposed in (62). The trajectory coordinates are then presorted once, at the beginning of the method, so that trajectory points can be gridded to their respective octant, concurrently, with the neighboring octant gridding processes. Splitting the grid through the origin, in this way, is particularly well suited for center out trajectories, since the largest sample density (in the center of k-space) is evenly divided across processes. The individual octants are then added to the full 3D grid (an operation that is also split into multiple processes).

The degriding operation is more easily parallelized since the convolution of each trajectory point is independent of neighboring points. The coordinate points are evenly divided by the number of desired parallel processes. The convolution at each point requires only read access to the 3D grid allowing multiple simultaneous convolutions.

5.3 Results

The simulation results are presented for the determination of the optimal grid oversample factor (R), benchmark comparisons between the grid and direct methods, comparisons between the grid and analytical methods, and of the

efficacy of various initial conditions.

5.3.1 Grid Oversample Factor

This section reviews the accuracy of the grid method at the tested levels of R followed by an assessment of how the accuracy of the grid method is affected by different image and trajectory combinations.

5.3.1.1 Accuracy and Convergence Properties

A few representative examples of how the oversample factor affects the convergence were generated using the SPI trajectory and T_1 brain image combination. Figure 5.5(a) shows that the algorithm converges on solutions at increasing levels of accuracy for increasing values of R . The relative difference in error between the final solutions for each of the chosen R is about an order of magnitude. However, for oversample factors of $R = 1.5$ and 2.1 , the algorithm converges to a solution with similar accuracy in the first 15 iterations. This is particularly useful (as will be seen in the Initial Conditions section) since the computation time is dependent on R^3 , as presented in the theory.

5.3.1.2 Trajectory and Image Variation

Figure 5.5(a) shows the final level of error for each of the 9 trajectory-image combinations after 100 iterations for various levels of oversampling. As can be seen in the plot, the accuracy of the final solution is most affected by the chosen level of oversampling. The final solution for each level of R has the same relative error for each of the 9 combinations (i.e. each plot has the same relative shape). This shows that the DCF produced by the proposed method, is minimally impacted by the input trajectory and, in turn, has a minimal impact on the reconstruction (i.e. the MTF) of the sampled data.

The plot also shows that there is a local minima of reconstruction error

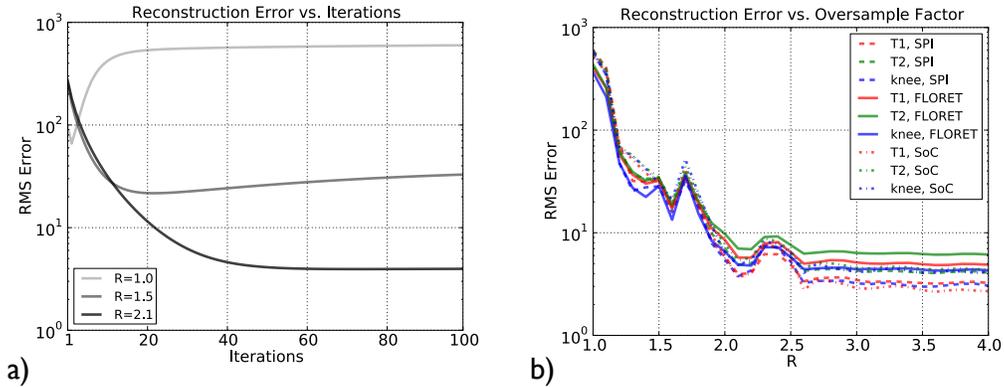


FIG. 5.5: Error convergence at various levels of oversampling. (a) insufficient R causes the algorithm to converge on a solution that is less than optimal. (b) the final error after 100 iterations is plotted for various oversampling factors R .

at $R = 2.1$. This local minima is maintained across image and trajectory variations, making it a good choice for achieving near maximal accuracy with the lowest number of operations. The next most effective oversample factor occurs at $R = 2.6$ and continues to improve in accuracy as R is increased, with diminishing returns.

5.3.2 Direct Method Comparison

5.3.2.1 Accuracy

The convergence rate for each image-trajectory combination was similar across trajectories for both the direct and grid methods. To summarize the results for each method, the reconstruction error was averaged over each image-trajectory combination, shown in Fig. 5.6(a). The plot shows that the direct method has a higher per-iteration convergence rate, however, each method and their parameter variants all converge to solutions of comparable error.

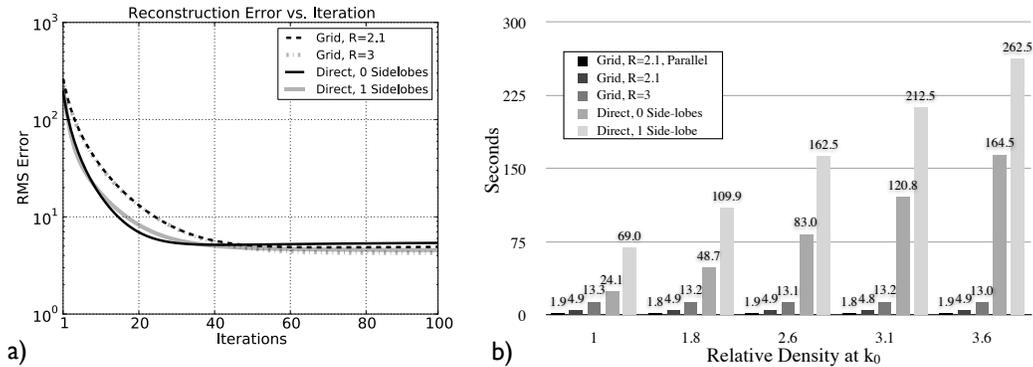


FIG. 5.6: The direct and grid method comparisons. (a) the average RMSE of the 9 image-trajectory combinations is plotted against the iteration number for each method. (b) per iteration computation time. The SPI trajectories containing 11, 21, 31, 41, and 51 spiral interleaves per projection have relative central densities of 1, 1.8, 2.6, 3.1, and 3.6 respectively. Each trajectory contains approximately the same total number of points.

5.3.2.2 Execution Time

The timing benchmarks in Fig. 5.6(b) are grouped in terms of the relative central density of each test trajectory. The per-iteration time for the direct method significantly increases with the central density and kernel size. The 0 side-lobe kernel configuration takes about 12 to over 85 times longer than the parallel grid method from the low to high density trajectories respectively. The 1 side-lobe kernel configuration requires a computation time that is 2.8 times longer than the 0 side-lobe configuration for the high density trajectory, and 1.5 times longer for the low density trajectory.

The per-iteration time of each of the tested grid methods is constant across trajectories. The parallelization reduces the computational time by a factor of about 2.6 (compared using $R = 2.1$). The time reduction between $R = 3$ and 2.1 is about a factor of 2.7.

5.3.3 Analytical DCF Comparison

The plot in Fig. 5.7(a) shows that grid method converges to a solution with similar accuracy to the analytical DCF used for the SPI trajectory. The grid method is also shown to converge on a solution with one tenth of the error produced by the analytical DCF for the SoC trajectory. As seen in the reconstructed images in Fig. 5.7(b), the error is mostly below the visual threshold in each set except for a slight shading of the ventricles and basilar pons, in the SoC reconstruction using the analytical DCF.

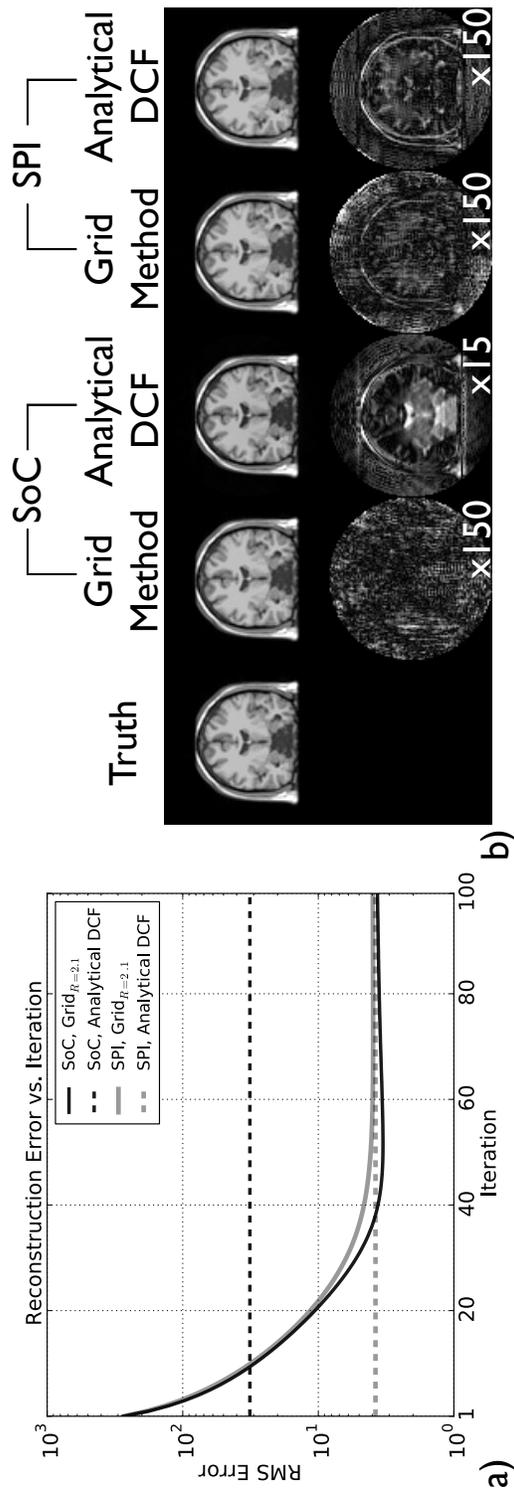


FIG. 5.7: The grid method is compared to the analytical DCF for the SPI and SoC trajectories. (a) the simulated data were reconstructed using the grid method for $R = 2.1$ (solid curves) and the analytical DCF (dashed lines). (b) reconstructed coronal slices and error residual images of the T_1 brain simulation. The residual images were scaled by a factor of 150 for the grid method and SPI DCF, and 15 for the SoC DCF.

5.3.4 Initial Conditions

The error convergence of the grid method ($R = 2.1$, parallel execution), with T_1 brain data, and each of the three trajectories were computed for various initial conditions. Figure 5.8 shows the reconstruction error versus the total computational time needed to achieve that error. The computational time was calculated based on the measured average iteration time and the number of iterations used to obtain each level of error.

Setting the initial condition to the analytical solutions that were specifically designed for SoC and SPI provided the shortest convergence time. In the SoC case, the first few iterations produce the most significant reduction in error, providing an optimal solution in less than 10 seconds. In the SPI case, the first iteration produces a slight reduction of error (compared to the ‘SPI, Analytical DCF’ in Fig. 5.7(a)) and converges on a solution, equivalent to the original level of error, and similar to that of the grid method without initial conditioning.

The next most significant time reduction is SPI specific, and is achieved by using the symmetry of the trajectory. The plot reflects how the pre-conditioning of a single 2D spiral takes about a tenth of a second per iteration. A distinct drop in reconstruction error is visible in the first second of the convergence, indicating the relative time spent on pre-conditioning. The resulting DCF for a single spiral was then duplicated for the 157 projections and multiplied by the radial density in the k_x, k_y plane.

The rough approximate analytical solution (W_{PR}) and cascaded operation both consistently make modest reductions in convergence time for each trajectory.

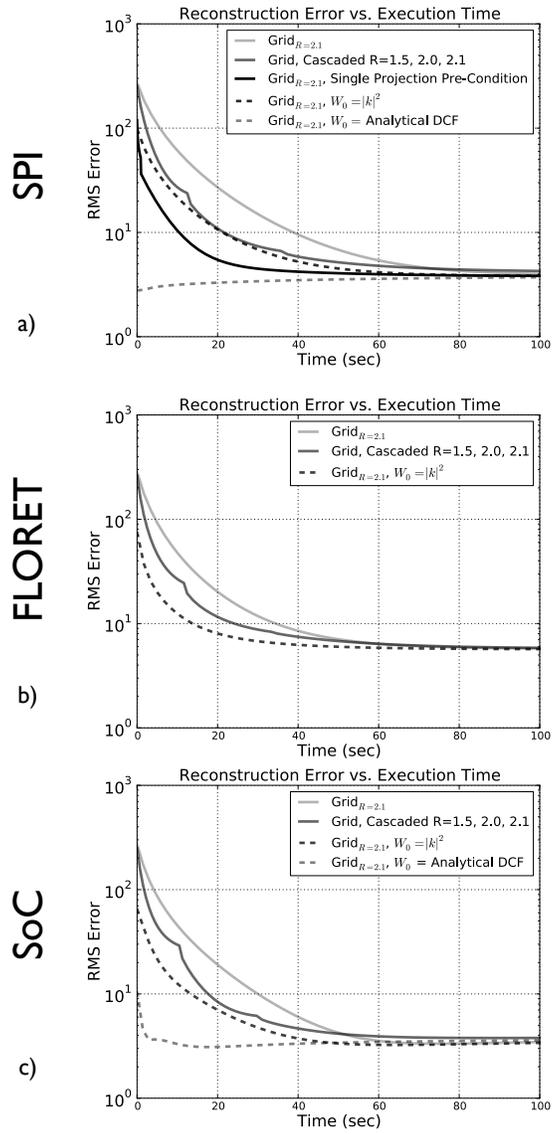


FIG. 5.8: The error convergence rate of various initial conditions for each trajectory using the grid method.

5.4 Discussion

The proposed method is shown to produce accurate density estimations across various image and trajectory combinations. The method provides more consistent results as compared to analytical methods and is shown to be less computationally demanding than previous implementations. The efficacy of initial conditions are shown to be related to the level of information they provide. These points are discussed in the same order they appear in the results section.

5.4.1 Grid Oversample Factor

Aliasing causes poor conditioning at low R , that builds (in error) over multiple iterations. Figures 5.3(c) and (d) show the reconstructed result of poor conditioning due to insufficient oversampling. As shown in these images the replicate FOVs overlap one another and contaminate the region of interest. Analogously, within the density estimation algorithm, this aliasing error occurs in the conditioning profile ($|r| \leq \zeta$) of c_{grid} , distributing error across the PSF. Over each iteration (of Eq. 5.2), the aliasing error is thought to coherently or incoherently compound, and in this case, create large side-lobes in the PSF within a ζ radius. Figure 5.5(b) shows that each trajectory-image combination converges to a local minima at $R = 2.1$. It is thought that, at this oversample factor, the aliasing lobes incoherently add in such a way that is benign to the iterative conditioning.

The kernel size, shape, and grid oversample factor can be designed to minimize aliasing error and computational time of the final gridding reconstruction stage (Eq. 5.1). In work by Beatty et al. (61), the authors derived a relationship between a parameterized kernel model, the grid oversample factor,

and resulting aliasing error. In that work, the kernel shape is variable and the spatial profile is accounted for when the gridded data are deapodized in the deconvolution step. In this work, the shape of the spatial profile is specifically designed to condition the PSF and is fixed by the shape of the FOV. This requirement significantly limits the potential for further reduction of the kernel size without placing more energy outside of the conditioning region (i.e. $|\vec{r}| > \zeta$), or changing the shape of the conditioning profile. For this reason, the level of aliasing error was minimized through the optimization of R , rather than the optimization of kernel shape.

A less obvious benefit of dividing the convolution of Eq. 5.2 into two convolutions is the reduction in kernel size due to the square root operation performed in Eq. 5.4. As shown in Fig. 5.2(a), the radius of the main lobe of C_{grid} is about 30% smaller than it is in C_{direct} . This provides an intrinsically compact kernel.

As discussed in (9), the minimum level of oversampling for the ideal conditioning profile is $R = 2$. The truncated kernel used here requires a relatively small increase in oversampling (i.e. $R = 2.1$) to achieve near optimal accuracy, as corroborated through the comparison to the direct method. Since the direct method is void of aliasing (which is inherent in the grid method), the comparison provides a measure of the level of relative aliasing error. Both methods use truncated kernels and will therefore also incur conditioning error due to their respective divergence from the ideal profile.

5.4.2 Direct Method Comparison

The direct method has a higher convergence rate, per iteration, than the grid method. Figure 5.6(a), shows that direct method converges in about 30% fewer

iterations on average over the 9 trajectory-image combinations. However, the execution time is highly dependent on the sample density as shown in Fig. 5.6(b), taking as much as 85 times longer than the grid method for the 0 side-lobe configuration and the high density SPI trajectory. The computational dependence on compartment size and sample density is also complex. As noted in the results, the execution time for the 0 side-lobe configuration, which utilizes more compartments, approaches the time required by the 1 side-lobe configuration, as the trajectory density is increased.

5.4.3 Analytical DCF Comparison

As shown by the two examples of analytical density estimations, the optimal is comparatively unpredictable. The analytical weighting function for SPI (Eq. 5.7) produces a solution with a comparable level of error to the grid method, and the weighting for SoC produces a solution with ten times the error. The common trait of the two analytical solutions is the residual image contrast shown in the error distributions below each reconstructed image in Fig. 5.7(b). This error is indicative of error in the sample density estimations in the low spatial frequencies, toward the center of the MTF. For center-out trajectories, such as those used here, the center of k-space is where density changes rapidly and is therefore difficult to analytically determine. The proposed method doesn't make any assumptions about the density and is therefore robust in these cases.

5.4.4 Initial Conditions & Convergence

Analytical solutions may not be easily calculated, nor work as effectively, if the underlying trajectory is corrected for motion, gradient delays, or warped to compensate for 1st order field inhomogeneity. In these situations, analytical

solutions may still potentially be used as an initial condition to the proposed sampled density estimation method.

The benchmarks for execution time show a time reduction between $R = 3$ and 2.1 is about a factor of 2.7; theoretically it is closer to 2.9 (i.e. $3^3/2.1^3$), however this shows that the implementation used has low level of unaccounted overhead. This concept is also the motivation for the cascaded operation. While the algorithm doesn't converge to the optimal accuracy when using an insufficient oversample factor, the algorithm effectively converges at the same rate for a smaller number of iterations, at a reduced iteration time. This configuration may be applied with no a priori knowledge of the trajectory and provides a modest reduction in overall execution time.

5.5 Conclusion

The proposed method was shown to significantly reduce the number of computations necessary to estimate the sample density of arbitrary 3D trajectories as compared to the method proposed in (9). The time reductions come at no cost to generality or accuracy. The method is robust in areas of trajectory overlap, where analytical methods tend to be inaccurate. No assumptions are made about the underlying trajectory and therefore the method may be applied after trajectory corrective techniques such as motion correction, system delays, or temporal filtering used in dynamic imaging. The method is also simple and can be effectively parallelized for processing on the latest multi-core computer platforms in a straightforward manner.

5.6 Code Release

A C implementation of the proposed method is available online at the ISMRM MRI-Unbound website: http://ismrm.org/mri_unbound. Wrappers (or gate-

way functions) for both AVS (Advanced Visual Systems, Waltham, MA) and MATLAB-MEX (The MathWorks Inc., Natick, MA) are also provided.

6 MULTI-DIRECTIONAL HIGH MOMENT ENCODING

Phase contrast MRI is a motion encoding technique that is used to provide quantitative velocity information of imaged spins. A component of the velocity is encoded through the use of a bipolar gradient pulse. The pulse encodes a phase in the image proportional to the gradient first moment and the component of the velocity oriented in the direction of the gradient. Three orthogonal components are typically measured to characterize the velocity in three dimensions. Each velocity sensitive acquisition is subtracted by a velocity insensitive acquisition (a reference set) to remove any baseline phase. The scan time relative to a non-velocity sensitive scan increases by a factor of the number of velocity sensitive directions plus the reference scan.

Improvements in the signal quality of this data are made by collecting more signal averages or higher moment encoding. Averages are collected sequentially or more efficiently by using a superposition technique such as Hadamard encoding or by encoding multiple directions (34). High moment encoded phase contrast methods must solve a phase aliasing problem produced when encoding a relatively high velocity. Solving phase aliasing on a per pixel basis typically requires the use of additional velocity encoded data. As MRI scan techniques have become increasingly more efficient at acquiring data these high moment techniques have become more viable (3, 4, 36). Speed improvements in data acquisition methods such as parallel imaging (14, 23) and compressed sensing (24) enable the acquisition of more encoding directions. Additionally, other methods use neighboring pixel information and a priori knowledge of flow to overcome other measurement related inconsistencies such as partial volume (63–65), which is increasingly problematic as resolution is decreased.

This work analyzes a method of acquiring and reconstructing high moment

encoded data to improve the SNR of phase contrast scans. Phase aliasing is estimated by using the correlated velocity information inherent in the measurements of nonorthogonal directions. The effect of the number of measurement directions is explored with reference to the level of phase aliasing.

6.1 Theory

In phase contrast MR, velocity sensitivity is achieved by encoding the change in gradient first moment Δ_{m_1} with a pair of bipolar gradient pulses (16). Spins that are moving during the application of the bipolar pulses are encoded with a phase (ϕ) proportional to the gyromagnetic ratio (γ), the applied gradient moment and the spin velocity (v).

$$\phi = \gamma \Delta_{m_1} v \tag{6.1}$$

Velocity encoding in any arbitrary direction requires at least two points of reference in order to remove any baseline phase on a per pixel basis. This is typically achieved by acquiring a non (or zero) moment encoded set as a reference to the baseline phase.

