

ÉCOLE NATIONALE SUPÉRIEURE DES MINES DE PARIS
ESCUELA TÉCNICA SUPERIOR DE INGENIEROS DE TELECOMUNICACIÓN
(UNIVERSIDAD POLITÉCNICA DE MADRID)

NONLINEAR RECONSTRUCTION IN
PARALLEL MAGNETIC RESONANCE IMAGING

CARLOS FERNÁNDEZ GRANDA

SEPTEMBER 20, 2008

Supervisors

DR. JULIEN SÉNÉGAS (PHILIPS RESEARCH HAMBURG)

PR. YVES ROUCHALEAU (ENSMP)

Abstract

This Master's thesis studies the application of nonlinear techniques to the reconstruction of parallel magnetic resonance imaging (MRI) data. The implementation and evaluation of a recently published algorithm introducing a nonlinear framework to the area of parallel MRI is presented. In addition, the adaptation of Compressive Sampling techniques to parallel MRI is explored. Finally, two novel Compressive Sampling algorithms for parallel MRI are proposed and evaluated using both simulated and measured MRI data.

Key words : Parallel MRI, nonlinear reconstruction, Compressive Sampling, wavelet transform, greedy algorithms, transform domain incoherence, polynomial model.

Résumé

Ce rapport de stage traite de l'application de techniques non-linéaires à la reconstruction d'images dans le domaine de l'imagerie par résonance magnétique (IRM) parallèle. En particulier, un algorithme publié récemment est décrit et les résultats de son évaluation sont analysés. En outre, l'adaptation des techniques de Compressive Sampling à l'IRM parallèle est explorée. Dans cette perspective, deux nouveaux algorithmes sont proposés et leurs performances évaluées à l'aide de données simulées et mesurées.

Mots-clés : IRM parallèle, reconstruction non linéaire, Compressive Sampling, décomposition en ondelettes, algorithme glouton, incohérence entre bases, modèle polynomial.

Acknowledgements

This Master's thesis was made possible by the guidance of my supervisor Dr. Julien S negas, who had the patience to direct my first steps in the field of medical imaging while always encouraging me to develop my own ideas. I would also like to thank Dr. Holger Eggers from Philips Research Hamburg for reading the first drafts and giving me very valuable advice, my supervisor at l' cole des Mines de Paris Professor Yves Rouchaleau for his support, Professor Yves Meyer from l' cole Normale Sup rieure at Cachan for introducing me to Compressive Sampling and for his kindness, and, finally, my fellow students at Philips Research Hamburg, in particular Stephan and Julien, who shared many a coffee and a kickerspiel with me.

I cannot end without expressing my gratitude to Hanna, who brought me to Hamburg, and to my parents, whose encouragement has never wavered. To them this work is dedicated.

Contents

Introduction	1
1 Parallel Magnetic Resonance Imaging	2
1.1 Magnetic resonance imaging	2
1.2 Spatial encoding and k-space formalism	3
1.3 Parallel MRI	6
1.3.1 Localised Sensitivity Functions	7
1.3.2 SENSE	8
1.3.3 Limitations of Linear Reconstruction Techniques	11
2 Compressive Sampling	14
2.1 Theory	15
2.1.1 Synthesis and Measurements	15
2.1.2 Signal Reconstruction	15
2.1.3 Lossy Reconstruction	16
2.2 Application to MRI	17
2.2.1 Sparsity of MR Images	17
2.2.2 Domain Incoherence in MRI	18
3 Joint Estimation of Image and Coil Sensitivity Functions	21
3.1 Theory	21
3.1.1 Nonlinear Optimization Problem	21
3.1.2 Implementation of JSENSE	23
3.1.3 Initialization	24
3.2 Methods	24
3.2.1 Algorithm Implementation	24
3.2.2 Data Sets	25
3.2.3 Experiments	26
3.3 Results	28
3.3.1 Polynomial Fitting of Sensitivity Functions	28

CONTENTS

3.3.2	JSENSE Reconstruction	28
3.3.3	Noisy Data Reconstruction	32
3.4	Discussion	47
4	Compressive Sampling of 2D Parallel Magnetic Resonance Data	48
4.1	Theory	48
4.1.1	Image Sparsity	48
4.1.2	K-Space Sampling Grid Design	49
4.1.3	Compressive Sampling Sum of Squares Image Reconstruction	51
4.1.4	Compressive SENSE Image Reconstruction	51
4.1.5	Compressive Sampling Joint Reconstruction	51
4.2	Methods	53
4.2.1	K-Space Sampling Grid Design	53
4.2.2	Algorithm Implementation	54
4.2.3	Data Sets	54
4.2.4	Experiments	55
4.3	Results	56
4.3.1	Signal Sparsity	56
4.3.2	Compressive Sampling SoS Image Reconstruction	57
4.3.3	CSENSE Image Reconstruction	60
4.3.4	Compressive Sampling Joint Reconstruction	65
4.4	Discussion	72
5	SHARK (Sparse incoherent Arbitrary Resampling in K-space)	74
5.1	Theory	75
5.1.1	Two-Step Reconstruction	75
5.1.2	K-space data interpolation	75
5.1.3	K-Space Sampling Grid Design	76
5.1.4	Motivations for SHARK	77
5.2	Methods	77
5.2.1	K-Space Sampling Grid Design	77
5.2.2	Algorithm Implementation	78
5.2.3	Data Sets	78
5.2.4	Experiments	78
5.3	Results	79
5.3.1	K-Space Interpolation Accuracy	79
5.3.2	SIM Phantom Reconstruction	79
5.3.3	Phantom Data	82
5.4	Discussion	82

CONTENTS

Conclusion	83
Bibliography	84
A K-Space Sampling Grid Design	88
A.1 Phase-Encoding Subsampling	88
A.2 SHARK Subsampling	92
B Compressive Sampling Regularization Parameters	96
B.1 Compressive Sampling SoS Reconstruction	97
B.2 CSENSE Image Reconstruction	106
B.3 CSENSE Noisy Data Reconstruction	118
B.4 Coil Sensitivity Reconstruction	124
B.5 SHARK Reconstruction	131

Introduction

The main subject of this Master's thesis is the development of nonlinear algorithms for the reconstruction of images from parallel magnetic resonance (MR) data. As opposed to current linear methods, nonlinear algorithms make it possible to exploit intrinsic characteristics of MR images, such as their sparsity in a certain transform domain, to improve their reconstruction. They also allow for more flexible reconstruction models, in which the sensitivity functions of the coils used to acquire the data can be incorporated as additional variables.

The first chapter in the thesis provides a brief introduction to parallel magnetic resonance imaging (MRI), with a special emphasis on the most popular current reconstruction methods and their main limitations. The second chapter is dedicated to Compressive Sampling, a recently proposed technique which permits reconstruction from highly undersampled data, and its applications to MRI. This sets the context for the analysis of several nonlinear algorithms which aim to improve the quality of parallel MR image reconstructions by addressing some of the limitations of current techniques.

More specifically, a recently published algorithm that introduces a nonlinear framework for the joint estimation of the image and the coil sensitivity functions is analysed in the third chapter. The combination of such a joint reconstruction procedure with Compressive Sampling is explored in the fourth chapter, along with some other applications of Compressive Sampling to parallel MRI. Finally, in the fifth chapter, a novel algorithm which combines k-space interpolation and Compressive Sampling is proposed. Each of the three last chapters includes a theoretical analysis of the corresponding MRI reconstruction algorithm, a detailed description of the employed methods, an account of the different experiments and a final section discussing the results.

Chapter 1

Parallel Magnetic Resonance Imaging

The first two sections of this chapter provide a brief introduction to magnetic resonance imaging (MRI). The main aim is to introduce MRI concepts which are used throughout the report and not to provide a detailed technical account. Such an account can be found, for example, in the main reference for these two sections [13]. The third section describes the motivation for the development of parallel acquisition, along with one of the most representative reconstruction procedures associated with this technique. Finally, several limitations of current reconstruction approaches are pointed out.

1.1 Magnetic resonance imaging

MRI is an imaging technique used primarily in medical applications to produce high quality images of cross sections of the human body. MRI is based on the principles of nuclear magnetic resonance (NMR), a spectroscopic technique that can be exploited to obtain microscopic chemical and physical information about molecules.

Magnetic resonance imaging is made possible by the reaction of certain nuclei in the human body to magnetic fields. Every nucleus with uneven atomic mass or uneven atomic number has a non-zero spin which induces a magnetic field. The magnetic moments produced by individual nuclei are usually randomly oriented and produce no net magnetic effect. It is possible, however, to align them in a certain direction by applying a static magnetic field. Each nucleus can be thought of as a tiny magnetic dipole oriented parallel or antiparallel to the external field.

The alignment of the magnetic moments with the magnetic field is not perfect. Instead, there is a precession movement called *Larmor precession* around the axis of the field. The rate of precession is determined by the specific characteristics of the nucleus involved and the strength of the external magnetic field. The expression for the angular frequency is:

$$\omega_0 = \gamma B_0 \tag{1.1}$$

γ is a constant called *gyromagnetic ratio*, which depends on the nucleus, while B_0 is the applied static magnetic field. ω_0 is the Larmor frequency, a resonance frequency at which energy can be efficiently transferred to the system.

The difference between the energy of the parallel (lower energy) and antiparallel (higher energy) orientation of the nuclei produce a net magnetization of the set of atoms affected by the field. This net magnetization depends on the strength of the external field and on the temperature.

If the nuclei are irradiated with a short burst of radio waves at the Larmor frequency, they absorb energy and are tilted away from the axis of net magnetization. After a certain time, they return to their former positions, emitting the absorbed energy in form of electromagnetic radiation. The resulting signal, which can be detected as induced currents by a receive coil, can then be processed to reconstruct an image. Some of the most relevant details of this procedure are explained in the next section.

1.2 Spatial encoding and k-space formalism

Image reconstruction from magnetic resonance signals is based upon what is known as *spatial encoding*. Spatial encoding can be implemented by exploiting the relation in Eq. 1.1 between the Larmor frequency at which nuclei resonate and the static magnetic field. By modifying the magnetic field it is possible to introduce variations in the resonance frequency which depend on the position of the resonating nuclei, thus encoding the location of the nuclei within the signal.

In two dimensional imaging, for instance, it is necessary to make sure that the signal received by the coil is irradiated only from protons belonging to a certain plane. In order to achieve this, an additional component is added to the static magnetic field. This component, the *slice selection gradient*, has a constant gradient that is perpendicular to the selected imaging plane. In MRI terminology, *gradient* frequently refers to a magnetic field with a constant gradient in a certain direction.

As a result of the application of the slice selection gradient, the different planes react only when irradiated at their specific Larmor frequency, which varies linearly in the slice selection axis. The measured signal for an irradiation at a given Larmor frequency corresponds consequently only to a certain plane. This procedure is illustrated in Fig. 1.1.

Within the imaging plane two additional mechanisms based on the same principle are employed: frequency and phase encoding. Frequency encoding involves the application of another magnetic




Figure 1.1: Selection of the imaging plane in the z direction by adding a constant gradient to the static magnet field B_0 .

gradient (in the MRI sense) in the direction of a certain axis x contained in the imaging plane during signal acquisition. This makes the oscillation frequency of the MR signal linearly dependent on its x coordinate. The applied magnetic field can be expressed in this form:

$$B = B_0 + G_x x \quad (1.2)$$

where G_x is the magnitude of the gradient. As a result, according to Eq. 1.1, the signal dS obtained from an infinitesimal interval dx at a time t will be:

$$dS(x, t) = \rho(x) e^{-i\gamma(B_0 + G_x x)t} dx \quad (1.3)$$

$\rho(x)$ designates the proton density for a certain value of x . It must be noted that this expression is simplified for notational convenience and does not include relaxation effects.

The second mechanism, which incorporates spatial information into the signal, is quite similar. In this case, it is the phase of the signal, rather than the frequency, which is modified by adding yet another gradient G_y , contained in the imaging plane and perpendicular to G_x . This gradient is turned on when the irradiation signal is emitted and then turned off after a short interval of time T . As a result, a linear dependence is induced between the y coordinate and the phase of the signal. After the interval, the signal will have accumulated a phase difference of $\gamma G_y y T$. Taking into account both spatial encoding techniques, the signal received at time t can be expressed as:

$$S(t) = \int \rho(x, y) e^{-i\gamma(B_0 t + G_x x t + G_y y T)} dx dy \quad (1.4)$$

$\rho(x, y)$ is the spatial dependent proton density function which corresponds to the two dimensional image that is to be reconstructed. After demodulation, which eliminates the $e^{-i\gamma B_0 t}$ factor, the following expression is obtained:

$$S(t) = \int \rho(x, y) e^{-i\gamma(G_x x t + G_y y T)} dx dy = \int \rho(x, y) e^{-i(k_x x + k_y y)} dx dy \quad (1.5)$$

$k_x = \gamma G_x t$ and $k_y = \gamma G_y T$ can be interpreted as spatial frequencies. The received signal is consequently equivalent to samples of the two dimensional Fourier transform of the proton density at those frequencies. In MRI this spatial frequency domain is usually called *k-space*. Figure 1.2 shows a two dimensional MR image of a human brain and its k-space representation. Most of the frequential information is concentrated around the centre of the k-space, so that in comparison the outer regions seem almost completely black.



Figure 1.2: An image of the brain obtained by MRI (left) and its equivalent representation in the spatial frequency domain or k-space (center). The k-space representation is also shown on a logarithmic scale (right), where high frequency data are more visible.

If enough samples of the k-space are measured, the image can be reconstructed by applying a two-dimensional inverse Fourier transform. According to Nyquist's theorem, we can determine the image field of view (an MRI term for the length of the image in a certain direction) x_{\max} and resolution Δx by adjusting the spacing between the frequency samples Δk_x and the maximal acquired frequency $(k_x)_{\max}$:

$$x_{\max} = \frac{1}{2\Delta k_x} \quad (1.6)$$

$$\Delta x = \frac{1}{2(k_x)_{\max}} \quad (1.7)$$

y_{\max} and Δy can obviously be adjusted in an analogous way.

It is important to note that the spatial encoding scheme just described corresponds to a what is usually denominated a *cartesian* sampling of the k-space, since the measurements are equivalent to sampling the spatial frequencies of the image on a regular cartesian grid. Other trajectories, such as radial or spiral, are possible, but their description is out of the scope of this brief introduction to spatial encoding.

In order to make clear the need for acceleration techniques such as those presented in the following section, it is useful to picture the trajectory we follow when sampling the MR signal after a radio frequency signal excitation. On the one hand, k_x changes linearly with time, on the other, k_y remains fixed, since it only depends on the initial period T . As a result, for every excitation it is only possible to cover a single straight line in k-space. In order to complete the sampling of the whole k-space it will thus be necessary to carry out many additional excitations. After each excitation, the read-out of a single signal takes approximately between 2 and 6 milliseconds. Since the nuclei must return to equilibrium before a new excitation can be carried out, there is an additional delay that ranges approximately between 10 milliseconds and 1 second. Even though more elaborate measuring techniques, such as for example using *spoiling* gradients, can help to reduce these delays, the repetition of excitations implied by phase encoding is the main time limitation for the whole process of MR data acquisition.

1.3 Parallel MRI

As explained at the end of the previous section, the acquisition of MR images is slowed down by the sequence of repeated measurements necessary for phase encoding. Long scans are inconvenient for the patient and may result in degraded image quality, especially in applications such as cardiac imaging. Reducing scan time is thus a major challenge in MR imaging research. This challenge has been tackled with remarkable success by parallel MRI.

Parallel MRI involves simultaneously sampling the signal with several coils, instead of just one. This makes it possible to obtain an image of similar quality with less samples, since the data from the different coils can be combined to make up for the missing information. If the data were reconstructed directly, the violation of Nyquist's limit would produce *aliasing artifacts* consisting of multiple superpositions of the image. As we will see, using n coils makes it possible, at least theoretically, to use n times less samples, eliminating the artifacts by exploiting redundant information from the different coils. In practice, this means that for arrays with 8 coils, for example, it is possible to obtain good quality images 5 or 6 times faster than in the case of conventional MRI. Since it is the phase-encoded direction that is more costly in terms of measuring time, the data is

usually only subsampled in this direction.

The effect of measuring with several coils is modeled in the following way. Each of the coils receives a different signal, which can be interpreted as the original signal multiplied by the particular *sensitivity* function of the coil $s_c(x, y)$. After demodulation, the signal S_c received by a certain coil c can be expressed by the formula:

$$S_c(k_x, k_y) = \int s_c(x, y) \rho(x, y) e^{-i(k_x x + k_y y)} dx dy \quad (1.8)$$

Figure 1.3 shows examples of sensitivity functions and the effect they have on the acquired data.




Figure 1.3: Coil sensitivity functions (left) and coil images of a human brain (right) obtained with a 6-coil array.

This can be interpreted as a yet another encoding of spatial information: sensitivity encoding. Algorithms for parallel MRI reconstruction use sensitivity encoded information to recover the original signal from subsampled data. Two examples of such algorithms are explained in the next two sections. The first is of limited practical utility, but can be used to illustrate the main ideas underlying this kind of reconstruction. The second is one of the most widespread and successful algorithms in parallel MRI.

