Universidad de Valladolid

Escuela Técnica Superior de Ingenieros de Telecomunicación

Dpto. de Teoría de la Señal y Comunicaciones e Ingeniería Telemática

Tesis Doctoral

Artifact reduction in Magnetic Resonance Imaging: noise modelling in 2D/3D GRAPPA accelerated acquisitions and motion–induced ghosting correction in multishot diffusion MRI

Presentada por **D. Iñaki Rabanillo Viloria** para optar al grado de Doctor por la Universidad de Valladolid

Dirigida por:

Santiago Aja Fernández y Carlos Alberola López

Abril 2019

Valladolid, España

He may look like an idiot and talk like an idiot but don't let this dissertation fool you. He really is an idiot.

Inspired by Groucho Marx

ACKNOWLEDGEMENTS

I would like to acknowledge all the people who have helped me throughout the process of completing this dissertation. First of all, I am deeply thankful to my advisors Carlos Alberola and Santiago Aja for giving me the opportunity to discover the world of research. I have nothing but gratitude for their kindness, as well as for their flexibility to discuss my ideas and the freedom they have given me to pursue them. Their advice and support have been key to the achievements of this dissertation.

I must also mention the different advisors I have had during my abroad experiences. I owe a debt of gratitude to Dr. Diego Hernando, whose brilliance, enthusiasm and charisma made him a role model I look up to. His guidance has been fundamental to my research and I am proud to say he has become a friend along the way. I want to thank Dr. Rachid Deriche for his gentleness welcoming me at his fantastic lab, where not only he has created a cutting edge research laboratory, but also an amazing group of people who made my time there so memorable. Finally, I would like to thank Dr. Lucilio Cordero, the person who fisrt introduced me to research when I was still an undergraduate student and who helped me conclude this thesis. He is one of the finest researchers I have met and I feel very fortunate to have been able to work and learn from him.

Last, I want to also thank everybody at the Image Processing Laboratory for making it a great and fun place to work. And of course, I want to express my love to my parents and my close friends. They all know who they are and how much I appreciate them. This thesis is dedicated to them.

ABSTRACT

Magnetic Resonance Imaging (MRI) is a powerful diagnostic imaging modality for numerous diseases due to its versatility and sensitivity to multiple tissue properties. Nevertheless, it is often limited by the lengthy scan times required to collect the data necessary to form an image and it suffers from different types or artifacts arising from multiple causes. For this reason, reducing the scanning time has been one of the most active areas of research, but it might come at the cost of aggravating the effects of certain artifacts or introducing new ones.

One way to reduce the acquisition time is by using Parallel Imaging techniques, which acquire only a portion of the data and rely on the availability of different coil channels to further reconstruct the images. However, these techniques require more complex reconstruction algorithms that result in the appearance of spatially varying noise maps. In order to mitigate the impact of noise degradation in subsequent parameter estimation it is important to characterize these noise maps. In particular, its exact characterization has been considered computationally infeasible under a widely used technique termed GRAPPA, which directly reconstructs the missing data in the sampled domain, the so called $\boldsymbol{k}$–space. The reason lies on the need to carry out a noise propagation analysis through the reconstruction pipeline that involves very large covariance matrices. In this thesis, we show how to overcome this computational load and obtain an exact noise characterization both for 2D and 3D GRAPPA acquisitions by exploiting the presence of extensive symmetries and the block separability in the reconstruction steps.

Another common approach to reduce the scan time is by means of Echo–Planar Imaging (EPI). In contrast to Spin–Warp Imaging, where one acquires one line of the $\boldsymbol{k}$–space per excitation, EPI segments the acquisition into multiple shots by collecting several lines within a single excitation. This modality offers major advantages over conventional Spin–Warp Imaging, which include reduced imaging time, decreased motion artifacts and the ability to image rapid physiologic processes of the human body. In particular, it has become the standard modality in Diffusion MRI (dMRI). However, since dMRI is aimed at capturing the microscopic movements of water molecules, it is sensitive as well to any kind of bulk motion from the patient. Due to the way dMRI sequences are designed, the molecules motion is encoded in the phase of the spins and consequently bulk motion results in phase corruption of

the images. If the motion differs from shot to shot, the resulting phase discrepancies lead into ghosting artifacts in the reconstructed images. In this thesis, we propose an algorithm based on a Maximum Likelihood formulation to iteratively reconstruct the images and estimate the phase–maps under the assumption of linearity or smoothness.

In this dissertation we include the theoretical derivation of our models and the description of the proposed algorithms to determine the parameters of interest. Finally, simulations, phantom and *in–vivo* experiments are included to provide empirical support of the properties of our methods, as well as to compare them to previous state of the art approaches.

RESUMEN

La Imagen por Resonancia Magnética (IRM) es una modalidad de imagen de gran importancia para el diagnóstico de numerosas enfermedades gracias a su versatilidad y sensibilidad a múltiples propiedades de los tejidos. Sin embargo, se ve habitualmente limitada por los largos tiempos requeridos para muestrear los datos necesarios para la formación de la imagen, así como por los distintos tipos de artefactos originados por diferentes causas. Por este motivo, reducir los tiempos de adquisición es uno de los principales objetivos de la investigación actual, pero su consecuención suele implicar la amplificación de los efectos indeseados de cirtos artefactos o la aparición de nuevos tipos de artefactos.

Una de las maneras para disminuir el tiempo de adquisición es mediante el uso de técnicas de Imagen en Paralelo, las cuales adquieren una fracción de los datos requeridos y aprovechan la disponibilidad de múltiples antenas para reconstruir las imágenes. No obstante, estas técnicas requieren algoritmos de reconstrucción más complejos que se traducen en la aparición de mapas de ruido no uniformes espacialmente. Para mitigar el efecto de la contaminación de ruido en la subsiguiente estimación de parámetros, es importante disponer de una caracterización de dichos mapas de ruido. En particular, su caracterización exacta se consideraba inviable computacionalmente para una de las técnicas más utilizadas, conocida como GRAPPA. La explicación reside en el tipo de reconstrucción realizada, que se basa en estimar los datos no muestreados directamente en el dominio de adquisición, el denominado $\boldsymbol{k}$–espacio. Sin embargo, estos datos deben ser transformados para obtener la imagen anatómica de interés, lo que obliga a llevar a cabo un análisis de la propagación del ruido a lo largo de todo el procedimiento de reconstrucción, involucrando matrices de covarianza muy grandes. En esta tesis se muestra cómo superar la limitación dada por la carga computacional de cara a obtener una caracterización exacta de los mapas de ruido para adquisiciones GRAPPA 2D/3D explotando la presencia de numerosas simetrías en las matrices de covarianza y la separabilidad en bloques de las mismas a lo largo de las distintas etapas de la reconstruccción.

Una alternativa utilizada frecuentemente para reducir el tiempo de adquisición es el uso de secuencias de Imagen Eco–Planar (EPI). A diferencia de la Imagen estándar por *"Spin–Warp"*, que se basa en la adquisición de una línea del $\boldsymbol{k}$–espacio por cada excitación, EPI segmenta la adquisición en múltiples disparos muestreando varias líneas en cada excitación. Gracias a

ello, EPI ofrece algunas ventajas notables respecto a *"Spin–Warp"*, tales como la reducción del tiempo de muestreo, la disminución de los artefactos relacionados con el movimiento del paciente o la capacidad para hacer imagen de procesos fisiológicos que ocurren en intervalos de tiempo muy breves. En concreto, EPI se ha convertido en la modalidad estándar en Difusión por RMI (DRMI). No obstante, dado que DMRI busca captar el movimiento microscópico de las moléculas de agua, es también sensible a cualquier tipo de movimiento del paciente. Debido al modo en que las secuencias de DMRI están diseñadas, el movimiento de las moléculas de agua se codifica en la fase de los espines y, consecuentemente, el movimiento del paciente se traduce en una corrupción de la fase de las imágenes. Concretamente, si el movimiento difiere de un disparo a otro, las discrepancias de fase resultantes darán lugar a la aparición de artefactos de solapamiento en las imágenes reconstruidas. En esta tesis se propone un algoritmo basado en una formulación de Máxima Verosimilitud para, de manera iterativa, reconstruir las imágenes y estimar los mapas de fase bajo la asunción de linearidad o suavidad.

En esta tesis se incluyen los fundamentos teóricos para la obtención de nuestros modelos así como la descripción de los algoritmos propuestos para determinar los parámetros de interés. Finalmente, se incluyen los experimentos realizados por medio de simulaciones, phantoms y adquisiciones *in–vivo* con el objetivo de proporcionar respaldo empírico de las propiedades de los métodos presentados, así como para compararlos frente a alternativas presentes en el estado del arte.

LIST OF ACRONYMS

ACS Auto Calibration Signal

CT Computarized Tomography

DFT Discrete Fourier Transform

dMRI Diffusion MRI

EPI Echo–Planar Imaging

FE Frequency–Encoding

FOV Field Of View

FT Fourier Transform

GRAPPA GeneRalized Autocalibrating Partial Parallel Acquisition

iFT Inverse Fourier Transform

IID Independent and Identically Distributed

LORAKS Low Rank Modeling of $\boldsymbol{k}$–space Neighborhoods

MB Multi–Band

MRI Magnetic Resonance Imaging

ms Multi–Shot

MUSE MUltiplexed Sensitivity Encoding

MUSSELS MUlti-Shot Sensitivity Encoded Diffusion Data Recovery using Structured Low Rank Matrix Completion

PE Phase–Encoding

PI Parallel Imaging

PROPELLER Periodically Rotated Overlapping Parallel Lines with Enhanced Reconstruction

RF Radio–Frequency

rFOV Reduced Field Of View

SENSE SENSitivity Encoding

SMS Simultaneous Multi–Slice

SNAILS Self-NAvigated InterLeaved Spiral

SNR Signal to Noise Ratio

SoS Sum of Squares

SPIRiT Iterative Self-consistent Parallel Imaging Reconstruction from Arbitrary $\boldsymbol{k}$–space

ss Single–Shot

TE Echo Time

TR Repetition Time

VD-GRAPPA Variable Density GRAPPA

TABLE OF CONTENTS

CHAPTER

ONE

INTRODUCTION

1.1 Motivation

Magnetic Resonance Imaging (MRI) is a discipline widely used in clinical practice. It is a powerful imaging modility because of its versatility and sensitivity to a broad range of tissue properties (Brown et al., 2014; Bernstein et al., 2004). Furthermore, a main thrust comes from safety since it does not use ionizing radiation, rendering it a non–invasive technique. However, the process of creating MR images results in the appearance of artifacts in the final display. Moreover, despite the advances in hardware, it remains a relatively slow modality compared to other alternatives such as Computarized Tomography (CT), and, in fact, slowness itself is a source of image artifacts while procedures to speed up acquisitions are, as well, the cause of additional distortions. For this reason, the knowledge of the causes and effects of artifacts is vital for eliminating them and reducing the probability of an incorrect diagnose. Better understanding can lead to better modelling of the image formation, and consequently to ameliorate the quality of the images. This is the purpose in the present Thesis.