Since the velocity is encoded as a phase, the range of values that are expressed fall between $-\pi$ and π . This limit is commonly defined in terms of the maximum expressible velocity magnitude ($VENC$).

$$VENC = \frac{\pi}{\gamma |\Delta_{m_1}|} \tag{6.2}$$

Velocity magnitudes that exceed the $VENC$ will result in ambiguous phase measurements or ‘phase aliased’ measurements. Phase aliased values are offset by an integer multiple of 2π . This corresponds to a velocity measurement that is offset by an integer multiple of $2VENC$. Each multiple of 2π or $2VENC$

is referred to as a ‘phase wrap’ or ‘wrap’.

The motivation for lowering the $VENC$ is described in the following equation from the work of Lee et al. (4). The noise (σ_v) in the velocity estimate is proportional to the $VENC$ which is inversely proportional to the gradient moment (equation 6.2). As opposed to signal averaging which increases the SNR with the square root of scan time, lowering the $VENC$ makes a proportional increase in SNR.

$$\sigma_v \approx \frac{\sqrt{2} VENC}{\pi SNR_{Mag}} \quad (6.3)$$

6.1.1 Unaliasing Phase

Lee et al. (4) have shown that phase aliasing can be unwrapped by using additional velocity measurements collected at high $VENC$ where $v < |VENC|$. For example, a spin moving along the ‘x’ direction with a velocity $|v_x| > VENC_1$ has a measured velocity $\tilde{v}_{x,1}$. The measured velocity is unaliased using an additional measurement where $|v_x| < VENC_2$. The estimated velocity \hat{v}_x is

$$\hat{v}_x = \tilde{v}_{x,1} + 2VENC_1 k \quad (6.4)$$

and

$$k = N.I. \left(\frac{\tilde{v}_{x,2} - \tilde{v}_{x,1}}{2VENC_1} \right) \quad (6.5)$$

where k (the number of wraps) is the factor of $2VENC$ that v_x is aliased by and $N.I.$ is a function that returns the nearest integer wrap. Since the SNR of the velocity estimate is inversely proportional to the $VENC$, the acquisition using

$VENC_2$ has a much lower SNR than the acquisition of $VENC_1$. The overall gain in SNR is driven by choosing a $VENC_1$ to be several times lower than $VENC_2$. The unaliased $VENC_1$ is used as the final data set. The $VENC_2$ set maybe combined in a weighted average with the unaliased $VENC_1$ set for a slight gain in SNR. This is called the ‘three-point’ method.

6.1.2 Proposed Method

In the proposed method, all velocity sensitive directions maintain the same low $VENC$. Measurement directions are oriented as orthogonal as possible. Figure 6.1 shows an example of measurement orientations for six directions representing the vertices of an icosahedron. The measurement vectors u_i measure the component of the spin velocity parallel to each measurement direction. Since the vectors u_i are not completely orthogonal, each measurement contains some shared information with its neighbors. This information is used to unwrap phase aliased measurements as follows.

In the absence of phase aliasing, the measured velocity components (\tilde{v}_i) are projections of the true velocity vector \vec{V} where $\vec{V} = [V_x, V_y, V_z]$.

$$\tilde{v}_i = \vec{V} \cdot u_i + \eta \tag{6.6}$$

η is the random noise in the measurement and u_i is the unit vector in the measurement direction (ie. $[u_{i,x}, u_{i,y}, u_{i,z}]$).

If the measurement vector has the possibility of being phase aliased, the velocity estimate (\hat{v}_i) is an integer number of wraps from the measured velocity, that is

$$\hat{v}_i = \tilde{v}_i + 2VENCk_i, \quad i = 0, 1, 2 \dots N. \tag{6.7}$$

Where k_i is the number of aliasing wraps for each of N measurement directions. The velocity estimates in each direction (\hat{v}_i) are also the projections of the estimated velocity vector (\vec{V}_{est}) onto the measurement direction (u_i), similar to equation 6.6.

$$\hat{v} = \vec{V}_{est} \cdot u_i \quad (6.8)$$

If $\vec{\hat{v}}$ is a vector containing the correct velocity estimates for all N directions and u is the matrix containing all direction vectors, then the true velocity vector \vec{V} is estimated by multiplying $\vec{\hat{v}}$ with u^+ , the Moore-Penrose pseudo inverse of u .

$$\vec{V}_{est} = u^+ \cdot \vec{\hat{v}} \quad (6.9)$$

The estimates are back projected onto the three primary axes.

The solution vector \vec{k} (ie. $[k_1, k_2 \dots k_N]$) is determined by forcing consistency between the estimated solutions \hat{v}_1 to \hat{v}_N . The solutions are consistent if the projected estimate of velocity \vec{V}_{est} equals the measurement plus any additional aliasing wraps. From equations 6.8 and 6.7, the correct solution

$$\vec{V}_{est} \cdot u_i = \tilde{v}_i + 2VENCk_i . \quad (6.10)$$

6.1.2.1 Velocity Encoding in 2D

For illustrative purposes a 2D example of the solution space for two measurement directions is shown in figure 6.2. Figures 6.2(a, b) show how the solution space moves relative to the measured velocity \tilde{v} . When the correct velocity estimate \hat{v} has a magnitude larger than $VENC$, phase aliasing occurs which pushes the correct answer to another position in the solution space (figure 6.2(b)).

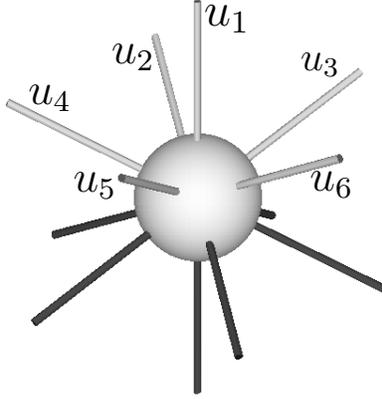


FIG. 6.1: Orientations for six velocity sensitive directions that make the vertices of a regular icosahedron. The white vertices represent the measured directions and the black are the complements. Other measurement configurations are determined by a charge repulsion algorithm (66).

The dotted lines in figures 6.2 and 6.3 represent the solution at integer multiples of phase aliasing in each measurement direction. Solutions with the lowest consistency error reside in positions where there are more solution-line intersections. If all solution lines intersect at a given point then equation 6.10 is satisfied. Solutions where only a few lines intersect will not balance the left hand side and right hand side of equation 6.10 to various degrees. The difference between the right and left hand side is a measure of the \vec{v} consistency. The consistency error (E) is a measure of the RMS proximity to the solution intersections in units of displacement per unit time. The solution vector \vec{k} is found by minimizing the consistency error calculated using equation 6.11.

$$E = \sum_{i=1}^N (\vec{V}_{est} \cdot u_i - (\tilde{v}_i + 2VENCK_i))^2 \quad (6.11)$$

Figure 6.3(a) shows a 2D example using three measurement directions positioned at equal angular distances. The greyed dot represents a solution of

high consistency error with a solution vector, $\vec{k} = [1, 2, 0]$. The black dots represent two solutions with the same consistency error. Both the correct solution ($\vec{k} = [0, 0, 0]$) and a neighboring solution ($\vec{k} = [1, 0, 1]$) are located at points where all solution lines intersect.

In the 2D case, a unique solution is not found until four measurements are collected. Figure 6.3(b) shows neighboring solutions that all have a consistency error greater than zero and equal to each other along concentric rings. The minimum error found within the first phase aliasing wrap (the inner circle) is greater than the minimum error found if a second aliasing wrap were added to the search (the outer circle). Adding phase wraps to the search allows more degrees of freedom to be fitted which increases the noise sensitivity.

The maximum velocity that can be estimated depends on the number of wraps allowed in the search space. The limit on the maximum speed, defined in equation 6.12, can be adjusted through $VENC$ or the number of allowable wraps which trade SNR or noise sensitivity respectively.

$$|V| \leq (2|k| + 1) \cdot VENC \quad (6.12)$$

6.1.2.2 Velocity Encoding in 3D

The 3D solution space works similarly to the 2D solution space. Instead of consistency error aligned in concentric level curves about the solution (2D), errors of the same solution geometry fall into concentric spheres. Figure 6.4 shows the minimum consistency error found within a given radius (in terms of velocity over $VENC$ ratio) in the solution space. If $|k_i|$ is less than or equal to 1, then the error threshold is relatively high for all velocity to $VENC$ ratios less than ~ 1.1 as opposed to velocity to $VENC$ ratios between 1.1 and 3.

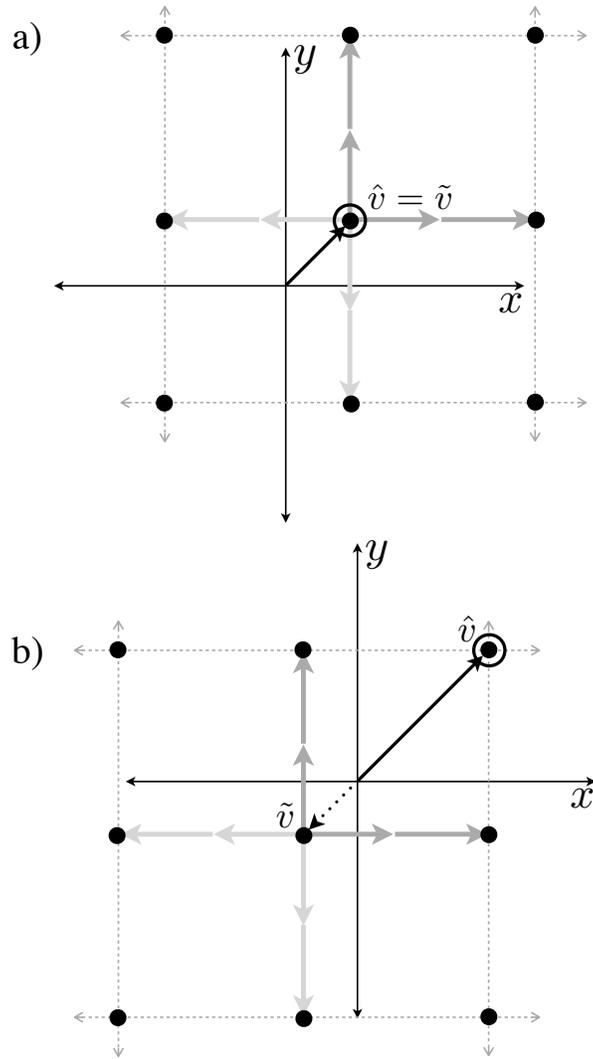


FIG. 6.2: A 2D example of the solution spaces for two measurement directions. (a) the magnitude of \tilde{v} and \hat{v} is 1, with $\vec{k} = [0, 0]$. (b) the magnitude of \hat{v} is 2, however, the measured \tilde{v} magnitude is ~ 0.8 in the opposite direction, with $\vec{k} = [1, 1]$. For (a) and (b) the solution space is $|k| \leq 1$. Each possible solution is represented by a black dot. The correct solution is circled. A $2VENC$ displacement in velocity is represented by dark grey arrows for positive k and light grey for negative k wraps.

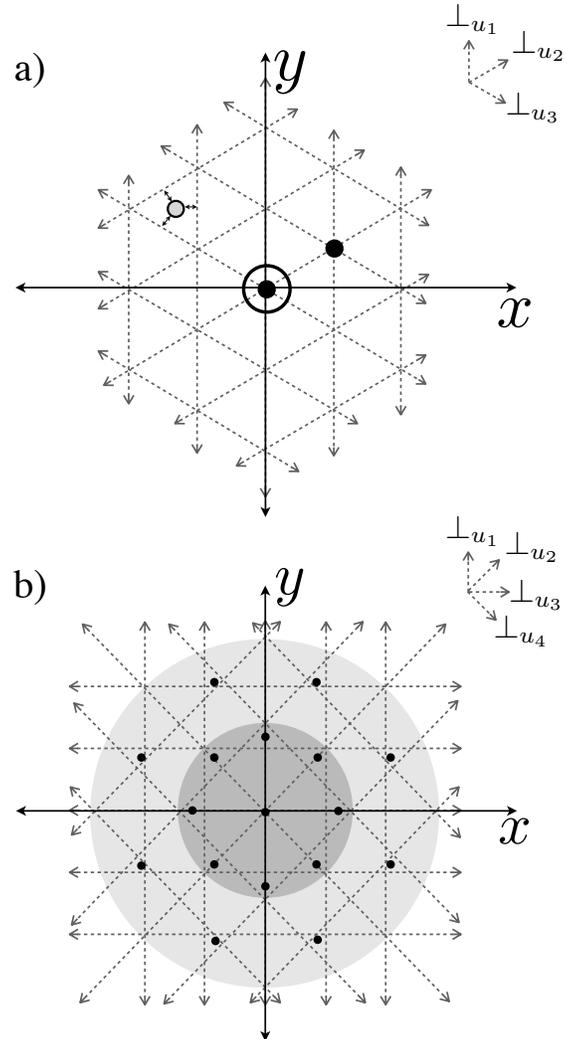


FIG. 6.3: A 2D example of the solution space for 3 and 4 measurement directions allowing for 2 wraps (ie. $|k_i| \leq 2$). (a) solutions of high consistency error (grey dot) are found further away from intersecting phase lines. If the number of measurement directions is inadequate there are multiple solutions found where all phase lines intersect (ie. the black dots). (b) solutions with equal consistency error lie in concentric circles around the correct solution. The inner (dark grey) circle marks the 1 wrap boundary. The outer (light grey) circle marks the 2 wrap boundary.

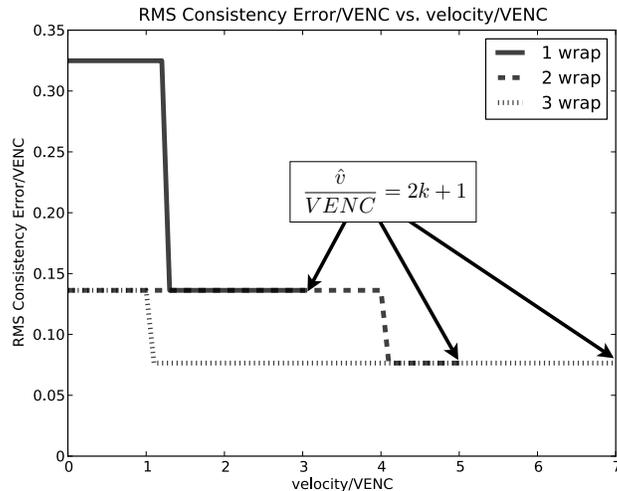


FIG. 6.4: Error sensitivity in terms of consistency error for 6 directions using 1, 2 and 3 wraps (ie. $|k| \leq 1$, $|k| \leq 2$ and $|k| \leq 3$ respectively). The input velocity and output error are normalized by $VENC$. The input velocity noise is zero. The maximum velocity that can be estimated is determined by the number of allowable wraps as indicated by the arrows.

Similarly for 2 and 3 wraps where $|k_i| \leq 2, 3$, there are velocities at which the reconstruction is more sensitive to noise.

The relative benefits and disadvantages of 6, 7, 8 and 9 measurement directions are explored in this work. The minimum number of measurement directions required to have a unique solution (barring any priori knowledge about the measured velocity) in 3D is six. The orientations of the measurement directions are found in the same way as they are for diffusion tensor scans. In order to maximize the orthogonality of the measurement information the measurement directions are oriented using a charge repulsion algorithm suggested in (66) for $N = 7, 8, 9$. For six directions the most homogeneous distribution corresponds to the vertices of a regular icosahedron.

For 2D and 3D, in the presence of noise, the solutions along concentric boundaries do not stay consistent. The error sensitivity thresholds in fig-

ure 6.4 hold relative to each other in the presence of noise and are used to determine the more sensitive velocity to $VENC$ ratios. Measurement noise translates to a shifting of the solution lines in the direction of the measurement (perpendicular to the solution line). The consistency error of each solution is therefore correlated to the error in all other solutions. Characterizing the noise properties of this system is relegated to simulation due to the complexity.

6.2 Experiments

The following work characterizes the proposed method in terms of noise sensitivity, sensitivity to signal loss (i.e. signal biased phase), and signal to noise ratio. The method characteristics are analyzed with respect to the current state of the art, dual-VENC method proposed in (3, 4). Each of the following sections focusses on these characteristics through the specific areas of design and limitation which are the noise sensitivity, voxel shape (signal bias sensitivity), gradient moment directional configuration (SNR), and reconstruction (combined limitations).

6.2.1 Noise Sensitivity

The work presented in this section explores the noise sensitivity characteristics of the proposed method through Monte-Carlo computer simulations (67). The simulations are single point based and do not include signal weighted averaging effects or dephasing. The specific methods and parameters are presented first, followed by an analysis of the results.

6.2.1.1 Methods

The system response to noise was characterized through the use of Monte-Carlo simulations. Individual pixel simulations were solved using equation 6.11 by brute force to eliminate any potential error attributed to a minimization

technique. The input noise was added to each simulated measurement as phase noise (35). This way, the correlated noise in each velocity measurement is consistent with using the same reference (zero moment encoded set). Each pixel simulation is also created with a random input velocity direction.

The velocity noise sensitivity was measured by finding the ratio of the number of incorrectly unwrapped pixels to the number of correctly unwrapped pixels. The number of trials used for each input level of velocity noise (σ_v) changed linearly from 10^7 to 10^5 over a range of $\sigma_v/VENC$ from 0 to 0.05 in increments of 0.001. A line of best fit was used to project the lower bound of the fraction of incorrectly unwrapped pixels. Then the range of $\sigma_v/VENC$ was refined to the limits of each method where their fraction of incorrectly unwrapped pixels lie between 10^{-3} and 10^{-6} . The Monte-Carlo simulation was then re-run over the new range with the same sliding scale of trials for each point.

The simulations were performed with the proposed method for an $N = 6, 7, 8, 9$ and a $|k_i| \leq 1$ and the three-point method proposed in (4). The most noise sensitive velocities were used for each simulation as determined by the data shown in figure 6.4 and similar data for $N = 7, 8, 9$ (data not shown). The velocities for $N = 6$ at $|k| \leq 1, 2, 3$ were $|V| = 2, 4.5, 4.5$ respectively. The velocities for $N = 7, 8, 9$ for $|k| \leq 1$ were $|V| = 2.5$. The high to low $VENC$ ratio used in the three-point method was $VENC_2/VENC_1 = 4$. An additional $N = 6$ was run with a relative $VENC$ of 1.275 to determine whether the noise threshold could be regained if the output SNR were matched to the three-point method. Also an additional three-point method was run with a relative $VENC$ of 0.785 to match the SNR of the $N = 6$ method at a $VENC$ of 1 to characterize the drop in noise threshold. The sensitivity of multiple

wraps was also determined for $N = 6$ at $|k| \leq 1, 2, 3$.

The output SNR was measured by finding the normalized RMS error ($\sigma_{out}/\frac{\sigma_v}{VENC}$) of the difference between the truth and the correctly unwrapped velocity estimate. The relative SNR efficiency is then the quotient of the SNR over the square root of the relative acquisition time for each method. The number of Monte-Carlo at each input noise level was 10^5 . Input noise ($\sigma_v/VENC$) was varied from 0 to 0.1 in increments of 0.01. Simulations were performed with the proposed method for an $N = 6, 7, 8, 9$ and the three-point method with a high to low $VENC$ ratio $VENC_2/VENC_1 = 4$.

For comparison against direct averaging a two-NEX method with a single reference point was also simulated. The $VENC$ for this method is the same as the high $VENC$ used in the three-point method which is $VENC_2 = 4$, since each of these methods require at least one measurement in each direction that is not phase aliased.

6.2.1.2 Results & Discussion

The input noise sensitivity for each method is plotted in figure 6.5. To obtain a pixel error rate no greater than one in one million the $N = 6$ set requires an input noise (σ_θ) of no greater than 6.3° . The threshold at the same error rate for the three-point method is 9.36° . Matching the SNR of the three-point method for $N = 6$ and relative $VENC$ of 1.275, the error threshold is 8.1° . To match the SNR of the $N = 6$ with the three-point method using a relative $VENC$ of 0.785, the error threshold is 7.2° .

The thresholds for $N = 7$ and 9 are both about 1° lower than the thresholds for even numbers of directions $N = 6$ and 8 respectively.

For $N = 6$, each additional increase in the number of wraps lowers the

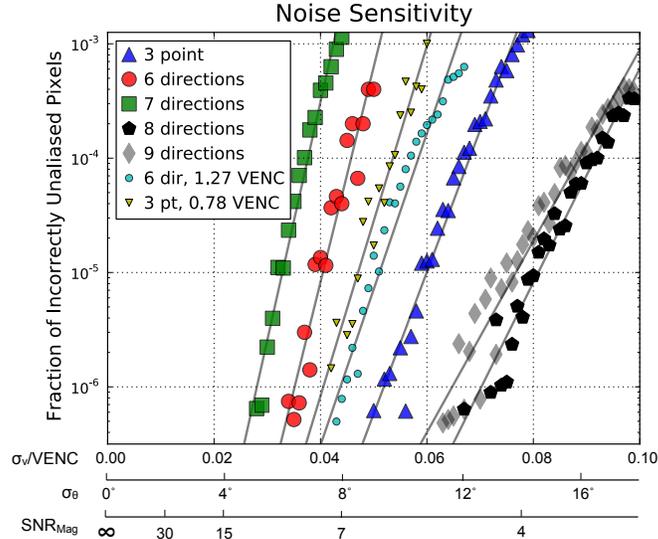


FIG. 6.5: The error sensitivity thresholds of 7 different measurement direction configurations versus input noise. The sensitivity is measured in fractions of incorrectly reconstructed pixels. The input velocity noise is normalized by $VENC$.

error threshold bound by a half of a degree per wrap (data not shown).