1.3.1 Localised Sensitivity Functions

This section focuses on the case of coil sensitivities that are restricted to a certain area of the image. If the sensitivity function of a coil is localised enough, the aliasing artifacts caused by subsampling will not affect the section of the image measured by that coil. This makes it possible to perform the reconstruction by directly combining the images from the different coils, which correspond to replicated localised sections of the original image. Figure 1.4, taken from reference [2], illustrates the PILS (Partially Parallel Imaging With Localized Sensitivities) algorithm proposed by Griswold *et al* [6], which implements such a procedure.




Figure 1.4: PMLS reconstruction of an MR image subsampled at an acceleration factor of 2 and measured with an array of 4 coils. (a) The four coil sensitivities. (b) Fully sampled reconstruction of each coil image. (c) Undersampled reconstruction of each coil image. (d) Selection of image segments which are free of aliasing. (e) Elimination of remaining artifacts. (f) Final reconstruction.

In general coil sensitivity functions are not strictly spatially localised, so parallel MRI algorithms must perform more sophisticated operations on the data. Some of these algorithms use an explicit expression of the sensitivity function, while others perform the reconstruction using only the subsampled data. The following section describes the most representative example of the first kind of algorithms.

1.3.2 SENSE

SENSE (SENSitivity Encoding), proposed by Pruessmann *et al* in 1999 [25], is the most widespread reconstruction method based on a previously available estimate of coil sensitivities. SENSE reconstruction *unfolds* the aliased coil images while maximizing the signal-to-noise ratio. In the case of regular cartesian k-space sampling, the unfolding can be significantly reduced and efficiently implemented [25]. In the more general case, the reconstruction process is equivalent to the solution of a

linear system of equations, which can be efficiently carried out by an iterative algorithm [24]. This subsection is dedicated to the second approach, which is of more relevance to the work presented further on.

From now on, the proton density function ρ will be replaced by its discrete approximation I , which is precisely the image we wish to reconstruct. The sensitivity function for each coil will be assumed to be known. In practice, they are either measured in a previous scan [25] or estimated from the central k-space data [22]. As we will see further on, this may introduce errors in the reconstruction.

Using a linear algebra framework, MRI undersampled data can be expressed as the result of applying two linear transformations to the original image. As implied by Eq. 1.8, these transformations are sensitivity and Fourier encoding. Sensitivity encoding can be represented by a matrix S that satisfies:

$$S_{(r,c),j} = s_c(r)\delta_{r,j} \quad (1.9)$$

where r is the pixel position, c the coil number, s_c the sensitivity function for coil c and $\delta_{i,j}$ the Kronecker delta, equal to one if $i = j$ and to zero otherwise. The matrix dimensions are $(n_I n_c) \times n_I$, where n_I is the number of pixels in the image and n_c the number of coils. A matrix containing the different coil images \bar{I} can thus be obtained by multiplying S and I (in this formulation I is a vector containing the n_I pixels of the image):

$$\bar{I} = S I \quad (1.10)$$

The discrete Fourier transform of \bar{I} can be expressed as a matrix product $F\bar{I}$, where $F_{k,r} = e^{i2\pi kr}$ and k and r respectively span the k-space and the image space. The matrix is consequently of size $(n_I n_c) \times (n_I n_c)$. Subsampling of the k-space, which results in missing values for certain k-space positions, can be represented as the effect of multiplying the complete Fourier transform of the data $F\bar{I}$ by a projection matrix P . P is obtained from a diagonal matrix with n_d ones in the positions that correspond to sampled k-space values and zeros elsewhere. If we eliminate the rows of the matrix which are equal to zero, we obtain the $n_d \times (n_I n_c)$ matrix P . n_d is the total number of k-space measurements for all coils. In this way, we obtain an expression of the data d measured by the MR coils as a linear transformation E of the image:

$$d = P F S I = E I \quad (1.11)$$

This system of equations is a discrete approximation of Eq. 1.8. The whole encoding scheme that models the linear transformation between the image and the subsampled data used for reconstruction is illustrated in Fig. 1.5.




Figure 1.5: Example of parallel MRI encoding for an array with 6 elements and an acceleration factor of 4. The image I is weighted by the coil sensitivities SI , transformed into the frequency domain (FSI) and subsampled ($PFSI$). SI is represented by the 6 coil images, while FSI and $PFSI$ are represented by the respectively fully sampled and undersampled k-space representations of the single coil images.

In the reconstruction process, there are n_I variables to be determined, namely the number of pixels in the image that correspond to the desired resolution and field of view, while the number of equations is $n_d n_c$. In order for this system not to be underdetermined, there is an obvious upper bound on the subsampling factor $R = \frac{n_I}{n_d}$:

$$R \leq n_c \quad (1.12)$$

In general, noise affecting real MR images is well simulated as Gaussian noise. Under such an assumption, the acquired data \hat{d} is corrupted by a vector η of zero mean:

$$\hat{d} = d + \eta = EI + \eta \quad (1.13)$$

If we consider a sample noise matrix $\hat{\Psi}$, describing the level of stochastic noise in signal samples, the weighted square distance between the measured data \hat{d} and the noiseless data d would be:

$$(\hat{d} - d)^H \hat{\Psi}^{-1} (\hat{d} - d) = (\hat{d} - EI)^H \hat{\Psi}^{-1} (\hat{d} - EI) \quad (1.14)$$

Deriving this expression and setting the derivative to zero we can obtain the image estimator which minimizes it:

$$\hat{I} = (E^H \hat{\Psi}^{-1} E)^{-1} E^H \hat{\Psi}^{-1} d \quad (1.15)$$

Under the assumption of equation 1.13, this estimator is also the maximum likelihood estimator of I . If we consider that the noise has a constant variance and is uncorrelated, the correlation matrix is equal to the identity matrix. In this case, the reconstruction process is equivalent to calculating the following least squares estimator:

$$\hat{I} = \arg \min_I \|\hat{d} - EI\|^2 \quad (1.16)$$

In [24] Pruessmann *et al* propose calculating estimate 1.15 for arbitrary k-space trajectories with a preconditioned conjugated gradient algorithm, which exploits the fact that the product EI can be implemented efficiently as a pixel by pixel multiplication of the sensitivity functions and the image followed by the application of the Fast Fourier Transform algorithm. This approach is commonly known as Generalized SENSE (GSENSE). Since its publication, SENSE has proven extremely successful and is at present widely used in clinic applications. There are, however, some limitations to the technique which will be explained in the following section.

1.3.3 Limitations of Linear Reconstruction Techniques

In this section, the principal limitations of the SENSE reconstruction will be outlined. Special attention will be paid to the fact that these limitations generalize to the main alternatives to SENSE, since they too are based on linear reconstruction schemes. This suggests that in order to

further improve reconstruction quality, it will be necessary to turn to nonlinear image reconstruction techniques.

Noise Amplification

For high accelerating factors, SENSE reconstruction suffers from noise amplification, which degrades image quality. The reason for this is that the system of equations 1.11 becomes increasingly ill-conditioned. As a result, any errors present in the data have a large influence on the reconstruction results. In this case, the reconstruction procedure is an ill-posed inverse problem, which must be solved by introducing *a priori* information in the form of regularization. An example of such an approach is Tikhonov regularization [14], which places a penalty on the l_2 norm of the solution. Assuming uncorrelated white Gaussian noise, the modified estimator is:

$$\hat{I} = \arg \min_I \|\hat{d} - EI\|^2 + \lambda \|I\|^2 \quad (1.17)$$

where λ is a constant which represents the relative importance of each of the terms in the reconstruction.

Sensitivity Estimation

A further drawback of SENSE is the need for a prior estimation of the coil sensitivity functions. As mentioned before, this estimate can be obtained from a previous low resolution scan [25], in which the signal of each coil is divided by a reference signal obtained with a body coil which has almost uniform sensitivity. The accuracy of the reference scan is logically critical for SENSE image reconstruction quality. If the coil array or the patient move between the reference scan and the actual acquisition, the sensitivity function will change and the estimation error will produce artifacts in the reconstruction. The additional reference scan not only takes time, but is consequently also a potential source of reconstruction errors.

So-called autocalibrated sensitivity estimates can be obtained directly from the MR data [22]. A central region of k-space, consisting of a certain number of autocalibration signal (ACS) lines, is completely sampled and then used to extract an approximation to the sensitivity by dividing each of the low frequency coil images by their sum of squares. Although the frequential representation of sensitivity functions is concentrated near the k-space center due to their regularity, this is not necessarily the case for their convolution with the image (the Fourier space equivalent of point to point multiplication in the image space). Furthermore, the absence of a reference body coil image also affects the quality of the estimation. Thus, this estimation procedure is far from optimal and is also a source of quality degradation in the reconstruction.

Finally, other parallel MRI algorithms try to exploit sensitivity encoding in an implicit way. The best known example is GRAPPA (Generalized autocalibrating partially parallel acquisitions) proposed by Griswold *et al* [7]. GRAPPA reconstructs the original k-space representation of the image. The basic underlying idea is that the missing points in the undersampled k-space data of each coil can be calculated by a linear combination of neighbouring samples from all the coils. The coefficients for the linear combination are fitted using a small number of ACS lines. Intuitively, it is mainly the ACS lines that are responsible for supplying information about the sensitivity functions. As a result, for similar reasons as in the self-calibrated approach, this information does not suffice to fully characterize the coil sensitivities.




Figure 1.6: GSENSE (top) and GRAPPA (bottom) reconstructions of a human brain image from k-space data acquired with an array with 6 elements and undersampled by a factor of 2 with 4 additional ACS lines (left) and by a factor of 4 with 12 additional ACS lines (right).

Figure 1.6 shows some examples of aliasing artifacts in linear reconstructions of data acquired with an array with 6 elements, caused by an inaccurate sensitivity estimation in the case of GSENSE and by an insufficient number of ACS lines in the case of GRAPPA. As we can see, for an acceleration factor of 4 the reconstruction is severely degraded. A way to overcome this limitation could be to estimate the coil sensitivities using all available k-space information. Such an estimate, however, is not possible in the framework of a linear reconstruction where sensitivities are estimated *a priori* (GSENSE), through low-frequency data (self-calibrated SENSE) or not at all (GRAPPA). This is the main motivation for the development of the nonlinear algorithms presented in further chapters.

Chapter 2

Compressive Sampling

Most of the signals we encounter in real life, such as photographs or videos, are highly compressible. This follows from the fact that it is possible to find a certain transform domain where their representation consists of a small number of nonzero coefficients. In most applications, data are first fully acquired and then compressed by transforming them into the *sparsifying* domain and discarding negligible coefficients. In this way they can be stored and processed more efficiently. A seemingly naïve question might be if we could not just acquire the relevant coefficients directly, instead of carrying out the usually large number of redundant measurements.

At a first glance this idea might seem of little practical use. In general, we do not know which of the coefficients carry the information and which do not. A random measurement of a small number of coefficients in the sparse base, for example, would almost surely yield an overwhelming majority of zeros and miss most of the valuable nonzero coefficients. Without knowing the position of the relevant coefficients beforehand, the task of reconstructing the whole signal by taking a small number of samples appears to be hopeless. In fact, this is not the case. The recently developed theory of *Compressive Sampling* or *Compressed Sensing* [4] describes how sparse signals can be sampled at rates significantly below the Nyquist limit and then accurately reconstructed.

In the first section of this chapter, a brief theoretical description of Compressive Sampling is presented. In the following section, the application of Compressive Sampling techniques to MRI is discussed, paying special attention to several difficulties that arise from practical considerations, some of which are yet to be overcome.

2.1 Theory

2.1.1 Synthesis and Measurements

The general framework of Compressive Sampling is based upon what Emmanuel Candès and Terence Tao, two of the pioneers of the theory along with Dave Donoho, call a *synthesis/measurement pair* [4]. Measurements will be considered to be linear transformations of the data. A matrix A represents the set of linear functionals responsible for extracting the measurement vector y from the original data x :

$$y = Ax \tag{2.1}$$

A defines the transformation of the data to the *measurement domain*. The *synthesis domain*, on the other hand, is a certain transform domain in which the original signal x is assumed to be sparse. This means that most of the coefficients θ_i of the representation of x in a basis of that transform domain $\theta(x)$ will be equal to zero. This condition can be expressed as a bound on the l_0 norm (which is equal to the number of nonzero components in a vector) of $\theta(x)$:

$$\|\theta(x)\|_{l_0} \leq K \tag{2.2}$$

where K is an upper bound for the number of nonzero coefficients.

The sparsity of the signal in the synthesis domain is the first condition for an accurate reconstruction. The second condition is that the synthesis/measurement pair must be sufficiently *incoherent*, in the sense that basis vectors of one domain cannot sparsely represent basis vectors of the other. From an intuitive point of view, the information which is concentrated in just a few coefficients in the synthesis domain must be as spread out as possible in the measurement domain, so it can be retrieved from just a few measurements. This second condition is formalized by Candès and Tao through the Exact Reconstruction Principle and the Uniform Uncertainty Principle [4].

2.1.2 Signal Reconstruction

Once the two conditions are satisfied, reconstructing the signal x amounts to finding a vector which is sparse in the synthesis domain and at the same time is coherent with the available data in the measurement domain. Compressive Sampling achieves this through an optimization procedure.

The ideal way of reconstructing a signal satisfying condition 2.2 from measurements modeled by Eq. 2.1 would be to minimize the l_0 norm of $\theta(x')$ with the additional constraint that $y = Ax'$. Unfortunately, this optimization problem is not computationally tractable. An alternative could be to minimize the synthesis representation of the signal in another norm, such as l_2 . Unfortunately,

l_2 norm minimization, despite having been widely and efficiently applied, is of little use for our problem since it does not favour sparse solutions.

An adequate compromise between solution sparsity and computational considerations is the l_1 norm. Unlike l_2 norm minimization, l_1 norm minimization favours sparsity. An intuitive explanation for this is that it penalises solutions with a few large components less than others with a large number of small components. l_2 minimization has exactly the opposite effect and consequently favours dense solutions. Another advantage of l_1 norm minimization, in particular when compared to l_0 norm minimization, is that it can be solved efficiently. It can, for instance, be reformulated as a convex quadratic program, which can then be tackled by several standard optimization techniques such as interior-point methods (for a recent implementation of such an approach, see [10]).

As a result, the Compressive Sampling reconstruction procedure is implemented by solving the following constrained l_1 norm optimization problem:

$$\arg \min_{x'} \|\theta(x')\|_{l_1} \quad \text{subject to} \quad y = Ax' \quad (2.3)$$

In [4], Candès and Tao prove that the solution of this problem is an exact reconstruction of the original signal x with overwhelming probability, when the number of samples exceeds the number of sparse coefficients in the synthesis base by a certain amount. If the original signal is sparse, this results in an accurate reconstruction from a number of samples well below Nyquist's sampling limit.

2.1.3 Lossy Reconstruction

Compressive sampling results can be extended to cases where the signal to be reconstructed is not completely sparse. It is enough for the signal to be *compressible*, which means that its energy must be concentrated in the largest coefficients of its synthesis domain representation.

This assumption is much better adapted to practical situations than strict sparsity, which is frequently too constraining a condition for real-life signals. The basic idea is the same as that underlying lossy compression algorithms such as JPEG [30]: the largest coefficients in a certain domain (discrete cosine transform or wavelet in the case of JPEG and JPEG 2000 [1]) can be selected and the rest discarded to yield a very accurate approximation of the signal. It is interesting to note that such an approximation is nonlinear, since it is the result of projecting the signal not on a subspace formed by preselected base vectors, but on a subspace which is dependent on the very structure of the signal. As discussed for example in Ch. 9 of [21], such simple nonlinear schemes are frequently more efficient than the best possible linear approximation.

A possible way of describing compressible signals is that the size of their synthesis domain representation coefficients must follow an exponential decay. If this condition is satisfied, the Compressive

Sampling reconstruction is surprisingly robust. In fact, in [4] the authors prove that it is hardly possible to do any better, in the sense that the reconstruction is practically as accurate, in terms of l_2 norm approximation error to the original signal, as choosing the largest sparse domain coefficients *a posteriori*. Intuitively, this means that Compressive Sampling is capable of automatically acquiring a compressed representation of the signal.

Practical applications of Compressive Sampling are increasingly varied and abundant. The next section focuses on its application to MRI, which has been particularly promising. There are, however, some practical problems, related to the specific characteristics and constraints of MR data reconstruction and acquisition, which must be overcome if Compressive Sampling MR imaging is to become a practical alternative to the reconstruction methods currently in use.

2.2 Application to MRI

The application of Compressive Sampling techniques to MRI was suggested from the very beginning (see for example Sec. 1.7 of [4]). The reason is that both key requirements for the Compressive Sampling reconstruction process are at least partially satisfied. On the one hand, MR images have been shown to be sparse in certain transform domains such as the finite differences or the wavelet domain. On the other hand, the k-space domain in which MR data is encoded has been proven to be incoherent with those transform domains for some sampling trajectories.

In the notation of Sec. 2.1.1, a Compressive Sampling reconstruction of MR data can be described as follows: The signal x is the image corresponding to the proton density of a certain area of the body, as explained in Chap. 1, and the data y is the undersampled k-space representation of the image measured by a MR scan. The measurement matrix A is determined by the k-space sampling scheme. It can consequently be represented by the product of the Fourier operator F and the projection matrix P (as in Sec. 1.3.2, P can be obtained from a diagonal matrix with ones at the positions that correspond to sampled k-space values by eliminating the zero rows): $A = PF$.