We define artifacts as a misrepresentation of the true anatomy or tissue characteristics (Henkelman, 1987). If not detected, they can introduce pseudo-lesions, degrade image quality or provoke the need to reacquire data. The observed distortion can have its origin in multiple sources related to hardware imperfections, physical/temporal constraints, patient motion or incorrect modelling of the acquisition process. Moreover, their effects might be subject to the particular imaging scenario, observing different types of degradation depending on the scanner, the imaging sequence or the reconstruction pipeline. For this reason, it seems likely that there will not be a universal methodology for correcting all artifacts in any kind of situation. Instead, a toolbox of multiple techniques developed for different scenarios seems a

more reasonable approach. This way, the choice of which tool to apply would be based on the particular circumstances of the acquisition at hand. Hereby, we provide a brief description of the main sources of artifacts in MRI:

1. **$\vec{\mathbf{B_0}}$ inhomogeneties**: the formation of the image in MR acquisitions is based on the codification of the position of the spins in their frequency and phase (see Fig.2.1.1). For that purpose, the magnet of the scanner is assumed to create a spatially homogeneous magnetic field, but this assumption is often violated. First of all, magnets are not ideal and the presence of different metal components within the scanner, the room or even the building can distort the constant magnetic field. Even though special shimming coils are used to compensate for these effects, slight variations of the $\vec{\mathbf{B_0}}$ field can still be observed within the field-of-view. Second, biologic tissues behave differently in the presence of a magnetic field, dispersing (diamagnetic materials) or concentrating (paramagnetic or ferromagnetic materials) the field. As a result of the magnetic field distortions, the oscillating frequency of the spins varies spatially. These frequency variations, in turn, result in a signal loss from T2*–dephasing and in a spatial mismapping of the spins in the excited object. One particularly challenging scenario arises from the presence of metal objects and implants, which can introduce notable distortions in the magnetic field that can lead to signal voids near them (Wiens et al., 2018).

   Related to this, we can also mention the problems associated to $\vec{\mathbf{B_1}}$ inhomogeneties. When one increases the strength of the magnet, the Larmor frequency increases as well, resulting in a reduction of the wavelegth of the Radio–Frequency (RF) pulse, also known as B1 pulse. If the wavelength becomes comparable to the sample being imaged, the intensity of the observed B1–field will vary spatially, which in turn modulates the measured signal. Furthermore, dielectric effects might appear causing the emergence of constructive and destructive interference patterns, which manifest as dark/shade areas at the center of the image (Collins et al., 2005). An approach used often to avoid these artifacts consists of using specially designed pulses, called adiabatic pulses, which are insensitive to changes in the magnitude of $\vec{\mathbf{B_1}}$ under certain conditions (Tannus and Garwood, 1997). On the other hand, a common strategy in quantitative MRI to deal with the $\vec{\mathbf{B_1}}$ inhomogeneities is to measure the $\vec{\mathbf{B_1}}$ map and further correct the intensity inhomogeneities it introduces (Sacolick et al., 2010).

2. **Chemical shift**: it refers to the small shifts in the resonant frequency of the spins due to the different chemical environments they observe (Liang and Lauterbur, 2000). Although chemical–shift artifacts can appear at the interface of any two substances having different Larmour frequencies, probably the most prevalent is the case of fat protons. The hydrogen H^1 protons of fat are nestled with long-triglycerides and covered by electron clouds which partially shield them from the external field. It is precisely

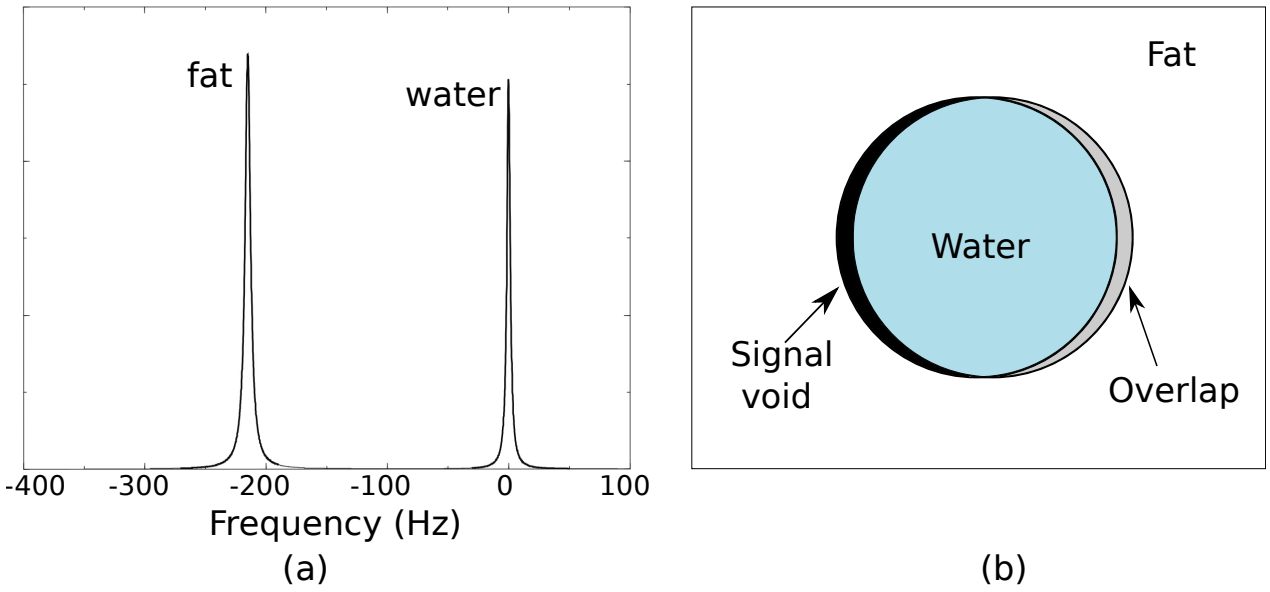

Figure 1.1: (a) Water and fat spectrum. (b) Effect of chemical shift misregistration showing the shift of the fat signal along the readout direction.

the different shielding that explains the frequency shift from fat protons to a lower frequency with respect to that of water protons. Once again, since we encode the spatial position in the oscillating frequency of the spins, this subtle difference will result in a displacement of fat protons with respect to water protons originally placed at the same location. Depending on the type of sequence, the spatial displacement could be observed in any encoding–direction. For example, for spin–warp imaging, where one line is acquired per excitation, the phase differences accumulate along the Frequency–Encoding (FE) direction. On the other hand, for Echo–Planar Imaging (EPI), where all the lines are collected after a single excitation, chemical–shift effects would be much larger along the Phase–Encoding (PE) direction. Since the water-fat displacement is inversely proportional to the bandwith, it can be reduced by increasing the second. However, this comes at the expense of reducing the Signal–to–Noise ratio (Signal to Noise Ratio (SNR)), so a compromise between noise enhancement and artifact removal needs to be reached. As an alternative, fat suppression techniques can be used, such as those based on inversion recovery pulses, where the signal originated from fat protons is nulled.

3. **Motion**: as opposed to CT or ultrasound imaging, where acquisition times are in the range of miliseconds to seconds, MR acquisition can take up to minutes. For this reason, motion is probably the main source of artifacts in MRI, resulting in different types of distortion depending on the direction of the motion and the imaging modality. Amongst the most problematic physiological motions we can cite blood flow, respiratory motion, cardiac motion or gross movements (rigid or elastics) of the body. However, some of

these artifacts might in turn be of interest and MRI can be used to provide image contrast sensitive to particle motion such as in diffusion or blood flow imaging. In order to mitigate its effects, the most basic approach consists of preventing the patient from moving by means of training, stabilization or sedation. Since this is not always possible, common approaches to minimize the impact of motion include (Zaitsev et al., 2015) supression of the signal from moving tissues through spatial saturation pulses; gradient moment nulling aimed at nulling the phase acquired by spins moving at a constant velocity; faster imaging sequences such as those using EPI or Parallel Imaging (PI): motion insensitive sequences, usually based on oversampling the center of $\boldsymbol{k}$–space such as Periodically Rotated Overlapping Parallel Lines with Enhanced Reconstruction (PROPELLER) (Pipe et al., 2002); and triggering/gating for cyclic types of motion, where different points are acquired within the same window of the cardiac cycle or breathing position. Finally, post–processing techniques can be used to estimate and correct for motion, either using callibration data termed "navigator echoes" or relaying on iterative approaches that try to minimize some measure of artifacts.

4. **Fourier Transform**: in MRI we do not acquire the data directly in the image space, but in the frequency domain, the so–called $\boldsymbol{k}$–space. In order to obtain the image, the Inverse Fourier Transform (iFT) needs to be applied and this processing can introduce undesired distortions (Brown et al., 2014). Probably the most well-known artifacts within this group are the Gibbs artifac and the aliasing artifact. The first, also called ringing artifact, arises from the truncation of the acquisition in $\boldsymbol{k}$–space. Theoretically, in order to represent a support–limited image with abrupt transitions, infinite frequency components are required. However, in practice this is unfeasible, and the need to truncate the number of acquired components manifests as undershoot and overshoot oscillations at high–contrast interfaces. This effects can never be entirely removed, but they can be minimized either by increasing the number samples or by filtering the acquired data with a smoothly decreasing window, as we observe in Fig.1.2.

