

ADVANCES IN NON-CARTESIAN PARALLEL
MAGNETIC RESONANCE IMAGING USING THE
GRAPPA OPERATOR

DISSERTATION ZUR ERLANGUNG DES
NATURWISSENSCHAFTLICHEN DOKTORGRADES
DER JULIUS-MAXIMILIANS-UNIVERSITÄT WÜRZBURG

vorgelegt von

Nicole Seiberlich

aus Milwaukee, WI, USA

Würzburg, Februar 2008

Eingereicht am: 29.02.08
bei der Fakultät für Physik und Astronomie

1. Gutachter: Prof. Dr. Peter Jakob
 2. Gutachter: Prof. Dr. Wolfgang Bauer
 3. Gutachter: Prof. Dr. Markus von Kienlin
- der Dissertation.

1. Prüfer: Prof. Dr. Peter Jakob
 2. Prüfer: Prof. Dr. Wolfgang Bauer
 3. Prüfer: Prof. Dr. Georg Reents
- im Promotionskolloquium.

Tag des Promotionskolloquiums: 21.07.08

Contents

1	Motivation	7
1.1	Magnetic Resonance Imaging	7
1.2	New Options for Fast MR Imaging	8
1.3	Goals of this Thesis	9
2	MRI Basics	11
2.1	Larmor Frequency	11
2.2	The Fourier Transform and Signal Sampling	12
2.3	Spatial Encoding using Gradients	14
2.3.1	Read Encoding Gradient	14
2.3.2	Phase Encoding Gradient	15
2.4	K-Space	16
2.5	Cartesian Imaging Experiment	18
2.6	Signal-to-Noise Ratio and Fast Imaging	21
2.7	Non-Cartesian Imaging	22
2.7.1	Trajectories	22
2.7.2	Applications of Non-Cartesian Imaging	26
2.8	Data Gridding	26
2.8.1	Convolution Gridding	27
2.8.2	Inversion Methods	28
2.8.3	Iterative Methods	28
2.8.4	Non-Uniform Fourier Transform	29
3	Parallel Imaging	31
3.1	Basics of Parallel Imaging	31
3.1.1	Historical Overview	31
3.1.2	General Remarks on Parallel Imaging	33
3.1.3	SNR Loss and Reconstruction Errors in Parallel Imaging	35
3.2	Cartesian Parallel Imaging	35
3.2.1	Generalized Autocalibrating Partially Parallel Imaging (GRAPPA)	36
3.2.2	GRAPPA Operator	39

3.2.3	GRAPPA vs. the GRAPPA Operator	41
3.2.4	Other Parallel Imaging Methods	42
3.3	Non-Cartesian Parallel Imaging	45
3.3.1	Non-Cartesian GRAPPA	45
3.3.2	Other Non-Cartesian Methods	49
3.3.3	Non-Cartesian GRAPPA vs. CG-SENSE	53
3.4	Applications of Non-Cartesian Parallel Imaging	53
3.5	The Current State of Non-Cartesian Parallel Imaging	54
4	GRAPPA Operator Gridding (GROG)	57
4.1	Introduction	57
4.2	Theory	58
4.2.1	Properties of the GRAPPA Operator	58
4.2.2	Gridding with the GRAPPA Operator	59
4.2.3	GROG Weight Determination	62
4.3	Methods	63
4.3.1	Simulations	63
4.3.2	In Vivo Experiments	64
4.4	Results	65
4.4.1	Simulation	65
4.4.2	In Vivo Experiments	68
4.5	Discussion	69
4.6	Conclusion	75
5	Self-Calibrating GRAPPA Operator Gridding	77
5.1	Introduction	77
5.2	Theory	78
5.2.1	GROG Review	78
5.2.2	Self-Calibrating GROG	79
5.3	Methods	83
5.3.1	Simulation	83
5.3.2	Experiments	83
5.4	Results	84
5.5	Discussion	84
5.6	Conclusion	88
6	Pseudo-Cartesian GRAPPA	89
6.1	Introduction	89
6.2	Theory	90
6.2.1	Cartesian and non-Cartesian GRAPPA	90

6.2.2	Pseudo-Cartesian GRAPPA	92
6.3	Methods	97
6.3.1	Data Reconstruction	97
6.3.2	Simulations	98
6.3.3	In vivo Experiments	99
6.4	Results	100
6.4.1	Simulations	100
6.4.2	In vivo Experiments	102
6.5	Discussion	103
6.6	Conclusion	107
7	GROG-Facilitated Bunched Phase Encoding	109
7.1	Introduction	109
7.2	Generalized Sampling Theorem	110
7.3	Bunched Phase Encoding	112
7.4	GROG-BPE	113
7.5	Simulation	114
7.5.1	Methods	114
7.5.2	Results	116
7.6	In Vivo Experiments	120
7.6.1	Methods	120
7.6.2	Results	123
7.7	Discussion	124
7.8	Conclusion	127
8	Summary	129
9	Zusammenfassung	133
	Bibliography	137
	Curriculum Vitae	149
	Publications	151
	Acknowledgements	155
	Erklärung	157

Chapter 1

Motivation

1.1 Magnetic Resonance Imaging

The field of medical imaging offers an abundance of options for obtaining anatomical or functional information about the human body. These methods include Computed Tomography (CT), ultrasound, Positron Emission Tomography (PET), Single Photon Emission Computed Tomography (SPECT), and Magnetic Resonance Imaging (MRI). While each of the techniques mentioned here is non-invasive, only MRI offers both high resolution and soft tissue contrast without employing ionizing radiation (unlike CT, PET, or SPECT). MRI allows the operator to select a desired contrast in an arbitrarily positioned slice, and can be used to generate dynamic image series which show anatomical motion. Angiography (MRA), functional imaging (fMRI), and spectroscopy (MRS) are also possible with MR, offering the opportunity to obtain, for instance, late-enhancement perfusion images, signals indicating brain activity, or information about chemical composition in tissues. These properties make MRI an attractive imaging modality.

However, MRI is plagued by a number of drawbacks. One major problem is that performing a full clinical MR protocol is relatively time-consuming in comparison to other imaging modalities, especially ultrasound and CT. This difficulty has many practical consequences. The first is that due to the high costs of building (i.e. buying) and maintaining an MR scanner, the long measurement times required for many studies lead to a high cost per patient. Secondly, the narrow bore of the MR scanner can potentially lead to feelings of discomfort or even claustrophobia. Thus, uncooperative patients, children, or even impatient test subjects often fidget or struggle when they must lie still in the scanner for a long time. This patient motion can lead to artifacts or misregistration problems in the subsequent MR images, which must then be reacquired, further increasing the total scan time. For these reasons, it would be highly desirable to reduce the total amount of time a patient spends in the MR scanner.

The long acquisition time for single images also causes problems in MRI. As stated above, patient motion can lead to unwanted effects in MR images. Thus, in addition to voluntary patient motion, breathing, cardiac or smooth muscle motion, and blood flow must be minimized during the scan to avoid blurring or other artifacts. While breath-hold techniques, cardiac gating, flow compensated or saturation methods, and/or segmentation can be used to improve image quality, these options cannot

be employed in all cases. Ideally, it would be possible to acquire data from a complete 2D slice or 3D volume fast enough to generate an image on a time scale not affected by the motion. This would allow one to depict, for instance, the movement of the beating heart at a real-time frame rate. However, even in cases where such a fast acquisition is possible, the image contrast options are then restricted, thereby compromising a major advantage of MRI. Another challenge is that distortions and signal cancellation can also occur when attempting to acquire an entire image especially quickly in what is known as a “single shot” method. For functional MRI (fMRI), the need for fast imaging is additionally motivated by the desire to capture changes in the brain quickly. Thus, it is advantageous to shorten MR scan times not only for cost effectiveness, patient comfort, and artifact reduction, but also to acquire dynamic information about structure or function at a frame rate fast enough to depict the relevant changes.

In the past, options for the reduction of scan time while maintaining a given image contrast were available, but limited. Unfortunately, the spatial resolution and time resolution are inversely related in MRI; a higher resolution image requires more time to acquire than a lower resolution image. Thus, to reduce the total imaging time, the resolution can be decreased. However, high resolution images are often needed for the detection and classification of disease, making this option less attractive. Another option is the use of advanced reconstruction techniques which rely on the conjugate symmetry of MR data. Such methods, while viable in many cases, allow for a maximum time reduction factor of 2. Thus, with the increasing use of MR in a clinical setting in the late 1990’s, the search for alternative acceleration options was well underway.

1.2 New Options for Fast MR Imaging

With advances in MR technology at the end of the last century, several new options for MR imaging which ameliorate some of the above mentioned difficulties have become available. The first major advance was the development of multi-coil arrays and multi-channel MR scanners. Such arrays and scanners allow the simultaneous acquisition of the image information at different spatial locations around the object. While the primary reason for building these multi-coil arrays was to increase the SNR in the resulting image, an interesting side-effect is that each coil in such an array “sees” the object to be imaged differently. It did not take long for many researchers to realize that this additional spatial information could be used to reduce the amount of spatial information which must be acquired by the MR scanner, in turn reducing the amount of time needed for the scan. The first major breakthrough in this new field of “parallel imaging” came in 1997 with the introduction of SMASH by Sodickson et al. [1]. Other techniques, such as AUTO-SMASH [2], SENSE [3], SPACE RIP [4], PILS [5], VD-AUTO-SMASH [6], and GRAPPA [7], closely followed. These techniques all allow one to generate an artifact-free image using an undersampled (i.e. not fully-sampled) dataset. The most commonly used techniques in the clinic today are SENSE and GRAPPA, both of which can allow for time reductions of up to a factor of 8, given an appropriate array coil set-up.

Another technological advance which led to new options for image acquisition is the increasingly robust gradient systems used in MR scanners. Although MR data is usually acquired in a rectilinear (or

Cartesian) fashion, advanced gradients allow one to acquire data along “non-Cartesian” trajectories. Such non-Cartesian datasets, which are generally made up of data acquired along a star- or spiral-shaped trajectory, have some speed advantages over standard Cartesian imaging. Some trajectories are also inherently motion insensitive, meaning that motion such as blood flow does not result in strong artifacts in the resulting images. In addition, specialized trajectories can be used to correct for patient motion, and can be helpful when working with uncooperative patients or children. Thus, non-Cartesian imaging is another tool which can be used in some applications to reduce imaging time and improve image quality.

The combination of parallel imaging and non-Cartesian imaging would allow for additional benefits. As an example, accelerated “single-shot” non-Cartesian datasets often contain fewer distortion artifacts than their Cartesian counterparts. However, a number of obstacles make this combination non-trivial. One major hurdle is the difficulty in general when working with non-Cartesian data. Because the datapoints do not lie on a rectilinear grid and a simple Fourier transformation cannot be used for image processing, they are much more difficult to work with non-Cartesian than Cartesian data. Similarly, many basic assumptions for Cartesian parallel imaging are not fulfilled in the non-Cartesian data, and each of the methods listed above must be modified before it can be applied to non-Cartesian data (if it can be used at all). Both of these challenges have been addressed by a number of researchers, and a plethora of different solutions have been proposed, many of which are difficult to employ without significant experience in non-Cartesian imaging. For instance, many algorithms exist which convert non-Cartesian to Cartesian data, known as gridding algorithms, but they require input parameters which, when selected non-optimally, can lead to image artifacts. Non-Cartesian parallel imaging reconstruction methods are also often time consuming and complicated, and require additional datasets which can be difficult or inconvenient to acquire. Thus, straight-forward and parameter-free methods for non-Cartesian data reconstruction, both for data gridding and for non-Cartesian parallel imaging, are currently the subject of intensive research.

1.3 Goals of this Thesis

This thesis gives a summary of my work combining non-Cartesian MRI and parallel imaging. Chapters 2 and 3 offer overviews of Magnetic Resonance Imaging, including non-Cartesian imaging and parallel imaging, respectively. Chapter 4 details my work in developing a novel data gridding technique using parallel imaging, which simplifies the gridding process for datasets acquired along non-Cartesian trajectories. This new method, GRAPPA Operator Gridding (GROG), has a number of advantages over standard gridding methods, making it both robust and easy to implement and use. Following the introduction to GROG, a method of determining the few parameters intrinsic to GROG using the non-Cartesian datapoints themselves is discussed in Chapter 5. Self-Calibrating GROG (SC-GROG) is an interesting development for non-Cartesian data gridding, because it requires no information besides the non-Cartesian trajectory. Other gridding methods are much more unwieldy, as many input parameters must be optimized in order to arrive at the final gridded image. Thus, with the introduction of

SC-GROG, one of the major difficulties in non-Cartesian imaging has been removed.

Once the concept of GROG has been developed, the thesis turns to the use of special properties of GROG to simplify non-Cartesian parallel imaging. Unlike Chapters 4 and 5, which discuss gridding fully-sampled non-Cartesian datasets, the next two chapters deal with reconstructing artifact-free images from undersampled non-Cartesian datasets. Chapter 6 introduces a new method of reconstructing arbitrarily undersampled datasets using a modified version of Cartesian GRAPPA. By first using GROG to grid the undersampled non-Cartesian data (a feat which cannot be accomplished by other gridding techniques), the GRAPPA reconstruction process is greatly simplified. In Chapter 7, the use of GROG to mimic a trajectory known as Bunched Phase Encoding is described. The basic idea is to generate bunches of points around each acquired point by shifting the acquired data in a manner similar to gridding with GROG. Unfortunately, shifts large enough to completely reconstruct the missing datapoints cannot be performed with GROG. However, by exploiting the Generalized Sampling Theorem of Papoulis, one can perform an iterative reconstruction using the artificial bunched points to generate artifact-free images. It is important to note that the two non-Cartesian image reconstruction methods described here depend on GROG, and cannot be performed without this technique.

Finally, Chapter 8 contains a summary of the work presented in this thesis, as well as a discussion of future directions that could be explored using the ideas presented here.

Chapter 2

MRI Basics

2.1 Larmor Frequency

The idea of nuclear spin and the resulting spin angular momentum is a familiar concept in nuclear physics [8]. The spin of an atomic nucleus can be characterized by its quantum number, which for protons (and many other nuclei, such as ^{13}C , ^{15}N , ^{19}F , and ^{31}P) is $s = \frac{1}{2}$. The relationship between the spin of a nucleus and the magnetic moment $\vec{\mu}$ that arises from it can be described as

$$\vec{\mu} = \gamma \hbar \vec{S} \quad (2.1)$$

where the parameter γ is known as the gyromagnetic ratio, which is specific to the nucleus in question (approximately 2.674×10^8 rad/Ts for protons), \hbar is Planck's constant divided by 2π , and \vec{S} is the total spin angular momentum operator.

Descriptions of the interaction between nuclei and an external magnetic field, i.e. Nuclear Magnetic Resonance (NMR), first appeared as early as 1946 [9, 10]. When an outside magnetic field, B_0 , is applied to a nucleus with a non-zero magnetic moment, the spin will seek to align itself either parallel or anti-parallel to the field in accordance with the quantization of angular momentum. For a nucleus of spin $\frac{1}{2}$, this creates two distinct energy levels:

$$E_{\pm} = \pm \frac{\gamma \hbar B_0}{2} \quad (2.2)$$

This phenomenon of energy splitting is known as the Zeeman effect [11]. If radiation (in NMR or Magnetic Resonance Imaging (MRI), generally in the form of a radio-frequency pulse) is applied to the aligned spins with an energy equal to the energetic difference between the two states, or $\Delta E = \gamma \hbar B_0$, a transition between the two energy levels occurs. The characteristic resonance frequency of this transition is known as the Larmor frequency:

$$\nu_L = \frac{\Delta E}{h} = \frac{\gamma \cdot B_0}{2\pi} \quad (2.3)$$

This resonance frequency can also be thought of as the frequency at which the spins precess around the axis of the magnetic field. As seen in the equations above, the Larmor frequency is directly proportional to the magnetic field strength, and is specific to a chemical species. For instance, the Larmor frequency for ^1H nuclei (i.e. protons) is 42.5781 MHz at a field strength of 1 T. Because of the large amount of water in the human body and their relatively high gyromagnetic ratio, protons are the most commonly observed nuclei in both NMR and MRI.

As stated above, the z-component of the angular momentum, S_z , of nuclei with $s = \frac{1}{2}$ aligns itself to be parallel or anti-parallel to an applied magnetic field; these spins are generally spoken of as being parallel or anti-parallel to the magnetic field. Because the spins which are parallel have a slightly lower energy than those that are anti-parallel, there is a slight excess of these parallel spins. The resulting difference in the population leads to a net macroscopic magnetic moment in the direction of the external magnetic field. In order to detect the signal from this macroscopic magnetization vector, it must be tipped away from the z-direction (i.e. the axis of the main magnetic field B_0) into the transverse plane. An RF excitation pulse, $B_1(t)$, of the appropriate resonance frequency applied perpendicular to the magnetic field B_0 will cause the magnetic moment to experience a torque, which in turn causes the magnetization to tip away from the axis of the field. The flip angle, α , is defined as the angle through which the magnetization is rotated by the radiofrequency (RF) pulse after it has been applied for a time τ :

$$\alpha = \gamma \int_0^{\tau} |B_1(t)| \cdot dt \quad (2.4)$$

Thus, the pulse forces the net magnetization to move away from the longitudinal direction and into the transverse plane, i.e. the x-y plane. The resulting precession of the spins leads to signal which can be detected by a receiver coil through magnetic induction.

2.2 The Fourier Transform and Signal Sampling

There is usually more than one resonance frequency in a sample due to either the presence of different chemical species or spatially dependent magnetic field gradients. In this case, the signal that results from the rotating magnetization is a combination of these frequencies weighted by their respective amplitudes. In order to obtain a spectrum which can be used to gather information about the sample, the signal must be deconvolved into its component frequencies. The Fourier transform is ideal for accomplishing this task.

The continuous Fourier Transform $f(t)$ of a function $F(\nu)$ and its reciprocal, the inverse Fourier transform, can be written as follows:

$$f(t) = \int_{-\infty}^{\infty} F(\nu) e^{-i2\pi\nu t} d\nu \quad (2.5)$$




Figure 2.1: An example of two frequencies sampled at different rates. The frequency of the signal shown in black is twice that of the signal shown in gray. When the sampling rate is chosen to be $\frac{1}{\Delta t_1}$, the points on the curves labeled as black circles are acquired. Note that the two frequencies cannot be distinguished from one another when this sampling rate is chosen, and a Fourier transform would represent them as a single peak. This phenomenon is called aliasing. However, when the sampling rate is chosen to be $\frac{1}{\Delta t_2}$, i.e. twice the highest frequency present, the points depicted as gray squares are also acquired. With this information, the two frequencies can be distinguished from one another, and no aliasing occurs after the Fourier transform.

$$F(\nu) = \int_{-\infty}^{\infty} f(t)e^{i2\pi\nu t} dt \quad (2.6)$$

Unlike as depicted in the equations above, the data collected by the spectrometer are discrete, and a number n data points are collected in intervals of Δt . The Shannon sampling theorem [12] states that a continuous signal can be properly reconstructed from a discretely sampled function given that the sampling rate is at least equal to two times the highest frequency in the signal being sampled:

$$|\nu_{max}| = \frac{1}{2\Delta t} \quad (2.7)$$

This theorem can be understood by examining the discrete Fourier transform:

$$F_D(\nu) = \sum_{n=0}^{N-1} f_D(n\Delta t)e^{2\pi i\nu \frac{n\Delta t}{N}} \quad (2.8)$$

where t has been replaced with $n\Delta t$ to denote the discrete nature of the sampled points. It is clear that only frequencies which fall between the range of $[-\nu_{max}, \nu_{max}]$ can be uniquely distinguished due to the periodicity of the exponential function, and any frequencies that appear outside of this range will be replicated, or aliased, in this interval. This maximum sampling frequency is often called the Nyquist sampling rate [13], and the sampling law is also known as the Nyquist criterion.

The consequences of violating the Nyquist criterion can be seen in Figure 2.1. As stated above, if the sampling frequency is less than the Nyquist rate, frequencies at fractional multiples of the original

frequency cannot be properly distinguished. This leads to signal aliasing, where two or more signal components appear at the same frequency due to the insufficient sampling rate. Aliasing is an important concept in MRI and specifically parallel imaging, and will be discussed in more detail in Chapter 3.

Although the Nyquist criterion is generally accepted as a strict rule for the sampling of MRI datapoints, exceptions and extensions to this rule are possible. For instance, some sampling schemes are less prone to aliasing artifacts; such trajectories are discussed in Section 2.7. In addition, research into the sampling of continuous functions has shown that there are cases in which the Nyquist criterion can be violated without aliasing. The Generalized Sampling Theorem of Papoulis [14] states that as long as the average sampling rate is equal to the Nyquist rate, the continuous signal can be reconstructed, given that the signal is bandlimited (which is the case in MRI). While this theorem is not often used when sampling MRI data, it is important to remember that the appearance of aliasing artifacts is not always dictated by the Nyquist criterion.

2.3 Spatial Encoding using Gradients

As described in Section 2.2, nuclei under the influence of different magnetic fields precess with different frequencies and lead to distinguishable peaks in a spectrum. This property can be used to make a distinction between nuclei present at different spatial locations. If the magnetic field varies in space, the spins will precess with frequencies that depend on their locations in the sample. Such a spatial dependency in the magnetic field can be created using magnetic field gradients; a typical gradient has the form

$$G_x = \frac{\partial B_z}{\partial x} \quad (2.9)$$

where G_x would be referred to as the gradient strength in the x-direction. The application of such a linear gradient yields the following spatially-dependent magnetic field:

$$B(x) = B_0 + xG_x \quad (2.10)$$

Thus, by applying a magnetic field gradient, spins at different locations can be differentiated by examining their resonance frequencies. The use of gradients for spatial encoding in this manner is known as Magnetic Resonance Imaging (MRI), and the basic concepts of gradient encoding are described below. For a more detailed explanation of MRI and gradient encoding, Reference [8] is recommended to the reader.

2.3.1 Read Encoding Gradient

A read gradient is a gradient that is applied during the acquisition time, or read-out time, of an experiment. In the simplest 1D imaging experiment (described by the Nobel laureate Paul Lauterbur in 1973 [15]), an RF pulse is applied to the magnetization, which rotates it into the x-y plane. At this point, given that there is only one chemical species present in the sample and that B_0 is constant over the FoV, all the spins would precess at the same frequency. However, if a linear, time-independent

gradient is applied during the read-out, spins will precess at frequencies which depend linearly on their location in the sample:

$$\omega(x) = \gamma(B_0 + xG_x) \quad (2.11)$$

or, moving into the “rotating frame,” where the x and y axes precess around the z-axis at a frequency ω_0 with respect to the laboratory frame:

$$\omega_{rot}(x) = \gamma x G_x \quad (2.12)$$

where x is the position of the magnetization on the x-axis, and G_x is defined as above. In standard Cartesian imaging, this gradient is applied during the entire acquisition time, which results in spatially dependent resonance frequencies along one axis.

After the data have been acquired, a Fourier transform can be performed to yield peaks whose frequencies show the one-dimensional spatial distribution of the spins in the sample. The amplitude of each of the peaks describes the amount of magnetization in each location.

2.3.2 Phase Encoding Gradient

For a one-dimensional experiment, encoding with a read gradient alone would suffice. However, to create two or three-dimensional images, a phase encoding gradient is a practical approach to encode the information in the second or third dimension. The phase encoding gradient is a gradient applied in a direction perpendicular to the read direction, after the excitation pulse but before the acquisition. Let x be the read encoding direction and y be the phase encoding direction. The application of such a phase encoding gradient, G_y , causes the spins to precess with spatially dependent frequencies in the y -direction, i.e. the effect is the same as that of the read gradient:

$$\omega_{rot}(y) = \gamma y G_y \quad (2.13)$$

When the gradient is turned off, the spins have acquired a phase, which depends on their location along the y -axis of the sample:

$$\phi_y = \gamma y G_y t \quad (2.14)$$

where t is the total amount of time the phase gradient has been applied. This phase modulation is not affected by gradients applied in perpendicular directions during the rest of the experiment, including the read gradient. However, unlike in one-dimensional read encoding, one data acquisition period no longer supplies enough information to create an image, as only one phase shift has been induced in the y -axis. Thus, the experiment must be repeated N_y times, where N_y is the desired number of points in the y direction. With each experiment, either the phase gradient strength, G_y , or the amount of time the phase gradient is applied, t , is changed, giving the spins a different phase shift for each experiment. Once the data is acquired, a 2D Fourier transform can be performed, yielding a two-dimensional image of spin density.

2.4 K-Space

Up to this point, only linear gradients have been discussed. For a generalized gradient, the signal induced by the spins at a time t during the application of the gradient can be written as:

$$S(t) \propto \int \rho(\vec{x}) \cdot e^{-i \int_0^t \gamma \cdot \vec{G}(t') \vec{x} dt'} d\vec{x} \quad (2.15)$$

where $\rho(\vec{x})$ is the spin density at the location \vec{x} , and $\vec{G}(t)$ is a gradient in an arbitrary direction with any functional form. A useful simplification is

$$\vec{k}(t) = \int_0^t \gamma \cdot \vec{G}(t') dt' \quad (2.16)$$

where, for the usual case of a linear (i.e. time-independent) encoding gradient:

$$\vec{k}(t) = \gamma \cdot \vec{G}t \quad (2.17)$$

With the parameter \vec{k} , the signal acquired during an experiment can be better understood. For one value of \vec{k} , the signal can be written as:

$$S(\vec{k}) = \int \rho(\vec{x}) \cdot e^{-i\vec{k}\vec{x}} d\vec{x} \quad (2.18)$$

which is exactly the Fourier transform of the spin density. Inversely,

$$\rho(\vec{x}) = \frac{1}{2\pi} \int S(\vec{k}) \cdot e^{i\vec{k}\vec{x}} d\vec{x} \quad (2.19)$$

Thus, for a Cartesian sampling of the values of $S(\vec{k})$, otherwise known as k-space, the spin density in each location can be calculated using the inverse fast Fourier transform, as discussed in Section 2.2. K-space has the same dimensional order as the image to be acquired, i.e. a two-dimensional image has a two dimensional k-space, one dimension for the read encoding direction and one for the phase encoding direction. The idea of employing k-space to describe gradient encoding was introduced by Twieg [16], and this convention greatly simplifies the concept of the time-domain signal in MRI.

In an actual experiment, only a finite number of k-space points can be acquired. Thus, a two-dimensional k-space is made up of a grid of points. The Nyquist criterion presented in Section 2.2 can be understood to imply that the distance between each point in k-space is inversely related to the field-of-view (FoV) of the experiment:

$$FoV \propto \frac{1}{\Delta k} \quad (2.20)$$




Figure 2.2: A representation of the relationships between parameters in k-space and in the image domain. The FoV in the image domain is proportional to the inverse of the distance between k-space points; when the points are closely spaced, the FoV is larger than when the k-space points are farther apart. If the k-space points sampled lie too far apart, the FoV becomes smaller than the object to be imaged, and aliasing occurs. The resolution in the image domain is proportional to the inverse of the maximum k-space point sampled. If only the center portion of k-space were sampled, the image that results would have the same FoV, but a lower resolution.

Any spin density that exists outside this range will lead to fold-over artifacts, or aliasing, in the image. This is due to the fact that the k-space domain must be sampled at a rate fast enough to separate each of the frequencies present in the acquired signal. Thus, for standard imaging experiments, the entire object must fall into the field-of-view defined by the gradients. An exception to this rule is discussed in depth in Chapter 3, where parallel imaging is examined. When using the reconstruction methods described there, it is possible to use a field-of-view that is not large enough to cover the entire object to be imaged.

The resolution, on the other hand, is given by the extent of the k-space sampled:

$$\Delta x \propto \frac{1}{2 \cdot k_{x,max}} \quad (2.21)$$

The relationship between the image dimensions and the k-space parameters is shown in Figure 2.2. The k-space location in an experiment can also be understood as the momentum of a gradient pulse; for a linear gradient, the gradient strength multiplied by the length of the gradient is applied. The idea of gradient momentum is important when examining refocusing and defocusing of magnetization using gradients, as described in the next section.




Figure 2.3: A schematic of the pulse sequence for a gradient echo experiment. The top line shows the RF and acquisition channel, and the following three lines show the phase, read, and slice gradient axes, respectively. The events in the gradient echo experiment, and their timing, are explained in detail in the text.

2.5 Cartesian Imaging Experiment

This section describes how the acquisition of an image is performed in standard 2D Cartesian MRI. The basic gradient echo sequence (i.e. spin warp [17, 18]) is used as an example, and is made up of the following fundamental steps:

- application of an α pulse coupled with the slice-selection gradient
- application of a slice rephasing gradient
- dephasing of the spins in the read direction by a negative read gradient for a time t_1
- application of the phase encoding gradient
- refocusing of the spins by a positive read gradient for a time t_2
- simultaneous acquisition

Each of these steps can be seen in the pulse sequence for the gradient echo experiment, shown in 2.3. This pulse sequence diagram shows what occurs on each gradient axis, as well as when RF pulses are given and when data is acquired during the sequence.

In a standard 2D imaging experiment, two dimensions of k-space must be acquired. However, before the magnetization can be encoded as described above, it must be first transferred from the longitudinal axis to the transverse axis. This is performed by applying an RF pulse, as described in Section 2.1. For example, a 90° excitation pulse brings the magnetization completely into the x-y plane, whereas a pulse with less energy, generally denoted as α , tips the magnetization vector only partially into the transverse plane. This excitation pulse has a specific bandwidth, i.e. it only excites frequencies in a certain interval. In order to select the 2D slice within the three-dimensional object to be imaged, a gradient (G_{SS}) is applied simultaneously with the pulse, such that only the spins within the selected slice will be in the appropriate frequency interval and therefore tip into the x-y plane. The slice thickness Δs is calculated from the pulse bandwidth Δf and the gradient strength G_{SS} by employing the following equation:

$$\Delta s = \frac{\Delta f}{\gamma \cdot G_{SS}} \quad (2.22)$$

In order to rephase the magnetization, a slice rephasing gradient (G_{SR}) with half the momentum of the original slice selection gradient is applied after the slice gradient has been completed.

Once the magnetization has been brought into the transverse plane, it is completely unencoded in the k_x and k_y directions (although it is encoded in the slice direction). Thus, it can be said to be at the center of k-space, where $k_x = k_y = 0$. The next portion of the experiment is to apply the first gradient along the read direction, namely G_D for time τ_1 , which causes the magnetization to dephase. This can also be seen as a movement away from the center of k-space in the k_x direction (the horizontal axis) as shown in Figure 2.4. A phase encoding gradient (G_P) is applied (in this pulse sequence, at the same time as the read dephasing gradient) in the phase encoding direction, which is orthogonal to the read encoding direction. This gradient can also be seen as a movement in k-space, this time in the k_y direction. The strength and length of the phase encoding direction determines the k-space location, as described in Section 2.4. The simultaneous application of these gradients leads to the movement in k-space between t_1 and t_2 shown in Figure 2.4.

At this point in the experiment, the magnetization has been properly prepared, and signal acquisition can take place. To this end, another gradient, G_R , is applied in the read direction, this time opposite in sign to the first gradient. This second gradient causes the magnetization to rephase, which can also be seen as a movement towards the center of k-space in the read direction. A complete rephasing of the spins results in a magnetization echo, and the continued application of the read gradient results in a further dephasing of the spins. The magnetization echo occurs at the point where the momentum of the first gradient is exactly cancelled by that of the second gradient; given that $\tau_1 = \frac{\tau_2}{2}$, and the magnitudes of the gradients are equal ($G_R = -G_D$), then the echo occurs in the middle of the second gradient. Thus, in this case, the echo appears in the center of the read-out line in k-space, and the echo time T_E can be defined as the time between the center of the excitation pulse and the appearance of the echo. The use of such an echo allows the acquisition of one half of a line of k-space during the rephasing of the magnetization and the other half of the line during the further dephasing of the spins.




Figure 2.4: A schematic of the k-space trajectory and the timings for the gradient echo experiment. At time t_1 , the magnetization has been brought into the transverse plane by the RF pulse, but no encoding has occurred. Both the read dephase and the phase gradients are applied between t_1 and t_2 , which can be represented as a movement in k-space in the k_x and k_y directions, respectively. The combination of the two yields a diagonal path in k-space. At time t_2 , the phase gradient is turned off, and the read gradient started, which is opposite in direction to the read dephase gradient. In k-space, this results in a movement from left to right. During this time, data is acquired. Given that the strengths of the read and the read dephasing gradients are equal but opposite, the gradient echo occurs at time t_3 , in the middle of k-space. The acquisition is continued until time t_4 , when the read gradient is turned off and the entire k-space line has been acquired. Note that one k-space line is acquired during this process, and other lines with different phase encoding gradient strengths must be acquired to completely fill the k-space.

As can be seen in Figure 2.4, only one line of k-space is acquired during such an experiment. In order to obtain a 2D image, multiple phase encoding steps must be applied to cover the entire k-space. Thus, for a final matrix size of 128×128 , 128 points must be acquired in the read-direction during the magnetization echo, and 128 phase encoding steps must be performed. Given that the repetition time, or the time between subsequent excitation pulses, of the experiment is denoted as T_R , the total experiment time is equal to $128 \cdot T_R$. Once the entire k-space has been acquired, a two-dimensional inverse Fourier transform is applied to arrive at the final image.

In a three-dimensional imaging experiment, two separate phase encoding directions are used, and the third dimension is encoded using a read gradient, as in the 2D case. Thus, if a final matrix of $128 \times 128 \times 128$ is desired, a total of 128×128 phase encoding steps must be performed. After the 3D k-space has been acquired, a three dimensional Fourier transform is used to obtain the 3D image.

2.6 Signal-to-Noise Ratio and Fast Imaging

An important indicator of the image quality in MRI is the Signal-to-Noise Ratio (SNR), which can be defined as the signal level divided by the standard deviation of the image noise:

$$SNR = \frac{S}{\sigma_N} \quad (2.23)$$

The SNR depends on a number of factors, the most important of which are discussed below.

The SNR in Equation 2.23 can broadly be written for a 2D experiment as:

$$SNR \propto \Delta x \cdot \Delta y \cdot \Delta z \cdot \sqrt{N_x \cdot N_y \cdot N_{rep}} / \sqrt{BW} \quad (2.24)$$

where Δx and Δy denote the size of the voxels in the x- and y-directions, Δz is the slice thickness, N_x and N_y the number of points acquired in the respective directions, N_{rep} the number of averages performed, and BW the acquisition bandwidth (equal to $\frac{1}{\Delta t}$). The SNR is seen as a fundamental limit in MRI, as it is generally not possible to accelerate the imaging process without decreasing the SNR of the resulting image. However, for applications where the SNR is high enough to accommodate such losses, possibilities for increasing the imaging speed can be examined.

As stated in the previous section, the total experiment time can be written as:

$$T_{exp} = T_R \cdot N_y \cdot N_{rep} \quad (2.25)$$

where N_y is the total number of phase encoding lines that must be acquired. Thus, the simplest way to shorten the total experiment time is to decrease the number of signal averages, N_{rep} . However, assuming that this value is already at a minimum, i.e. $N_{rep} = 1$, another option must be chosen. The next parameter which one can reduce is the number phase encoding steps, N_y . This can be done in two ways; either the distance between adjacent phase encoding lines can be increased, or the outer phase encoding lines can be left unsampled. In the former case, the FoV of the image in the phase encoding direction is decreased, and in the latter case, the resolution of the image will be decreased (see Figure 2.2). However, if the FoV is reduced such that it is smaller than the object to be imaged, fold-over artifacts will appear as a result of the violation of the Nyquist criterion described in Section 2.2. Conversely, a reduction in resolution can make it more difficult to distinguish small features in the image. Thus, before the advent of parallel imaging, decreasing the number of phase encoding steps was not a viable way of significantly decreasing the total experiment time. The topic of parallel imaging is discussed in great detail in Chapter 3.

Finally, one could also decrease the repetition time, T_R . For the gradient echo sequence shown in above, the shortest T_R possible is dictated by the gradient strengths and lengths. In order to shorten the minimum T_R by a factor of two and keep the same FoV and resolution, the gradient lengths must be decreased, and the strengths increased, each by a factor of two. In addition, the gradient ramp times, the so-called switching times, must be increased by a factor of four. This has indeed been done; while gradient strengths of 10mT/m and switching times of more than 1 ms were common in the 1980s,

modern scanners work with gradient strengths of 120 mT/m and switching times of less than $50\mu\text{s}$. However, further increases in the area of gradient technology appear unfeasible due to the massive power needed to ramp and apply such gradients. In addition, physiological stimulation of the patient caused by fast gradient switching forms an absolute boundary for the development of powerful gradient systems. Current MR scanners often operate at the limit of acceptable gradient speed, which means that a further decrease in the total experiment time using faster gradients is not expected.

Naturally, one can also change the way the data are acquired in order to accelerate the collection of data. One possibility is to acquire multiple echoes during the same acquisition window, as in the RARE sequence [19]. The most extreme example is the single shot experiment, where the entire k-space is acquired using only one excitation, which is performed when using EPI [20]. However, these experiments necessarily change the contrast of the image and are often plagued by undesirable relaxation effects and image distortions, although the total experiment time can indeed be decreased.

However, there are a number of other options for increasing the speed of an imaging experiment. These methods generally rely on decreasing the number of phase encoding steps needed to acquire an image without changing the resolution or FoV. One class of methods for reconstructing an unaliased image from an incomplete dataset, namely partial Fourier methods [21, 22], is an option for the acceleration of data acquisition. These methods are generally limited to time reductions of less than one-half. However, three other methods which allow for larger time reductions are examined and combined in this work; they are non-Cartesian trajectories, parallel imaging, and the Generalized Sampling Theorem of Papoulis. The first of these options, non-Cartesian imaging, is discussed below, and parallel imaging is the subject of Chapters 3 and 6. The Generalized Sampling Theorem, and its implementation as Bunched Phase Encoding in MRI [23, 24, 25], is treated in detail in Chapter 7.

2.7 Non-Cartesian Imaging

2.7.1 Trajectories

As stated above, one possibility for further decreasing imaging time is the use of non-Cartesian trajectories. Such trajectories do not traverse k-space as described in Section 2.5, i.e. in a line-by-line fashion, but instead take different paths which are dictated by the gradient applied. One can imagine altering the gradient in Equation 2.17 from a linear gradient to a time varying waveform, which would result in a k-space trajectory which differs from the straight read-out trajectories used in Cartesian imaging. Examples of the most often used non-Cartesian trajectories and their gradient waveforms are shown in Figure 2.5. For reference, the standard Cartesian acquisition scheme is shown at the top of this figure. While these trajectories can sometimes be used to decrease imaging time, each also has additional characteristics which can be exploited to perform tasks which are not possible using Cartesian trajectories.

The most common non-Cartesian trajectory (and first MR trajectory implemented) is the radial trajectory [15, 26] (see Figure 2.5), which is generated using the gradient waveform depicted in the second row, center. The radial (otherwise known as projection reconstruction (PR)) trajectory is ad-




Figure 2.5: Depictions of several important non-Cartesian trajectories, their gradient waveforms, and typical undersampling artifacts.

vantageous because the center of k-space is sampled multiple times during the acquisition, making such a dataset relatively robust against motion or flow artifacts. In addition, although $\frac{\pi D}{2}$ radial projections are strictly required to fulfill the Nyquist criterion at the edges of k-space (where D is the diameter of the radial k-space), undersampling artifacts appear incoherent, i.e. as noise-like structures, or at worst as streaks in the image domain. Thus, the number of repetitions needed can be greatly reduced, especially in applications such as angiography, where the image domain contains sparse information. The most common 3D non-Cartesian trajectory, 3D radial [27], is based on the radial trajectory.

Another commonly used non-Cartesian trajectory is the spiral trajectory [28, 29], also shown in Figure 2.5. The advantage of this trajectory is that the complete k-space can be covered in a single shot, i.e. only one excitation is required to obtain the entire image. Although the spiral trajectory can be used for rapid image acquisitions (50 to 100 ms), the long read-out of this trajectory means that the resolution can be limited by the T_2^* relaxation parameters of the sample. If this parameter is short, high resolution images cannot be obtained due to the lack of signal in the outer portions of k-space, and image distortion can occur due to off-resonance effects. Another possibility is to employ a multi-shot spiral sequence, where each spiral arm has a shorter read out time than the single shot spiral, although more excitations are needed. In addition to the number of shots, the actual acquisition of the points along the spiral can also be varied. Two different spiral trajectories are commonly employed: the constant-linear-velocity trajectory and the constant-angular-velocity trajectory. While the constant-linear-velocity spiral has signal-to-noise advantages [30], the constant-angular-velocity spiral contains a quasi-radial symmetry [28], which is useful when post-processing steps are required for image reconstruction. One disadvantage of the spiral trajectory is that undersampling leads to coherent artifacts throughout the image, unlike the undersampled radial trajectory. These strong artifacts can be reduced by using a dual-density spiral, where the center of k-space is sampled at the Nyquist rate, and the outer portions are undersampled by increasing the spiral gradient velocity.

A third non-Cartesian trajectory is the PROPELLER or BLADE trajectory [31, 32]. This trajectory is made up of a number of Cartesian blades which are rotated to cover the entire k-space. Because multiple blades run through the center of k-space, and each blade can be reconstructed separately to form a low-resolution image, this trajectory can be used for motion correction. Similar to the radial trajectory, a large degree of undersampling can be performed with PROPELLER without obtaining coherent aliasing artifacts. While this trajectory can be more demanding in terms of acquisition time, its added motion correction advantages make it a popular acquisition scheme.

Two trajectories which are used in specific fast imaging applications, namely the zig-zag trajectory [24, 33] and the 1D non-Cartesian (1D NC) trajectory [34], are also depicted in Figure 2.5. The zig-zag trajectory, or Bunched Phase Encoding (BPE), while not widely used, has been demonstrated to have properties which make it unique. This trajectory allows the acquisition of groups of closely spaced points in the phase encoding direction, although the k-space overall is strictly undersampled. By making use of the Generalized Sampling Theorem discussed in Section 2.2, an unaliased image can be reconstructed despite the violation of the Nyquist criterion. This trajectory and the use of the Generalized Sampling Theorem for image reconstruction are the topics of Chapter 7, and will not be further discussed here.

The 1D NC, or variable density, trajectory is similar to the zig-zag trajectory in that it is often used in conjunction with parallel imaging. This trajectory is advantageous because it exhibits artifacts which are more incoherent than a purely Cartesian trajectory when it is undersampled; because the trajectory has few uses outside of parallel imaging, it will be discussed in more detail in Chapter 3.

Two final trajectories, rosette [35, 36, 37] and stochastic [38], have similar properties. The rosette trajectory is shown at the bottom of Figure 2.5, and the stochastic trajectory, which is random, is not depicted. Both of these trajectories can be used for spectral selection due to the overlapping points in k-space. While these acquisitions schemes are not as commonly used as the other non-Cartesian trajectories depicted here, they are interesting due to their lack of simple radial symmetry, and relatively incoherent aliasing artifacts. They are added as examples of exotic trajectories which are generally difficult to post-process, as will be seen in the following sections.