2.2.1 Sparsity of MR Images

In order to apply Compressive Sampling to MR data, it is necessary for the images to be sparse in a certain transform domain Θ . In Ref. [18], angiograms (images of blood vessels) are shown to be sparse in the spatial finite-difference domain (they even look quite sparse in the image domain), while brain images are shown to be sparse in the wavelet and discrete cosine transform domains. Applications to dynamic MR imaging which exploit sparsity in the time domain have also been developed [19].

It is important to emphasize the difficulty of evaluating the quality of compressed medical images.

The main problem is that the basic criterion, whether the images are useful for clinical diagnosis, is rather subjective. This makes it particularly complicated to decide if a certain approximation, obtained by truncating the representation of the image in a certain base, is acceptable or not. In any case, the diagnostic quality of MR images obtained by Compressive Sampling has already proven comparable and in some cases even superior to those obtained by other reconstruction methods [18].

2.2.2 Domain Incoherence in MRI

The second step in applying Compressive Sampling to MR image reconstruction to ensure that the sparsifying transform domain is incoherent with the measuring scheme. In Ref. [4], random projections of Fourier transform matrices are proven to be incoherent with the image domain. Intuitively, the reason for this is that whereas regular undersampling produces a low number of high energy coherent aliasing artifacts, random undersampling aliasing artifacts are highly irregular and similar to background noise. Figure 2.1 illustrates the effects of regular and random undersampling of a brain MR image on the shape and distribution of the aliasing artifacts.




Figure 2.1: Aliasing artifacts caused by a twofold undersampling in the case of regular cartesian undersampling (left) and random irregular undersampling (right).

Measuring Incoherence

The notion of incoherence between the image and frequential domains can be formalized and quantified by the point spread function (PSF) of the Fourier sampling scheme. For a certain basis vector e_i in the image domain, equivalent to a pixel, this function is defined as:

$$\text{PSF}(i, j) = e_j^* (P F)^H P F e_i \quad (2.4)$$

This function measures the effect of the aliasing generated by the pixel e_i on another pixel e_j . For a regular undersampling by a factor R , the point spread function consists of R regularly spaced delta functions, so that it is impossible to identify the original point. For a random undersampling scheme

of the same factor, the point spread function spreads irregularly in a noise-like fashion through all of the image space. If R is not too high, the original point will still be recognisable. Figure 2.2, taken from [20], shows several point spread functions which illustrate the varying degrees of incoherence between the image and k-space domains for different sampling schemes.




Figure 2.2: PSFs of various sampling trajectories: (a) random lines in 2-D, (b) random points in 2-D or cross section of random lines in 3-D, (c) radial, (d) uniform spirals, (e) variable density spirals, and (f) variable density perturbed spirals. The height of the red lines measures coherence.

In [18] and [20] Lustig *et al* present a generalization of the point spread function: the Transform Point Spread Function (TPSF). The TPSF measures how each coefficient in the sparsifying transform representation of the image affects the other coefficients. If we denote by Θ the sparsifying transform matrix, the formula of the TPSF which quantifies how a certain coefficient i affects another coefficient j , would be given by:

$$\text{TPSF}(i, j) = e_j^* \Theta (P F)^H P F \Theta^H e_i \quad (2.5)$$

In order to ensure an effective reconstruction of the MR data, the projection matrix P must sample k-space such that TPSF coherence is minimized. Ideally the TPSF should be as similar to a delta function as possible, so that the energy of each nonzero sparse base coefficient leaks as little as possible into the other coefficients. This optimal sampling scheme would be very similar to random measurements, and indeed many of the initial results in Compressive Sampling reconstruction assumed a completely random sampling scheme.

In practice, k-space sampling is greatly constrained by hardware and physiological considerations. In general, sampling trajectories must follow relatively smooth lines and curves and be robust,

since non-cartesian sampling can be highly sensitive to system imperfections. This makes it almost impossible to apply purely random trajectories. It is, however, still possible to achieve significant incoherence as can be seen in Fig. 2.2 through the use of non-cartesian sampling schemes such as radial or spiral trajectories.

Chapter 3

Joint Estimation of Image and Coil Sensitivity Functions

As seen in Sec. 1.3.3, one of the main sources of image degradation in linear reconstruction algorithms is an inaccurate estimation of coil sensitivities. This chapter presents the theoretical aspects, implementation and reconstruction results of an algorithm proposed by Ying and Sheng in 2007 to address this problem: Joint Image and Sensitivity Estimation in SENSE (JSENSE) [33].

3.1 Theory

3.1.1 Nonlinear Optimization Problem

In order to improve the accuracy of coil sensitivities used for image reconstruction, JSENSE seeks to jointly estimate the image and the sensitivity functions during the reconstruction process. The reconstruction problem is thus generalized to estimating not only the image I , but also the different sensitivities represented by S , as in Eq. 1.10. The model which associates the unknown signals to be reconstructed to the data is no longer linear:

$$d = P F M(S, I) \tag{3.1}$$

M is a function which represents the pixel by pixel multiplication of the image and the coil sensitivities. Applying a least squares approximation to obtain S and I in Eq. 3.1 would not yield acceptable estimates of the image or the sensitivity functions, as the ensuing optimization problem is not convex and has many local optima. This is easy to see, since we can generate any number of useless solutions by simultaneously multiplying the image and dividing the coil sensitivities by an arbitrary function.

The fact that this inverse problem is not well posed should not come as a surprise, since the number of parameters to be estimated is much larger than the available data. If the undersampling

factor is R and the total number of measurements is equal to n_d , $R n_d$ sensitivity pixels and $\frac{R n_d}{n_c}$ image pixels must be reconstructed from just n_d samples. This means that there are $\frac{R(n_c+1)}{n_c}$ more parameters than data points, making calibration impossible without additional *a priori* information.

Coil sensitivities are known to be regular functions, so that they are well represented by polynomial models (see for example [25], where polynomial fitting is used to denoise sensitivity maps). Ying and Sheng propose to use such a model to reduce the number of parameters in the inverse problem. More precisely, the sensitivity s_c of a coil c is represented by a two-dimensional polynomial function centered at (\bar{x}, \bar{y}) :

$$s_c(x, y) = \sum_{i=0}^{n_p} \sum_{j=0}^{n_p} a_{i,j,c} (x - \bar{x})^i (y - \bar{y})^j \quad (3.2)$$

The highest power n_p is chosen to be the same for both dimensions. The unknown parameters are no longer the pixels of the coil sensitivities, but the coefficients of the polynomial functions. The nonlinear transformation that models the sampled data is modified to:

$$d = P F M_p(a, I) \quad (3.3)$$

where M_p represents the pixel by pixel product of the image I and the n_c polynomial functions, whose coefficients are represented by a vector a . Just as in the case of SENSE, the fact that noise in MRI is well modeled as additive white Gaussian motivates a reconstruction based on least squares. We will assume that the noise is uncorrelated and uniform (as we saw in Sec. 1.3.2, it is always possible to introduce prior noise information in the form of a sample noise correlation matrix), so that the estimator is the least squares solution to Eq. 3.3. Such a reconstruction is thus equivalent to solving the following nonlinear optimization problem:

$$\arg \min_{a, I} \|\hat{d} - P F M_p(a, I)\|^2 \quad (3.4)$$

From a theoretical point of view, the problem is overdetermined as long as the number of data samples n_d is superior to the number of parameters $n_I + (n_p + 1)^2 n_c$, which places a constraint on the acceleration factor $n_f = \frac{n_I}{n_d}$:

$$n_f \leq \frac{n_I n_c}{n_I + (n_p + 1)^2 n_c} \quad (3.5)$$

Because of the regularity of the sensitivity functions, low order polynomials are usually sufficient for an adequate representation. The solution to 3.4 is a joint estimation of the image and polynomial sensitivity functions that minimize the l_2 norm error between the model and the data. The following section explains how Ying and Sheng propose to calculate such a solution.

3.1.2 Implementation of JSENSE

In order to solve Eq. 3.4, Ying and Sheng resort to a greedy iterative algorithm that alternates between two separate steps. In the first step the image is updated according to the current value of the coil sensitivities. In the second step, the coil sensitivities are updated according to the current value of the image. As we will see in this section, this simplifies the problem to two subproblems that are linear and consequently much easier to solve.

Image Update

The image update step is carried out by applying the Generalized SENSE reconstruction algorithm explained in Sec. 1.3.2. If we fix the sensitivity values in Eq. 3.3, the equation is equivalent to Eq. 1.16 and can be solved efficiently, for instance, by a conjugate gradient algorithm.

Sensitivity Update

The sensitivity update is a least squares fitting of each coil sensitivity given the undersampled data from the corresponding coil and the estimated image. If we fix every variable in Eq. 3.4 except for the sensitivity polynomial coefficients, we obtain the following linear problem:

$$\arg \min_a \|\hat{d} - B a\|^2 \quad (3.6)$$

B is a matrix that combines a multiplication by the different powers of x and y , a multiplication by the image, a transformation into the frequential domain and k-space undersampling. It can be seen as a linear transformation which maps the polynomial coefficient representation to the measured data. Its dimensions are thus $n_d n_c \times (n_p + 1)^2 n_c$. Since the transformation maps the coefficients of a coil sensitivity function exclusively to the data measured by that coil, the problem can be decoupled into n_c separate subproblems:

$$\arg \min_{a_c} \|\hat{d}_c - B_c a_c\|^2 \quad (3.7)$$

In order to do this, B is split into n_c submatrices B_c of dimension $n_d \times (n_p + 1)^2$. The elements of these matrices are of the form:

$$B_c [(k_x, k_y), (p, q)] = \sum_{x=1}^{n_x} \sum_{y=1}^{n_y} I(x, y) x^p y^q e^{i2\pi(k_x x + k_y y)} \quad (3.8)$$

k_x and k_y are the k-space coordinates of the sampled data, x and y the image coordinates and p and q the powers in dimensions x and y associated to each coefficient. Eq. 3.6 can therefore be solved by using the n_c B_c matrices to separately fit the polynomial coefficients of each coil sensitivity, as in Eq. 3.7. Many methods can be applied to obtain a solution, including conjugate gradients and SVD decomposition.

In the implementation carried out in this work, the method of choice to solve 3.7 was performing a Cholesky decomposition of the $(B_c)^H B_c$ matrices to solve $(B_c)^H B_c a_c = (B_c)^H d_c$. For high values of n_p , this system of equations is close to singular, due to the fact that there is more than one combination of polynomial coefficients that fit the data equally well. This is a common problem in least squares coefficient fitting (see for example Sec. 15.4 of [23]) and indicates that once a certain order is reached, there is not much to be gained in terms of representation accuracy by increasing the order of the polynomial functions. In order to alleviate the numerical instabilities caused by this situation, a constant diagonal regularization matrix with a single value λ was added to the matrix $(B_c)^H B_c$, so that the minimized cost function was:

$$\arg \min_a \|\hat{d} - B_c a\|^2 + \lambda \|a\|^2 \quad (3.9)$$

Intuitively, the extra term avoids the appearance of fitted coefficients with very large magnitudes that cancel each out producing unstable solutions [23]. The global cost function is thus slightly modified to incorporate the regularizing term:

$$\arg \min_{a, I} \|\hat{d} - P F M_p(a, I)\|^2 + \lambda \|a\|^2 \quad (3.10)$$

3.1.3 Initialization

The greedy algorithm used to solve Eq. 3.10 must be initialized by an estimate either of the image or of the coil sensitivity functions. Coil sensitivities obtained from a low resolution scan or a self-calibrated approach, which are usually used to perform SENSE reconstructions, can both be used as an initialization for JSENSE.

3.2 Methods

3.2.1 Algorithm Implementation

The method described in Sec. 3.1.2 was implemented in C. The GSENSE preconditioned conjugate gradient algorithm proposed in Ref. [24] was implemented for the image update. The sensitivity update was implemented by solving Eq. 3.9, as described in Sec. 3.1.2.

In the case of the image update, the inner iterative optimization step was stopped when the magnitude of the cost function gradient, normalized by the magnitude of the data vector, fell below a threshold of 10^{-4} or when the number of iterations reached 100. In the last outer iteration the stopping threshold was lowered to 10^{-7} . For some of the reconstructions higher acceleration factors, the final image was affected by the noise amplification problems mentioned in Sec. 1.3.3. In these cases, Tikhonov regularization, as in Eq. 1.17, was applied to improve the final reconstruction by performing a last regularized GSENSE iteration with the final estimate of the coil sensitivities.

3.2.2 Data Sets

In order to test the performance of the algorithm, tests were carried out with three different data sets. In each case, the k-space was fully sampled and then artificially subsampled before applying the reconstruction algorithm.

SIM Data Set

The first data set, which from now on will be referred to as SIM, was obtained by simulation. A measured reference phantom image with size 256×256 was multiplied by a set of simulated complex coil sensitivity functions and transformed into the frequential domain. The simulated coil array consisted of $n_c = 6$ circular receive coils, surrounding the cylindrical phantom without overlap. The complex sensitivities of these coils were calculated using the Biot-Savart law [9].

BR Data Set

The second data set, which from now on will be referred to as BR, consisted of *in vivo* measurements of a human brain performed with a 1.5 T scanner. A custom-made six-element head phased-array coil was used. A spin echo sequence was performed with the following parameters: FOV= 250 mm, matrix= 256×256 , TE= 5 ms, TR= 1000 ms, slice thickness= 6 mm. Fat suppression was applied.

A low resolution reference scan was performed to obtain low resolution images I_l from the body coil and $I_{l,c}$ from the 6-element phase array coil with a Cartesian acquisition. Coil sensitivity estimates S_c were computed pixel by pixel in the following way:

$$S_c(x, y) = \frac{I_{l,c}(x, y)}{I_l(x, y)} \quad (3.11)$$

PH Data Set

The third data set, which from now on will be referred to as PH, consisted of phantom measurements performed with a 1.5 T scanner. A commercial eight-element head phased-array coil was used. The eight elements were regularly spaced around the phantom. A spin echo sequence was performed with the following parameters: FOV= 250 mm, matrix= 256×256 , TE= 10 ms, TR= 195 ms, slice thickness= 5 mm.

Coil sensitivity estimates were obtained from a low resolution reference scan, as for the BR data set.

3.2.3 Experiments

Polynomial Fitting of Sensitivity Functions

The accuracy of the polynomial representation of coil sensitivities was evaluated by performing a least squares fit of the fully sampled coil images for different values of the maximum polynomial order n_p . Accuracy was quantified by the normalized root mean square error (NRMSE) between the fully sampled coil images I_c and the product of the fitted polynomial sensitivity functions S_c and the image I used for the fitting:

$$\text{NRMSE} = \sqrt{\frac{\sum_{c=1}^{n_c} \|I_c - S_c I\|^2}{\sum_{c=1}^{n_c} \|I_c\|^2}} \quad (3.12)$$

In the case of SIM, the image I was the original one used to produce the data. In the case of BR and PH, it was that of a SENSE reconstruction of the fully sampled data (in this case, the algorithm is usually called CLEAR).

JSENSE Reconstruction

A JSENSE reconstruction was carried out for acceleration factors of 2 and 4, in the case of SIM and BR, and of 2, 4 and 8 in the case of PH. The number of extra ACS lines was 4 for the acceleration factor of 2, 6 and 12 lines for the factor of 4, and 14 and 28 lines for the factor of 8. Reconstructions without ACS lines were also carried out. GSENSE with Tikhonov regularization and GRAPPA reconstructions were carried out for each subsampling scheme (except for schemes without ACS lines in the case of GRAPPA) as representative examples of standard linear reconstruction algorithms.

The coil sensitivities used to initialize the reconstruction process in the case of the measured data sets BR and PH were the estimates obtained from the low resolution scan following Eq. 3.11. In the case of the simulated data SIM, the simulated sensitivity functions used to produce the data set were modified by translation and rotation to simulate a faulty estimation. Figure 3.1 shows both the original and modified coil sensitivity functions.

No convergence criterion was fixed to stop the reconstruction algorithm. The l_2 norm of the error does not necessarily correspond to visual quality of the reconstruction, so it is problematic to determine a systematic stopping criterion. The algorithm was simply allowed to run for 100 iterations for all the reconstructions.

The regularization term λ in Eq. 3.10 was set heuristically by initializing it to a numerically negligible value, trying to perform the Cholesky decomposition and then doubling it each time the decomposition proved to be unstable.




Figure 3.1: Original coil sensitivities used to generate the SIM data (left) and modified coil sensitivities used to initialize the JSENSE reconstruction for this data set.