   On the other hand, aliasing artifacts arise from the violation of the Nyquist sampling theorem, which states that the digital sampling rate must be at least twice the largest frequency component contained within the sampled signal. This is the case when the defined Field Of View (FOV) is smaller than the object being imaged, and it results in a folding over of the anatomic parts into the area of interest (see Fig.1.3 for a schematic depiction in 1D). Common approaches to correct for this artifact consist of either increasing the FOV and further discarding the undesired samples (No Phase Wrap) or using spatial saturations pulses outside the FOV to eliminate the signal from those tissues. Alternatively, Reduced Field Of View (rFOV) sequences aim at directly exciting a certain portion of the object by careful design of the imaging gradients in combination with the RF pulse.

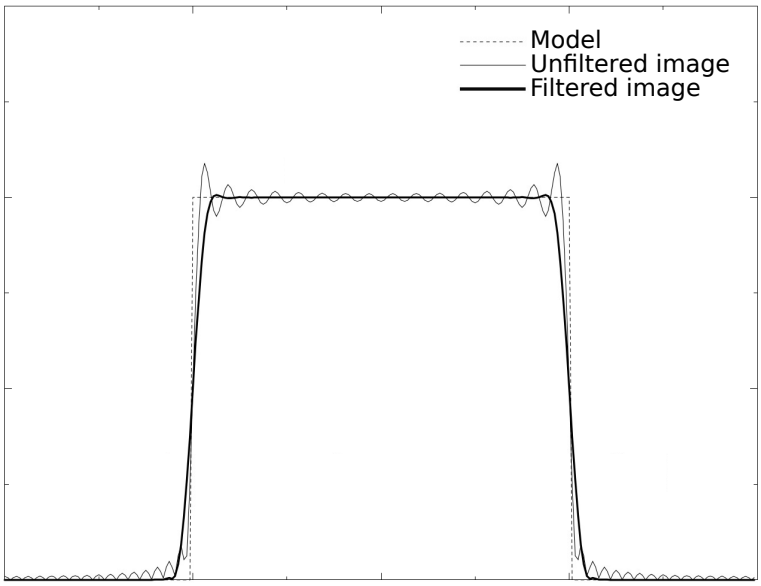

Figure 1.2: The ideal model is a rectangular function. However, the reconstructed image presents ringing artifact due to truncation. By applying a Hanning filter, it can be notably reduced af the expense of losing resolution. *Figure obtained from (Brown et al., 2014).*

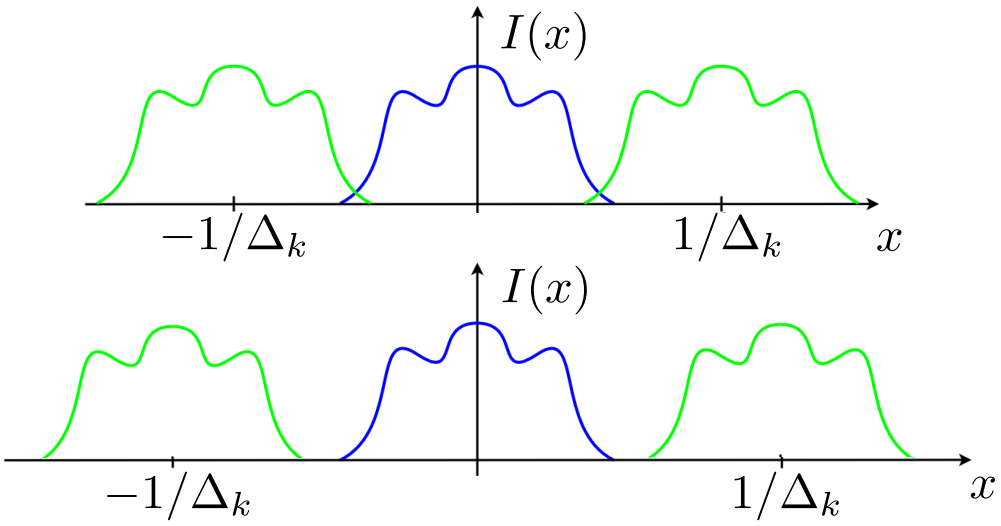

Figure 1.3: The effect of respecting or violating the Nyquist sampling theorem on Fourier reconstructions.

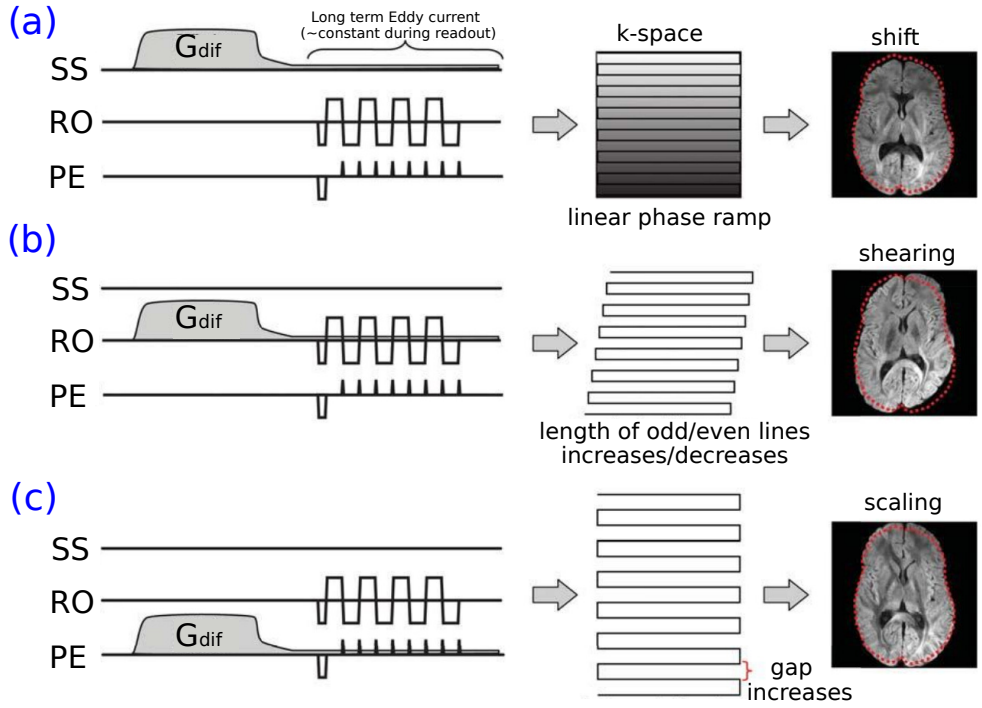

Figure 1.4: Eddy currents effects during the EPI readout. (a) An eddy current in the slice selection direction is equivalent to a constant offset in the B0–field, which introduces a linear phase ramp along the phase–encoding dimension in the acquired $\boldsymbol{k}$–space and results in a shift in the image–space. (b) An eddy current in the FE direction introduces an asymmetry in the positive–negative lobes of the FE gradients. This results in a shearing in the image. (c) An eddy current in the phase–encoding direction increases the size of the phase–encoding blips and changes the sampling density along that dimension. This results in a scaling in the image.

5. **Eddy currents**: according to the Faraday-Lenz Law of electromagnetism, the varying magnetic field originated by turning on and off the gradient pulses induces currents in the surrounding conductors, which in turn generates their own magnetic fields that may persist after the primary gradients are switched off (Jezzard et al., 1998). Furthermore, eddy currents scale up with the strength of the magnetic gradient. For this reason, Diffusion MRI (dMRI) is a particular scenario where they can cause severe distortions. First of all, since they modify the effective diffusion gradients, the b–value might be slightly different from expected (K. Jones, 2011), as seen in Fig.1.5.

   Second, the residual eddy current from the diffusion gradients may persist during the data acquisition window (Le Bihan et al., 2006). This can not only introduce signal losses due to improper spin rephasing, but might also deviate the imaging gradients used for frequency and phase–encoding from their original value. This mismatch would result in different types of distortion such as scaling, shearing or shifts, which need to be taken into account before further processing. Otherwise, since computing the diffusion coefficients requires several diffusion–sensitized images subject to different distortions,

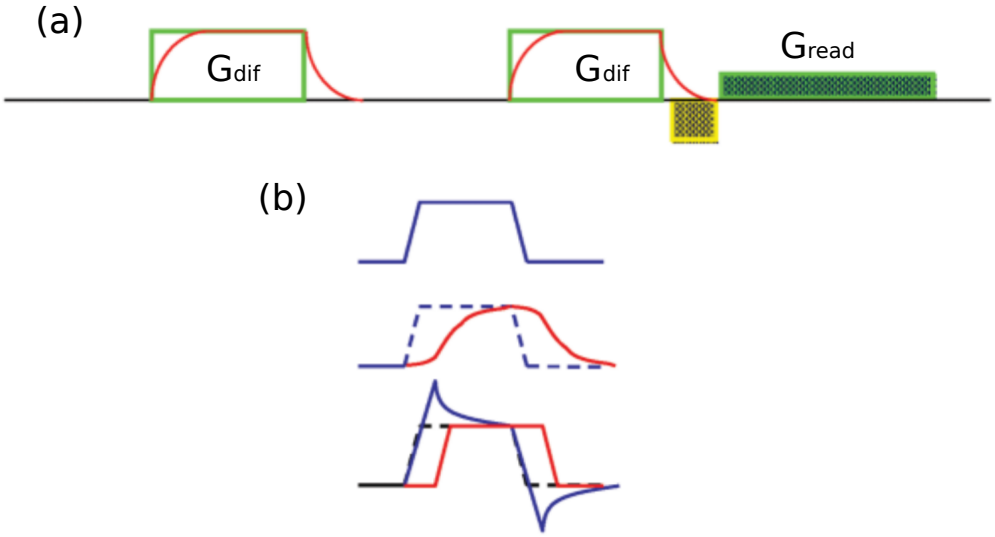

Figure 1.5: (a) Eddy currents modify the shape of the diffusion gradients (red) with respect to the desired one (green). (b) In order to obtain the expected shape one can make use of the pre–emphasis approach modifying the applied gradient (blue) to obtain the desired response (red). *Figure obtained from (Le Bihan et al., 2006).*

one can incur in misestimating them.