The list of non-Cartesian trajectories presented here is not meant to be comprehensive; many other trajectories such as Lissajou [39], twirl [40], and STAR [41] (all 2D trajectories), as well as cones [42], 3D stack of spirals or stars, and spherical shells [43] (all 3D) also exist. However, this work will focus primarily on the trajectories described here.

On the right-hand side of Figure 2.5, examples of the images resulting from Nyquist undersampling for each of the trajectories described are shown. It is important to note that the standard aliasing artifacts described in Section 2.2 do not appear in undersampled non-Cartesian data. This is because the undersampling present along such trajectories changes in severity and direction throughout k-space. These complex undersampling patterns lead to unusual aliasing in the image domain, as each pixel in the image can potentially alias with every other pixel, unlike in standard Cartesian imaging [44, 45]. The effects of undersampling in non-Cartesian data generally appear as a blurring or streaking in both (or all three) directions in the image, especially for trajectories with a radial symmetry. Specifically for the radial, rosette, and PROPELLER trajectories, very few artifacts are present because the center of k-space is still Nyquist sampled, leaving the contrast and structure of the image intact. Trajectories with a more Cartesian structure, namely 1D NC or zig-zag, lead to more traditional fold-over artifacts; in the case of 1D NC, such artifacts are not as pronounced as in Cartesian imaging, while in zig-zag they are spread into both directions. The unusual artifacts that are present in undersampled non-Cartesian images, and the fact that they are the result of irregular undersampling in k-space, are important ideas for non-Cartesian parallel imaging, discussed in detail in Section 3.3 and Chapter 6.

Although non-Cartesian trajectories often have advantages over Cartesian acquisitions, namely reduced scan time, insensitivity to motion artifacts and undersampling, and additional information such as spectral resolution, they also have disadvantages. The largest difficulty with such acquisition schemes is the need to perform data gridding, discussed in the next section. However, other difficulties such as ill-defined FoVs, location dependent resolution and signal-to-noise ratios, and off-resonance effects are problems that must be considered when choosing to use a non-Cartesian trajectory. However, for certain applications, the advantages of employing non-Cartesian imaging outweigh the difficulties inherent in such trajectories.

2.7.2 Applications of Non-Cartesian Imaging

As stated above, each non-Cartesian trajectory has specific properties which can be exploited to improve certain imaging applications. Several common clinical uses of non-Cartesian trajectories are described here to demonstrate the advantages of selecting the appropriate trajectory for a given application.

Radial trajectories are often implemented due to the fact that no gradient encoding is performed before the data acquisition begins, making such sequences valuable for imaging tissues with fast relaxation parameters. For instance, imaging of the lung parenchyma using radial trajectories has been demonstrated as early as 1993, and this research has been followed up by many groups [46, 47, 48]. 3D radial trajectories in the form of ultra-short echo time (UTE) sequences can also be used to image bone or cartilage tissues. The visualization of articular cartilage [49], the spine [50], entheses (the point at which a tendon inserts into the bone) [51], periostium [52], and the Achilles tendon [53] have all been demonstrated using UTE. In addition, UTE can be used to image sodium in the heart [54] as well as in the brain and skeletal muscle [55]. Functional sodium imaging has also been performed with UTE [56].

Spiral trajectories are used in order to generate images quickly due to their fast coverage of large portions of k-space. Functional MRI is an application where such fast image acquisition is desirable; thus, the spiral trajectory is often implemented for fMRI [57, 58, 59, 60].

Radial and spiral trajectories are also often used in fast cardiac imaging. One advantage of these trajectories is that they can be employed to acquire real-time cardiac data without requiring k-space segmentation (i.e. images from all heart phases can be acquired in a single R-R interval) [61, 62, 63, 64]. For instance, radial and spiral trajectories have been used to detect cardiac masses without ECG triggering [65], to quantify cardiovascular flow [66], to examine coronary vessel walls [67], etc. A three dimensional version of radial imaging is often employed in MR angiography due to the effective k-space coverage per time and potential for isotropic voxel sizes [68, 69, 70].

The PROPELLER trajectory has the advantage that significant motion artifacts can be corrected by examining the low resolution images made from each of the blades. Thus, when working with children or uncooperative patients, this trajectory can be employed, and any movement in the plane of the image can be subsequently removed with dedicated reconstruction algorithms [71, 72]. Additionally, PROPELLER in various forms is well-suited to diffusion tensor imaging (DTI) because this sequence displays less severe distortion artifacts than EPI [73, 74, 75].

2.8 Data Gridding

The most obvious difficulty with non-Cartesian trajectories is that the points acquired no longer lie on a Cartesian grid. Thus, before a fast Fourier Transform (FFT) can be applied to the k-space to generate an image, the non-Cartesian datapoints must be assigned to Cartesian grid points. There are a number of methods available for performing this data gridding; the most common will be examined here.




Figure 2.6: A 1D schematic of the convolution used in standard convolution gridding. The left-hand side shows the result of convolving a single non-Cartesian point, shown as the black circle in the center, with the Kaiser-Bessel (KB) function (the dotted line indicates that the KB function is centered around the non-Cartesian point). The resampling intensities at each of the regularly spaced Cartesian points, shown as white squares, can be calculated by multiplying the value of the non-Cartesian point with the appropriate value of the KB function (shown as solid vertical lines). The Cartesian point shown as a gray square receives some contribution from the non-Cartesian point because it is within the KB resampling window. The right-hand side shows all of the contributions to the Cartesian point shown in gray from the non-Cartesian points (white). The sum of these contribution yields the total value of the Cartesian point. These non-Cartesian points contribute to the Cartesian point because they fall in the KB window centered around the Cartesian point (solid line)

2.8.1 Convolution Gridding

The gold-standard gridding method in non-Cartesian MRI is convolution gridding [76]. This method involves convolving each non-Cartesian sampling point with a dedicated kernel and resampling the result at the appropriate Cartesian grid locations (see Figure 2.6). While O’Sullivan et al. [77] concluded that the optimal convolution kernel is a sinc function of infinite extent, the use of such a kernel is computationally unfeasible. Thus, the standard kernel used in MRI is a Kaiser-Bessel window, usually with a width of 3 or 5, extended to two (or three) dimensions [76]. The convolution gridding method also requires an oversampled gridding matrix in order to accurately determine the values of the grid points; grid oversampling factors of at least two, but as much as 16, are usually employed, depending on the time available and accuracy needed for the reconstruction. After the data have been resampled, a Fourier transform of the k-space data yields the image multiplied by the convolution kernel. A roll-off correction, which involves dividing the intermediate image by the Fourier transform of the convolution kernel, must then be performed to arrive at the final gridded image. A schematic of the convolution gridding method is shown in Figure 2.7.

Another requirement for convolution gridding is a density compensation function (DCF), which takes the variations in sampling density into account for the gridding process. While this DCF can be analytically calculated for closed form trajectories, such as radial or spiral, DCFs for more exotic trajectories must be determined iteratively or numerically. Thus, for trajectories such as rosette or stochastic sampling, convolution gridding is non-trivial to perform, and the choice of the DCF can greatly influence the final image.

While the convolution gridding method is the gold-standard for data resampling in MRI, it has a number of drawbacks. The first is, as stated above, the need for a DCF. Secondly, there are a number of parameters that must be determined before the data can be gridded, such as the optimal convolution kernel, the kernel width, and the grid oversampling size. In addition, the gridding process is computationally intensive, especially if large kernels or oversampling factors are chosen. Finally, it is not possible to accurately grid non-Cartesian data that have not been acquired at the Nyquist rate, as the convolution method assumes that the region of support is at least Δk . Any areas of data density which do not meet the Nyquist criterion will lead to errors in the roll-off correction discussed above. Due to these difficulties, a number of additional methods have also been proposed for data gridding, some of which will be discussed below.

2.8.2 Inversion Methods

URS/BURS (Uniform Resampling/Block Uniform Resampling) [78] are methods which perform data resampling by transforming the gridding problem into a linear equation which can be solved using singular value decomposition (SVD). An advantage of this method is that no grid oversampling is employed for the gridding process and no DCF is required. However, this family of methods has several drawbacks. In the URS method, the large number of data samples leads to an inconveniently large linear equation. The BURS method is somewhat more practical, although it is highly sensitive to noise due to the need for a matrix inversion in the SVD. The extension to these approaches, rBURS (regularized Block Uniform Resampling) [79], addresses this noise sensitivity problem, although the results are strongly dependent on the parametrization of the matrix inversion problem, i.e. the regularization and the size of the region of support.

2.8.3 Iterative Methods

Several iterative methods have also been proposed for the data gridding problem. One such approach is INNG (Iterative Nearest Neighbor Gridding) [80], which differs from convolution gridding in that a multiplication in the image domain is substituted for a convolution in k-space. In this way, an effective sinc convolution can be performed in k-space (through the multiplication of a box-car function with the image in the frequency domain) instead of using an approximation (such as a Kaiser-Bessel window). However, a number of iterations of the INNG algorithm are required before an artifact-free image results from the non-Cartesian data. Each iteration involves one Fourier transformation and one inverse Fourier transformation, which becomes quite time-consuming when large oversampling factors are used. A second iterative gridding algorithm is DING (Deconvolution/Interpolation Gridding) [81], which also formulates the gridding problem as a linear equation. Unlike the URS family, this linear equation is solved using a conjugate gradient optimization method which is discussed in Section 3.3.2. Because no grid oversampling is used, this method is less computationally intensive than INNG, although each iteration requires the same number of Fourier transformations. However, the convolution window used in DING must be optimized, again leading to potential difficulties in parametrization.

2.8.4 Non-Uniform Fourier Transform

Another option for the reconstruction of non-Cartesian MRI data is the group of non-uniform Fourier Transform methods. Because the non-uniform Fourier Transform still requires hours of computational time, even with modern computers [82], the focus in MRI reconstructions has been the non-uniform fast Fourier transform (NUFFT). A one-dimensional version of the NUFFT was introduced by Dutt and Rokhlin in 1993 [83], which employs approximation theory to allow the non-uniform case to be reduced to a combination of simple fast Fourier transforms. Because of this approximation, the 1D NUFFT leads to errors in the final result, although the magnitude of these errors can be minimized by selecting appropriate values of several parameters which appear in the approximation. The corresponding two-dimensional NUFFT, otherwise known as the generalized fast Fourier transform (GFFT), was described by Sarty [82]. In his paper on the GFFT, Sarty also noted that this method is equivalent to convolution gridding using a Gaussian kernel, instead of the Kaiser-Bessel kernel generally used in convolution gridding (see Section 2.8.1). Thus, there is no fundamental difference between the use of the non-uniform Fourier transform and a convolution gridding method; both have similar strengths and weaknesses, including the need for additional gridding parameters, the variability of the DCF, and the inability to grid non-Nyquist sampled data. Other methods which attempt to perform the NUFFT without these additional parameters, such as the least-squares NUFFT (LS_NUFFT) [84], have been shown to be less accurate than the Kaiser-Bessel convolution gridding [85]. Thus, although the NUFFT method is available for the gridding of non-Cartesian data, it has been shown to be equivalent to convolution gridding, which is still the gold-standard for data gridding in MRI.

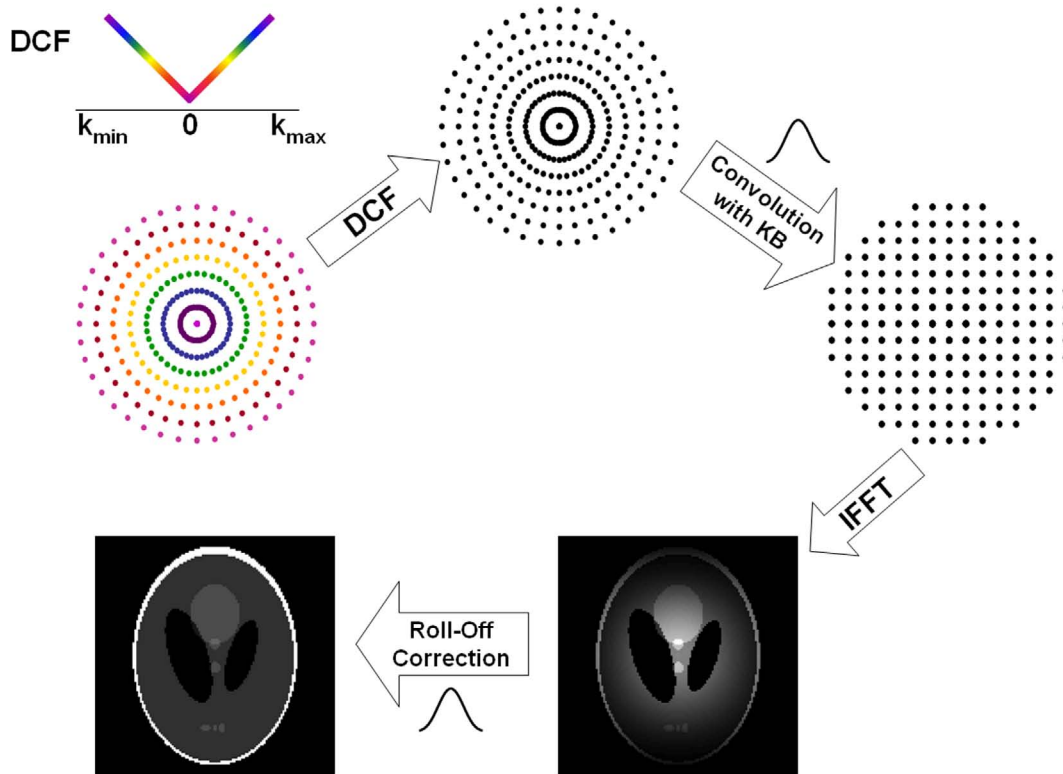


Figure 2.7: A schematic depiction of convolution gridding. Center left: The acquired radial dataset, shown in color to emphasize to the decreasing k-space sampling density as the distance between sampled points in the center of k-space increases, must be multiplied by a DCF, as explained in the text. Top: The result is a radial dataset which has a constant sampling density. Center right: These points are then multiplied with a convolution kernel and resampled at the appropriate Cartesian points. This sampling grid is generally chosen to be finer than the k-space grid which corresponds to the actual FoV in order to increase the gridding accuracy. Bottom right: An inverse FFT is performed to move to the image domain. Bottom left: The roll-off correction corresponding to the inverse of the convolution kernel is applied, yielding the final gridded image. If a grid oversampling factor was employed, the actual FoV can be cropped out of the larger image at this stage.

Chapter 3

Parallel Imaging

3.1 Basics of Parallel Imaging

3.1.1 Historical Overview

Due to the constraints imposed on MRI discussed in the previous chapter, increasing imaging speed while maintaining the same contrast and resolution seemed impossible in the 1990's. However, at the same time, parallel coil array technology matured, offering the possibility for much larger coil arrays. While single volumetric coils were generally used in the last century, standard clinical arrays currently in use are made up of between four and twelve elements. This trend has continued with the development of arrays with 23 [86], 32 [87], or even 90 [86] independent receiver coils, and these new arrays are poised to become common in a clinical setting. An example of data acquired using a phased array with eight elements is shown in Figure 3.1. While these arrays can increase the SNR at the surface of the object due to their higher sensitivities in these areas, they also allow the localization of the signal in a way that was not possible with standard volumetric coils. The use of this new source of spatial information from independent coils to decrease scan time in MRI is known as parallel imaging (or pMRI).

In parallel imaging, the amount of k-space data acquired is reduced, as discussed in Section 2.6. If the distance between each line of k-space in a Cartesian acquisition is increased while the extent of the k-space (i.e. the resolution) is maintained, the total number of lines needed to fill the k-space is reduced, but the Field-of-View (FoV) of the image is also reduced. Given that only half the k-space lines are acquired, the FoV of the resulting image will also be halved, resulting in an image with aliasing artifacts. This can be understood by examining the Nyquist criterion; if the sampling rate in one direction of the k-space is reduced, certain frequencies will be indistinguishable from one another, leading to the aliasing present in Figure 3.2. At this point, the goal of shorter imaging time has been achieved, although the resulting image cannot be used because features cannot be easily distinguished due to the artifacts. Thus, a parallel imaging reconstruction algorithm is required in order to generate an unaliased image.

The concept of using spatial information from coil arrays to replace some or all gradient encoding was introduced more than 25 years ago [88, 89, 90]. However, due to the lack of availability of either




Figure 3.1: An eight-channel receiver array, and representative single channel images. Each channel has its own receiver chain and acquires data simultaneously and independently of the other channels (as indicated by the arrows in the image). The different sensitivity profiles of the channels lead to different views of the object to be imaged, and used together, the object is covered completely. The additional spatial information inherent in such an array can be used to decrease the imaging time using parallel imaging methods.

receiver array coils or multi-channel MR systems systems, the acceleration of in vivo images could not be accomplished robustly with these early algorithms. However, in 1997, Sodickson et al. introduced a method known as SMASH [1]. This reconstruction algorithm uses linear combinations of coil sensitivity variations to mimic the spatial harmonics normally generated by the phase encoding gradients. At the same time, multi-channel receiver coils and systems were also becoming more commonplace and robust, which allowed new parallel imaging techniques to be applied reliably in vivo. Thus, with the arrival of SMASH and these array coils, many other parallel imaging methods followed, such as AUTO-SMASH [2], SENSE [3], SPACE RIP [4], PILS [5], VD-AUTO-SMASH [6], and GRAPPA [7].

In general, there are two classes of methods for reconstructing the unaliased image. The first class operates in the image domain, and employs coil sensitivity maps to “unfold” the aliased image. The most commonly used algorithm in this class is generalized SENSE [3]. Although this method is optimal in cases where the coil sensitivity map can be determined, it can only be used when the undersampling is regular, i.e. Cartesian sampling. The second class of algorithms seeks to reconstruct the missing k-space data using coil sensitivity variations in place of gradient encoding. Once the data in k-space have been reconstructed, a Fourier Transform can be performed on the original and reconstructed data to arrive at the unaliased image. The GRAPPA [7] method, which is based on SMASH [1] and was introduced in Würzburg in 2002, is one of the most flexible and often employed k-space reconstruction algorithms. The work in this thesis has been based primarily on GRAPPA and its extensions, and for this reason, GRAPPA is discussed in detail in later sections.




Figure 3.2: Top left: An example of a $R=4$ undersampled k-space, where every fourth line (solid) has been acquired, and all other k-space lines (dotted) have not been measured. Top right: The corresponding $R=4$ image containing aliasing artifacts, which appear because the Nyquist criterion has not been fulfilled. Note that because the undersampling in k-space is in a single direction (the phase encoding direction), the fold-over artifacts appear exclusively in this direction. Bottom left: The parallel imaging method GRAPPA works by reconstructing missing k-space lines using the coil sensitivity variations to mimic gradient encoding. Once the k-space has been reconstructed, an IFFT is performed to arrive at the reconstructed image. Bottom right: In SENSE, on the other hand, the image reconstruction is performed in the image domain using a coil sensitivity map.

3.1.2 General Remarks on Parallel Imaging

Before discussing specific reconstruction algorithms, a few general remarks about parallel imaging must be made (see [91, 92] for more detailed comments on parallel imaging). The following points are important to keep in mind:

- Parallel imaging algorithms are not new imaging sequences, but merely techniques for the reconstruction of undersampled data
- Parallel imaging requires data acquired simultaneously and independently with multiple receiver channels with differing sensitivity profiles
- The acceleration factor, R , cannot exceed the number of receiver channels used to collect the undersampled data
- Most imaging experiments can be accelerated using parallel imaging

- Parallel imaging generally does not change the contrast behavior of the underlying imaging sequence
- Parallel imaging can be used to decrease the amount of time needed to perform an imaging sequence, increase the resolution given a specific measurement time, or perform a combination of the two
- The SNR in the reconstructed image is generally reduced by a factor greater than or equal to the square root of the reduction factor, as R -times fewer datapoints are acquired (see Equation 2.24)

As a convention, the amount of data acquired as compared to the amount needed for a fully-sampled image is described as the acceleration factor, R . For example, if 256 k-space lines are required for a given FoV and resolution, and only 64 are actually acquired, the acceleration factor of the dataset is said to be $R=4$, which is depicted in Figure 3.2. This acceleration occurs in the phase encoding direction, as the total time of an MR experiment is dependent on the number of phase encoding steps used. Thus, for 2D Cartesian imaging, undersampling is performed in a single direction; for 3D Cartesian imaging, two acceleration directions are available. Acceleration factors of between 2 and 4 are currently used in a clinical setting, although factors between 9 and 12 are possible for 3D imaging in high SNR applications using appropriate receiver arrays [91].

The acceleration factor is limited by the sensitivity variations inherent in the receiver coil used to make the measurement. At the absolute limit, the acceleration factor cannot be larger than the number of coils employed, although most array configurations cannot accomplish such high acceleration factors due to the distribution of the sensitivity variations in two (or even three) spatial directions. However, the trend towards greater numbers of elements in coil arrays is influenced by the fact that more independent coils lead to higher possible parallel imaging acceleration factors.

As stated in the bullet points above, parallel imaging generally does not change the contrast in the accelerated image. This does not hold true in a number of special cases, most importantly for multi-echo or single shot acquisitions. In these cases, the use of parallel imaging can improve the SNR and decrease susceptibility artifacts, as well as change the contrast, due to shorter echo train lengths. Such effects are more common and relevant in non-Cartesian parallel imaging, and are discussed specifically when they apply.

Finally, the optimal method for the combination of multi-channel data to form a single image has been shown to be sum-of-squares (SoS) [93]. While this combination method yields a higher signal value in those areas where the coil sensitivity is higher, the SNR is maximized in all portions of the image. An additional disadvantage of the SoS method is the loss of phase information which is present in the single channel images. Other methods for combining multi-channel images to create a single image include adaptive reconstruction [94], which uses correlation analysis to compute the coil sensitivities. With this information, an image with the desired characteristics can be produced, i.e. a constant signal level over the entire image, constant noise level, or optimal SNR. In addition, phase information in the image domain is retained. Due to these advantages, the adaptive reconstruction method is often employed despite its relative complexity as compared to the SoS method.

3.1.3 SNR Loss and Reconstruction Errors in Parallel Imaging

The acceleration factor describes not only the time savings, but is also an indication of the minimum SNR loss for an image. Because R -times fewer points are acquired in an accelerated scan, the SNR of the reconstructed image will be at least \sqrt{R} times lower than that of the unaccelerated image. An additional factor which describes the encoding efficiency of the receiver coil also influences the SNR. This parameter, known as the g -factor, is a measure of how easily the inversion of the coil sensitivity matrix required for parallel imaging (see for instance Equations 3.9 and 3.23) can be performed. Given completely independent coil sensitivity profiles, the g -factor is equal to one; if the sensitivity of each coil is identical, the g -factor is equal to infinity, as an inversion of the matrix made up of the sensitivity values would be impossible.

The SNR for an image reconstructed using parallel imaging can then be described by the following equation:

$$SNR_{pMRI} = \frac{SNR}{g \cdot \sqrt{R}} \quad (3.1)$$

Thus, parallel imaging can only be implemented when the resulting loss in SNR does not lead to unacceptably degraded image quality. Due to the inclusion of the g -factor, the actual SNR loss is determined by the coil array used to acquire the data. Arrays with significant sensitivity variations allow for higher acceleration factors, as parallel imaging employs these sensitivity variations as additional agents of spatial encoding. As the acceleration factor increases, either the loss in SNR in the resulting image also increases, or residual aliasing artifacts remain.

A commonly employed measure of the quality of the reconstructed image is the root-mean-square error (RMSE) of the new image as compared to an unaccelerated image:

$$RMSE = \sqrt{\frac{\sum (x_i - x_{i,ref})^2}{\sum x_{i,ref}^2}} \quad (3.2)$$

where x and x_{ref} denote the pixel values in the reconstructed and reference images, respectively. The RMSE takes all deviations of the reconstructed image from the reference image, whether they are due to noise enhancement, artifacts, loss of resolution, or intensity changes, into account, and thus must be used with caution.

3.2 Cartesian Parallel Imaging

Although the work in this thesis focuses on non-Cartesian parallel imaging, certain Cartesian parallel imaging concepts are essential to understanding their non-Cartesian counterparts. While there are a number of reconstruction algorithms that can be used for Cartesian parallel imaging, as seen in Section 3.1.1, it is beyond the scope of this thesis to describe each method in detail. Thus, this section briefly outlines the basics of the Cartesian parallel imaging methods used in this work, as well as several other important algorithms.

3.2.1 Generalized Autocalibrating Partially Parallel Imaging (GRAPPA)

The most commonly used and robust k-space based parallel imaging algorithm is GRAPPA [7]. This algorithm is the result of the natural development from SMASH [1] to AUTO-SMASH [2], which in turn evolved into GRAPPA. Because GRAPPA can be seen as a generalization of these earlier methods, only GRAPPA will be discussed in this section.

The basic idea of GRAPPA is that coil sensitivity variations can be used to generate missing spatial harmonics in the undersampled k-space data. This can be understood by examining the 1D Fourier Transform:

$$S(k_y) = \int dy \rho(y) \cdot e^{ik_y y} \quad (3.3)$$

In standard imaging experiments, the spin density $\rho(y)$ is modulated by the appropriate harmonic, $e^{ik_y y}$, through the application of an encoding gradient, as described in Section 2.3. The integral, or in MRI, the sum, over these products yields the signal $S(k_y)$ at that particular location in k-space. Including the influence of the coil sensitivity leads to the following equation for the signal from a coil L :

$$S_L(k_y) = \int dy \rho(y) \cdot C_L(y) \cdot e^{ik_y y}$$

In parallel imaging, when the data are undersampled to reduce scan time, some of these k-space lines are skipped. By examining the encoding equations for two lines separated by a distance m in k-space, one can see that the only difference between the two is the spatial harmonic applied, i.e. for the signal at a k-space location $k_y + m\Delta k_y$:

$$S_L(k_y + m\Delta k_y) = \int dy \rho(y) \cdot C_L(y) \cdot e^{i(k_y + m\Delta k_y)y} \quad (3.4)$$

Thus, if one can recreate the additional spatial harmonic of the missing k-space line, or $e^{im\Delta k_y y}$, the missing lines can be reconstructed. In GRAPPA, this reconstruction involves using a linear combination of the coil sensitivity variations of all the coils in the array to mimic the gradient encoding:

$$C_L(y) \cdot e^{im\Delta k_y y} \approx \sum_{K=1}^{NC} n_{K,L,m} \cdot C_K(y) \quad (3.5)$$

where both K and L run from 1 to the number of coils in the array, or NC , and for a complete reconstruction of the missing k-space lines, m runs from 1 to $R - 1$. By substituting Equation 3.5 into Equation 3.4, the signal from an acquired line of k-space can be used to generate an approximation of the signal from a missing line:

$$S_L(k_y + m\Delta k_y) \approx \sum_{K=1}^{NC} n_{K,L,m} \cdot S_K(k_y) \quad (3.6)$$

or, in matrix form:

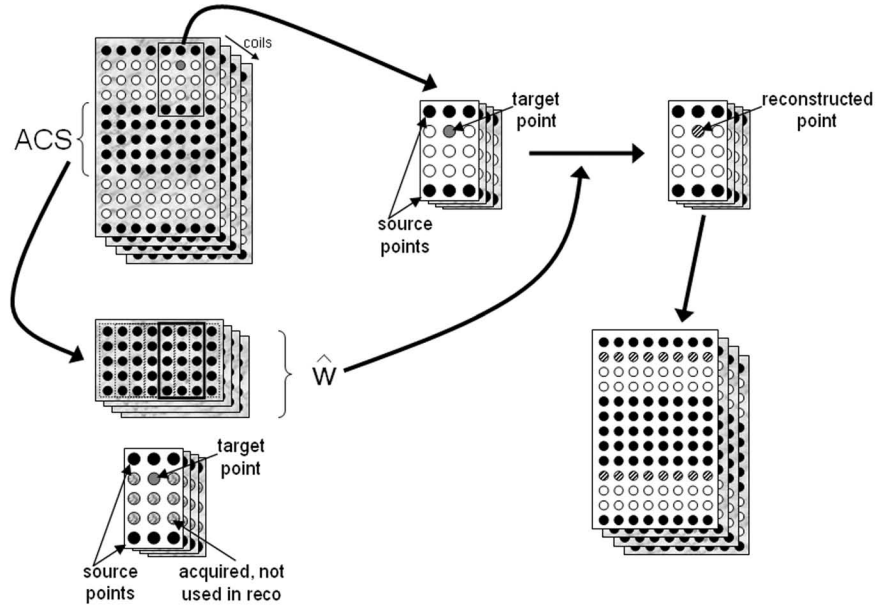


Figure 3.3: A schematic depiction of GRAPPA. Top left: The undersampled $R=4$ multi-channel k-space data, where the black points have been acquired, and the white points have not. The auto-calibration signal (ACS), which is required for the calibration of the GRAPPA weights, can be seen as the extra acquired data in the center of k-space. Bottom left: The first step in GRAPPA is to determine the weights using the ACS data. A 2×3 kernel is shown here, where six source points (two in the phase encoding direction and three in the read direction) from all the coils (shown in black) are used to fit the target point (solid gray) in one coil (the first, shown in white), and thus the weight set needed has a size of 4×6 (coils \times source points). This kernel appears multiple times in the ACS data (the dotted lines around the groups of points), and all repetitions are used to determine the the weight set \hat{w} . Top center: The same kernel appears in the undersampled portion of k-space, and when the weights \hat{w} are applied to the source points (top right), the target point (black stripes) can be reconstructed. Bottom right: The same weight set can be used to reconstruct all of the missing points in the first coil with the same relationships to the acquired points. In order to reconstruct the entire missing k-space, weight sets for each missing spatial harmonic and coil must be calculated, which means that 3 (for the three missing spatial harmonics) different $4 \times 4 \times 6$ weight sets are needed.

$$\vec{S}(k_y + m\Delta k_y) \approx \hat{w}_m \cdot \vec{S}(k_y) \quad (3.7)$$

Equation 3.6 is simply the matrix version of Equation 3.7, and both state that a linear combination of acquired signals from all receiver coils yields an approximation of the missing signal for a single coil. The elements of the matrix \hat{w} , referred to as the GRAPPA weight set, are the values of n for each of the coil combinations, and the subscript m remains to denote that a different weight set is required for each missing spatial harmonic, i.e. each of the missing k-space lines.

The equations above describe the simple case of employing one source point, or acquired point, in each coil to reconstruct one target point, or missing point, in a single coil. However, GRAPPA generally employs a reconstruction kernel made up of more than one source point for the approximation of each target point, as shown in Figure 3.3. In this schematic, two source points in the phase encoding direction and three in the read direction (a 2×3 kernel) from all coils are used to fit a single target point in a single coil. Because a coil-by-coil reconstruction is performed in GRAPPA, the total size of each weight set is $(NC \times NC \times \text{source points})$. Thus, for a $R=4$ accelerated dataset acquired with a 4 channel coil array where a 2×3 kernel is used for GRAPPA, $R - 1$ or 3 different $(4 \times 4 \times 6)$ weight sets would be required for the reconstruction of the entire k-space. The use of such a reconstruction kernel makes GRAPPA more robust than methods that use only a single source point to reconstruct a single target points, as the extra degrees of freedom allow the missing spatial harmonics to be better mimicked.

It is important to note that the difference in spatial harmonics between two adjacent k-space lines is the same in all portions of k-space. Thus, once the weight set for a given pattern (for instance, the 2×3 kernel and $R=4$ shown in Figure 3.3) is determined, it can be applied to reconstruct missing points in all portions of k-space which have the same pattern of acquired/missing points. This idea can also be seen in the figure, which shows that missing k-space points both above and below the ACS dataset can be reconstructed (the black striped points) using the appropriate source points and a single weight set, \hat{w} .

Naturally, before GRAPPA can be employed for a reconstruction, the weight sets must be determined. This can be accomplished, for example, by acquiring a few additional lines in the center of k-space (where the signal level is the highest), as shown on the upper left-hand side of Figure 6.1. These auto-calibration signal lines, or ACS, can be used in conjunction with Equation 3.7 to generate the values of the weight set:

$$\vec{S}_{ACS}(k_y + m\Delta k_y) = \hat{w}_m \cdot \vec{S}_{ACS}(k_y) \quad (3.8)$$

This GRAPPA kernel (and therefore the harmonic relationship) appears many times in the auto-calibration signal, although the weight set for each equation is the same throughout the entire k-space (as shown in Equation 3.5). Thus, the many appropriate signal and target vectors can be written together in the form of a matrix. The weight sets can then be determined with the help of the pseudo-inverse, or Moore-Penrose inverse, which allows the non-square matrix made up of the source signals to be inverted and results in the least-squares solution to this linear problem:

$$\hat{S}_{ACS}(k_y + m\Delta k_y) \cdot \text{pinv}(\hat{S}_{ACS}(k_y)) = \hat{w}_m \quad (3.9)$$

In this equation, the pseudo-inverse is denoted as “pinv”, and is equivalent to the following:

$$\text{pinv}(\hat{A}) = (\hat{A}^H \hat{A})^{-1} \hat{A}^H \quad (3.10)$$

where \hat{A}^H is the conjugate transpose (otherwise known as the Hermetian conjugate) of the matrix \hat{A} . Once the weight sets have been determined, they can be used to reconstruct the portions of the dataset where data are missing in order to arrive at the final image.

3.2.2 GRAPPA Operator

The GRAPPA operator formalism [95] is an extension of GRAPPA which describes the special case of a square GRAPPA weight set, or in other words, an equal number of source and target points. For simplicity, the case where one source and one target point are employed will be examined (the so-called one-to-one operator), although the same formalism can be used to describe larger GRAPPA operators (two source points and two target points, etc.), and all the properties described here are valid for such operators. Both GRAPPA and the GRAPPA operator can be used to generate a missing point with a specific spatial harmonic using a linear combination of acquired points, and thus the mathematics (i.e. calibration and application) for the GRAPPA operator are identical to GRAPPA. The basic GRAPPA operator equation for a shift of Δk with a single source and target point is identical to Equation 3.7:

$$\vec{S}(k_x, k_y + \Delta k_y) = \hat{G}_y \cdot \vec{S}(k_x, k_y) \quad (3.11)$$

In this equation, \hat{G}_y is a $NC \times NC$ matrix containing the GRAPPA weights for a shift in the y-direction, and will be further referred to as the GRAPPA operator, and \vec{S} is a vector containing the k-space points at the proper location for each of the receiver coils. The GRAPPA operator \hat{G} can be derived in the same fashion as standard GRAPPA weights; namely, a fit of points with the appropriate relationship is performed with the use of the pseudo-inverse:

$$\hat{S}(k_x, k_y + \Delta k_y) \cdot \text{pinv}(\hat{S}(k_x, k_y)) = \hat{G}_y \quad (3.12)$$

While the basic mathematics are identical to those of standard GRAPPA, the GRAPPA operator has additional properties which makes this concept interesting for parallel imaging. The defining feature of the GRAPPA operator (and that which distinguishes it from GRAPPA) is that if \hat{G} is applied twice successively to the signal at k_y , the result is the signal at $k_y + 2\Delta k_y$:

$$\vec{S}(k_x, k_y + 2\Delta k_y) = \hat{G}_y \cdot \vec{S}(k_x, k_y + \Delta k_y) = \hat{G}_y \cdot \hat{G}_y \cdot \vec{S}(k_x, k_y) = \hat{G}_y^2 \cdot \vec{S}(k_x, k_y) \quad (3.13)$$

In other words, the application of the square of the GRAPPA operator leads to a shift twice the size of the shift from the original operator. The name “GRAPPA operator” derives from this feature, which it shares with ladder or propagator operators in quantum mechanics. In principle, the same is true in

reverse; applying the square root of the operator leads to a k-space shift of half the size [96]:

$$\vec{S}(k_x, k_y + \frac{1}{2}\Delta k_y) = \hat{G}_y^{\frac{1}{2}} \cdot \vec{S}(k_x, k_y) \quad (3.14)$$

Thus, an arbitrarily small shift can be derived by taking the proper root of the base GRAPPA Operator, \hat{G}_y :

$$\vec{S}(k_x, k_y + \delta k_y) = \hat{G}_y^\delta \cdot \vec{S}(k_x, k_y) \quad (3.15)$$

It is important to note that the root operation for a matrix \hat{A} can be defined as follows:

$$\hat{A} = \hat{V} \cdot \hat{D} \cdot \hat{V}^{-1} \quad (3.16)$$

$$\hat{A}^\delta = \hat{V} \cdot \hat{D}^\delta \cdot \hat{V}^{-1} \quad (3.17)$$

where \hat{V} is a matrix containing the eigenvectors of the matrix \hat{A} , and \hat{D} is a diagonal matrix of the eigenvalues of \hat{A} .

Another property of the GRAPPA operator is that a shift in one direction followed by a shift in the opposite direction must yield the original point:

$$\vec{S}(k_x, k_y) = \hat{G}_y^{-m} \cdot \vec{S}(k_x, k_y + m\Delta k_y) = \hat{G}_y^{-m} \cdot \hat{G}_y^m \cdot \vec{S}(k_x, k_y) \quad (3.18)$$

and thus

$$\hat{G}_y^{-m} \cdot \hat{G}_y^m = \hat{I} \quad (3.19)$$

or

$$\hat{G}_y^{-m} = (\hat{G}_y^m)^{-1} \quad (3.20)$$

This equation implies that the operator for a shift in the $-k_y$ direction is simply the inverse of the operator for a shift in the k_y direction. Finally, GRAPPA operators for shifts in orthogonal directions can be consecutively applied to generate a multi-dimensional shift:

$$\vec{S}(k_x + m\Delta k_x, k_y + n\Delta k_y, k_z + p\Delta k_z) = \hat{G}_{mx,ny,pz} \cdot \vec{S}(k_x, k_y, k_z) = \hat{G}_x^m \cdot \hat{G}_y^n \cdot \hat{G}_z^p \cdot \vec{S}(k_x, k_y, k_z) \quad (3.21)$$

Thus, using the GRAPPA operator formalism, shifts in any direction by any amount in k-space can be calculated out of a set of orthogonal unit shift operators. In theory, these orthogonal operators always commute, as the final result does not depend on the order in which the shifts were applied:

$$\vec{S}(k_x + m\Delta k_x, k_y + n\Delta k_y) = \hat{G}_x^m \cdot \hat{G}_y^n \cdot \vec{S}(k_x, k_y) = \hat{G}_y^n \cdot \hat{G}_x^m \cdot \vec{S}(k_x, k_y) \quad (3.22)$$

It is important to keep in mind that these shift matrices are always GRAPPA weight sets, and that the accuracy of the spatial harmonic depends on the coil encoding characteristics; the more varied the coil sensitivities, the larger the possible shift in k-space.

The following bullet points outline the basic properties of the GRAPPA operator:

- An operator can be applied multiple times to generate increasingly larger shifts in k-space (Equation 3.13)
- Operators for smaller shifts can be generated using operators from larger shifts (Equation 3.16)
- The operator for a shift in one direction is equal to the inverse of the operator for a shift in the opposite direction (Equation 3.20)
- The operator for a shift in an arbitrary direction can be separated into operators for shifts in orthogonal directions (Equation 3.21)
- Operators for orthogonal directions commute (Equation 3.22)

It was hoped that the GRAPPA operator could be used to generate GRAPPA weight sets without the need for an ACS dataset. In [96], the authors attempted to do this by using R=2 undersampled data to generate GRAPPA operator weight sets for a $\Delta k=2$ shift in k-space, and then taking the square root of this weight set to arrive at the $\Delta k = 1$ weight set using Equation 3.14. However, this approach was not generally feasible because the proper sign for the square roots of the eigenvalues, needed as shown in 3.16, had to be determined empirically. For R=2, 2^{NC} signs must be determined, and this value increases rapidly as the acceleration factor increases. This difficulty generating GRAPPA operator weight sets for smaller shifts using non-Nyquist sampled data led to the discontinuation of research in this area.

However, the GRAPPA operator has been shown to be useful in increasing the speed for 2D and 3D image reconstructions, although operators with more than one source and target point were employed for these algorithms [97]. In addition, this formalism has been used to increase the resolution of PROPELLER images by increasing the extent of the k-space [98]. Although the GRAPPA operator has been demonstrated to have properties which could potentially improve and/or simplify image reconstruction, it is not a method which is frequently employed (see Section 3.2.3). However, some novel uses of the GRAPPA operator, specifically for data gridding, are discussed in Chapters 4 and 5.

3.2.3 GRAPPA vs. the GRAPPA Operator

In the previous sections, two k-space based parallel imaging methods, namely GRAPPA and the GRAPPA operator, were examined. As stated above, GRAPPA operator is a special case of GRAPPA where weight sets are square matrices; when one source point is used to fit a single target point, the appropriate GRAPPA operator \hat{G} is then a matrix of size $NC \times NC$. GRAPPA, on the other hand, uses multiple source points to fit one target point. These two methods, while similar, are used in applications which take advantage of their different properties.

In order to understand why GRAPPA and the GRAPPA operator are used for different purposes, one must examine how GRAPPA weight sets behave with increasing numbers of source points. When a large reconstruction kernel is used, the fits of the spatial harmonics needed for data reconstruction

(see Equations 3.4 and 3.5) are more accurate for a given coil arrangement than for a small kernel. Thus, when fitting large shifts in k-space, it is beneficial to employ many source points. In addition, when the value of the target point is based on many source points, the reconstructed point is less sensitive to the noise in the individual source points, and therefore more accurate. However, because reconstruction times increase with increasing numbers of source points, a balance must be found between accuracy and reconstruction time. In standard GRAPPA implementations, kernels of size 4×5 or 2×3 are employed; larger kernels generally do not significantly improve the reconstruction quality. Thus, because the weight sets in GRAPPA are not limited in terms of the numbers of source points, GRAPPA is preferred over the GRAPPA operator for clinical image reconstruction.

For the special case of the GRAPPA operator that employs only one source point per target point, the weight sets generated cannot fit higher spatial harmonics accurately. This is because a small number of parameters is available for the fit; only 64 different values can be chosen for the one-to-one GRAPPA Operator with 8 receiver channels (the weight set is 8×8), whereas 384 different values can be chosen to fit the desired spatial harmonic in the case of a 2×3 GRAPPA kernel (the weight set is $8 \times 8 \times 6$). Such a GRAPPA operator is therefore only appropriate for small shifts in k-space, which correspond to low frequency and easy-to-fit harmonics. Similarly, the reconstructed points are highly sensitive to any inaccuracies due to noise in the source points. Thus, GRAPPA operator reconstructions tend not to be as robust as GRAPPA reconstructions. However, the one-to-one GRAPPA operator has other properties which make it useful (see the bullet points in Section 3.2.2). While GRAPPA weight sets require a specific spatial relationship between the source points and the target points, the GRAPPA operator can be applied to any single point to shift it to a nearby location, because only one source point is involved. As stated in Section 3.2.2, the operator can be applied multiple times to generate many k-space points. The weight set can be scaled to fit different spatial harmonics (see Equation 3.16), and thus the weight sets for different shift sizes can be generated out of a single basis weight set. These properties of the one-to-one GRAPPA operator, which are not shared by GRAPPA, allow this method to be used for the exotic applications discussed in later chapters of this thesis.