In every case, the criteria to evaluate the reconstruction results included the final value of the cost function in Eq. 3.4 and the regularized cost function in Eq. 3.10, both normalized by the magnitude of the data vector. Two further measures were used to try to quantify the quality of the reconstructions. The first measure was the normalized root mean square error (NRMSE) between the normalized magnitude of the reconstructed coil images $(S_c I)_n$ and the magnitude of the normalized fully sampled coil images $(I_c)_n$. In both cases the subindex n indicates normalization by the maximum value of the magnitude, as in $x_n = \frac{x}{\max_i(|x_i|)}$. The reconstructed coil images were calculated as the product of the final reconstructed image I and sensitivity functions S_c :

$$\text{NRMSE}_{\text{coil images}} = \sqrt{\frac{\sum_{c=1}^{n_c} \left\| |(I_c)_n| - |(S_c I)_n| \right\|^2}{\sum_{c=1}^{n_c} \|(I_c)_n\|^2}} \quad (3.13)$$

The second quality measure was the NRMSE between the magnitude of the final reconstructed image I_{ref} and the magnitude of a reference image I . This measure was also used in Ref. [33] to evaluate the quality of reconstructions:

$$\text{NRMSE}_{\text{reference image}} = \frac{\| |I_{\text{ref}}| - |I| \|}{\|I_{\text{ref}}\|} \quad (3.14)$$

Once again, in the case of SIM, the reference image was the original one used to produce the data. In the case of BR and PH, it was that of a CLEAR reconstruction of the fully sampled data.

It is important to point out that measures based on l_2 norm errors, which are widely used, do not necessarily correspond to visual quality, so that it is unavoidable to also perform visual examinations of the reconstructed images when assessing the performance of image reconstruction algorithms.

Finally, a further experiment was carried out with the SIM data set to evaluate the robustness of the reconstruction method in the presence of noise. The simulated data from each coil was corrupted with uncorrelated Gaussian noise with constant variance σ^2 . Different values of σ^2 were used: 10^{10} , 10^{11} and 10^{12} . The signal to noise ratio in the k-space that corresponds to these noise levels is respectively: 21.55 dB, 11.55 dB and 1.55 dB. JSENSE, GSENSE and GRAPPA reconstructions were carried out for each of these simulated data sets. The criteria described in the previous paragraph were again used to analyse the results.

In order to counteract the effect of noise amplification due to a bad conditioning of the underlying linear system of equations, in the case of JSENSE and GSENSE, Tikhonov regularization, as defined by Eq. 1.17, was applied. The value of the regularization parameter was the same for the two algorithms, but was empirically varied for the different SNR values and undersampling schemes: 0.7 (SNR of 21.55 dB), 0.9 (SNR of 11.55 dB) and 1 (SNR of 1.55 dB) for an acceleration ratio of 2 with 4 ACS lines, and 0.3 (SNR of 21.55 dB), 0.4 (SNR of 11.55 dB) and 0.9 (SNR of 1.55 dB) for an acceleration ratio of 4 with 12 ACS lines.

3.3 Results

3.3.1 Polynomial Fitting of Sensitivity Functions

Figure 3.2 shows the normalized root mean square error for the polynomial fitting of the coil sensitivities of the three data sets with different maximum polynomial orders. For all three data sets, higher polynomial orders yield more accurate fits. The fitting error for the simulated data reaches a lower level (1%) than that of the two real data sets (around 5%), possibly due to the presence of noise and to the fact that the real coil sensitivity functions are not as smooth as the simulated ones.

3.3.2 JSENSE Reconstruction

Cost Functions

Figure 3.3 shows the final cost function and regularized cost function values for the JSENSE reconstruction of the three data sets with different values of maximal polynomial order in the sensitivity model and different undersampling schemes. In general higher polynomial orders yield lower final values for both the regularized and non-regularized cost function, especially for the SIM data set. Undersampling schemes with less data are fitted more exactly by the optimization algorithm, so that higher acceleration factors with less number of ACS lines yield lower cost function values. It is worth noting that this is not the case for the reconstruction of the SIM data set with an acceleration factor of 4 with 6 ACS lines. This is possibly due to stability problems in the convergence of the optimization procedure.




Figure 3.2: Normalised RMS error of the coil sensitivity function polynomial representations for the SIM (top), BR (center) and PH (bottom) data sets using different values of maximal polynomial order.

Reconstruction Quality Measures

Figure 3.4 shows the final value of the two NRMSE criteria defined in Eq. 3.13 and 3.14 for the JSENSE reconstruction of the three data sets. Reconstruction is carried out for different values of different undersampling schemes and different maximal polynomial orders in the sensitivity model. There does not seem to be a general pattern in the NRMSE measures for the different values of maximum polynomial order. There is, however, a clear pattern which relates different undersampling patterns to NRMSE measures. Lower acceleration factors yield lower NRMSE values. Also, for the same acceleration factor, higher number of ACS lines correspond to lower NRMSE measures. As an example, in the reconstruction of the SIM data, better results were obtained for an




Figure 3.3: Normalised cost function (left) and regularized cost function (right) values for the JSENSE reconstruction of the SIM (top), BR (center) and PH (bottom) data sets for different reduction factors (RF) and number of ACS lines. For the SIM and BR data sets: RF 2 ACS 0 (green), RF 2 ACS 4 (blue), RF 4 with no ACS lines (magenta), RF 4 ACS 6 (black) and RF 4 ACS 12 (red). For the PH data set: RF 2 ACS 4 (green), RF 4 ACS 6 (blue), RF 4 ACS 12 (magenta), RF 8 ACS 14 (black) and RF 8 ACS 28 (red).

acceleration factor of 4 with 12 ACS lines (which corresponds to a total acceleration factor of 3.36) than for an acceleration factor of 2 with no ACS lines.

In general, the results thus seem to suggest that denser grids with more ACS lines generate better reconstructions. In some cases, such as acceleration factor 4 with no ACS lines for the SIM data and both schemes with acceleration factor 8 for the PH data, high NRMSE values seem to indicate very low quality reconstructions.

Visual Results

The image reconstruction results for data sets SIM, BR and PH are shown in Figs. 3.5, 3.7 and 3.9. In each case the JSENSE reconstruction is compared both to a reference image and to regularized GSENSE and GRAPPA reconstructions.

The visual results for JSENSE seem to confirm the NRMSE criteria results. Lower acceleration factors result in better reconstructions, where aliasing artifacts are less visible. The same is true of subsampling patterns with more ACS lines. For the SIM data set, in particular, the number of ACS lines has a considerable impact on the suppression of aliasing artifacts. The number of ACS lines necessary for the suppression seems to augment for higher acceleration factors. For a factor of 2, 4 ACS lines suffice to eliminate almost all artifacts, while for a factor of 4 6 ACS lines are not sufficient, but 12 are. In fact, the visual result for an acceleration factor of 4 with 6 ACS lines confirms that the optimisation algorithm does not converge adequately, resulting in an amplification of the signal at one side of the image that completely degrades the reconstruction.

GSENSE results present more aliasing artifacts than JSENSE reconstructions for almost all accelerating factors. As for JSENSE, reconstruction quality is improved by lower undersampling ratios. The impact of ACS lines, however, is much less important, and in no case leads to an artifact suppression effect comparable to that of JSENSE. Results for the SIM data are of especially low quality, due to the erroneous sensitivity function estimate used in the reconstruction.

Since GRAPPA does not rely on a sensitivity estimate, its results for the SIM data set present almost no aliasing artifacts. For the BR and PH data, the reconstruction is of quite good quality for an acceleration factor of 2, but presents heavy aliasing for higher factors. In all reconstructions, when compared to those of the two algorithms using explicit sensitivity estimates, a certain inhomogeneity effect is noticeable. This is possibly a consequence of the low number of ACS lines used in the reconstruction and of the sum of squares reconstruction of the image.

Coil sensitivity reconstruction results are shown in Figs. 3.6, 3.8 and 3.10, where the coil sensitivity reconstructions are compared to the low resolution scan estimate and in the case of the SIM data

set to the erroneous and original coil sensitivities. In the case of the SIM data, high quality image reconstructions seem to correspond to well corrected sensitivity functions, which are very similar to the original ones. This is the case for a reduction factor of 2 with 4 ACS lines and 4 with 12 ACS lines. In the two other data sets, the corrected sensitivities are in general similar to the reference scan estimates, with large differences only for undersampling schemes that yield low quality image reconstructions (such as for an acceleration factor of 4 with no ACS lines for the BR data) and small, yet detectable, differences for schemes that generate good quality images.

Convergence

Figure 3.11 shows the evolution of the different result evaluation criteria during the first 100 iterations of the JSENSE reconstruction for a subsampling scheme with an acceleration factor of 4 and 12 ACS lines. In all three data sets, the cost function and regularized cost function values decrease sharply at first and then converge to a stable value after 4 or 5 iterations, except for the functions corresponding to the simulated data which continued to decrease even after 100 iterations.

The behaviour of the NRMSE values for the SIM data is similar to that of the cost functions. It decreases sharply at first and then more slowly, without completely converging to a stable value. In the case of the BR and PH data, the NRMSE values do not seem to decrease noticeably. In fact, in the case of the BR data they actually increase during the first iterations.

The seemingly puzzling behaviour of the NRMSE values can perhaps be explained by visual inspection of the intermediate results in Fig. 3.12. While in the SIM reconstruction no noise is present, both the BR and PH data reconstructions suffer from noise amplification. This is especially evident on the BR image. At the same time, the intermediate images show substantial aliasing artifact suppression taking place in the first iterations. It is possible that both effects cancel out each other, yielding NRMSE values which are not very representative of reconstruction quality. It is interesting to note that the noise amplification effect seems to be restricted to the first iterations, so that further JSENSE iterations do not seem to contribute to it. All of the presented images are intermediate updates in the JSENSE reconstruction, so Tikhonov regularization has not been applied to control noise amplification in any of them.

3.3.3 Noisy Data Reconstruction

Figure 3.13 compares the normalized cost function, regularized cost function and NRMSE results at different values of k-space SNR for a reduction factor of 2 with 4 ACS lines and a reduction factor of 4 with 12 ACS lines. The maximum polynomial order used in the sensitivity model was 9 in both cases. All of the cost functions and quality measures degrade at higher noise levels.

The corresponding image reconstruction results are shown in Figs. 3.14 and 3.15, where they are compared to GSENSE with Tikhonov regularization and GRAPPA reconstructions. In each case, the same regularization value was applied to GSENSE and to the final image update of JSENSE. In general JSENSE presents less artifacts than GSENSE for every noise level. The visual quality of the reconstructions is worse at lower SNR values, where noise mainly affects the central regions of the image. Tikhonov regularization alleviates this effect, but also lowers the intensity in those regions.

Once again, since GRAPPA does not use a prior estimate of coil sensitivities, it presents substantially less aliasing artifacts than the GSENSE reconstruction with the faulty sensitivity estimate. In any case, some artifacts are present, in particular for an acceleration factor of 4, possibly due to the lack of regularization and the reduced number of ACS lines. The inhomogeneity effect, possibly caused by the sum of squares reconstruction and the reduced number of ACS lines, is again apparent. However, when comparing it to the GSENSE and JSENSE reconstructions at high noise levels, the effect of regularization and the degradation of the sensitivity estimate in JSENSE produces similar inhomogeneities so that, in this case, the reconstruction quality of JSENSE and GRAPPA is quite similar.

The final coil sensitivity reconstructions are shown in Fig. 3.16. In general, the estimate is still corrected by the algorithm, although at higher noise levels the final estimate is less sharp and does not approximate the original coil sensitivity functions as accurately as in the noiseless case. In fact, the signal level towards the center of each coil sensitivity image becomes higher, relative to the signal level at the borders, as the noise level increases. This may be one of the causes that generate lower signal levels at the center of the final JSENSE image, producing the inhomogeneous effect mentioned in the previous paragraph.




Figure 3.4: NRMSE comparison to the fully sampled coil images (left) and to the reference image (right) for the JSENSE reconstruction of the SIM (top), BR (center) and PH (bottom) data sets for different reduction factors (RF) and number of ACS lines. For the SIM and BR data sets: RF 2 ACS 0 (green), RF 2 ACS 4 (blue), RF 4 with no ACS lines (magenta), RF 4 ACS 6 (black) and RF 4 ACS 12 (red). For the PH data set: RF 2 ACS 4 (green), RF 4 ACS 6 (blue), RF 4 ACS 12 (magenta), RF 8 ACS 14 (black) and RF 8 ACS 28 (red).




Figure 3.5: Image reconstruction of the SIM data set by JSENSE (second row), regularized GSENSE (third row) and GRAPPA (bottom row) for different reduction factors (RF) and number of ACS lines. The original image used to generate the data set is shown in the top row.




Figure 3.6: The top two images show the original coil sensitivities used to generate the SIM data set (left) and the erroneous coil sensitivities used to initialize the reconstruction (right). The remaining images show JSENSE reconstructions for different reduction factors (RF) and number of ACS lines.




Figure 3.7: Image reconstruction of the BR data set by JSENSE (second row), regularized GSENSE (third row) and GRAPPA (bottom row) for different reduction factors (RF) and number of ACS lines. At the top the fully sampled SENSE or CLEAR reconstruction is shown.




Figure 3.8: The top left image shows the low resolution scan estimate of the coil sensitivities for the BR data set. The rest of the images show coil sensitivity reconstructions by JSENSE for different reduction factors (RF) and number of ACS lines.




Figure 3.9: Image reconstruction of the PH data set by JSENSE (second row), regularized GSENSE (third row) and GRAPPA (bottom row) for different reduction factors (RF) and number of ACS lines. At the top the fully sampled SENSE or CLEAR reconstruction is shown.




Figure 3.10: The top left image shows the low resolution scan estimate of the coil sensitivities for the PH data set. The rest of the images show coil sensitivity reconstructions by JSENSE for different reduction factors (RF) and number of ACS lines.




Figure 3.11: Evolution of the cost function (top left), regularized cost function (top right), NRMSE comparison to the fully sampled coil images (bottom left) and NRMSE comparison to the reference image (bottom right) during the first 100 iterations of a JSSENSE reconstruction of the SIM (black), BR (red) and PH (blue) data sets subsampled at a reduction factor of 4 with 12 ACS lines.




Figure 3.12: Evolution of the image reconstruction for some intermediate iterations of JSENSE for the SIM (top row), BR (middle row) and PH (bottom row) data sets subsampled at a reduction factor of 4 with 12 ACS lines.




Figure 3.13: Normalised cost function (top left), regularized cost function (top right), NRMSE comparison to the full sampled coil images (bottom left) and NRMSE comparison to the reference image (bottom right) of the JSENSE reconstruction of the SIM data set with different values of k-space SNR for a reduction factor of 2 with 4 ACS lines (red) and 4 with 12 ACS lines (blue).




Figure 3.14: Regularized JSENSE (top row), regularized GSENSE (middle row) and GRAPPA reconstructions of the SIM data set with different values of k-space noise variance for a reduction factor of 2 with 4 ACS lines.




Figure 3.15: Regularized JSENSE (top row), regularized GSENSE (middle row) and GRAPPA reconstructions of the SIM data set with different values of k-space noise variance for a reduction factor of 4 with 12 ACS lines.




Figure 3.16: The top row shows the original coil sensitivities used to generate the SIM data set (left) and the erroneous coil sensitivities used to initialize the reconstruction (right). The other figures show JSENSE coil sensitivity reconstructions for a reduction factor of 2 with 4 ACS lines (left column) and a reduction factor of 4 with 12 ACS lines (right column) at different values of k-space SNR: 21.55 dB (second row), 11.55 dB (third row) and 1.55 dB (bottom row).

3.4 Discussion

JSENSE reconstruction relies on a nonlinear procedure that iteratively corrects the coil sensitivity functions and the image of a parallel MRI SENSE model. For this, a two-dimensional polynomial sensitivity model is used, which seems to provide a good equilibrium between representation accuracy and an implicit regularization effect, which constrains potential solutions to be smooth functions. JSENSE can also be interpreted as a unifying framework for reconstruction algorithms relying on coil sensitivity estimates, like SENSE, and those that use ACS lines to exploit sensitivity encoding, like GRAPPA.

The experimental results for a variety of examples presented in the previous section seem to show that the method effectively exploits extra information encoded in the center of the k-space to improve the accuracy of the sensitivity function estimate and suppress aliasing artifacts. When compared to SENSE, JSENSE is able to perform better quality reconstruction due to sensitivity correction, but needs a certain number of ACS lines to achieve this. When compared to GRAPPA, JSENSE yields more accurate and homogeneous reconstructions for a similar number of ACS lines, since it is able to exploit the whole measured k-space to correct the explicit sensitivity estimate. Finally, sensitivity correction seems not to be severely degraded by the presence of noise. This is possibly due in part to the regularization implicit in the polynomial model, which hinders the noise from propagating and being amplified between JSENSE iterations.

A major limitation of the JSENSE algorithm is its computational complexity. Although the algorithm usually attains good image quality after a few iterations, each of the iterations is very costly, as it effectively includes a GSENSE reconstruction and a large polynomial fitting problem. Another limitation is the fact that a moderately accurate initial sensitivity estimate is necessary to initialize the reconstruction (this limitation is common to all algorithms based on a SENSE model). For inaccurate initial estimates, a certain number of ACS lines are needed for adequate convergence. Recently, the application of a variable projection method to the solution of the JSENSE nonlinear optimization problem has been proposed to alleviate this problem [32].