Third, since the shape of gradient pulses is slightly modified because of eddy currents, the acquisition window might not be completely centered around the echo, as in Fig.1.6. If we are using an EPI sampling scheme, as is often the case in dMRI, the back–and–forth trajectory in $\boldsymbol{k}$–space creates a shift between the position of the echo between the odd and even lines. As a consequence, we introduce a modulation along the phase–encoding direction in $\boldsymbol{k}$–space, which will manifest as a ghosting halfway across the image once the iFT is applied. Finally, commons strategies to reduce the eddy currents include using "self–shielded" gradient coils designed to reduce the amount of conductive surfaces, using gradient pre–emphasis to modify in advance the shape of the gradients so they have the desired waveform after eddy currents using a reference scan to estimate the phase shifts and further correct them.

6. **Noise**: as any other measurement process, MR imaging is subject to the undersired presence of a randomly distributed noise component which is uncorrelated with the original signal. The main sources for these perturbations are thermal noise, mainly generated from the imaged subject, and electronic noise, originated by the stochastic motion of free electrons in the RF coils. One of the objectives of MRI is precisely to provide sufficient signal relative to noise, which is measured by the SNR. It has been proven (Macovski, 1996) that the SNR is affected by certain factors as in the following equation:

$$\text{SNR} \propto \frac{V_h \cdot \sqrt{n_s}}{\sqrt{\text{BW}}}. \tag{1.1}$$

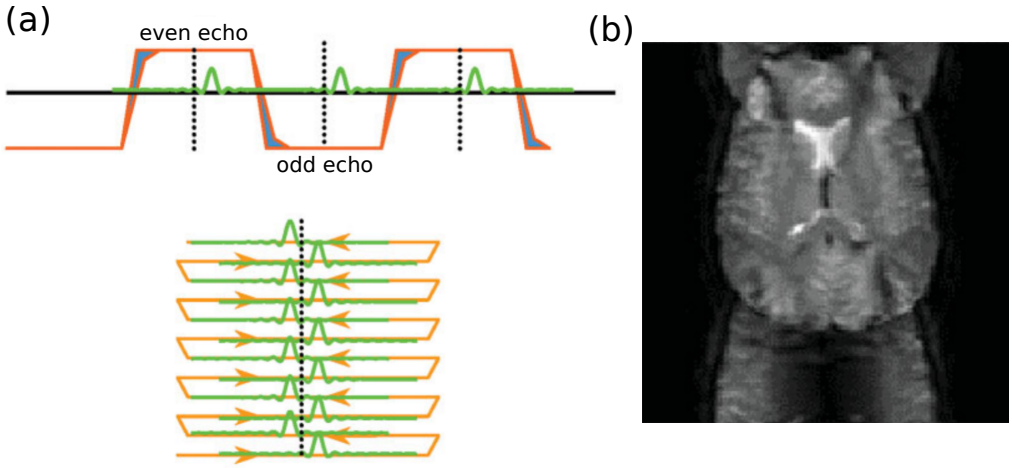

Figure 1.6: (a) Eddy currents modify the shape of the diffusion gradients (red) with respect to the desired one (green). (b) In order to obtain the expected shape one can make use of the pre–emphasis approach modifying the applied gradient (blue) to obtain the desired response (red). *Figure obtained from (Le Bihan et al., 2006).*

First, we observe that the SNR increases with the voxel size due to the larger amount of signal present at each location of the image. Second, we observe that it is proportional to the square root of the acquisition time $t_s = n_s/\text{BW}$. The bandwith of the receiver, BW, indicates the rate at which the received temporal signal is sampled. In order to respect the aforementioned Nyquist Theorem, the signal is filtered previous to the analog–digital conversion, setting the cut-off frequency to half the bandwidth. As a consequence, the amount of allowed noise into the system increases with the bandwidth BW of the filter. On the other hand, n_s refers to the number of acquired samples. As we have previously described, the data is acquired in the $\boldsymbol{k}$–space and the FT is applied to obtain the final image. Because of the averaging performed by the FT, increasing the number of points will in turn reduce the amount of noise. Finally, in order to cope with noise, the usual processing consists in finding a suitable probability distribution to model it, estimating the parameters that fully characterize that distribution and applying some sort of filtering to remove or reduce the noise in the final image (Aja Fernández and Vegas Sánchez-Ferrero, 2016).

1.2 Objectives

The main objective of this Thesis is **to better characterize the distortion introduced by two different types of artifacts in accelerated acquisitions and to develop the tools to deal with them in a feasible computational time. Specifically we will focus on (a) the estimation of exact g–factor maps in 2D/3D GeneRalized Autocalibrating Partial Parallel Acquisition (GRAPPA) and (b) the reconstruction of multi–shot dMRI suffering from intra–shot motion.**

The main objective can be divided in the following individual objectives:

1. To propose an exact and time–efficient method for the estimation of g–factor maps in 2D GRAPPA acquisitions under the assumptions of stationarity and uncorrelation in the original undersampled $\boldsymbol{k}$–space acquisition. The proposed algorithm will carry out a noise propagation analysis through all the steps in GRAPPA reconstruction and will exploit the symmetries in the involved matrices as well as the separability into blocks.

2. To extend the method to the more challenging case of 3D acquisitions due to the greater flexibility offered by the possibility of undersampling along the two phase–encoding directions. Particularly, memory requirements can be a limiting factor since the 3D GRAPPA reconstruction introduces $\boldsymbol{k}$–space correlations along the three dimensions.

3. To propose a joint image reconstruction and phase map estimation method for Multi–Shot (ms)–EPI dMRI under the assumption of intra–shot rigid motion. Exploiting the well known fact that rigid motion results in linear phase maps corruption, we propose a maximum likelihood formulation that iteratively estimates both the parameters characterizing the linear phase corruption and the reconstructed image.

4. To extend the previous algorithm to the more complex scenario of non–rigid subject motion. In order to model smooth non–linear phase variations on top of the linear component, a phase–map defined with a B–spline will be proposed. Similarly, a joint algorithm will alternate between the estimation of the linear component parameters, the parameters characterizing the B–spline and the reconstruction of the image.

1.3 Methodology

In this section we describe the methodology followed throughout this thesis, consisting in the following steps.

1. **Identification of a limitation or possible source of inaccuracies**: in the case of noise map estimation for GRAPPA reconstruction, it was observed that analytical methods were based on an approximate reconstruction that was only equivalent in the case of uniform subsampling patterns, rendering inexact the noise map obtention. On the other hand, state–of–the–art methods for navigator–free ms–EPI have explored different assumptions about the phase corruption maps, such as smoothness, low–rankness or limited support. However, to the best of our knowledge, the linearity associated to rigid–motion has only been considered in navigated strategies, either using a pre–scan (extra–navigated sequences) or introducing redundancy in the $\boldsymbol{k}$–space sampling (self–navigated sequences).

2. **Theoretical modelling**: the aim si to design a framework where the observed limitation is taken into acount and overcome. Regarding the g–factor estimation in GRAPPA acquisitions, the inaccuracy derived from the pseudo–equivalent formulation can be overcome by a noise–propagation analysis through all the steps in the reconstruction pipeline. In order to deal with the big size of the matrices involved in the process and provide a feasible solution with respect to computation time and memory requirements, the symmetries and the block separability are deeply exploited. As for the ms–EPI dMRI images under rigid–motion corruption, a joint maximum likelihood formulation is posed that alternates between the estimation of the parameters characterizing the linear phase maps and the reconstruction of the image following a greedy fashion. Key to successful reconstruction is the initialization to avoid local minima, which we propose to obtain based on a shot–by–shot SENSE reconstruction.

3. **Validation**: in order to test our methods, we first partly validated the individual modules that composed them, and later validated them globally, which in somes cases required to insert the developed algorithm as a segment of a multi-step reconstruction pipeline that comprised different post-processing stages. To do so, we have used three different types of data–sets, whose particular details are thoroughly described in the corresponding chapters:

    - Synthetic data–sets: simulated data–sets obtained from open databases or generated using the XCAT tool (Segars et al., 2010; Wissmann et al., 2014). In these scenarios we could exploit the knowledge of both the exact image to reconstruct and the undesired corruption, either regarding noise degradation or patient motion. Under this ideal conditions, one can test the range of validity of the developed algorithm, but it is important to notice that real acquisitions might deviate from our simplified assumptions.
    - Real phantom data–sets: in an attempt to isolate the source of artifacts under study while recreating a realistic acquisition, we have scanned several phantoms varying different parameters of the acquisition with the aim of observing to what extent our methods can provide a reasonable reconstruction. By proceeding this way, we were still able to obtain a silver ground–truth we could compare our reconstructions to.
    - In–vivo data–sets: last, in order to test the clinical applicability of the proposed algorithms, we acquired data from human subjects, again varying the acquisition parameters to test the soundness of our proposals. These datasets will be described in detail in the corresponding chapters.

4. **Diffusion of our contributions**: with the aim of sharing with the community our improvements with respect to the state–of–the–art, the current disertation is presented

describing the motivation for our studies, the theoretical framework, the detailed algorithms and the experiments carried out. Likewise, derived from the work described in this thesis, several publications have been submitted and presented in different journals and conferences.