It must be noted that larger GRAPPA operators can be defined which fit two source points to two target points, etc. These larger operators can also be scaled by taking the appropriate roots, and have the advantage that they can fit larger spatial harmonics more accurately. However, such GRAPPA operators can only be employed when the relationship between the source and target points has the correct pattern, as in GRAPPA. This limitation makes such GRAPPA operators less flexible than the one-to-one operator discussed above, and only the one-to-one case is examined in later chapters of this work.

3.2.4 Other Parallel Imaging Methods

Although GRAPPA and the GRAPPA operator are the two methods which are discussed and employed the most throughout this work, there are other Cartesian parallel imaging methods, for example PILS [5] and SENSE [3], which should be mentioned in any treatment of parallel imaging. Thus, a brief description of these methods will be provided before moving on to non-Cartesian parallel imaging.




Figure 3.4: A schematic description of PILS. Top: When using a coil with a localized and homogeneous sensitivity profile, only a small portion of the object to be imaged appears in the single coil image. When this single coil image is undersampled by a factor of 4, the outer portions of the image fold into the center portion, as in Figure 3.2. However, because these outer portions contain no signal due to the lack of coil sensitivity in this area, only noise folds into the visible portion of the image. Bottom: $R=4$ single coil images for four such localized coils, each of which sees a unique portion of the object. These images can simply be combined with, for example, the sum-of-squares (SoS) algorithm to arrive at the final image. Note that PILS only works so easily when the receiver coils used “see” completely different portions of the image, i.e. do not overlap, and have relatively homogeneous sensitivity profiles.

The most basic parallel imaging method is PILS (Parallel Imaging with Localized Sensitivities), which does not require any advanced data processing. In order to perform PILS, localized coil sensitivities are required, as shown in the top center of Figure 3.4. Given that each coil “sees” a unique portion of the image, and that the FoV is completely covered by these coils, the effective FoV for the imaging process can be reduced. As in any imaging sequence, the reduction of the FoV will lead to fold-over artifacts, but these are weighted by the coil sensitivity, which is zero outside of the sensitive area. Thus, only noise folds into the reduced FoV, and a simple sum-of-squares reconstruction of the individual coil images can be performed. This process is described schematically in Figure 3.4.




Figure 3.5: A schematic depiction of the components of the SENSE Equation 3.23. The object, and four specific voxels ρ in locations 1 through 4 are shown at the left. When a $R=4$ undersampled image of this object is acquired with coil A, whose sensitivity map is shown in the center, the result is a folded image, where the pixels which fall on top of one another are simply the pixels 1 through 4 multiplied by the coil sensitivities C_{A1} through C_{A4} . When four receiver coils, with sensitivity maps C_A through C_D , are used for the acquisition, four folded images with different coil sensitivity information result, and the spin density can be calculated using the inverse the matrix \hat{C} .

SENSE (SENSitivity Encoding) is the most common image-based parallel imaging method, and it functions by using coil sensitivity maps to separate aliased pixels. By examining Figure 3.5, it is clear that the following equation can be written to describe the relationship between the actual image and the aliased image:

$$\begin{bmatrix} I_A \\ I_B \\ I_C \\ I_D \end{bmatrix} = \begin{bmatrix} C_{A1} & C_{A2} & C_{A3} & C_{A4} \\ C_{B1} & C_{B2} & C_{B3} & C_{B4} \\ C_{C1} & C_{C2} & C_{C3} & C_{C4} \\ C_{D1} & C_{D2} & C_{D3} & C_{D4} \end{bmatrix} \cdot \begin{bmatrix} \rho_1 \\ \rho_2 \\ \rho_3 \\ \rho_4 \end{bmatrix} \quad (3.23)$$

This equation can be written more efficiently using a matrix formulation:

$$\vec{I} = \hat{C} \cdot \vec{\rho} \quad (3.24)$$

In both of these equations, the values in \vec{I} represent the aliased voxels from the coils A through D, $\vec{\rho}$ are the actual voxel values at locations 1 through 4, and the matrix \hat{C} contains the coil sensitivity values for the four different coils at the appropriate locations. If the coil sensitivity matrix is known, then the aliased pixels \vec{I} can be used to determine the actual spin density values $\vec{\rho}$ by taking the inverse of the matrix \hat{C} :

$$\hat{C}^{-1} \cdot \vec{I} = \hat{C}^{-1} \cdot \hat{C} \cdot \vec{\rho} = \vec{\rho} \quad (3.25)$$

If the number of coils is greater than the acceleration factor, which is often the case, the sensitivity matrix \hat{C} is not square, and the Moore-Penrose pseudo-inverse (see Equation 3.10) must be employed in place of the inverse.

SENSE has been shown to deliver the best possible image reconstruction with optimal SNR given an accurate sensitivity map [91], although this constraint often cannot be met. In certain imaging

modalities, i.e. lung imaging, dynamic imaging, or single-shot methods, the acquisition of a sensitivity map is either time consuming or difficult due to SNR limitations. In addition, the SENSE algorithm cannot be used to reconstruct images containing aliasing in the FoV (so-called pre-folded images) [99], which are often acquired in cardiac and abdominal MRI. Thus, although SENSE is a simple and generally fast parallel imaging method, the acceleration of some MR applications is only practical when working with a k-space based method.

3.3 Non-Cartesian Parallel Imaging

In the previous section, Cartesian k-space based parallel imaging methods were explained and discussed. However, these methods can only be applied when the k-space undersampling is regular, i.e. patterns appear in the undersampled k-space data. Such regular undersampling leads to well-defined aliasing characteristics in the image domain. Thus, given an acceleration factor R in a single direction, R voxels must be separated from one another; that this is the case can be seen in Equation 3.23. However, when non-Cartesian k-space data are undersampled, aliasing artifacts appear in all directions, and each voxel in the image domain can potentially alias with all of the other voxels [44, 45]. The complex aliasing behaviors of different non-Cartesian trajectories can be seen in Figure 2.5. Thus, non-Cartesian parallel imaging is considerably more complicated than Cartesian parallel imaging. This section focuses on non-Cartesian parallel imaging methods, especially non-Cartesian GRAPPA.

3.3.1 Non-Cartesian GRAPPA

As can be seen in the basic GRAPPA equation, Equation 3.7, as well as the schematic shown in Figure 3.3, GRAPPA can be applied efficiently only when the undersampling in k-space leaves regular patterns of missing datapoints. If this is not the case, a separate GRAPPA weight set is required for each missing point in k-space, a time-consuming and computationally intensive task. A schematic of the irregular undersampling found in non-Cartesian trajectories, in this case radial, can be seen on the far left hand side of Figure 3.6. It is clear that if the GRAPPA weight set was determined for the specific pattern shown in green, the application of the weight set to the source points shown in blue or red would lead to errors in the reconstructed points. Thus, the regular undersampling patterns in k-space that are crucial for a successful GRAPPA reconstruction are not present in undersampled non-Cartesian datasets, and standard Cartesian GRAPPA cannot be applied to such datasets. However, one can see that in cases where the non-Cartesian trajectory is highly symmetric, similar patterns do exist. Non-Cartesian GRAPPA takes advantage of these similar patterns to reconstruct the missing k-space points.

Radial GRAPPA

The left-hand side of Figure 3.6 shows a schematic of an undersampled radial dataset. As can clearly be seen, different GRAPPA kernels, or patterns, are required to reconstruct different portions of this k-space data. A pattern that would work well in one section of k-space, such as the blue points, would

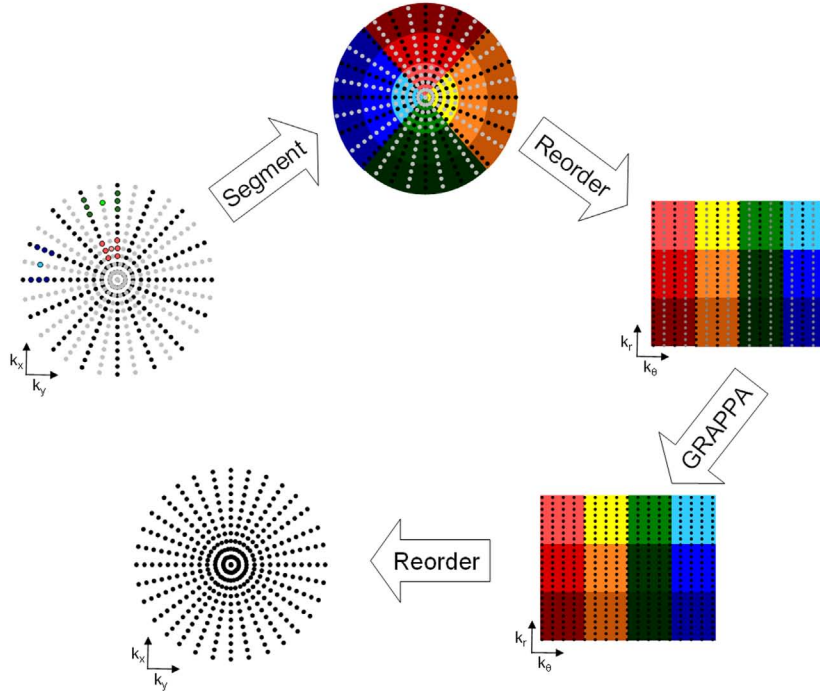


Figure 3.6: The radial GRAPPA algorithm. Left: A $R=2$ undersampled radial dataset, where the black points have been acquired and the gray points have not, and the coil dimension has been neglected for simplicity. One could imagine using the points highlighted in dark green to reconstruct the missing point in light green. However, the weight set for this reconstruction would not be applicable for the red points, where the source and target points are not as far away from each other (the spatial harmonic to be fitted is not as large). Similarly, the weights for the green pattern could not be applied to the blue pattern, because the direction of the undersampling is not the same. Strictly speaking, different weight sets for each missing point must be defined. Top center: To limit the number of weight sets which must be calculated, segments are defined in k -space where the undersampling degree and direction are approximately the same. In this schematic, 12 segments are defined, but more could be used in practice if necessary. Right: The data are reordered from the $k_x - k_y$ plane into the $k_r - k_\theta$ plane, where the segments can be observed to be blocks of k -space. Weight sets are determined for each segment using a reordered fully-sampled dataset. Bottom right: When the weights have been applied to the appropriate segments, a completely reconstructed dataset results. Bottom left: After reordering the reconstructed data back to the $k_x - k_y$ plane, the reconstructed radial dataset can be gridded and Fourier transformed to arrive at the reconstructed image.

be completely inappropriate for other k-space sections, such as the green or the red points. This is because both the direction and the degree of undersampling depends on the location in k-space, unlike for Cartesian acquisitions. Standard radial GRAPPA [100] assumes that in some portions of k-space, the undersampling pattern is similar enough to allow the use of the same GRAPPA weight set. For instance, the red highlighted section at the top of Figure 3.6 shows such a segment; the data points are not parallel, but a weight set that is determined for the average distance and undersampling direction will approximate the points in the segment. Thus, although the source points used with the weight set do not have the proper relationship in k-space with the target point, the approximation allows for the reconstruction of the missing radial rays. Thus, the steps involved in radial GRAPPA are as follows:

- The undersampled k-space data are segmented such that the missing data in each of the segments can be approximated using the same GRAPPA kernel
- The segments are reordered from the $k_x - k_y$ plane onto the $k_r - k_\theta$ plane
- Using weight sets for each segment (see below), each segment is reconstructed separately, yielding the reconstructed radial dataset
- The reconstructed radial dataset are converted to the image domain using the radon transform or using the gridding techniques described in Section 2.8, yielding the final image

These steps are shown schematically in Figure 3.6. Because each segment of the data in the $k_r - k_\theta$ plane appears Cartesian, the application of the weight sets can be performed as in Cartesian GRAPPA.

While radial GRAPPA yields reconstructed images that no longer contain the streaking and blurring artifacts present in the undersampled image, this algorithm has a number of drawbacks. The first is that radial GRAPPA is only an approximation to the actual solution, as the angular weight set which is used in a given segment does not reconstruct the spatial harmonics which are missing. This is due to the fact that the distances between the source and target points change by a small amount even within a single segment. The second drawback is that a fully-sampled radial dataset is required in order to determine the weight sets for each angular segment, as the acceleration factor and direction, and therefore the weight sets, are different for each segment. Thus, the use of radial GRAPPA is generally limited to dynamic studies where a single fully-sampled dataset can be acquired and used for the reconstruction of subsequent undersampled datasets. There are a number of methods which have been proposed to bypass this disadvantage, for instance [101], although none of them is currently in clinical use. In addition, because low SNR areas at the edges of k-space must be used to determine the weights for these portions of the undersampled k-space, these weights are less robust than those determined for the high SNR center portion of k-space. The inaccuracy of these outer k-space segment weights often leads to high-frequency reconstruction errors in the final image. However, despite these drawbacks, radial GRAPPA is one of the fastest and most commonly used non-Cartesian parallel imaging methods.

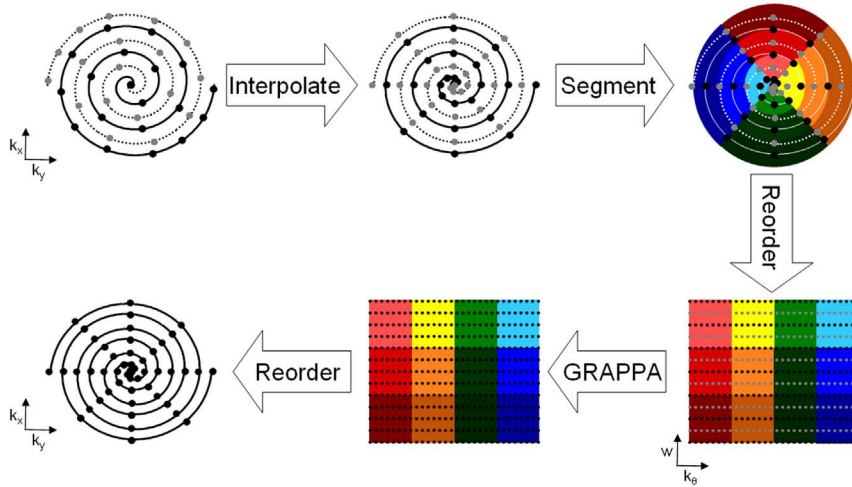


Figure 3.7: The spiral GRAPPA algorithm. Top left: A $R=2$ undersampled spiral dataset, where the data have been acquired using the constant-linear-velocity acquisition scheme. Top center: The spiral data must first be interpolated to the constant-angular-velocity trajectory using a 1D interpolation in the read direction. This trajectory has nearly the same symmetry as the radial trajectory, and the GRAPPA reconstruction after the interpolation is almost identical to radial GRAPPA. Top right: The data are segmented, and (bottom right) reordered into the $w - k_\theta$ plane, where the weights for each segment can be calculated (again using a fully-sampled spiral dataset) and applied. Bottom center: The application of the weights for each segment results in a completely reconstructed dataset, which can be reordered into the $k_x - k_y$ plane (bottom left), and gridded.

Spiral GRAPPA, Zig-Zag GRAPPA, and 1D Non-Cartesian GRAPPA

Although the undersampling present when working with the spiral trajectory is in a direction perpendicular to the undersampling in the radial trajectory, the underlying symmetry of these trajectories is similar. Thus, for spiral GRAPPA [102], the same concept of segmentation, reordering, reconstruction, and gridding can be employed. Because only the choice of the reconstruction kernel changes between the two methods, spiral GRAPPA will not be further discussed here. Figure 3.7 depicts the steps of spiral GRAPPA, which are essentially identical to those of radial GRAPPA, except for the additional interpolation step needed to transform the spiral data into quasi-radial data.

Similarly, zig-zag GRAPPA [33] also works by segmenting the undersampled data into groups for which the same reconstruction kernel can be employed. Again, because this method is similar in concept to radial GRAPPA, it is not discussed in further detail.

1D non-Cartesian GRAPPA [103] is slightly different from radial, spiral, or zig-zag GRAPPA. The 1D non-Cartesian acquisition scheme is shown in Figure 2.5, and one can see that a segmentation of the data would require many different GRAPPA weight sets which would have to be determined from the Cartesian ACS portion in the center of k -space. Instead of determining a weight set for each local acceleration factor, standard GRAPPA weight sets are determined for several acceleration factors

($R=2$, $R=3$, etc). Then, in order to reconstruct the missing lines in k-space, the GRAPPA weights for the actual local acceleration factors ($R=2.5$, for instance) are interpolated from the weight sets which have been actually calculated. This interpolation means that far fewer weight sets must be calculated, thereby speeding up the reconstruction process. However, this interpolation method only works when the relationships between the source and target points, as well as the between source points themselves, change by a constant factor. In other words, only changes in the relative distances are allowed, making this method only applicable to 1D non-Cartesian trajectories. For radial and spiral trajectories, the interpolation of weights as in 1D non-Cartesian GRAPPA would not lead to accurate reconstructions due to the significant changes in the undersampling factor and direction.

Although the non-Cartesian methods described in this section are all similar, the fact that different formulations of GRAPPA must be used for each different trajectory highlights the major problem with non-Cartesian GRAPPA, namely that no generalized GRAPPA algorithm has been defined which allows the reconstruction of arbitrarily undersampled datasets. This problem of non-Cartesian GRAPPA will be addressed in great detail in Chapter 6, where a generalized GRAPPA method which works with any sampling scheme is proposed.

3.3.2 Other Non-Cartesian Methods

In addition to the non-Cartesian GRAPPA methods discussed above, there are a plethora of other methods which can be used for the reconstruction of undersampled non-Cartesian data. Two of these methods, CG-SENSE [104] and PARS [105, 106], are included to offer the reader insight into the current status of non-Cartesian parallel imaging.

Conjugate-Gradient SENSE

Conjugate-Gradient SENSE (CG-SENSE) [104] (similar to [107]) was one of the first parallel imaging methods to be described for non-Cartesian image reconstruction. Although it shares a name with Cartesian SENSE, the principles of CG-SENSE are different than its Cartesian counterpart. As stated above, pixels from all over the FoV can be aliased in an undersampled non-Cartesian image, which makes the direct solution of the SENSE equation nearly impossible for non-Cartesian data because of the sizes of the matrices involved. Such a solution would be extremely computationally intensive due to the need to invert the sensitivity matrix describing the aliasing of all pixels with each other. The CG-SENSE method instead relies on the fact that the multi-channel data, combined with information about the coil sensitivities, are redundant, even when the k-space data themselves are undersampled. The relationship between the image and the acquired non-Cartesian k-space data can be written as a matrix equation:

$$\hat{E} \cdot \vec{v} = \vec{m} \quad (3.26)$$

$$(\hat{E}^H \hat{E}) \cdot \vec{v} = \hat{E}^H \cdot \vec{m} \quad (3.27)$$

where \vec{m} is a vector containing the acquired k-space points for each coil, \vec{v} is a vector containing the unknown image voxel values, and \hat{E} and \hat{E}^H are the matrices which represent a combination of coil

and gradient encoding and decoding, respectively. The vector \hat{m} has a length of $NC \cdot n_k$, where NC is the number of receiver coils, and n_k is the number of k-space points acquired for each channel. The vector \vec{v} has a length N^2 , where N is the image matrix size. The encoding matrix, which necessarily has a size of $NC \cdot n_k \times N^2$, can be written as:

$$E_{L,\vec{k},\vec{r}} = e^{i\vec{r}\cdot\vec{k}} \cdot C_L(\vec{r}) \quad (3.28)$$

where L runs from 1 to NC , \vec{r} is the pixel position in the image domain, and \vec{k} is the k-space sampling point. Thus, given that $NC \cdot n_k$ is greater than N^2 , it should be possible to reconstruct the missing pixel values \vec{v} using Equation 3.27. As stated above, solving Equation 3.26 directly, i.e. by employing the inverse of the matrix $\hat{E}^H \hat{E}$, would require immense amounts memory and computation time due to the large sizes of the matrices and vectors involved. However, Equation 3.26 is a linear equation of the form

$$\hat{A}\vec{x} = \vec{b} \quad (3.29)$$

where the values in the matrix \hat{A} and vector \vec{b} are known, and the values in the vector \vec{x} are unknown. This means that although it would be computationally challenging to directly solve Equation 3.26, other solutions for this basic type of linear system can be used. In CG-SENSE, Pruessmann et al. chose to solve the equation using the well-known iterative Conjugate Gradient method [108], which is often employed when solving large systems of linear equations. This method is particularly efficient because it seeks to minimize the error in the residual by taking steps in conjugate directions, as opposed to methods such as Steepest Descent, which can take multiple steps in the same direction. Thus, given a well-behaved system of equations, it can be shown that the CG algorithm converges predictably after a small number of iterations [109, 110], which is important for fast and robust image reconstructions.

One of the important features of the CG-SENSE method is that the encoding matrix depicted in Equation 3.28 is made up mostly of Fourier terms, and one can use a Fast Fourier transform instead of calculating the \hat{E} and \hat{E}^H matrices explicitly. However, if one is to use the FFT, the non-Cartesian data, i.e. the values in the vector \hat{m} , must be gridded as described in Section 2.8. Although the original work of Pruessmann et al. used the convolution gridding approach described in Section 2.8.1, other authors have combined the method with INNG [23] (see Section 2.8.2), which is less computationally intensive. Thus, the vast majority of the operations in the CG-SENSE method are involved in either the Fourier transform or the gridding procedure. The implementation of the CG-SENSE method is depicted in Figure 3.8, which also describes each of the individual steps.

The CG-SENSE method is important for parallel imaging because it allows one to reconstruct arbitrary undersampled trajectories. Thus, unlike GRAPPA, CG-SENSE can reconstruct undersampled radial, spiral, zig-zag, etc. data without requiring a modification in its formulation. However, like SENSE, CG-SENSE requires a coil sensitivity map, which means that CG-SENSE has the same limitations as Cartesian SENSE. Although the coil map can sometimes be extracted from the undersampled non-Cartesian data, errors in this map lead to reconstruction artifacts which cannot be removed. In addition, the use of an iterative method can be more problematic than using a direct method due to

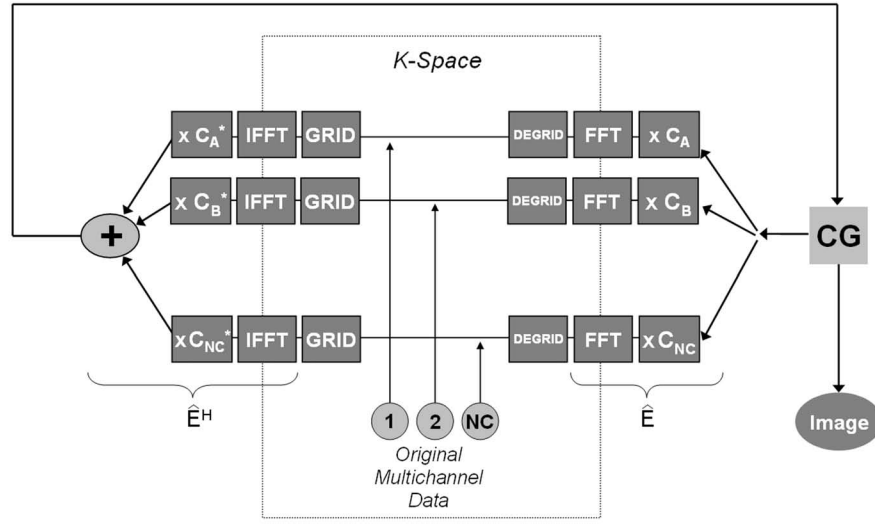


Figure 3.8: A schematic depiction of the Conjugate-Gradient SENSE algorithm. The acquired multi-channel undersampled non-Cartesian data are fed into the algorithm and then gridded, Fourier transformed, multiplied by the conjugate coil sensitivities and combined. These operations make up the decoding of the data, which is represented by \hat{E}^H matrix. The image that results is the “first guess” for the actual image, and is fed into the conjugate gradient algorithm, which yields a “second guess” for the actual image by attempting to minimize the residual. This “second guess” is then separated into multi-channel images by multiplying it with the individual coil sensitivities and performing a Fourier transform, yielding the k-space once again. These operations are equivalent to the \hat{E} matrix. This new k-space data is resampled along the original k-space trajectory to obtain the updated k-space, and the entire process is repeated. The CG-SENSE algorithm ends when the image calculated in the $(n - 1)^{th}$ iteration differs from the image from the n^{th} only by a small defined amount.

uncertainties in the stopping criteria. Finally, when performing dynamic imaging, CG-SENSE must be performed separately for each time frame, making the reconstruction process much more time consuming than k-space based methods, which typically require only one weight set for all time frames. Thus, the non-Cartesian GRAPPA methods described above are often employed, despite the fact that they are specific to the trajectory.

Parallel Imaging with Adaptive Radius (PARS)

Another method for reconstructing undersampled datasets with non-Cartesian trajectories is PARS [105, 106]. Unlike CG-SENSE, PARS is a direct method which reconstructs missing points in k-space, although coil maps are employed in both methods. Instead of generating the missing non-Cartesian points, as in the non-Cartesian GRAPPA methods, PARS directly reconstructs the Cartesian k-space points. In order to accomplish this, source points are chosen from the non-Cartesian data that falls within a specified radius k_r , the so-called local neighborhood, of the "missing" Cartesian point, $\vec{S}(k_x, k_y)$. These source points $\vec{S}(k_x + \delta x, k_y + \delta y)$, for which $\sqrt{\delta x^2 + \delta y^2} \leq k_r$, are combined using weight sets equivalent to those in GRAPPA to reconstruct the Cartesian points:

$$\sum_{\Delta k} \sum_{L=1}^{NC} n_{\Delta k, L, K} \cdot S_L(k_x + \delta x, k_y + \delta y) \approx S_K(k_x, k_y) \quad (3.30)$$

where Δk indicates all the points within the local neighborhood that are used for the fit, and L and K run from 1 to NC . In contrast to GRAPPA, the weights $n_{\Delta k, L, K}$ for PARS are determined using coil sensitivity maps instead of an auto-calibration dataset:

$$\sum_{\Delta k} \sum_{L=1}^{NC} n_{\Delta k, L, K} \cdot e^{i \cdot (-\Delta k) \cdot r} \cdot C_L(r) \approx C_K(r) \quad (3.31)$$

By comparing Equation 3.31 to Equation 3.4, one can see that the PARS weights should be identical to those that would be calculated using GRAPPA, although they are determined using the sensitivity map and the appropriate spatial harmonics instead of an auto-calibration dataset. Indeed, in the limiting case where only one source point falls within the local neighborhood, PARS is equivalent to SMASH, and when only one receiver coil is employed, PARS is equivalent to BURS, the gridding method discussed in Section 2.8.

PARS is advantageous for several reasons: the first is that any undersampled trajectory can be reconstructed using this method, as no symmetry is employed for the reconstruction. Secondly, because PARS determines the values of the Cartesian points using parallel imaging, no gridding step is needed, as the data are purely Cartesian after the reconstruction. Finally, PARS is a direct method, and not iterative, like CG-SENSE. However, PARS does have a number of disadvantages. The first is the need for coil sensitivity maps, which, as stated above, can be difficult to acquire in some situations. The second problem is the large number of weight sets which must be calculated for the reconstruction. In the worst case scenario, each Cartesian k-space location requires a different weight set; the calculation

of 256^2 distinct weight sets for a 256×256 image takes a considerable amount of time and calculating power. Similarly, in order to calculate each weight set, large matrix inversions must be performed (the relevant matrices have dimensions of $N^2 \times \Delta k$), and the sizes of these matrices increase as the k_r , and thus the number of source points in the local neighborhood, increases. Thus, small values of k_r must be employed to maintain acceptable reconstruction times, which can lead to less than optimal reconstruction quality due to the small number of source points in some locations in k-space. Although this method is highly parallelizable, it is not commonly used due to its general efficiency problems.

3.3.3 Non-Cartesian GRAPPA vs. CG-SENSE

The most commonly used non-Cartesian parallel imaging methods, non-Cartesian GRAPPA and CG-SENSE, have been examined in the previous sections. Both of these methods have advantages and disadvantages. CG-SENSE, while capable of reconstructing arbitrary undersampled trajectories, requires coil sensitivity maps, which can be difficult to generate in some cases. Non-Cartesian GRAPPA, on the other hand, must be formulated separately for each different trajectory, as it relies on symmetry in k-space for an accurate reconstruction. Methods such as radial and spiral GRAPPA also suffer because a fully-sampled k-space is required for the generation of the GRAPPA weight sets, making these methods primarily useful for dynamic imaging. A method which combines the advantages of these two reconstruction techniques, namely the ability to reconstruct undersampled datasets with arbitrary trajectories without the need for a coil map, would be advantageous.

3.4 Applications of Non-Cartesian Parallel Imaging

Due to the complexity and time requirements of non-Cartesian parallel imaging, these methods have not yet become commonplace in a clinical setting. However, there are a number of applications which have benefited from the combination of parallel imaging and non-Cartesian trajectories. For instance, CG-SENSE has been applied to accelerated spiral data to measure real time values of blood flow velocity in the aorta and through the aortic valve [111, 112]. The use of parallel imaging allowed the temporal resolution to be increased by a factor of three over the conventional imaging while maintaining a high spatial resolution. Similarly, accelerated 4D (time-resolved 3D) radial images of the coronary artery have been reconstructed using CG-SENSE with an acceleration factor of two in a single breath-hold [113].

As stated in Section 2.7.2, spiral trajectories are often employed in fMRI, where a fast acquisition is of the utmost importance. Thus, acceleration of the spiral trajectory could yield not only better temporally resolved response information, but also a higher resolution or reduced distortions (which result from the relatively long spiral read out). CG-SENSE has been used in conjunction with spiral fMRI in order to accomplish this [114, 115], although only an acceleration factor of two was employed in both cases.

CG-SENSE has also been employed to accelerate both the PROPELLER trajectory and a modified spiral trajectory for diffusion-weighted imaging [116]. A slightly modified version of the procedure is

used to simultaneously correct the phase differences between the different PROPELLER blades. This combination allows acceleration factors of up to $R=3$, which allows for the acquisition of high resolution diffusion-weighted images.

3.5 The Current State of Non-Cartesian Parallel Imaging

As can be seen in Section 3.4, there are few actual applications of non-Cartesian parallel imaging, and they are limited to CG-SENSE applied to the most common non-Cartesian trajectories. This deficit can be traced to a few sources. The first is the long reconstruction times of the CG-SENSE algorithm. In many of the references cited in Section 3.4, the reconstruction times were on the order of hours; although there are possibilities to speed up the reconstruction process, these options are not often put into practice. The reason for this is the second difficulty with non-Cartesian parallel imaging, namely its complexity. In order to employ either non-Cartesian GRAPPA or CG-SENSE, one must have a fair amount of experience in non-Cartesian parallel imaging, as they are complex methods which require many input parameters and corrections. This makes both methods challenging to implement automatically, and no major MR developer currently includes non-Cartesian parallel imaging as part of a product package. As an addendum, it is often difficult to work with fully-sampled non-Cartesian trajectories, due to the need for data gridding. For instance, exotic trajectories such as rosette [35] do not have a DCF which can be analytically described, making gridding of such datasets tricky and often subjective. Thus, developing straight-forward and parameter-free gridding and parallel imaging algorithms for non-Cartesian data are of the utmost importance. Once this is accomplished, parallel imaging could be combined with all of the non-Cartesian imaging applications described in Section 2.7.2.

This work focuses on the development of a non-Cartesian GRAPPA method which does not rely on k -space symmetry, making it applicable to all undersampled trajectories. However, in order for such a method to be practical, the undersampled non-Cartesian data must first be gridded. As can be seen in the methods described above, the gridding step never comes before the parallel imaging step; they are either simultaneous, or the gridding occurs after the reconstruction. This is due to the difficulty in gridding undersampled data, as discussed in Section 2.8. All standard gridding methods, including convolution gridding [76, 77], URS/BURS/rBURS [78, 79], INNG [80], DING [81], etc, require that the Nyquist criterion has been fulfilled; if this is not the case, errors appear in k -space areas where the region of support is less than Δk . An additional problem with performing the gridding step first is that one must know which k -space points must be reconstructed using GRAPPA. If gridding were performed first, undersampling would lead to areas of low, but not no, signal, and distinguishing points that have been “sampled” from those which must be reconstructed is impossible (see Section 6.2.2). Thus, before a new non-Cartesian GRAPPA method could be developed, a new gridding method with these properties had to be introduced.

This new gridding method, which is based on the GRAPPA Operator and uses parallel imaging concepts to shift non-Cartesian datapoints to their nearest Cartesian locations, is named GROG

(GRAPPA Operator Gridding). The basics of this technique and its properties are discussed in detail in Chapter 4. A modification to GROG which ensures that it is self-calibrating, i.e. the parallel imaging weights can be determined from the data points themselves, for several important trajectories is described in Chapter 5. The primary advantage of GROG is that it can be used to grid undersampled datasets. These gridded undersampled datasets can be further processed using a generalized non-Cartesian GRAPPA method for all trajectories, namely pseudo-Cartesian GRAPPA, which is discussed in Chapter 6. This method has the advantages of GRAPPA, i.e. that it is auto-calibrating and a direct k-space reconstruction is performed, without the need for a different formulation of GRAPPA for each trajectory.

Finally, an additional use of GROG to reconstruct undersampled datasets without using other parallel imaging methods is discussed in Chapter 7. This method uses GROG to mimic the so-called bunched phase encoding acquisition scheme of Moriguchi et al. [23, 24, 25] The use of GROG in conjunction with the generalized sampling theory of Papoulis [14] allows one to reconstruct images even when the k-space data are not strictly sampled according to the Nyquist criterion. This method, similar in function to pseudo-Cartesian GRAPPA, allows one to reconstruct undersampled datasets with arbitrary trajectories, although no GRAPPA is employed.

Chapter 4

GRAPPA Operator Gridding (GROG)

4.1 Introduction

As discussed in Section 2.7, non-Cartesian imaging has advantages over standard Cartesian imaging due to, for example, efficient k-space coverage or suppression of flow, undersampling, or off-resonance artifacts [15, 26, 28, 29, 31, 35, 38]. However, all acquired points do not necessarily fall onto a grid and must be resampled onto a Cartesian matrix before a fast Fourier transform can be performed, which is a major obstacle to non-Cartesian imaging (see Section 2.8). The convolution gridding approach of Jackson et al. [76] is the gold-standard method to prepare the raw data prior to the FFT, but it requires a density compensation function (DCF) to even out the sampling density throughout the trajectory. Unlike the radial trajectory, computing an analytic DCF is not trivial for trajectories with exotic geometries [117], such as PROPELLER [31], rosette [35], or stochastic trajectories [38]. Other gridding methods [78, 79, 80, 81] can be used to grid data acquired along such trajectories without a DCF, although these methods have other disadvantages such as noise enhancement, strenuous parametrization, or iterative reconstructions.

Instead of using convolution gridding, parallel imaging concepts can be applied [1, 2, 3, 4, 5, 6, 7, 105, 106] to reassign non-Cartesian data to the nearest Cartesian grid points. As discussed in Chapter 3, the basic idea behind parallel imaging is to use several coils at different spatial locations to acquire Nyquist undersampled data, and then to reconstruct the missing points by combining the coil information using a reconstruction algorithm such as GRAPPA [7]. In the case of using parallel imaging for gridding, the missing points are the Cartesian points, and the acquired data are the non-Cartesian points. By forming linear combinations of the multi-channel non-Cartesian datapoints, the appropriate spatial harmonics of the missing Cartesian points can be synthesized. This chapter describes GRAPPA Operator Gridding (GROG), a method for the gridding of non-Cartesian datapoints using parallel imaging, and validates this method using data collected along radial, spiral, PROPELLER, and rosette trajectories. This method has been published as a Full Paper in the journal *Magnetic Resonance in Medicine* [118].

4.2 Theory

4.2.1 Properties of the GRAPPA Operator

A parallel imaging method which is well-suited to shift non-Cartesian points to nearby Cartesian locations is the GRAPPA operator gridding (GROG) approach [118, 119]. Because this section describes the use of the GRAPPA operator for gridding, only those properties which are relevant for this task are reviewed here, and only one-to-one operators are examined. It is important to note, however, that larger GRAPPA operators which employ more source and target points can indeed be defined and utilized. These larger GRAPPA operators have the advantage that they act more like GRAPPA weight sets (see Section 3.2.3) and allow for larger k-space shifts. However, they also require a specific source point pattern, making the use of such weight sets limited, although they also have the properties described below. For a more extensive treatment of the GRAPPA operator, and a description of the uses of these larger GRAPPA operators, please refer to Section 3.2.2.

As recently described by Griswold et al [95], the GRAPPA operator is a formulation of GRAPPA where the weight set is a square matrix. The name ‘‘GRAPPA Operator’’ was chosen because this matrix increases or decreases the k-space position or energy of a point in k-space, similar to a raising or lowering operator (i.e. ladder or propagator operator) in quantum mechanics. As is the case for the corresponding raising and lowering operators in quantum mechanics, parallel imaging reconstructions performed using the GRAPPA Operator require only one basis weight set which can be applied consecutively to generate the appropriate spatial harmonics. While the similarities between the two end here, it is often helpful when examining the properties of the GRAPPA Operator to keep this comparison in mind.

The GRAPPA operator formalism has several features which make it useful for gridding. First, a GRAPPA operator can be used to shift a point in k-space by an arbitrary amount $n \cdot \Delta k_y$, where n is not restricted to integer values:

$$\vec{S}(k_x, k_y + n \cdot \Delta k_y) \approx \hat{G}_n \cdot \vec{S}(k_x, k_y) \quad (4.1)$$

where $\vec{S}(k_x, k_y)$ is a vector containing the values of the acquired point for each receiver coil, \hat{G}_n is a matrix containing the appropriate coil weighting factors (weights) for the desired shift, and $\vec{S}(k_x, k_y + n \cdot \Delta k_y)$ is the vector containing the signal from each receiver coil at the desired location. It is important to note that this formulation is equivalent to GRAPPA with a single source point and a single target point. As mentioned above, the weight set is simply a square matrix of size $NC \times NC$, where NC is the number of coil elements used for the acquisition.

A second important property of the GRAPPA operator is that an operator \hat{G}_δ for a small shift δ can be derived from an operator for a larger shift of size $n \cdot \delta$, $\hat{G}_{n \cdot \delta}$, by taking the n^{th} root of the larger operator [96]:

$$\hat{G}_\delta = (\hat{G}_{n \cdot \delta})^{\frac{1}{n}} \quad (4.2)$$

This implies that the calibration step for an arbitrary shift in one direction must be performed only once, as any smaller shift can be derived from a larger shift. The operation above can be performed by first diagonalizing the GROG weight set for a base shift \hat{G}_1 , and using this diagonal form to calculate the smaller shift operator \hat{G}_δ by employing the δ^{th} power of the eigenvalues:

$$\hat{G}_1 = \hat{E} \cdot \hat{V} \cdot \hat{E}^{-1} \quad (4.3)$$

and

$$\hat{G}_\delta = \hat{E} \cdot \hat{V}^\delta \cdot \hat{E}^{-1} \quad (4.4)$$

where \hat{E} is a matrix containing the eigenvectors of the base shift matrix, and \hat{V} is the diagonal matrix containing the eigenvalues. Because the weights are calculated using a Nyquist sampled calibration dataset, the $\Delta k = 1$ GROG weights are completely determined, and the principle roots of the appropriate weight set (i.e. those where the eigenvalue has a non-negative real part) can always be used to calculate the weight sets for smaller shifts. Thus, no time-consuming root determination must be performed as would be the case if employing undersampled calibration data.

Finally, GRAPPA operators for shifts in orthogonal directions can be derived and subsequently applied to generate a multi-dimensional shift:

$$\vec{S}(k_x + \delta_x, k_y + \delta_y, k_z + \delta_z) = \hat{G}_{\delta_x, \delta_y, \delta_z} \cdot \vec{S}(k_x, k_y, k_z) = \hat{G}_{\delta_x} \cdot \hat{G}_{\delta_y} \cdot \hat{G}_{\delta_z} \cdot \vec{S}(k_x, k_y, k_z) \quad (4.5)$$

Only unit shifts along the logical k-space axes (read, phase and partition for 3D imaging) need to be calibrated rather than every possible shift, and a shift in any direction by any amount can be calculated from this single group of orthogonal basis weights. It is important to note at this point that the GRAPPA Operator formalism assumes that the shifts along orthogonal directions can be considered to be independent from one another. In other words, orthogonal operators commute with one another in theory, and the order of their application should not impact the result of the shift calculation.

4.2.2 Gridding with the GRAPPA Operator

GROG makes use of these properties of the GRAPPA operator to rapidly map non-Cartesian datapoints to the nearest Cartesian grid location. This is performed through the following steps (for a 2D image):

- \hat{G}_x and \hat{G}_y weights are determined using the read lines (\hat{G}_x) and the phase encoding lines (\hat{G}_y) of a low-resolution Cartesian calibration signal by fitting each point to the point adjacent to it in the appropriate direction (this calibration step is described in Section 4.2.3). In the case of Nyquist-sampled Cartesian data, the base weights \hat{G}_x and \hat{G}_y are the GRAPPA operators for a shift of $\Delta k = 1$ in the read and phase encoding directions, respectively.
- The distance needed to shift a non-Cartesian point (analogous to the GRAPPA source point) to the nearest Cartesian point (analogous to the GRAPPA target point) is calculated using the k-

space trajectory. For example, if a datapoint lies at $[k_x, k_y] = [19.77, 5.83]$, the nearest Cartesian location is $[k_x, k_y] = [20, 6]$, and the distance to be shifted is $[\Delta k_x, \Delta k_y] = [+0.23, +0.17]$.

- As shown in Equation 4.2, the appropriate operators for the small shifts in the k_x - and k_y -directions are calculated using \hat{G}_x and \hat{G}_y as the base weights. For instance, if a point must be shifted by $\Delta k_x = +0.23$ and $\Delta k_y = +0.17$ and the base weights are calibrated for $\Delta k = 1$ shifts, the appropriate weight set would be calculated as shown in Equation 4.5:

$$\hat{G}_{\delta k_x, \delta k_y} = \hat{G}_x^{0.23} \cdot \hat{G}_y^{0.17} \quad (4.6)$$

- The shifted points are calculated by applying the appropriate operators to the non-Cartesian datapoints as shown in Equation 4.1 and deposited in the proper Cartesian location in k-space:

$$\vec{S}(20, 6) = \hat{G}_x^{0.23} \cdot \hat{G}_y^{0.17} \cdot \vec{S}(19.77, 5.83) \quad (4.7)$$

- In cases where multiple non-Cartesian points map to the same Cartesian grid point, a simple average is performed to arrive at the final value at that Cartesian location

A schematic which depicts the shifting of non-Cartesian k-space points to their nearest Cartesian locations, and the proper weight calculation, is shown in Figure 4.1.