Apart from increasing the computational efficiency of the nonlinear optimization problem, an interesting extension to JSENSE would be to study the effects of applying Tikhonov regularization within the inner image update, since this may yield more robust sensitivity estimates, and in general the compatibility of JSENSE with other forms of regularization, which can be applied to the SENSE reconstruction. Another possibility would be to use alternative models for the coil sensitivities within the JSENSE framework. Regularization of sensitivity solutions could then be performed through the use of Sobolev norms as in [29], or the l_1 norm of the sensitivity functions in a transform domain, as is proposed in the following chapter.

Chapter 4

Compressive Sampling of 2D Parallel Magnetic Resonance Data

In this chapter, several Compressive Sampling approaches to the reconstruction of 2D parallel MR images are presented and discussed. In particular, the possibility of improving coil sensitivity function estimates by exploiting their sparse representation in the wavelet domain is explored.

4.1 Theory

4.1.1 Image Sparsity

The first step towards the application of Compressive Sampling to image reconstruction in MRI is to make sure that the reconstructed signal has a sparse representation in a certain domain. The sparsifying transforms which seem to be well adapted to MR images are spatial finite differences and the wavelet transform [18].

Spatial finite differences are calculated as the difference between neighbouring pixels of an image. The l_1 norm of this transform is also known as Total Variation (TV), as it is the sum of the absolute values of all the local variations in the image. Sparse images in this domain tend to be piecewise constant. With the exception of special cases like angiograms, *in vivo* MR images are not well suited for a piecewise constant representation. However, adding a term of total variation regularization has been shown to be useful in reducing artifacts while preserving image edges for SENSE reconstruction of parallel MR data [15].

The wavelet transform W can be interpreted as a multiresolution analysis of the image, in which each coefficient carries both spatial frequency and spatial position information. As we can see in Fig. 4.1, some coefficients, placed in the upper left corner of the wavelet representation, contain coarse details, while others chiefly correspond to fine-scale details. In most images, many of the fine-

scale coefficients are practically negligible. As a result, the images are compressible in the wavelet domain. This is indeed the case for many MR images.




Figure 4.1: Wavelet representation (center) of an MR phantom image (left). The wavelet representation is also shown on a logarithmic scale (right), where fine-scale coefficients are better visible.

MR images can consequently be considered as at least moderately sparse both in the wavelet and in the spatial finite differences domain. In order to exploit this, the Compressive Sampling optimization problem of Eq. 2.3 can be slightly modified to incorporate both sparsifying transforms. A parameter α is introduced to calibrate the relative importance of each term. The resulting extended optimization problem solved to perform the reconstruction is of the form:

$$\arg \min_I \|W(I)\|_{l_1} + \alpha \text{TV}(I) \quad \text{subject to} \quad \hat{d} = P F I \quad (4.1)$$

where \hat{d} is the MR data, F the Fourier transform matrix and P the projection matrix that describes the k-space sampling pattern.

4.1.2 K-Space Sampling Grid Design

As discussed in Sec. 2.2.2, for a k-space sampling grid to be adequate for Compressive Sampling reconstruction there must be sufficient incoherence between the measurement and the synthesis domain. This condition can be tested via the TPSF of the sparsifying transform coefficients. In the case of 2D Cartesian sampling, only the position of the phase-encoded lines must be chosen, as data in the frequency-encoded direction are fully sampled. The distribution of these lines could be adapted to maximize TPSF incoherence.

Such an approach would, however, not take into account the specific structure of image representations in k-space, which do not display a regular energy distribution. Quite on the contrary, the energy is very concentrated in the center of k-space. Missing samples consequently have an effect on

the final reconstruction which is strongly dependent on their position. This places an additional constraint on the design of the k-space sampling scheme, which should be significantly denser towards the center. Fig. 4.2 shows the effects of uniform and variable density random k-space sampling in the final reconstruction.




Figure 4.2: Compressive Sampling results for the reconstruction of a phantom for two different k-space random sampling grids: a uniform random grid (top left) which produces a reconstruction with low frequency irregularities (top right) and a variable density grid (bottom left) which generates a more faithful reconstruction (bottom right). In both grids white points represent samples and black points represent unsampled k-space positions.

Designing an optimal sampling grid in both respects is computationally not tractable. However, following an approach proposed in Ref. [18], it is possible to find adequate sampling patterns by a Monte Carlo approach. The basic idea is to generate a large number of different grids, analyse them and choose the one which is most suitable for Compressive Sampling reconstruction. Let G be an operator indicating whether a phase-encoded line y is sampled ($G(y) = 1$) or not ($G(y) = 0$). In the design process, grids are generated by sampling each $G(y)$ value from a probability density function of the form:

$$P\{G(y) = 1\} \begin{cases} \propto \frac{1}{|y-\bar{y}|^p} & \text{if } |y| > r \\ = 1 & \text{if } |y| \leq r \end{cases} \quad (4.2)$$

where \bar{y} is the y coordinate of the k-space center, p a real-valued parameter which determines how much denser the sampling is towards the center and r the radius of the central fully sampled k-space. For each sampled grid, the TPSF of the sparsifying transform is calculated. Finally the most incoherent grid among all those generated in this way is selected.

4.1.3 Compressive Sampling Sum of Squares Image Reconstruction

Probably the simplest and most direct way of applying Compressive Sampling to parallel MRI data is separately reconstructing each of the coil images I_c by solving problem 4.1. The resulting n_c images can be combined to produce a Sum of Squares (SoS) reconstruction of the image, where each pixel $I_{\text{SoS}}(x, y)$ is obtained as a combination of the corresponding coil image pixels $I_c(x, y)$ in the following way:

$$I_{\text{SoS}}(x, y) = \sqrt{\sum_{c=1}^{n_c} I_c(x, y)^2} \quad (4.3)$$

4.1.4 Compressive SENSE Image Reconstruction

The SoS reconstruction described in the previous section takes into account coil sensitivity information neither implicitly nor explicitly. However, if an estimate of the coil sensitivities is available, it is possible to exploit it by incorporating the SENSE encoding model presented in Sec. 1.3.2 into the Compressive Sampling reconstruction. Such an approach has been independently proposed by several authors [11], [31], [34]. If both the wavelet and the spatial finite difference transforms are selected as sparsifying transform domains, the resulting optimization problem solved during the reconstruction is:

$$\arg \min_I \|W(I)\|_{l_1} + \alpha \text{TV}(I) \quad \text{subject to} \quad \hat{d} = E I \quad (4.4)$$

where E is the SENSE encoding matrix defined in Eq. 1.11, which includes the explicit expression of the coil sensitivity functions, and α a constant that controls the relative importance of the two sparsifying transforms in the reconstruction. From now on, we will refer to this algorithm as Compressive SENSE or CSENSE.

4.1.5 Compressive Sampling Joint Reconstruction

In Chapter 3, we described JSENSE, an algorithm performing a simultaneous estimate of both the MR image and the coil sensitivity functions. For some examples, JSENSE obtained better quality images than algorithms dependent on linear estimates of MR images. In particular, it managed to

correct inaccurate prior estimates of the coil sensitivities, which severely degraded the performance of the SENSE algorithm. JSENSE depends on a polynomial model to represent the sensitivity functions and on a greedy algorithm to perform the reconstruction. In this section, we propose an algorithm which, like JSENSE, performs a joint reconstruction of the image and the coil sensitivities by solving a nonlinear optimisation problem.

Unlike JSENSE, the smoothness constraint on the possible solutions for coil sensitivities is not applied through a preimposed model, but through a constrained l_1 optimization problem, as in Compressive Sampling. The main motivation for this approach is that coil sensitivities are generally smoothly varying functions, which can be compressed in a wavelet representation. This has been exploited, for instance, to improve sensitivity estimates by applying wavelet denoising techniques [12]. The resulting nonlinear optimisation problem, which must be solved to perform the joint reconstruction, is of the form:

$$\arg \min_{S,I} \|W(I)\|_{l_1} + \alpha \text{TV}(I) + \lambda_S \|W(S)\|_{l_1} \quad \text{subject to} \quad \hat{d} = P F M(S, I) \quad (4.5)$$

where $M(S, I)$ represents the pixel by pixel product of the image I and the sensitivity functions S , and parameters λ_S and α weigh the different l_1 terms to be minimized. The proposed method to solve Eq. 4.5 is a greedy approach, which iteratively updates the image and sensitivity estimates.

Image Update

The image update is calculated by applying the CSENSE algorithm presented in Sec. 4.1.4. Equation 4.4 is solved for a matrix E formed with the current coil sensitivity function estimate.

Sensitivity Update

In order to calculate an estimate of the sensitivities, an image estimate is supposed to be known, as in the sensitivity update of Sec. 3.1.2. An encoding matrix H can then be defined to represent the linear transformation that maps the coil sensitivity function to the sampled data of that particular coil. This transformation includes a pixel by pixel multiplication with the image I , an application of the Fourier transform F and subsampling in the k-space domain, as defined by a projection matrix P . For each coil c , the transformation can consequently be expressed as follows:

$$d_c = P F I S_c = H S_c \quad (4.6)$$

where, for convenience of notation, I denotes a diagonal matrix containing the image pixel values. Exploiting the sparse representation of the sensitivity function in the wavelet domain, a Compressive

Sampling framework can be applied to calculate the update. The sensitivity function of each coil is consequently calculated as the solution of the following constrained optimization problem:

$$\arg \min_{S_c} \|W(S_c)\|_{l_1} \quad \text{subject to} \quad \hat{d} = H S_c \quad (4.7)$$

In this case, it would be counterproductive to add a Total Variation regularization term in the optimization problem, since this would favour solutions with sharp edges instead of smooth functions.

4.2 Methods

4.2.1 K-Space Sampling Grid Design

In order to generate the k-space sampling grids, the Monte Carlo design technique described in Sec. 4.1.2 was implemented. The TPSF of the wavelet transform was used to measure the incoherence of the different grids. The incoherence for a given coefficient wavelet coefficient x_i was quantified by the normalized difference between the TPSF maximum and the highest secondary lobe:

$$\text{Inc}(x_i) = \frac{|\text{TPSF}(x_i, x_i) - \max_{i \neq j} \text{TPSF}(x_i, x_j)|}{|\text{TPSF}(x_i, x_i)|} \quad (4.8)$$

Calculating this value for each of the n_I wavelet coefficients of each grid would be extremely inefficient. It was empirically ascertained, however, that coefficients at the same resolution level had very similar TPSF functions. It was thus decided to test only one of the coefficients for each level in order to reduce the computational complexity of the design procedure.

The Monte Carlo procedure can be further optimized by exploiting the fact that finer level coefficients are more limiting, as to domain coherence, than coarser level coefficients. This is a consequence of the variable density sampling scheme. Coarse levels of detail tend to consist of low frequency information, which is sampled more completely as described by Eq. 4.2. Fine level details correspond to information encoded primarily in high frequency regions, which are undersampled more severely. This suggests using Eq. 4.8 for an arbitrary coefficient at the finest detail level of the wavelet multiresolution decomposition as a criterion to measure the incoherence of a given grid.

The optimized Monte Carlo procedure which was finally implemented consists of calculating the measure given by Eq. 4.8 for each sampled grid at a certain fine level coefficient. When a sampled grid proves to be more incoherent at that level than the previously sampled ones, lower coefficients from each wavelet transform scale are tested to confirm that there is indeed less coherence at those levels. This was repeated 10^5 times. Finally, the most incoherent grid was selected. In this way, three grids were generated with 133, 73 and 36 fully sampled k-space lines. This corresponds to

acceleration factors of 1.92, 3.51 and 7.11 respectively. More details about the selected grids for the different acceleration factors can be found in App. A.1.

4.2.2 Algorithm Implementation

The different algorithms and methods described above were implemented in MATLAB. Wavelet transforms were calculated using the WAVELAB library [5]. Daubechies 4 wavelets were selected to implement the sparsifying wavelet transform for all reconstructions.

For the SoS reconstruction of coil images, Michael Lustig's Sparse MRI software [17] implementing a nonlinear conjugate gradient algorithm with a backtracking line-search was directly applied. Optimization problems 4.4 and 4.7 were solved by adding two operators to the implementation. In the case of CSENSE, the new operator implemented the sensitivity encoding matrix E , as in Eq. 1.11, and thus included pixel by pixel multiplication with the coil sensitivities before performing the undersampled Fourier transform. In the case of sensitivity estimation, the operator implemented the encoding matrix H , as in Eq. 4.6, and included pixel by pixel multiplication with the current image estimate.

The optimization algorithm does not solve Eq. 2.3 directly, but rather minimizes a Lagrangian representation of the problem, where the Lagrangian multiplier λ must be fixed a priori. This effectively transforms the constrained problem of Eq. 4.1 into an unconstrained least squares problem with an l_1 regularization term:

$$\arg \min_I \|\hat{d} - PFI\|_{l_2} + \lambda \|W(I)\|_{l_1} + \alpha \text{TV}(I) \quad (4.9)$$

λ and α can be interpreted as regularization parameters which quantifies the compromise between l_1 norm minimization of the reconstruction in the synthesis domain and approximation accuracy to the sampled measurements in the measurement domain. In the case of the Compressive Sampling estimation of the coil sensitivity functions, α was set to zero. A more detailed description of the optimization algorithm can be found in Appendix I of Ref. [18].

The joint estimation scheme was implemented by an iterative application of CSENSE and the Compressive Sampling sensitivity estimation algorithm. In this case, three different l_1 norm regularization parameters had to be calibrated.

4.2.3 Data Sets

The three data sets described in section 3.2.2 were used for the experiments.

4.2.4 Experiments

Signal Sparsity

In order to test the degree of sparsity of both MR images and coil sensitivity functions, several experiments were done using the WAVELAB library. An MR brain image from data set BR was used to quantify the effect of truncating its wavelet representation in a nonlinear way. More precisely, the images generated by keeping the 20%, 10%, 5%, 2.5% and 1% of the largest Daubechies 4 wavelet coefficients were visually evaluated. This truncation process is nonlinear because coefficients are chosen depending on the specific representation of the signal and not *a priori*.

In the case of the sensitivity functions, the estimates obtained from a low resolution scan, as explained in section 3.2.2, for the data sets BR and PH were analysed in a similar way. Their representation in different wavelet transform domains, such as Daubechies, Symmlet, Battle, Coiflet, Vaidyanathan and Beylkin, were truncated to different extents. For each truncated reference, the normalised l_2 error of the representation was calculated.

Image Reconstruction

Compressive Sampling SoS and CSENSE reconstructions of the SIM, BR and PH data sets were performed for the three sampling grids with acceleration factors of 3.51 and 7.11. In the case of CSENSE, reconstructions were carried out using both the erroneous and correct SIM coil sensitivity functions (see Sec. 3.2.3) and the low resolution scan estimates of the BR and PH coil sensitivities (see section 3.2.2). The criterion for evaluating the reconstruction was the normalised root mean square error (NMRSE) of the magnitude of the reconstructed image, as defined in Eq. 3.14. In the case of SIM, the reference image was the one used to generate the data set, for BR and PH it was the fully sampled SENSE reconstruction.

CSENSE was also applied on simulated noisy data. The same simulated data as in the case of JSENSE were used (see section 3.2.3). Reconstructions were carried out for the grids with acceleration factors of 1.92, 3.51 and 7.11. The same criteria as in the case of the noiseless data were applied to evaluate the reconstruction.

In order to compare the results with image reconstructions at lower acceleration factors, additional CSENSE reconstructions were performed for the grid with an acceleration factor of 1.92 for both the noisy and noiseless SIM data and the BR and PH data sets.

Joint Image and Coil Sensitivity Reconstruction

In order to evaluate the Compressive Sampling coil sensitivity estimation, the algorithm was applied to the SIM, BR and PH data sets for the sampling grids with acceleration factors of 3.51

and 7.11. The reference images used were the original SIM image and the fully sampled SENSE reconstructions of the BR and PH images. The criterion for evaluating the reconstruction was the normalised root mean square error (NMRSE) between the product of the sensitivity estimate and the reference image of the fully sampled coil images, as defined in Eq. 3.13. The reconstruction results were also visually evaluated.

The joint estimation algorithm was applied to the SIM data set in the case of the erroneous prior sensitivity estimate for an acceleration factor of 1.92. In this case the wavelet and TV regularization parameters for the image were fixed to 10^4 and the wavelet regularization parameter for the sensitivity to 10^9 . In the case of the BR and PH data sets, the image reconstructions for acceleration factors of 1.92 and 3.51 were of high quality, so the algorithm was only applied for an acceleration factor of 7.11. In both cases, visual inspection was used to evaluate the performance of the algorithm.

4.3 Results

4.3.1 Signal Sparsity

Figure 4.3 shows the result of truncating the wavelet representation of the MR measured brain image to different extents. The position of the coefficients used for the representation are also shown. Eliminating up to 95% of the coefficients does not produce significant differences in the image. Coefficients are principally eliminated from the finer scales of the wavelet representation, indicating that they may correspond to noise or high frequency fluctuations that have no visible influence on the image quality. When 97.5% or more coefficients are discarded, the quality of the approximation is substantially degraded, although the general structure of the image is nevertheless well conserved. In any case, degradation only occurs at very high compression rates, so that we can conclude that the image is indeed sparse in the wavelet domain.