1.4 List of publications

In pursuit of the aforementioned objectives, the following contributions have been published:

- Publications in indexed international journals:
  - Rabanillo, I., Aja-Fernández, S., Alberola-López, C., and Hernando, D. (2018a). Exact calculation of noise maps and g -factor in GRAPPA using a k–space analysis. *IEEE Transactions on Medical Imaging*, 37(2):480–490.
  - Rabanillo, I., Zhu, A., Aja-Fernández, S., Alberola-López, C., and Hernando, D. (2018c). Computation of exact g-factor maps in 3D GRAPPA reconstructions. *Magnetic Resonance in Medicine*.
- Publications in Conference and Workshop Proceedings:
  - Rabanillo, I., Hernando, D., and Aja-Fernández, S. (2016a). Variation of noise in multi-echo MRI acquisitions using Parallel Imaging. *Proceedings of the 33rd European Society for Magnetic Resonance in Medicine and Biology Scientific Meeting (ESMRMB '16), Vienna, Austria,*, 29:400.
  - Pieciak, T., Rabanillo, I., and Aja-Fernández, S. (2018). Bias correction for non-stationary noise filtering in MRI. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI)*
  - Rabanillo, I., Holmes, J., Guidon, A., Aja-Fernández, S., Alberola-López, C., Reeder, S., and Hernando, D. (2016b). ADC measurement accuracy in quantitative diffusion phantoms using reduced field-of-view and multi-shot acquisitions. In *Proceedings of the 33rd International Society of Magnetic Resonance in Medicine Workshop on Breaking the Barriers of Diffusion MRI.*
  - Zhang, Y., Holmes, J., Rabanillo, I., Guidon, A., Wells, S., and Hernando, D. (2018). Quantitative diffusion MRI using reduced field-of-view and multi-shot acquisition techniques: Validation in phantoms and prostate imaging. *Magnetic Resonance Imaging*, 51:173–181.

- Rabanillo, I., Sanz-Estebanez, S., Aja-Fernández, S., Hajnal, J., Alberola-López, C., and Cordero-Grande, L. (2018b). Joint image reconstruction and phase corruption maps estimation in multi-shot echo planar imaging. In *Proceedings of the 21st International Conference on Medical Image Computing & Computer Assisted Intervention. International Workshop on Computational Diffusion MRI (MICCAI/CDMRI 2018)*.
- Sanz-Estébanez, S., Rabanillo, I., Royuela-del Val, J., Aja-Fernández, S., and Alberola-López, C. (2018). Joint groupwise registration and ADC estimation in the liver using a b-value weighted metric. *Magnetic Resonance Imaging*, 37:1–9

1.5 Structure of the thesis

This dissertation includes a chapter dedicated to present the backfround, which will provide context for the core of our work. Our contributrions are centered on the problem of noise map estimation in GRAPPA acquisitions, both for 2D (Chapter 3) and 3D (Chapter 4), and the reconstruction of ms–EPI in dMRI under shot–to–shot phase corruption (Chapter 5). It is organized as follows:

- Chapter 2 briefly describes the fundamentals of MRI, PI and dMRI. The objective is to present an overview of the image formation and the underlying physical process in order to better understand the impact of both noise corruption and patient motion in the reconstructed final image.
- Chapter 3 provides an overview of the noise ang g-factor estimation problem in 2D GRAPPA reconstructions, describing the sources of inexactitude in the current state–of–the art methods. Additionally, we provide the details of the proposed algorithm to overcome these limitations and present analytical, simulation and *in–vivo* results to demonstrate the ability of our method to exactly estimate the noise maps under the assumption of stationarity and uncorrelation in the original $\boldsymbol{k}$–space.
- Chapter 4 presents an extension of the previous method to the more complicated case of 3D GRAPPA reconstructions, where the presence of two phase–encoding dimensions offers greater flexibility for undersampling and poses an increased challenge regarding memory requirements and computation time. An analytical proof of the exactitude of the method is provided together with simulations and *in–vivo* experiments to support the mathematical results.
- Chapter 5 introduces the problem of image reconstruction in ms–EPI for dMRI, where different shots present phase corruption discrepancies that result in ghosting artifacts.

Initially, we present an iterative algorithm that jointly estimates the phase maps and the final image under the assumption of rigid–motion, or equivalently, under the assumption of linear phase maps. Additionally, we extend this philosophy to the more complex case of non–linear phase corruption derived from non–rigid motion, such as pulsatile motion. Finally, we present simulations to test the ability of the method to reconstruct diffusion images under varying levels of noise corruption, number of shots and phase distortion.

- Chapter 6 concludes the dissertation with the discussion and final conclusions of the presented work, including the main contributions of this thesis, the limitations of our proposals and suggestions for future lines of research.

CHAPTER

TWO

BACKGROUND

2.1 Magnetic Resonance Imaging Fundamentals

Magnetic Resonance (MR) is based on the quantum property of the nuclear spin and its interaction with an external magnetic field termed $\vec{\mathbf{B_0}}$. However, even though the spin is a quantum effect, the phenomena of MR can be fully explained by classical mechanics (Hanson, 2008). In the presence of a static magnetic field, the spins precess around $\vec{\mathbf{B_0}}$ at an angular frequency known as the Larmor frequency and given by:

$$\omega_0 = \gamma \left\| \vec{\mathbf{B_0}} \right\|_2, \tag{2.1}$$

where γ is a nucleus–specific constant known as the gyromagnetic ratio. When no external magnetic field is applied, spins point in all the directions with equal probability. Nevertheless, when the polarizing $\vec{\mathbf{B_0}}$–field is used, the spins show a slight tendency to point along the direction given by the main magnetic field, creating a net magnetization aligned with $\vec{\mathbf{B_0}}$, as shown in Fig.2.1.

In order to generate a measurable signal, we must tip the magnetization vector away from the external field direction. This is done by applying, during a short time, a time–varying RF pulse, $\vec{\mathbf{B_1}}(t)$. This pulse is perpendicular to the $\vec{\mathbf{B_0}}$–field and its radiofrequency ω_{RF} is tuned to the Larmor frequency, being modelled as (Liang and Lauterbur, 2000):

$$\vec{\mathbf{B_1}}(t) = B_1^e(t) \cdot \left[\cos(\omega_{RF} \cdot t)\vec{\mathbf{x}} - \sin(\omega_{RF} \cdot t)\vec{\mathbf{y}} \right], \tag{2.2}$$

where $B_1^e(t)$ defines the envelope of the pulse. The effect of the net magnetization vector $\vec{\mathbf{M}} = [M_x, M_y, M_z]$ under the application of a magnetic field $\vec{\mathbf{B_1}}(t)$ is modelled by the Bloch equation:

$$\frac{d\vec{\mathbf{M}}}{dt} = \gamma \vec{\mathbf{M}} \times \vec{\mathbf{B}} - \frac{M_x\vec{\mathbf{x}} + M_y\vec{\mathbf{y}}}{T_2} - \frac{(M_z + M_z^0)\vec{\mathbf{z}}}{T_1}, \tag{2.3}$$

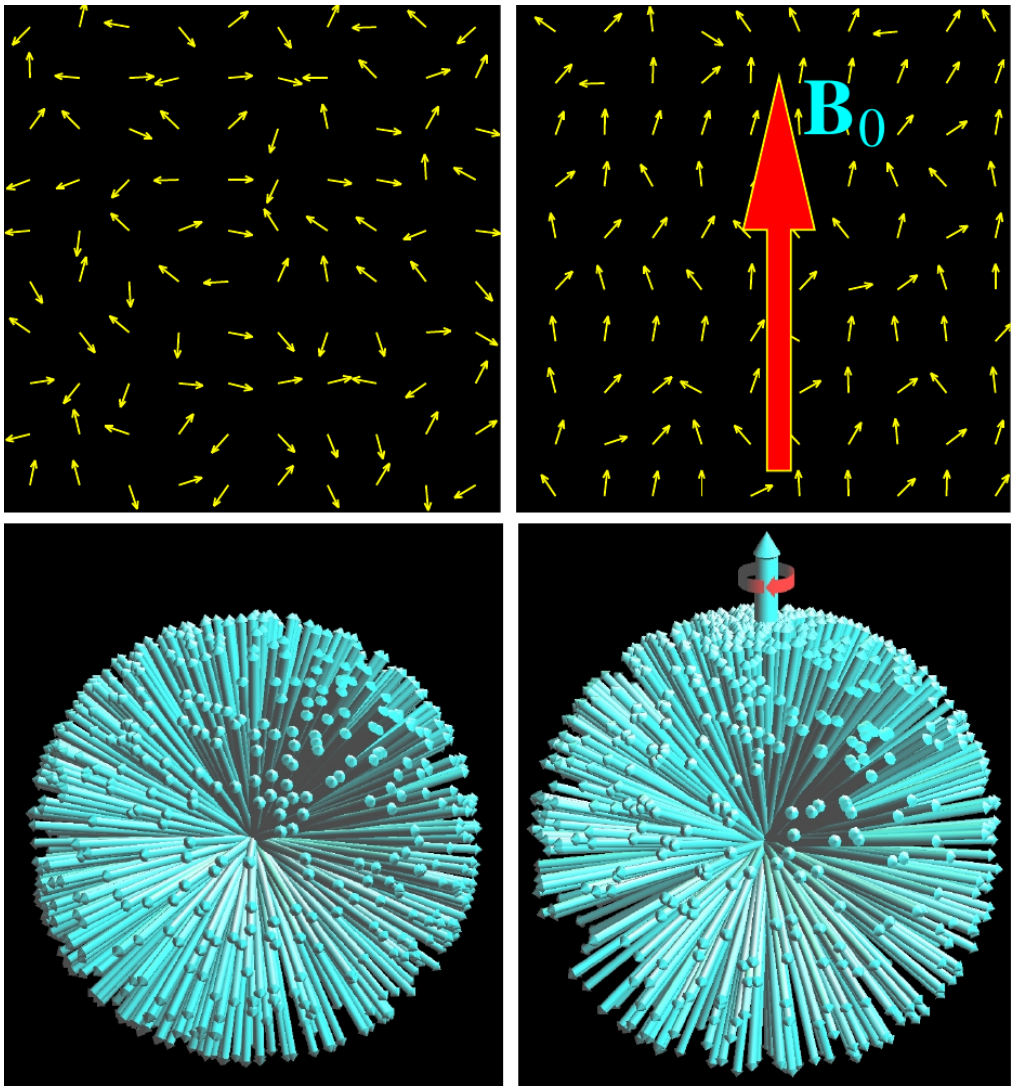

Figure 2.1: The figure shows the same situations in both 2D and 3D. In the 3D case, to better illustrate their directions, all the spins share their origin. In the absence of an external magnetic field, the spins point in all directions with equal probability. However, when the $\vec{\mathbf{B_0}}$ is applied, they tend to align with it, although all directions remain possible. *Figure obtained from (Hanson, 2008).*