The increase in SNR efficiency of the proposed method at $N = 6$ is about 30% higher than the three-point method, as shown in table 6.1. Between $N = 6$ and $N = 7, 8, 9$ the relative SNR efficiency drops by an additional 4% for each added measurement.

Both the proposed and three-point methods exhibit a hard noise sensitivity

Table 6.1: The normalized and effective SNR (time-normalized) using the relative acquisition time for each method.

Method	Relative SNR	Relative Time	Relative SNR_{eff}
2 NEX	1.00	7	1.00
3 point	2.81	7	2.81
6 dir	3.60	7	3.60
7 dir	3.69	8	3.45
8 dir	3.75	9	3.31
9 dir	3.83	10	3.21

threshold, beyond which, each method will precipitously increase in producing erroneous results. However, this region of noise sensitivity lies in an image SNR that is relatively low for all the methods shown in figure 6.5. The $N = 7$ method requires a high relative image SNR of at least 11 and the $N = 8$ method requires a low SNR of at least 4.

The $N = 6$ method is the most time efficient method, of those simulated, in terms of output SNR. This is true for any number of wraps since output SNR is unaffected by the number of wraps in the search space. The three-point method with a $VENC_2/VENC_1 = 4$ allows for 2 wraps to be detected. The input noise sensitivity of the $N = 6$, $|k| \leq 2$ method is increased compared the aforementioned three-point point method by about 3.5° . For these parameters the minimum required image SNR increases by 3.5 over that of the three-point method.

If the output SNR of both the three-point and $N = 6$ methods are set equal, the $N = 6$ method will have increased noise sensitivity by at least 1° (for a wrap of 1). However, this will also reduce the relative moment required by the $N = 6$ method which is expected to increase reconstruction robustness in pixels with signal biased phase.

The SNR efficiency diminishes as the number of measurements increase. At eight measurements the SNR efficiency is still $\sim 14\%$ higher than the three-point method and lowers the noise insensitivity by about 3.2° . Measurements made with $N = 7, 9$ show a higher noise sensitivity which is expected to be related to the orientation of the measurement vectors. Increasing SNR may be better (in terms of lowering noise sensitivity) if a second reference set were collected instead of adding another direction to $N = 6$ or 8.

The brute force reconstruction time for $N = 6$ and $|k| \leq 1$ takes 7 minutes

and a $|k| \leq 2$ takes 138 minutes for a volume of 512^3 pixels on an 8-core Intel Xeon 3.2GHz Mac Pro. The algorithm reconstructs on a per pixel basis which allows straightforward parallelization of pixel volumes. Additionally a smoothness criteria between neighboring \vec{k} could be used to speed up the reconstruction (by limiting the search space) and filter out noise.

6.2.2 Optimal Voxel Shape

This section explores the optimization of the voxel shape in terms of minimizing phase bias through the modification of the modulation transfer function (MTF). The MTF has fixed limits based on the underlying scan trajectory. This means that the resolution cannot be extended beyond the k-space extent collected by the scan. In this case, k-space is spherically supported in 3D, which means the base MTF is as well. The point spread function (PSF) (or voxel shape) resulting from a spherical MTF is a radially symmetric sinc-like function, where a voxel value is the weighted combination of the spin density within the immediate vicinity and neighboring voxels that fall within the side-lobes of the sinc pulse. The affect of such a PSF is manifested as Gibbs ringing around image structures with sharp transitions. However, additional windowing and tapering (i.e. low pass filtering) can be applied to the MTF to change the level of ringing. This is of particular interest in phase contrast where the encoded velocity gradients cause signal loss in some voxels while others straddle the boundary between flow lumen and wall, encapsulating volumes of differing signal levels and phase distributions (68). This section begins by first identifying some physical limitations and then proceeds to corroborate these observations with computer simulations of various MTF and flow profile combinations.

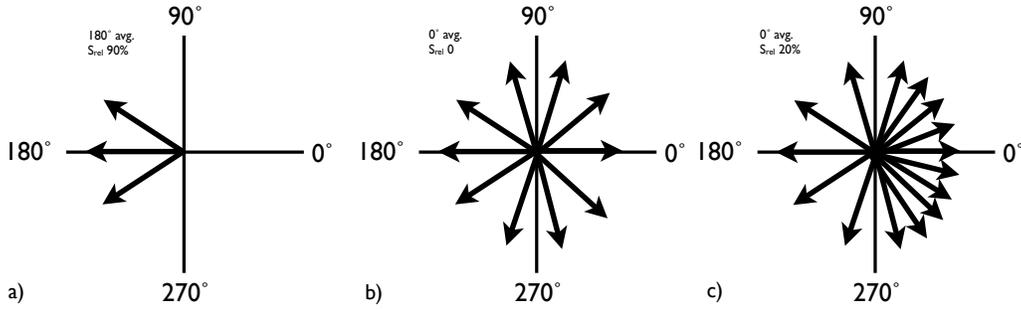


FIG. 6.6: Intra-voxel phase induced signal loss. (a) a 90° linear phase distribution across a voxel will cause 10% signal loss. (b) a 360° distribution causes complete signal loss. (c) a 450° distribution causes a point of coherence, which regains 20% of the signal, however, the average is now centered around zero, causing a misleading estimate of velocity.

A fundamental limitation to recovering the average phase encapsulated by a voxel is the level of signal available to adequately preserve the phase. Intra-voxel phase distributions inherently lower the voxel signal due to phase cancellation. A linearly distributed phase across a voxel will lose signal (S) according to (16),

$$S = \left| \text{sinc} \left(\frac{\Delta\theta}{2} \right) \right|, \quad (6.13)$$

where $\Delta\theta$ is the relative minimum and maximum phase difference.

Figure 6.6 shows three distributions of linearly varying phase of 90° , 360° , and 450° each centered around an example average phase of 180° . While the signal loss in the 90° distribution is only 10%, the remaining signal in the 360° distribution is zero because all of the phase vectors cancel. An additional 90° of phase will again increase the signal however, the average phase will then be centered around zero yielding a type of phase aliasing that is unaccounted for by the previously described unaliasing method. This effectively makes the intra-voxel phase distribution limit less than 360° .

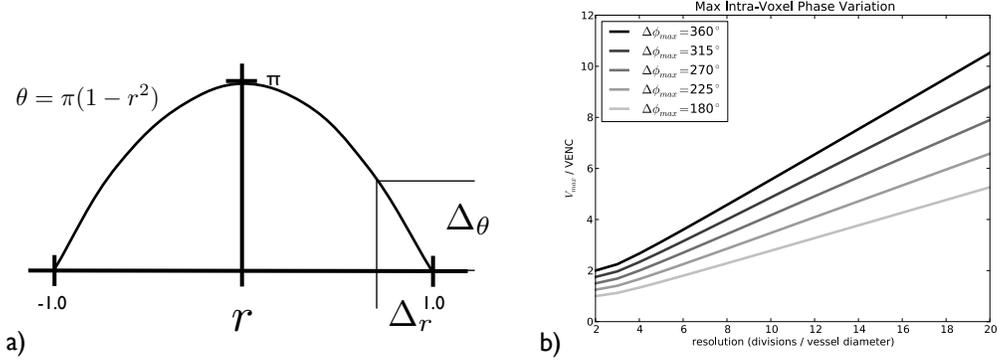


FIG. 6.7: Intra-voxel phase distribution model and limits. (a) a parabolic model of a laminar flow distribution. (b) level curves relating maximum velocity to VENC ratio and resolution to the maximum phase contained in the edge voxel.

In 1D, this limit can be easily calculated for a parabolic flow model (figure 6.7(a)),

$$\theta = \pi \frac{V_{max}}{VENC} (1 - r^2) , \quad (6.14)$$

where V_{max} is the velocity at the parabolic apex, N is the number of divisions across the lumen diameter,

$$N = 2r_{max} / \Delta_r , \quad (6.15)$$

r_{max} is the radius of the profile, and Δ_r is the physical resolution. Combining equations 6.14 and 6.15 yields the maximum phase distribution,

$$\Delta\theta_{max} = \frac{\frac{V_{max}}{VENC}}{(1 - (1 - 2/N)^2)} , \quad (6.16)$$

which occurs at the edges of the flow lumen.

Setting equation 6.16 to a specific phase allows level curves to be drawn

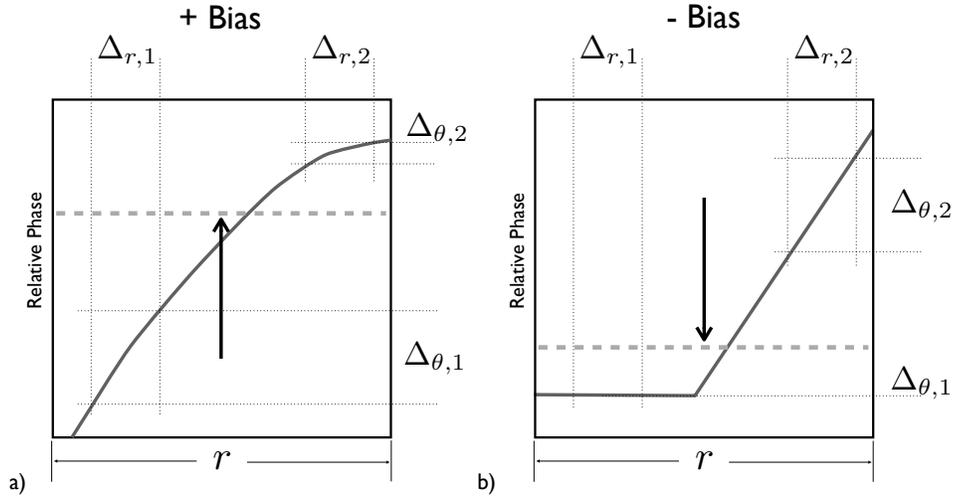


FIG. 6.8: Additive and subtractive signal biased phase. (a) a concave down segment, with respect to the sign of the velocity, causes an increase in estimated speed. (b) a concave up profile causes a decrease in estimated speed.

for various $V_{max}/VENC$ ratio and resolution combinations, as shown in figure 6.7(b). While this model is only 1 dimensional it yields an approximate limit on the $V_{max}/VENC$ ratio, and provides some intuition into the sensitivity of the relationship between the ratio and resolution. As the resolution increases the maximum intra-voxel distribution decreases rapidly, but not until a resolution of 4 divisions over the lumen diameter is reached.

Signal bias occurs when partial volumes of differing signal and phase distributions are contained within a single voxel. Figure 6.8(a) shows an intra-voxel distribution similar to what would be found near the apex of a parabolic flow distribution. If the voxel is further subdivided (Δ_r) it is apparent that within division $\Delta_{r,1}$ and $\Delta_{r,2}$ the phase distributions ($\Delta_{\theta,1}$, $\Delta_{\theta,2}$) are not equal. In this case $\Delta_{\theta,2}$ has a smaller distribution and according to equation 6.13, will lose less signal than $\Delta_{\theta,1}$. The average phase within this voxel is a signal weighted average based on the shape of the PSF and the spatial signal loss,

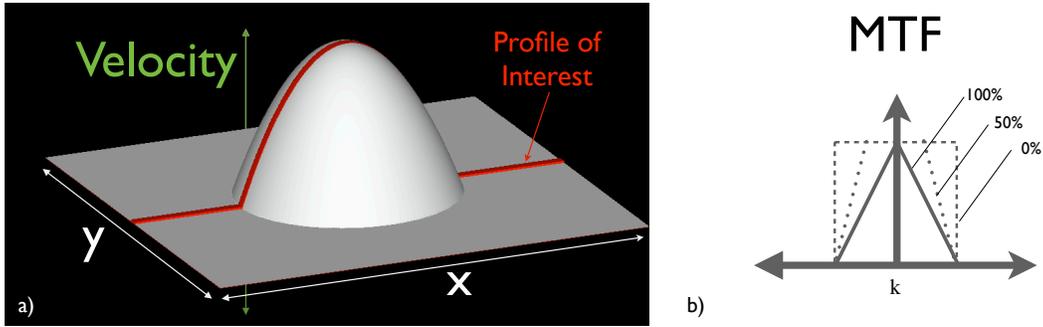


FIG. 6.9: A 2D cross-section of a 3D laminar flow model (a) and corresponding MTF filters (b) used in simulation. The profile of interest is a 1D cross-section that spans the center of the 2D flow profile in y . The 1D profile of the spherically symmetric window functions is set to one of three levels.

therefore, the average phase within $\Delta_{r,2}$ will impart a larger contribution to the voxel average than the average phase within $\Delta_{r,1}$. Since the average phase in $\Delta_{r,2}$ is higher than $\Delta_{r,1}$, the absolute average phase ($|\theta|$) for this voxel will be inflated, in comparison to the phase that the voxel is centered upon.

An example of a voxel encompassing flow lumen and wall partial volumes is shown in figure 6.8(b). This is similar to the first example except that the wall signal is based on the acquisition parameters and not on flow velocity. The encoded velocity in the wall volume is zero and therefore causes a negative bias in the estimated absolute phase ($|\theta|$).

6.2.2.1 Methods

In order to observe the combined effects of signal weighting and voxel shape, a 3D simulation was used to aid in identifying these characteristics. This section covers the design parameters used in this simulation.

The simulation is that of laminar (parabolic) flow through a 3D tube oriented parallel to the direction of flow along the z axis. Figure 6.9 shows a 2D cross section of the 3D simulation volume. The flow is a radially symmetric

parabola using equation 6.14 where $r = \sqrt{(x^2 + y^2)}$.

The parabolic phase profile is combined with a constant signal profile across the flow lumen and a constant signal profile across the wall volume that is a variable percentage of the lumen signal. This is done at a sampling factor 10 times the chosen diameter of the lumen onto an isometric 3D data volume. The complex data volume is then fourier transformed (via FFT) into k-space and the MTF is filtered by a spherically symmetric linear taper as shown in figure 6.9(b). The filter was set to three levels 100°, 50° and 0° which corresponds to no additional filter, only the symmetric rectangle window that corresponds to the acquisition MTF. A linear phase is added across the kx direction to provide sub pixel shifts in the image. Since the MTF created by the acquisition trajectory is spherically symmetric the 3D volume is shifted only along the x axis to simulate partial volume effects. The matrix is then zero padded before fourier transforming back.

A 1D profile, shown in red in figure 6.9(a), of a mid level slice is searched for the maximum signal bias. The original k-space simulation is kept so that linear phases corresponding to one tenth of a pixel shift are successively added and subsequent profiles searched for the max of the maximum biases. This search is used for each parameter configuration, such as $V_{max}/VENC$ ratio, resolution, relative wall signal, and filter.

6.2.2.2 Results & Discussion

The potential for minimizing bias by changing the voxel shape, through low pass filtering, is limited to a short range in resolution, velocity to VENC ratio, and specific relative wall signal. The plot in figure 6.10(a) shows that for resolutions between 5 and 7, and a velocity to VENC ratio of 2, the bias may

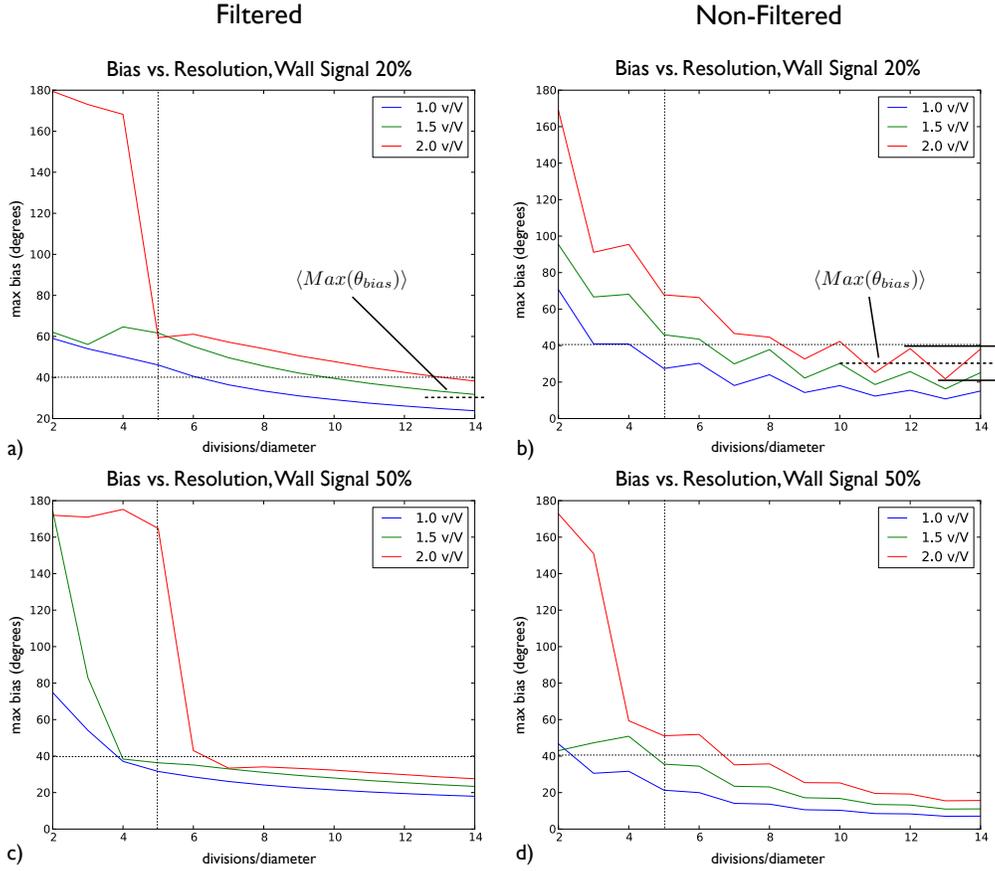


FIG. 6.10: The bias effects of voxel shape in terms of resolution and low pass filtering. (a) and (b) show max bias using filtered and non-filtered simulations with a wall signal set to 20% of the flow signal. (c) and (d) show filtered and non-filtered simulations with a wall signal set to 50%. Each plot contains bias curves at three different levels of V_{max}/V_{ENC} ratios indicated as v/V . The data used in these sets were zero-padded by a factor of 2 in all 3 directions.

be moderately mitigated as compared to the non-filtered case (figure 6.10(b)). Filtering is shown to increase bias at the advantage of reducing the amount of Gibbs ringing (69). However, filtering also significantly raises the bias for resolutions below 5 pixels for a velocity to VENC ratio of 2.

The profiles in figures 6.10(c) and (d) have an elevated relative wall signal compared to (a) and (b). An increase in wall signal appears to increase bias in low resolutions (< 6 divisions per diameter) and decreases bias in high resolutions (> 8 divisions per diameter). This is an expected result considering the mechanics illustrated in figure 6.8. low res At low resolution the ringing is higher, due to reduced spatial frequencies and higher signal losses (which works to square the high and low edges of the profile). If the wall signal is lowered, then the bias due to partial volume is lower. The edges of the profile are not as sharp as a result, which also leads to less ringing. The less wall signal there is, the lower the amount of signal bias due to partial volume, and Gibbs ringing becomes the dominant biasing factor. As the wall signal is increased the ringing is decreased and the signal bias is increased. high res At high resolution the profile retains its shape. There is less signal loss because the edge voxels do not span large phase distributions. There is less signal bias because the edge voxels do not cover as much partial volume. These factors allow the profile to remain smooth, which reduces ringing.

For high spatial resolutions, filtering does not appear to be advantageous. The average max bias for a non-filtered, 20% relative wall signal, is shown as $\langle Max(\theta_{bias}) \rangle$ in figures 6.10 (b). Translating this measurement to figure 6.10(a) shows that although filtering equalizes the maximum bias between neighboring resolutions (between 10 to 14), it increases the average maximum bias to the max maximum bias.

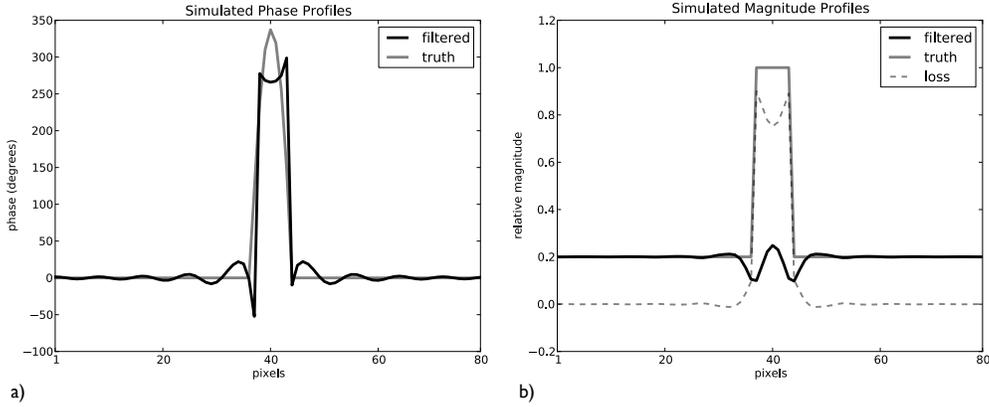


FIG. 6.11: 3D simulated profile cross sections. (a) a 1D cross section of a 3D phase profile both filtered and analytically calculated (truth). (b) 1D cross sections of the corresponding magnitude profiles. The edge pixels span a 360° phase distribution. The profile is shifted by a sub-pixel amount causing higher levels of bias on the lower left and upper right corners of the parabolic profile. The areas of high bias in (a) are demarcated by corresponding areas of high signal loss shown in (b). The data used in these sets were zero-padded by a factor of 4 in all 3 directions.