The graphs in Fig. 4.4 quantify the precision of truncated wavelet representations for the low resolution scan coil sensitivity estimates of data sets BR and PH. Vaidyanathan and Beylkin wavelets seem not to be well adapted to the compression of coil sensitivities, whereas Battle, Daubechies, Coiflet and Symmlet wavelets perform much better. The BR coil sensitivity estimate is particularly well compressed. For this data set, normalised approximation errors of under 1% are attained with only 5% of the wavelet coefficients. In the case of the PH coil sensitivities, 18% of the coefficients are needed to obtain such low approximation errors. In general, for both data sets a quite accurate sparse representation of the sensitivity estimate is obtained.




Figure 4.3: Sparsifying approximation of the CLEAR reconstruction of the BR data set image obtained by keeping a certain percentage of the largest wavelet coefficients. At the right of each image, the position of those coefficients in the wavelet representation of the image is shown.

4.3.2 Compressive Sampling SoS Image Reconstruction

All of the images obtained by the Compressive Sampling SoS reconstruction of the three test data sets are presented in App. B.1. In each case, the reconstructed image visually selected as being of best quality corresponded to the second largest value (10^4 for SIM and 5 in the case of BR and PH) of wavelet regularization and an equal value of TV regularization. This selected set of reconstructions is shown in Fig. 4.5. In general, the images are very inhomogeneous due to the sum of squares reconstruction. The reconstruction of the three data sets for the lowest acceleration factor presents no aliasing artifacts and is quite sharp. The reconstruction for an acceleration factor of 3.51 is similar but of lower quality. The highest acceleration factor, however, yields reconstructions with blurred details and traces of aliasing artifacts.




Figure 4.4: Normalised deviation from the original BR (top) and PH (bottom) coil sensitivity functions for wavelet representations truncated to different extents. The type of wavelet used were *Battle* (red), *Daubechies* (blue), *Coiflet* (magenta), *Symmlet* (black), *Vaidyanathan* (yellow) and *Beylkin* (green).




Figure 4.5: Selected Compressive Sampling sum of squares reconstructions of the SIM (top row), BR (middle row) and PH (bottom row) data sets subsampled at reduction factors (RF) of 1.92 (left column), 3.51 (middle column) and 7.11 (right column). For the SIM data set, the value of the wavelet and TV regularization parameters was 10^4 . For the BR and PH data sets their value was 5.

4.3.3 CSENSE Image Reconstruction

Figure 4.6 shows the NMRSE results for the CSENSE reconstruction of the SIM (using the correct coil sensitivity function), BR and PH data sets. This measure is minimised at an intermediate level of wavelet regularization for three of the TV regularization levels. In the case of the higher level of TV regularization, the NMRSE just increases monotonously with the magnitude of the wavelet regularization parameter. The combinations of TV and wavelet regularization which yield a better match to the reference image are either a small value for the wavelet regularization parameter combined with a large value (100 times larger) of the TV regularization parameter, or an intermediate wavelet regularization level with the same value for the TV regularization parameter.

All of the images are presented in App. B.2, including reconstructions using the erroneous sensitivity of the SIM data set. In each case, the reconstructed image visually selected as being of best quality corresponded to the second largest value (10^4 in the case of SIM and 0.3 in the case of BR and PH) of wavelet regularization and an equal value of TV regularization. This selected set of reconstructions is shown in Fig. 4.7. It is interesting to note that for acceleration factors of 1.92 and 3.51 the erroneous sensitivity estimate in the SIM data set does not have as much an impact as in the SENSE reconstruction (see Sec. 3.3.2). For the remaining data sets, the reconstruction was of good quality, although some aliasing artifacts are visible in the BR reconstruction for an undersampling factor of 7.11. It is worth noting that, in the case of SIM and BR where the number of coils is equal to 6, this undersampling factor is superior to the theoretical acceleration factor that can be achieved by linear reconstruction methods.

Figure 4.8 shows the NMRSE error for the CSENSE reconstruction of the SIM data set with the three different levels of added Gaussian noise. Lower error values were obtained for an intermediate level of wavelet and TV regularization.

All of the visual results for the reconstruction of the noisy SIM data are presented in App. B.3. As predicted by the comparison to the reference image, the best images for the lowest noise level were obtained for an intermediate value of wavelet and TV regularization ($2 \cdot 10^3$ for both). In the case of lower SNR values (11.55 and 1.55 dB), however, the best reconstructions were yielded by same level of wavelet regularization combined with a higher level of TV regularization ($2 \cdot 10^5$). The selected set of reconstructions is shown in Fig. 4.9. For the two lower noise levels most of the noise is eliminated without significantly degrading the image quality. For the highest noise level, the noise in the center of the image cannot be effectively suppressed without introducing more aliasing artifacts.




Figure 4.6: NRMSE comparison to the reference image values of the SIM (top row), BR (middle row) and PH (bottom row) data set CSENSE reconstruction for two subsampling schemes with reduction factors of 7.11 (left column) and 3.51 (right column). Reconstructions were carried out with different values of the l_1 wavelet regularization parameter λ and for several values of the total variation parameter: 0 (red), 10^{-2} (blue), 1 (green) and 10^2 (black) times λ .




Figure 4.7: Selected CSENSE reconstructions of the SIM data set with the erroneous sensitivity estimate (top row), the SIM data set with the correct sensitivity estimate (second row), the BR data set (third row) and the PH data set (bottom row) subsampled at reduction factors (RF) of 1.92 (left column), 3.51 (middle column) and 7.11 (right column). For the SIM data set, the value of the wavelet and TV regularization parameters was of 10^3 . For BR and PH their value was 0.3.




Figure 4.8: NRMSE comparison to the reference image of the SIM data set CSENSE reconstruction for different noise levels and two subsampling schemes with reduction factors of 7.11 (left column) and 3.51(right column). Reconstructions were carried out with different values of the l_1 wavelet regularization parameter λ and for several values of the total variation parameter: 0 (red), 1 (blue) and 10^2 (green) times λ . The k-space SNR values were: 21.55 dB (top row), 11.55 dB (middle row) and 1.55 dB (bottom row).




Figure 4.9: Selected CSENSE reconstructions of the SIM data set for k-space SNR values of 21.55 dB (left column), 11.55 dB (middle column) and 1.55 dB (right column) using the correct sensitivity estimate. The data set was subsampled at reduction factors of 3.51 (top row) and 7.11 (bottom row). The value of the wavelet regularization parameter was of $2 \cdot 10^3$. The value of the Total Variation regularization parameter was of $2 \cdot 10^5$.




Figure 4.10: NRMSE of the Compressive Sampling coil sensitivity function estimation for the SIM (top row), BR (middle row) and PH (bottom row) data sets. The results correspond to two subsampling schemes with reduction factors of 7.11 (red) and 3.51 (blue).

4.3.4 Compressive Sampling Joint Reconstruction

Figure 4.10 shows the results for the sensitivity estimation using the reference images. For the two subsampling schemes analysed in the graph, the NRMSE error reaches its lowest values for intermediate wavelet regularization levels.

The different sensitivity estimates corresponding to different values of the wavelet regularization parameter are presented in App. B.4. As suggested by the graphs, the best approximations were obtained for intermediate wavelet regularization levels ($5 \cdot 10^9$ for SIM, 2 for BR and 50 for PH). The corresponding set of coil sensitivity estimates is shown in Figs. 4.11, 4.12 and 4.13. The estimates are quite smooth, hardly incorporate image detail and resemble the low resolution scan estimates in the cases of PH and BR and the original sensitivity functions in the case of SIM quite closely.




Figure 4.11: Original simulated sensitivity (top left) and coil images (top right) of the SIM data set compared to the reconstructed sensitivities for acceleration factors 3.51 (middle row) and 7.11 (bottom row). The estimated sensitivities are shown on the left column and their product with the reference image in the right column.




Figure 4.12: Low resolution sensitivity estimate (top left) and fully sampled coil images (top right) of the BR data set compared to the reconstructed sensitivities for acceleration factors 3.51 (middle row) and 7.11 (bottom row). The estimated sensitivities are shown on the left column and their product with the reference image in the right column.




Figure 4.13: Low resolution sensitivity estimate (top left) and fully sampled coil images (top right) of the PH data set compared to the reconstructed sensitivities for acceleration factors 3.51 (middle row) and 7.11 (bottom row). The estimated sensitivities are shown on the left column and their product with the reference image in the right column.

Figure 4.14 shows the first three iterations of the joint estimation in the case of the SIM data set at an acceleration factor of 1.92. Initialization was done with the erroneous SIM sensitivity estimate. The sensitivity estimates for the three first iterations are shown in Fig. 4.16. The algorithm corrects the sensitivity estimate and as a result aliasing artifacts disappear from the reconstructed image.




Figure 4.14: The three first image updates in the joint reconstruction of the SIM data set image and sensitivity function for an acceleration factor of 1.92.

Figure 4.15 shows the first three image updates of the joint estimation when applied to the BR data set at an acceleration factor of 7.11. It seems that some of the aliasing in the first reconstruction of the image is eliminated in later iterations. The sensitivity estimate, shown in Fig. 4.17, seems to be degraded by the reconstruction process. On the one hand, it presents some aliasing artifacts, on the other hand, it is not as smooth as the original estimate, in particular at areas where the image has sharp transitions.




Figure 4.15: The three first image updates in the joint reconstruction of the BR data set image and sensitivity function for an acceleration factor of 7.11.




Figure 4.16: The initial erroneous sensitivity estimate and the two first image updates in the joint reconstruction of the SIM data set image and sensitivity function for an acceleration factor of 1.92.




Figure 4.17: The initial low resolution sensitivity estimate and the two first image updates in the joint reconstruction of the BR data set image and sensitivity function for an acceleration factor of 7.11.

4.4 Discussion

The image results for CSENSE were of good quality at each of the tested acceleration factors. Reconstructions for an acceleration factor superior to the number of coils used for data acquisition, which is the traditional limit for parallel MRI reconstruction, were successfully carried out. Moreover, the reconstruction seemed to be quite robust in the presence of noise. In the case of the noisy reconstruction data, however, results need to be interpreted with care, since the piecewise constant structure of the phantom which was used is ideally suited for TV regularization based denoising. *In vivo* MR images, with some exceptions like angiograms, do not usually have such a simple structure and it is possible that strong TV regularization may deteriorate the quality of their reconstruction by eliminating small details, as was observed in the case of the BR data set at high acceleration factors. In any case, linear reconstruction techniques failed to even perform the reconstruction at those factors.

The results for Compressive Sampling SoS reconstructions were of lower quality, as they did not exploit sensitivity-encoded information. The method could, however, be useful to calculate an initial coil sensitivity or image estimate within an iterative joint estimation framework.

The results obtained for the estimation of coil sensitivity functions from accurate image estimates through l_1 norm minimization are promising. When integrated in the joint estimation algorithm, however, coil sensitivity estimates are significantly degraded in the case of *in vivo* data. The reason for this could be that, even though regular functions like coil sensitivities are sparse in the wavelet domain, the opposite statement is not necessarily true, as there are many irregular images that admit sparse wavelet representations. As a consequence, minimizing the l_1 norm in the wavelet domain does not suffice to assure smooth estimates. The results for the application of the Compressive Sampling joint reconstruction, when compared to JSENSE results, suggest that favouring smooth solutions is essential in order to effectively correct coil sensitivity functions using inaccurate image estimates. It is consequently necessary to have a regularizing term that imposes function regularity and this is not the case for the l_1 norm minimization of the wavelet transform. Possible solutions to this problem could be the use of other transforms which are better adapted to smooth functions, such as ridgelets [3], or adding a further regularization term based on a Sobolev norm, as in [29].

Obtaining an accurate approximation of the coil sensitivity functions from an image estimate through Compressive Sampling could also be useful in a setting different from joint reconstruction. It could be used in the estimation of coil sensitivities from a low resolution scan. Instead of simply dividing the whole body scan image by each of the coil images, as explained in section 3.2.2, Compressive Sampling of the coil sensitivities could be applied to yield a more robust estimate. In this framework, compressive sampling of the sensitivities would play a similar role as wavelet

denoising [12].

A major difficulty in the application of the Compressive Sampling techniques presented in this chapter is the empirical calibration of the l_1 regularization parameters. Developing an automatic criterion for the selection of adequate values for the different parameters involved would be highly desirable. Such a criterion could for example take into account the level of estimated noise, which can be used to determine an upper bound for the l_2 norm error of the data approximation as suggested in Ref. [18]. The difficulty of determining the value of the regularization parameters for the reconstruction of parallel MR data is, however, far from unique to the Compressive Sampling framework. In fact, it is encountered in the vast majority of current reconstruction methods, as almost all of them rely on some form of regularization.

Chapter 5

SHARK (Sparse incoherent Arbitrary Resampling in K-space)

As mentioned in Chap. 2, one of the main difficulties in the application of Compressive Sampling methods to MR imaging is the physical implementation of adequate k-space trajectories. In 3D imaging this can be alleviated by combining subsampling in both phase-encoding directions. Similarly, in dynamic imaging, irregular subsampling in the time direction can be applied to accomplish the necessary incoherence. However in the case of 2D imaging, subsampling only makes sense along the one phase encoding direction. Additionally, hardware considerations limit the possible trajectories which can be accurately implemented. To make matters worse, most MR scans in the clinic today are based on Cartesian trajectories, which hinder Compressive Sampling reconstruction due to the coherence between the synthesis and measurement domains (see Sec. 2.2.2).

A solution to this problem could be to resample k-space data from an undersampled Cartesian grid to an arbitrary grid. This is obviously not possible unless there is an additional source of information with which to carry out the interpolation. The main idea underlying the Sparse incoherent Arbitrary Resampling in K-space (SHARK) algorithm presented in this chapter is that sensitivity encoding can be exploited to obtain such a resampling.

An approach which bears certain similarities to SHARK has recently been presented [27]. This algorithm uses the so called GRAPPA Operator Gridding (GROG) to produce an irregular sample distribution. Unaliased images are then reconstructed from the undersampled data employing a conjugate gradient algorithm. The point of applying the resampling step is to accelerate acquisition as justified by the generalized sampling theorem. This theorem states that unaliased images can be reconstructed as long as the average sampling rate is equal to the Nyquist rate, even if it is violated in portions of k-space. GROG resampling is used to facilitate the suppression of aliasing artifacts. Consequently this approach is equivalent to SHARK without l_1 minimization and thus

without exploiting the sparsity of MR images.

5.1 Theory

5.1.1 Two-Step Reconstruction

SHARK is a two-step algorithm. The input data can be any parallel MR data, including data sampled on a Cartesian grid. After the initial regular MR parallel acquisition, k-space interpolation is carried out by combining the signals acquired by the different coils. The interpolated k-space samples belong to a previously determined irregular grid which is optimized for a certain synthesis domain. This step of the algorithm uses coil sensitivity information to compute k-space data distributed on an irregular grid that is adequate for Compressive Sampling reconstruction.

The second step basically consists of applying Compressive Sampling to the resampled data. If the original image is sufficiently sparse in the synthesis domain for which the resampling grid is designed, this will produce an accurate reconstruction, as described in Chap. 2. In this way, the incoherence, which is artificially generated with the help of sensitivity encoding, makes it possible for the Compressive Sampling approach to exploit image sparsity. Fig. 5.1 is a graphical description of the reconstruction procedure.




Figure 5.1: Description of the different steps in the SHARK algorithm.

5.1.2 K-space data interpolation

K-space interpolation can be carried out based on the two main approaches to parallel MR data processing: using an explicit sensitivity function estimate or implicitly exploiting the sensitivity information encoded in ACS lines. In the first case, it would be necessary to carry out a SENSE-like reconstruction in k-space. Examples of such algorithms are kSPA [16] and direct SENSE

(dSENSE) [28]. dSENSE was the chosen approach in the current implementation of SHARK.

In order to perform image reconstruction in k-space, dSENSE approximates each k-space sample by a linear combination of only a subset of the acquired data. The procedure is equivalent to calculating an approximation for the estimator of the image based on the SENSE encoding described in Sec. 1.3.2, only in k-space instead of the image domain.

Pixel by pixel multiplication in the image domain is equivalent to convolution in k-space. The effect of sensitivity encoding on the signal in k-space can consequently be described as a convolution between the k-space representation of the image $\tilde{I} = F I$ and the coil sensitivities $\tilde{S}_c = F S_c$. A linear operator C can be defined, which represents this convolution, and thus maps the k-space representation of the image \tilde{I} to the k-space representation of the coil images \tilde{I}_c . Introducing a projection matrix P to represent k-space undersampling, we obtain an expression in k-space equivalent to Eq. 1.11:

$$\tilde{I}_c = P C \tilde{I} \quad (5.1)$$

The interest in this alternative approach stems from the fact that Eq. 5.1 can be truncated to only take k-space data into account that is situated in a local neighbourhood. This makes it possible to calculate \tilde{I} by direct inversion of matrices of small dimensions, consequently avoiding the application of iterative methods as in GSENSE (hence the name of Direct SENSE). A more detailed account of the algorithm can be found in Ref. [28].