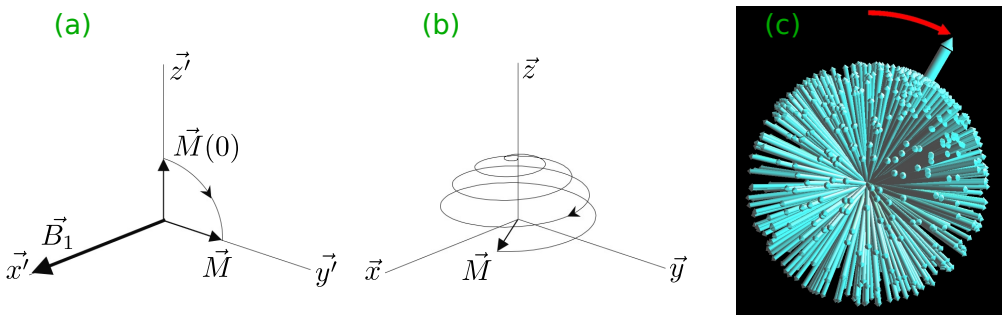

Figure 2.2: Illustration of how the application of an RF pulse along the $\vec{x}$ direction tips down the net magnetization around , both in the rotating frame of reference (a) and the static frame of reference (b). We also show how this precession around the RF pulse affects the whole distribution of spins in (c). *Figure obtained from (Brown et al., 2014) and (Hanson, 2008).*

where $\vec{\mathbf{B}} = \vec{\mathbf{B_0}} + \vec{\mathbf{B_1}}(t)$ is the overall magnetic field, M_z^0 is the thermal equilibrium value of $\vec{\mathbf{M}}$ when only $\vec{\mathbf{B_0}}$ is present, and T_1 and T_2 are the time constants characterizing the relaxation process we will explain later on. Due to the shortness of the RF pulse, the last two terms can be dropped for the study of its effects. Thus, solving the previous equation shows that the effect of the B_1–field is to rotate the spins distribution as a whole around it. The amount of rotation, characterized by the flip angle α, depends on the strength and duration of the RF pulse. A common approach at this point is to get rid of the effect of the Larmor-precession, which is done by defining a rotating frame of reference synchronized with the spins precession frequency. In Fig.2.2 we show the effect of the RF pulse in both the static and the rotating frame of reference.

After the perturbation from its equilibrium, the magnetization spin system returns to that state following the laws of thermodynamics. This process, known as relaxation, see Fig.2.3, is modelled by the aforementioned Bloch equation given in Eq.2.3. Operating in the rotating frame of references $[x', y', z']$, the last terms explain the two different phenomena that take place:

$$\begin{cases} \frac{dM_{z'}}{dt} = -\frac{M_{z'} + M_z^0}{T_1} \\ \frac{d\vec{M_{x'y'}}}{dt} = -\frac{M_{x'y'}}{T_2}. \end{cases} \tag{2.4}$$

First we have the longitudinal or spin–lattice relaxation by whom the longitudinal net magnetization recovers in order to realign with the static B_0–field:

$$M_{z'}(t) = M_z^0 \cdot \left(1 - e^{-\frac{t}{T_1}}\right) + M_{z'}(t=0) \cdot e^{-\frac{t}{T_1}}. \tag{2.5}$$

Second we have the transverse or spin–spin relaxation, which refers to the decrease of the transverse component of the net magnetization due to the dephasing of the spins by the small local contributions to the magnetic field that vary from spin to spin:

$$M_{xy'}(t) = M_{x'y'}(t=0) \cdot e^{-\frac{t}{T_2}}. \tag{2.6}$$

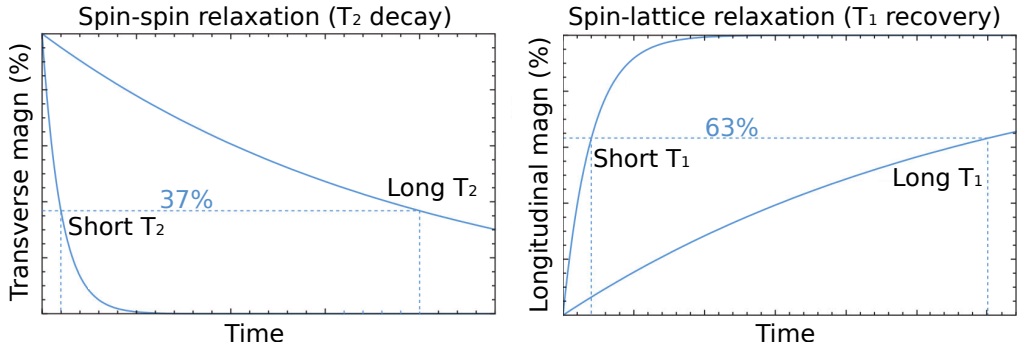

Figure 2.3: Relaxation of both the transverse and longitudinal components after the application of a 90°–RF pulse for different constant decays. *Figure obtained from (Suetens, 2002).*

Besides the loss of coherence caused by the different nuclear environment observed by each spin, $B0$–field inhomogeneities will introduce further dephasing, which can be modelled by a decay constant $T_2' = (\gamma \cdot \Delta B_{\text{inhom}})^{-1}$. This gives rise to a modified version of the previous equation:

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \frac{1}{T_2'} \rightarrow M_{xy'}(t) = M_{x'y'}(t=0) \cdot e^{-\frac{t}{T_2^*}}. \tag{2.7}$$

These relaxation constants, as well as the proton density, are parameters characteristic of each tissue. By properly setting two parameters from the MR acquisition sequence, the Echo Time (TE) and the Repetition Time (TR), we can weight our image in order to provide contrast sensitive to any of these three parameters.

2.1.1 Signal localization

Using RF pulses allows us to create a time–varying signal in the transverse plane that we can measure. However, the receptor will observe the superposition (the integral) of the signal coming from all the excited spins. For this reason, it is necessary to localize the position of the spins in order to differentiate local contributions from different parts of the object. There are two key types of spatial localization: selective excitation and spatial encoding. Both of them are based on the use of a gradient field and exploit the Fourier relationship between the image–space and the $\boldsymbol{k}$–space. A gradient B_G can be defined as a spatial linearly varying magnetic field

$$B_G(\vec{\mathbf{R}}) = \langle \vec{\mathbf{G}}, \vec{\mathbf{R}} \rangle = G_x X + G_y Y + G_z Z. \tag{2.8}$$

The application of such a gradient modifies the precession frequency of the spins of the object along its direction

$$\omega(\vec{\mathbf{R}}) = \omega_0 + \gamma B_G(\vec{\mathbf{R}}), \tag{2.9}$$

and as a consequence, the received signal would be

$$s(t) = \int_\Omega M_{x'y'}(\vec{\mathbf{R}}, t) e^{-i\gamma \int_0^t \langle \vec{\mathbf{G}}(t'), \vec{\mathbf{R}} \rangle dt'}. \tag{2.10}$$

Defining the $\boldsymbol{k}$–space trajectory $\vec{\mathbf{k}}(t)$ as

$$\vec{\mathbf{k}}(t) = \frac{\gamma}{2\pi} \int_0^t \vec{\mathbf{G}}(t')dt', \tag{2.11}$$

we can express the received signal as:

$$s(t) = \int_\Omega M_{x'y'}(\vec{\mathbf{R}}, t) e^{-2\pi i \langle \vec{k}(t), \vec{\mathbf{R}} \rangle}. \tag{2.12}$$

We observe that the signal $s(t)$ is nothing else but the three–dimensional Fourier Transform (Fourier Transform (FT)) evaluated at location $\vec{\mathbf{k}}(t)$. The idea behind spatial encoding is to use gradients to go through the different locations of $\boldsymbol{k}$–space while we sample the received signal. Further, we will need to solve the inverse problem of reconstructing the MR image from the sampled trajectory in $\boldsymbol{k}$–space, whose difficulty will precisely depend on the way we sampled it. Equivalently, this philosophy can be applied in excitation (Pauly et al., 1989), where one would make use of the gradients to travel through $\boldsymbol{k}$–space during the application of the RF pulse. It can be proved that the three–dimensional Fourier Transform of the RF pulse defined along the travelled $\boldsymbol{k}$–space would define the excitation profile. This is the basis for slice selection, but it can be used for much more complex patterns, such as in rFOV sequences (Saritas et al., 2008).

2.2 Accelerated Imaging

As we have already mentioned, MRI is a slow modality, which causes the resulting images to suffer from different types of artifacts. For this reason, reducing the scanning time has always been one of the major goals of the research community. In this section we present two of the most successful and widely used strategies for that purpose: Parallel Imaging (PI) and Echo–Planar Imaging (EPI).

2.2.1 Parallel Imaging

PI (Larkman and Nunes, 2007; Deshmane et al., 2012) has been possible thanks to the hardware advances in multichannel phased arrays. These arrays have been available since the late 1990s (Roemer et al., 1990), although they were initially developed to improve the SNR (Constantinides et al., 1997). By arranging the different coils (channels or antennas) in such a way that covers the desired FOV, the images resulting from each receiver can later be combined. As a result of the image combination, the noise can be reduced and thus the SNR of the final image is improved.