The signal mechanics presented in figures 6.6, 6.7 and 6.8, are exemplified in simulation using the limits presented in figure 6.7.

Figure 6.11 shows a profile with traits characteristic of the signal bias as seen in-vivo. For a relative wall signal of 20% and a resolution of 2 divisions across the lumen diameter, the signal loss, partial volume, and Gibbs ringing transforms the parabolic profile to a lower resolution square profile. The profile was zero-padded, in 3D, to a factor of 4 to accentuate these features.

The profiles show how signal bias and Gibbs ringing positively influence each other. At the base of the profile, the signal bias subtracts from the phase value bringing the value closer to the phase value of the wall. While the limit of this effect would result in an estimated edge voxel value equal to the wall value (i.e. zero), the edge voxel has a phase of -50° . The extra subtracted phase is due to ringing. Similarly for the top of the profile, the limit of the

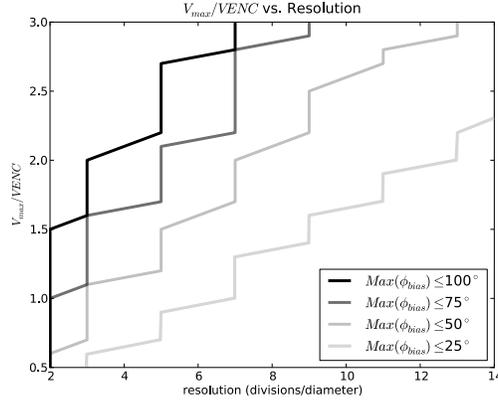


FIG. 6.12: Level curves of maximum bias in the 3D simulated profiles. The $V_{max}/VENC$ resolution is 1/10 versus integer numbers of divisions per lumen diameter. The stair stepped pattern is due to the integer pixel resolution used to minimize the search. The relative wall signal is 20% of the lumen signal.

bias on the edge pixel would be the value of its neighboring pixel (closer to the apex), however it extends higher because Gibbs ringing is additive to offset the negative lobe at the lower edge. The shift in the profile pronounces this affect in both the top right and bottom left edges. This is because the shift focusses the signal drop to be partially volumed with either high average phase distributions or low average phase distributions respectively.

The signal loss profile (figure 6.11(b)) shows that there is high loss across the simulated flow region. The highest losses are concomitant with the largest phase biases, which appear at the edges.

Figure 6.12 shows the relationship between the velocity to VENC ratio and the resolution in terms of maximum signal bias using a non-filtered MTF. This shows that velocity unaliasing techniques that extend past one wrap (i.e. $|k| > 1$) must be robust toward potentially high levels of bias if the resolution is relatively low. This further advocates the idea of signal loss based replacement strategies (3), for techniques a that operate in the low resolution high velocity

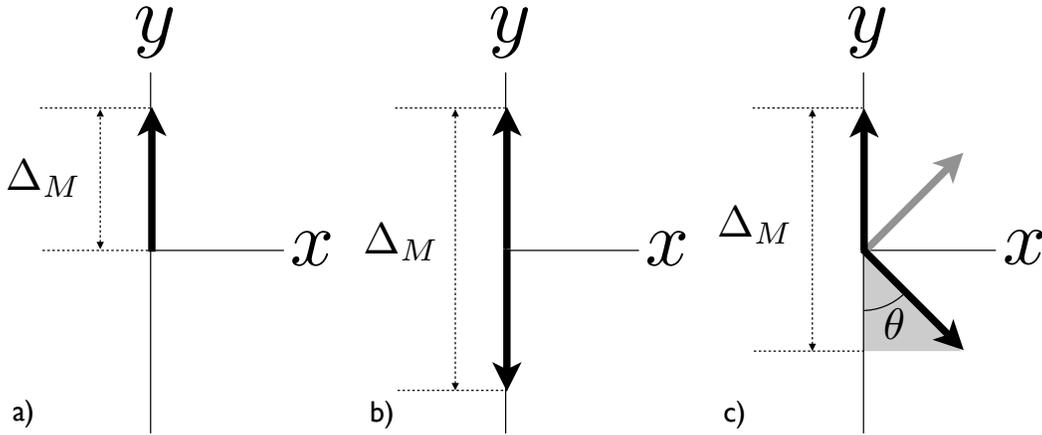


FIG. 6.13: An illustration of moment addition in balanced and semi-balanced cases. (a) a single referenced velocity measurement (two measurements required). (b) two moment balanced measurements (two measurements required). (c) partially balanced measurements (three measurements required).

regime. The bias cusps shown in the profile of figure 6.11 also indicates that full replacement may be necessary in areas of high bias such as these.

6.2.3 SNR Optimized Moment Balancing

This section provides background on the basic theory of moment balancing in phase encoded imaging and how this effect benefits the proposed 6-direction encoding scheme. Methods for the optimization of moment balancing in the proposed 6-direction scheme are discussed and explored in the following subsections.

Moment balancing is a method of effectively extending the gradient area of the bipolar encoding preparatory pulses by making two separate measurements of opposing polarity (34, 70). Figure 6.13(a) graphically depicts the moment difference (Δ_M) between a velocity encoded measurement and a reference measurement. This is the basic moment encoding technique which requires at least two measurements, resulting in a noise reduction by a factor of $\sqrt{2}$ due to the signal averaging of the two measurements (4, 35). This will be referred to as

a ‘referenced measurement’. Figure 6.13(b) shows the effective difference in moment if two measurements are made with opposing bipolar gradient areas. In this case, where the moments are directly opposing each other, a reference measurement is not necessary, however if one is provided, then its noise contribution cancels when the two referenced phases are subtracted from one another (70). The total moment is extended by a factor of two, with signal averaging the VNR is twice that of the single referenced measurement. Figure 6.13(c) is a hybrid of the two aforementioned cases and more closely exemplifies the proposed method. The black arrows represent partially balanced moments. A reference measurement is necessary to resolve the encoded phase for each measurement, but the correlated noise is partially canceled by an amount relative to the cosine of the angle θ . If the measurement direction indicated by the grey arrow was used instead of its reflected counterpart, there would be little to no cancelation of the noise in the reference measurement used by both.

In the proposed method, the angular distance between each measurement direction is maximized in order to maximize the level of independent information provided by each measurement. A relative quantification of independence can be determined by finding the condition number of the measurement direction matrix u (66). As detailed by Hasan et al., the optimal condition number can be determined through the minimization of force on bound charged particles. If the measurement directions u_i are considered charged, the relative energy E to be minimized can be calculated as,

$$E = \sum_{i=1}^{2N} \sum_{f>i}^{2N} \frac{1}{\|u_i - u_j\|} \quad (6.17)$$

where N is the number of measurement directions. Since the u matrix is fixed,

the moments must be balanced for each principle direction, simultaneously.

Moment balancing in the proposed method is analogous to providing the solution directions with the most measurement coverage. Similar to the ideas presented in (34), there are areas of the measurement geometry that are less supported, meaning the effective VENC in that direction is the highest. In this case there are two aspects to optimize: 1) ensure the solution direction is well covered by the measurements and 2) ensure the solution direction is centered on the maximum amount of reference noise cancellation (figure 6.13(c)).

6.2.3.1 Methods

This section covers the algorithms used to optimize measurement coverage and moment balancing as well as the experiments used for validation.

Similar to equation 6.17, the distance between the solution vectors and the measurement vectors must be minimized while maintaining the relationship encoded in u by minimizing the net moment in each of the principle axes. The net moment vector $\Sigma_M^{\vec{}}$ on each of the principle axes is the row-wise summation of u calculated as,

$$\Sigma_M^{\vec{}} = \sum_{i=1}^N \vec{u}_i . \quad (6.18)$$

The minimization of the net moment can be calculated in a number ways to optimize different aspects of the solution. Here we explore the minimization of the root (M_{RSS}) sum of squares of the net x, y and z moments,

$$M_{RSS} = \sqrt{\Sigma_{M,x}^2 + \Sigma_{M,y}^2 + \Sigma_{M,z}^2} , \quad (6.19)$$

where $\Sigma_{M,\{x,y,z\}}$ are the net x, y and z moments, and the max absolute net gradient moment ($Max(|\Sigma_M|)$). Each moment configuration is added to the

energy minimization of equation 6.17 as an additional cost function make the solution,

$$Min \left(\sum_{i=1}^{2N} \sum_{f>i}^{2N} \frac{1}{\|u_i - u_j\|} + \alpha \cdot M_{RSS} \right) , \quad (6.20)$$

or,

$$Min \left(\sum_{i=1}^{2N} \sum_{f>i}^{2N} \frac{1}{\|u_i - u_j\|} + \alpha \cdot Max(|\Sigma_M|) \right) , \quad (6.21)$$

where α determines the amount of ‘flexing’ a measurement is allowed, which will be discussed later. Forcing the net moment to zero ensures that the gradient area in each direction on the solution axis is balanced, automating the trade-off between projected moment and reference noise cancellation.

As previously discussed, the u measurement matrix has fixed angular distances between each measurement. The component that is not fixed is the measurement sign. The sign of each measurement must be set to maximize the moment balancing and the ability to adequately cover each solution direction. For 6 directions, the number of sign combinations is 64 (2^6). Starting the minimization of equation 6.20 or 6.21 with a particular sign pattern will allow the optimal configuration to emerge in fewer iterations.

Table 6.2 shows two solutions (one for each equation 6.20 and 6.21) and an unbalanced moment encoding matrix. The unbalanced moment set is chosen to group all the measurements in one direction (as shown in figure 6.14(a)). The moments are hand selected by choosing vertices of an icosahedron directly (66). In this case the net moment is perfectly balanced in x and becomes increasingly unbalanced for y and z. The net moment resulting from equation 6.21 is fairly balanced between each principle axis. The cost function forces

Table 6.2: Measurement direction matrices (u) and moment sums for each principle direction. Each row represents an ordered triplet corresponding to the measurement direction. The x, y and z components also represent the relative gradient moments used in each of the physical gradient directions.

Unbalanced ^{1,2}			~Balanced			Min(RSS)		
x	y	z	x	y	z	x	y	z
0	$1/\alpha$	θ/α	0.641	-0.393	-0.659	0.916	-0.019	-0.400
$1/\alpha$	θ/α	0	-0.406	0.881	0.241	0.194	0.694	-0.693
θ/α	0	$1/\alpha$	-0.587	-0.643	-0.490	-0.595	-0.761	-0.256
0	$-1/\alpha$	θ/α	-0.207	0.146	0.967	0.437	-0.891	-0.121
$-1/\alpha$	θ/α	0	0.318	0.795	-0.516	-0.797	0.212	-0.565
$-\theta/\alpha$	0	$1/\alpha$	0.966	0.007	-0.257	-0.030	0.313	0.949
Relative moment sum for x, y, and z columns ($ \sum \gamma \Delta_M $).								
0	1.701	2.753	0.724	0.7931	0.603	0.124	0.450	1.087

¹ $\theta = 1 + \sqrt{5}/2$
² $\alpha = \sqrt{1 + (1 + \sqrt{5}/2)^2}$

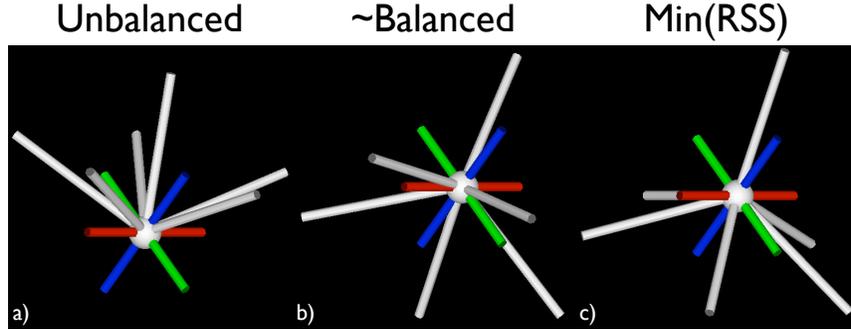


FIG. 6.14: Moment balanced measurement vector configurations. (a) the unbalanced configuration uses two physical gradients per projection. (b) the nearly balanced gradient moments has an equal net moment in each principle direction. (c) the minimum root sum of squared moments provides asymmetric moments in each principle direction.

the maximum moment down causing each of the net moments to balance. The choice of α was empirically determined to be 0.1 and corresponds to a maximum deviation of 1° from the original u matrix. The resulting moments from minimizing equation 6.20 are not equally balanced, however, they are significantly lower than those of the unbalanced set.

The moment sets calculated in table 6.2 were used in phantom experiments to determine optimality and demonstrate the characteristics of each method.

A flow phantom and pump combination were used to simulate laminar flow conditions for each acquisition. The phantom used in these experiments consisted of a polyimide tube 10cm in length, 10mm in diameter and approximately a 0.1mm thick wall. The pump used was a Compu-Flow 1000 MR, from Shelly Medical Imaging Technologies, London, Ontario. A 60% glycerine and water mix (by volume) was pumped through the phantom at a rate of 5mL/s to create a parabolic flow peak velocity of approximately 35cm/s . The flow direction was aligned with the main magnetic field (B_0), which is in the direction of the z axis gradient moment. The scanning platform used was a 3Tesla General Electric Signa Excite Twin Speed MR scanner. The phase noise was measured in each experiemnt at the same position within the glycerine bath area (71).

Acquisitions of each of the gradient moment configurations and the dual-VENC method were scanned using a VENC of 20cm/s , 40cm/s , 60cm/s and 80cm/s . Image reconstructions for the dual-VENC set were conducted using a high-VENC 80cm/s acquisition. The proposed 6-direction method was reconstructed three ways: 1) using no phase unaliasing algorithm, 2) using phase unaliasing, and 3) velocity thresholded phase unaliasing with $V_{max} = 35\text{cm/s}$. The imaging scan used was a basic 3D stack of spirals with slab selective RF.

Each TR acquired a single Archimedean spiral (72). The same spiral was collected 7 times, one for each moment encoded direction including a reference set, before moving to the next spiral trajectory. The scan parameters were as follows: 12 z-phase encodes, 33 spiral interleaves per plane, $TR = 16ms$, $TE = 5.5ms$, and a flip angle of 15° . The receive gain and linear shim were held constant over all acquisitions. The receive coil used was a standard quadrature head coil.

6.2.3.2 Results & Discussion

The following section covers the results of the proposed phantom experiments and provides a small discussion about each of the findings.

Noise measurements from the flow phantom experiments are plotted in figure 6.15. Plots from sub-figures 6.15(a) to (d) show the noise measured in each solution direction (x, y and z). The 6-direction plots include the average noise of the 3-direction-referenced acquisitions as a black fit line to serve as a reference for the relative noise produced in the dual-VENC method. In comparison of the three moment configurations, the unbalanced shows the largest distribution of noise in each solution while the semi-balanced method produces the lowest average distribution.

The relative noise distributions within each solution also corresponds to the relative net moment. As seen from table 6.2 the x, y, and z net moments for the unbalanced configuration are 0, 1.7 and 2.7. The x net moment is exactly balanced and results in roughly a factor of 2 reduction of noise as compared to the average 3-direction referenced set. The z moment of 2.7 produces the highest amount of noise which is a little more than twice the noise in the x solution. The y moment is almost half of the z moment and produces a noise

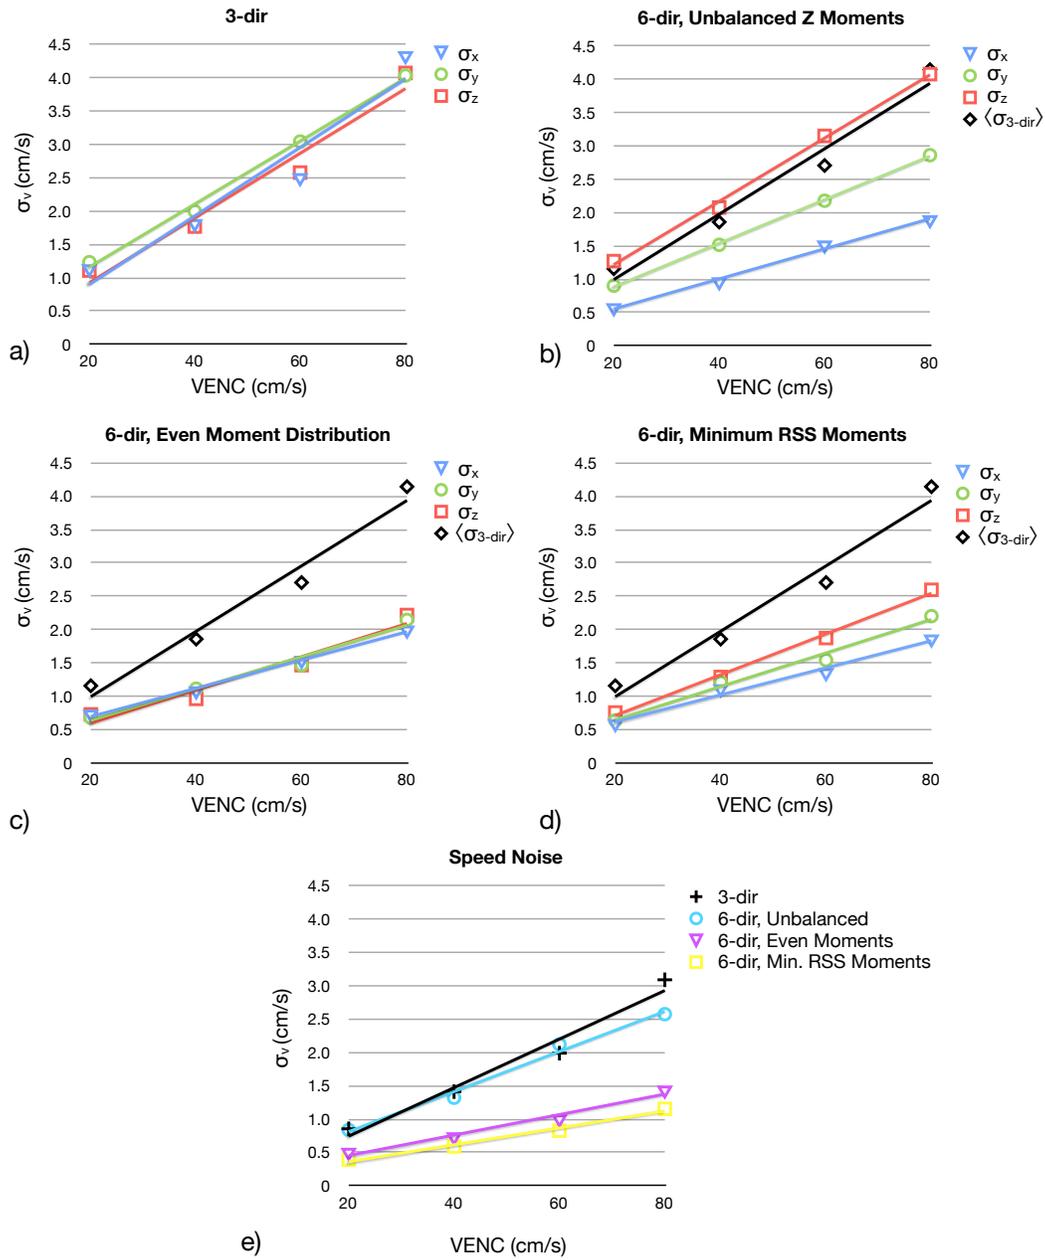


FIG. 6.15: Measured noise produced in various moment balancing strategies. (a) the output noise for a 3-direction Dual-VENC acquisition at several VENC levels. (b) unbalanced, minimum gradient, 6-direction acquisition. (c) nearly balanced even moment distribution across the 3 principle axes. (d) minimum RSS moment balancing. (e) noise measured in the speed image.

that is also roughly half the noise in the z solution.

The semi-balanced moment configuration shows the best overall performance with noise reductions of about 2 compared to the referenced measurements, in each solution. This is particularly interesting as a comparison to the dual-VENC method. Assuming the high-VENC is much larger than the low-VENC in the dual-VENC setup, the high-VENC will contribute a negligible amount of signal to the solution, the dual-VENC solution will produce the same level of noise as the referenced set plotted in figure 6.15(a). This assumption is appropriate for a high flow distribution where the high to low VENC ratio is high enough to unalias phase. The achievable gain in VNR made by the proposed method, in comparison, is as high as a factor of 2. The gain is also limited to a factor of 2. In event that the VENC is greater than the maximum flow velocity (i.e. there is no phase aliasing), the measurement vectors can be placed in perfectly moment balanced pairs on each solution axis (figure 6.13(b)). As shown in (70), this results in maximal cancellation of the reference noise and maximal moment addition.