Alternatively, a GRAPPA-like approach can be used to perform k-space interpolation. In Ref. [8] Griswold *et al.* argue that the GRAPPA reconstruction can be reformulated as a matrix operator that shifts data in k-space. Using this formalism, they prove that there exists an infinitesimal GRAPPA operator that shifts data in k-space by arbitrarily small amounts. Resampling the k-space data on the chosen arbitrary grid for Compressive Sampling reconstruction could then be accomplished through repeated applications of this infinitesimal GRAPPA operator. The idea is further developed in [26]

5.1.3 K-Space Sampling Grid Design

As already discussed in Sec. 4.1.2, designing an adequate grid for Compressive Sampling of MR data can be done through a Monte Carlo procedure. In the case of SHARK, an additional constraint must be taken into account when designing the sampling grid. K-space interpolation by dSENSE is more accurate at locations which are near measured k-space samples [28]. As a result, resampling should favour points lying near those locations. The uniform probability density function, which determines with which probability each point of the grid operator $G(x, y)$ (in this case the grid

operator takes values for each pixel, and not for each phase-encoded line) is sampled or not when generating the grids in the Monte Carlo procedure, can be adapted to do so:

$$P\{G(x, y) = 1\} \begin{cases} \propto \alpha(x^2 + y^2)^{-\frac{p}{2}} + (1 - \alpha)|y - y_s|^{-q} & \text{if } \sqrt{x^2 + y^2} > r \\ = 1 & \text{if } \sqrt{x^2 + y^2} \leq r \end{cases} \quad (5.2)$$

where (x, y) are the coordinates relative to the k-space center, y_s the position of the nearest sampled phase-encoded line, p and q real-valued parameters which determine how much denser the sampling is near the sampled lines and the k-space center respectively, and r the radius of the fully sampled k-space area. Parameter α calibrates the relative importance between performing denser sampling near the center and denser sampling near the sampled locations. As in Sec. 4.1.2, the TPSF of the sparsifying transform for each of the grids generated in this way is calculated and used to select the most incoherent grid.

5.1.4 Motivations for SHARK

The point of carrying out a Compressive Sampling reconstruction after having applied k-space interpolation could be questioned in the following way: Why not interpolate all of k-space and reconstruct the image directly as in GRAPPA or dSENSE? There are two possible answers to this question which correspond to two potential applications of the algorithm.

The first answer is that at high subsampling factors, a full resampling with GRAPPA or dSENSE may be degraded by k-space interpolation errors. However, it may still be possible to interpolate accurately enough on the arbitrary grid to make a Compressive Sampling reconstruction possible. This imposes a double constraint on the position of the resampled points. They must be located close enough to the original samples to ensure an accurate interpolation, but at the same time their relative positions must be incoherent enough for the Compressive Sampling algorithm to produce a good reconstruction.

The second answer is that a Compressive Sampling reconstruction may be more robust against noise than a SENSE or GRAPPA reconstruction. This is caused by the fact that Compressive Sampling exploits the sparsity of the original image to eliminate noise. The application of SHARK would then be close in spirit to performing l_1 regularization denoising.

5.2 Methods

5.2.1 K-Space Sampling Grid Design

In order to generate the k-space sampling grids, the Monte Carlo design technique described in Sec. 5.1.3 was implemented. The practical details of the implementation are essentially the same

as in Sec. 4.2.1. The original regular grids used to measure the data follow three different regular subsampling patterns: acceleration factor of 4 with 12 ACS lines and acceleration factor of 8 with 14 and 28 ACS lines. For each of these original grids, an arbitrary resampling grid was designed by the Monte Carlo procedure. An additional resampling grid was generated in each case by combining all the samples of each regular grids with those of its corresponding Monte Carlo grid. More details about the designed grids can be found in Sec. A.2.

5.2.2 Algorithm Implementation

An implementation of dSENSE in C was used for the k-space interpolation step of the algorithm. For the Compressive Sampling step, Michael Lustig’s Sparse MRI software [17] was used. As commented in Sec. 4.2.2, this software solves an unconstrained least squares problem with l_1 norm regularization terms, where the regularization constants must be previously determined. The chosen sparsifying transforms were the Daubechies 4 wavelet and spatial finite differences.

5.2.3 Data Sets

The data sets used for testing the algorithm were SIM and PH (see Sec. 3.2.2).

5.2.4 Experiments

K-Space Interpolation Accuracy

An important hypothesis in SHARK reconstruction is that the accuracy of interpolation of a k-space point depends on its distance to the nearest sampled k-space location. In order to verify this, the whole k-space of the PH data set image was reconstructed using dSENSE for three regular subsampling schemes: acceleration factor of 4 with 12 ACS lines and acceleration factor of 8 with 14 and 28 ACS lines. The normalised root square error (NRSE) of each k-space position (k_x, k_y) was calculated in order to quantify k-space interpolation accuracy:

$$\text{NRSE}(k_x, k_y) = \sqrt{\frac{[\tilde{I}(k_x, k_y) - \tilde{I}_{\text{full}}(k_x, k_y)]^2}{\tilde{I}_{\text{full}}(k_x, k_y)^2}} \quad (5.3)$$

where \tilde{I}_{full} is the k-space representation of an image reconstructed with dSENSE using the fully sampled data.

Total Variation Reconstruction

As an initial simple experiment to explore the viability of SHARK reconstruction, reconstructions of the SIM data set were performed for acceleration factors 2 and 4 by reinterpolating on completely random grids. The second step of the reconstruction was carried out by applying Compressive Sampling with Total Variation regularization. The normalised root mean square error (NMRSE)

approximation to the magnitude of the SIM reference image, as defined in Eq. 3.14, was used to evaluate the quality of the reconstruction.

Reconstruction of Phantom Data

The six grids designed with the Monte Carlo procedure were used to reinterpolate the PH data from the three original regular undersampling grids. The SHARK reconstructions were then visually compared to complete dSENSE and regularized SENSE reconstructions of the data.

5.3 Results

5.3.1 K-Space Interpolation Accuracy

Fig. 5.2 shows the normalised k-space error defined by Eq. 5.3 for each of the k-space positions in the complete k-space dSENSE reconstruction of the PH data from the three regular undersampling schemes. Sampled locations are marked with a red cross. For the three grids, most of the largest errors occur at unsampled locations, suggesting that the hypothesis that interpolation is less accurate away from measured k-space positions is correct for these examples.

5.3.2 SIM Phantom Reconstruction

Fig. 5.3 shows the NMRSE of the SHARK reconstruction for an acceleration factor of 2. At first, the NMRSE decreases when more l_1 regularization is applied. At a certain point, further increasing the regularization has the opposite effect, an increase of the NMRSE. This suggests the existence of an intermediate regularization parameter value, which is an adequate compromise between l_1 Total Variation minimization and l_2 data approximation.

The effect of the different values of the TV regularization parameter on the reconstruction can be seen in Fig. 5.4, where the reconstruction of the SIM data set for the random grid with acceleration factor 4 and three different values of the parameter is shown. In the case of the lowest value, imposing a more precise approximation to the data results results in the presence of aliasing artifacts. In the case of the highest value, too much TV regularization results in a loss of sharpness and detail. Intermediate values suppress artifacts and result in images with a good level of detail.




Figure 5.2: Normalised square root error of the dSENSE reconstruction at each k-space positions for three grids with undersampling factors 4 (with 12 ACS lines) and 8 (with 14 and 28 ACS lines). Red crosses indicate sampled k-space locations.




Figure 5.3: NMRSE for different values of the Total Variation regularization parameter in the SHARK reconstruction of the SIM data after reinterpolation on a random grid with undersampling factor 2.




Figure 5.4: SHARK reconstructions of the SIM data for TV regularization parameter values of 100, 10^5 and 10^{11} (acceleration factor of 4).

5.3.3 Phantom Data

All the results from the reconstruction of the PH data set subsampled with three different regular subsampling schemes and resampled onto the designed grids described in A.2 can be found in App. B.5. Fig. 5.5 shows a selected reconstruction for the arbitrary grid and the grid which combines the arbitrary and regular schemes, compared to dSENSE and regularized SENSE reconstructions. For an acceleration factor of 4 with 12 ACS lines and an acceleration factor of 8 with 28 lines, the SHARK reconstructions are superior to the other alternatives in terms of the suppression of aliasing artifacts and noise amplification. The best quality reconstructions are obtained for the resampling grid that combines the arbitrary and regular sampling schemes.

5.4 Discussion

In this chapter, the Sparse incoherent Arbitrary Resampling in K-space (SHARK) algorithm was presented. This algorithm uses sensitivity encoding to facilitate the exploitation of image sparsity by a non-linear Compressive Sampling optimization scheme. As opposed to other algorithms which also employ Compressive Sampling techniques, SHARK can be applied to regular Cartesian input data.

The results of the first experiments were promising, but not completely conclusive. In the mentioned examples, the algorithm does in fact improve the reconstruction quality with respect to a complete dSENSE reconstruction. As a result, it seems plausible that performing l_1 minimisation in the chosen synthesis domains to limit noise amplification could be an interesting alternative to the l_2 based approach on which Tikhonov regularization is based. Further work will include more

detailed tests to determine the denoising performance of the algorithm and its application to k-space trajectories other than Cartesian.




Figure 5.5: Different reconstructions for data originally sampled at an acceleration factor of 4 with 12 ACS lines (top row), 8 with 14 lines (middle row) and 8 with 28 lines (bottom row). The column at the left shows regularized SENSE reconstructions, the second column shows dSENSE reconstructions, the third row shows SHARK reconstruction for the resampling grid generated by the Monte Carlo procedure and the column at the right shows SHARK reconstructions for the combination of the resampling and regular grids.

Conclusion

In general, nonlinear parallel MRI reconstruction models are more complex than their linear alternatives. This makes it a challenge to define corresponding reconstruction problems which are well-posed and computationally tractable, but at the same time it allows to address the main limitations of linear approaches by exploiting *a priori* information about the acquisition process and the desired solutions in a more flexible way. The results presented in this work suggest that the application of nonlinear techniques could provide significant improvements in parallel MR image reconstruction. On the one hand, the joint coil sensitivity and image estimation model can improve the explicit sensitivity estimate of the SENSE framework, thereby yielding more accurate image reconstructions. On the other hand, Compressive Sampling can be effectively applied to parallel MR data, using coil sensitivity information either to adapt the Compressive Sampling model or to resample the data onto a grid especially adapted for sparse reconstruction. In both cases, noise amplification is reduced and good reconstruction quality is achieved for high acceleration factors in several examples.

There are two main challenges to address in the further development of the nonlinear reconstruction techniques presented here. On the one hand, it would be very useful to analyze the characteristics of both images and coil sensitivities in more detail, so that they can be incorporated in the reconstruction models in a more effective way. This includes, for instance, the problem of finding a transform domain in which the sparsity of coil sensitivities can be used to yield more accurate and robust estimates through l_1 norm minimisation. On the other hand, methods that solve the optimization problems arising from those models in a more efficient way could be developed. Examples of such methods would include alternatives to the *greedy* algorithm of Chap. 3 and to the empirical calibration of the l_1 norm regularization parameters in Chaps. 4 and 5.

Bibliography

- [1] JPEG 2000 image coding system, ISO/IEC International Standard 15444-1. ITU Recommendation T.800, 2000.
- [2] M. Blaimer, M. Mueller, M. A. Griswold, P. M. Jakob, and R. M. Heidemann. SMASH, SENSE, PILS, GRAPPA: how to choose the optimal method. *Top Magn Reson Imaging*, 15(4):223–236, 2004.
- [3] E. Candès. Ridgelets and the representation of mutilated Sobolev functions. *SIAM J. Math. Anal.*, 33:347–368, 1999.
- [4] E. J. Candès and T. Tao. Near-optimal signal recovery from random projections: universal encoding strategies. *IEEE Trans. Inform. Theory*, 52:5406–5425, 2004.
- [5] D. Donoho, A. Maleki, and M. Shahram. WAVELAB www-stat.stanford.edu/~wavelab/.
- [6] M. A. Griswold and P. M. Jakob et al. Partially parallel imaging with localized sensitivities (PILS). *Magn Reson Med.*, 44:602–609, 2000.
- [7] M. A. Griswold, P. M. Jakob, and R. M. Heidemann et al. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn Reson Med.*, 47:120210, 2002.
- [8] M. A. Griswold, M. Blaimer, P. M. Jakob, and R. M. Heidemann et al. Parallel magnetic resonance imaging using the GRAPPA operator formalism. *Magn Reson Med.*, 54:1553–1556, 2005.
- [9] J. D. Jackson. *Classical Electrodynamics*. John Wiley and Sons, 1998.
- [10] S.J. Kim, K. Koh, M. Lustig, S. Boyd, and D. Gorinevsky. A Method for Large-Scale l_1 -Regularized Least Squares. *IEEE Trans. Inform. Theory*, 52:5406–5425, 2004.
- [11] K. F. King. Combining Compressed Sensing and Parallel Imaging. In *Proc. Intl. Soc. Mag. Reson. Med. 16 (2008)*.
- [12] Z. Liang, R. Bammer, J. Ji, N. J. Pelc, and G. H. Glover. Improved image reconstruction from sensitivity-encoded data by wavelet denoising and Tikhonov regularization. In *2002 IEEE International Symposium on Biomedical Imaging*.

BIBLIOGRAPHY

- [13] Z. P. Liang and P. C. Lauterbur. *Principles of Magnetic Resonance Imaging*. IEEE Press, 2000.
- [14] F. H. Lin, K. K. Kwong, J. W. Belliveau, and L. L. Wald. Parallel Imaging Reconstruction Using Automatic Regularization. *Magn Reson Med.*, 51:559 – 567, 2002.
- [15] B. Liu, L. Ying, M. Steckner, J. Xie, and J. Sheng. Regularized SENSE Reconstruction using Iteratively Refined Total Variation Method. In *2007 IEEE International Symposium on Biomedical Imaging*.
- [16] C. Liu, R. Bammer, and M. E. Moseley. Parallel imaging reconstruction for arbitrary trajectories using k-space sparse matrices (kSPA). *Magn Reson Med.*, 58:1171 – 1181, 2007.
- [17] M. Lustig. Sparse MRI, www.stanford.edu/~mlustig/.
- [18] M. Lustig, D. L. Donoho, and J. M. Pauly. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magn Reson Med.*, 58:1182–95, 2007.
- [19] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly. k-t SPARSE: High frame rate dynamic MRI exploiting spatio-temporal sparsity. In *Proc. Intl. Soc. Mag. Reson. Med. 13 (2006)*.
- [20] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly. Compressed Sensing MRI. *Signal Processing Magazine, IEEE*, 25:72–82, 2008.
- [21] S. Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, 1999.
- [22] C. M. McKenzie, E. N. Yeh, and M. A. Ohliger et al. Self-Calibrating Parallel Imaging with Automatic Coil Sensitivity Extraction. *Magn Reson Med.*, 47:529538, 2002.
- [23] W. H. Press, P. C. Vetterling, S. A. Teukolsky, and B. P. Flannery. *Numerical Recipes in C*. Cambridge University Press, 1999.
- [24] K. P. Pruessmann, M. Weiger, P. Börnert, and P. Boesiger. Advances in Sensitivity Encoding With Arbitrary k-Space Trajectories. *Magn Reson Med.*, 46:638–651, 2001.
- [25] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger. SENSE: Sensitivity Encoding for Fast MRI. *Magn Reson Med.*, 42:952962, 1999.
- [26] N. Seiberlich, F. A. Breuer, R. Heidemann, M. Blaimer, M. A. Griswold, and P. Jakob. Reconstruction of undersampled non-Cartesian data sets using pseudo-Cartesian GRAPPA in conjunction with GROG. *Magn Reson Med.*, 59:1127 – 1137, 2008.
- [27] N. Seiberlich, P. Ehses, F. A. Breuer, M. Blaimer, P. M. Jakob, and M. A. Griswold. Reconstruction of Undersampled Non-Cartesian Data using GROG-Facilitated Random Blipped Phase Encoding. In *Proc. Intl. Soc. Mag. Reson. Med. 16 (2008)*.

BIBLIOGRAPHY

- [28] J. Senegas, T. Knopp, and H. Eggers. dSENSE: Direct k-space Reconstruction for Non-Cartesian Parallel MRI. In *Proc. Intl. Soc. Mag. Reson. Med. 13 (2006)*.
- [29] M. Uecker, T. Hohage, K. T. Block, and J. Frahm. Image reconstruction by regularized nonlinear inversion—joint estimation of coil sensitivities and image content. *Magn Reson Med.*, 60:674–82, 2008.
- [30] G. K. Wallace. The JPEG still picture compression standard. *Communications of the ACM*, 34:30–44, 1991.
- [31] B. Wu, R. P. Millane, R. Watts, and P. Bones. Applying Compressed Sensing in Parallel MRI. In *Proc. Intl. Soc. Mag. Reson. Med. 16 (2008)*.
- [32] L. Ying and J. Sheng. A Variable Projection Method to JSENSE. In *Proc. Intl. Soc. Mag. Reson. Med. 16 (2008)*.
- [33] L. Ying and J. Sheng. Joint Image and Sensitivity Estimation in SENSE (JSENSE). *Magn Reson Med.*, 57:1196–1202, 2007.
- [34] C. Zhao, T. Lang, and J. Ji. Compressed Sensing Parallel Imaging. In *Proc. Intl. Soc. Mag. Reson. Med. 16 (2008)*.