The effective DCF for GROG, i.e. the simple averaging of the shifted points that are mapped to the same Cartesian location, can be understood as follows: each non-Cartesian point that is shifted to the same Cartesian point should ideally have the same value (given perfect GROG weights, no noise, and an absence of experimental factors such as relaxation). In other words, after shifting the non-Cartesian points with GROG, the resulting dataset can be treated as a dataset of Cartesian points, some of which have been acquired multiple times. Thus, the shifted points belonging to the same Cartesian location generated from different non-Cartesian points can simply be averaged.

It is important to note that for a unit sampled base grid (i.e. equal and isotropic image and grid FOV), the maximum shift magnitude in one direction that must be performed for gridding is $0.5 \cdot \Delta k$. As discussed in Section 3.2.3, the use of the GRAPPA Operator leads to less accurate reconstructions because only a single source point is used to fit the target point. Thus, the GRAPPA Operator is only practical for small shifts in k-space; larger shifts lead to larger errors in the reconstructed Cartesian points. Naturally, as in all parallel imaging methods, more coils in the receiver array allow for a more reliable determination of the weight sets, and therefore a larger shift distance. While it is possible to perform a step of $0.5 \cdot \Delta k$ accurately with GROG using most commercially available coils (see Figure 4.5), an array with a small number of detector elements or poorly arranged coils can result in artifacts due to insufficient coil sensitivity variation along that direction. Although the trajectory affects the distribution of shifts along a given direction, most shifts are smaller than this maximum. Since the shifts are variable in size and direction, artifacts that result from GROG are expected to be incoherent and appear as a noise enhancement in the image.




Figure 4.1: GROG gridding of non-Cartesian points. Cartesian destinations are at the intersections of the straight, finely dotted lines, and consecutive samples of an arbitrary trajectory are represented by solid circles. GROG grids a non-Cartesian data point by shifting it to its nearest Cartesian location via an appropriate weight set. For example, the GROG weights $\hat{G}_x^{0.22} \cdot \hat{G}_y^{0.29}$ are applied to the upper sampling location. When applied to each point along the non-Cartesian trajectory, GROG results in a purely Cartesian k-space.




Figure 4.2: An example of the GROG weight determination for the PROPELLER trajectory. The top left shows the k-space points for the non-Cartesian PROPELLER trajectory. Two of the four blades present are purely Cartesian. By selecting one (or both) of the Cartesian blades, the \hat{G}_x weights can be calculated by fitting adjacent points in the k_x -direction to each other, as described in Equation 4.8 and shown at the bottom of the figure. The same can be performed in the k_y -direction to calculate the \hat{G}_y weights. Using these GROG weights, the non-Cartesian portions of the PROPELLER trajectory can be gridded, yielding a purely Cartesian dataset (top right).

4.2.3 GROG Weight Determination

In order to employ GROG to grid non-Cartesian points, a calibration dataset must be used to determine the weights for steps in the \hat{G}_x and \hat{G}_y directions (for two-dimensional imaging). For arbitrary k-space trajectories, this calibration signal can be a low-resolution Cartesian dataset where steps of $\Delta k = 1$ (or less) are performed (i.e. \hat{G}_x is calculated from the points in the read direction, and \hat{G}_y from the points along the phase encoding direction). The GRAPPA operator can be derived in the same fashion as standard GRAPPA weights; namely, a fit of points with the appropriate relationship is performed with the use of the pseudo-inverse:

$$\hat{S}_{ACS}(k_x, k_y + \Delta k_y) \cdot \text{pinv}(\hat{S}_{ACS}(k_x, k_y)) \approx \hat{G}_y \quad (4.8)$$

In these equations, the signal matrices \hat{S}_{ACS} are made up of a collection of signal vectors from the Cartesian auto-calibration signal with the appropriate distance relationship, analogous to the weight set determination in GRAPPA (described schematically in Figure 3.3).

For some trajectories, this low-resolution Cartesian dataset is automatically present, as in the PROPELLER trajectory, where one blade can be used to calculate the base weights. An example showing how one blade in this trajectory can be used as the calibration dataset for the non-Cartesian

k-space points is shown in Figure 4.2. The same method can be employed to determine the orthogonal weight sets from an additional Cartesian dataset for other non-Cartesian trajectories, such as for the radial and spiral datasets shown here.

4.3 Methods

4.3.1 Simulations

In order to examine the effects of noise on the gridded images, simulated datasets were constructed to compare convolution gridding and GROG. A standard Shepp-Logan phantom was employed with an eight-element one-ring head coil array, where sensitivities were derived using an analytic integration of the Biot-Savart equations. The Cartesian data were resampled as radial data (200 projections, 256 read-out points, base matrix of 128×128) by sinc interpolation. Progressively larger levels of normal complex noise (real and imaginary standard deviations (σ_N) of 0.1, 0.5, 1, and 2.5, corresponding to approximately 0.1%, 0.5%, 0.95%, and 2.36% of the maximum amplitude of the DC signal) were applied to the radial k-space data. The data were then gridded with both GROG and the standard convolution gridding in order to compare the root mean square error (RMSE) of the resulting images as compared to the noiseless Shepp-Logan phantom. Note that the region of k-space support for the Cartesian phantom was reduced with a radial mask of the same extent as the radial trajectory to enforce isotropic image resolution; this will slightly affect the appearance of the standard noiseless phantom. Additionally, a 1D interpolation along the Nyquist sampled read direction (i.e. the radial arms) of each of the noisy radial datasets was performed in order to increase the data oversampling in the read direction to a factor of four. This is expected to improve the GROG reconstruction, as the additional radial points are densely sampled and therefore more non-Cartesian points contribute to each individual Cartesian location (i.e. each Cartesian point is reconstructed from non-Cartesian points which must be shifted in different directions, which has an averaging effect on the resulting Cartesian points).

In order to examine the effect of different numbers of array elements on the RMSE of the GROG-gridded radial images, several coils were simulated with similar geometries but different numbers of coil elements. As in the previous example, coil sensitivities were calculated using an analytic integration of the Biot-Savart equations. The geometry chosen was that of a head coil, i.e. the elements were arranged in a circle around the object. Following the generation of a coil sensitivity map for each of the different coils, radial datasets for each coil were simulated as stated above, and these datasets were gridded using GROG. The RMSE of each image as compared with the noiseless Shepp-Logan phantom was then calculated. In addition, a comparison was made between the inherent DCF resulting from GROG gridding of fully- and undersampled radial datasets and the corresponding DCFs proposed by Pipe [120]. To this end, a fully-sampled radial dataset with one Cartesian arm along the x-direction was simulated as described above (128 projections, 82 read-out points, base matrix of 82×82 , corresponding to a dataset which is sampled according to the Nyquist criterion in the azimuthal direction). The dataset was gridded using GROG, and the contribution of the Cartesian arm to the GROG k-space, i.e. the effective DCF along this radial arm, was examined. In this way, the effect of the averaging of multiple

non-Cartesian points which end up at the same Cartesian location could be examined. The effective GROG DCF was then compared to the standard radial Ram-Lak DCF. To examine the effective GROG DCF of an undersampled dataset, three-fourths of the radial arms were removed, and the remaining arms (32 projections, $R=4$) were gridded using GROG. The contribution of the radial arm along the x-direction was then compared to the undersampled DCF shown to have the optimal SNR and lowest artifact energy by Pipe.

Finally, in order to demonstrate that GROG can be used to grid data acquired along exotic trajectories without the use of a DCF, a rosette dataset with a base matrix size of 128×128 was simulated using the 8-channel coil described above. The trajectory is made up of 21600 points (divided into 180 read points for each of 120 phase encoding steps) and is shown on the left-hand side of Figure 4.7. Before gridding, the data were interpolated in the read direction to yield a four-fold oversampled dataset, as discussed in the paragraphs above.

4.3.2 In Vivo Experiments

In order to compare GROG with the gold-standard of convolution gridding, in vivo radial, spiral, and PROPELLER datasets were acquired with parameters as given in Table 4.1. Informed consent from the volunteers was obtained before each study. Since radial and spiral are the most commonly applied trajectories in non-Cartesian imaging, their well-known artifacts from convolution gridding can be readily compared to the GROG reconstruction results. The PROPELLER trajectory was selected for two reasons. The first is that at least one of the blades can be used to calculate the base weights as described in Section 4.1. In addition, GROG completely avoids the computation of an analytic DCF for the PROPELLER trajectory, which is a non-trivial but necessary process prior to convolution gridding.

	Radial	Spiral	PROPELLER
Read Points	512	7289 ^a	256
Phase Encoding Steps	256	4	410 ^b
Number of Coils	12	8	17
Base Matrix Size	256	192	256
Scanner	1.5T Espree	3T Trio	1.5T Espree
Sequence	FLASH	Gradient Echo	FLASH
TR	20 ms	2500 ms	20 ms
TE	5 ms	30 ms	5 ms
DCF	Ram-Lak	modified Ram-Lak	iterative ^c
Base Grid	2	2	2
Kernel Width	5	3	5

Table 4.1: A list of the measurement and reconstruction parameters for the in vivo datasets gridded using GROG and convolution gridding. ^aA center-out constant-linear-velocity spiral trajectory was used. ^bThe 410 phase encoding steps for the PROPELLER trajectory are comprised of 41 parallel lines for each of 10 blades. ^cSee [121] for a detailed description of the iterative PROPELLER DCF.

Before gridding each dataset with GROG, as described above, the data were interpolated along the read direction to an oversampling factor of four, which is expected to improve the reconstruction quality. No additional DCF was applied to the data before gridding with GROG. In addition, each dataset was also gridded with standard convolution gridding (parameters shown in Table 4.1). A Ram-Lak DCF was applied to the radial data before performing convolution gridding, and a modified Ram-Lak to the constant-angular-velocity spiral data [122]. An iteratively-calculated PROPELLER density compensation function [121] was applied prior to convolution gridding since an analytical DCF has yet to be described.

4.4 Results

4.4.1 Simulation

The results from the simulations are shown in Figures 4.3 to 4.7. As can be seen in the profiles and the RMSE in Figure 4.3, images gridded with GROG without oversampling (dashed gray line) contain elevated noise levels when compared to the convolution gridding images (solid gray line). At low noise levels ($\sigma_N < 1$), the profiles of both reconstructions are nearly identical to that of the Shepp-Logan phantom (solid black line), implying that these two methods offer similar image quality for high SNR values as the additional noise due to GROG is insignificant. However, at higher noise levels, the profile in the lower signal areas of the GROG images deviates more strongly from the standard profile than that of the convolution gridding image, which is also reflected in the higher RMSE of the GROG images. These SNR losses can be traced to the application of a weight set to a noisy data point; the underlying noise of a given datum can be enhanced by the shifting process, especially for large shifts. However, the RMSE values of the images reconstructed from data interpolated to an oversampling factor of four (thin black line) are similar to those of the convolution gridding images. This effect is to be expected, as the higher oversampling factor in the read direction indicates that more non-Cartesian points are shifted to the same Cartesian point, thus effectively averaging out the errors that occur when the weights are applied to noisy data points. Thus, when gridding low SNR data with GROG, a 1D interpolation along the read direction must be used to increase the read oversampling in order to avoid SNR losses. Alternatively, the oversampling factor in the read direction can be increased, which would yield the same results. Figure 4.4 shows the $\sigma_N = 1$ images gridded with the standard convolution gridding (left) and with GROG (interpolation factor 4, right). As can be seen in this figure, no additional noise enhancement is observed in the image gridded with GROG.

Figure 4.5 is a plot of the results of the RMSE calculation for radial datasets simulated with different numbers of coil elements and gridded using GROG. As can be seen in the figure, GROG is not effective when two coil elements are used, as there are only coil sensitivity variations which can be employed to determine the GROG weights in one direction. This leads to the appearance of artifacts in the gridded image (inset, left image) and a high RMSE. Thus, when a small number of coils is employed, the RMSE of the GROG-gridded image as compared to the standard Shepp-Logan phantom indicates that additional noise enhancement is present. However, as the number of coil elements increases, the

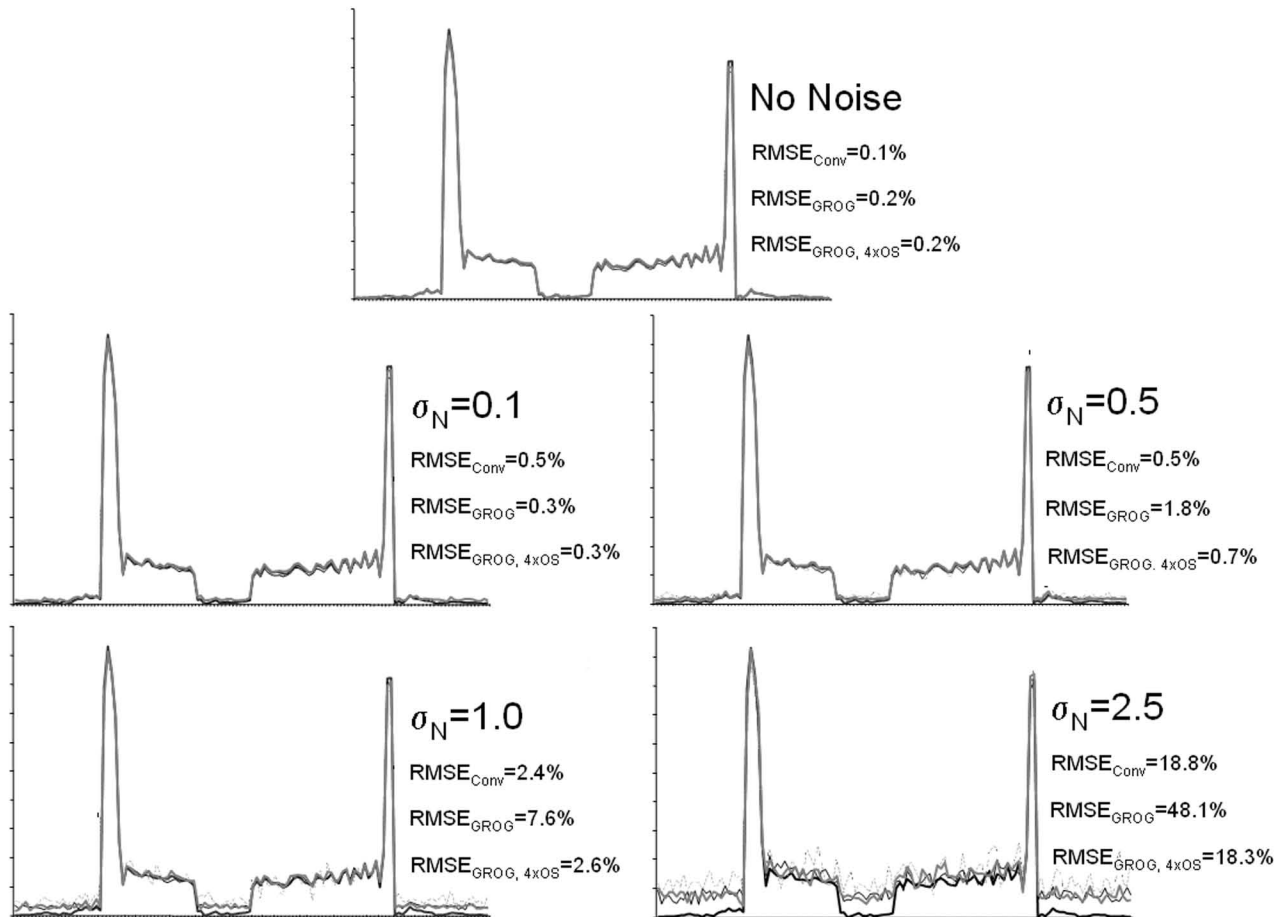


Figure 4.3: Representative profiles of the results of convolution gridding (solid gray line), GROG without oversampling (dashed gray line), and GROG interpolated to yield an oversampling factor of 4 (thin black line) with different noise levels (σ_N), compared to the noiseless Shepp-Logan phantom (thick black line). As the noise level of the dataset increases, the GROG without oversampling shows a considerable increase in noise, as indicated by the RMSE values, whereas the GROG images with oversampling have approximately the same noise level as the convolution gridding images.




Figure 4.4: Examples of a radial dataset with a noise level of $\sigma_N = 1$ gridded with convolution gridding (left) and GROG (right). As can be seen in the RMSE values shown in Figure 4.3, there is no appreciable difference between the convolution gridding and the GROG image.

RMSE values of the images approach a constant value; with the simulated coil arrangement, there is no appreciable difference between the use of six and twenty elements. The images in the center and right of the inset were reconstructed using GROG with six and twelve channels, and appear virtually identical. Because most commercially available coil arrays have at least six elements (with an increasing tendency towards larger number of elements), GROG can be performed with these coils on clinical systems without concern for additional noise enhancement due to insufficient coil sensitivity variations.

In order to examine the appearance of the density compensation in GROG, the GROG effective DCFs for fully- and undersampled radial data are compared with the Ram-Lak and undersampled DCF proposed by Pipe in Figure 4.6. The GROG effective DCFs are shown as solid lines, and the standard DCFs as dotted lines. As can be seen in the figure, the effective DCFs of GROG closely follow the standard radial DCFs, especially in the center of k-space, where the data weighting is essential to account for correlations between data points. It is important to note that the effective GROG DCF does not change after the data have become more sparse than dictated by the Nyquist criterion, as evidenced by the R=4 undersampled data (the arrow in Figure 4.6 indicates the point at which the Nyquist criterion is no longer fulfilled). This is consistent with the idea of the DCF as a tool for normalizing data correlation; as soon as the data no longer fulfill the Nyquist criterion in the azimuthal direction, compensation for varying density is no longer necessary.

Finally, to demonstrate the usefulness of GROG for the reconstruction of data collected along an exotic trajectory, the image resulting from the rosette trajectory data gridded using GROG is shown in Figure 4.7. The RMSE of this image in comparison to the noiseless Shepp-Logan phantom was calculated to be 0.17%, and the images are visually indistinguishable. It is important to note that the rosette data were gridded without a DCF or any additional gridding parameters.




Figure 4.5: Images and the corresponding RMSE values reconstructed using GROG from radial datasets simulated with different numbers of coil elements. As can be seen in the graph, the use of coils with too few elements leads to a high RMSE, indicating that the image contains artifacts due to the GROG weights. The inset image on the left shows such a case; performing GROG on radial data acquired with two coil elements leads to artifacts in the reconstructed image. As the number of coils increases, the RMSE also decreases, until a minimum is reached. The use of six coil elements offers approximately the same image quality as coils with more elements (dotted line in chart). The middle and right inset images, reconstructed using six and twelve elements, respectively, are visually indistinguishable from each other and from the reference image.

4.4.2 In Vivo Experiments

The results of the gridding of the in vivo images acquired using the radial, spiral, and PROPELLER trajectories are shown in Figures 4.8 and 4.9. In Figure 4.8, the results of convolution gridding are shown on the left, and the GROG images on the right; in Figure 4.9, the convolution gridding image is on top, and the GROG image is below. Upon visual inspection of the results, GROG and convolution gridding yield equivalent results, where any contrast or resolution differences are not evident. As discussed for the simulations in Section 4.4.1, no parameters besides the trajectories are required for the application of GROG. This is especially important for the PROPELLER trajectory, where the DCF required for most gridding methods, including convolution gridding, is complicated and non-analytical.

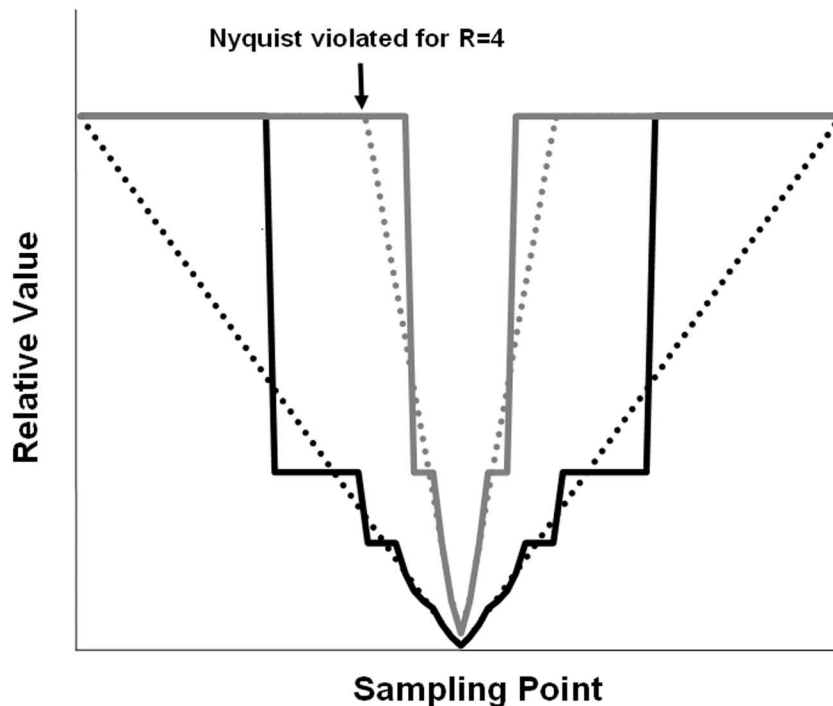


Figure 4.6: The effective GROG DCFs (solid lines) and the standard DCFs (dotted lines) for fully-sampled (black) and $R=4$ undersampled (gray) radial datasets. For the central portion of k -space, the GROG DCFs follow the standard radial DCFs closely. In addition, in the $R=4$ case, the GROG DCF follows the optimal SNR DCF for undersampled radial data as proposed by Pipe; at the point where the Nyquist criterion is violated in the azimuthal direction, the DCF becomes constant.

4.5 Discussion

GRAPPA operator gridding (GROG) has been demonstrated for radial, spiral, rosette, and PROPELLER trajectories. One advantage of GROG is that no pre-calculated density compensation functions or other parameters are required for gridding, whereas convolution gridding requires a DCF in addition to other parameters. While this is not a difficulty for the radial or spiral trajectory, the ability to grid PROPELLER, rosette, or stochastic data without having to calculate a DCF is highly advantageous. It is important to note that the effective DCF used in GROG, i.e. the averaging of shifted points which map to the same Cartesian location, cannot be used for other gridding techniques, because GROG explicitly calculates the values of the Cartesian points. Thus, after applying the appropriate GROG weights to each non-Cartesian point, the resulting dataset is made up of purely Cartesian points, which can simply be averaged. It would also be possible to weight the shifted points with scaling factors which depend on the distances of their GROG shifts, although this method of calculating the DCF has not been examined. In addition, for undersampled datasets, GROG automatically performs an




Figure 4.7: An example of an image simulated using the rosette trajectory and gridded with GROG. The trajectory used is shown in the image on the left, where the black path delineates one phase encoding step, and the reference Shepp-Logan phantom is shown in the center. The image resulting from the use of GROG to grid the simulated rosette data is on the right, and is visually indistinguishable from the reference image (RMSE = 0.17%).

approximation of the high SNR, low artifact energy DCF proposed by Pipe for undersampled datasets. Thus, undersampled data are also correctly density compensated without the need for considerations about the degree of undersampling present in the dataset.

Despite the fact that no additional parameters are needed for GROG, the *in vivo* GROG images presented here are visually indistinguishable from the convolution gridding images in terms of contrast and resolution. GROG is also less computationally intensive than convolution-based methods. For a kernel size of five, convolution gridding must process 25 gridding points for each actual non-Cartesian datapoint acquired. GROG, however, uses in essence a kernel size of one, which means that the gridding operation can be accomplished in much less time and with fewer computational requirements.

There are other methods besides GROG which perform data gridding without the need for a DCF. For instance, URS/BURS [78] is a method which performs data resampling by transforming the gridding problem into a linear equation which can be solved using singular value decomposition (SVD). As in GROG, no subsampling is employed for the gridding process and no DCF is required. However, this family of methods has several drawbacks. In the URS method, the large number of data samples leads to an inconveniently large linear equation. The BURS method is somewhat more practical, although it is highly sensitive to noise due to the need for a matrix inversion in the SVD. The extension to these approaches, rBURS [79], addresses this noise sensitivity problem, although the results are strongly dependent on the parametrization of the matrix inversion problem, i.e. the regularization and the size of the region of support. In comparison with GROG, which requires no parametrization, the URS family is much more difficult to employ, and is more computationally intensive.

Several iterative methods have also been proposed which operate without a DCF. One such approach is INNG [80], which differs from convolution gridding in that a multiplication in the image domain is substituted for a convolution in *k*-space. In this way, an effective sinc convolution can be performed in *k*-space (through the multiplication of a box-car function with the image in the frequency domain) instead of using an approximation (such as a truncated Kaiser-Bessel window). However, a





Figure 4.8: In vivo examples of images gridded with the standard convolution gridding (left) and with GROG (right) for the datasets acquired using the radial (top) and spiral (bottom) trajectories.



PROPELLER Dataset

Figure 4.9: An example of an in vivo image gridded with the standard convolution gridding (top) and with GROG (bottom) for the PROPELLER trajectory.

number of iterations of the INNG algorithm are required before an artifact-free image results from the non-Cartesian data. Each iteration involves one Fourier transformation and one inverse Fourier transformation, which becomes quite time-consuming when large oversampling factors are used. A second iterative gridding algorithm is DING [81], which also formulates the gridding problem as a linear equation. Unlike the URS family, this linear equation is solved using a conjugate gradient optimization. Because no subsampling is used, this method is less computationally intensive than INNG, although each iteration requires the same number of Fourier transformations. However, the convolution window used in DING must be optimized, again leading to potential difficulties in parametrization.

The advantage of GROG over such methods is that GROG is a direct method, and requires no iteration or parametrization (i.e. subsampling, iteration numbers, window widths, or window shapes). In addition, all of the methods described above assume that the non-Cartesian data fulfill the Nyquist criterion, which is not an assumption of GROG. However, it is important to note that unlike convolution gridding, URS/BURS/rBURS, INNG, or DING, GROG can only be used to grid data acquired with a multi-element coil; this is not a requirement for most other gridding methods. The coil array employed must offer sufficient sensitivity variation to allow the calculation of robust weights. Because most modern scanners employ multi-element coils due to their SNR advantages, besides the option of using parallel imaging, this condition is not a serious constraint. In addition, as demonstrated in the examples shown here, standard coils which are readily available in a clinical setting possess the sensitivity variations needed to reconstruct images using GROG. While the 18-channel abdomen coil used to acquire the PROPELLER data may not be currently commonplace, the 12-channel and 8-channel head coils used to acquire the radial and spiral data, respectively, are both standard and commercially available. Thus, the need for multi-channel data from appropriate coils is not a major obstacle to the use of GROG.

For datasets with low SNR, the use of GROG without interpolation may increase the noise in the final images as compared with convolution gridding, as shown in Figure 4.3. This SNR loss can be attributed to the application of weight sets to the individual noisy points, which amplifies the noise in the shifted point. However, if the data are first interpolated along the read direction to increase the read oversampling, more non-Cartesian points can be used to reconstruct a single Cartesian point. The use of more non-Cartesian points, which must be shifted in different directions, acts as an effective averaging, which combats the noise enhancement. As demonstrated in Figure 4.3, interpolation of the noisy data to an oversampling factor of four before gridding with GROG leads to images with RMSE values which do not differ greatly from those of the images resulting from convolution gridding. Thus, noise enhancement is not observed when using GROG to grid non-Cartesian data collected with clinical multi-element coils, even in datasets with lower SNR.

In the implementation of GROG described here, a calibration dataset is required. For some non-Cartesian trajectories, such as PROPELLER, this calibration dataset is automatically acquired as a part of the trajectory. However, for many others, the trajectories do not contain this Cartesian portion, and they are not self-calibrating using the method described here. However, a method of obtaining the weights from the non-Cartesian data themselves for radial and spiral trajectories has been developed

during the course of this thesis, and is described in detail in Chapter 5, as well as in several publications by the author [123, 124, 125]. For trajectories which are not inherently self-calibrating, such as the rosette trajectory, a pre-scan can be employed to generate the Cartesian calibration dataset required for the use of GROG.

As stated in Section 4.2.1, the use of the GRAPPA Operator formalism for gridding assumes implicitly that the matrices for unit shifts in the logical directions commute. The basis for this assumption is that the shifts in orthogonal directions should be independent of one another; a shift in the k_x -direction followed by a shift in the k_y -direction should be the same as performing the same shift in the opposite order. However, for actual receiver coils, this assumption may not always hold true, and the values of the final k-space points may depend on the order of the shifting. This phenomenon has not yet been examined in detail, although it is likely due to noise- and data-correlation between the receiver coils. Indeed, the k-space values of datasets gridded when changing the order of the application of the \hat{G}_x and \hat{G}_y weights in Equation 5.3 are not identical to the values obtained when using the stated order, although the image artifacts are similar in terms of appearance and severity. Despite the fact that the images gridded using GROG appear not to suffer from the fact that the weight sets do not commute, it must be kept in mind that this assumption is not necessarily fulfilled in actual multi-channel non-Cartesian datasets.

Finally, the GRAPPA operator gridding approach is less computationally burdensome than other possible parallel imaging-based gridding techniques. GROG requires only one weight set for each orthogonal direction in k-space, i.e. two dimensional data require two weight sets. Using the method described in Section 4.2.3, the specific weight sets needed to shift the non-Cartesian points to the nearest Cartesian locations are then calculated out of these two weight sets. Thus, both the calibration and reconstruction processes are fast due to the small number of weight sets. In contrast, one could imagine using a technique such as PARS [105, 106] for data gridding. PARS examines the local vicinity of a missing point to find possible source points, and then uses these points together with a sensitivity map to calculate a weight set which allows one to reconstruct the missing point. PARS is clearly inferior in terms of computational complexity to GROG due to the repeated search operations. In addition, because GROG requires such a small number of weight sets, these weights can be calculated from simple Cartesian calibration data. Thus, GROG does not require coil maps, weight interpolation, or differently spaced Cartesian calibration data, which simplifies the reconstruction with respect to other possible parallel imaging-based gridding schemes.

An extremely important feature of GROG which has not been discussed in detail in this chapter is that GROG can be used to grid undersampled data as well as fully-sampled data, which is not the case in convolution gridding. Convolution gridding methods assume that the non-Cartesian data fulfill the Nyquist criterion throughout the trajectory. Despite acceptable streak artifacts for angularly undersampled radial trajectories, gridding with convolution-based methods essentially “fills up” more k-space locations than were actually measured. In the extreme case of a single point, GROG can be used to shift this point to the nearest Cartesian location, yielding a single point in k-space surrounded by zeros, a feat which cannot be accomplished with convolution gridding. The ability of GROG to

grid undersampled non-Cartesian datasets to yield Cartesian datapoints near the sampled locations and zeros in all other locations is be advantageous for other types of non-Cartesian parallel imaging reconstructions. For instance, non-Cartesian GRAPPA algorithms [100, 102, 126], as discussed in 3.3.1, first reconstruct missing datapoints, and then grid the reconstructed data. However, as demonstrated in Chapter 6 as well as in several manuscripts by the author [127, 128, 129], the undersampled data can instead first be gridded with GROG followed by a Cartesian GRAPPA reconstruction of the resultant Cartesian data points. This method, namely pseudo-Cartesian GRAPPA, is the first application of GROG for performing tasks other than the gridding of Nyquist sampled datasets, and is discussed in Chapter 6.

Similarly, the ability of GROG to shift acquired points by small amounts is advantageous for a number of other applications [130, 131, 132, 133, 134]. One could, for instance, use GROG to mimic the Bunched Phase Encoding trajectories of Moriguchi et al. [24, 23, 25], a topic discussed in great detail in Chapter 7 of this thesis. Thus, while GROG is introduced here as an alternative to standard gridding algorithms, it can also be used for a variety of purposes.

4.6 Conclusion

The GRAPPA Operator Gridding (GROG) technique proposed here is an alternative method for the gridding of non-Cartesian datasets. Instead of employing a convolution window as in the gold-standard gridding, the GRAPPA operator is instead used to shift the non-Cartesian datapoints to their nearest Cartesian locations. The calculation and application of the GROG weights is computationally efficient, as only one weight set is needed per orthogonal direction in k-space. Gridding of data acquired using radial, spiral, rosette, and PROPELLER trajectories has been demonstrated, where visually equivalent contrast and resolution were obtained with GROG and standard convolution gridding. Simulations were conducted which show that noise enhancement resulting from the use of GROG is minimal when a 1D interpolation is performed to increase the oversampling in the read direction. Only the trajectory, a small Cartesian calibration dataset and multi-channel data are needed to perform GROG. This simplicity is advantageous for trajectories, such as PROPELLER, rosette, or stochastic sampling, that would otherwise require the calculation of a non-trivial DCF prior to convolution gridding. Additionally, single points or undersampled data can also be gridded and properly density compensated with GROG. In summary, GROG demonstrates a novel use of parallel imaging to perform tasks other than the reconstruction of undersampled data.

Chapter 5

Self-Calibrating GRAPPA Operator Gridding

5.1 Introduction

GRAPPA Operator Gridding (GROG), a method to grid non-Cartesian data, has been introduced in Chapter 4. This method has many advantages over other gridding methods. For instance, some gridding methods, namely convolution gridding [76] or URS/BURS/rBURS [78, 79], are computationally intensive and/or highly sensitive to gridding parameters. Other methods, such as INNG [80] or DING [81], are iterative methods which can also be both time consuming and difficult to parameterize. In addition, many of these methods require a density compensation function (DCF) to counteract the uneven sampling density of non-Cartesian acquisition schemes.

GROG, unlike the methods described above, is not iterative and does not require grid oversampling, making the algorithm fast and efficient. In addition, no DCF must be employed with GROG, which is advantageous for exotic trajectories which have non-trivial acquisition density. GROG uses a completely new approach for gridding; the method is similar to the parallel imaging method GRAPPA [7] in that it employs a combination of coil sensitivities as an alternative to gradient encoding to reconstruct missing points. In contrast to conventional parallel MRI, coil sensitivity variations are employed in GROG to shift non-Cartesian data points onto a Cartesian grid.

However, before one can use GROG to grid non-Cartesian points to obtain a 2D image, several requirements must be met. First, the data must be acquired with a coil array which provides enough sensitivity variations in both the x- and y-directions to allow the calculation of accurate base weight sets for GROG, i.e. \hat{G}_x and \hat{G}_y (and in the case of three-dimensional imaging, \hat{G}_z as well). In addition, as described in Section 4.2.3, a calibration dataset is required to determine the values of \hat{G}_x and \hat{G}_y (the Δk_x and Δk_y shift operators). This calibration dataset usually takes the form of Cartesian data, where the datapoints have a constant distance from each other in the k_x and k_y directions. Using such data, \hat{G}_x and \hat{G}_y can be calculated directly from the data as in GRAPPA. Some trajectories, such as PROPELLER or 1D non-Cartesian, contain a Cartesian portion which can be used for calibration. However, for most trajectories, namely those which do not contain a Cartesian portion, a separate

Cartesian dataset must be acquired for this calibration step. The need for an extra calibration dataset is a major disadvantage, as it requires extra scan time and can introduce errors due to image mismatches, and greatly reduces the usefulness of GROG.

The innovation described in this chapter, self-calibrating GROG (SC-GROG), is that the values of \hat{G}_x and \hat{G}_y can be calculated directly using the non-Cartesian data points. In order to these values, the relationship between many pairs of multi-channel points along the trajectory must be known. However, for this calibration step, these points do not necessarily need to lie on the Cartesian grid. This chapter outlines the determination of the base weight sets \hat{G}_x and \hat{G}_y using radial (or spiral) data themselves. This extension of GROG eliminates the need for an additional Cartesian dataset for many trajectories, thereby restoring the advantages of GROG as a gridding method. In addition, the ability to use the non-Cartesian datapoints themselves for calibration eliminates potential mismatch errors between the non-Cartesian dataset and the calibration dataset, making gridding with GROG more robust. This method has been described in a Technical Note in the journal *Magnetic Resonance in Medicine* [123].

5.2 Theory

5.2.1 GROG Review

GRAPPA Operator Gridding (GROG), introduced in Chapter 4, is briefly reviewed here. GROG performs the gridding process by shifting each acquired non-Cartesian point to its nearest Cartesian location using a GRAPPA-like weight matrix of size $NC \times NC$, where NC is the number of coils in the receiver array. Because a large shift is equivalent to the consecutive applications of several smaller shifts, the same idea can be used to derive weight sets for smaller shifts given the weight set for a larger shift. Specifically, given the appropriate weight set for a unit shift of Δk in a single logical direction, the weight set for a smaller jump, such as $0.25\Delta k$, can be found by taking the appropriate root of the weight set for the larger jump, as described in detail in [96]:

$$\hat{G}_1 = \hat{G}_{0.25}^4 \quad (5.1)$$

and consequently

$$\hat{G}_{0.25} = \hat{G}_1^{0.25} \quad (5.2)$$

By using properties of the GRAPPA Operator, a general equation for shifting points in arbitrary directions and distances in k-space can be written:

$$\vec{S}(k_x + \delta_x, k_y + \delta_y) = \hat{G}_x^{\delta_x} \cdot \hat{G}_y^{\delta_y} \cdot \vec{S}(k_x, k_y) \quad (5.3)$$

where $\vec{S}(k_x, k_y)$ is a vector containing the acquired signal from each of the receiver coils at the k-space location $[k_x, k_y]$ to be shifted by an amount $[\delta_x, \delta_y]$; \hat{G}_x and \hat{G}_y are the appropriate GROG weight sets for unit shifts in the k_x - and k_y - directions, respectively; and $\vec{S}(k_x + \delta_x, k_y + \delta_y)$ is the shifted signal

at location $[k_x + \delta_x, k_y + \delta_y]$. GROG works by using the appropriate root of the base weight sets to shift each non-Cartesian point to the nearest Cartesian location on a grid using Equation 5.3. In this case, the source points, or $\vec{S}(k_x, k_y)$ in Equation 5.3, are the non-Cartesian points, which are shifted by the application of the appropriate weight sets to $\vec{S}(k_x + \delta_x, k_y + \delta_y)$, the nearest Cartesian locations. A schematic of data gridding with GROG is shown in Figure 4.1. It is important to note that in this formulation, the order of the shifts, and thus the order of the application of the \hat{G}_x and \hat{G}_y weights, is irrelevant. In other words, the GROG weight matrices along orthogonal directions commute, because a shift in one logical direction should be independent of a shift along another direction.

In cases where multiple non-Cartesian points are mapped to the same Cartesian location, the results of the GROG shifts are simply averaged. Because the points are purely Cartesian after this shift is performed, no density compensation function or other corrections are required after the gridding step. It is also important to note that because only one source point (the non-Cartesian point) is employed, most standard coil arrays can accurately perform maximum shifts of $\Delta k = 0.5$. For that reason, each non-Cartesian point is mapped only to its nearest Cartesian neighbor; thus, GROG cannot be used for the reconstruction of unaliased images from undersampled data, as is commonly expected from conventional parallel imaging. However, unlike most other gridding methods, GROG is capable of accurately determining the Cartesian k-space values for undersampled non-Cartesian data because GROG does not implicitly assume that the Nyquist criterion has been fulfilled.

As stated above, when employing a Cartesian calibration dataset, the GROG weights can be determined as in GRAPPA. However, most non-Cartesian datasets (with the exception of trajectories such as PROPELLER [31]) do not have enough Cartesian-like data to allow the GROG weights to be determined directly, and an additional Cartesian calibration dataset must be acquired. The method describes in this chapter seeks to remove the need for such a Cartesian calibration dataset for radial and spiral trajectories.

5.2.2 Self-Calibrating GROG

The simplest way of calculating GROG weights directly from radial data would be to use the read points from 1 to $L - 1$ (where L is the number of read-out points) along two radial rays which are orthogonal to one another (see Figure 5.1). This method, while straightforward, suffers from the fact that few datapoints in the dataset are actually used for the calibration; low SNR values in these two rays will lead to less stable GROG weights and poor GROG reconstructions.

A more complicated, but stable, method is to determine the weights using all of the radial rays. In order to accomplish this, intermediate “angular” weights are first calculated for each radial ray from θ_1 to θ_{N_y} (as shown in Figure 5.1):

$$\vec{S}(\theta, r + 1) = \hat{G}_\theta \cdot \vec{S}(\theta, r) \quad (5.4)$$

$$\hat{G}_\theta = \hat{S}(\theta, r) \cdot \text{pinv}(\hat{S}(\theta, r + 1)) \quad (5.5)$$




Figure 5.1: For self-calibrating GROG using radial data, one could imagine using only the radial rays which are orthogonal to each other (here, the gray rays) in order to obtain the GROG weights. However, very little data is employed, which results in unstable weight sets (see text). Instead, weights along each projection are calculated (the \hat{G}_{θ} matrices shown schematically for four rays) as described in Equation 5.5 in the text. These angular weights can then be separated into their \hat{G}_x and \hat{G}_y components (as shown on the right-hand side of the figure) using Equation 5.6, thereby allowing an accurate calculation of these base weights as described in the text.

using the pseudo-inverse (pinv) of this relationship and all the points along the ray. In Equations 5.4 and 5.5, θ denotes the angle, or “phase encoding step” of the specific ray, r is a point in the read direction and runs from 1 to $L - 1$, $r + 1$ is the next read point along the ray, and \hat{G}_θ is the $NC \times NC$ weight set which describes a shift along the ray (see Figure 5.1 for a better understanding of the notation). In Equation 5.4, the vector \vec{S} is the multi-channel signal for a point in the non-Cartesian k-space, and in Equation 5.5, $\hat{S}(\theta, r)$ is a matrix containing the signal values from 1 to $L - 1$ along a ray at angle θ for all of the receiver channels. To clarify, for a radial dataset with 200 projections and 128 read points, both signal matrices \hat{S} in Equation 5.5 would have dimensions of $NC \times 127$, and this equation appears Ny , or 200, times.