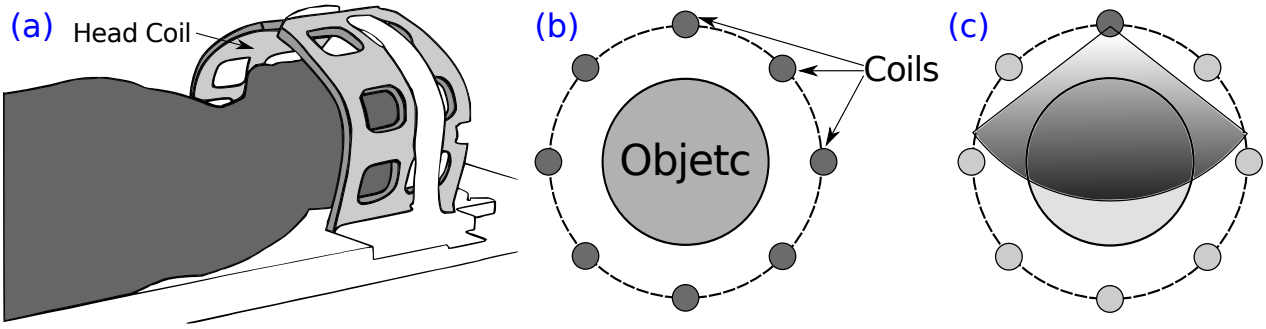

Figure 2.4: (a) Example of a head coil for MRI acquisitions. (b) Distribution of an 8–coil system around the object. (c) Spatial sensitivity of a single coil. *Figure obtained from (Aja Fernández and Vegas Sánchez-Ferrero, 2016).*

Data Sampling and Image Reconstruction

In traditional MRI one samples a single $\boldsymbol{k}$–space line per TR and thus multiple excitations are required to acquire the whole slice. Moreover, this needs to be repeated for as many slices we need for the imaged 3D volume. As a consequence, the total acquisition time is proportional to the number of lines and slices needed to cover the object. Nevertheless, soon after the introduction of phased arrays, researchers noticed that their use could be exploited to reduce the scan time as well (Hutchinson and Raff, 1988; Carlson and Minemura, 1993).

In PI data is sampled simultaneously at L different coils distributed around the object. In such a way, the signal they measure varies as a function of position (the closer they are to a location, the stronger the signal they record), as we depict in Fig.2.4.

Mathematically, the coils can be characterized by their sensitivity map $C_l(\boldsymbol{x})$, allowing to model the signal $S_l(\boldsymbol{x})$ they receive as:

$$S_l(\boldsymbol{x}) = C_l(\boldsymbol{x}) \cdot S(\boldsymbol{x}), \qquad l = 1, \ldots, L \tag{2.13}$$

where $S(\boldsymbol{x})$ denotes the original image. In order to reduce the scan time, one would skip the acquisition of a portion of the lines in the original $\boldsymbol{k}$–space, which as we have already mentioned, results in an aliasing in the image domain (see Fig.2.5). However, the availability of the differently weighted data at each antenna allows us to recover the missing data. Multiple sampling and reconstruction strategies have been proposed, but probably the most well-known are SENSE and GRAPPA.

In the SENSitivity Encoding (SENSE) Cartesian philosophy the acquired signal is modelled by means of a matrix operation: (Pruessmann et al., 2001)

$$\vec{\mathbf{y}} = \mathbf{E}\vec{\mathbf{x}} = \mathbf{A}\mathcal{F}\mathbf{S}\vec{\mathbf{x}} \tag{2.14}$$

where $\vec{\mathbf{x}}$ is the true MR image and $\vec{\mathbf{y}}$ is the undersampled $\boldsymbol{k}$–space data acquired by all the coils, defined as single column vectors. The encoding operator $\mathbf{E}$ consists of several steps.

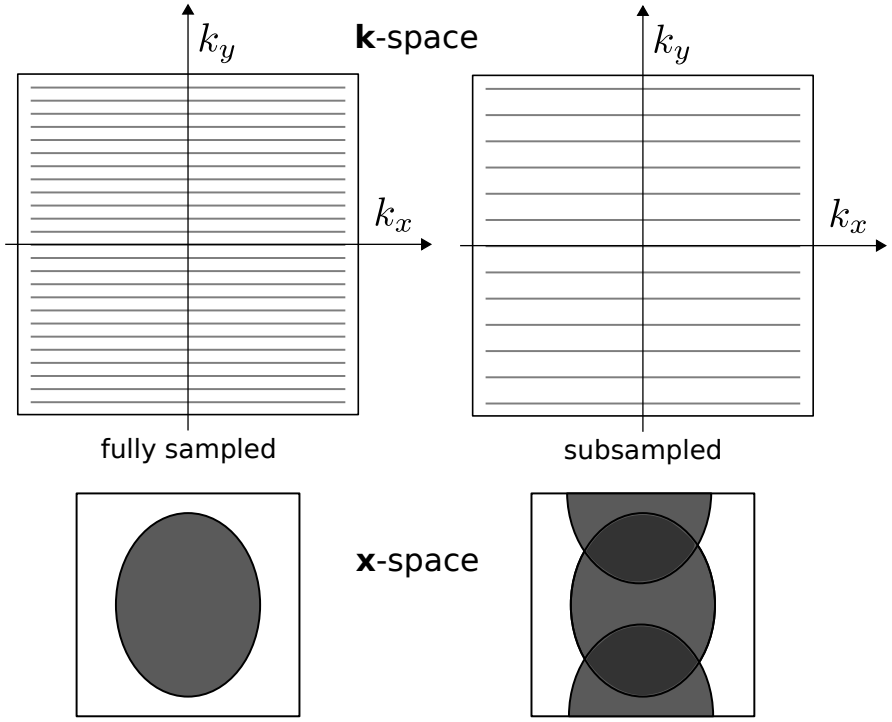

Figure 2.5: One way to accelerate the acquisition is by skipping some lines in the $\boldsymbol{k}$–space. In the example, we apply an acceleration rate $R = 2$, which in the image space introduces aliasing. *Figure obtained from (Aja Fernández and Vegas Sánchez-Ferrero, 2016).*

First we multiply by the coil sensitivities $\mathbf{S} = [\mathbf{S}_1 \dots \mathbf{S}_L]^T$. Then we transform into the Fourier domain using $\mathcal{F} = \mathbf{I}_L \otimes \mathbf{F}_I$, where $\mathbf{I}_L$ refers to the identity matrix of size $L \times L$, $\mathbf{F}_I$ is the Fourier transform matrix and $\otimes$ denotes the Kronecker product. Last we undersample using a mask $\mathbf{A}$ selecting the acquired k-space positions, which are the same for all coils. In case we are using non–Cartesian trajectories an extra operator to perform the degridding operation needs to be incorporated into $\mathbf{A}$. The acceleration factor R will be given by the ratio of the acquired and reconstructed k-space positions.

The reconstruction problem can then be posed as the solution of the previous linear system and given that the acceleration rate R is smaller than the number of available channels L, it can be solved in a minimum squared error sense. Under regular subsampling pattern, the reconstruction can be highly accelerated by operating directly in the image space, where several replicas of the image overlap. Therefore, we can solve pixel–wise the system for the overlapping locations, as depicted in Fig. 2.6.

This can be done either analytically or iteratively, but in both cases there is an exact way to compute the noise maps in the reconstructed image. Moreover, since the sensitivity profiles of the channels vary across the image, the noise standard–deviation in the final image will not necessarily be stationary anymore.

However, SENSE requires the knowledge of the sensitivity profiles $C_l(\boldsymbol{x})$, which are mea-

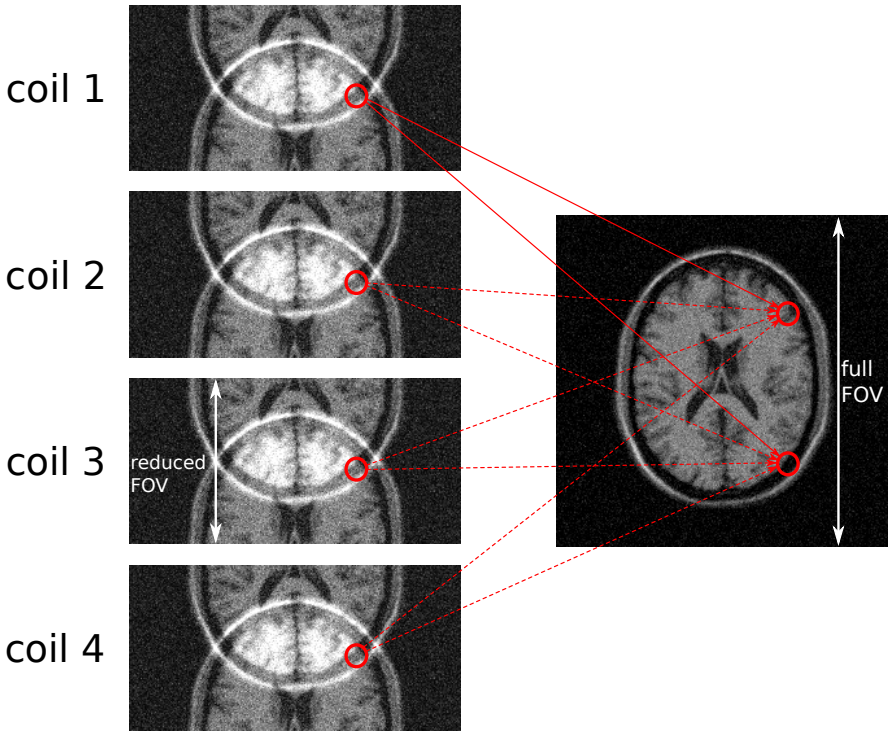

Figure 2.6: Example of a regular subsampling pattern with acceleration rate $R = 2$ and how each pair of pixels overlap in the acquired data. This can be exploited for fast reconstruction by operating pixel–wise directly in image–domain. *Figure obtained from (Aja Fernández and Vegas Sánchez-Ferrero, 2016).*