The ‘flexing’ of the measurement directions as a result of choosing an α greater than zero in equations 6.20 and 6.21, appears to be inconsequential in the improvement of SNR. As figure 6.15(b) shows the x solution achieved a maximal noise reduction without the need for changing the relative phase angles between measurements. Comparing the net moments between the unbalanced and semi-balanced methods it can be deduced that the level of noise in each solution affects the level of noise in all solutions. It is hypothesized that the unbalanced method sacrifices the potential noise reducing contribution of the y and z solutions to perfectly balance x , resulting in the maximum noise reduction. While the semi-balanced configuration offers relatively high net

moments, they are almost equal, producing higher distributed noise reduction characteristics.

Measuring the noise in the speed images yielded some insight as to the definition of optimality. As shown in 6.15(e) the unbalanced method makes a marginal reduction in maximum solution noise but is hindered by the high level of noise in the y and z solutions. The RSS net moment solution is shown to have slightly lower noise than the semi-balanced method. Since the speed noise is more Rician due to the square operation in its calculation, its difficult to determine by these means whether the RSS set is higher in quality.

Mid level slices of the reconstructed flow phantom experiments are shown in figure 6.16. The Dual VENC reconstruction was omitted between VENC of 80cm/s to 40cm/s since they were equivalent to the single direction referenced reconstruction. The relative noise level between the single referenced direction and the proposed method is visibly apparent. This figure displays only the semi-balanced moment configuration of the proposed method, so the relative noise level between the two methods is a factor of 2. As quantified in figure 6.15(c), the velocity noise to VENC ratio (as measured by the proposed method) is about 0.05. The noise sensitivity threshold estimated in the previous section (figure 6.5) shows that this level of noise should result in unwrapping errors of roughly 1 in every 1,000 pixels for any of the chosen VENC. This corresponds to roughly 10 erroneously unwrapped pixels over the shown 100^2 image. The unaliased reconstructions in the second to last row of figure 6.16 show 1 to 2 incorrectly reconstructed pixels in the 60cm/s and 80cm/s images respectively. The 20cm/s reconstruction shows 11. This suggests that the estimated noise threshold is accurate and that it may be further improved upon by thresholding out solutions that exceed a maximum speed (as shown

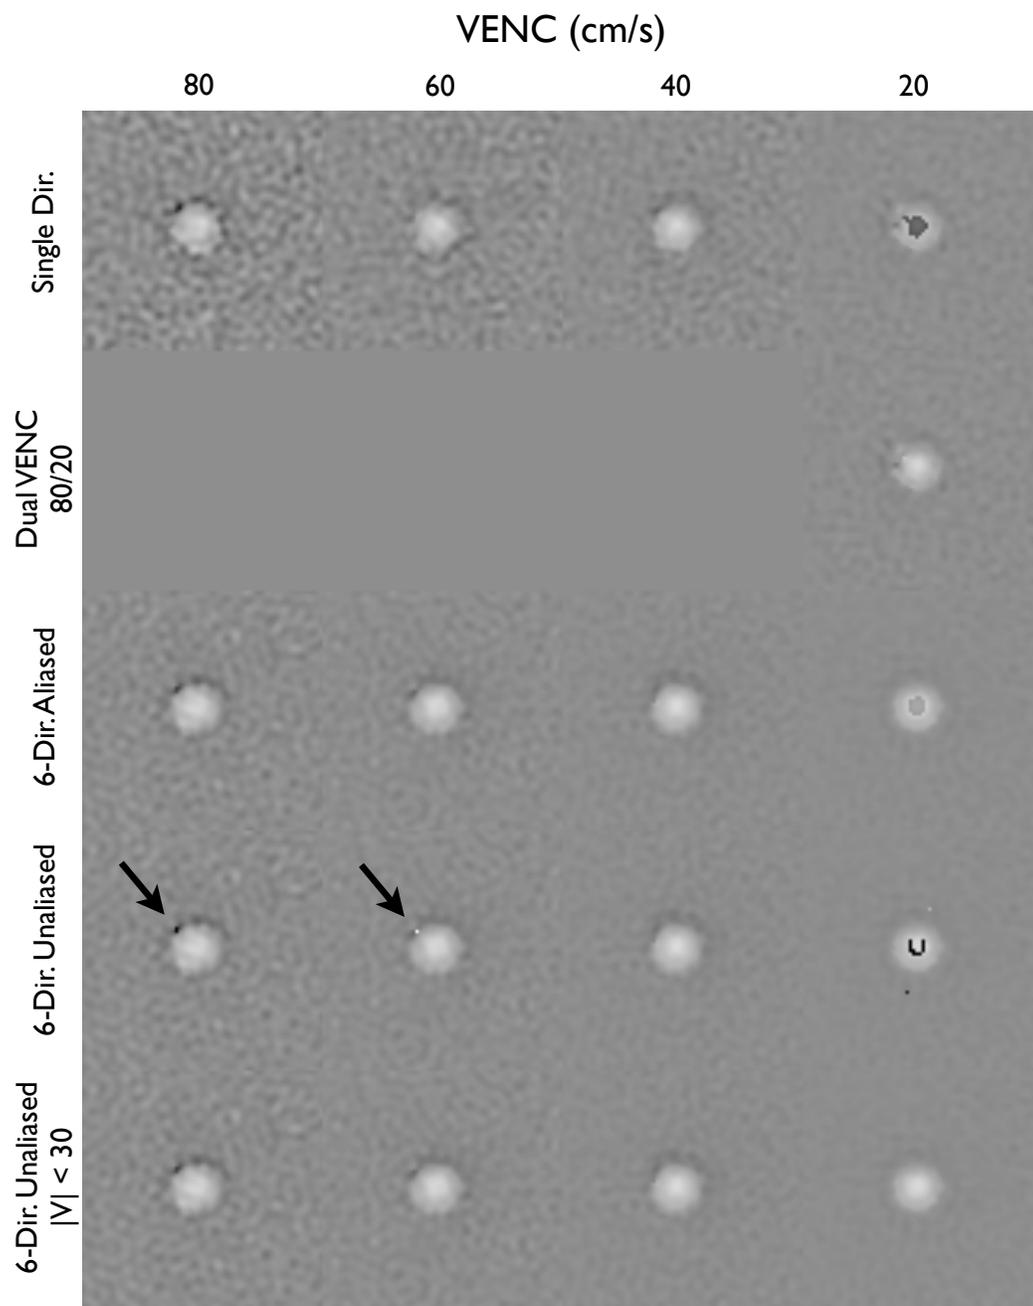


FIG. 6.16: Flow phantom experiments comparing dual-VENC and 6-direction reconstructions.

in the bottom row).

The aliased reconstructions shown in the far right column of figure 6.16 give some indication of the level of redundancy in the 6-direction method. Compared to the aliased single direction referenced set on the top right, the 6-direction aliased set contains an average of the aliased directions (more parallel to the flow direction) and more orthogonal projections of neighboring measurements. This is more noticeable in the 1D cross section of the unaliased reconstructions shown in figure 6.17. The velocity profiles match more closely in areas of low dephasing, and become more disparate in areas of high dephasing next to the lumen edge.

The sensitivity of the dual VENC method is independent in all three directions while the reconstruction of all three directions in the proposed method are interdependent. This means that unwrap errors are caused in all solution directions. This is another reason why it is potentially more valuable to distribute the net moment evenly.

The SNR gain afforded through moment balancing allows the proposed 6-direction method to be acquired at half the gradient moment compared to the dual-VENC method. As shown in the previous section, this can provide a significant reduction in the amount of phase bias (figures 6.12 and 6.7(b)).

6.2.4 Constraints & Reconstruction

The previous sections have outlined the limitations of the proposed method using the 6 direction moment configuration. These limits have been demonstrated using computer simulations and flow phantom experiments. The work presented in this section unites these concepts with real world limitations on SNR and resolution using the rapid 3D SPI MR imaging technique. A neuro-

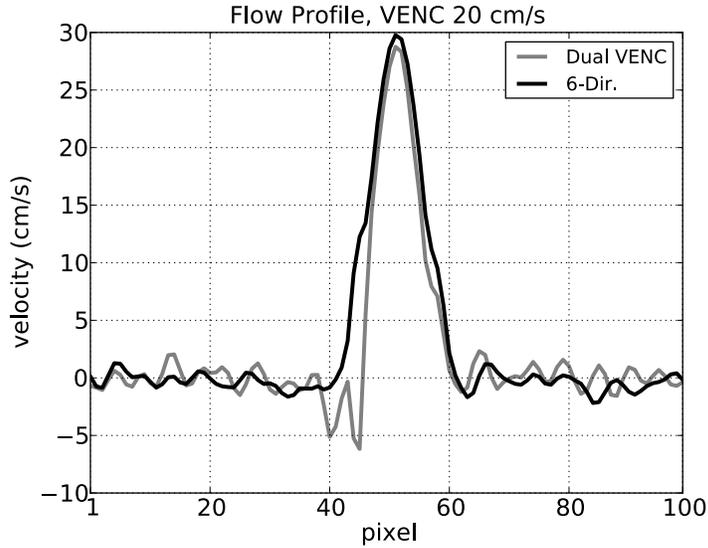


FIG. 6.17: A 1D cross section of flow for both 6-direction and dual-VENC acquisitions at a 20cm/s VENC.

angiography exam will be the application used to demonstrate these real world limitations. This section examines the error sensitivity in the proposed unaliasing algorithm and proposes constraints and reconstruction methods to reduce the error in the final reconstructed image.

As shown in the previous sections, the signal bias phase can easily exceed the phase noise in a high moment velocity encoded measurement.

As previously discussed, the bias can either add or subtract from the true average intra-voxel phase distribution. Figure 6.18(a) is a digram of the solution space for a 2D signal biased example. In 2D, four measurements are required to resolve the phase aliasing. In this example the bias subtracts from the average phase causing an underestimation of the velocity magnitude. The correct solution lies at the point indicated by the red arrow (the example velocity vector). Other solutions are indicated by a solid triangle shaded blue for high error solutions and red for lower error solutions. As the bias increases

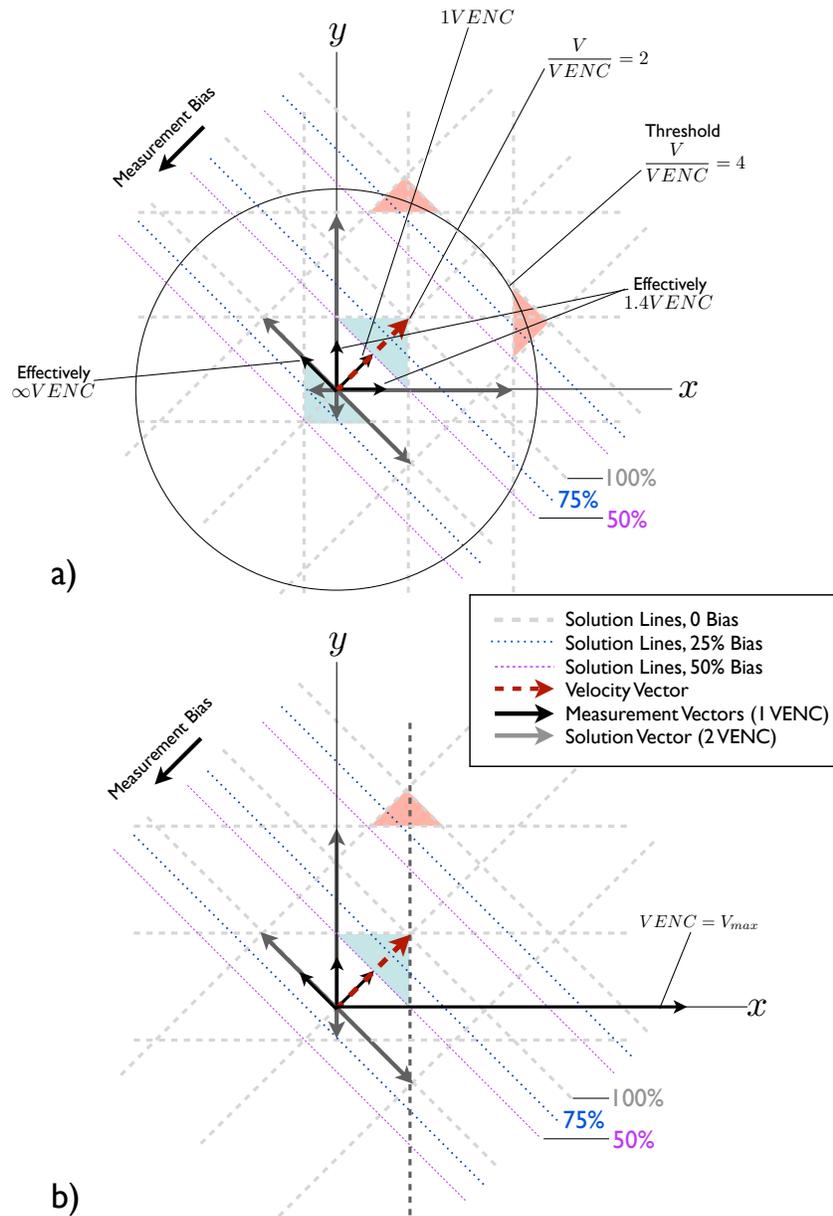


FIG. 6.18: The sources of phase bias sensitivity in 2D low moment encoding. (a) the measurement bias in one direction forces solutions first at the boundaries of the solution space (red) and finally to a solution of the opposite direction (blue). (b) increasing the VENC of one of the measurements increases stability by removing one of the red solutions and one of the blue solutions.

the measurement that is parallel to the velocity vector will experience high signal loss and will start to become underestimated. By following the blue and magenta solution lines it can be seen that the underestimated velocity causes the solution error to increase near the correct solution, and decrease near other possible solutions. The first solutions to become more consistent than the correct solution are shaded in red, which correspond to solutions that are roughly twice the velocity magnitude of the correct solution. These areas become comparable between losses of 25% to 50% phase bias. After a 50% bias in phase (magenta line), which, in this case, corresponds to a 180° phase difference, the most consistent solution becomes the blue triangle on the lower left, causing the estimated velocity magnitude to be half of the original in the opposite direction. The low angular distance between the neighboring measurements increases the threshold that the bias must overcome to make the solution in the opposing direction (the lower blue triangle) more consistent .

The solution space is limited by the maximum number phase aliasing wraps that are expected to occur. This solution space can be further constrained by choosing a maximum velocity V_{max} outside of which solutions are ignored. The circle, centered at the origin (zero velocity), has a radius of 4 times the VENC level. This radius bisects the low consistency error solutions marked in red. A slight decrease in the radius will block out the red shaded solutions completely, making system robust to higher levels of bias. In this case, since only one wrap is allowed ($|k| \leq 1$), the highest isotropic maximum speed would be $3 \cdot VENC$. This would omit solutions that might occur in directions that angularly bisect the measurement directions (where the effective moment is lowest). Thresholding makes the algorithm more robust to both bias and noise for the same reason. The maximum level of tolerable bias is lower in

practice due to the addition of thermal noise.

In 3D using 6 directions, the sensitivity of a corresponding case, where the high magnitude velocity vector is parallel to one of the measurement directions, is higher between each measurement because the angle between is larger. This lowers the level correlated information. However, there are more neighboring measurements which increases the average level of information. This is analogous to moment balanced presented in the previous section. This is advantageous in 3D since the signal bias is stronger in measurements that are closer to being parallel with the velocity vector.

As a result of equation 6.13, and the intra-voxel signal distribution (due the level of partial volume), the neighboring measurements will not be biased by an amount equal to the cosine of the angle between. Assuming a laminar model and 3D spherical MTF, the relative levels of bias for each measurement can be approximated using the plot in figure 6.12. The effective VENC of neighboring measurements are proportional to the cosine of the angle between. The bias sensitivity is dependent on the measurement geometry, the solution velocity threshold, velocity to VENC ratio, and resolution.

Proportional biasing between measurements is only true in a few cases, such as the one for the 2D setup (figure 6.18) where the measurement direction is exactly parallel with the flow, the neighboring measurements are symmetrically distributed around the velocity vector, and the model, PSF and wall signal are of that specified for the data in figure 6.12. The V_{max} of a parabolic flow profile can be approximated using equation 6.14. The V_{max} to VENC ratio can then be used to estimate the bias in the parallel and neighboring measurement directions using the level curves plotted in figure 6.12. The solution to the 2D problem is then consistent under all of these parameter assumptions, if

the resolution is low enough to maintain the phase bias relationship between a parallel $V_{max}/VENC = 4$ and an angled $V_{max}/VENC = 4/\cos(\theta) \approx 2.82$. This type of mutual bias coherence is more likely to occur if the measurements are symmetric about the measured velocity vector (and in the more linear $V_{max}/VENC$ to resolution regime) since it increases the probability that the maximum bias between each measurement is approximately proportional (figure 6.12).

In the dual-VENC method, assuming the high-VENC set is high enough that the level of bias is negligible, the bias in the low-VENC set must exceed 180° before an unaliasing error occurs. As shown in figure 6.12 this varies based on the velocity to VENC ratio and resolution. This relatively high threshold makes the dual-VENC method robust to bias, because the solution (i.e. high-VENC profile) is effectively known.

Since the proposed unaliasing method relies on correlated information to resolve aliasing wraps (as opposed to the known solution provided by a high-VENC set) the trade-off is an increased sensitivity to measurement errors, such as bias or noise. The relatively increased SNR of the proposed method can be traded for a reduced error sensitivity by either reducing the gradient moment in all directions, as previously suggested, or by reducing the moment of individual measurement directions. Figure 6.18(b) shows how the solution space is reduced by making one moment lower (shown as an increased VENC) than the rest. One of the red and blue solutions has been removed because the only valid solutions now must lie on the solution line parallel to the y axis.

Measurements that are corrupted by bias continue to be inconsistent with phase estimates of neighboring pixels, even in the event that they are correctly unaliased. This is because the bias itself is not estimated and corrected. Bias

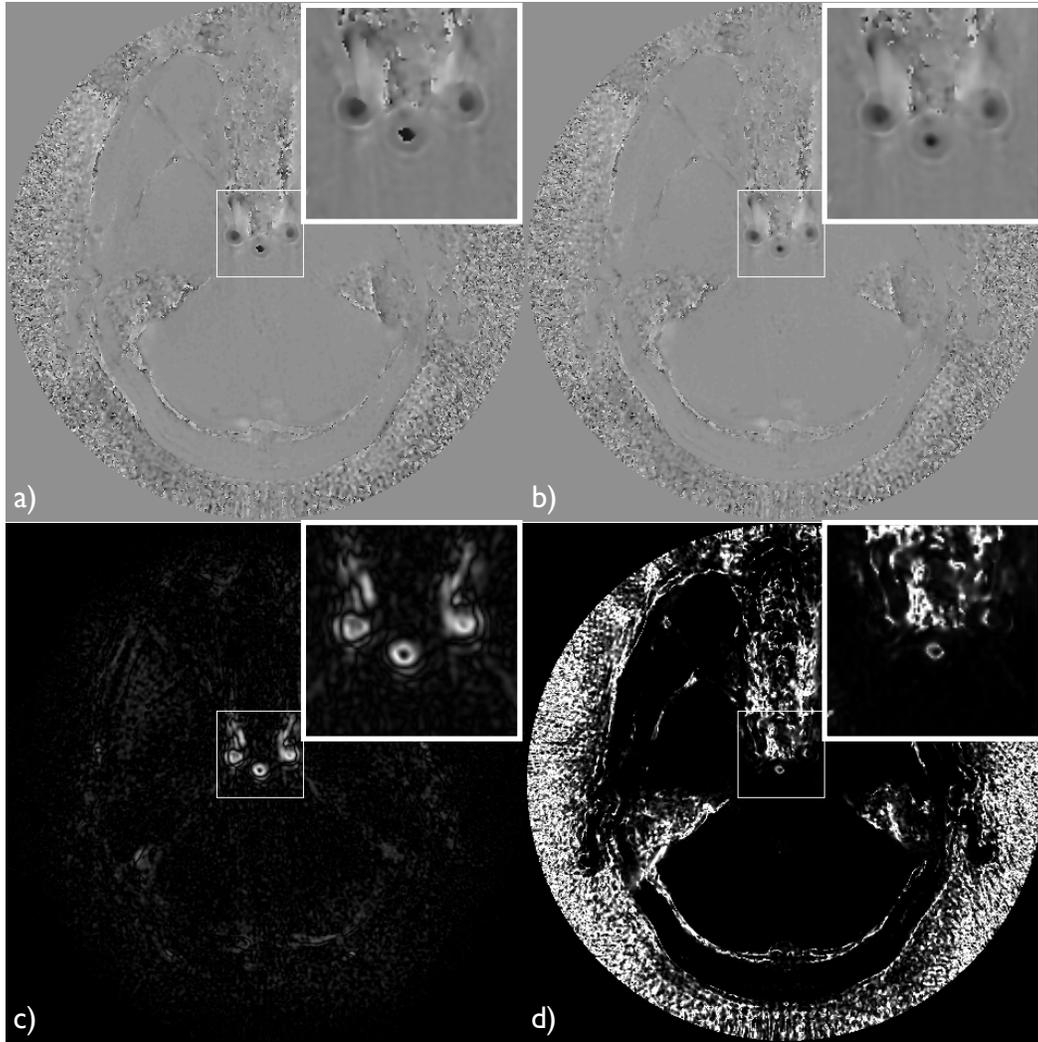


FIG. 6.19: Axial PCMRA reconstructions showing signal loss and consistency error specificity. (a) a reconstruction of the basilar artery that contains rough edges where the solution is suboptimal. (b) a thresholded reconstruction showing improvement at the basilar lumen edges. (c) a signal loss map. (d) a solution consistency error map.

estimation (as depicted in figures 6.8 and 6.11) requires knowledge of the neighboring voxel averages (to determine bias sign) and intra-voxel distribution (to calculate the bias magnitude). The latter is made more difficult by the spatially variable PSF that straddles the biased pixels.