Since a direct shift along a ray can be described by two consecutive shifts along the x- and y-directions, each angular weight set \hat{G}_θ can then be written as a function of the \hat{G}_x and \hat{G}_y weights:

$$\hat{G}_\theta = \hat{G}_x^n \cdot \hat{G}_y^m \quad (5.6)$$

where n and m are the shifts in the k_x - and k_y -directions, respectively, and \hat{G}_x and \hat{G}_y also have the dimensions $NC \times NC$. Once the \hat{G}_θ values have been calculated for each ray using Equation 5.5 (from the first ray, θ_1 , to the last ray, θ_{Ny} , and thus Ny different values of n and m), the following equations containing the unknowns \hat{G}_x and \hat{G}_y can be written:

$$\hat{G}_{\theta_1} = \hat{G}_x^{n_1} \cdot \hat{G}_y^{m_1} \quad (5.7)$$

$$\hat{G}_{\theta_2} = \hat{G}_x^{n_2} \cdot \hat{G}_y^{m_2} \quad (5.8)$$

⋮

$$\hat{G}_{\theta_{Ny}} = \hat{G}_x^{n_{Ny}} \cdot \hat{G}_y^{m_{Ny}} \quad (5.9)$$

Although this set of equations is non-linear, it can be simplified by taking the matrix logarithm of both sides:

$$\ln(\hat{G}_{\theta_1}) = n_1 \cdot \ln(\hat{G}_x) + m_1 \cdot \ln(\hat{G}_y) \quad (5.10)$$

$$\ln(\hat{G}_{\theta_2}) = n_2 \cdot \ln(\hat{G}_x) + m_2 \cdot \ln(\hat{G}_y) \quad (5.11)$$

⋮

$$\ln(\hat{G}_{\theta_{Ny}}) = n_{Ny} \cdot \ln(\hat{G}_x) + m_{Ny} \cdot \ln(\hat{G}_y) \quad (5.12)$$

The previous step assumes that the weight sets \hat{G}_x and \hat{G}_y commute, an assumption based on the fact that the weight sets perform shifts in orthogonal directions. This condition may not be exactly fulfilled in real imaging scenarios, due to noise, correlated data, etc. However, deviations from this assumption are small and in first approximation this condition is fulfilled for many applications. A more detailed discussion of this assumption and its implications is given in Section 5.5.

Equations 5.10 to 5.12 can then be rearranged into NC^2 independent matrix equations, where a and b both run from 1 to NC (i.e. over each coil):

$$\vec{V}(a, b) = [\vec{n}, \vec{m}] \cdot \left[\ln(\hat{G}_x(a, b)), \ln(\hat{G}_y(a, b)) \right] \quad (5.13)$$

In this equation, $\vec{V}(a, b)$ is the vector made up of the appropriate elements (namely in position $[a, b]$) of the logarithms of the angular weight sets on the left-hand side of Equations 5.10 to 5.12 and consequently has a length Ny ; \vec{n} and \vec{m} are vectors of length Ny , which together form a matrix of size $Ny \times 2$; and the two unknowns $\ln(\hat{G}_x(a, b))$ and $\ln(\hat{G}_y(a, b))$ form a vector of length 2. This equation is solved using the pseudo-inverse of the distance matrix, i.e. $[\vec{n}, \vec{m}]$ (which must be calculated only once), for each value of a and b (i.e. coil by coil), and the values in \hat{G}_x and \hat{G}_y are calculated by taking the matrix exponents of the resulting matrices (which must be rearranged again into the appropriate $NC \times NC$ matrices after the calculation):

$$\text{pinv}([\vec{n}, \vec{m}]) \cdot \vec{V}(a, b) = \left[\ln(\hat{G}_x(a, b)), \ln(\hat{G}_y(a, b)) \right] \quad (5.14)$$

and

$$\hat{G}_x = \exp(\ln(\hat{G}_x)) \quad (5.15)$$

$$\hat{G}_y = \exp(\ln(\hat{G}_y)) \quad (5.16)$$

In this way, the base weights \hat{G}_x and \hat{G}_y can be calculated from the radial datapoints themselves, using all of the available radial data.

This principle can also be applied to spiral trajectories. While most spiral data are acquired using the constant-linear-velocity-trajectory due to a lower demand on the gradients and SNR advantages [30], these data can be converted to the constant-angular-velocity trajectory using a simple 1D interpolation in the time domain. The advantage of the constant-angular-velocity spiral trajectory is that datapoints acquired using such a trajectory can be resorted into radial data [28], as the number of sampling points per winding is constant. Thus, using the radial self-calibration scheme described above, the base weights \hat{G}_x and \hat{G}_y can be calculated from spiral data, making this trajectory self-calibrating as well.

5.3 Methods

5.3.1 Simulation

Radial datasets both with and without noise were simulated in order to examine the differences between convolution gridding, self-calibrating GROG where only two radial arms are employed, and self-calibrating GROG with all available radial arms. To that end, an eight-element one-ring head coil array was simulated, and the sensitivities derived for this coil array using an analytic integration of the Biot-Savart equations were applied to the standard Shepp-Logan phantom. This simulated coil was chosen because previous work [118] has shown that the arrangement and number of coils allows accurate k-space shifts of $0.5\Delta k$ in both the x- and the y-direction using GROG without the appearance of artifacts. The simulated phantom image was Fourier transformed to produce Cartesian k-space data, which were then resampled as radial data (200 projections, 256 read-out points, base matrix of 128×128) by sinc interpolation. Datasets both without and with noise (real and imaginary standard deviation (σ_N) of 1, corresponding to approximately 0.95% of the maximum amplitude of the DC signal applied to the radial k-space data) were gridded using the methods described above, and the root mean square error (RMSE) of the resulting images were compared to the noiseless Shepp-Logan phantom. As a comparison, the datasets were also gridded using the standard convolution gridding with a Kaiser-Bessel window of width 5, grid oversampling factor of 2, and a Ram-Lak DCF. Note that the region of k-space support for the Cartesian phantom was reduced with a radial mask of the same extent as the radial trajectory to enforce isotropic image resolution, which slightly affects the appearance of the standard noiseless phantom. After the gridding processes, the resulting multi-coil k-space datasets were Fourier transformed into the image domain coil-by-coil and combined using a Sum-of-Squares method to arrive at the final image.

5.3.2 Experiments

In vivo radial data (12 channel, 256 projections, 512 read-out points, base matrix size of 256×256) were acquired using a 1.5T Siemens Avanto scanner (Siemens Medical Solutions, Erlangen, Germany), and in vivo spiral data (32 channel [87], 4 spiral arms, 7289 read-out points, base matrix size of 192×192) were acquired using a 3T Trio scanner (Siemens Medical Solutions) using a constant-linear-velocity trajectory. The base GROG weights were determined from the radial data as described in Section 5.2.2 without further processing; the spiral data were first reordered onto a constant-angular-velocity trajectory using a 1D Fourier interpolation in the time domain (i.e. read-out direction). This reordering yielded 448 pseudo-projections, each containing 128 points which could be used for the calibration of the GROG weights. After the weights were determined using the method described above, the datasets were gridded using GROG. In addition, convolution gridding was performed as a comparison. As in the simulations, Kaiser-Bessel windows of widths 5 and 3 for the radial and spiral data, respectively, were employed for the convolution gridding. A Ram-Lak DCF was applied to the radial data before performing convolution gridding, and a modified Ram-Lak to the constant-angular-velocity spiral data [122].

5.4 Results

The results of the radial simulations with and without noise are shown in Figure 5.2. The top row shows images gridded using the standard convolution gridding, 2-arm self-calibrating GROG, and SC-GROG with all 200 arms for the noiseless radial dataset. In addition, the upper right hand side of each reconstructed image shows the difference between the reconstruction and the noiseless Shepp-Logan phantom multiplied by a factor of 5. Although all three images appear quite similar, the weights obtained from only two orthogonal radial arms result in incoherent image artifacts (appearing here as an increased noise level), indicated by the slightly higher RMSE of 0.3%. However, both the convolution gridding image and the self-calibrated GROG image with all 200 arms show similar artifact energy (0.1%). The result can be seen more clearly when examining the radial datasets with noise, shown in the lower row of Figure 5.2. Again, the difference images multiplied by a factor of 5 are shown in the upper right hand corners. As in the noiseless dataset, the images gridded with conventional convolution gridding and the 200-arm self-calibrating GROG exhibit similar RMSE values (2.5% and 2.7%, respectively). GROG calibration using only two of the radial arms results in a much higher RMSE value (9.3%), which can be seen as additional noise in the gridded image. Thus, although it is possible to calibrate the GROG weights using just two orthogonal radial arms, the weights which result are unstable; for all further images shown here, all radial rays were used for the calibration. As evidenced by the RMSE values of the self-calibrating GROG images as compared to those for the convolution gridding images, little to no noise enhancement can be seen when gridding with self-calibrating weights (given adequate coil sensitivity variations). A slight increase in background noise is noticeable in the difference images for SC-GROG, although the fidelity of the edges of the Shepp-Logan phantom appears to be slightly better for these images than for the convolution gridding reconstructions.

Figures 5.3 and 5.4 show images gridded using convolution gridding (left) and self-calibrating GROG (right) for the in vivo radial and spiral acquisitions, respectively. As can be seen in these figures, the results of the self-calibrating GROG are visually indistinguishable from the convolution gridding images. In addition, there is no observable noise enhancement, and the contrast and resolution of the self-calibrating GROG images are impossible to visually differentiate from that of the standard images.

5.5 Discussion

As shown in the simulations and in vivo images, self-calibrating GROG yields images that are visually indistinguishable from images gridded using the gold-standard convolution gridding. For both the noiseless simulations and those with noise, the RMSE values of the convolution gridding images and the self-calibrating GROG images were similar, and the images themselves cannot be distinguished with the naked eye. Similarly, the in vivo images gridded with GROG appear visually equivalent to the convolution gridding images, as shown in Figures 5.3 and 5.4. Thus, the ability to calibrate the GROG weights from the non-Cartesian data themselves removes the main drawback of GROG, namely the need for an additional Cartesian calibration dataset, as required in standard GROG. Self-calibrating GROG maintains the other advantages of standard GROG, namely the lack of gridding parameters such as




Figure 5.2: Radial simulations using the Shepp-Logan phantom without (top) and with (bottom) noise. Images reconstructed using convolution gridding are shown on the left, those generated using self-calibrating GROG and only two orthogonal radial arms for calibration are shown in the center, and those created with GROG weights calibrated using the entire 200 radial arms are shown on the right. Difference images multiplied by a factor of five for each reconstruction as compared to the noiseless Shepp-Logan phantom are shown in the upper right-hand corners. As can be seen both visually in the reconstructions, RMSE values, and difference images, the images gridded using GROG weights from two radial arms are much noisier than the convolution gridding images. However, when all radial arms are employed for the GROG calibration, the RMSE values are comparable to those for the convolution gridding images, and the difference images show only a slight noise increase in the background of the images. Thus, when employing weights calibrated using the entire radial dataset, the images do not exhibit a significant noise enhancement or artifact level in comparison to convolution gridding.

convolution shapes, window widths, oversampling factors, or density compensation functions. The only information that is required for self-calibrating GROG is a knowledge of the sampling positions. In addition, because the kernel width for GROG is one (as opposed to the standard kernel width of 3 or 5 for convolution gridding), no grid oversampling is required, and time intensive Fourier transformations are avoided, GROG is much less computationally intensive than methods such as URS/BURS/rBURS, INNG, or DING. However, it is important to note that unlike other gridding methods, GROG requires a multi-element coil for signal detection due to the parallel imaging reconstruction process. As has been shown in previous work [118], most standard coils available on modern clinical scanners provide the needed sensitivity variations. Indeed, because the shift sizes required for GROG (a maximum of $0.5\Delta k$ in each logical direction) are small in comparison to the shift sizes in GRAPPA, less sensitivity variation is needed, indicating that this requirement does not present a major obstacle to the use of GROG.

Another advantage of self-calibrating GROG is that there is no longer a potential mismatch between the calibration data and the non-Cartesian data. In cases where resolution, Field-of-View, or coil sensitivity mismatches (for example, image distortion in spiral acquisitions) could occur in the two images, or motion takes place between the acquisition of the two datasets, the weights determined for GROG using the Cartesian dataset will be suboptimal, leading to a noise enhancement or even artifacts in the GROG image. For this reason, in addition to the measurement time reduction, the application of weights determined using self-calibrating GROG is clearly superior to using those determined as in standard GROG. It is important to note that the amount of time needed to calculate the GROG weights using the self-calibrating method with 200 radial arms is slightly longer than the standard method, although both are on the order of a second using non-optimized Matlab code for typical datasets, such as those shown here.

As stated in Section 5.2.2, the self-calibrating GROG method relies on the assumption that the matrices for unit shifts in the logical directions commute. This assumption is based on the basic properties of the GRAPPA Operator, as discussed in Section 4.2.1, and is fundamental to self-calibrating GROG. As seen in both the self-calibrating and Cartesian GROG, the weight sets often do not commute for actual receiver coils, and the values of the final k-space points depend on the order of the application of the weights. However, as demonstrated in the examples shown here, the GROG weight sets determined under this assumption, and the images reconstructed with them, do not suffer in quality even when the weight sets do not commute. Although the images gridded using self-calibrating GROG appear visually identical to those generated using convolution gridding despite the fact that the base weight sets do not commute, it must be kept in mind that the self-calibrating calculation of the GROG weights relies on this assumption.

The efficient self-calibration method has been demonstrated here using radial and spiral data. Other non-Cartesian trajectories, namely the PROPELLER [31] trajectory, are inherently self-calibrating, as one blade can always be used as a Cartesian calibration dataset in order to grid the data along the other blades. However, trajectories which do not contain such a Cartesian portion or have a radial symmetry cannot be calibrated using either the standard GROG calibration or the method proposed




Figure 5.3: : In vivo radial images gridded using convolution gridding (left) and self-calibrating GROG (right). These images are visually equivalent, and no difference can be seen in contrast or resolution. In contrast to the convolution gridding image, the self-calibrating GROG image required only knowledge of the trajectory, as no additional parameters are necessary for the gridding process.




Figure 5.4: In vivo spiral images gridded using convolution gridding (left) and self-calibrating GROG (right). As in the radial case, the two images are indistinguishable. The self-calibrating GROG process did not require a density compensation function or other parameters which were required for the convolution gridding method.

here. This is because a large group of points which have the same k_x and k_y distances are needed for the self-calibration method (due to the need for angular weight sets), as can be seen in Equation 5.5. This means that the GROG weights for trajectories such as rosette [35], zig-zag [23, 24], or stochastic [38], cannot currently be found from the datapoints themselves, although work on this topic is currently in progress. However, because the number of unknowns (the \hat{G}_x and \hat{G}_y matrix values) is much smaller than the number of possible knowns (adjacent points in the trajectory), it is expected that the GROG self-calibration for other exotic trajectories will involve a numerical optimization process, and should be realizable given sufficient computing power.

5.6 Conclusion

The self-calibrating GROG method presented here removes the major obstacle to the use of GROG, namely the need for an additional Cartesian calibration dataset. For radial and spiral trajectories, the non-Cartesian dataset can be used to generate the orthogonal GROG base weights, which are then employed to grid the datapoints themselves. Simulations and in vivo examples of data gridded using this method have been compared to images gridded with convolution gridding, and the results are indistinguishable in both contrast and resolution. Self-calibrating GROG maintains the advantages of GROG, in that gridding can be accomplished without DCFs or other gridding parameters such as convolution shapes, window widths, or oversampling factors, while maintaining a low computational complexity. The only requirement for the use of GROG to grid radial and spiral datasets is multi-channel data which is needed for the application of this parallel imaging based method.

Chapter 6

Pseudo-Cartesian GRAPPA

6.1 Introduction

As explained in Section 3.3, the application of parallel imaging techniques to undersampled non-Cartesian trajectories is non-trivial. This is due to the unusual point spread functions of nonuniformly sampled k-space which results from undersampled non-Cartesian trajectories [44, 45]. However, non-Cartesian sampling schemes offer some inherent advantages due to their better coverage of the center of k-space and/or faster acquisition times, among other qualities (see Section 2.7 for more properties of different k-space trajectories). In addition to faster scan times, different trajectories have added advantages, such as the motion-insensitivity of radial trajectories, the ability to perform motion correction using the PROPELLER trajectory [31, 32], or the possibility of obtaining spectral information using the rosette [35] or stochastic [38] trajectory. Another feature of non-Cartesian trajectories is that undersampling leads to incoherent artifacts, in contrast to Cartesian trajectories, where the undersampling artifacts are clearly identifiable. These properties make such sampling schemes highly desirable, and the combination of non-Cartesian trajectories and parallel imaging would be ideal in cases where a further reduction in scan time is advantageous.

However, as seen in Section 3.2.1, most k-space-based parallel imaging reconstruction techniques, such as Generalized Autocalibrating Partially Parallel Acquisitions (GRAPPA), necessitate the acquisition of regularly sampled Cartesian k-space data in order to reconstruct a non-aliased image efficiently. Unfortunately, non-Cartesian sampling schemes have the disadvantage that the points acquired generally do not lie on a grid and have complex k-space sampling patterns. Several authors have shown that the combination of GRAPPA and radial [100] and spiral [102, 126] trajectories can be performed more efficiently, although these techniques require radial segmentation of the data before the application of the reconstruction algorithm and large amounts of auto-calibration data (ACS), as described in Section 3.3.1. 1D non-Cartesian [103] and zig-zag [33] trajectories can also be combined with GRAPPA, although they require weight interpolation and segmentation, respectively, due to their lack of Cartesian symmetry. However, some trajectories, such as rosette [35] or TwiRL [135], do not possess the symmetry required for the segmentation approach; this implies that neither Cartesian nor non-Cartesian GRAPPA techniques are able to reconstruct images from undersampled data along these trajectories.

There are other techniques besides GRAPPA which can be used to reconstruct undersampled non-Cartesian datasets. Pruessmann et al. have demonstrated a method for the reconstruction of undersampled non-Cartesian trajectories using the SENSE algorithm in combination with a conjugate-gradient iteration approach [104]. Other authors have shown that the combination of other parallel imaging methods and non-Cartesian trajectories is possible, although time-consuming due to the large system of linear equations that must be solved [105, 106, 107]. However, these methods suffer from either computational complexity, long reconstruction times, a large number of optimizable parameters, or the need for coil sensitivity maps.

This chapter introduces a simple, novel method for performing Cartesian GRAPPA reconstructions on undersampled non-Cartesian k-space data gridded using GROG (GRAPPA Operator Gridding) to arrive at a non-aliased image. This method is known as pseudo-Cartesian GRAPPA. In contrast to other GRAPPA methods, the method proposed here does not rely on the radial symmetry of non-Cartesian trajectories, but instead uses several Cartesian GRAPPA patterns to reconstruct a non-aliased image using gridded undersampled datasets [127, 128, 129]. In addition, no coil sensitivity maps or additional parameters are required, which reduces the computational complexity of the reconstruction. The primary advantage of this pseudo-Cartesian method over previously proposed non-Cartesian reconstruction schemes is its similarity to the typical Cartesian GRAPPA procedure and general applicability to many different non-Cartesian sampling schemes. This flexibility in terms of both the appearance and number of patterns allows this pseudo-Cartesian GRAPPA to be used with undersampled datasets acquired with any non-Cartesian trajectory, making this reconstruction scheme more general than other non-Cartesian GRAPPA reconstruction techniques. A comparison of this novel method with non-Cartesian GRAPPA and CG-SENSE is described here, and the successful implementation of the pseudo-Cartesian GRAPPA using several different trajectories, including radial, rosette, spiral, 1D non-Cartesian, and zig-zag trajectories, is demonstrated. This method has been published as a Full Paper in the journal *Magnetic Resonance in Medicine* [127].

6.2 Theory

6.2.1 Cartesian and non-Cartesian GRAPPA

Before describing the use of GRAPPA for the reconstruction of non-Cartesian datasets, a simple 2D Cartesian case is reviewed. A more detailed treatment of GRAPPA is given in Section 3.2.1. In general, an undersampled dataset is reconstructed using coil weighting factors derived from a reference dataset, or Auto-Calibration Signal (ACS), which is usually a low-resolution portion of k-space sampled at the Nyquist rate. Representative examples of both for an acceleration factor $R=4$ and a single coil are shown in Figure 6.1 (the remaining coils, which have the same data structure, have been omitted from the figure for simplicity). An appropriate GRAPPA pattern, in this case, a 2×5 kernel (where two points in the phase encoding direction and 5 points in the read direction are selected for the pattern), could be as shown in the inset of Figure 6.1. Using this pattern, weighting factors are obtained on a coil-by-coil basis by examining the relationship between the acquired points from all coils and a single




Figure 6.1: A schematic Cartesian $R=4$ acquisition scheme shown for a single receiver coil for simplicity. The black circles show data that have been acquired, and the empty circles data that must be reconstructed using the GRAPPA algorithm and a 2×5 kernel. The fully-sampled central portion of k-space is used as the ACS, and it is important that the undersampled portions maintain the same symmetry. In the smaller figure to the right, the GRAPPA source points (black) are shown, as well as the points to be reconstructed (gray) using these source points.

missing point in the ACS. These weighting factors are subsequently applied to the source points in the undersampled portion of the data to find the values of the target points. It is important to note that these weights can be applied everywhere in k-space where this pattern appears because they describe the general harmonic shift for the given arrangement of source and target points, which does not change in different areas of k-space. Because the undersampled points fall on a regular pattern in the Cartesian dataset, i.e. are highly symmetric, only one GRAPPA pattern with three target points, and thus one set of coil weighting factors, is needed to reconstruct the missing points.

Due to the fact that the GRAPPA technique necessitates matching patterns in both the undersampled source data and the reference data, most common non-Cartesian GRAPPA methods make use of the radial symmetry inherent in the trajectory. Radial GRAPPA and spiral GRAPPA, the most common non-Cartesian GRAPPA methods, have been described in Section 3.3.1. In these methods, the undersampled radial or spiral data are first reordered onto a new Cartesian-like grid, where the axes are the read-out (r) and projection angle (θ) [100, 102, 126]. The missing data are then reconstructed using a segmented GRAPPA procedure which calculates the coil weights for different portions of k-space separately; a simple Cartesian GRAPPA method cannot be used because the distance and direction between sampled points changes depending on their location in k-space. Finally, the entire reconstructed Cartesian-like data must be gridded to arrive at the final k-space, upon which a fast Fourier transformation is performed to obtain an image. It is clear that these methods necessitate a regular radial symmetry as well as several different GRAPPA kernels in order to reconstruct the missing points in the appropriate locations, and can thus be used only with specific trajectories.

6.2.2 Pseudo-Cartesian GRAPPA

The following sections describe a more general case of non-Cartesian GRAPPA, namely pseudo-Cartesian GRAPPA. This method is made up of three distinct steps:

- Gridding the acquired data onto a Cartesian grid using GROG
- Determination of the appropriate pseudo-Cartesian patterns
- Reconstruction of the missing points using GRAPPA and the selected patterns

Each of these three steps is discussed below.

GRAPPA Operator Gridding (GROG)

GRAPPA Operator Gridding (GROG), described in Chapter 4 and [118], is a method which uses parallel imaging to grid non-Cartesian datasets. GROG is a special case of GRAPPA where the number of source and target points is the same, usually one. GROG works by using the appropriate root of the base weight sets to shift each non-Cartesian point to the nearest Cartesian location on a grid using the following equation for a 2D image:

$$\vec{S}(k_x + \delta_x, k_y + \delta_y) = \hat{G}_x^{\delta_x} \cdot \hat{G}_y^{\delta_y} \cdot \vec{S}(k_x, k_y) \quad (6.1)$$

where $\vec{S}(k_x, k_y)$ is a vector containing the non-Cartesian signal from each of the receiver coils (1 to NC , where NC is the number of coils) at the k-space location $[k_x, k_y]$ to be shifted by an amount $[\delta_x, \delta_y]$; \hat{G}_x and \hat{G}_y are the appropriate GROG weight sets of size $NC \times NC$ for unit shifts in the x and y directions, respectively; and $\vec{S}(k_x + \delta_x, k_y + \delta_y)$ is the shifted signal at the Cartesian location $[k_x + \delta_x, k_y + \delta_y]$. The GROG weights \hat{G}_x and \hat{G}_y can be determined either from the non-Cartesian data themselves (see Section 5.2.2) or using a separate low-resolution Cartesian dataset (see Section 4.2.3).

It is important to note that GROG alone cannot be used to reconstruct undersampled non-Cartesian k-space data because only one source point (non-Cartesian) is used for each target point (Cartesian). The small number of source points implies that only small shifts can be accomplished; depending on the coil array, shifts of more than $\Delta k = 0.5$ in each direction can lead to errors which result in artifacts in the image domain. This limitation does not generally hamper gridding, because the largest shift required to grid is $\Delta k = 0.5$, although it means that GROG cannot be used alone for the reconstruction of images from undersampled non-Cartesian datasets.

GROG is advantageous for gridding because the process can be performed quickly, as only two weight sets are required for the needed k-space shifts, and no density compensation function is required. In addition, it has been shown in Chapter 4 that GROG can be used to grid fully-sampled datasets without a visible increase in the noise in the final image when working with clinical coil arrays. Specifically, symmetric array coils with six or more elements can be used to perform GROG shifts of $\Delta k = 0.5$ in both logical directions without significant noise enhancement; thus, only such coils were employed in this work.




Figure 6.2: : An $R=4$ radial dataset gridded with the standard convolution gridding (left) and GROG (right). The center of k -space is correctly gridded using both methods because the Nyquist criterion is fulfilled in this area. However, as soon as the region of support of the data is less than Δk , i.e. the Nyquist criterion is no longer met, convolution gridding spreads acquired points into areas of k -space where no data have been acquired (the dark areas of k -space between the acquired radial arms) by the convolution with the gridding kernel (see Section 2.8). GROG, however, reconstructs only those grid points nearest to actually acquired k -space points. Thus, undersampled data gridded using convolution gridding cannot be reconstructed using GRAPPA, as there are no “holes” in the k -space to be filled. In addition, due to the insufficient region of support, even the grid points along the acquired spokes are incorrect. However, GROG is ideal for use with GRAPPA due to the fact that the areas of k -space which must be reconstructed are easily identifiable (i.e. the “holes”), and the gridded datapoints are purely Cartesian (see Chapter 4).

However, in pseudo-Cartesian GRAPPA, GROG is used for a different reason. Most other gridding techniques, including the standard convolution gridding method [76], assume that the non-Cartesian data fulfill the Nyquist criterion at all locations in k-space. When this is not the case, gridding errors appear. GROG, however, only shifts each non-Cartesian data point to its nearest Cartesian neighbor using parallel imaging weight sets. This procedure leaves empty spaces where the shift necessary to reach the next Cartesian data points exceeds the above mentioned value of $\Delta k = 0.5$. In comparison, no empty spaces are left in the undersampled k-space data if a conventional gridding kernel is used. Due to the convolution of the non-Cartesian points with this gridding kernel, a single non-Cartesian data point is spread out over several Cartesian locations. Thus, pseudo-Cartesian GRAPPA, which requires these “holes” in k-space for the reconstruction, would not be possible using conventional gridding methods. The difference between an undersampled dataset gridded using convolution gridding and GROG can be seen in Figure 6.2.

Determination of Patterns

Unfortunately, the gridded undersampled data do not contain one single Cartesian pattern that is repeated regularly throughout the data and that can be used with the standard Cartesian GRAPPA algorithm, i.e. there is no simple Cartesian symmetry in the non-Cartesian dataset. In spiral data, for instance, the 2×5 GRAPPA kernel forms a pattern that would be optimal for a pie-shaped sector running 0° from the center of k-space would be completely unsuccessful for the sector along the 90° radius (Figure 6.3, right), because the selected Cartesian pattern does not exist in these data. For this reason, several simple Cartesian GRAPPA patterns must be employed to cover the entire undersampled dataset, such as those shown at the bottom of Figure 6.4. Thus, the next step in the pseudo-Cartesian GRAPPA reconstruction is to determine the patterns needed to reconstruct the missing points in the gridded, undersampled dataset. There are many different possibilities for determining these patterns, and the method used to reconstruct the images in this work is explained here. This method is used because it quickly finds appropriate patterns for the trajectories employed in this chapter, but it would not necessarily find the optimal patterns for every undersampled trajectory (and there are surely more efficient methods of determining the appropriate and optimal patterns). It is also important to note that the sole function of the pattern selection process is to determine the patterns for which the GRAPPA weight sets must be determined; weights could be calculated for all of the possible patterns, although most of the weight sets would not be needed for the reconstruction.

A constant number of source points to be used was chosen; for all reconstructions shown here, this number was set to six. More source points could be employed, although this would lead to longer reconstruction times due to the increased number of weights and the increased time needed to apply them, as well as an increase in the number of patterns required. Starting with a standard 2×3 Cartesian kernel, such as the R=3 pattern shown at the top left of Figure 6.4, each source point is assigned an index variable which can be set to either 0, -1 or 1. Then the source points (black circles) are shifted according to the values of the index for each respective point. First, the lowest value indices [-1 -1 -1 -1 -1 -1] are assigned to each point (as denoted by the lower case letters by the R=3 base pattern

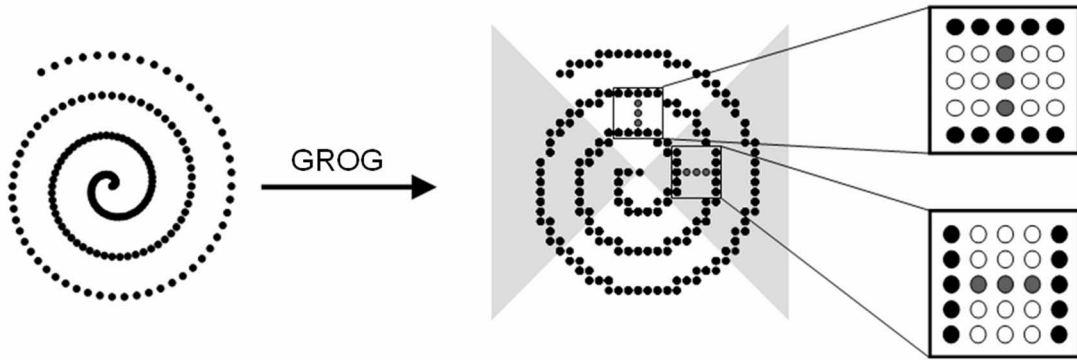


Figure 6.3: : Left: A typical $R=4$ spiral sampling scheme. It is clear that the acquired points do not fall on a Cartesian grid, and must be gridded before the FFT can be calculated. Middle: The gridded spiral data, resampled using GROG. Portions of this data are pseudo-Cartesian; for example, the simple Cartesian GRAPPA 2×5 kernel for $R=4$ (far right, top) can be used to reconstruct some missing points in the white areas of the under-sampled k-space. This pattern would be inappropriate for the reconstruction of missing points in the gray shaded area, and a different pseudo-Cartesian GRAPPA pattern, such as the one on the far right, bottom, must be selected.

in Figure 6.4). The indices for the top points are added to their locations, the indices for the lower points are subtracted from their locations (such that the index -1 always indicates motion towards the target points). This leads to the second pattern shown at the top of Figure 6.4. For this pattern, the possible source point is indicated by the white triangle. This pattern is then applied to a single channel of the masked out undersampled k-space in order to see if missing points (white triangles) can be reconstructed with this pattern of source points. All points that can be reconstructed using this pattern are marked as reconstructed, and then the pattern is rotated 90° and the same process is performed. The $R=3$ pattern with indices $[-1 -1 -1 -1 -1 -1]$, however, does not allow one to reconstruct any missing points because the target point is the same as one of the source points, and therefore would not be selected for weight calculation. The indices are then increased to $[0 -1 -1 -1 -1 -1]$, and the new pattern is examined in order to determine if it allows for the reconstruction of missing points. As an additional example, the patterns created using the indices $[-1 0 1 0 0 0]$ and $[1 1 1 0 0 0]$ and the $R=3$ base pattern are shown as the third and fourth patterns at the top of Figure 6.4.

As in all non-Cartesian parallel imaging, defining the actual acceleration factor is not trivial. For the purposes of pseudo-Cartesian GRAPPA, the nominal acceleration factor is determined by the difference between the top center and lower center point in the pattern. Thus, for the patterns shown at the top of Figure 6.4, the nominal acceleration factors are $R=3$, $R=1$, $R=3$, and $R=4$, although these patterns have been created using the Cartesian $R=3$ as a basis. It is important to note that a maximum acceleration factor R_{max} must be selected by the user. This maximum acceleration factor is based on the coil geometry, and the undersampling of the trajectories must be performed with this maximum acceleration factor in mind so that larger “holes” are not present in k-space. In order to assure that the

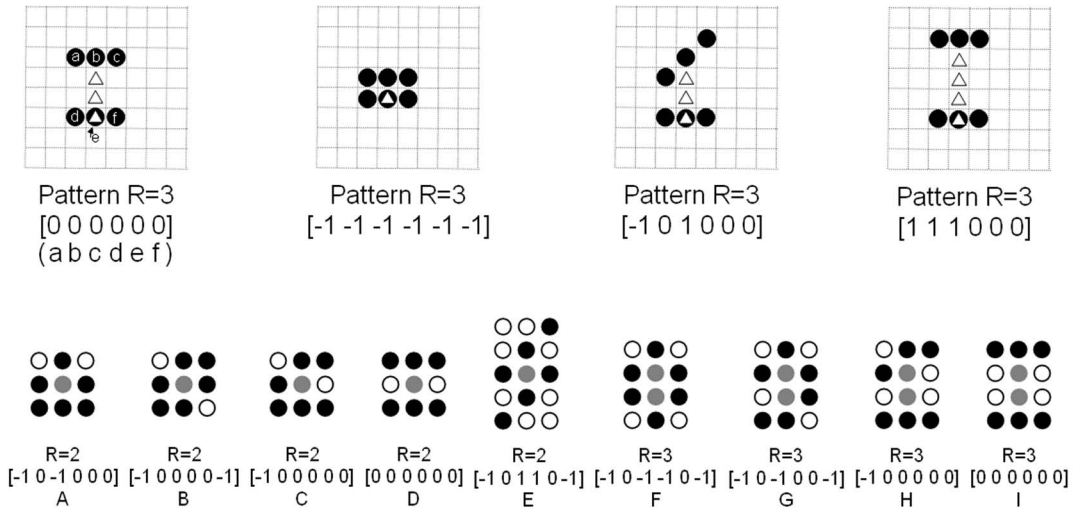


Figure 6.4: Examples of patterns used in a pseudo-Cartesian GRAPPA reconstruction. The top portion of the figure shows an example of the basic R=3 pattern (far left), and the letters a through f indicate the source point with which each index value is associated. Additional patterns generated by the search algorithm and index values $[-1 -1 -1 -1 -1 -1]$, $[-1 0 1 0 0 0]$, and $[1 1 1 0 0 0]$ are also shown (top left, right, far right). The black circles denote the source points, and the triangles are target points. It should be noted that although these patterns are generated using the basic R=3 pattern, they have different nominal acceleration factors (namely R=3, R=1, R=3, and R=4). The bottom portion of the figure shows several different patterns (source points shown in black and target points in gray) and their indices in the order they would be applied in a pseudo-Cartesian GRAPPA reconstruction.

pattern search algorithm does not select patterns with larger nominal acceleration factors, the indices run only from -1 to 0 for R_{max} and $R_{max} - 1$, such that the larger acceleration factors that arise due to the index 1 are avoided. When performing the pattern selection, all of the R=2 patterns are examined first, followed by the R=3 patterns, up until the maximum acceleration factor established for the reconstruction is reached. This prevents using patterns with higher nominal acceleration factors to reconstruct missing points which could be reconstructed with a pattern with a lower acceleration factor.

As seen in the previous discussion, for some arrangements of patterns and source points, such as the $[0 0 0 -1 -1 -1]$ arrangement for the standard R=2 GRAPPA pattern, no points can be reconstructed because one of the source points is the same as the target point. Such patterns were discarded. Similarly, some patterns are identical, such as the R=2 $[1 1 1 0 0 0]$ pattern and the R=3 $[0 0 0 0 0 0]$ pattern. In this case, the first formulation of the pattern would be selected, and the second would not be needed for the later reconstruction.

Examples of several possible and sensible (i.e. the source and target points are not the same, and the patterns are not repeated) kernel patterns and the order in which they would be applied are shown at the bottom of Figure 6.4. It is important to keep in mind that the 90° rotation of each of these patterns is checked before moving on to a new pattern. For an undersampled radial dataset, patterns C, D, E, H, and I would be selected, as they allow different missing points to be reconstructed. Patterns A, B, F, and G would not exist in such a gridded dataset, and would not be selected. However, in order to allow this implementation to be used for all other trajectories, such patterns must be checked.

Reconstruction of the Missing Points

Once the appropriate patterns for the reconstruction have been determined, a modified GRAPPA technique with adapted kernels is applied using these patterns to reconstruct the missing points and arrive at the reconstructed k-space. An example of an undersampled dataset gridded using GROG is shown in Figure 6.5 (center, top); the “holes” where the data are undersampled can be clearly seen. These missing spaces in a k-space gridded with GROG can be reconstructed using a standard GRAPPA method. This involves generating the weight set needed for the specific pattern using an ACS dataset and the standard GRAPPA method, and then applying the weight set to the undersampled dataset gridded with GROG. For some of the datasets, such as the radial and 1D non-Cartesian data, the weights can be calculated using the fully-sampled central portion of k-space. It should be noted that the reconstruction time directly depends on the number of patterns chosen; the use of fewer patterns means that less time must be spent calibrating the weight sets and searching for the appropriate k-space locations where the pattern must be applied. For this reason, the reconstruction of trajectories which were made up of more regular undersampling patterns, such as the zig-zag trajectory and 1D non-Cartesian trajectory, were much less time consuming than reconstructions using the radial, spiral, or rosette trajectories in this implementation.

6.3 Methods

6.3.1 Data Reconstruction

All fully-sampled and undersampled datasets shown in this manuscript were gridded using GROG. As stated above, this method has the advantage that single non-Cartesian points are shifted to the nearest Cartesian point using the GRAPPA Operator, leaving empty spaces in k-space where data must be reconstructed. For the radial, spiral, and 1D non-Cartesian datasets, the GROG weights were determined from the data themselves using the self-calibrating GROG method described in Chapter 5; a separate low-resolution Cartesian k-space was used to calibrate the weights for the rosette and zig-zag trajectories.

After the data were gridded, appropriate patterns for each undersampled trajectory were determined using the algorithm described above. In order to avoid leaving gaps in the reconstructed k-space, a large number of patterns was chosen (i.e. between 4 to 102, depending on the trajectory and undersampling factor, see Table 6.1). Each pattern was arranged in a different configuration with six source points (see

Figure 6.4 for examples of these patterns). Once the patterns were determined, the pseudo-Cartesian GRAPPA reconstruction was performed. First, the weight set for each pattern is found using the ACS in the same way that the weight sets are derived for GRAPPA, and then this weight set is applied in those places in the undersampled k-space where the specific pattern fits. For the radial dataset, the ACS was taken from the central fully-sampled portion of the gridded undersampled k-space; for the remaining datasets, the central portion of the fully-sampled and gridded data was used.

6.3.2 Simulations

In order to test the feasibility of the method described above, the standard Shepp-Logan phantom was employed with a simulated eight-element one-ring head coil array, where sensitivities were derived using an analytic integration of the Biot-Savart equations. The Cartesian data were resampled as radial data (50 projections, 256 read-out points, base matrix of 128×128) by sinc interpolation, which corresponds to an undersampling factor of 4 in comparison to the fully-sampled case. The GROG weights were determined from the radial arms themselves using the self-calibrating GROG method of Chapter 5, and the undersampled radial data were gridded. Using the pattern-finding algorithm described above, it was determined that the use of 67 different pseudo-Cartesian GRAPPA patterns would allow for a complete reconstruction of the missing k-space points given the arrangement of the undersampled data points. Examples of these patterns are given in Figure 6.4. The reconstruction was then carried out using the 67 patterns. As can be seen in Figure 6.5 (center, top), the central portion of k-space is fully sampled, and this block (25×25) can be used to determine the weights for the pseudo-Cartesian GRAPPA patterns. A complete depiction of this process can be found in Figure 6.5.

In addition, a standard radial GRAPPA reconstruction [100] with 8 read-out and 10 angular segments was performed using the same undersampled radial dataset (and the fully-sampled dataset as the ACS) in order to investigate the differences in the images between the standard and proposed reconstruction techniques. For the standard radial reconstructions shown here, the fully-sampled radial k-space data were used for the calibration, although self-calibrating methods have been proposed for radial GRAPPA [101]. The use of such a method is not expected to significantly change the image quality in this reconstruction. A conjugate-gradient SENSE (CG-SENSE) reconstruction [104] was also performed using 30 iterations of the conjugate gradient loop. Because this method requires a coil sensitivity map, such a map was derived from the radial dataset using array correlation statistics [94]. The CG-SENSE algorithm yields a single image instead of multi-channel images, and thus the multi-channel images resulting from the pseudo-Cartesian GRAPPA and radial GRAPPA reconstructions, as well as the reference image, were combined using R=1 SENSE [3] for comparison purposes. After the reconstructions were performed, the results were subtracted from the reference image to yield difference images which highlight the reconstruction errors.

Finally, an undersampled simulated rosette dataset was generated by resampling an in vivo Cartesian head image along this trajectory using sinc interpolation in order to investigate the performance of pseudo-Cartesian GRAPPA for such a trajectory. The head coil array employed was the same as that used for the radial simulations. The complete dataset was made up of 120 phase encoding steps, each

	Radial	Rosette	Spiral	1D NC	Zig-Zag
Acceleration	4	4	2.7	3	3
Patterns for Reco	67	92	102	4	5
Number of Coils	8	8	32	12	12
Matrix Size	128 x 128	128 x 128	256 x 256	256 x 256	256 x 256
Phase ^a	50 / 200	30 / 120	3 / 8	88 / 256	88 / 256
Read	256	360	6671 ^b	256	256
Scanner	NA	NA	3T Trio	3T Trio	1.5T Avanto

Table 6.1: Reconstruction parameters for the simulated and in vivo images shown in this chapter. ^aThe first phase encoding number indicates the number of “phase encoding” steps used for each accelerated trajectory, and the second is the number of phase encoding steps needed to fulfill the Nyquist criterion. ^bA center-out constant-linear-velocity spiral trajectory was used.

with 360 read-out points, and the R=4 accelerated rosette dataset was created by reducing the number of phase encoding steps to 30. A total of 92 patterns were employed for the pseudo-Cartesian GRAPPA reconstruction, including those shown in Figure 6.4. Both the GROG weights and the weights for the pseudo-Cartesian patterns were derived from a separate low-resolution (25×25) Cartesian k-space.

6.3.3 In vivo Experiments

Three images were acquired with two different scanners, a Siemens 1.5T Avanto and a Siemens 3T TIM Trio (Siemens Medical Solutions, Erlangen, Germany), in order to test the pseudo-Cartesian GRAPPA reconstruction with different non-Cartesian trajectories, namely spiral, 1D non-Cartesian, and zig-zag. Informed consent from the volunteers was obtained before each study. Information about the data acquisition parameters, scanners, and coils used is shown in Table 6.1.