Appendix A

K-Space Sampling Grid Design

A.1 Phase-Encoding Subsampling

Three grids were designed for the tests in Chap. 4. In the first grid, parameters p and r in equation 4.2 took the values 5 and 20% (r is expressed as a percentage of k-space length). For the second and third grids, p was 5 and 7.5, and r the values 20% and 1.5%. The sampling probability of each phase-encoding line for each of these grids is shown in Fig. A.1.




Figure A.1: Sampling probability for each phase encoding lines in three grid design schemes which aim to generate grids with acceleration factors of 1.9 (left), 3.5 (center) and 7.1 (right)).

The three grids obtained in this way are displayed in Figs. A.2, A.3 and A.4. In each case the grids are compared to two other grids of similar acceleration factor, obtained respectively with regular and uniform random subsampling. The comparison is based on the Transform Point Spread Function shape and secondary lobe level, and on the aliasing artifacts generated on the phantom of the SIM data set by performing the following transformation: $I_{\text{aliased}} = (P_g F)^H P_g F I$. I is the phantom image and P_g the projection matrix for the grid. In every case, we can see how the phase-encoding pseudo-random grid is a compromise between the coherence of the regular grid and the incoherence of the totally random grid.




Figure A.2: Three alternative grid designs (top row, white points indicate samples), with their respective TPSF for a fine-scale Daubechies 4 wavelet coefficient (middle row) and the aliasing pattern they produce on a phantom image (bottom row). In the graphs, the TPSF maximum is marked in red and a green line indicates the maximum secondary lobe value. For the regular grid with ACS lines, the acceleration factor is 1.94. For the other two grids the acceleration factor is 1.92.




Figure A.3: Three alternative grid designs (top row, white points indicate samples), with their respective TPSF for a fine-scale Daubechies 4 wavelet coefficient (middle row) and the aliasing pattern they produce on a phantom image (bottom row). In the graphs, the TPSF maximum is marked in red and a green line indicates the maximum secondary lobe value. For the regular grid with ACS lines, the acceleration factor is 3.36. For the other two grids the acceleration factor is 3.51.




Figure A.4: Three alternative grid designs (top row, white points indicate samples), with their respective TPSF for a fine-scale Daubechies 4 wavelet coefficient (middle row) and the aliasing pattern they produce on a phantom image (bottom row). In the graphs, the TPSF maximum is marked in red and a green line indicates the maximum secondary lobe value. For the regular grid with ACS lines, the acceleration factor is 5.56. For the other two grids the acceleration factor is 7.11.

A.2 SHARK Subsampling

Three grids were designed for the application of SHARK, following the method described in Secs. 5.1.3 and 5.2.1. They were respectively based on regular grids with acceleration factor 4 and 12 ACS lines, with acceleration factor 8 and 14 and with acceleration factor 8 and 28 ACS lines. The sampling probability of each k-space position for each of these grids is shown in Fig. A.5.




Figure A.5: Sampling probability for each pixel in three sampling schemes based on regular grids with an acceleration factor of 4 and 12 ACS lines (left), an acceleration factor of 8 and 14 ACS lines (center) and an acceleration factor of 8 and 28 ACS lines (right). Probability values are represented by shades of grey ranging from black, which indicates zero probability, to white, which indicates probability 1.

The three grids obtained by the Monte Carlo design procedure are displayed in Figs. A.6, A.7 and A.8. For each designed grid, another grid was generated by combining its samples with those of the regular grid used to generate it. In the three cases, both grids are compared to the regular grid used to generate them in terms of the Transform Point Spread Function shape and secondary lobe level, and in terms of the aliasing artifacts generated on the phantom of the SIM data set as explained in the previous section.

In general, the incoherence is more limited than in the case of random sampling in the phase-encoding direction, and of course than in the case of totally random sampling. In all three examples, both resampling grids are better in terms of incoherence than the regular scheme and the arbitrary grid is in turn more incoherent than the combination of both.




Figure A.6: Regular grid with an acceleration factor of 4 and 12 ACS lines (top left), compared to the designed grid based on the same undersampling pattern (top center and right). The comparison is based on their respective TPSF for a fine-scale Daubechies 4 wavelet coefficient (middle row) and the aliasing pattern they produce on a phantom image (bottom row). In the graphs, the TPSF maximum is marked in red and a green line indicates the maximum secondary lobe value.




Figure A.7: Regular grid with an acceleration factor of 8 and 14 ACS lines (top left), compared to the designed grid based on the same undersampling pattern (top center and right). The comparison is based on their respective TPSF for a fine-scale Daubechies 4 wavelet coefficient (middle row) and the aliasing pattern they produce on a phantom image (bottom row). In the graphs, the TPSF maximum is marked in red and a green line indicates the maximum secondary lobe value.




Figure A.8: Regular grid with an acceleration factor of 8 and 28 ACS lines (top left), compared to the designed grid based on the same undersampling pattern (top center and right). The comparison is based on their respective TPSF for a fine-scale Daubechies 4 wavelet coefficient (middle row) and the aliasing pattern they produce on a phantom image (bottom row). In the graphs, the TPSF maximum is marked in red and a green line indicates the maximum secondary lobe value.

Appendix B

Compressive Sampling Regularization Parameters

In this chapter, reconstruction results for the different Compressive Sampling algorithms presented in previous chapters are shown for different values of wavelet and TV l_1 regularization parameters. In general, the effect of the different regularization terms is similar for the different data sets and subsampling schemes.

When both regularization parameters have small values, aliasing artifacts are present in the reconstructed images. When the TV regularization parameter value is large, the solution is forced to resemble a piecewise constant function, suppressing some of the details in the image. This is more obvious in the case of the BR data set, since both the phantoms of the SIM and PH images are much better approximated by piecewise constant functions. When the wavelet regularization parameter value is large, the signal level in some areas of the image is strongly amplified. When both parameters are large, intermediate effects are noticeable. The best quality images are obtained for parameter values which are an adequate compromise between data approximation and sparsity in the two transform domains.

Sec. B.1 shows the results of the Compressive Sampling SoS reconstruction (see Sec. 4.1.3) for the three data sets at the different acceleration factors. In general, all of the results present inhomogeneities caused by the SoS reconstruction. Sec. B.2 shows CSENSE results for noiseless data and Sec. B.3 for noisy data. Sec. B.4 shows sensitivity estimation results. Finally, Sec. B.5 shows image reconstructions obtained with SHARK.

B.1 Compressive Sampling SoS Reconstruction




Figure B.1: Compressive Sampling sum of squares reconstruction of the SIM data set subsampled at a reduction factor of 1.92 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 1 (third row) and 10^2 (bottom row) times λ . In the first row the original image used to generate the data set is shown for comparison.




Figure B.2: Compressive Sampling sum of squares reconstruction of the SIM data set subsampled at a reduction factor of 3.51 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 1 (third row) and 10^2 (bottom row) times λ . In the first row the original image used to generate the data set is shown for comparison.




Figure B.3: Compressive Sampling sum of squares reconstruction of the SIM data set subsampled at a reduction factor of 7.11 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 1 (third row) and 10^2 (bottom row) times λ . In the first row the original image used to generate the data set is shown for comparison.




Figure B.4: Compressive Sampling sum of squares reconstruction of the BR data set subsampled at a reduction factor of 1.92 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 1 (third row) and 10^2 (bottom row) times λ . In the first row a CLEAR reconstruction is shown for comparison.




Figure B.5: Compressive Sampling sum of squares reconstruction of the BR data set subsampled at a reduction factor of 3.51 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 1 (third row) and 10^2 (bottom row) times λ . In the first row a CLEAR reconstruction is shown for comparison.




Figure B.6: Compressive Sampling sum of squares reconstruction of the BR data set subsampled at a reduction factor of 7.11 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 1 (third row) and 10^2 (bottom row) times λ . In the first row a CLEAR reconstruction is shown for comparison.




Figure B.7: Compressive Sampling sum of squares reconstruction of the PH data set subsampled at a reduction factor of 1.92 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 1 (third row) and 10^2 (bottom row) times λ . In the first row a CLEAR reconstruction is shown for comparison.




Figure B.8: Compressive Sampling sum of squares reconstruction of the PH data set subsampled at a reduction factor of 3.51 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 1 (third row) and 10^2 (bottom row) times λ . In the first row a CLEAR reconstruction is shown for comparison.




Figure B.9: Compressive Sampling sum of squares reconstruction of the PH data set subsampled at a reduction factor of 7.11 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 1 (third row) and 10^2 (bottom row) times λ . In the first row a CLEAR reconstruction is shown for comparison.

B.2 CSENSE Image Reconstruction




Figure B.10: CSENSE reconstruction of the SIM data set subsampled at a reduction factor of 1.92 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 1 (third row) and 10^2 (bottom row) times λ . In the first row the original image used to generate the data set is shown for comparison.




Figure B.11: CSENSE reconstruction of the SIM data set subsampled at a reduction factor of 3.51 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 10^{-2} (third row), 1 (fourth row) and 10^2 (bottom row) times λ . In the first row the original image used to generate the data set is shown for comparison.




Figure B.12: CSENSE reconstruction of the SIM data set subsampled at a reduction factor of 7.11 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 10^{-2} (third row), 1 (fourth row) and 10^2 (bottom row) times λ . In the first row the original image used to generate the data set is shown for comparison.




Figure B.13: CSENSE reconstruction based on an erroneous sensitivity function estimate of the SIM data set subsampled at a reduction factor of 1.92 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 1 (third row) and 10^2 (bottom row) times λ . In the first row the original image used to generate the data set is shown for comparison.




Figure B.14: CSENSE reconstruction based on an erroneous sensitivity function estimate of the SIM data set subsampled at a reduction factor of 3.51 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 1 (third row) and 10^2 (bottom row) times λ . In the first row the original image used to generate the data set is shown for comparison.




Figure B.15: CSENSE reconstruction based on an erroneous sensitivity function estimate of the SIM data set subsampled at a reduction factor of 7.11 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 1 (third row) and 10^2 (bottom row) times λ . In the first row the original image used to generate the data set is shown for comparison.




Figure B.16: CSENSE reconstruction of the BR data set subsampled at a reduction factor of 1.92 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 1 (third row) and 10^2 (bottom row) times λ . In the first row a CLEAR reconstruction is shown for comparison.




Figure B.17: CSENSE reconstruction of the BR data set subsampled at a reduction factor of 3.51 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 10^{-2} (third row), 1 (fourth row) and 10^2 (bottom row) times λ . In the first row a CLEAR reconstruction is shown for comparison.




Figure B.18: CSENSE reconstruction of the BR data set subsampled at a reduction factor of 7.11 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 10^{-2} (third row), 1 (fourth row) and 10^2 (bottom row) times λ . In the first row a CLEAR reconstruction is shown for comparison.




Figure B.19: CSENSE reconstruction of the PH data set subsampled at a reduction factor of 1.92 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 1 (third row) and 10^2 (bottom row) times λ . In the first row a CLEAR reconstruction is shown for comparison.




Figure B.20: CSENSE reconstruction of the PH data set subsampled at a reduction factor of 3.51 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 10^{-2} (third row), 1 (fourth row) and 10^2 (bottom row) times λ . In the first row a CLEAR reconstruction is shown for comparison.




Figure B.21: CSENSE reconstruction of the PH data set subsampled at a reduction factor of 7.11 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 10^{-2} (third row), 1 (fourth row) and 10^2 (bottom row) times λ . In the first row a CLEAR reconstruction is shown for comparison.

B.3 CSENSE Noisy Data Reconstruction




Figure B.22: CSENSE reconstruction of the SIM data set with added Gaussian noise (k-space SNR of 21.11 dB) subsampled at a reduction factor of 3.51 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (first row), 1 (second row) and 10^2 (bottom row) times λ .




Figure B.23: CSENSE reconstruction of the SIM data set with added Gaussian noise (k-space SNR of 21.11 dB) subsampled at a reduction factor of 7.11 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (first row), 1 (second row) and 10^2 (bottom row) times λ .




Figure B.24: CSENSE reconstruction of the SIM data set with added Gaussian noise (k-space SNR of 11.11 dB) subsampled at a reduction factor of 3.51 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (first row), 1 (second row) and 10^2 (bottom row) times λ .




Figure B.25: CSENSE reconstruction of the SIM data set with added Gaussian noise (k-space SNR of 11.11 dB) subsampled at a reduction factor of 7.11 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (first row), 1 (second row) and 10^2 (bottom row) times λ .




Figure B.26: CSENSE reconstruction of the SIM data set with added Gaussian noise (k-space SNR of 1.11 dB) subsampled at a reduction factor of 3.51 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (first row), 1 (second row) and 10^2 (bottom row) times λ .




Figure B.27: CSENSE reconstruction of the SIM data set with added Gaussian noise (k-space SNR of 1.11 dB) subsampled at a reduction factor of 7.11 for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (first row), 1 (second row) and 10^2 (bottom row) times λ .

B.4 Coil Sensitivity Reconstruction




Figure B.28: Original simulated sensitivity of the SIM data set (top left), low resolution scan estimates for the BR (center left) and PH (bottom left). In the right column full sampled coil images for the three data sets are shown. These images serve as references to which the next figures can be compared to.




Figure B.29: Compressive Sampling sensitivity reconstructions for the SIM data set (left column) and their product with the reference image (right column) for a sampling grid with an acceleration factor of 3.51. The wavelet l_1 regularization parameter values were: 10^5 (top row), $5 \cdot 10^7$ (second row), $5 \cdot 10^9$ (third row) and 10^{12} (bottom row.)




Figure B.30: Compressive Sampling sensitivity reconstructions for the SIM data set (left column) and their product with the reference image (right column) for a sampling grid with an acceleration factor of 7.11. The wavelet l_1 regularization parameter values were: 10^5 (top row), $5 \cdot 10^7$ (second row), $5 \cdot 10^9$ (third row) and 10^{12} (bottom row).




Figure B.31: Compressive Sampling sensitivity reconstructions for the BR data set (left column) and their product with the reference image (right column) for a sampling grid with an acceleration factor of 3.51. The wavelet l_1 regularization parameter values were: 0.1 (top row), 2 (second row), 50 (third row) and 10^3 (bottom row).




Figure B.32: Compressive Sampling sensitivity reconstructions for the BR data set (left column) and their product with the reference image (right column) for a sampling grid with an acceleration factor of 7.11. The wavelet l_1 regularization parameter values were: 0.1 (top row), 2 (second row), 50 (third row) and 10^3 (bottom row).




Figure B.33: Compressive Sampling sensitivity reconstructions for the PH data set (left column) and their product with the reference image (right column) for a sampling grid with an acceleration factor of 3.51. The wavelet l_1 regularization parameter values were: 0.1 (top row), 2 (second row), 50 (third row) and 10^3 (bottom row).




Figure B.34: Compressive Sampling sensitivity reconstructions for the PH data set (left column) and their product with the reference image (right column) for a sampling grid with an acceleration factor of 7.11. The wavelet l_1 regularization parameter values were: 0.1 (top row), 2 (second row), 50 (third row) and 10^3 (bottom row).

B.5 SHARK Reconstruction

CLEAR Image





Figure B.35: SHARK reconstruction of the PH data set subsampled at a reduction factor of 4 with 12 ACS lines and resampled onto a variable density grid designed by a Monte Carlo procedure (center column in Fig. A.6). The images show reconstructions for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 1 (third row) and 10^2 (bottom) times λ . In the first row the CLEAR reconstruction is shown for comparison.




Figure B.36: SHARK reconstruction of the PH data set subsampled at a reduction factor of 4 with 12 ACS lines and resampled onto a grid generated as a combination of a variable density grid designed by a Monte Carlo procedure and the regular grid (right column in Fig. A.6). The images show reconstructions for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 1 (third row) and 10^2 (bottom) times λ . In the first row the CLEAR reconstruction is shown for comparison.




Figure B.37: SHARK reconstruction of the PH data set subsampled at a reduction factor of 8 with 14 ACS lines and resampled onto a variable density grid designed by a Monte Carlos procedure (center column in Fig. A.7). The images show reconstructions for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 1 (third row) and 10^2 (bottom) times λ . In the first row the CLEAR reconstruction is shown for comparison.




Figure B.38: SHARK reconstruction of the PH data set subsampled at a reduction factor of 8 with 14 ACS lines and resampled onto a grid generated as a combination of a variable density grid designed by a Monte Carlo procedure and the regular grid (right column in Fig. A.7). The images show reconstructions for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 1 (third row) and 10^2 (bottom) times λ . In the first row the CLEAR reconstruction is shown for comparison.




Figure B.39: SHARK reconstruction of the PH data set subsampled at a reduction factor of 8 with 28 ACS lines and resampled onto a variable density grid designed by a Monte Carlos procedure (center column in Fig. A.8). The images show reconstructions for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 1 (third row) and 10^2 (bottom) times λ . In the first row the CLEAR reconstruction is shown for comparison.




Figure B.40: SHARK reconstruction of the PH data set subsampled at a reduction factor of 8 with 28 ACS lines and resampled onto a grid generated as a combination of a variable density grid designed by a Monte Carlo procedure and the regular grid (right column in Fig. A.8). The images show reconstructions for several values of the wavelet regularization parameter λ and the total variation parameter: 0 (second row), 1 (third row) and 10^2 (bottom) times λ . In the first row the CLEAR reconstruction is shown for comparison.