The spatial location and relative level of bias, within the biased voxels, are highly resolved by the proposed unaliasing method. In low noise cases, the consistency error calculated in 6.11 is mostly affected by bias. While the affected measurements within u are not resolved by this information alone, the spatial position and relative bias information are contributed to by all of the measurements. Accurate relativity between the consistency error and level of bias depends on correct phase aliasing resolution.

The additions to the dual-VENC method proposed in (3) were created to address highly biased voxels in the low-VENC data by a weighted replacement with the high-VENC only in areas of high bias. In this algorithm the areas, level of bias, and subsequently the weights, are determined by the level of signal loss for each voxel. A similar method of replacement can be employed in the reconstruction of the proposed method. However, unlike the dual-VENC method, the proposed method does not contain a high-VENC set to be used as a replacement. The biased areas as previously discussed (figure 6.11) are a combination of profile squaring (caused by signal loss) and Gibbs ringing (caused by the squaring). This effectively means the PSF changes spatially across the profile narrowing in the areas of high signal loss.

A narrowing of the PSF is analogous to widening of the MTF. This is explained by the following logic. As previously discussed, the biased regions of the flow lumen lie at the edges of the lumen where the sharpest spatial transitions occur. Biasing makes the edge sharper by elevating the relative

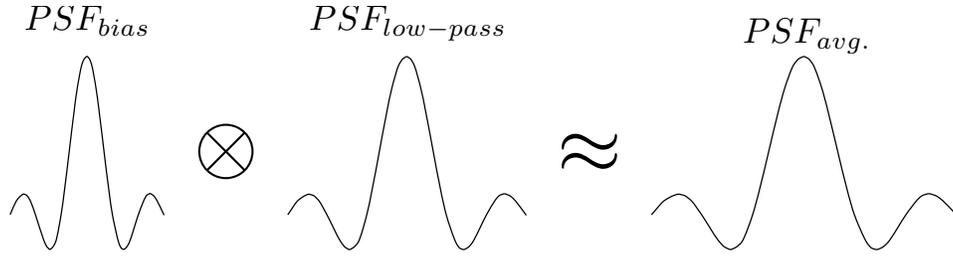


FIG. 6.20: An approximate relationship between the corrupted PSF in phase biased areas and the average image PSF.

contribution of elements (through the intrinsic signal weighted averaging) of low signal loss over the elements of high loss. The areas of lower signal loss by definition contain more slowly varying velocity distributions and therefore are mostly comprised of low spatial frequencies. Elevating their relative contribution (in the signal weighted average) effectively stretches or pushes the low spatial frequency signal out into higher spatial frequency areas, thus widening the MTF. The frequency domain support provided by the acquisition trajectory essentially filters out the high spatial frequencies required to fully restore the PSF in these areas. Additionally, these areas are presumed to be more effected by spatial aliasing due to undersampling since their spectral (k-space) profiles are wider.

In order to estimate the level of bias in these areas, the individual reconstructed velocity images (\hat{v}) are low pass filtered as real valued images. As shown in figure 6.20 the level of filtering must create a PSF ($PSF_{low-pass}$) of the correct width so that when convolved with the PSF of the biased region (PSF_{bias}) results in the average image PSF ($PSF_{avg.}$). This is an approximate correction to the PSF in highly biased areas, and blurs the PSF in the rest of the image.

As seen in figure 6.19(c) and (d), the consistency error has a higher speci-

ficity than the signal loss for the biased pixels shown as a sharp edge around the basilar artery in Fig. figCONSTcsseMaps(a). To increase the specificity of the relative signal loss, the consistency error is multiplied by each measurement of loss (S_c). The areas of bias in the unaliased velocity images \hat{v} are then corrected by weighted replacement of the blurred velocity images where the unaliased weighting factor (W_1) is,

$$W_1 = \begin{cases} 1 - \frac{1}{\beta}S_c, & 1 - \frac{1}{\beta}S_c > 0 \\ 0, & \textit{else} \end{cases}, \quad (6.22)$$

where β is a selectivity factor, and the blurred imaged weighting is,

$$W_2 = 1 - W_1. \quad (6.23)$$

The selectivity is controlled by varying β where the minimum value is equal to $Max(S_c)$. Each velocity image \hat{v} is then combined with its blurred counterpart \hat{v}_b to make a composite \hat{v}_c via,

$$\hat{v}_c = \hat{v} \cdot W_1 + \hat{v}_b \cdot W_2. \quad (6.24)$$

Since the signal loss is relative to the level of bias, the measurements with low loss receive little or not replacement while the measurements with high loss receive high or complete replacement.

6.2.4.1 Methods

Since the level of sensitivity to bias depends on numerous conditions including relative resolution and SNR, the robustness of the proposed method was tested on a neuroangiographic acquisition. This section covers the parameters used in the acquisition and reconstruction of the 6-direction method.

SPI was used as the underlying imaging technique. The base spiral planes consisted of 33 fully sampled achimedean interleaves. The number of planes rotated around the k_z axis was 125 for a matrix diameter of 240 points. This corresponds to a radial undersampling factor of 3 in the k_x, y plane. The sampling period (ADC) within each TR was $7.3ms$. The flip angle was 15° with $TR = 17ms$ and $TE = 2.8ms$ for a total scan time of 8 minutes and 30 seconds. A VENC of $40cm/s$ was used for the proposed method and a VENC of $80cm/s$ was used in a subsequent 3-direction referenced scan.

A 3D conjugate gradient SENSE reconstruction was used to reconstruct each measurement set and the reference set. The number of SENSE iterations used was 3 per set. The lowpass filter used to generate the blurred replacement set was applied in the spectral domain of the real valued velocity image. The filter was a spherically symmetric linear taper extending from a k-space radius of zero to k_{max} , effectively reducing the resolution by 2. The β parameter used in replacement selectivity was set to $Max(S_c)$.

Three different reconstruction were performed for a comparison of the constraints. The unaliasing algorithm was performed without thresholding, with a threshold of $|V_{max}| = 70cm/s$, and combined thresholding and signal based replacement. The noise outside of the head was masked by an ellipsoidal shape before the maximum intensity projection was made for each set in the axial, coronal, and sagittal directions.

6.2.4.2 Results & Discussion

This section examines the efficacy of the proposed reconstruction and constraints.

Figure 6.21 shows the maximum intensity projections (MIPs) for each of

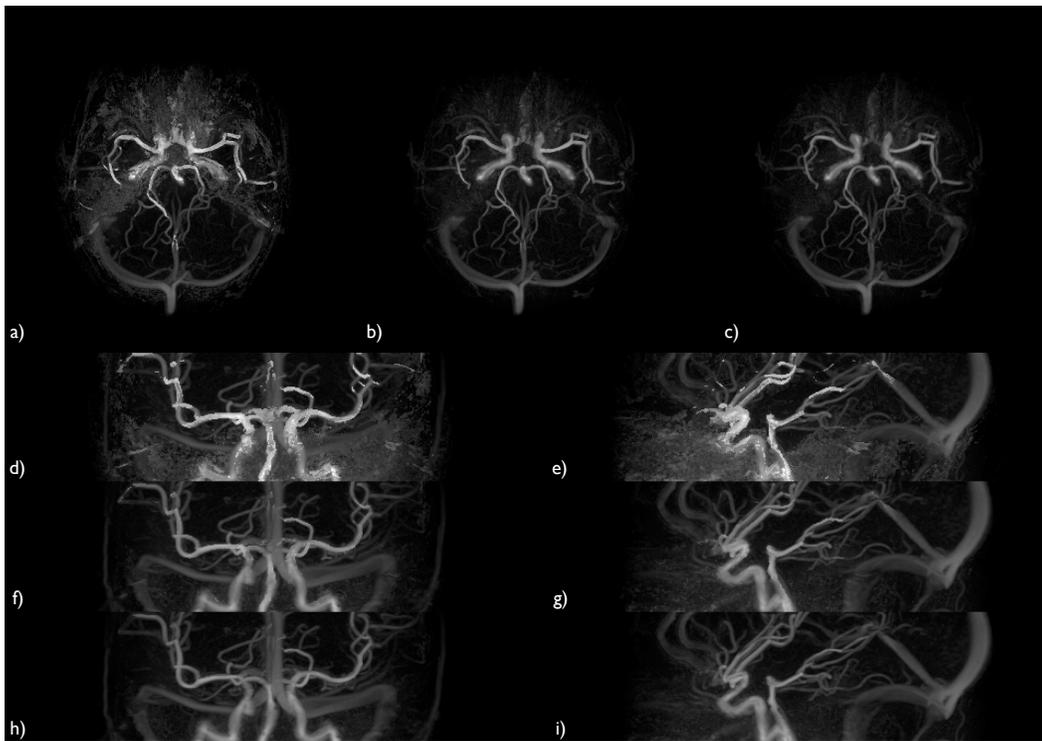


FIG. 6.21: Axial, coronal, and sagittal maximum intensity projections demonstrating each constraint in the proposed reconstruction process.

the reconstruction methods. The subfigures (a), (d) and (e) show the results of the unaliasing algorithm without the proposed constraints. These images show unwrapping errors and bias errors mostly located at the vessel walls similar to the error shown in figure 6.19(a), around the basilar artery. The sagittal MIPs show the carotid syphons and basilar are obfuscated by noise originating in the sinus and ear canal regions.

Comparing these figures to the subfigures (b), (f) and (g) shows that the thresholding effectively reduces the number of unwrap errors at the edges, clearing most of the discrete erroneously high velocity estimations. The noise is also not as over-fit, which lowers the relative intensity through the MIP. The remaining discontinuities at the vessel edges are due to bias and are more easily visible at the edges of the left and right cerebral arteries in the axial MIP.

In comparing both sets to the subfigures (c), (h) and (i) show that most of the remaining biased edges are returned to their true resolution. The noise is also further mitigated in the sinuses and ear canal because of the consistency error based replacement. As shown in figure 6.19(c) and (d) the signal loss is low in signal void areas such as the sinuses. However, the consistency error in these regions is high targeting the solutions for areas of blurred replacement. Since the real valued velocity map is blurred, the zero mean noise will be averaged. The low signal loss in these areas does not produce full replacement, but the effect is visible.

For a closer comparison of reconstruction, a 1D cross section of the basilar artery was plotted for the thresholded and composite reconstructions in figure 6.22. The biased measurements on the left side of the profile are most destructive, with peaks that are over and under estimated by approximately 20cm/s

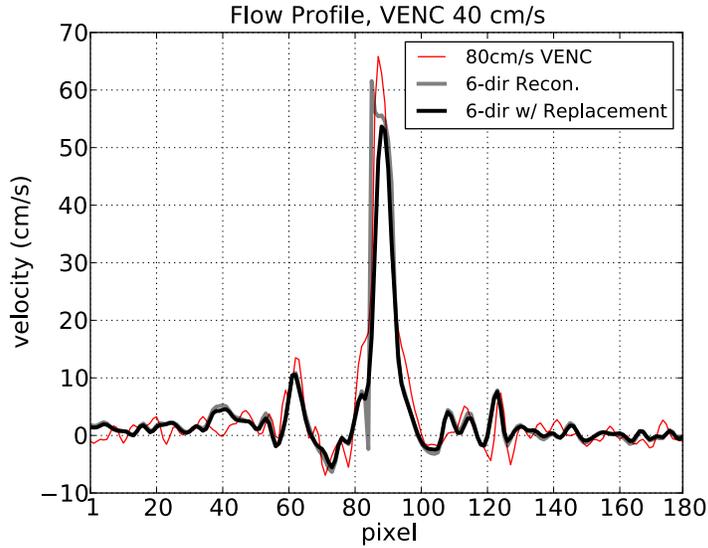


FIG. 6.22: Measured velocity profiles of the basilar artery for constrained reconstruction compared to an 80cm/s VENC referenced reconstruction.

as compared to the 80cm/s reference set. This level of bias corresponds to a 90° phase offset. This distribution roughly falls into the a $V_{max}/VENC$ ratio of 1.6, which, according to the simulations (figure 6.12), estimates the number of divisions per diameter is about 3-4. The mostly parabolic portion of the profile in figure 6.22 is roughly 10-12 points wide at a zero-padding factor of 2, making the resolution 5-6 divisions per diameter. Since the basilar is in close proximity the sinus cavity, the blurring due to field inhomogeneity is assumed to be relatively high. This blurring can be seen in the MIPs around the carotid and basilar up to the Circle of Willis. It is assumed that this type of blurring is having an effect on the profile in a way that is unaccounted for by the simulation.

As previously shown in simulation, the maximum noise level tolerated by the unaliasing algorithm is about 9° . In-vivo, for this set of parameters, the tolerable bias for phase unwrapping appears to be at least 90° as mentioned

before. While this is only half of the dual-VENC tolerance (i.e. 180°), it is significantly higher than the predicted noise threshold. This reinforces the idea that the inconsistency in bias across measurements is relatively high.

Between the threshold reconstructed and composite reconstructed profiles (figure 6.22) it can be seen that the replacement is highly specific. Most of the replacement works directly on the biased estimates while the velocity estimates in for the stationary tissue remain mostly untouched. This indicates that the resolution is in tact. Since the replacement sets are blurred, the replaced values will regain some of the signal lost to dephasing, making these estimates relatively lower in SNR, which is a similar problem in the dual-VENC, high-VENC replacement method.

6.2.5 Dual VENC Method Comparison

This section compares the proposed method and dual-VENC method each using signal loss replacement reconstruction. The two methods are compared using the same gradient moment since the moment level will induce the same amount of bias in each. Both are compared to a conventional 3 direction referenced set that uses half the gradient moment. Each configuration is also compared to a deblurred set which is used to emphasize the effects of signal bias and resolution.

6.2.5.1 Methods

The weighting functions for both replacement methods were determined empirically using the linear relationship shown in equation 6.22, where S_c represents the combined consistency error and signal loss in the proposed method, and represents only the signal loss in the dual-VENC method. The β parameter is empirically determined for each set by finding the lowest beta necessary to

force the heavily biased velocity estimates out of the reconstructed images.

Each set was collected with the SPI imaging technique using 33 spiral interleaves, 125 projections (for a radial undersampling factor of $R = 3$), 240 diameter matrix, 18ms TR, 10° flip angle, 3.8ms TE, and 7.3ms sampling window for a total scan time of 8min and 50sec. All scans were performed on a GE 3Tesla Signa Excite scanner using a standard 8-channel head coil.

Radial undersampling alone was used to compensate for the minimal variance in z-axis coil sensitivity. A 3D CG-SENSE parallel reconstruction technique was used to reconstruct each measurement volume before PCMRA specific reconstruction was applied.

The dual-VENC method was prescribed with a high and low VENC of 80cm/s and 40cm/s respectively. The proposed 6 direction method used a 40cm/s VENC. The 80cm/s VENC was reconstructed separately as a conventional 3 direction referenced set to be used for additional comparisons.

The previously described semi-balanced gradient moment configuration was used in the proposed method.

Two extra SPGR sequences were collected at the same time as the two velocity encoded sets. The extra sequences maintained the same trajectories, TR and flip angle as the velocity encoded sets. A TE of 2.8ms and 1.8ms were used in conjunction with the reference scan ($TE = 3.8ms$) of each velocity encoded set, to generate off-resonance (B_0) maps. The base images were first reconstructed using SENSE, then blurred using a low-pass k-space filter, and were then interpolated using signal weighted average of the change in phase between the reference and $TE = 2.8ms$ and the $TE = 2.8ms$ and the $TE = 1.8ms$ sets. Spiral deblurring was performed using the method proposed in (33).

6.2.5.2 Results & Discussion

Figure 6.23 shows the resulting maximum intensity projections for each method. The top row of axial images and two columns of coronal (left) and sagittal right (right) correspond, in order, with the conventional 80cm/s 3 direction referenced set, the dual-VENC set, and the proposed method set. The vessel conspicuity increase with each method, with the proposed method offering the highest. This is mostly due to the increased VNR created by moment balancing.

The signal based replacement appears to increase the noise, in the dual-VENC reconstruction, and blurs the proposed method reconstruction. Both of which are expected to non-uniformly lower the vessel conspicuity in each reconstruction. This effect is mitigated as the specificity of the signal based replacement is increased. The measure of consistency error produced in the proposed method reconstruction provides an advantage in the level of specificity as previously shown (figures 6.19 and 6.22). The dual-VENC signal based replacement has only the signal loss maps which are lower in spatial resolution. This specificity produces a limit to the amount of VNR gain created through gradient moment increases.

The signal bias produces an artificial increase the resolution that makes the true vessel diameter difficult to resolve. The conventional high-VENC and dual-VENC sets produce sharper edge transitions than the proposed method possibly indicating that the proposed method is losing resolution. According the simulation results shown in figure 6.12, an approximate vessel granularity of 6 divisions will produce around 25° of bias at a $V_{max}/VENC$ of 0.8. This is roughly 10cm/s for a flow of 80cm/s at an 80cm/s VENC, which potentially

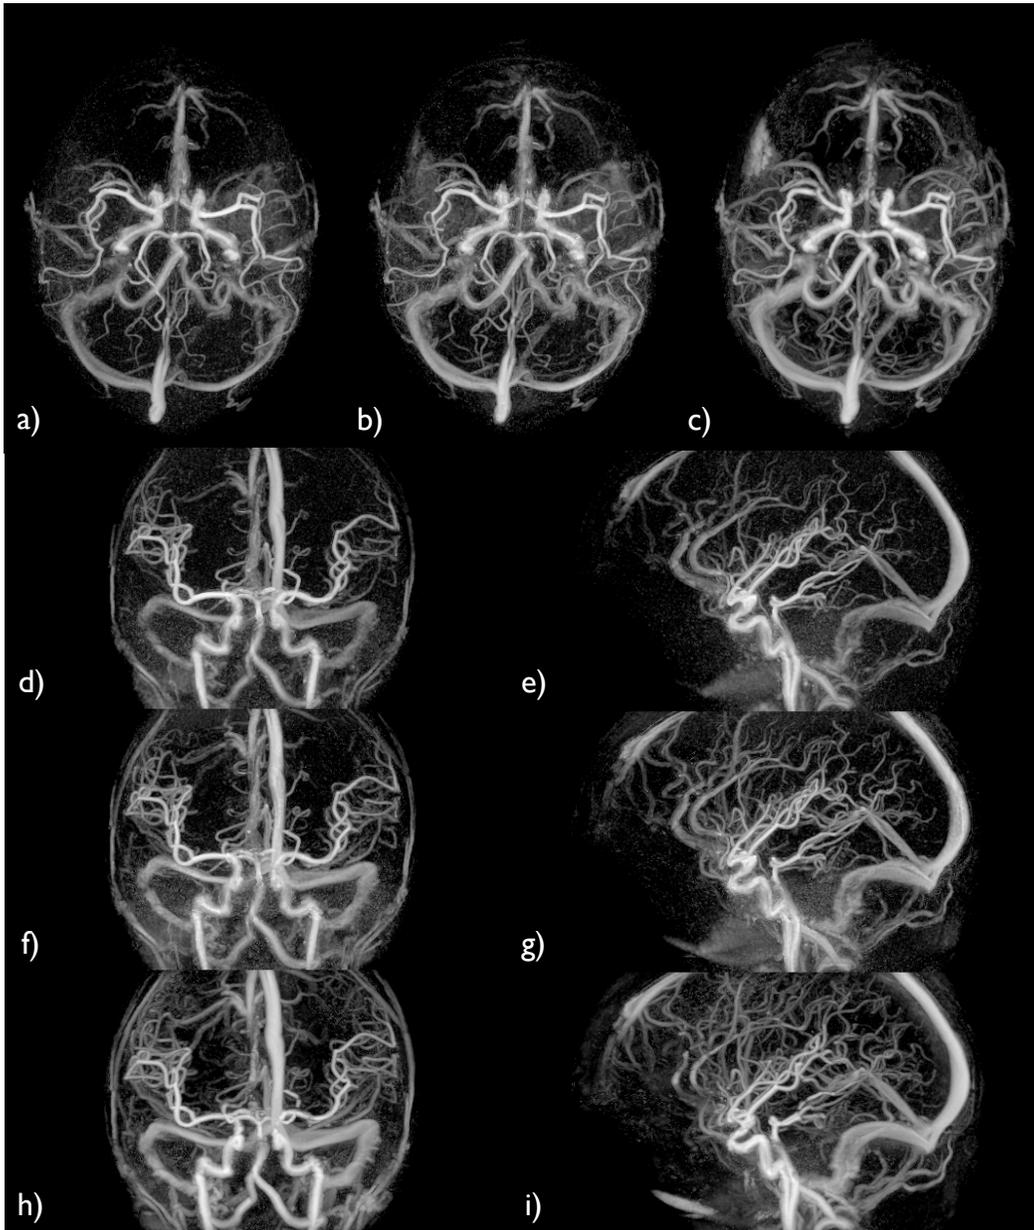


FIG. 6.23: Maximum intensity projections in the axial, coronal, and sagittal planes. MIPs (a), (d) and (e) correspond to the conventional 80cm/s set, (b), (f) and (g) to the dual-VENC set, and (c), (h) and (i) to the proposed method.