The spiral trajectory was used to acquire an image of the brain accelerated by a factor of 2.7 (3 of 8 spiral arms needed to fulfill the Nyquist criterion were used). The GROG weights were calculated using a fully-sampled spiral dataset, and the 102 different pseudo-Cartesian GRAPPA weights used were determined using a 30×30 portion from the center of the fully-sampled k-space. As stated above, each pattern was made up of at least six source points; these patterns were similar to those shown in Figure 6.4 and specifically, patterns D, E, H, and I were actually employed for the reconstruction. In addition to the pseudo-Cartesian GRAPPA reconstruction, a CG-SENSE reconstruction was also performed, where the coil sensitivity map was generated using array correlation statistics [94]. In this CG-SENSE implementation, 50 iterations were performed and a grid oversampling factor of 4 was employed.

An image of a phantom was acquired using the 1D non-Cartesian trajectory (shown in Figure 2.7) accelerated by a factor of 3. The central portion of k-space was completely acquired (yielding a matrix size of 21×256), and was used for the calibration of both the GROG and the pseudo-Cartesian weights. Only 4 patterns were required, as the gridded k-space is purely Cartesian in the y-direction.

In vivo data were acquired using the R=3 zig-zag trajectory [33] that blipped over approximately $2\Delta k_y$ in 32 cycles of 16 points. These data were then gridded using GROG, where the weights were determined using a low-resolution k-space (40×40 vs. the original 256×256 data matrix). Pseudo-Cartesian GRAPPA with 5 patterns was employed to reconstruct the missing data points.

6.4 Results

6.4.1 Simulations

The k-space and image reconstructed using R=4 undersampled radial data and the pseudo-Cartesian GRAPPA method are shown in Figure 6.5 (center right). The root-mean-square-error (RMSE) of the reconstructed image as compared to the fully-sampled image was calculated to be 0.06%, which is similar to that of the radial GRAPPA image (0.08%, Figure 6.1, center, far right), and lower than that of the CG-SENSE image (0.47%). Each reconstruction technique yielded excellent results, as can be seen in the RMSE values as well as the quality of the reconstructed images. In addition, the difference images at the bottom of Figure 6.5 show that the pseudo-Cartesian GRAPPA method yields a more homogeneous noise distribution in the reconstructed image, as compared to the radial GRAPPA reconstruction, where the differences are concentrated around the edges. This indicates that the radial GRAPPA image is blurrier than the pseudo-Cartesian GRAPPA image, more than likely due to errors in the outer portions of k-space, although it has less noise enhancement overall. The CG-SENSE reconstruction is significantly worse than the two GRAPPA reconstructions, which is due to errors in the determination of the coil sensitivity map.

The results seen in the difference images are to be expected. The pseudo-Cartesian reconstruction yields more overall noise due to errors from the GROG gridding which are propagated in the pseudo-Cartesian GRAPPA reconstruction. However, radial GRAPPA, which approximates segments of the radial data as Cartesian blocks, leads to reconstruction errors in portions of k-space where this approximation is less accurate. In addition, calibration in areas of low signal, such as the outer portions of k-space, tends to be less stable, potentially resulting in high frequency artifacts in radial GRAPPA due to the larger error at the edges of k-space. The CG-SENSE method, which yields an optimal least-mean-squared error reconstruction, is limited by the quality of the sensitivity map. In addition, because this method yields a single image, and not coil-by-coil reconstructions, the results obtained are more difficult to compare to the multi-coil reference image.

It is important to note that the most time consuming portion of the reconstruction is the determination of the patterns to be used and the k-space locations where they must be employed. However, the pattern determination can be performed independently of the reconstruction, given that the trajectory and undersampling factor to be used are known. Thus, only the actual determination of the weights and the reconstruction are relevant for the speed of the algorithm. For the Matlab implementations used here, the pseudo-Cartesian reconstruction required approximately four times the calculation time of the radial GRAPPA algorithm. This is due to the larger number of patterns required, as well as the fact that the pattern for radial GRAPPA is much easier to work with in Matlab than the unusual patterns




Figure 6.5: A schematic of pseudo-Cartesian GRAPPA with GROG. The fully sampled k-space and image are shown on the left top and bottom as comparison. The R=4 under-sampled radial data are first gridded with GROG, yielding the gridded k-space shown in the center, top and the undersampled image (center, bottom). The undersampled k-space data is reconstructed to generate the fully-sampled k-space (right, top) using the pseudo-Cartesian GRAPPA method explained in the text. As a comparison, images reconstructed using radial GRAPPA and CG-SENSE are shown on the far right-hand side in the center of the figure. The bottom row shows the results of subtracting the reconstructed images from the reference image (12.5 times magnified).

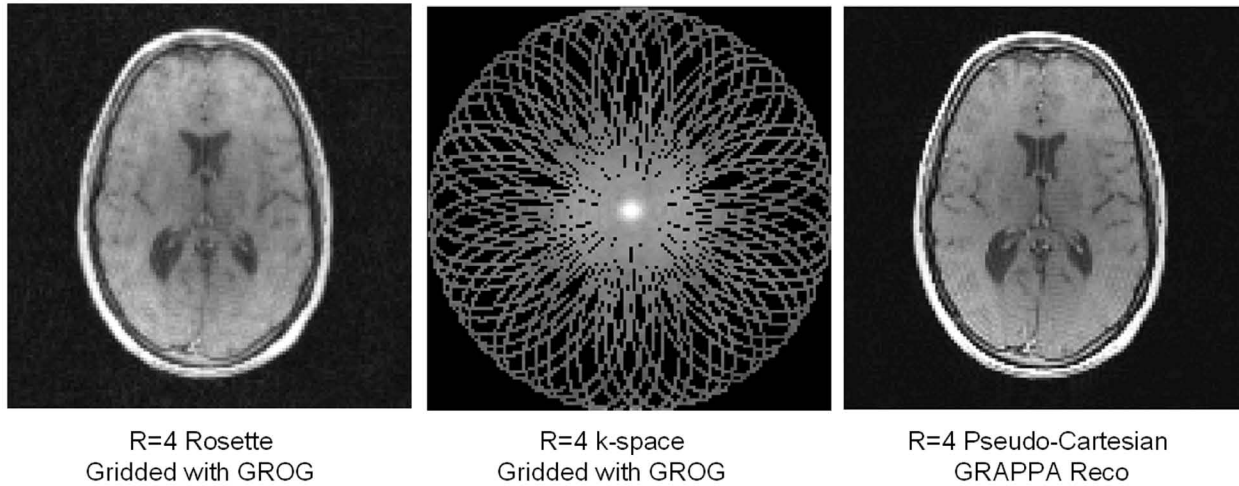


Figure 6.6: The R=4 undersampled simulated rosette image (left) and the corresponding k-space gridded using GROG (middle). After pseudo-Cartesian reconstruction, the image shown on the right results. This image has a RMSE of 1.73% in comparison to the fully-sampled image.

needed for the pseudo-Cartesian reconstruction. The CG-SENSE algorithm required approximately the same amount of time as the pseudo-Cartesian reconstruction due to the need to grid, fast Fourier transform, and degrid the non-Cartesian data for each iteration. The calculation of the sensitivity map is not included in this time estimate.

The reference, R=4 undersampled k-space, and pseudo-Cartesian reconstruction of the rosette dataset are shown in Figure 6.6. As can be seen in the center of this figure, the undersampled image is both blurred and mottled by irregular undersampling at the outer portions of k-space. After the pseudo-Cartesian reconstruction, this blurring is removed, and the mottling greatly reduced. The RMSE of the reconstruction in comparison to the reference image is 1.73%. The reconstruction of undersampled rosette data cannot be accomplished with other GRAPPA techniques due to their irregular undersampling pattern; thus, no standard reconstruction can be compared to this pseudo-Cartesian reconstruction.

6.4.2 In vivo Experiments

The results of the reconstructions performed with experimental data are shown in Figures 6.7 through 6.9. Figure 6.7 (left and middle) shows the image and k-space which result from gridding the 2.7-times accelerated spiral data with GROG, respectively. After the application of pseudo-Cartesian GRAPPA with 102 patterns, the aliasing artifacts are no longer evident. There is a slight loss in SNR in the center of the image due to the extreme fold-over artifacts in this region. The CG-SENSE reconstruction still contains significant artifacts in the center of the image; these artifacts are most likely due to errors in the determination of the coil sensitivity map. No improvement was seen when increasing the number




Figure 6.7: The $R=2.7$ undersampled in vivo spiral image (left) and the corresponding k-space gridded using GROG (middle left). After pseudo-Cartesian reconstruction with 102 patterns, the undersampling artifacts are no longer visible (middle right). As a comparison, the CG-SENSE algorithm was also used for the reconstruction of the undersampled spiral data; the resulting image (far right) contains residual artifacts due to errors in the coil sensitivity map.

of conjugate gradient iterations.

Figure 6.8 (left and middle) shows the image and k-space which result from gridding the three-fold accelerated 1D non-Cartesian data with GROG, respectively. This trajectory is interesting because the non-Cartesian points must be shifted only in the y-direction, and the resulting gridded data is purely Cartesian. The acceleration factor takes on a value of one in the center of k-space and then increases up to a value of 5 at the edges. Thus, for the pseudo-Cartesian reconstruction, only 4 different patterns are required, and all are typical Cartesian patterns for acceleration factors 2 to 5. As can be seen on the right in Figure 6.8, the application of pseudo-Cartesian GRAPPA removes the aliasing artifacts, although there is a slight noise increase in the center of the image.

Figure 6.9 (left and middle) shows the image and k-space which result from gridding the three-fold accelerated zig-zag data with GROG, respectively. After the application of pseudo-Cartesian GRAPPA with 5 patterns, the blurring and ghosting artifacts evident in the image on the left are removed, yielding the unaliased image (right).

6.5 Discussion

The pseudo-Cartesian GRAPPA reconstruction method has been shown to allow the reconstruction of unaliased images from undersampled non-Cartesian datasets. While examples using the radial, rosette, spiral, 1D non-Cartesian and zig-zag trajectories have been shown, this method can be used to reconstruct images from any arbitrary trajectory. The reconstruction method requires only that the coil array used has sufficient sensitivity variations to mimic the missing spatial harmonics, as in all other parallel imaging reconstruction techniques. By first gridding the undersampled data using GRAPPA




Figure 6.8: The R=3 undersampled phantom 1D non-Cartesian image (left) and the corresponding k-space gridded using GROG (middle). The undersampling factor starts at R=1 in the center of k-space and increases to R=5 at the edges, which can be seen in the middle image. After pseudo-Cartesian reconstruction using 4 patterns, the blurring and ghosting artifacts are no longer visible (right).

Operator Gridding (GROG), only the Cartesian points nearest the acquired non-Cartesian points are reconstructed, leaving “holes” which can later be filled in using the pseudo-Cartesian GRAPPA reconstruction. Other gridding methods, such as the standard convolution gridding, assume implicitly that the data to be gridded fulfills the Nyquist criterion, and errors can occur when gridding undersampled data with such methods. In addition, these methods do not leave the “holes” in k-space necessary for the application of pseudo-Cartesian GRAPPA. Thus, the gridding of the undersampled non-Cartesian data with GROG is an essential part of the pseudo-Cartesian GRAPPA reconstruction.

However, it is important to note that because the algorithm described here uses two different types of parallel imaging during the course of the reconstruction, an additional noise amplification can be expected over traditional GRAPPA reconstructions, i.e. any errors in the GROG gridded data will propagate due to the use of the pseudo-Cartesian GRAPPA algorithm. For this reason, receiver coils that exhibit sufficient sensitivity variations must be used, such that the GROG weights allow for accurate gridding. Although it has been shown in Chapter 4 that standard clinical coils allow the use of GROG without significant noise amplification, these potential errors stemming from GROG must be kept in mind. This noise amplification depends on a number of factors, including the coil arrangement, trajectory, and undersampling factor, and is difficult to characterize in general. In order to combat such effects, the pseudo-Cartesian reconstruction was only performed with patterns that had at least six source points. In addition, work on a version of GROG which uses multiple non-Cartesian points for the gridding process is currently underway, which could potentially improve the reconstruction process.

After the data have been gridded, standard Cartesian GRAPPA with different patterns can be used to reconstruct the data points missing in the undersampled gridded k-space. An appropriate choice




Figure 6.9: The R=3 undersampled in vivo zig-zag image (left) and the corresponding k-space gridded using GROG (middle). After pseudo-Cartesian reconstruction using 5 patterns, the image shown on the right results.

of patterns can be determined using a pattern-finding algorithm, which automates the process. In addition, when multiple images are acquired using the same trajectory, the patterns are the same for all of those k-space datasets and must be determined only once. Thus, the patterns are independent of the object to be imaged and the coil used; they depend only on the underlying trajectory. It is important to note that the patterns employed in this work used only the original acquired and then gridded data points as source points for the reconstruction, which means that a relatively large number of patterns must be used. Iterative methods, which allow the use of points reconstructed in one step to be used as source points in subsequent reconstruction steps, would require fewer patterns, but could lead to unexpected reconstruction artifacts due to error propagation (i.e. from employing a reconstructed point as a target point for further reconstructions). For this reason, no iterative processes were employed here.

Each pattern used in the reconstructions shown in this chapter was made up of at least six source points. Reducing the number of source points for each pattern would also reduce the number of patterns necessary for the reconstruction, but would also lead to a reduced reconstruction quality. Increasing the number of source points could improve the reconstructed image, although more patterns, and thus reconstruction time, would be required. Many Cartesian and non-Cartesian GRAPPA reconstruction algorithms employ six source points per coil for each target point, which is why this value was chosen for the pseudo-Cartesian reconstructions. In addition, the patterns for the reconstructions shown here were chosen such that all missing points in k-space were reconstructed. A method which allows selection of the optimal patterns, i.e. patterns which employ the optimal source points for the reconstruction of each specific missing point, has yet to be developed, and could further improve the image quality.

In addition to the patterns, an ACS dataset is also needed for the calibration of the weight sets. For some trajectories, namely radial, dual-density spirals, 1D non-Cartesian, or PROPELLER, both the GROG and the pseudo-Cartesian GRAPPA portions of the reconstruction are self-calibrating; for other trajectories, a low-resolution calibration dataset, identical to that used in Cartesian GRAPPA, is needed for the calibration steps. Alternatively, for some trajectories, a dual-density implementation could also be used for the weight determination for both GROG and the pseudo-Cartesian reconstruction; for instance, the central Nyquist-sampled portion of the rosette dataset can be gridded and then used to determine the necessary weight sets.

The pseudo-Cartesian GRAPPA method proposed here is much more general than other non-Cartesian GRAPPA methods. Such GRAPPA methods, i.e. radial GRAPPA [100], spiral GRAPPA [102, 126], 1D non-Cartesian GRAPPA [103], and zig-zag GRAPPA [33], rely on a high degree of symmetry in the undersampled data, and cannot be used for arbitrary trajectories (see Section 3.3.1). The segmentation approach has the additional disadvantage that weights for the outer portions of k-space are determined using these low SNR outer portions, and not the center of k-space, as in Cartesian GRAPPA. Thus, the weights for these segments are less accurate, leading to errors in the non-Cartesian GRAPPA reconstruction, especially in 3D reconstructions (see Section 3.3.1). In addition, other trajectories such as TwiRL or rosette cannot be reconstructed using these methods, as they lack the radial symmetry needed for segmented GRAPPA. The pseudo-Cartesian GRAPPA method, in contrast, can be applied to any undersampled non-Cartesian dataset, thus making it appropriate for a larger number of non-Cartesian trajectories. As the weights are always determined using the high SNR center of k-space, they tend to be more accurate than the outer k-space weights in segmented non-Cartesian GRAPPA. The method could also easily be extended to a third dimension, allowing for the reconstruction of unusual 3D non-Cartesian datasets. Thus, although the algorithm proposed here is not always as exact as previously demonstrated non-Cartesian GRAPPA reconstruction schemes, it is far more general and can be used in a variety of applications.

Another option for the reconstruction of arbitrary trajectories is CG-SENSE [104], which uses the knowledge of the coil sensitivity maps and the undersampled data to derive unaliased images. This method has three drawbacks: the first is that sensitivity maps are required, which can be difficult to generate from the non-Cartesian data, and the second is that the image is reconstructed using an iterative conjugate gradient method. Methods have been demonstrated which help determine the stopping criteria for CG-SENSE [110], although many effects influence the selection of this value, including the trajectory, undersampling factor, coil array used, preconditioning, gridding method, and initial image. Thus, a direct reconstruction method which requires few additional parameters, such as that proposed here, can be advantageous. In addition, in dynamic imaging the iterative reconstruction has to be performed for each time frame. In pseudo-Cartesian GRAPPA, the pattern and weight determination has to be performed only once and can then be applied to the following time-frames, leading to a fast data reconstruction.

A second reconstruction method which has been proposed for arbitrarily sampled data is PARS [105, 106], which performs a form of non-Cartesian SMASH using sensitivity maps and the appropriate spatial harmonics to determine the weight set needed to calculate each Cartesian point. The main drawback of PARS is that an extremely large number of weight sets must be calculated in order to reconstruct the image, as each non-Cartesian point generally requires a different weight set, which can become quite time-consuming. Also, similar to CG-SENSE, coil sensitivity maps must be derived, which can be difficult in some imaging scenarios.

A third method is BOSCO [136], which first grids the undersampled data using standard convolution gridding, and then convolves the data a second time with a GRAPPA kernel, leading to an unaliased image. BOSCO has only been demonstrated using spiral data, and suffers from the use of a single kernel, which could hamper the reconstruction of other undersampled trajectories. Pseudo-Cartesian GRAPPA suffers from none of these drawbacks. It is not iterative, requires neither a sensitivity map nor an inordinately large number of weight sets, and the number of kernels used for the reconstruction is not limited to one.

Finally, it is important to note that this method is based directly on the conventional Cartesian GRAPPA method. Because pseudo-Cartesian GRAPPA applies this standard GRAPPA to reconstruct undersampled data acquired along arbitrary trajectories, it is much more straightforward than other non-Cartesian GRAPPA methods, which require dedicated algorithms. The combination of its simplicity and general applicability to many non-Cartesian trajectories makes the pseudo-Cartesian method in conjunction with GROG demonstrated here advantageous.

6.6 Conclusion

The pseudo-Cartesian GRAPPA reconstruction method proposed here is a simple application of the basic Cartesian GRAPPA algorithm to non-Cartesian datasets. As the algorithm uses the Cartesian patterns in gridded undersampled data for reconstruction, no additional segmentation or interpolation is necessary. In addition, for radial acquisitions, dual-density spirals, 1D non-Cartesian, or PROPELLER data, no additional ACS data must be acquired as is needed for the standard GRAPPA reconstructions. However, the main strength of the pseudo-Cartesian algorithm is that it can be applied to any undersampled non-Cartesian dataset; this includes not only radial and spiral trajectories, but also trajectories which cannot be reconstructed using other non-Cartesian GRAPPA techniques, such as TwiRL and rosette. Because this algorithm can be applied to other non-Cartesian trajectories, the need for a separate reconstruction procedure for each accelerated non-Cartesian trajectory is eliminated.

Chapter 7

GROG-Facilitated Bunched Phase Encoding

7.1 Introduction

As discussed in all earlier chapters of this thesis, non-Cartesian imaging can be advantageous due to the unusual properties of certain trajectories [15, 28, 29, 59, 35, 38, 31], and acceleration of such non-Cartesian acquisitions is currently a subject of much interest. One possibility to reconstruct artifact-free images from undersampled non-Cartesian data is parallel imaging, as seen in Section 3.3 as well as Chapter 6. While the image-based conjugate gradient SENSE (CG-SENSE) method of Pruessmann et al. [104] is able to reconstruct undersampled arbitrary trajectories in an efficient manner, this method requires a coil sensitivity map as well as many parameters, which in some cases can be difficult to determine. K-space based parallel imaging techniques for non-Cartesian trajectories [33, 100, 102, 103, 105, 106, 126, 129] have also been proposed. Most of these methods make use of specific symmetries in the k-space trajectories, such as radial or spiral GRAPPA [100, 102, 126], and cannot be applied to non-Cartesian trajectories without the requisite symmetry (see Section 3.3.1). Although some k-space based methods such as PARS [105, 106] or pseudo-Cartesian GRAPPA (discussed in Chapter 6 and [127]) are able to reconstruct images from undersampled arbitrary k-space data, these methods can be time consuming when using many source points for the reconstruction. Thus, a straightforward reconstruction method which is applicable to arbitrary trajectories and does not require a sensitivity map would be advantageous in many situations.

However, other methods have also recently been introduced for the reconstruction of undersampled data which are not based on parallel imaging. For instance, the work of Moriguchi et al. [23, 24, 25] exploits the Generalized Sampling Theorem of Papoulis [14] to generate images using non-uniformly sampled data. This theorem and its implications for accelerated MR imaging are discussed in Sections 7.2 and 7.3, respectively.

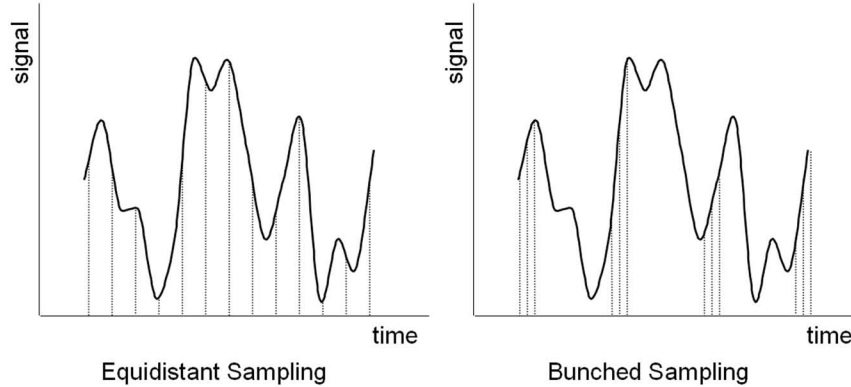


Figure 7.1: Left: A signal sampled equidistantly at Nyquist rate. Using the discrete datapoints, the underlying signal can be reconstructed. Right: The same signal sampled at $1/3$ of the Nyquist rate with bunches of 3 points. Because the average sampling rate is equal to the Nyquist sampling rate, the Generalized Sampling Theorem states that the underlying signal can be accurately reconstructed using these bunched points although the Nyquist criterion has been violated in some portions of the acquisition.

7.2 Generalized Sampling Theorem

The Generalized Sampling Theorem, first introduced by Papoulis in 1977, states that unaliased images can be reconstructed even when the Nyquist criterion is violated in portions of k -space, as long as the average sampling rate is equal to the Nyquist rate. In other words, if n bunched samples are acquired at $\frac{1}{n}$ th of the Nyquist rate, the original signal can be recovered despite the fact that the Nyquist criterion has not been met. This idea is shown schematically in Figure 7.1. Certain criteria must be fulfilled in order for the reconstruction to be successful, such as a bandwidth limited signal, although these criteria are generally met in MR images. In addition, such a reconstruction is highly sensitive to noise in the signal [137], meaning that it cannot be implemented in some cases. It is important to note that the idea of bunched sampling is novel in MRI, as it is generally assumed that the Nyquist criterion must be met in all portions of k -space to avoid aliasing artifacts (Section 2.2).

A discussion of the reconstruction of equidistantly sampled points in the time domain (i.e. k -space) from the bunched sampled points is excluded from this work due to the highly non-intuitive forms of the interpolation functions needed. However, the reconstruction can be understood by examining the fold-over artifacts that occur in the frequency domain (i.e. image) due to the violation of the Nyquist criterion. For simplicity, it is assumed that the signal consists of bunches of two points sampled at $\frac{1}{\alpha}$ th the Nyquist rate (where α is less than one, and the sampling rate exceeds the Nyquist rate), where each group of bunched points is sampled at $\frac{1}{2}$ the Nyquist rate. In k -space, the distance between the two points in the bunch is $\alpha \cdot \Delta k$, and the distance between the groups is $2 \cdot \Delta k$. This sampling schema can be seen in Figure 7.2, where the black points are separated from each other by $2 \cdot \Delta k$, and the black and gray points by $\alpha \cdot \Delta k$.

If only the black points in Figure 7.2 are used to generate an image, the typical $R=2$ fold-over




Figure 7.2: Left: The sampling schema described in the text, where the bunched groups (black and gray points) are sampled at one-half the Nyquist rate, but the average overall sampling rate is equal to the Nyquist rate. Performing an inverse Fourier transformation on the black points leads to an aliased image (right, top), where the signal values at locations $+L/2$, 0 , and $-L/2$ have the same phase. However, when an inverse Fourier transformation is calculated for the gray points, an additional location-dependent phase term appears (right, bottom). Using these phase differences, the unaliased image can be reconstructed (see text).

artifacts appear (see Sections 2.2 and 3.1.1, especially Figure 3.2); given that the original image appears between $-L/2$ and $+L/2$, a replica of the original appears shifted by $L/2$. As can be seen on the right-hand side, the phases of the signal at locations $+L/2$, 0 , and $-L/2$ (and all other locations) are identical. However, if only the gray points are used to generate an image, an additional location-dependent phase term is evident (Figure 7.2, bottom right). This phase term is directly related to the k -space shift of the gray points, namely $\alpha \cdot \Delta k$, and is a result of the Fourier Shift Theorem. This phase shift between the image generated from the black points, $I_0(x)$, and that generated from the gray points, $I_\alpha(x)$, can be related to the original, non-aliased image, $I(x)$ and $I(x - \frac{L}{2})$:

$$\begin{bmatrix} I_0(x) \\ I_\alpha(x) \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 1 & 1 \\ 1 & e^{-i\alpha\pi} \end{bmatrix} \cdot \begin{bmatrix} I(x) \\ I(x - \frac{L}{2}) \end{bmatrix} \quad (7.1)$$

This equation can be solved by taking the inverse (or pseudo-inverse) of the phase matrix:

$$2 \cdot \text{pinv} \left(\begin{bmatrix} 1 & 1 \\ 1 & e^{-i\alpha\pi} \end{bmatrix} \right) \cdot \begin{bmatrix} I_0(x) \\ I_\alpha(x) \end{bmatrix} = \begin{bmatrix} I(x) \\ I(x - \frac{L}{2}) \end{bmatrix} \quad (7.2)$$

thereby allowing the reconstruction of the non-aliased image. This method, which performs essentially a pixel-by-pixel unfolding of the aliased images, is strikingly similar to SENSE [3], which was discussed in detail in Section 3.2.4.

For higher acceleration factors, i.e. larger bunches of points with a larger distance between them, a general formula for the reconstruction can be written as:

$$\begin{bmatrix} I_0 \\ I_{\alpha_1} \\ \vdots \\ I_{\alpha_{m-1}} \end{bmatrix} = \frac{1}{p} \begin{bmatrix} 1 & e^{i2\pi\alpha_0/p} & \dots & e^{i2\pi\alpha_0(p-1)/p} \\ 1 & e^{i2\pi\alpha_1/p} & \dots & e^{i2\pi\alpha_1(p-1)/p} \\ \vdots & \vdots & \vdots & \vdots \\ 1 & e^{i2\pi\alpha_{m-1}/p} & \dots & e^{i2\pi\alpha_{m-1}(p-1)/p} \end{bmatrix} \cdot \begin{bmatrix} I(x) \\ I(x - \frac{L}{p}) \\ \vdots \\ I(x - \frac{L(p-1)}{p}) \end{bmatrix} \quad (7.3)$$

given that m bunched data points are sampled at a sampling rate of $p \cdot \Delta k$ (and $m \geq p$, such that the average sampling rate is greater than or equal to the Nyquist rate). Thus, an unaliased image can indeed be generated using bunched points, even when the sampling rate is lower than the Nyquist rate in some portions of the signal sampling.

7.3 Bunched Phase Encoding

The application of the Generalized Sampling Theorem for the acceleration of MRI data was first introduced by Moriguchi et al. [23, 24, 25]. In Moriguchi's work, a so-called bunched sampling technique is employed, where datapoints are collected along a zig-zag shaped trajectory such that some areas of k-space are more densely sampled than dictated by the Nyquist criterion, and fewer phase encoding lines than strictly necessary are acquired (see Figure 7.3). When the bunched points are acquired along a Cartesian-like trajectory, as shown on the left hand side of 7.3), the matrix method described in Section 7.2 can be employed [24]. However, when the bunched points are irregular, which is the case when using a non-Cartesian zig-zag trajectory (such as the radial bunched phase encoding (PR-BPE) trajectory shown on the right hand side of Figure 7.3), it is not possible to write the phase matrix in a simple form, as in Equation 7.3. As in SENSE, where the non-Cartesian case cannot be easily solved using a simple matrix inversion, the conjugate gradient (CG) method is employed (see Sections 3.2.4 and 3.3.2). Thus, using bunched data points along an undersampled non-Cartesian trajectory, unaliased images can be reconstructed with a CG algorithm [23, 25], despite the fact that the Nyquist criterion has been violated in some portions of k-space. This BPE method in conjunction with CG has been demonstrated using undersampled radial (projection reconstruction, or PR) data. It is important to note that the CG algorithm used for the bunched phase encoding (BPE) method does not require knowledge of coil sensitivity maps, as in CG-SENSE, and thus no parallel imaging is involved in this reconstruction.

While the BPE method has been shown to produce unaliased images with Nyquist undersampled datasets, the method requires non-standard trajectories and therefore high-performance gradients. Consequently, special pulse sequences and potentially gradient systems must be employed to generate such datasets. In addition, because the gradient execution is not always exactly what is desired, trajectory




Figure 7.3: Left: The Cartesian bunched phase encoding scheme of Moriguchi et al, where $r\Delta k_y$ denotes the width of the k-space covered by a bunched phase encoding step, and $q\Delta k_y$ the distance between two phase encoding lines. The value of q is set to be greater than one to decrease measurement time. Given N phase encoding lines in a standard measurement, and an acceleration of R , the total number of points in one bunched phase encoding line must be $N \cdot R$ for the reconstruction. Right: A schematic of the BPE-PR acquisition scheme, where the zig-zag trajectory is used along each projection. As shown in [25], fewer projections are required for an unaliased imaging when using the BPE-PR trajectory in conjunction with a CG reconstruction than when using a standard radial trajectory.

measurements are often required in order to perform the reconstruction. Thus, although the bunched phase encoding trajectories show promise for faster imaging, they often are difficult to implement and employ.

7.4 GROG-BPE

However, instead of generating these bunched datapoints using gradient encoding, the bunched trajectory can be mimicked using a standard “straight” trajectory in conjunction with parallel imaging. In this work, the BPE trajectory is replaced with points generated using the GRAPPA Operator Gridding (GROG) method (introduced in Chapter 4 and [118]), which replicates the function of these bunched gradients. Using GROG, each acquired datapoint can be shifted by a small amount in the x-direction, y-direction, or both simultaneously, thereby generating a cloud of additional points surrounding the original acquired point. These new bunched datapoints can then be used with the CG optimization to reconstruct accelerated images. In this chapter, several different strategies for generating bunched points are examined, namely using systematic [130] or random bunched points [131] as the BPE points. A schematic of the GROG-facilitated BPE reconstruction (GROG-BPE) and examples of the blipping strategies employed in this chapter are shown in Figure 7.4.

This method of generating bunched points using GROG is advantageous for several reasons. The first is that the reconstruction technique can be applied to arbitrary trajectories, because no symmetry is required, as in many k-space based non-Cartesian reconstruction schemes. No sensitivity maps are needed, and the computation times are similar to CG-SENSE, because both are based on the conjugate gradient algorithm. In addition, because the GROG weights can often be calculated from the under-

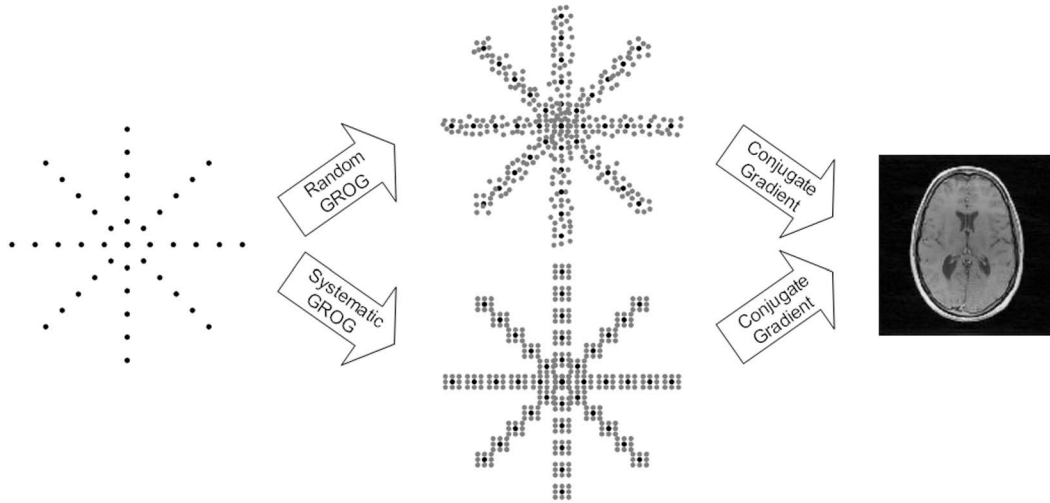


Figure 7.4: A schematic of the GROG-BPE procedure. The leftmost image depicts the standard undersampled radial dataset. Each datapoint is “bunched” using GROG; the upper middle image shows an example of random bunched points, and the lower image shows a systematic blipping scheme. Once the additional bunched points have been generated, a coil-by-coil conjugate gradient reconstruction algorithm is applied to the datapoints, resulting in the final unaliased image, as shown on the right.

sampled non-Cartesian datapoints themselves, this method can provide a completely self-calibrating reconstruction of undersampled non-Cartesian data [124]. For instance, the GROG weights needed to generate the bunched points can be calculated from undersampled radial and rosette datasets. For spiral datasets, the weights must be calculated from the fully-sampled data, although an undersampled dual-density or interleaved dynamic spiral acquisition would also be self-calibrating. The only requirement for the GROG-BPE method is that the receiver coils employed must provide sufficient sensitivity variations to generate the bunched points, a condition met by most clinical array coils. In this chapter, simulations are performed in order to examine the effects of different parameters on the reconstruction quality, and examples of accelerated in vivo radial, spiral and rosette trajectories are shown.

7.5 Simulation

7.5.1 Methods

In order to examine the effects of different shift sizes and number of additional GROG-BPE points on the reconstruction quality, simulated undersampled radial data were reconstructed using the method proposed here. To this end, an eight-element one-ring head coil array was simulated, and the sensitivities derived for this coil array using an analytic integration of the Biot-Savart equations were applied to a contrast-modified Shepp-Logan phantom. This simulated coil was chosen because it is similar to commercially available array coils. In addition, previous work [118] has shown that this arrangement and number of coils allows k -space shifts of at least $\Delta k = 0.5$ in both the x - and the y -direction using

GROG without the appearance of artifacts. The Cartesian data were resampled as radial data (200 projections, 256 read-out points, base matrix of 128×128) by sinc interpolation. This dataset was undersampled retrospectively to yield datasets with acceleration factors of 2 to 6 (100 projections for $R=2$, 67 for $R=3$, 50 for $R=4$, 40 for $R=5$, 34 for $R=6$). The GROG weights were determined from the undersampled radial datasets [124], and GROG-BPE datasets for each acceleration factor were created using these weights. The effects of different GROG jump sizes on the reconstruction quality were examined by increasing the maximum blip amplitude from $\Delta k_{max} = 0.1$ to $\Delta k_{max} = 1.0$ in increments of $\Delta k = 0.1$. In addition, the effect of the density of bunched points on the reconstruction was examined by increasing the total number of GROG-BPE points. The minimum number of points used for a reconstruction was 9 GROG-BPE bunched points per acquired point, and the maximum was 225. Finally, two different blipping patterns were employed: the first was systematic blipping, and the second was random blipping. For the systematic blipping, the GROG-BPE points were regularly distributed over the maximum blip amplitude in both the k_x - and the k_y -directions. For instance, if a total of 49 bunched points are to be used with a maximum amplitude of $\Delta k_{max} = 0.6$, the k_x -locations of the bunched points would be:

$$k_x + (-0.6, -0.4, -0.2, 0, 0.2, 0.4, 0.6)$$

With the same distribution of k_y points, the desired number of total points (49) has been achieved. For the random blipping, the direction and distance of the bunched points was determined randomly. As in the systematic blipping, the maximum blip amplitude and the number of BPE points was given, but the actual location of the bunched points was random and uniformly distributed over the range from $-\Delta k_{max}$ to $+\Delta k_{max}$. A schematic depiction of both the systematic and the random blipping is shown in Figure 7.4. After the BPE-mimicked points were generated using GROG, the resulting k-space data was processed using the CG-INNG algorithm described in [23], although without using sensitivity maps (i.e. no CG-SENSE is employed). This algorithm is similar to that used in [25] and is made up of the following steps (performed independently for each coil):

- The k-space data (made up of the bunched points generated using GROG) are distributed onto an oversampled Cartesian grid using a nearest neighbor algorithm (i.e. the non-Cartesian datapoints are placed onto the nearest rounded Cartesian grid point)
- An inverse FFT is performed, yielding an image with an increased FoV (corresponding to the oversampling factor)
- The central portion of the image is cropped out and used to initialize the CG algorithm
- The new image from the CG algorithm is placed into a larger matrix filled with zeros (again corresponding to the oversampling factor)
- An FFT is performed on the large FoV image, resulting in an oversampled updated k-space
- The k-space data at the original locations is retained, and all other areas in k-space are set to zero

- The inverse FFT is performed on the updated k-space data, yielding an updated image
- The updated image is cropped out of the larger FoV, and fed back into the CG algorithm
- After a given number of iterations have been performed, or the stopping criteria have been met, the cropped updated image is accepted as the final reconstructed image

A schematic of CG-INNG process is shown in Figure 7.5. No density compensation function is required for this gridding algorithm. For all datasets, 10 iterations were performed using an oversampling factor of $s = 2$, followed by 10 iterations with an oversampling factor of $s = 10$. No improvement in the image quality was observed when using more iterations or larger grid oversampling factors. The final reconstructed images for all acceleration factors, maximum blip amplitudes, and number of blip points were compared to the noiseless Shepp-Logan phantom using the root mean square error (RMSE) value in order to determine which parameters offered the best image quality for each acceleration factor. Note that the region of k-space support for the Cartesian phantom was reduced with a radial mask of the same extent as the radial trajectory to enforce the circular FoV, which slightly affects the appearance of the standard noiseless phantom.

7.5.2 Results

The results of the simulations for the acceleration factor $R=4$ and different GROG-BPE parameters are shown in the graph in Figure 7.6. Due to the large number of simulated images, only a few graphs are shown to demonstrate the trends in RMSE values for different parameters. Tables 7.1 and 7.2 show the lowest RMSE errors and parameters used for each of the reconstruction types (i.e. systematic or random) for each acceleration factor. As seen in Figure 7.6, for systematic blipping (the black lines and points), an increase in the maximum amplitude of the bunched points leads first to a decrease, and then an increase in the RMSE value for a given acceleration factor and number of bunched points. The reason for this can be seen in Figure 7.7, which shows the images resulting from a $R=4$ using 9 BPE points and an increasing maximum blip amplitude. When the maximum blip is small ($\Delta k_{max} = 0.1$), the blurring evident in the undersampled image is still present. By increasing the blip amplitude, this blurring decreases. However, because the GROG weights become more unstable with increasing blip distance, the error in the BPE points grows as the blurring decreases. A RMSE minimum (0.94%) is seen in Figure 7.6 when the maximum amplitude is $\Delta k_{max} = 0.4$, and the use of a larger amplitude leads to a darkening of portions of the image which can be attributed to errors in the center of k-space from inaccurate bunched points. The results of employing different numbers of systematic bunched points are also shown in Figure 7.6. For larger blip amplitudes, a greater number of bunched points leads to a larger RMSE. This is again due to the increasing inaccuracy of GROG weights as the maximum amplitude increases; when more bunched points are generated further away from the acquired point, the reconstruction error increases. It is interesting to note that there are no significant differences in RMSE values when employing small maximum amplitudes but different numbers of bunched points for a given acceleration value; this is due to the fact that these points have smaller GROG errors. It would be expected from the theory of Papoulis that a larger number of points for a given maximum

	R=2	R=3	R=4	R=5	R=6
Minimum RMSE	0.21%	0.44%	0.94%	1.65%	2.40%
Bunched/Acquired Point	9	9	9	49	49
Maximum Blip Amplitude	$\Delta k_{max} = 0.2$	$\Delta k_{max} = 0.3$	$\Delta k_{max} = 0.4$	$\Delta k_{max} = 0.5$	$\Delta k_{max} = 0.5$

Table 7.1: Minimum RMSE values and reconstruction parameters for images reconstructed from simulated undersampled radial data using systematically blipped GROG-BPE for acceleration factors from R=2 to R=6.

blip amplitude would lead to a better reconstruction than a smaller number; in this implementation of the method, however, no evidence for that is seen for systematic blipping.

More interesting are the results of the simulations with random blipping. As evidenced in Figure 7.6, the RMSE for images reconstructed with random bunched points (gray lines) are lower than those reconstructed using systematic blipping (black lines). This statement is true especially for large maximum blip amplitudes ($\Delta k_{max} > 0.5$). This phenomenon can be explained by the incoherent artifacts which are a result of the random blipping patterns. By arbitrarily choosing a shift direction and amount for each acquired point, the errors which result from the GROG weights tend to average each other out. In the systematic blipping case, these errors, while still present as incoherent noise, tend to reinforce each other, thereby yielding larger RMSE errors for the reconstructed images. It should be noted that the minimum RMSE for the random blipping method appears at a larger maximum blip amplitude ($\Delta k_{max} = 0.5$) than the minimum for the systematic blipping, and this minimum is not as narrow. This indicates that the choice of the maximum blipping amplitude is not as crucial for the random blipping as it is for the systematic blipping. Finally, the RMSE of the reconstructions with random blipping with more points always yielded a lower RMSE than reconstructions with fewer points (dotted vs. solid gray lines). This result fits with the generalized sampling theorem of Papoulis, in that more bunched points should lead to a better reconstruction of the actual signal.

Examples of images reconstructed with the random blipping for an acceleration factor of 4, a total of 225 random GROG points per acquired datapoint, and differing maximum blip amplitudes are shown in Figure 7.8. Using a maximum blip amplitude of $\Delta k_{max} = 0.5$, the reconstruction yields an image which no longer contains the blurring evident in the R=4 undersampled image (with a RMSE of 0.64%). In addition, the darkening of the center of the image present in the systematic blipping reconstructions does not appear in the random blipping using this maximum amplitude; it is, however, present when

	R=2	R=3	R=4	R=5	R=6
Minimum RMSE	0.11%	0.24%	0.64%	1.42%	2.31%
Bunched/Acquired Point	225	225	225	225	225
Maximum Blip Amplitude	$\Delta k_{max} = 0.3$	$\Delta k_{max} = 0.5$	$\Delta k_{max} = 0.5$	$\Delta k_{max} = 0.6$	$\Delta k_{max} = 0.6$

Table 7.2: Minimum RMSE values and reconstruction parameters for images reconstructed from simulated undersampled radial data using randomly blipped GROG-BPE for acceleration factors from R=2 to R=6.




Figure 7.5: A schematic of the CG-INNG algorithm, as described in detail in the text. The portion on top (the initial gridding of the GROG-bunched data) is performed only once, and the portion in the dotted area is performed multiple times until the CG algorithm converges. The schematic k-spaces represent the acquired data non-Cartesian data assigned to the oversampled grid ($s = 2$ in this case), while the small and large images show the correct FoV and zero-padded FoV, respectively.

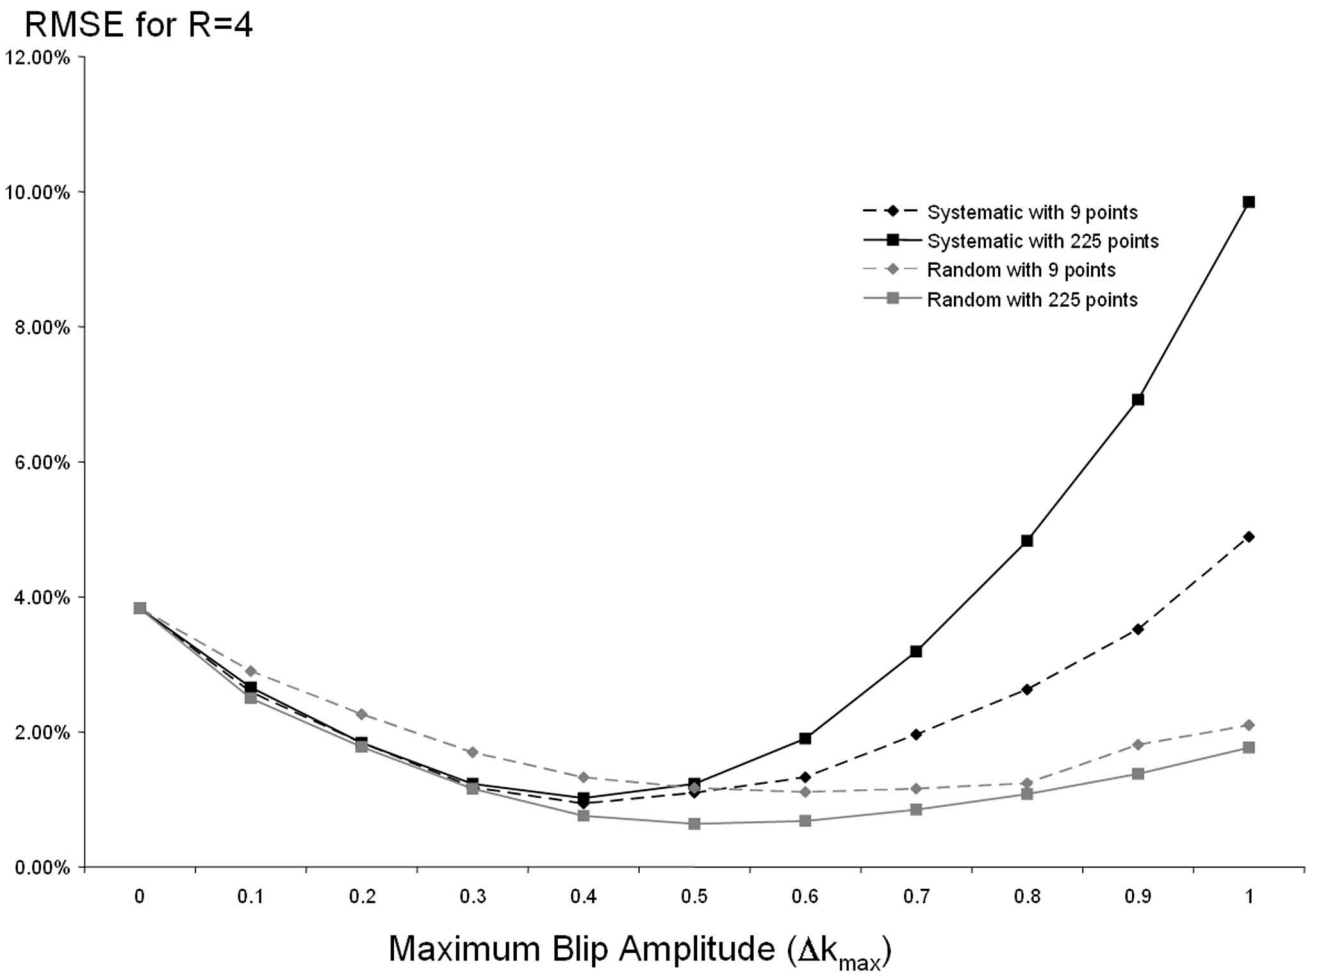


Figure 7.6: A graph showing the RMSE values for simulated R=4 radial images reconstructed with GROG-BPE performed using different parameters (black: systematic GROG blipping, gray: random GROG blipping; diamond: 9 GROG bunched points/acquired point, square: 225 bunched points/acquired point).




Figure 7.7: Reconstructions of the $R=4$ simulated data using systematic GROG-BPE and 9 bunched points/acquired point for different maximum blip amplitudes. The top left image, with blip amplitude $\Delta k_{max} = 0$, corresponds to the $R=4$ undersampled image. As the maximum blip amplitude increases, the blurring evident in the undersampled image decreases, but the center of the image darkens when using blip amplitudes larger than $\Delta k_{max} = 0.5$. This phenomenon can also be seen in the plot (black diamonds) in Figure 7.6.

using larger blip amplitudes and random blipping ($\Delta k_{max} = 0.9$), thereby increasing the RMSE of the reconstruction (RMSE=1.38%). Finally, as apparent in the graph shown in Figure 7.6, the quality of images reconstructed using random blipping is not as dependent on the maximum blip amplitude as those reconstructed using systematic blipping. Thus, in general, random blipping leads to both better reconstruction quality (i.e. lower RMSE values), and more stable reconstructions (i.e. less dependent on the choice of maximum blip amplitude).

7.6 In Vivo Experiments

7.6.1 Methods

Given the results for simulations performed with different blip amplitudes, different numbers of bunched points, and the two blipping patterns, it was determined that all further in vivo reconstructions should be performed with the following reconstruction parameters:




Figure 7.8: Reconstructions of the R=4 simulated data using random GROG-BPE and 225 bunched points/acquired point for different maximum blip amplitudes. The top left image again corresponds to the R=4 undersampled image. The R=4 image with the lowest RMSE (0.64%) is the $\Delta k_{max} = 0.5$ image shown on the lower left.

- Random Blipping
- Large oversampling matrices for the conjugate gradient reconstruction ($s = 10$)
- Large numbers of bunched points
- Maximum amplitudes determined by the number of receiver coils and geometry

Several comments must be made about the choice of parameters. Random blipping was chosen because the RMSE of such reconstructions was lower than the RMSE for the systematic blipping. The number of bunched points used for each reconstruction was limited by memory constraints; the generation of 225 additional points for a dataset with 512 read out points, 50 radial rays, and 12 channels was not always feasible with the memory available. Thus, a smaller number was chosen for the in vivo reconstructions. In addition, coil arrays with a large number of elements allow for larger GROG shifts without substantial artifacts. Thus, several maximum amplitude values were examined, and the sharpest image chosen for each in vivo reconstruction. Finally, the same oversampling matrix sizes were used as for the simulations. The different non-Cartesian trajectories employed are described below, and the specific measurement parameters are given in Table 7.3.

	Radial	Spiral	Rosette
Read Points	512	7289 ^a	1000
Phase Encoding Steps	50	2	15
Number of Coils	12	32	12
Base Matrix Size	256	192	128
Undersampling Factor	5	2	4
Maximum Blip Amplitude	$\Delta k_{max} = 0.3$	$\Delta k_{max} = 1.0$	$\Delta k_{max} = 1.0$
Bunched/Acquired Points	121	121	50

Table 7.3: Dataset sizes and reconstruction parameters for the in vivo accelerated radial, spiral, and rosette trajectories used to demonstrate GROG-BPE. ^aA center-out constant-linear-velocity spiral trajectory was used.

In vivo radial data (12 channel, 256 projections, 512 read-out points, base matrix 256×256) were acquired using a 1.5T Siemens Espree scanner (Siemens Medical Solutions, Erlangen, Germany). The data were retrospectively undersampled by removing all but 50 of the radial arms, yielding an acceleration factor of $R=5$ with respect to Cartesian sampling (and approximately $R=8$ with respect to radial sampling). The GROG weights were determined from the undersampled radial data themselves without further processing using the GROG self-calibration method described in Chapter 5 and [124]. Using the random blipping method described above with 121 bunched points for each acquired point and a maximum blip amplitude of $\Delta k_{max} = 0.3$, GROG-BPE reconstructions were performed.

In vivo spiral data (32 channel [87], 4 spiral arms, 7289 read-out points, base matrix 192×192) were acquired using a 3T Trio scanner (Siemens Medical Solutions) using a constant-linear-velocity trajectory. As stated in Section 7.1, the undersampled spiral trajectory is currently not self-calibrating for GROG. Thus, in order to determine the GROG weights, the full spiral data were first reordered onto a constant-angular-velocity trajectory using a 1D Fourier interpolation in the time domain (i.e. read-out direction). This reordering yielded 448 pseudo-projections, each containing 128 points which could be used for the calibration of the GROG weights [124]. The spiral data were then retrospectively undersampled to $R=2$, i.e. two spiral arms were removed from the acquisition, and the random blipping method was employed, also with 121 bunched points per acquired point and a maximum blip amplitude of $\Delta k_{max} = 1.0$. This larger blip amplitude was made possible by the 32 independent receiver channels, which lead to more stable GROG weights for larger shifts.

Finally, in vivo rosette data were acquired on a 1.5T Avanto scanner (Siemens Medical Solutions). The rosette trajectory employed differed slightly from that proposed in [35], in that the following formula was used to describe a single shot of the trajectory:

$$k = 4 \cdot \cos(t) - \cos(4 \cdot t) + i \cdot 4 \cdot \sin(t) - i \cdot \sin(4 \cdot t)$$

where t runs from 0 to 2π . This parametrization leads to a trajectory comprised of three large petals which overlap 15 times in the course of a single read out. Using this trajectory, the entire k -space for a 128×128 matrix can be covered with 60 shots (by rotating k from 0° to 59° in increments of 1°).




Figure 7.9: Left: The radial reference image as described in the text. Center: The R=5 undersampled radial image. Right: The GROG-BPE reconstruction of the R=5 radial data.

The trajectory was implemented on the scanner and measured using the method of Duyn et al. [138]. Accelerated in vivo data were acquired with 15 repetitions of the above trajectory (by rotating k from 0° to 55° in increments of 4° , corresponding to an acceleration factor of $R=4$), 1000 read points, and a base matrix of 128×128 , using a 12-channel receiver coil. The GROG weights cannot currently be calculated using raw rosette trajectory data, and thus the central portion of k -space (which is fully-sampled despite the $R=4$ undersampling) was gridded using the CG-INNG algorithm. The resulting low-resolution Cartesian dataset (20×20) was used to determine the base GROG weights. Using these weights, random blipping was performed with a total of 50 GROG points per acquired point and a maximum blip amplitude of $\Delta k_{max} = 1.0$. Finally, a reconstructed image was created using the GROG-BPE dataset and CG-INNG.

7.6.2 Results

The reference images, undersampled images, and results of the reconstructions of the accelerated radial, spiral and rosette datasets using GROG-BPE with CG-INNG are shown in Figures 7.9 through 7.11. For comparison purposes, the reference and undersampled images were also reconstructed using CG-INNG but without the extra bunched points.

In the radial reconstruction of the R=5 undersampled dataset shown on the right-hand side of Figure 7.9, the blurring and streaking artifacts present in the undersampled image are removed when employing the BPE-GROG algorithm. Similarly, the undersampling artifacts present in the R=2 accelerated spiral image shown in the center of Figure 7.10 are no longer evident in the reconstruction shown on the right, despite the relatively malicious artifacts which result from missing points in the center of k -space. Finally, the reconstruction of the R=4 undersampled rosette dataset shown in Figure 7.11 also appears less blurry than the undersampled image, although a noise increase can be observed as compared to the reference image depicted on the left.

7.7 Discussion

As shown in the simulations and in vivo reconstructions, GROG-facilitated bunched phase encoding with conjugate gradient reconstruction can be used to reconstruct artifact-free images from undersampled non-Cartesian datasets. According to the Generalized Sampling Theorem of Papoulis [14], the sampling of all points in k -space does not necessary have to fulfill the Nyquist criterion; given that the average sampling rate is equal to the Nyquist rate, it is possible to reconstruct the original signal using methods such as the iterative conjugate gradient algorithm. This is a novel concept in MR imaging, where it is generally assumed that aliasing artifacts result when the Nyquist criterion is violated in any single portion of k -space. While it has been shown that modified bunched phase encoding trajectories, which allow the acquisition of clusters of points which are bunched in k -space, can be exploited to reduce sampling time, this method requires novel trajectories and gradient performance above and beyond that required for the standard trajectory. Thus, a method which allows the utilization of the Generalized Sampling theorem to reconstruct undersampled datasets without the need for such modifications to standard non-Cartesian trajectories, such as GROG-BPE, is advantageous.

Using GROG, either additional systematic or random bunched points can be generated from points acquired along a standard non-Cartesian trajectory. These points can then be used in conjunction with the CG algorithm, exactly in the same way Moriguchi et al. [23, 25] employ the bunched points, to reconstruct artifact-free images from undersampled non-Cartesian data. This manuscript demonstrates the application of this method for generating images from undersampled simulated radial data and in vivo radial, spiral, and rosette data.

As shown in Section 7.5, random blipping outperforms systematic blipping in terms of stability and image reconstruction quality (as measured by the RMSE of the reconstructed images). As described in previous work [118], the use of GROG can lead to errors in the generated bunched points, especially when the shifts are large. When random blipping is employed, these errors are also random in both direction and magnitude, and appear as less coherent artifacts in the reconstructed images. Systematic blipping results in more coherent artifacts, which lead to higher RMSE values. Thus, one can conclude that random blipping is preferable to systematic blipping, and the in vivo reconstructions shown here all employed random blipping. In addition, the effects of parameters such as maximum blip amplitude and number of bunched points were also examined. As expected from the Generalized Sampling theorem, the use of a greater number of bunched points in conjunction with the CG algorithm yields a better image quality than using a smaller number of bunched points for all acceleration factors examined.

Finally, as shown in Figures 7.6 through 7.8, the reconstruction quality increases and then decreases when the maximum blip amplitude is increased. This can be explained by the use of GROG to generate the bunched points. As mentioned above, large k -space shifts using GROG weight sets are less accurate than small shifts. Because of this, larger maximum blip amplitudes in k -space lead to larger errors in the bunched points generated using GROG. On the other hand, Moriguchi et al. [24] have shown that larger blip amplitudes lead to better reconstructions as a consequence of the Generalized Sampling Theorem. Thus, a balance must be found between removing the aliasing artifacts (blurring, streaking, fold-over), and introducing errors due to noise-prone GROG points. The trade-off depends on the quality of the




Figure 7.10: Left: The spiral reference image as described in the text. Center: The R=2 undersampled spiral image. Right: The GROG-BPE reconstruction of the R=2 spiral data.

GROG weights and the errors they introduce in the bunched points. However, the additional error introduced in the bunched points depends on a number of factors, such as the number of receiver coils used, their orientation, the SNR of the reference and undersampled dataset, etc. In addition, the maximum blip amplitude required depends on the undersampling factor and the trajectory. For instance, although the same receiver coil was employed for the radial and rosette datasets, the optimal maximum blip amplitude was found to be $\Delta k_{max} = 0.3$ for the R=5 radial dataset and $\Delta k_{max} = 1.0$ for the R=4 rosette dataset. This is due to the differences in the degree of undersampling in different portions of k-space for the two trajectories; larger “holes” are found closer to the center of k-space for the R=4 rosette trajectory than for the R=5 radial trajectory, requiring larger blip amplitudes for the reconstruction. Because the weights derived from the receiver coil can be used to generate more accurate $\Delta k = 0.3$ GROG shifts than $\Delta k = 1.0$ shifts, the resulting radial reconstruction appears to have a noise level similar to that of its reference image, while the rosette reconstruction is notably noisier than the rosette reference. Thus, the choice of maximum blip amplitude must be made in consideration of the above factors. Luckily, one does not have to decide on the maximum blip amplitude before acquiring the data; different amplitudes can be examined and an appropriate value chosen in the course of the reconstruction.

There are, of course, other options for the reconstruction of accelerated non-Cartesian datasets, such as parallel imaging methods. The most commonly used is CG-SENSE [104], described in Section 3.3.2, which is advantageous because it can be used with arbitrary non-Cartesian trajectories. This reconstruction algorithm also relies on the CG method to generate artifact-free images from undersampled datasets. The primary difference between the method proposed here and CG-SENSE is that CG-SENSE requires sensitivity maps. While such sensitivity maps can in theory be generated whenever the central portion of k-space is fully-sampled, in cases such as the radial and rosette trajectories shown here, these maps are difficult to calculate in low SNR areas or given pre-folding. GROG-BPE, however, does not require sensitivity maps and is also self-calibrating for the same undersampled trajectories




Figure 7.11: Left: The rosette reference image as described in the text. Center: The R=4 undersampled rosette image. Right: The GROG-BPE reconstruction of the R=4 rosette data.

as CG-SENSE. The GROG weights for the radial dataset can be determined using all of the acquired datapoints [124], while the central, Nyquist-sampled portion of the undersampled rosette dataset can be gridded and used as a Cartesian block of data for the GROG weight determination. While the R=2 spiral dataset demonstrated here cannot currently be used to determine the GROG weights, possibilities for a self-calibrating alternative include the use of dual-density spirals, which would be self-calibrating.

The other major parallel imaging alternative is the general class of non-Cartesian GRAPPA methods, including radial GRAPPA [100], spiral GRAPPA [102, 126], 1D NC GRAPPA [103], etc. which are detailed in Section 3.3.1. Non-Cartesian GRAPPA methods generally rely on symmetrical trajectories, which can be used to generate and apply weight sets using entire segments of the undersampled k-space. In cases of trajectories which do not contain the necessary symmetry, such as the rosette trajectory used here, non-Cartesian GRAPPA cannot be employed. In addition, these methods often require a fully-sampled calibration dataset in order to determine the GRAPPA weights for each of the segments. GROG-BPE, in contrast, is self-calibrating for those trajectories which include a fully-sampled central k-space portion, such as radial, rosette, dual-density spiral, or PROPELLER. Thus, the GROG-BPE method presented here is far more general than other non-Cartesian GRAPPA methods and requires no additional calibration information. However, the current implementation of GROG-BPE requires more computation time than other iterative methods due to the large number of points which must be generated with GROG and gridded with CG-INNG. Fortunately, once the additional bunched points have been determined using the GROG weights, the following CG algorithm can be performed coil-by-coil and in parallel. This greatly decreases the time to required perform the reconstruction, although significant memory requirements still remain a challenge.

7.8 Conclusion

The new method described here, namely GROG-facilitated Bunched Phase Encoding, or GROG-BPE, can be used as an alternative to non-Cartesian parallel imaging techniques and to the standard BPE described by Moriguchi et al [23, 24, 25]. By exploiting the Generalized Sampling Theorem of Papoulis [14], GROG-BPE allows the reconstruction of artifact-free images from undersampled non-Cartesian data using additional points generated using GROG in conjunction with an iterative CG algorithm. This method does not require specialized k-space trajectories or additional gradient performance, as does the method of Moriguchi et al, because the “bunched” points are generated using GROG and the standard undersampled non-Cartesian datasets. In addition, it can be used for arbitrary k-space trajectories without the need for sensitivity maps, unlike CG-SENSE. Finally, for many undersampled datasets, GROG-BPE is completely self-calibrating and requires no additional dataset. As demonstrated with in vivo datasets, GROG-BPE results in reduced artifacts and blurring for undersampled radial, spiral, and rosette images.

Chapter 8

Summary

Magnetic Resonance Imaging (MRI) is an imaging modality which provides anatomical or functional images of the human body with variable contrasts in an arbitrarily positioned slice without the need for ionizing radiation. In MRI, data are not acquired directly, but in the reciprocal image space (otherwise known as k-space) through the application of spatially variable magnetic field gradients. The k-space is made up of a grid of data points which are generally acquired in a line-by-line fashion (Cartesian imaging). After the acquisition, the k-space data are transformed into the image domain using the Fast Fourier Transformation (FFT).

However, the acquisition of data is not limited to the rectilinear Cartesian sampling scheme described above. Non-Cartesian acquisitions, where the data are collected along exotic trajectories, such as radial and spiral, have been shown to be beneficial in a number of applications, such as real-time cardiac imaging [61, 62, 63, 64], MR angiography [68, 69, 70], lung imaging [46, 47, 48], motion correction [71, 72], and fMRI [57, 58, 59, 60]. However, despite their additional properties and potential advantages, working with non-Cartesian data can be complicated. The primary difficulty is that non-Cartesian trajectories are made up of points which do not fall on a Cartesian grid, and a simple and fast FFT algorithm cannot be employed to reconstruct images from non-Cartesian data. In order to create an image, the non-Cartesian data are generally resampled on a Cartesian grid, an operation known as gridding, before the FFT is performed. There are many different gridding algorithms which can be used to perform this task, and most involve many parameters which must be optimized, such as a convolution function, kernel width, oversampling factor, Density Compensation Function (DCF), regularization parameters, etc. In addition, such algorithms cannot be used to grid Nyquist undersampled datasets. Thus, a simple method of gridding which does not require many parameters to be optimized and can be used to grid undersampled data would simplify the gridding process and make non-Cartesian imaging more accessible.

Another challenge for non-Cartesian imaging is the combination of unusual trajectories with parallel imaging. Because the acceleration factor changes in both degree and direction in an undersampled non-Cartesian k-space, standard Cartesian parallel imaging methods such as SENSE [3] or GRAPPA [7] cannot be directly employed. The most commonly used and most flexible non-Cartesian parallel imaging reconstruction algorithm is CG-SENSE, developed by Pruessmann et al. [104]. While CG-

SENSE can be applied to any undersampled non-Cartesian trajectory, the quality of the reconstructed image depends on the accuracy of the coil sensitivity map, which can be difficult to acquire in many cases. Non-Cartesian GRAPPA methods [100, 102, 103, 126] do not require sensitivity maps, which makes them advantageous for lung and cardiac imaging, but they must be formulated differently for each non-Cartesian trajectory. This is because the missing non-Cartesian data must be reconstructed before gridding is performed due to the difficulties in gridding undersampled datasets. Prior to this thesis, a generalized non-Cartesian GRAPPA method which can be applied to arbitrary trajectories had yet to be developed.

This thesis has presented several new non-Cartesian parallel imaging methods which simplify both gridding and the reconstruction of images from undersampled data. In Chapter 4, a novel approach which uses the concepts of parallel imaging to grid data sampled along a non-Cartesian trajectory called GRAPPA Operator Gridding (GROG) is described. GROG shifts any acquired k-space data point to its nearest Cartesian location, thereby converting non-Cartesian to Cartesian data. Unlike other parallel imaging methods, GROG synthesizes the net weight for a shift in any direction from a single basis set of weights along the logical k-space directions. Given the vastly reduced size of the basis set, GROG calibration and reconstruction requires fewer operations and less calibration data than other parallel imaging methods for gridding. Instead of calculating and applying a DCF, GROG employs only local averaging, as the reconstructed points fall upon the Cartesian grid. In addition, because the effective gridding kernel is one, local gridding can be performed, making the gridding of undersampled data possible. The only requirements for GROG are a multi-channel acquisition and a calibration dataset for the determination of the GROG weights. Simulations demonstrating that the root mean square error values of images gridded with GROG are similar to those for images gridded using the gold-standard convolution gridding were shown. Finally, GROG was compared to the convolution gridding technique using data sampled along radial, spiral, rosette, and PROPELLER trajectories.

Chapter 5 discusses an extension of GRAPPA Operator Gridding, namely Self-Calibrating GRAPPA Operator Gridding (SC-GROG). SC-GROG is a method by which non-Cartesian data can be gridded using spatial information from a multi-channel coil array without the need for an additional calibration dataset, as required in standard GROG. Using self-calibrating GROG, the non-Cartesian datapoints are shifted to nearby k-space locations using parallel imaging weight sets determined from the datapoints themselves. SC-GROG also employs the GRAPPA Operator, a special formulation of the general reconstruction method GRAPPA, to perform these shifts. The innovation introduced in Self-Calibrating GROG allows the shift operators to be calculated directly out of the non-Cartesian data themselves. This eliminates the need for an additional calibration dataset, which reduces the imaging time, and also makes the GROG reconstruction more robust by removing possible inconsistencies between the calibration and non-Cartesian datasets. Simulated and in vivo examples of radial and spiral datasets gridded using self-calibrating GROG are compared to images gridded using the standard method of convolution gridding.

Although GROG can be used to grid undersampled datasets, it is important to note that this method uses parallel imaging only for gridding, and not to reconstruct artifact-free images from under-

sampled data. As stated above, simple GRAPPA cannot be used with non-Cartesian datasets because it necessitates the acquisition of regularly sampled Cartesian k-space data in order to reconstruct a non-aliased image efficiently. Thus, the extension of Cartesian GRAPPA to non-Cartesian sequences is non-trivial. However, GROG can be combined with GRAPPA in order to simplify the reconstruction process for undersampled non-Cartesian datasets. Chapter 6 introduces a simple, novel method for performing modified Cartesian GRAPPA reconstructions on undersampled non-Cartesian k-space data gridded using GROG to arrive at a non-aliased image. Because the undersampled non-Cartesian data cannot be reconstructed using a single GRAPPA kernel, several Cartesian patterns are selected for the reconstruction. This flexibility in terms of both the appearance and number of patterns allows this pseudo-Cartesian GRAPPA to be used with undersampled non-Cartesian datasets acquired with any non-Cartesian trajectory. The successful implementation of the reconstruction algorithm using several different trajectories, including radial, rosette, spiral, 1D non-Cartesian, and zig-zag trajectories, has been demonstrated. Pseudo-Cartesian GRAPPA is important because it is the first GRAPPA reconstruction method which can be applied to any arbitrary non-Cartesian trajectory. In addition, it is the first method to make use of the fact that GROG can be employed to grid undersampled data.

Finally, Chapter 7 discusses a novel method of using GROG to mimic the bunched phase encoding (BPE) scheme. It has been shown that the use of BPE in conjunction with a conjugate gradient reconstruction algorithm can decrease scan time by reducing the number of phase encoding lines needed to generate an unaliased image [23, 24, 25]. In MRI, it is generally assumed that an artifact-free image can be reconstructed only from sampled points which fulfill the Nyquist criterion. However, the BPE reconstruction is based on the Generalized Sampling Theorem of Papoulis [14], which states that a continuous signal can be reconstructed from sampled points as long as the points are on average sampled at the Nyquist frequency. Thus, the BPE method of Moriguchi et al. takes advantage of the idea that bunched points (which are sampled at a frequency higher than dictated by Nyquist) can be used to reconstruct k-space data which is sampled at a lower rate in some locations. However, the acquisition of such bunched data at the scanner requires both modified pulse sequences and extremely high and stable gradient performance. A novel method of generating the “bunched” data using GRAPPA Operator Gridding (GROG), which shifts datapoints by small distances in k-space using the GRAPPA Operator instead of employing zig-zag shaped gradients, is presented in this chapter. Many bunched points can be generated random distances and directions (within a maximal shift value) from the original point using GROG, which removes the limitations imposed by the gradient system and the sampling time available. With the conjugate gradient reconstruction method, these additional “bunched” points can then be used to reconstruct an artifact-free image from undersampled data. This method is referred to as GROG-facilitated Bunched Phase Encoding, or GROG-BPE. In order to better understand how the patterns of bunched points, maximal blip size, and number of bunched points affect the reconstruction quality, a number of simulations using the GROG-BPE approach are performed. Finally, to demonstrate that this method can be combined with any trajectory, examples of artifact-free in vivo images reconstructed from undersampled radial, spiral, and rosette data are shown.

With the introduction of GROG and SC-GROG, non-Cartesian datasets can be gridded without the need for additional parameters, greatly simplifying and speeding up the gridding process. The further combination of GROG and either GRAPPA or the Generalized Sampling Theorem allows for the reconstruction of artifact-free images from undersampled non-Cartesian data sampled along an arbitrary trajectory. For many trajectories, such as radial, dual-density spiral, PROPELLER, and rosette, these methods are completely self-calibrating, which can lead to both imaging time reduction and more stable reconstructions. Thus, the methods developed in this thesis make both non-Cartesian imaging and parallel imaging more accessible to the typical MR operator, and can be employed to accelerate the data acquisition for many non-Cartesian applications.

Indeed, work with both fully-sampled and undersampled non-Cartesian datasets is currently underway using the methods described here. The gridding of undersampled radial data with GROG has been shown to be useful in so-called Compressed Sensing (CS) reconstructions [133, 134], which are gaining more attention in the MR community due to the possibility for very high acceleration factors. It is expected that GROG will play a role in making these CS methods more efficient, as it is the only method which can be used to grid undersampled data, and the combination of GROG-BPE and CS is currently a subject of research. In addition to gridding, GROG has also been used to shift points in k-space in order to perform motion correction [132], a novel use of the GRAPPA Operator. Both pseudo-Cartesian GRAPPA and GROG-BPE are currently being examined as alternatives to accelerate radial acquisitions, especially in cardiac and lung imaging. These methods require no calibration data in the form of extra radial lines or a coil sensitivity map, which can be difficult to acquire in cardiac data due to motion and in lung data due to the low signal intensities. Similarly, GROG and GROG-BPE for the reconstruction and acceleration of 3D radial data have been proposed due to the simplicity and time-effectiveness of the GROG reconstruction method [125]. The possibility of combining metabolite detection or fMRI using an accelerated rosette trajectory and pseudo-Cartesian GRAPPA is also being considered, and ideas for developing new trajectories which can be realized only using GROG have been discussed. The variety of different topics under investigation using methods developed as part of this thesis indicate that both gridding and data acceleration using the GRAPPA Operator have and will continue to expand non-Cartesian imaging to make additional Magnetic Resonance Imaging applications possible.

Chapter 9

Zusammenfassung

Die Magnetresonanztomographie (MRT) ist ein nichtinvasives bildgebendes Verfahren ohne Strahlenbelastung und eignet sich zur biomedizinischen Darstellung verschiedener Gewebetypen mit hoher räumlicher Auflösung und sehr gutem Kontrastverhalten. In der MRT erfolgt die Datenaufnahme im reziproken Bildraum – auch k-Raum genannt - welcher typischerweise entlang eines diskreten kartesischen Gitters abgetastet wird. Ein Bild erhält man schließlich durch eine schnelle Fouriertransformation der aufgenommenen k-Raum-Daten. Neben den kartesischen Akquisitionsschemata haben sich in den letzten Jahren auch vereinzelt nichtkartesische MRT-Verfahren in der klinischen Routine durchgesetzt, welche den k-Raum beispielsweise entlang radial- oder spiralförmiger Trajektorien abtasten. Solche nichtkartesischen Trajektorien erreichen eine hohe Abtasteffizienz, was zu einer Verkürzung der Messzeit führt. Ein weiterer Vorteil ist ihre Robustheit gegenüber Bewegung oder pulsatilem Blutfluss. Die Schwierigkeit im Umgang mit nichtkartesischen Trajektorien liegt vor allem in der Tatsache begründet, dass nichtkartesisch akquirierte Datensätze vor Anwendung der schnellen Fouriertransformation auf ein kartesisches Gitter transformiert werden müssen („Gridding“). Hierzu gibt es eine Vielzahl von Verfahren, die von zahlreichen Parametern abhängen, womit ein hoher Aufwand und hohe Fehleranfälligkeit verbunden sind. Ein weiterer Nachteil dieser Gridding-Methoden ist, dass sie auf unvollständig aufgenommene Datensätze nicht angewendet werden können.

Alternativ zu den konventionellen MR-Verfahren haben sich in den letzten Jahren die sogenannten parallelen Bildgebungsmethoden (beispielsweise SENSE oder GRAPPA) in der klinischen MRT etabliert, die mittlerweile von nahezu allen Herstellerfirmen kommerziell zur Verfügung gestellt werden. Die parallele Bildgebung erlaubt es, die Bildmesszeiten um einen Faktor 2 bis 4 zu verkürzen und lässt sich prinzipiell auf jede beliebige Bilgebungsmethode anwenden ohne dabei das Kontrastverhalten zu beeinflussen. In der klinischen Routine ist diese Technik allerdings lediglich auf kartesische MRT-Verfahren beschränkt, und es ist bisher noch nicht gelungen, die Vorteile der nichtkartesischen MRT-Verfahren optimal mit den Leistungsmerkmalen der parallelen MRT zu verknüpfen.

Ziel dieser Arbeit war es, neue und effiziente Strategien zu entwickeln, um die nichtkartesische Magnetresonanztomographie für ein breiteres Anwendungsspektrum in der klinischen Praxis zu etablieren. Neben der Rekonstruktion von herkömmlich aufgenommenen nichtkartesischen Datensätzen sollten auch Verfahren entwickelt werden, die eine Kombination mit Messzeitverkürzungen durch par-

allele MRT-Verfahren erlauben.

In Kapitel 4 wird ein neues paralleles Bildgebungsverfahren zum Gridding nichtkartesischer Datensätze namens „GRAPPA Operator Gridding“ (GROG) vorgestellt. GROG benutzt GRAPPA-ähnliche Gewichtungsfaktoren, um die nichtkartesischen Punkte auf ein kartesisches Gitter zu schieben. Im Gegensatz zu anderen Gridding-Methoden (wie beispielsweise dem „Convolution-Gridding“) werden bei der Anwendung von GROG Parameter wie Faltungskerne, Regularisierungswerte oder Dichtekompensations-Funktionen nicht benötigt. Dies führt nicht nur zu einer erheblichen Vereinfachung des Griddingprozesses, sondern auch zur deutlichen Reduktion der Rechenoperationen. In Simulationen konnte nachgewiesen werden, dass trotz dieser erheblichen Vereinfachung des Griddingprozesses die Bildqualität im Vergleich zu den aktuellen Goldstandard-Methoden nahezu identisch ist. Besonders deutlich wird dies im Vergleich zum „Convolution-Gridding“ unter Verwendung verschiedener nichtkartesischer Trajektorien (Spiral, Radial, Rosette, PROPELLER). Ein weiterer großer Vorteil von GROG ist darin begründet, dass es mit dieser Methode möglich ist, unvollständig aufgenommene Daten zu Gridden, was andere Methoden wie beispielsweise das „Convolution-Gridding“ nicht erlauben.

In Kapitel 5 wird eine Erweiterung des GROG-Algorithmus vorgestellt, welche ohne Kalibrierungsdatensätze auskommt („Self-Calibrating GROG“, SC-GROG). Die Gewichtungsfaktoren für die Verschiebungen der Datenpunkte werden in dieser Methode aus den akquirierten Punkten selbst gewonnen. Neben der Verkürzung der Aufnahmezeit für die Datensätze führt dies auch zu einer deutlichen Erhöhung der Stabilität des GROG-Algorithmus, da Inkonsistenzen zwischen den Kalibrierungsdaten und den Daten aus den beschleunigten Experimenten entfallen. Um die Vorteile dieser Methode gegenüber den Standardmethoden zu demonstrieren, wurde ein Vergleich zum „Convolution-Gridding“ sowohl in der Simulation als auch im in-vivo Experiment durchgeführt.

Die erste Anwendung von GROG zur Vereinfachung der Rekonstruktion unvollständig aufgenommener nichtkartesischer Datensätze ist in Kapitel 6 beschrieben. Die Verwendung von GROG zur Transformation der unvollständig aufgenommenen nichtkartesischen Daten auf ein kartesisches Gitter erlaubt es, anschließend einen modifizierten GRAPPA-Algorithmus anzuwenden, und somit nichtkartesische Datensätze aus beschleunigten Experimenten zu rekonstruieren. Der besondere Vorteil dieser Methode liegt darin, dass sie zur Rekonstruktion unvollständig aufgenommener Datensätze beliebiger Trajektorien angewendet werden kann, was auf herkömmliche nichtkartesische GRAPPA-Methoden nicht zutrifft. Die große Flexibilität dieses „Pseudo“-kartesischen GRAPPA-Verfahrens konnte in der Anwendung auf verschiedene Trajektorien (Radial, Spiral, Rosette, 1D NC, zick-zack) demonstriert werden.

Schließlich wurde GROG in Kapitel 7 auf die „Bunched Phase Encoding“ (BPE)-Methode angewendet. Bereits zuvor wurde gezeigt, dass das BPE-Verfahren in Verbindung mit einem „Conjugate Gradient“ Rekonstruktionsverfahren eine deutliche Verkürzung der Messzeit gestattet. Basierend auf dem verallgemeinerten Abtasttheorem nach Papoulis werden die Daten bei diesem Verfahren entlang einer extrem schnell oszillierenden Trajektorie aufgenommen. Nach Papoulis ermöglicht die lokal höhere Datendichte eine artefaktfreie Bildrekonstruktion trotz Unterabtastung in anderen k -Raumbereichen. Allerdings werden dabei erhebliche Ansprüche an die Gradienten-Hardware des Tomographen gestellt, wodurch das Konzept auf geringe Beschleunigungsfaktoren beschränkt wird. Im Rahmen dieser Arbeit

konnte jedoch gezeigt werden, dass es möglich ist, auf dieses aufwändige Abtastschema zu verzichten, indem lediglich entlang einer regulären nicht-oszillierenden Trajektorie akquiriert wird und die höhere Datendichte nachträglich mittels GROG erreicht wird (GROG-BPE). In einer Reihe von Simulationen wurde zunächst ein tiefergehendes Verständnis für diese Methode entwickelt, um dieses Verfahren anschließend auf diverse nichtkartesische Trajektorien in in-vivo Experimenten anzuwenden.

Zusammenfassend lässt sich sagen, dass durch die Entwicklung von GROG die Rekonstruktion nichtkartesischer Datensätze erheblich vereinfacht und beschleunigt wurde. Die Verbindung von GROG mit kartesischem GRAPPA oder dem verallgemeinerten Abtasttheorem ermöglicht erstmals eine präzise und stabile k -Raum basierte Bildrekonstruktion von beschleunigten nichtkartesischen MRT-Experimenten mit beliebigen Trajektorien. Die in dieser Arbeit entwickelten Strategien stellen als effiziente und robuste Verfahren somit einen wichtigen Schritt dar, um die Vorteile nichtkartesischer Trajektorien in Kombination mit stark verkürzten Messzeiten für die Praxis nutzbar zu machen.

Bibliography

- [1] D. K. Sodickson and W. J. Manning. Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays. *Magn Reson Med*, 38(4):591–603, Oct 1997.
- [2] P. M. Jakob, M. A. Griswold, R. R. Edelman, and D. K. Sodickson. AUTO-SMASH: a self-calibrating technique for SMASH imaging. SiMultaneous Acquisition of Spatial Harmonics. *MAGMA*, 7(1):42–54, Nov 1998.
- [3] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger. SENSE: sensitivity encoding for fast MRI. *Magn Reson Med*, 42(5):952–962, Nov 1999.
- [4] W. E. Kyriakos, L. P. Panych, D. F. Kacher, C. F. Westin, S. M. Bao, R. V. Mulkern, and F. A. Jolesz. Sensitivity profiles from an array of coils for encoding and reconstruction in parallel (SPACE RIP). *Magn Reson Med*, 44(2):301–308, Aug 2000.
- [5] M. A. Griswold, P. M. Jakob, M. Nittka, J. W. Goldfarb, and A. Haase. Partially parallel imaging with localized sensitivities (PILS). *Magn Reson Med*, 44(4):602–609, Oct 2000.
- [6] R. M. Heidemann, M. A. Griswold, A. Haase, and P. M. Jakob. VD-AUTO-SMASH imaging. *Magn Reson Med*, 45(6):1066–1074, Jun 2001.
- [7] Mark A Griswold, Peter M Jakob, Robin M Heidemann, Mathias Nittka, Vladimir Jellus, Jianmin Wang, Berthold Kiefer, and Axel Haase. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn Reson Med*, 47(6):1202–1210, Jun 2002.
- [8] E.M. Haacke, R.W. Brown, M.R. Thompson, and R. Venkatesan. *Magnetic Resonance Imaging, Physical Principles and Sequence Design*. John Wiley and Sons, 1999.
- [9] E.M Purcell, H.C. Torrey, and R.V. Pound. Resonance absorption by nuclear magnetic moments in a solid. *Physical Review*, 69(1-2):37–38, 1946.
- [10] F. Bloch, Hansen W.W., and M. Packard. Nuclear Induction. *Physical Review*, 69(3-4):127, 1946.
- [11] P. Zeeman. The Effect of Magnetisation on the Nature of Light Emitted by a Substance. *Nature*, 55:347, 1897.
- [12] C. E. Shannon. Communication in the presence of noise. *Proc. Institute of Radio Engineers*, 37:1:10–21, 1949.

- [13] H. Nyquist. Certain topics in telegraph transmission theory. *Trans. AIEE*, 47:617–644, April 1928.
- [14] A. Papoulis. Generalized sampling expansion. *IEEE Trans Circ Syst*, CAS-24:652–654, 1977.
- [15] P. C. Lauterbur. Image formation by induced local interactions. Examples employing nuclear magnetic resonance. 1973. *Clin Orthop Relat Res*, (244):3–6, Jul 1989.
- [16] D. B. Twieg. The k-trajectory formulation of the NMR imaging process with applications in analysis and synthesis of imaging methods. *Med Phys*, 10(5):610–621, 1983.
- [17] A. Kumar, D. Welte, and R.R. Ernst. NMR Fourier zeugmatography. *J Magn Reson*, 18:69–83, 1975.
- [18] W. A. Edelstein, J. M. Hutchison, G. Johnson, and T. Redpath. Spin warp NMR imaging and applications to human whole-body imaging. *Phys Med Biol*, 25(4):751–756, Jul 1980.
- [19] J. Hennig, A. Nauerth, and H. Friedburg. RARE imaging: a fast imaging method for clinical MR. *Magn Reson Med*, 3(6):823–833, Dec 1986.
- [20] P. Mansfield. Multi-planar image-formation using NMR spin echoes. *Journal of Physics C-Solid State Physics*, 10(3):L55–L58, 1977.
- [21] D.C. Noll, D.G. Nishimura, and A. Macovski. Homodyne detection in magnetic resonance imaging. *IEEE Trans Med Imaging*, MI-10(2):154–163, 1991.
- [22] E. M. Haacke, E. D. Lindskog, and W. Lin. A Fast, Iterative Partial-Fourier Technique Capable of Local Phase Recovery,. *J Magn Reson*, 92:126–145, 1991.
- [23] H. Moriguchi, J.L. Sunshine, and J.L. Duerk. Further scan time reduction with bunched phase encoding using sensitivity encoding. In *Proc. Intl Soc Magn Reson Med 13*, page 287, 2005.
- [24] Hisamoto Moriguchi and Jeffrey L Duerk. Bunched phase encoding (BPE): a new fast data acquisition method in MRI. *Magn Reson Med*, 55(3):633–648, Mar 2006.
- [25] H. Moriguchi, M.A. Griswold, J.L. Sunshine, and J.L. Duerk. Bunched phase encoding in projection reconstruction. In *Proc. Intl Soc Magn Reson Med 14*, page 694, 2006.
- [26] G. H. Glover and J. M. Pauly. Projection reconstruction techniques for reduction of motion effects in MRI. *Magn Reson Med*, 28(2):275–289, Dec 1992.
- [27] C. M. Lai and P. C. Lauterbur. True three-dimensional image reconstruction by nuclear magnetic resonance zeugmatography. *Phys Med Biol*, 26(5):851–856, Sep 1981.
- [28] C.B. Ahn, J.H. Kim, and Z.H. Cho. High-speed spiral-scan echo planar NMR imaging. *IEEE Trans Med Imaging*, 5:2–7, 1986.