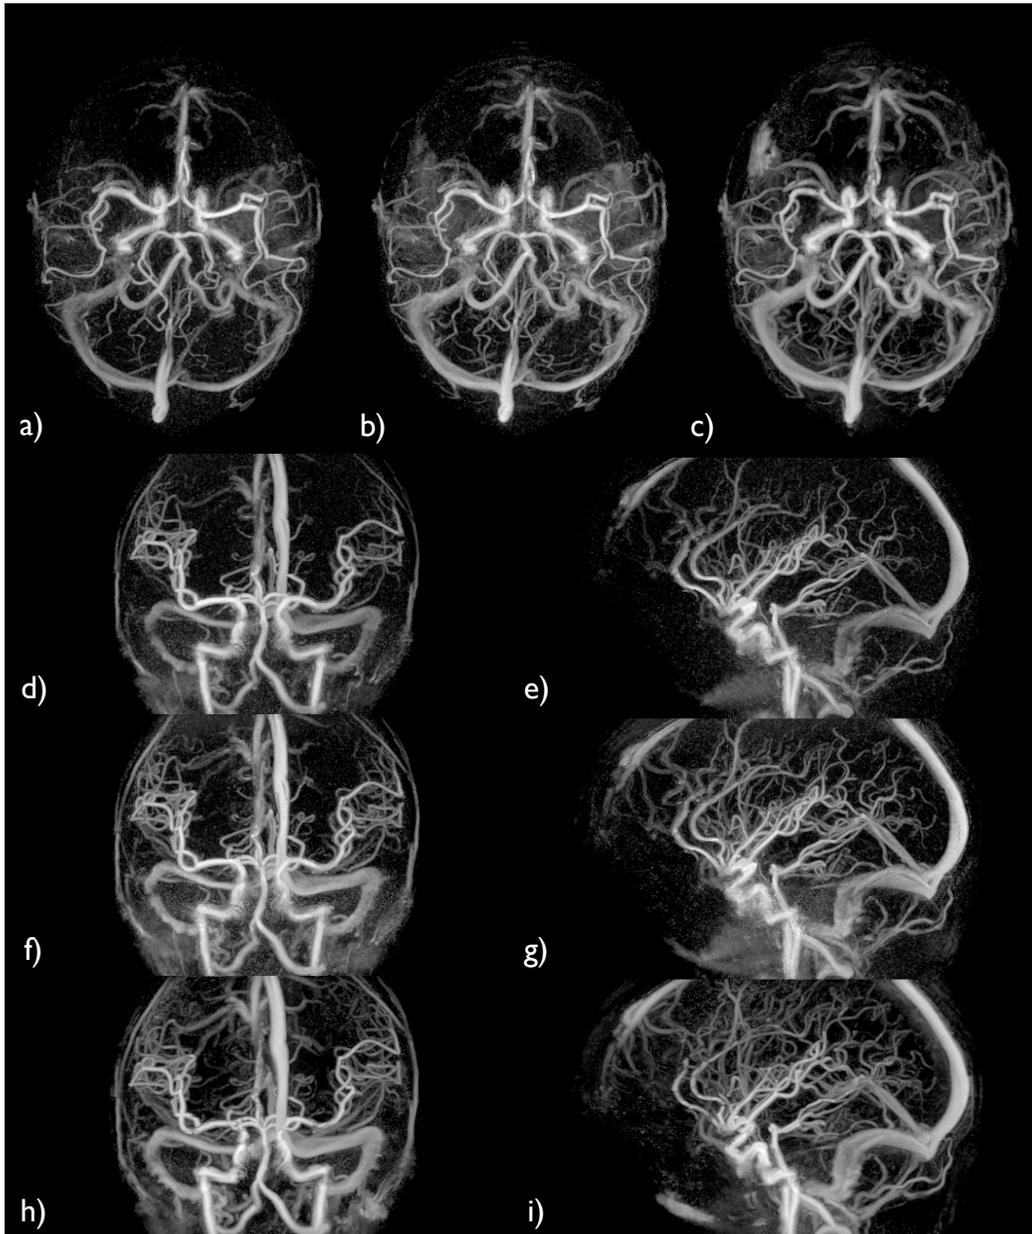


FIG. 6.24: Maximum intensity projections in the axial, coronal, and sagittal planes of deblurred data sets. MIPs (a), (d) and (e) correspond to the conventional 80cm/s set, (b), (f) and (g) to the dual-VENC set, and (c), (h) and (i) to the proposed method.

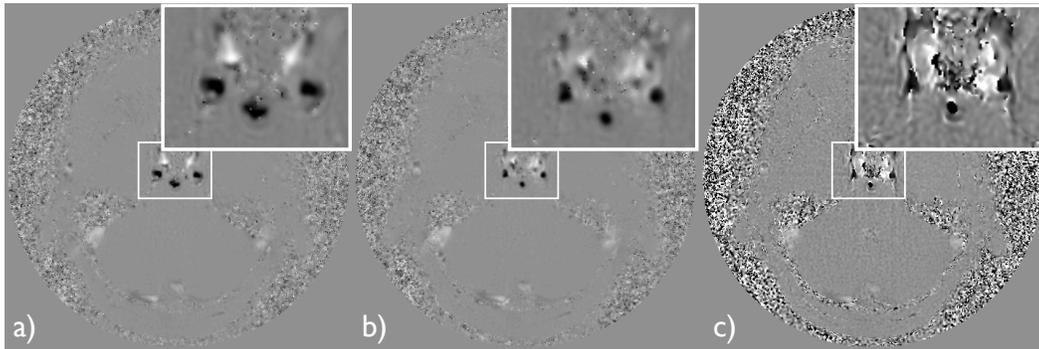


FIG. 6.25: Blurred and deblurred basilar artery cross sections. (a) a 6-direction reconstruction with blurred replacement. (b) a deblurred 6-direction reconstruction with blurred replacement. (c) a deblurred 80cm/s VENC referenced reconstruction.

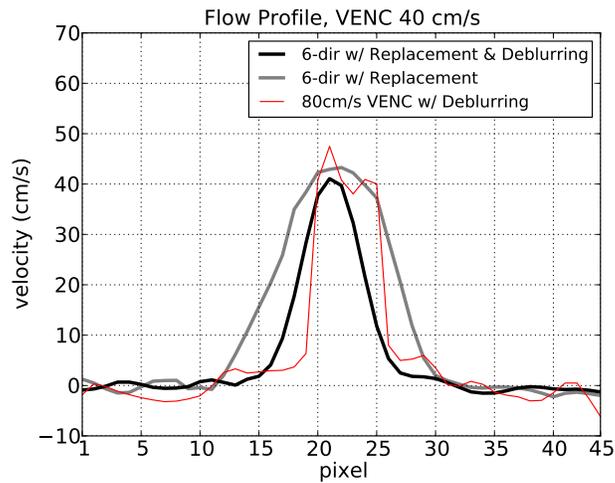


FIG. 6.26: A 1D cross section of the basilar artery flow lumen.

means that even the high-VENC set, which is used in replacement of low-VENC data, is also artificially narrowing the reconstructed vessel diameter.

According to figure 6.15(e), the RSS moment configuration may provide an additional moderate gain in vessel conspicuity over the used semi-balanced method. This comes at the cost of an even noise distribution, but may provide a 20% increase in VNR as shown in-vitro.

Blurring plays a significant role in the vessels closest to the sinuses (i.e. internal carotid arteries, basilar artery, and anterior cerebral arteries). The most noticeable area is where the anterior cerebral arteries meet the circle of Willis as shown in the sagittal MIPs of figure 6.23. This segment of the vessel appears to be almost missing in the figure 6.23(i).

The reconstruction of the deblurred MIPs are shown in figure 6.24. The diameter of the anterior cerebral arteries is reduced and the estimated speed is increased. The connection between the artery and the circle of Willis is also more clearly represented.

Cross sections of the basilar artery, similar to those shown in figures 6.19, 6.22 and 6.11, are shown in figures 6.25 and 6.26. As shown in figure 6.26, the profiles of the reconstructed 6-direction method are smooth due to signal weighted replacement. The deblurred profile has a diameter that is consistent with that of the 80cm/s VENC referenced reconstruction. The 80cm/s VENC reconstruction shows the characteristic squaring of the flow lumen resulting from a high $V_{max}/VENC$ to resolution ratio as shown in 6.11. Using the 6-direction method with deblurring reconstruction as a basis, the actual vessel diameter is roughly 15 points wide which corresponds to 7 voxels at the scanned resolution. Using the estimated bias in figure 6.12 the approximate bias at the lumen wall is 25° which corresponds to a deviation of roughly 10cm/s . This

indicates that a VENC of 80cm/s is insufficient for adequately reconstructing the flow profile at this position within the subject. The squared edges are also noticeable in the anterior cerebral arteries in figures 6.24(e) and (g) of the dual-VENC reconstructed sets. Since the high-VENC replacement set is also corrupted by inadequate resolution for this VENC, the signal based replacement maintains the rigid lumen edge in the final reconstruction.

6.3 Summary Discussion

This section summarizes the main points from each individual experiment presented in the previous section.

6.3.1 Noise Sensitivity

The proposed method provides a base improvement in SNR efficiency over the previously proposed dual-VENC (three-point) method proposed in (4), which is determined through measurement averaging. The measurement direction geometry is shown to be a significant factor in the efficacy of phase unaliasing. Using a charge repulsion technique to find the most even distribution of measurements is suboptimal for odd numbers of measurements, and optimal for even numbers. Noise sensitivity is shown to be only 3° more sensitive in phase noise compared to the dual-VENC method. Noise sensitivity is also shown to be inversely proportional to the VENC, allowing sensitivity to be traded for gradient moment.

6.3.2 Optimal Voxel Shape

The optimal voxel shape is shown to be the point spread function that produces the highest resolution supported by the acquired data. The work in this section demonstrates through 3D simulation that the application of a filter to the MTF in k-space will, in most situations, increase the level of signal bias. Using

a parabolic flow model, the level of bias is shown to increase rapidly at low resolutions (< 6 divisions over the lumen diameter) and with moderate velocity to VENC ratios (> 1.5).

6.3.3 SNR Optimized Moment Balancing

The relative orientation of the measurement geometry to the solution geometry can produce significant gains in SNR due to gradient moment balancing. The theoretical limit of SNR gain, over a typical referenced velocity measurement, for 6 directions is shown to be a factor of 2. Through in-vitro experiments using a flow phantom, proper gradient moment balancing is achieved providing the near maximum gain in SNR of 2. In comparison to the dual-VENC method (assuming $VENC_{low} \ll VENC_{high}$), which provides the same SNR as the typical referenced measurement in the absence of phase aliasing, the 6-direction method is potentially twice as SNR efficient.

6.3.4 Constraints & Reconstruction

The proposed method is shown to be robust in-vivo using the current state of the art in rapid 3D imaging sequences, with the proposed solution thresholding constraint and the signal based weighted replacement bias corrupted data. The method is shown to be robust to in-vivo biases of at least 90° . Signal based replacement is shown to have high spatial specificity, which maintains resolution and SNR.

6.3.5 Dual VENC Method Comparison

The proposed method is shown to provide higher vessel conspicuity and VNR for the same gradient moment as compared to the dual-VENC method. This is advantageous, in two significant ways: 1) the signal loss induced phase bias is minimized due to lower gradient moments and 2) the method is more time

efficient in terms of VNR to acquisition duration.

6.3.6 Conclusion

The proposed method is shown to produce an SNR improvement of 2 for the 6-direction configuration, over referenced methods such as the dual-VENC method, in-vitro. This makes the proposed method twice as SNR efficient in 3D as dual-VENC. The method is also robust to moment encoding related errors such as phase bias and phase aliasing, in the face of practical limitations such as spatial resolution, field inhomogeneity, and noise. The improvements in SNR are demonstratable in-vivo, resulting in higher vessel conspicuity and VNR. The added VNR allows the proposed method to be acquired using lower gradient moments (compared to dual-VENC) which reduces the number of reconstructed voxels affected by signal biased phase.

7 DISCUSSION

The purpose of this work was to develop a rapid 3D PCMRA acquisition and reconstruction technique. Conceptually, the requirement for achieving this goal is to improve vessel conspicuity through the increase the velocity to noise ratio. The means for achieving this goal were explored through two different avenues that constitute the relationship shown in equation 2.25. The first being image SNR, which is a function the imaging technique. The second, is the applied moment, which is a function of the gradient area.

Initially (chapter 3), the imaging aspects were explored by making use of the intrinsic SNR advantages provided by both 3D and spiral techniques through spiral projection imaging (11, 12). Existing velocity encoding techniques (3, 4) were explored within this framework resulting in a PCMRA technique that is robust to spatial aliasing at low reduction factors (via uniformly distributed undersampling, sample density correction and high-moment dual-VENC encoding), and signal biased phase (via composite high-VENC replacement). This motivated the exploration of parallel imaging reconstruction (14, 73) as a means of reducing aliasing at higher reduction factors as discussed in chapter 4. While reducing the acquisition time of an imaging technique is an imperative, reductions in the reconstruction time are also advantageous as they are more portable across platforms of varying cost and performance. The work presented in chapter 5 covered a method which provided a significant reduction in the computation of 3D sampling density (8, 34, 42). Finally, a new technique for resolving phase aliasing was explored as a method for allowing further increases in the applied gradient moment (chapter 6).

Spiral projection imaging has been shown to be a fitting imaging trajectory for this work. Its highly configurable nature provides a multifaceted platform

for evaluating new parallel imaging techniques (for various coil geometries) and undersampling patterns that produce incoherent spatial aliasing making it robust without parallel imaging at modest reduction factors. The high SNR efficiency also provides an advantage to phase contrast techniques reducing the level of required gradient moment.

The PEPI technique has been demonstrated to provide modest reduction factors at reduced reconstruction durations as compared to SENSE, for 3D imaging techniques. Currently, the PEPI technique falls short of SENSE in maximum achievable reduction factor. SENSE has therefore been used as the platform from which phase contrast methods have been continued to be explored.

The proposed sample density estimation method is shown to provide a significant enhancement to the reconstruction performance of both RSS coil combined and parallel imaging reconstruction methods. The time savings make the proposed method viable for practical 3D reconstruction which, in turn, allows the convergence of more accurate solutions. This method has also demonstrated stability over multiple image and trajectory combinations. This stability was shown to be consistent where analytical solution fail.

The most novel contribution of this work is in the proposed method for solving phase aliasing in low moment phase contrast acquisitions. This method extends the SNR efficiency beyond previously proposed methods by fully utilizing the SNR potential all of the measurements made (compared using the time footprint of the dual-VENC method). This is accomplished by employing high gradient moments in every velocity sensitive acquisition and gradient moment balancing. Phase aliasing is solved by minimizing the error of a non-linear system, which is conditioned through measurement geometry and a mixed integer

solution space. Symmetric measurement geometries are shown to be the most robust due to the fact that they provide the most independent information.

8 FUTURE WORK

Since the inception of this project, new imaging methods have emerged which offer advantageous properties where spiral projection imaging is limited. Specifically, in the ability to prescribe anisotropic fields of view and efficiency as measured by k-space undersampling uniformity. These methods include FLORET (45) and stack of spirals. While stack of spiral is not new per se, it provides the ability to minimize the field of view orthogonal to the spiral axis allowing for greatly reduced scan times. A basic implementation of this method was employed in the in-vitro experiments used in assessing moment balancing efficacy in chapter 6. The FLORET method was used in the characterization of the PEPI parallel imaging method in chapter 4. Adding the proposed phase contrast technique, which is compatible with both of these base imaging techniques, would further improve state of 3D rapid PCMRA, through the aforementioned advantages.

For the proposed multi-point phase contrast techniques as well as existing phase contrast techniques, the amount of data acquired is relatively large. This presents a challenge for 3D reconstruction techniques in general, and is particularly so for parallel imaging reconstruction because of the associated long computation times. PEPI provides a simple, computationally less expensive reconstruction that currently needs fine tuning in the construction of coil sensitivity maps. These needs may be addressed by looking to methods that intrinsically solve for coil sensitivity maps such as GRAPPA (44) or ESpirit(74, 75). Additionally, stopping criteria would make the reconstruction more robust by allowing a variable number of iterations to be automatically determined, taking full advantage of the image SNR (76).

The proposed phase unaliasing method is a category of non-linear systems

known as mixed-integer non-linear programming (MINLP). Minimization techniques are available for these applications such as BONMIN (77, 78). The implementation of this framework may provide a conduit for additional problem constraints or objective functions that can be simultaneously solved with the unaliasing problem in a practical time window.

9 SUMMARY CONCLUSIONS

The methods proposed in this work have demonstrated the viability of rapid 3D phase contrast magnetic resonance angiography in practical use. This has been achieved through the combination of efficient 3D imaging techniques, parallel imaging reconstruction, and novel phase contrast measurement and reconstruction methods. These areas have been addressed specifically with new imaging and reconstruction techniques called spiral projection imaging (SPI), sample density correction (SDC), pretty easy parallel imaging (PEPI) and multi-directional high moment encoding.

The most significant contributions are in the investigation of new high-moment phase contrast method that enables the use of larger gradient areas by resolving phase aliasing errors and reduces the need for large gradient areas through moment balancing. The method has been shown in-vitro to provide up to twice the SNR efficiency as compared to a well known dual-VENC technique (3, 4). In-vivo the method produces higher vessel conspicuity at an equivalent moment to the dual-VENC technique allowing lower moments to be used to avoid artifacts caused by phase bias errors.

A sample density correction technique which was comprised of two previously proposed methods has been shown to provide significant reconstruction time reductions for 3D non-Cartesian trajectories. The time reductions improve reconstruction accuracy by allowing full solution convergence within practical time windows. This work has also resulted in the public release of code for edification and utilization in the scientific community (chapter 5).

The pretty easy parallel imaging (PEPI) technique has been pursued, characterized and evaluated with respect to the well known SENSE parallel imaging technique. The successful and ineffective areas have been discovered and doc-

umented and an insight necessary to advance the method has been proposed.

REFERENCES

- [1] Zwart NR, Pipe JG. Dual VENC phase contrast MRA using spiral projection imaging. In ISMRM Workshop on Cardiovascular Flow, Function & Tissue Mechanics. Sintra, Portugal, 2009; .
- [2] Zwart NR, Pipe JG. Separable gridding reconstruction for spiral trajectories. In ISMRM Workshop on Data Sampling and Image Reconstruction. Sedona, Arizona, 2009; .
- [3] Pipe JG. A simple measure of flow disorder and wall shear stress in phase contrast MRI. *Magn Reson Med* 2003;49(3):543–50.
- [4] Lee AT, Pike GB, Pelc NJ. Three-point phase-contrast velocity measurements with increased velocity-to-noise ratio. *Magn Reson Med* 1995; 33(1):122–6.
- [5] Zwart NR, Pipe JG. Rapid 3D parallel imaging of non-Cartesian data. In ISMRM 18th Scientific Meeting & Exhibition. Stockholm, Sweden, 2010; .
- [6] Zwart NR, Pipe JG. 3D dual VENC PCMRA using spiral projection imaging. In ISMRM 18th Scientific Meeting & Exhibition. Stockholm, Sweden, 2010; .
- [7] Wajer F, KP P. Major speedup of reconstruction for sensitivity encoding with arbitrary trajectories. In Proc. Intl. Soc. Mag. Reson. Med. 9th Annual Meeting. 2001; .
- [8] Zwart NR, Pipe JG. Efficient sample density estimation by combining gridding and an optimized kernel. *Magn Reson Med* 2011;Online.

- [9] Johnson KO, Pipe JG. Convolution kernel design and efficient algorithm for sampling density correction. *Magn Reson Med* 2009;61(2):439–447.
- [10] Pipe JG. Basic spin physics. *Magn Reson Imaging Clin N Am* 1999; 7(4):607–627.
- [11] Johnson KO, Robison RK, Pipe JG. Rigid body motion compensation for spiral projection imaging. *IEEE Transactions on Medical Imaging* 2011; 30(3):655–665.
- [12] Irarrazabal P, Nishimura DG. Fast three dimensional magnetic resonance imaging. *Magn Reson Med* 1995;33(5):656–662.
- [13] Pipe JG. A fast 3D trajectory with orthogonal oversampling. In *ISMRM 18th Annual Meeting & Exhibition*. Stockholm, Sweden, 2010; .
- [14] Pruessmann KP, Weiger M, Bornert P, Boesiger P. Advances in sensitivity encoding with arbitrary k-space trajectories. *Magn Reson Med* 2001; 46(4):638–651.
- [15] Liang ZP. *Principles of Magnetic Resonance Imaging*, volume 4. IEEE Press, 2000.
- [16] Bernstein MA, King KF, Zhou XJ. *Handbook of MRI Pulse Sequences*, volume 32. Elsevier Academic Press, 2004.
- [17] Macovski A. Noise in MRI. *Magnetic Resonance in Medicine* 1996; 36(3):494–497.
- [18] Levitt MH. *Spin Dynamics: Basics of Nuclear Magnetic Resonance*, 2nd Edition. John Wiley & Sons, 2007.

- [19] Stanisz GJ, Odobina EE, Pun J, Escaravage M, Graham SJ, Bronskill MJ, Henkelman RM. T1, T2 relaxation and magnetization transfer in tissue at 3T. Proceedings 14th Scientific Meeting International Society for Magnetic Resonance in Medicine 2005;512(April):507–512.
- [20] Bernstein MA, Grgic M, Brosnan TJ, Pelc NJ. Reconstructions of phase contrast, phased array multicoil data. Magnetic Resonance in Medicine 1994;32(3):330–334.
- [21] Forbes KP, Pipe JG, Karis JP, Heiserman JE. Improved image quality and detection of acute cerebral infarction with PROPELLER diffusion-weighted MR imaging. Radiology 2002;225(2):551–555.
- [22] Zwart NR, Pipe JG. Motion corrected intracranial MRA using PROPELLER with RF quadratic encoding. Magn Reson Med 2009; 61(6):1405–1414.
- [23] Pipe JG. Non-cartesian parallel imaging, the easy way. In ISMRM Workshop on Data Sampling and Image Reconstruction. Sedona, Arizona, 2009; .
- [24] Lustig M, Donoho D, Pauly JM. Sparse MRI: The application of compressed sensing for rapid MR imaging. Magn Reson Med 2007;58(6):1182–1195.
- [25] Pipe JG. Asymmetric sampling along k-(slice-select) in two-dimensional multislice MRI. Magn Reson Med 1998;39(4):625–634.
- [26] Pipe JG. Analysis of localized quadratic encoding and reconstruction. Magn Reson Med 1996;36(1):137–146.

- [27] Noll DC, Boada FE, Eddy WF. A spectral approach to analyzing slice selection in planar imaging: optimization for through-plane interpolation. *Magn Reson Med* 1997;38(1):151–160.
- [28] Pruessmann KP, Weiger M, Scheidegger MB, Boesiger P. Sense: sensitivity encoding for fast mri. *Magn Reson Med* 1999;42(5):952–962.
- [29] Zwart NR, Pipe JG. 3D dual VENC PCMRA using spiral projection imaging. In *MR Angiography Club 21st International Conference on MRA*. East Lansing, MI, 2009; .
- [30] Pipe JG, Ahunbay E, Menon P. Effects of interleaf order for spiral MRI of dynamic processes. *Magn Reson Med* 1999;41(2):417–22.
- [31] Pipe JG. An optimized center-out k-space trajectory for multishot MRI: comparison with spiral and projection reconstruction. *Magnetic Resonance in Medicine* 1999;42(4):714–20.
- [32] Pipe JG. Reconstructing MR images from undersampled data: data-weighting considerations. *Magn Reson Med* 2000;43(6):867–875.
- [33] Ahunbay E, Pipe JG. Rapid method for deblurring spiral mr images. *Magn Reson Med* 2000;44(3):491–4.
- [34] Johnson KM, Markl M. Improved SNR in phase contrast velocimetry with five-point balanced flow encoding. *Magn Reson Med* 2010;63(2):349–355.
- [35] Conturo TE, Smith GD. Signal-to-noise in phase angle reconstruction: dynamic range extension using phase reference offsets. *Magn Reson Med* 1990;15(3):420–437.

- [36] Lamothe MJ, Rutt BK. Multistep phase difference phase contrast imaging. *J Magn Reson Imaging* 1997;7(5):838–842.
- [37] Salfity MF, Huntley JM, Graves MJ, Marklund O, Cusack R, Beauregard DA. Extending the dynamic range of phase contrast magnetic resonance velocity imaging using advanced higher-dimensional phase unwrapping algorithms. *Journal of the Royal Society Interface the Royal Society* 2006; 3(8):415–427.
- [38] Spagnolini U. 2D phase unwrapping and instantaneous frequency estimation. *IEEE Trans Geosci Remote Sens* 1995;33:579–589.
- [39] Larsson EG, Erdogmus D, Yan R, Principe JC, Fitzsimmons JR. Snr-optimality of sum-of-squares reconstruction for phased-array magnetic resonance imaging. *Journal of Magnetic Resonance* 2003;163(1):121–123.
- [40] Roemer PB, Edelstein WA, Hayes CE, Souza SP, Mueller OM. The NMR phased array. *Magnetic Resonance in Medicine* 1990;16(2):192–225.
- [41] Zwart NR, Pipe JG. Comparing gridding and masking in 3D parallel reconstruction. In *ISMRM 19th Scientific Meeting & Exhibition*. Montreal, Ontario, Canada, 2011; .
- [42] Pipe JG. Flow effects in localized quadratic, partial fourier MRA. *Magn Reson Med* 1999;41(2):309–314.
- [43] Ohliger MA, Sodickson DK. An introduction to coil array design for parallel MRI. *NMR in Biomedicine* 2006;19(3):300–315.
- [44] Griswold MA, Jakob PM, Heidemann RM, Nittka M, Jellus V, Wang J,

- Kiefer B, Haase A. Generalized autocalibrating partially parallel acquisitions (grappa). *Magnetic Resonance in Medicine* 2002;47(6):1202–1210.
- [45] Pipe JG, Aboussouan EA, Devaraj A, Johnson KO, Robison RK, Zwart NR. A new design and rationale for 3D orthogonally oversampled k-space trajectories. *Magn Reson Med* 2011;Online.
- [46] Noll DC. Multishot rosette trajectories for spectrally selective MR imaging. *IEEE Transactions on Medical Imaging* 1997;16(4):372–377.
- [47] Moriguchi H, Wendt M, Duerk JL. Applying the uniform resampling (URS) algorithm to a lissajous trajectory: fast image reconstruction with optimal gridding. *Magnetic Resonance in Medicine* 2000;44(5):766–781.
- [48] Pipe JG. Motion correction with PROPELLER MRI: application to head motion and free-breathing cardiac imaging. *Magn Reson Med* 1999; 42(5):963–969.
- [49] Robison RK, Devaraj A, Pipe JG. Fast, simple gradient delay estimation for spiral MRI. *Magn Reson Med* 2010;63(6):1683–1690.
- [50] Barger AV, Block WF, Toropov Y, Grist TM, Mistretta CA. Time-resolved contrast-enhanced imaging with isotropic resolution and broad coverage using an undersampled 3D projection trajectory. *Magn Reson Med* 2002;48(2):297–305.
- [51] Hoge RD, Kwan RK, Pike GB. Density compensation functions for spiral MRI. *Magn Reson Med* 1997;38(1):117–128.
- [52] Rasche V, Proksa R, Sinkus R, Bornert P, Eggers H. Resampling of data

- between arbitrary grids using convolution interpolation. *IEEE Trans Med Imaging* 1999;18(5):385–392.
- [53] Sedarat H, Nishimura DG. On the optimality of the gridding reconstruction algorithm. *IEEE Trans Med Imaging* 2000;19(4):306–317.
- [54] Bydder M, Samsonov AA, Du J. Evaluation of optimal density weighting for regridding. *Magn Reson Imaging* 2007;25(5):695–702.
- [55] Jackson JJ, Meyer CH, Nishimura DG, Macovski A. Selection of a convolution function for fourier inversion using gridding [computerised tomography application]. *IEEE Trans Med Imaging* 1991;10(3):473–478.
- [56] Mitsouras D, Mulkern RV, Rybicki FJ. Fast, exact k-space sample density compensation for trajectories composed of rotationally symmetric segments, and the SNR-optimized image reconstruction from non-Cartesian samples. *Magn Reson Med* 2008;60(2):339–349.
- [57] Meyer CH, Hu BS, Nishimura DG, Macovski A. Fast spiral coronary artery imaging. *Magn Reson Med* 1992;28(2):202–213.
- [58] Pipe JG, Menon P. Sampling density compensation in MRI: rationale and an iterative numerical solution. *Magn Reson Med* 1999;41(1):179–86.
- [59] Zwart NR, Pipe JG. Rapid sample density estimation for 3D trajectories. In *ISMRM 19th Scientific Meeting & Exhibition*. Montreal, Ontario, Canada, 2011; .
- [60] Gurney PT, Hargreaves BA, Nishimura DG. Design and analysis of a practical 3D cones trajectory. *Magn Reson Med* 2006;55(3):575–582.

- [61] Beatty PJ, Nishimura DG, Pauly JM. Rapid gridding reconstruction with a minimal oversampling ratio. *IEEE Trans Med Imaging* 2005;24(6):799–808.
- [62] Gregerson A. Implementing fast MRI gridding on GPUs via CUDA. Technical report, Nvidia Technical Reports on Medical Imaging using CUDA, Available at: http://www.nvidia.com/object/medical_imaging.html, 2008.
- [63] Sundareswaran KS, Frakes DH, Fogel MA, Soerensen DD, Oshinski JN, Yoganathan AP. Optimum fuzzy filters for phase-contrast magnetic resonance imaging segmentation. *J Magn Reson Imaging* 2009;29(1):155–165.
- [64] Frakes DH, Smith MJT, Parks J, Sharma S, Fogel SM, Yoganathan AP. New techniques for the reconstruction of complex vascular anatomies from mri images. *J Cardiovasc Magn Reson* 2005;7(2):425–432.
- [65] Frakes DH, Conrad CP, Healy TM, Monaco JW, Fogel M, Sharma S, Smith MJT, Yoganathan AP. Application of an adaptive control grid interpolation technique to morphological vascular reconstruction. *IEEE Trans Biomed Eng* 2003;50(2):197–206.
- [66] Hasan KM, Parker DL, Alexander AL. Comparison of gradient encoding schemes for diffusion-tensor MRI. *J Magn Reson Imaging* 2001;13(5):769–80.
- [67] Zwart NR, Pipe JG. Multi-directional high moment encoding in phase contrast MRI. In *ISMRM 19th Scientific Meeting & Exhibition*. Montreal, Ontario, Canada, 2011; .

- [68] Shaaban AM, Duerinckx AJ. Wall shear stress and early atherosclerosis: a review. *AJR Am J Roentgenol* 2000;174(6):1657–65.
- [69] Lagerstrand KM, Lehmann H, Starck G, Vikhoff-Baaz B, Ekholm S, Forssell-Aronsson E. Method to correct for the effects of limited spatial resolution in phase-contrast flow MRI measurements. *Magnetic Resonance in Medicine* 2002;48(5):883–889.
- [70] Conturo TE, Robinson BH. Analysis of encoding efficiency in mr imaging of velocity magnitude and direction. *Magnetic Resonance in Medicine* 1992;25(2):233–247.
- [71] Kaufman L, Kramer DM, Crooks LE, Ortendahl DA. Measuring signal-to-noise ratios in MR imaging. *Radiology* 1989;173(1):265–267.
- [72] Glover GH. Simple analytic spiral k-space algorithm. *Magnetic Resonance in Medicine* 1999;42(2):412–415.
- [73] Pipe JG, Zwart N. Turboprop: improved propeller imaging. *Magnetic Resonance in Medicine* 2006;55(2):380–385.
- [74] Lustig M. An Eigen-Vector Approach to AutoCalibrating Parallel MRI, where sense meets GRAPPA. In *ISMRM 19th Scientific Meeting & Exhibition*. Montreal, Ontario, Canada, 2011; .
- [75] Lustig M. Post-cartesian calibrationless parallel imaging reconstruction by structured low-rank matrix completion. In *ISMRM 19th Scientific Meeting & Exhibition*. Montreal, Ontario, Canada, 2011; .
- [76] Qu P, Zhong K, Zhang B, Wang J, Shen GX. Convergence behavior of

iterative SENSE reconstruction with non-cartesian trajectories. *Magnetic Resonance in Medicine* 2005;54(4):1040–1045.

[77] Czyzyk J, Mesnier MP, More JJ. The neos server. *Agenda* 1998;5(3):68–75.

[78] Bonami P, Lee J. Bonmin users manual. *Source* 2009;(November):1–38.

APPENDIX A
SPATIAL OFF-RESONANCE

Spatial off resonance occurs at the air-water interfaces, where the spatially varying magnetic flux rapidly changes density to accommodate materials of differing magnetic permeability. To measure the level of off resonance, two SPGR scans are taken at different echo times and phase subtracted from one another. The change in phase is related to the frequency of spin precession. The pixel by pixel phase difference provides a map of the spatial change in the frequency of precession (B_0 inhomogeneity). The different echo times provide more or less phase accrual in these areas within the low spatial frequencies providing a low resolution map of the changes in magnetic field amplitude.

A.1 Off-Resonance Vs. Head Position

A general level of dephasing was assessed through the creation of a B_0 inhomogeneity map. This map provides a spatial distribution of the level of off-resonance. The areas and level of off-resonance can then be interpreted as corresponding areas of blurring in a spiral based imaging technique and a relative level of impact.

The inhomogeneity maps were generated for different head positions using two spoiled gradient echo sequence with the common parameters: 64^3 matrix size, $4mm^3$ voxel size, $25.6cm$ FOV, fat saturation pulse, TR of $13.2ms$, $128kHz$ receive bandwidth, and flip angle of 10° , $TE_1 = 0.9ms$, and $TE_2 = 0.712ms$. The two acquisitions used for each head position differed in TE by $0.188ms$ allowing for just over $5kHz$ off-resonance to be resolved.

Figure A.1 shows the level of off-resonance at various head positions. These maps indicate that the highest levels of blurring in spiral reconstruction will be concentrated in the sinuses and potentially near the left and right side of the head depending on the relative orientation.

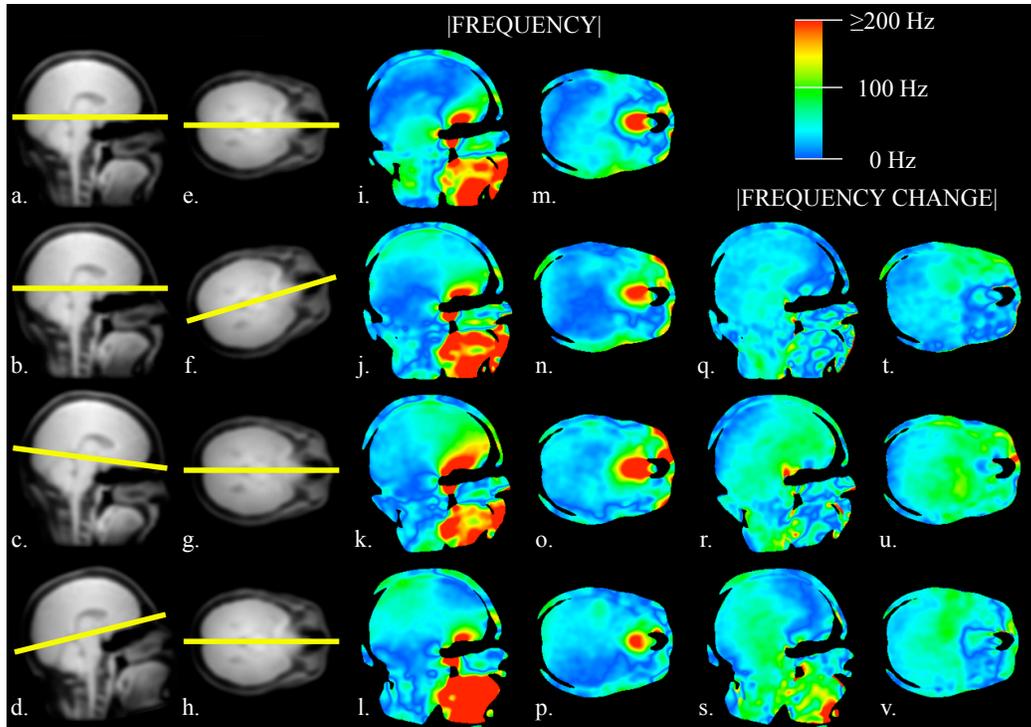


FIG. A.1: Spatial off resonance at air tissue interfaces for various head positions. Row (a), centered head position. Row (b), axial rotation in (X,Y) of 16.6° , Shift of 44mm (X) and -37mm (Y). Row (d), sagittal rotation in (Y,Z) of 14.0° , Shift of -36mm (Y) and 37mm (Z). Row (c), sagittal rotation in (Y,Z) of -7.5° , Shift of 17mm (Y) and -7mm (Z).

APPENDIX B
DYNAMIC 2D PEPI RECONSTRUCTION

The PEPI parallel imaging technique was used in the reconstruction of a dynamic, synthesized x-ray angiographic time series, as an entry in the ISMRM 2010 Data Reconstruction Challenge held in Stockholm Sweden during the 18th Annual Scientific Meeting and Exhibition. The entry was awarded first place in the ‘Need For Speed’ contest. This section briefly describes the methods used in the reconstructed entry.

B.1 ISMRM 2010 Data Reconstruction Challenge

The Need For Speed contest was focused on the reconstruction of a dynamic 2D, time resolved, neuro-angiograph of an arteriovenous malformation. The supplied simulated MR data was based off of an x-ray angiograph of 31 time points and 512x512 image resolution. Synthetic 8-channel coil information was supplied for parallel imaging reconstruction. The resulting data were significantly undersampled with only 200 trajectories and 2000 points each, covering k-space.

A 2D variable density spiral trajectory was chosen as the base imaging method. The spirals were designed with linearly varying undersampling starting with critical sampling at the center of k-space and increasing towards the edge of sampled k-space. Each spiral was rotated by the golden angle. After 13 TRs the approximate undersampling factor at the edges of k-space was 15.

A 2D PEPI (14, 23, 34) and RMS coil reconstruction were produced to show the aliasing reductions provided by PEPI. Figure B.1(b) and (c) show the swirling undersampled areas of k-space result in swirling aliasing artifact in the coil combined image. PEPI was shown to significantly reduce these swirling patterns in each of the reconstructed time points within a moderate reconstruction time of approximately 16 minutes.

Contest materials and data are available online at: http://www.ismrm.org/mri_unbound/simulated.htm

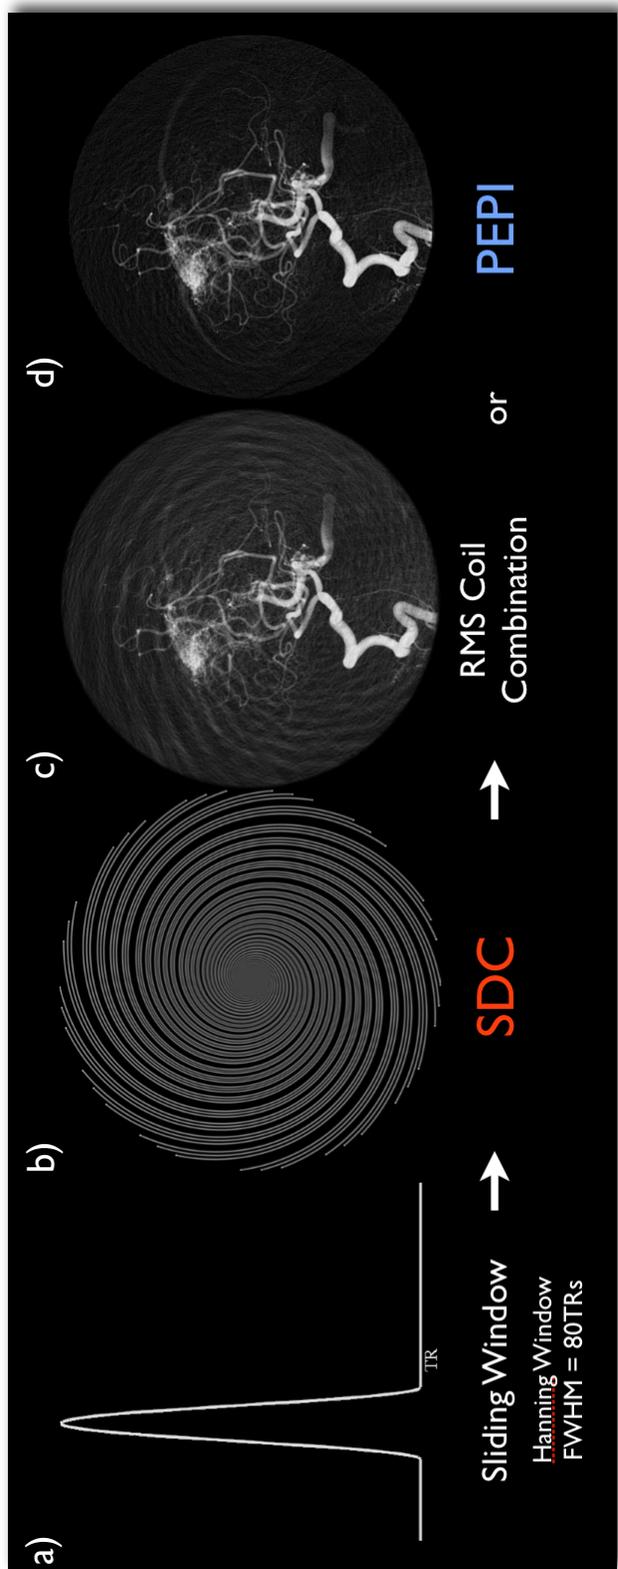


FIG. B.1: 2D dynamic PEPI reconstruction. (a) a sliding (hanning) window with an effective width of 80TRs. (b) relative k-space coverage of any given time point. (c) basic RMS coil combination (non-parallel) reconstruction. (d) 2D PEPI reconstruction.