- [29] C. H. Meyer, B. S. Hu, D. G. Nishimura, and A. Macovski. Fast spiral coronary artery imaging. *Magn Reson Med*, 28(2):202–213, Dec 1992.
- [30] A. Macovski and C. Meyer. A novel fast-scanning system. In *Proceedings of the 5th Annual Meeting of SMRM, Montreal, Canada*, pages 156–157, 1986.
- [31] J. G. Pipe. Motion correction with PROPELLER MRI: application to head motion and free-breathing cardiac imaging. *Magn Reson Med*, 42(5):963–969, Nov 1999.
- [32] M. Busch, A. Bornstedt, M. Wendt, J. L. Duerk, J. S. Lewin, and D. Grönemeyer. Fast "real time" imaging with different k-space update strategies for interventional procedures. *J Magn Reson Imaging*, 8(4):944–954, 1998.
- [33] F.A. Breuer, H. Moriguchi, N. Seiberlich, M. Blaimer and J. Duerk, P.M. Jakob, and M.A. Griswold. Zig-zag sampling for improved parallel imaging. In *Proc ESMRMB*, page 301, 2006.
- [34] C. M. Tsai and D. G. Nishimura. Reduced aliasing artifacts using variable-density k-space sampling trajectories. *Magn Reson Med*, 43(3):452–458, Mar 2000.
- [35] D. C. Noll. Multishot rosette trajectories for spectrally selective MR imaging. *IEEE Trans Med Imaging*, 16(4):372–377, Aug 1997.
- [36] D. C. Noll, S. J. Peltier, and F. E. Boada. Simultaneous multislice acquisition using rosette trajectories (SMART): a new imaging method for functional MRI. *Magn Reson Med*, 39(5):709–716, May 1998.
- [37] S. J. Peltier and D. C. Noll. Systematic noise compensation for simultaneous multislice acquisition using rosette trajectories (SMART). *Magn Reson Med*, 41(5):1073–1076, May 1999.
- [38] K. Scheffler and J. Hennig. Frequency resolved single-shot MR imaging using stochastic k-space trajectories. *Magn Reson Med*, 35(4):569–576, Apr 1996.
- [39] R.S. Likes. Moving gradient zeugmatography. Technical report, U.S. Patent 4307343, 1981.
- [40] H. T. Nielsen, E. W. Olcott, and D. G. Nishimura. Improved 2D time-of-flight angiography using a radial-line k-space acquisition. *Magn Reson Med*, 37(2):285–291, Feb 1997.
- [41] Gordon E Sarty. Single Trajectory Radial (STAR) imaging. *Magn Reson Med*, 51(3):445–451, Mar 2004.
- [42] P. Irarrazabal and D. G. Nishimura. Fast three dimensional magnetic resonance imaging. *Magn Reson Med*, 33(5):656–662, May 1995.
- [43] Edward Brian Welch, Armando Manduca, Roger C Grimm, Heidi A Ward, and Clifford R Jack. Spherical navigator echoes for full 3D rigid body motion measurement in MRI. *Magn Reson Med*, 47(1):32–41, Jan 2002.

- [44] M. L. Lauzon and B. K. Rutt. Effects of polar sampling in k-space. *Magn Reson Med*, 36(6):940–949, Dec 1996.
- [45] K. Scheffler and J. Hennig. Reduced circular field-of-view imaging. *Magn Reson Med*, 40(3):474–480, Sep 1998.
- [46] C. J. Bergin, G. M. Glover, and J. Pauly. Magnetic resonance imaging of lung parenchyma. *J Thorac Imaging*, 8(1):12–17, 1993.
- [47] S. L. Gewalt, G. H. Glover, L. W. Hedlund, G. P. Cofer, J. R. MacFall, and G. A. Johnson. MR microscopy of the rat lung using projection reconstruction. *Magn Reson Med*, 29(1):99–106, Jan 1993.
- [48] M. D. Shattuck, S. L. Gewalt, G. H. Glover, L. W. Hedlund, and G. A. Johnson. MR microimaging of the lung using volume projection encoding. *Magn Reson Med*, 38(6):938–942, Dec 1997.
- [49] G. E. Gold, D. R. Thedens, J. M. Pauly, K. P. Fechner, G. Bergman, C. F. Beaulieu, and A. Macovski. MR imaging of articular cartilage of the knee: new methods using ultrashort TEs. *AJR Am J Roentgenol*, 170(5):1223–1226, May 1998.
- [50] M. A. Hall-Craggs, J. Porter, P. D. Gatehouse, and G. M. Bydder. Ultrashort echo time (UTE) MRI of the spine in thalassaemia. *Br J Radiol*, 77(914):104–110, Feb 2004.
- [51] Michael Benjamin and Graeme M Bydder. Magnetic resonance imaging of entheses using ultrashort TE (UTE) pulse sequences. *J Magn Reson Imaging*, 25(2):381–389, Feb 2007.
- [52] Ines L H Reichert, Michael Benjamin, Peter D Gatehouse, Karyn E Chappell, Joanne Holmes, Taigang He, and Graeme M Bydder. Magnetic resonance imaging of periosteum with ultrashort TE pulse sequences. *J Magn Reson Imaging*, 19(1):99–107, Jan 2004.
- [53] M. D. Robson, M. Benjamin, P. Gishen, and G. M. Bydder. Magnetic resonance imaging of the Achilles tendon using ultrashort TE (UTE) pulse sequences. *Clin Radiol*, 59(8):727–735, Aug 2004.
- [54] R. Jerecic, M. Bock, S. Nielles-Vallespin, C. Wacker, W. Bauer, and Lothar R Schad. ECG-gated ^{23}Na -MRI of the human heart using a 3D-radial projection technique with ultra-short echo times. *MAGMA*, 16(6):297–302, May 2004.
- [55] Sònia Nielles-Vallespin, Marc-André Weber, Michael Bock, André Bongers, Peter Speier, Stephanie E Combs, Johannes Wöhrle, Frank Lehmann-Horn, Marco Essig, and Lothar R Schad. 3D radial projection technique with ultrashort echo times for sodium MRI: clinical applications in human brain and skeletal muscle. *Magn Reson Med*, 57(1):74–81, Jan 2007.
- [56] M-A. Weber, S. Nielles-Vallespin, M. Essig, K. Jurkat-Rott, H-U. Kauczor, and F. Lehmann-Horn. Muscle Na^+ channelopathies: MRI detects intracellular ^{23}Na accumulation during episodic weakness. *Neurology*, 67(7):1151–1158, Oct 2006.

- [57] A. T. Lee, G. H. Glover, and C. H. Meyer. Discrimination of large venous vessels in time-course spiral blood-oxygen-level-dependent magnetic-resonance functional neuroimaging. *Magn Reson Med*, 33(6):745–754, Jun 1995.
- [58] G. H. Glover, S. K. Lemieux, M. Drangova, and J. M. Pauly. Decomposition of inflow and blood oxygen level-dependent (BOLD) effects with dual-echo spiral gradient-recalled echo (GRE) fMRI. *Magn Reson Med*, 35(3):299–308, Mar 1996.
- [59] Gary H Glover and Moriah E Thomason. Improved combination of spiral-in/out images for BOLD fMRI. *Magn Reson Med*, 51(4):863–868, Apr 2004.
- [60] Allen Song, Karen Emberger, Charles Michelich, and Gregory McCarthy. FMRI signal source analysis using diffusion-weighted spiral-in acquisition. *Conf Proc IEEE Eng Med Biol Soc*, 6:4417–4420, 2004.
- [61] T. Schaeffter, S. Weiss, H. Eggers, and V. Rasche. Projection reconstruction balanced fast field echo for interactive real-time cardiac imaging. *Magn Reson Med*, 46(6):1238–1241, Dec 2001.
- [62] A. Shankaranarayanan, O. P. Simonetti, G. Laub, J. S. Lewin, and J. L. Duerk. Segmented k-space and real-time cardiac cine MR imaging with radial trajectories. *Radiology*, 221(3):827–836, Dec 2001.
- [63] A. C. Larson and O. P. Simonetti. Real-time cardiac cine imaging with SPIDER: steady-state projection imaging with dynamic echo-train readout. *Magn Reson Med*, 46(6):1059–1066, Dec 2001.
- [64] Daniel T Boll, Elmar M Merkle, Danielle M Seaman, Robert C Gilkeson, Andrew P Larson, Orlando P Simonetti, Jeffrey L Duerk, and Jonathan S Lewin. Comparison of ECG-gated rectilinear vs. real-time radial K-space sampling schemes in cine True-FISP cardiac MRI. *J Cardiovasc Magn Reson*, 6(4):793–802, 2004.
- [65] Elmar Spuentrup, Andreas H Mahnken, Harald P Köhl, Gabriele A Krombach, Rene M Botnar, Alexander Wall, Tobias Schaeffter, Rolf W Günther, and Arno Buecker. Fast interactive real-time magnetic resonance imaging of cardiac masses using spiral gradient echo and radial steady-state free precession sequences. *Invest Radiol*, 38(5):288–292, May 2003.
- [66] Joao L A Carvalho and Krishna S Nayak. Rapid quantitation of cardiovascular flow using slice-selective fourier velocity encoding with spiral readouts. *Magn Reson Med*, 57(4):639–646, Apr 2007.
- [67] Marcus Katoh, Elmar Spuentrup, Arno Buecker, Tobias Schaeffter, Matthias Stuber, Rolf W Günther, and Rene M Botnar. MRI of coronary vessel walls using radial k-space sampling and steady-state free precession imaging. *AJR Am J Roentgenol*, 186(6 Suppl 2):S401–S406, Jun 2006.

- [68] Andrew C Larson, Orlando P Simonetti, and Debiao Li. Coronary MRA with 3D undersampled projection reconstruction TrueFISP. *Magn Reson Med*, 48(4):594–601, Oct 2002.
- [69] Xiaoming Bi, Jaeseok Park, Andrew C Larson, Qiang Zhang, Orlando Simonetti, and Debiao Li. Contrast-enhanced 4D radial coronary artery imaging at 3.0 T within a single breath-hold. *Magn Reson Med*, 54(2):470–475, Aug 2005.
- [70] Ty A Cashen, Hyun Jeong, Maulin K Shah, Hem M Bhatt, Wanyong Shin, James C Carr, Matthew T Walker, H. Hunt Batjer, and Timothy J Carroll. 4D radial contrast-enhanced MR angiography with sliding subtraction. *Magn Reson Med*, 58(5):962–972, Nov 2007.
- [71] Kirsten P Forbes, James G Pipe, John P Karis, Victoria Farthing, and Joseph E Heiserman. Brain imaging in the unsedated pediatric patient: comparison of periodically rotated overlapping parallel lines with enhanced reconstruction and single-shot fast spin-echo sequences. *AJNR Am J Neuroradiol*, 24(5):794–798, May 2003.
- [72] Ashok Panigrahy and Stefan Blüml. Advances in magnetic resonance neuroimaging techniques in the evaluation of neonatal encephalopathy. *Top Magn Reson Imaging*, 18(1):3–29, Feb 2007.
- [73] Osamu Abe, Harushi Mori, Shigeki Aoki, Akira Kunimatsu, Naoto Hayashi, Tomohiko Masumoto, Haruyasu Yamada, Yoshitaka Masutani, Hiroyuki Kabasawa, and Kuni Ohtomo. Periodically rotated overlapping parallel lines with enhanced reconstruction-based diffusion tensor imaging. Comparison with echo planar imaging-based diffusion tensor imaging. *J Comput Assist Tomogr*, 28(5):654–660, 2004.
- [74] Fu-Nien Wang, Teng-Yi Huang, Fa-Hsuan Lin, Tzu-Chao Chuang, Nan-Kuei Chen, Hsiao-Wen Chung, Cheng-Yu Chen, and Kenneth K Kwong. PROPELLER EPI: an MRI technique suitable for diffusion tensor imaging at high field strength with reduced geometric distortions. *Magn Reson Med*, 54(5):1232–1240, Nov 2005.
- [75] Jie Deng, Frank H Miller, Riad Salem, Reed A Omary, and Andrew C Larson. Multishot diffusion-weighted PROPELLER magnetic resonance imaging of the abdomen. *Invest Radiol*, 41(10):769–775, Oct 2006.
- [76] J.I. Jackson, C.H. Meyer, D.G. Nishimura, and A. Macovski. Selection of a convolution function for Fourier inversion using gridding. *IEEE Trans Med Imaging*, 10:473–478, 1991.
- [77] J.D. O’Sullivan. A fast sinc function gridding algorithm for Fourier inversion in computer tomography. *IEEE Trans Med Imaging*, 4:200–207, 1985.
- [78] D. Rosenfeld. An optimal and efficient new gridding algorithm using singular value decomposition. *Magn Reson Med*, 40(1):14–23, Jul 1998.
- [79] Daniel Rosenfeld. New approach to gridding using regularization and estimation theory. *Magn Reson Med*, 48(1):193–202, Jul 2002.

- [80] Hisamoto Moriguchi and Jeffrey L Duerk. Iterative Next-Neighbor Regridding (INNG): improved reconstruction from nonuniformly sampled k-space data using rescaled matrices. *Magn Reson Med*, 51(2):343–352, Feb 2004.
- [81] Refaat E Gabr, Pelin Aksit, Paul A Bottomley, Abou-Bakr M Youssef, and Yasser M Kadah. Deconvolution-interpolation gridding (DING): accurate reconstruction for arbitrary k-space trajectories. *Magn Reson Med*, 56(6):1182–1191, Dec 2006.
- [82] G. E. Sarty, R. Bennett, and R. W. Cox. Direct reconstruction of non-Cartesian k-space data using a nonuniform fast Fourier transform. *Magn Reson Med*, 45(5):908–915, May 2001.
- [83] A Dutt and V Rokhlin. Fast Fourier transforms for nonequispaced data. *SIAM J Sci Comput*, 14:1368–1383, 1993.
- [84] Liewei Sha, Hua Guo, and Allen W Song. An improved gridding method for spiral MRI using nonuniform fast Fourier transform. *J Magn Reson*, 162(2):250–258, Jun 2003.
- [85] Jeffrey A Fessler. On NUFFT-based gridding for non-Cartesian MRI. *J Magn Reson*, 188(2):191–195, Oct 2007.
- [86] G.C. Wiggins, A. Potthast and C. Triantafyllou, F.-H. Lin, T. Benner, C.J. Wiggins, and L.L.Wald. A 96-channel MRI system with 23- and 90-channel phase array head coils at 1.5 Tesla. In *Proc Intl Soc Magn Reson Med 13*, page 671, 2005.
- [87] G. C. Wiggins, C. Triantafyllou, A. Potthast, A. Reykowski, M. Nittka, and L. L. Wald. 32-channel 3 Tesla receive-only phased-array head coil with soccer-ball element geometry. *Magn Reson Med*, 56(1):216–223, Jul 2006.
- [88] J. Kelton, R.M. Magin, and S.M. Wright. An algorithm for rapid image acquisition using multiple receiver coils. In *Proc Intl Soc Magn Reson Med 8*, page 1172, 1989.
- [89] J. B. Ra and C. Y. Rim. Fast imaging using subencoding data sets from multiple detectors. *Magn Reson Med*, 30(1):142–145, Jul 1993.
- [90] J. W. Carlson and T. Minemura. Imaging time reduction through multiple receiver coil data acquisition and image reconstruction. *Magn Reson Med*, 29(5):681–687, May 1993.
- [91] Martin Blaimer, Felix Breuer, Matthias Mueller, Robin M Heidemann, Mark A Griswold, and Peter M Jakob. SMASH, SENSE, PILS, GRAPPA: how to choose the optimal method. *Top Magn Reson Imaging*, 15(4):223–236, Aug 2004.
- [92] Robin M Heidemann, Nicole Seiberlich, Mark A Griswold, Katrin Wohlfarth, Gunnar Krueger, and Peter M Jakob. Perspectives and limitations of parallel MR imaging at high field strengths. *Neuroimaging Clin N Am*, 16(2):311–20, xi, May 2006.

- [93] P. B. Roemer, W. A. Edelstein, C. E. Hayes, S. P. Souza, and O. M. Mueller. The NMR phased array. *Magn Reson Med*, 16(2):192–225, Nov 1990.
- [94] D. O. Walsh, A. F. Gmitro, and M. W. Marcellin. Adaptive reconstruction of phased array MR imagery. *Magn Reson Med*, 43(5):682–690, May 2000.
- [95] Mark A Griswold, Martin Blaimer, Felix Breuer, Robin M Heidemann, Matthias Mueller, and Peter M Jakob. Parallel magnetic resonance imaging using the GRAPPA operator formalism. *Magn Reson Med*, 54(6):1553–1556, Dec 2005.
- [96] M. Blaimer, F.A. Breuer, R.M. Heidemann, M. Mueller, M.A. Griswold, and P.M. Jakob. Is parallel MRI without a prior information possible? In *Proc. Intl Soc Magn Reson Med 12*, page 2417, 2004.
- [97] Martin Blaimer, Felix A Breuer, Matthias Mueller, Nicole Seiberlich, Dmitry Ebel, Robin M Heidemann, Mark A Griswold, and Peter M Jakob. 2D-GRAPPA-operator for faster 3D parallel MRI. *Magn Reson Med*, 56(6):1359–1364, Dec 2006.
- [98] M. Blaimer, K. Barkauskas, S. Kannengiesser, F. Breuer, P.M. Jakob, J.L. Duerk, and M.A. Griswold. Artifact reduction in undersampled BLADE/PROPELLER MRI by k-space extrapolation using parallel imaging. In *Proc. Intl Soc Magn Reson Med 14*, page 5, 2006.
- [99] James W Goldfarb. The SENSE ghost: field-of-view restrictions for SENSE imaging. *J Magn Reson Imaging*, 20(6):1046–1051, Dec 2004.
- [100] M.A. Griswold, R.M. Heidemann, and P.M. Jakob. Direct parallel imaging reconstruction of radially sampled data using GRAPPA with relative shifts. In *Proc Intl Soc Meg Reson Med 11*, page 2349, 2003.
- [101] Arjun Arunachalam, Alexey Samsonov, and Walter F Block. Self-calibrated GRAPPA method for 2D and 3D radial data. *Magn Reson Med*, 57(5):931–938, May 2007.
- [102] Robin M Heidemann, Mark A Griswold, Nicole Seiberlich, Gunnar Krüger, Stephan A R Kannengiesser, Berthold Kiefer, Graham Wiggins, Lawrence L Wald, and Peter M Jakob. Direct parallel image reconstructions for spiral trajectories using GRAPPA. *Magn Reson Med*, 56(2):317–326, Aug 2006.
- [103] Robin M Heidemann, Mark A Griswold, Nicole Seiberlich, Mathias Nittka, Stephan A R Kannengiesser, Berthold Kiefer, and Peter M Jakob. Fast method for 1D non-cartesian parallel imaging using GRAPPA. *Magn Reson Med*, 57(6):1037–1046, Jun 2007.
- [104] K. P. Pruessmann, M. Weiger, P. Börnert, and P. Boesiger. Advances in sensitivity encoding with arbitrary k-space trajectories. *Magn Reson Med*, 46(4):638–651, Oct 2001.

- [105] Ernest N Yeh, Charles A McKenzie, Michael A Ohliger, and Daniel K Sodickson. Parallel magnetic resonance imaging with adaptive radius in k-space (PARS): constrained image reconstruction using k-space locality in radiofrequency coil encoded data. *Magn Reson Med*, 53(6):1383–1392, Jun 2005.
- [106] Alexey A Samsonov, Walter F Block, Arjun Arunachalam, and Aaron S Field. Advances in locally constrained k-space-based parallel MRI. *Magn Reson Med*, 55(2):431–438, Feb 2006.
- [107] S.A.R. Kannengiesser and T.G. Noll. Towards a Practical Generalized Image Reconstruction Method for MRI. In *Proc Intl Soc Magn Reson Med 10*, page 155, 2002.
- [108] MR Hestenes and E Stiefel. Methods of conjugate gradients for solving linear systems. *Natl Bur Standards J Res*, 49:409–436, 1952.
- [109] J Stoer and R Bulirsch. *Numerische Mathematik 2*. Springer, 3rd ed. edition, 1990.
- [110] Peng Qu, Kai Zhong, Bida Zhang, Jianmin Wang, and Gary X Shen. Convergence behavior of iterative SENSE reconstruction with non-Cartesian trajectories. *Magn Reson Med*, 54(4):1040–1045, Oct 2005.
- [111] Reza Nezafat, Peter Kellman, John Derbyshire, and Elliot McVeigh. Real time high spatial-temporal resolution flow imaging with spiral MRI using auto-calibrated SENSE. *Conf Proc IEEE Eng Med Biol Soc*, 3:1914–1917, 2004.
- [112] Reza Nezafat, Peter Kellman, J. Andrew Derbyshire, and Elliot R McVeigh. Real-time blood flow imaging using autocalibrated spiral sensitivity encoding. *Magn Reson Med*, 54(6):1557–1561, Dec 2005.
- [113] Jaeseok Park, Andrew C Larson, Qiang Zhang, Orlando Simonetti, and Debiao Li. 4D radial coronary artery imaging within a single breath-hold: cine angiography with phase-sensitive fat suppression (CAPS). *Magn Reson Med*, 54(4):833–840, Oct 2005.
- [114] Markus Weiger, Klaas P Pruessmann, Robert Osterbauer, Peter Börnert, Peter Boesiger, and Peter Jezzard. Sensitivity-encoded single-shot spiral imaging for reduced susceptibility artifacts in BOLD fMRI. *Magn Reson Med*, 48(5):860–866, Nov 2002.
- [115] Trong-Kha Truong and Allen W Song. Single-shot dual-z-shimmed sensitivity-encoded spiral-in/out imaging for functional MRI with reduced susceptibility artifacts. *Magn Reson Med*, Nov 2007.
- [116] Chunlei Liu, Michael E Moseley, and Roland Bammer. Simultaneous phase correction and SENSE reconstruction for navigated multi-shot DWI with non-cartesian k-space sampling. *Magn Reson Med*, 54(6):1412–1422, Dec 2005.
- [117] J. G. Pipe and P. Menon. Sampling density compensation in MRI: rationale and an iterative numerical solution. *Magn Reson Med*, 41(1):179–186, Jan 1999.

- [118] N Seiberlich, FA Breuer, M Blaimer, K Barkauskas, PM Jakob, and MA Griswold. Non-Cartesian data reconstruction using GRAPPA operator gridding (GROG). *Magn Reson Med*, 58(6):1257–1265, Dec 2007.
- [119] N. Seiberlich, M. Blaimer, K.J. Barkauskas, F.A. Breuer, P.M Jakob, and M.A. Griswold. Non-Cartesian Data Gridding using GRAPPA Operator Gridding (GROG). In *Proc. ESMRMB*, page 300, 2006.
- [120] J. G. Pipe. Reconstructing MR images from undersampled data: data-weighting considerations. *Magn Reson Med*, 43(6):867–875, Jun 2000.
- [121] K.J. Barkauskas, M.A. Griswold, J.L. Sunshine, and J.L. Duerk. Density Compensation for TrueFISP BLADE Imaging in the Steady-State (TrueBLISS). In *Proc Intl Soc Magn Reson Med 14*, page 2949, 2006.
- [122] R. D. Hoge, R. K. Kwan, and G. B. Pike. Density compensation functions for spiral MRI. *Magn Reson Med*, 38(1):117–128, Jul 1997.
- [123] N Seiberlich, FA Breuer, M Blaimer, PM Jakob, and MA Griswold. Self-Calibrating GRAPPA Operator Gridding for Radial and Spiral Trajectories. *Magn Reson Med*, 59(4):930–5, Apr 2008.
- [124] N. Seiberlich, F.A. Breuer, M. Blaimer, P.M. Jakob, and M.A. Griswold. Self-Calibrated GRAPPA Operator Gridder (SC-GROG). In *Proc Intl Soc Magn Reson Med 15*, page 153, 2007.
- [125] N Seiberlich, P Ehses, S Nilles-Vallespin, FA Breuer, M Blaimer, PM Jakob, and MA Griswold. Self-Calibrating Gridding for 3D Radial Trajectories Using GRAPPA Operator Gridding (GROG). In *Proc Intl Soc Magn Reson Med 16*, page 1281, 2008.
- [126] Keith Heberlein and Xiaoping Hu. Auto-calibrated parallel spiral imaging. *Magn Reson Med*, 55(3):619–625, Mar 2006.
- [127] N Seiberlich, FA Breuer, RM Heidemann, M Blaimer, MA Griswold, and PM Jakob. Reconstruction of Undersampled Non-Cartesian Datasets using Pseudo-Cartesian GRAPPA in Conjunction with GROG. *Magn Reson Med*, 59(5):1127–37, May 2008.
- [128] N. Seiberlich, R.M. Heidemann, F.A. Breuer, M. Blaimer, M.A. Griswold, and P.M. Jakob. Pseudo-Cartesian GRAPPA Reconstruction of Undersampled Non-Cartesian Data. In *Proc Intl Soc Magn Reson Med 14*, page 2463, 2006.
- [129] N. Seiberlich, F.A. Breuer, R.M. Heidemann, M. Blaimer, M.A. Griswold, and P.M. Jakob. Reconstruction of Arbitrary Non-Cartesian Trajectories using Pseudo-Cartesian GRAPPA in Conjunction with GRAPPA Operator Gridding (GROG). In *Proc Intl Soc Magn Reson Med 15*, page 334, 2007.

- [130] N. Seiberlich, F.A. Breuer, H. Moriguchi, P.M. Jakob, and M.A. Griswold. Fully Autocalibrated Parallel Imaging for Arbitrary Trajectories Using a Combination of GRAPPA Operator Gridding and Conjugate-Gradient Optimization. In *Proc Intl Soc Magn Reson Med 15*, page 1742, 2007.
- [131] N Seiberlich, P Ehses, FA Breuer, M Blaimer, PM Jakob, and MA Griswold. Reconstruction of Undersampled Non-Cartesian Data using GROG-Facilitated Random Blipped Phase Encoding. In *Proc Intl Soc Magn Reson Med 16*, page 7, 2008.
- [132] CA Bookwalter, MA Griswold, JL Sunshine, and JL Duerk. Multiple Overlapping k-space Junctions for Investigating Translating Objects (MOJITO). In *Proc. Intl Soc Magn Reson Med 15*, page 3425, 2007.
- [133] A. Fischer, F. Breuer, M. Blaimer, N. Seiberlich, , and P. M. Jakob. Introduction of a Nonconvex Compressed Sensing Algorithm for MR Imaging. In *Proc. Intl Soc Magn Reson Med 16*, page 1487, 2008.
- [134] A. Fischer, F. Breuer, M. Blaimer, N. Seiberlich, , and P. M. Jakob. Accelerated Dynamic Imaging by Reconstructing Sparse Differences using Compressed Sensing. In *Proc. Intl Soc Magn Reson Med 16*, page 341, 2008.
- [135] H. T. Nielsen, G. E. Gold, E. W. Olcott, J. M. Pauly, and D. G. Nishimura. Ultra-short echo-time 2D time-of-flight MR angiography using a half-pulse excitation. *Magn Reson Med*, 41(3):591–599, Mar 1999.
- [136] P. Hu and C.H. Meyer. BOSCO: Parallel Image Reconstruction Based on Successive Convolution Operations. In *Proc. Intl Soc Magn Reson Med 14*, page 10, 2006.
- [137] Jr. Brown, J.L., Jr. Brown, J.L., and S.D. Cabrera. On well-posedness of the Papoulis generalized sampling expansion. 38(5):554–556, 1991.
- [138] J. H. Duyn, Y. Yang, J. A. Frank, and J. W. van der Veen. Simple correction method for k-space trajectory deviations in MRI. *J Magn Reson*, 132(1):150–153, May 1998.

CURRICULUM VITAE

Nicole Seiberlich

PERSONAL INFORMATION

Address: Marktgasse 1 *Telephone:* +49 176 225 18525
97070 Würzburg *Email:* nseiberlich@hotmail.com
Germany *Work Email:* neseiber@physik.uni-wuerzburg.de

Date of birth: 22 January 1980
Nationality: American
Languages: English (native), German (fluent), French (conversational)

EDUCATION

Graduate Education 10/2003-02/2008 University of Würzburg, Würzburg, Germany
Undergraduate Education 08/1997-05/2001 Yale University, New Haven, CT, USA
GPA: 3.93 / 4.0
Bachelor of Science in Chemistry
Graduated Summa cum Laude, 2001
Distinction in the Chemistry Major, 2001
Fleischer Award, Excellence in Chemistry, 2001
Phi Beta Kappa Honors Society, 2001

High School Education 08/1993-05/1997 Divine Savior Holy Angels, Milwaukee, WI, USA
GPA: 3.97 / 4.0
Valedictorian, 1997

RESEARCH AND WORK EXPERIENCE

Prof. Peter Jakob 11/2004-02/2008 Graduate Studies, Topic: Non-Cartesian Parallel Imaging
University of Würzburg 10/2003-10/2004 Pre-Graduate Studies (Diplom)
Thesis, "Sensitivity Advantages of Chemical Shift
Imaging in High-Resolution MRI"

Oliver, Wyman & Co. 07/2001-09/2003 Management Consultant
Frankfurt 08/2002-09/2003 Specialized in Risk Quantification and Management
New York 07/2001-08/2002 in major European and US Universal banks

Prof. Gary Brudvig 08/2000-05/2001 Bachelor's Thesis, "Protein Structural Determination using
Yale University Relaxation Effects in Electron Paramagnetic
Resonance Spectroscopy"

Prof. David Cory 06/2000-08/2000 NMR applications of Quantum Information Processing,
Massachusetts Institute of Technology specifically quantum teleportation using alanine

Prof. Kurt Zilm 01/1999-05/2000 Coupled amino acids using solid-phase synthesis to form
Yale University tri-peptides, built a solid-state NMR probe

Würzburg, Germany, 27 February 2008

SELECTED PUBLICATIONS

PAPERS

Breuer FA, Hisamoto M, Seiberlich N, Blaimer M, Jakob PM, Duerk JL, Griswold MA.
Zigzag sampling for improved parallel imaging.
Magn Reson Med. 2008 Jul 29;60(2):474-478.

Seiberlich N, Breuer FA, Heidemann RM, Blaimer M, Griswold MA, Jakob PM.
Reconstruction of Undersampled Non-Cartesian Datasets using Pseudo-Cartesian GRAPPA in Conjunction with GROG.
Magn Reson Med. 2008 May;59(5):1127-37.

Seiberlich N, Breuer FA, Blaimer M, Jakob PM, Griswold MA.
Self-Calibrating GRAPPA Operator Gridding for Radial and Spiral Trajectories.
Magn Reson Med. 2008 Apr;59(4):930-5.

Kaufmann I, Seiberlich N, Haase A, Jakob P.
Diffusion Generated T1 and T2 Contrast.
J Magn Reson. 2008 May;192(1):139-50.

Seiberlich N, Breuer FA, Blaimer M, Barkauskas K, Jakob PM, Griswold MA.
Non-Cartesian Data Reconstruction using GRAPPA Operator Gridding (GROG).
Magn Reson Med. 2007 Dec;58(6):1257-65.

Pracht ED, Arnold JFT, Seiberlich N, Kotas M, Flentje M, Jakob PM.
"Oxygen-enhanced proton magnetic resonance imaging of the human lung"
CANCER IMAGING Vol. 1 edited by M. A. Hayat, Elsevier Science/Academic, 2008 Pg. 267-289

Heidemann RM, Griswold MA, Seiberlich N, Nittka M, Kannengiesser SA, Kiefer B, Jakob PM.
Fast method for 1D non-cartesian parallel imaging using GRAPPA.
Magn Reson Med. 2007 Jun;57(6):1037-46.

Blaimer M, Breuer FA, Mueller M, Seiberlich N, Ebel D, Heidemann RM, Griswold MA, Jakob PM.
2D-GRAPPA-operator for faster 3D parallel MRI.
Magn Reson Med. 2006 Dec;56(6):1359-64.

Heidemann RM, Griswold MA, Seiberlich N, Kruger G, Kannengiesser SA, Kiefer B, Wiggins G, Wald LL, Jakob PM.
Direct parallel image reconstructions for spiral trajectories using GRAPPA.
Magn Reson Med. 2006 Aug;56(2):317-26.

Blaimer M, Breuer FA, Seiberlich N, Mueller MF, Heidemann RM, Jellus V, Wiggins G, Wald LL, Griswold MA, Jakob PM.

Accelerated volumetric MRI with a SENSE/GRAPPA combination.

J Magn Reson Imaging. 2006 Aug;24(2):444-50.

Heidemann RM, Seiberlich N, Griswold MA, Wohlfarth K, Krueger G, Jakob PM.

Perspectives and Limitations of Parallel MR Imaging at High Field Strengths.

Neuroimaging Clin N Am. 2006 May;16(2):311-20.

Breuer FA, Blaimer M, Mueller MF, Seiberlich N, Heidemann RM, Griswold MA, Jakob PM.

Controlled aliasing in volumetric parallel imaging (2D CAIPIRINHA).

Magn Reson Med. 2006 Mar;55(3):549-56.

PRESENTATIONS

Seiberlich N, Ehses P, Breuer FA, Blaimer M, Jakob PM, Griswold MA.

Reconstruction of Undersampled Non-Cartesian Data using GROG-Facilitated Random Blipped Phase Encoding.

Proc Intl Soc Mag Reson Med 16 (2008), Pg. 7

Breuer F, Blaimer M, Seiberlich N, Jakob PM, Griswold MA.

A general formulation for quantitative g-factor calculation in GRAPPA reconstructions.

Proc Intl Soc Mag Reson Med 16 (2008), Pg. 10

Fischer A, Breuer F, Blaimer M, Seiberlich N, Jakob PM.

Accelerated Dynamic Imaging by Reconstructing Sparse Differences using Compressed Sensing.

Proc Intl Soc Mag Reson Med 16 (2008), Pg. 341

Seiberlich N, Breuer FA, Blaimer M, Jakob PM, Griswold MA.

Self-Calibrated GRAPPA Operator Gridder (SC-GROG).

Proc Intl Soc Mag Reson Med 15 (2007), Pg. 153

Seiberlich N, Breuer FA, Heidemann RM, Blaimer M, Griswold MA, Jakob PM.

Reconstruction of Arbitrary Non-Cartesian Trajectories using Pseudo-Cartesian GRAPPA in Conjunction with GRAPPA Operator Gridding (GROG).

Proc Intl Soc Mag Reson Med 15 (2007), Pg. 334

Seiberlich N, Blaimer M, Barkauskas K, Breuer FA, Jakob PM, Griswold MA.

Non-Cartesian Data Gridding using GRAPPA Operator Gridding (GROG).

Proceedings of the ESMRMB (2006), Pg. 300

Bauer S, Seiberlich N, Breuer FA, Blaimer M, Griswold MA, Lin FH, Wald LL, Jakob PM.
Parallel imaging with high acceleration factors using a 90-channel array coil.
Proceedings of the ESMRMB (2006), Pg. 305

Breuer FA, Moriguchi H, Seiberlich N, Blaimer M, Duerk J, Jakob PM, Griswold MA.
Zig-zag sampling for improved parallel imaging.
Proceedings of the ESMRMB (2006), Pg. 301

Seiberlich N, Breuer FA, Blaimer M, Speier P, Griswold MA, Jakob PM.
3D Cylindrical GRAPPA.
Proc Intl Soc Mag Reson Med 14 (2006), Pg. 7

POSTERS

Seiberlich N, Ehses P, Nielles-Vallespin S, Breuer FA, Blaimer M, Jakob PM, Griswold MA.
Self-Calibrating Gridding for 3D Radial Trajectories Using GRAPPA Operator Gridding (GROG).
Proc Intl Soc Mag Reson Med 16 (2008), Pg. 1281

Fischer A, Breuer F, Blaimer M, Seiberlich N, Jakob PM.
Introduction of a Nonconvex Compressed Sensing Algorithm for MR Imaging.
Proc Intl Soc Mag Reson Med 16 (2008), Pg. 1487

Ehses P, Seiberlich N, Nordbeck P, Fidler F, Jakob PM, Bauer WR.
Multi-spiral MRI for cardiac T2-star determination
Proc Intl Soc Mag Reson Med 16 (2008), Pg. 1015

Seiberlich N, Breuer FA, Moriguchi H, Jakob PM, Griswold MA.
Fully Autocalibrated Parallel Imaging for Arbitrary Trajectories Using a Combination of GRAPPA Operator Gridding and Conjugate-Gradient Optimization.
Proc Intl Soc Mag Reson Med 15 (2007), Pg. 1742

Lopez ML, Breuer F, Seiberlich N, Gareis D, S Voll, PM Jakob.
4 Channel Mouse Array at 7 Tesla.
Proc Intl Soc Mag Reson Med 15 (2007), Pg. 1026

Seiberlich N, Heidemann RM, Breuer FA, Blaimer M, Griswold MA, Jakob PM.
Pseudo-Cartesian GRAPPA Reconstruction of Undersampled Non-Cartesian Data.
Proc Intl Soc Mag Reson Med 14 (2006), Pg. 2463

Breuer FA, Blaimer M, Seiberlich N, Griswold MA, Jakob PM.
A 3D GRAPPA Algorithm for Volumetric Parallel Imaging.
Proc Intl Soc Mag Reson Med 14 (2006), Pg. 286

Breuer FA, Ebel D, Ruff J, Blaimer M, Seiberlich N, Griswold MA, Jakob PM.
Parallel 2D and 3D Spectroscopic Imaging Using GRAPPA.
Proc Intl Soc Mag Reson Med 14 (2006), Pg. 3653

Breuer FA, Blaimer M, Seiberlich N, Griswold MA, Jakob PM.
The Use of Coil Sensitivity Variations in the Read-Direction for Improved Parallel Imaging.
Proc Intl Soc Mag Reson Med 14 (2006), Pg. 2460

Seiberlich N, Haddad D, Griswold MA, Webb A, Jakob PM.
Sensitivity Advantages of Chemical Shift Imaging in Magnetic Resonance Microscopy.
Proceedings of the ICMRM 2005, Pg. P14

ACKNOWLEDGEMENTS

This work would not have been possible without a great number of people, including...

...Prof. Dr. Peter Jakob, for taking a chance on an unknown chemistry student, and inviting me to work on my PhD in EP5. His insight into MRI and parallel imaging have had an enormous impact on my research, and I am grateful to have worked with him.

...Prof. Dr. Mark Griswold, who has supported and guided me, both professionally and personally, throughout my time in Wuerzburg. Mark's enthusiasm and interest in my work have pushed me to perform at my best, and for that I cannot thank him enough. I am sincerely looking forward to continuing our work together in Cleveland.

...the PPA team in EP5, including Felix Breuer, Martin Blaimer, Robin Heidemann, Andre Fischer, Matthias Mueller, Daniel Gareis, Marcos Lopez, and Simon Bauer, for helping me refine my ideas and work out the bugs.

...Eberhard Pracht, for his support and friendship and knowledge of the German language.

...Dr. Peter Nordbeck, Philipp Ehses, and Daniel Gareis, who have been incredibly important in helping me keep my sanity both in and out of the office (but especially out). It will be hard not having them around on a day-to-day basis.

...my past and future colleagues at Case Western Reserve University, specifically Prof. Jeff Duerk, Kestutis Barkauskas, Hisamoto Moriguchi, and Candice Bookwalter.

Sabine.

...Siemens Medical Solutions in Erlangen, especially Stephan Kannengiesser, Sonia Nilles-Vallespin, Peter Speier, Berthold Kiefer, and Matthias Nittka, for the productive collaboration.

...my countless roommates, including David Balla, Robin Heidemann, Matthias Mueller, Ilja Kaufmann, Esra Lang, Michael Ledwig, Simon Bauer, Olga Trichtchenko, Peter Nordbeck, Philipp Ehses, Marcos Lopez, and Marcus Warmuth, for providing a productive and friendly working environment.

...my proof readers and listeners, namely Felix Breuer, Martin Blaimer, Andre Fischer, Philipp Ehses, Eberhard Pracht, Peter Nordbeck, and Daniel Gareis, for smoothing out the wrinkles and attempting to fix all of the inaccuracies in this thesis.

...the Kaffee-Task-Force (KTF), without whom nothing would ever have been accomplished.

...EP5 and all its members, for welcoming me to the University of Wuerzburg and making my four years here enjoyable.

...Der Kommissar, der Esel, Gagi, Phipse-Phipse, Mr. Meggy Knife, Fid, Marwin, Gerry, der Pinche, Zieni, Benni Zahnfleisch, Apu, Jackass, Dr. Willi, der Siebzig, und der Vektor. They know what they did.

Erklärung

Gemäß §5, Abs. 2, Ziff. 2 und 5
der Promotionsordnung der
Fakultät für Physik und Astronomie der
Universität Würzburg

Hiermit erkläre ich an Eides statt, dass ich die Dissertation eigenständig, d.h. insbesondere selbstständig und ohne Hilfe eines kommerziellen Promotionsberaters angefertigt und keine anderen als die von mir angegebenen Quellen und Hilfsmittel benutzt habe.

Die Dissertation wurde bisher weder in gleicher noch in anderer Form in einem anderen Prüfungsfach vorgelegt.

Im Mai 2001 wurde mir von Yale University in New Haven, CT, USA der akademische Grad “Bachelor of Science” verliehen. Andere akademische Grade hab ich weder erworben noch versucht zu erwerben.

Würzburg, den 27.02.2008

(Nicole Seiberlich)