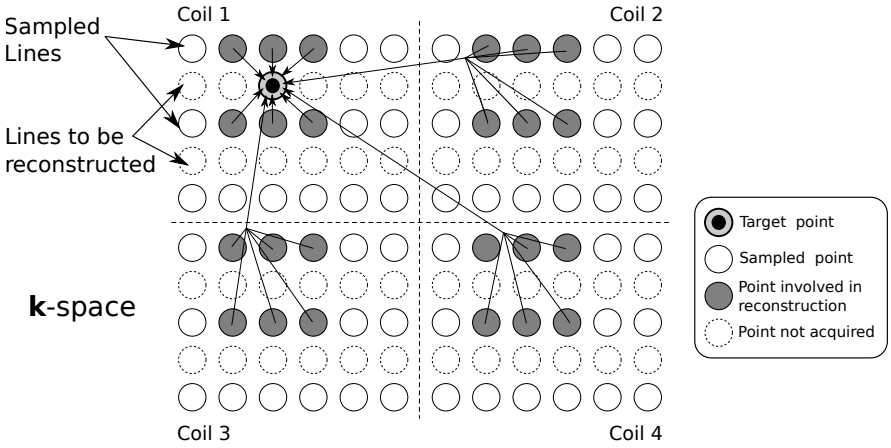

Figure 2.7: Illustration on how the GRAPPA kernel is applied to reconstruct the missing data from the acquired points from all the channels for the case of an acceleration rate $R = 2$ with 4 coils. *Figure obtained from (Aja Fernández and Vegas Sánchez-Ferrero, 2016).*

sured usually from a pre-scan. Since this can be time inefficient and subject to variability during scanning, there was interest in developping a self–calibrated strategy that can reconstruct the image from the data itself. To this end, GRAPPA proposes to solve the reconstruction problem directly in $\boldsymbol{k}$–space (Griswold et al., 2002). In order to do so, the authors propose to estimate the skipped $\boldsymbol{k}$–space lines and then apply the inverse Fourier Transform to obtain the coil images, which will be further combined to form the final reconstructed image. The idea behind GRAPPA comes from the observation that the coil multiplication in the forward model described in Eq.2.15 becomes a convolution in $\boldsymbol{k}$–space:

$$s_l(\boldsymbol{k}) = c_l(\boldsymbol{k}) \circledast s(\boldsymbol{k}), \tag{2.15}$$

where the previous signals are the Fourier transform of those described in Eq.2.13. Therefore, they assume that every point in $\boldsymbol{k}$–space can be restored as a combination from its neighbours weighted with a stationary kernel across $\boldsymbol{k}$–space (see Fig.2.7), which is somehow related to the sensitivity maps.

Therefore, GRAPPA reconstruction takes place directly in the $\boldsymbol{k}$–space, where each point in the missing lines is computed as an interpolation from its neighborhood $\eta(\boldsymbol{k})$ in all the coils (Griswold et al., 2002):

$$s_l^R(\boldsymbol{k}) = \sum_{m=1}^{L} \sum_{\boldsymbol{c} \in \eta(\boldsymbol{k})} s_m^S(\boldsymbol{k} - \boldsymbol{c}) \cdot \omega_m(l, \boldsymbol{c}), \tag{2.16}$$

where $s_m^S(\boldsymbol{k})$ is the sampled $\boldsymbol{k}$–space signal from coil m, $\omega_m(l, \boldsymbol{c})$ are the complex reconstruction weights for coil l, and $s_l^R(\boldsymbol{k})$ is the reconstructed $\boldsymbol{k}$–space signal from coil l. The weights $\omega_m(l, \boldsymbol{c})$ are usually calculated from a fully sampled low-frequency region in the $\boldsymbol{k}$–space, called the Auto Calibration Signal (ACS) lines (Griswold et al., 2002), which in turn is most often collected during the accelerated scan, as illustrated in Fig.2.8. This is the reason why GRAPPA is considered an autocalibrated approach.

In order to obtain the final image $S_T(\boldsymbol{x})$, it is necessary to combine the data from all the channels $S_l^R(\boldsymbol{x})$ obtained by Fourier transformation of $s_l^R(\boldsymbol{x})$. This can be done using the Sum of Squares (SoS) as in (Aja-Fernández et al., 2011), or using a properly weighted linear combination as described by (Roemer et al., 1990). Complex imaging, as proposed by (Prah et al., 2010; Eichner et al., 2015; Sprenger et al., 2016), provides certain advantages such as avoiding noise floor effects or exploiting the information contained in the phase, as shown by (Hernando et al., 2013; Ghugre et al., 2006) and it also preserves the gaussian behaviour in the final image due to the linearity of the whole process. For these reasons, we chose this type of linear coil–combination. Specifically, we have chosen the linear method proposed in (Walsh et al., 2000), where the combination vector is computed to maximize the SNR ratio in the final image, which is expressed as:

$$S_T(\boldsymbol{x}) = \vec{\mathbf{m}}(\boldsymbol{x}) \cdot \vec{\mathbf{S}^R}(\boldsymbol{x}) = \sum_{l=1}^{L} m_l(\boldsymbol{x}) \cdot S_l^R(\boldsymbol{x}), \tag{2.17}$$

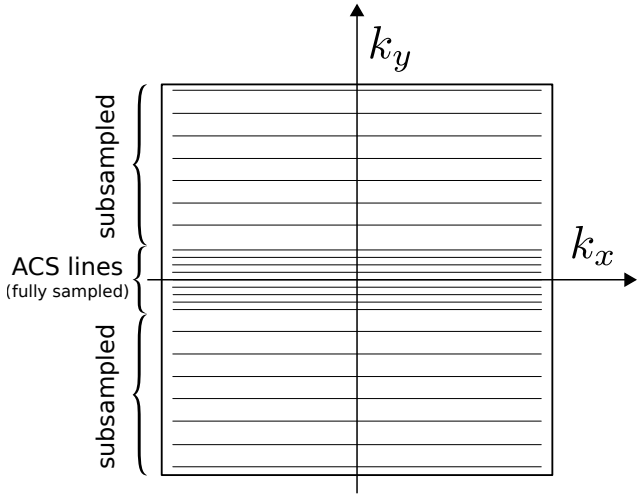

Figure 2.8: Example of GRAPPA sampling pattern in order to include the ACS lines: the center of the $\boldsymbol{k}$–space is fully–sampled in order to estimate from this low–frequency region the weights used for reconstruction. *Figure obtained from (Aja Fernández and Vegas Sánchez-Ferrero, 2016).*

where $\vec{\mathbf{m}}(\boldsymbol{x}) = [m_1(\boldsymbol{x}) \cdots m_L(\boldsymbol{x})]^T$ is a vector combining the information from each coil, and the $\boldsymbol{x}$ dependence indicates that the operation is pixel–wise.

Noise amplification

One of the problems associated with PI is the resulting noise amplification caused by the image reconstruction and the availability of fewer samples. In order to carry out a noise study, it is a common assumption to model the data acquired across the L coils $\vec{\mathbf{s}}(\boldsymbol{k}) = [s_1(\boldsymbol{k}) \cdots s_L(\boldsymbol{k})]$ as in (Aja Fernández and Vegas Sánchez-Ferrero, 2016; Thünberg and Zetterberg, 2007):

$$\vec{\mathbf{s}}(\boldsymbol{k}) = \vec{\mathbf{a}}(\boldsymbol{k}) + \vec{\mathbf{n}}(\boldsymbol{k}; \boldsymbol{\Gamma}_a, \mathbf{C}_a), \tag{2.18}$$

where $\vec{\mathbf{a}}(\boldsymbol{k}) = [a_1(\boldsymbol{k}) \cdots a_L(\boldsymbol{k})]$ is the noiseless signal in the $\boldsymbol{k}$–space and $\vec{\mathbf{n}}(\boldsymbol{k}; \boldsymbol{\Gamma}_a, \mathbf{C}_a)$ is the acquisition noise, which comes from a L-variate complex zero-mean normal distribution (Goodman, 1963) for every point. This distribution is characterized by the covariance matrices $\boldsymbol{\Gamma}_a = E\{\vec{\mathbf{n}} \cdot \vec{\mathbf{n}}^H\}$ and $\mathbf{C}_a = E\{\vec{\mathbf{n}} \cdot \vec{\mathbf{n}}^T\}^*$, where the operator ^{T} denotes the transpose of a matrix, whereas the operator ^{H} refers to the Hermitian of a matrix, i.e., its conjugate transpose. Furthermore, noise is assumed to be stationary, which implies that $\boldsymbol{\Gamma}_a$ and $\mathbf{C}_a$ do not depend on $\boldsymbol{k}$.

*These matrices are related to the variance matrices of the real and imaginary parts in the following manner: $\boldsymbol{\Gamma} = \boldsymbol{\Sigma}_\mathbf{x} + \boldsymbol{\Sigma}_\mathbf{y} + i(\boldsymbol{\Sigma}_\mathbf{yx} - \boldsymbol{\Sigma}_\mathbf{xy})$, $\mathbf{C} = \boldsymbol{\Sigma}_\mathbf{x} - \boldsymbol{\Sigma}_\mathbf{y} + i(\boldsymbol{\Sigma}_\mathbf{yx} + \boldsymbol{\Sigma}_\mathbf{xy})$. Usually it is assumed that the real and imaginary parts are equally distributed and uncorrelated, i.e., $\boldsymbol{\Sigma}_\mathbf{x} = \boldsymbol{\Sigma}_\mathbf{y}$ and $\boldsymbol{\Sigma}_\mathbf{yx} = \boldsymbol{\Sigma}_\mathbf{xy} = 0$, which results in $\boldsymbol{\Gamma} = 2\boldsymbol{\Sigma}_\mathbf{x}$ and $\mathbf{C} = 0$. For this reason, it suffices to analyse $\boldsymbol{\Sigma}_\mathbf{x}$. However, in order to keep our analysis as general as possible, we chose not to impose these constraints.

If the image reconstruction consists of a simple inverse Fourier Transform, as it is the case in conventional fully–sampled acquisitions, the noise will be spatially stationary and incorrelatad through the image. However, if more complex operations are performed, this may no be longer apply. This is the case for accelerated acquisitions, where one can observe noise amplification with respect to fully–sampled schemes. Moreover, in general they will introduce spatially varying noise maps in the final image. This noise amplification is typically described in terms of the so-called g-factor, as defined by (Breuer et al., 2009):

$$g_T(\boldsymbol{x}) = \frac{\mathrm{SNR}_{\mathrm{full}}(\boldsymbol{x})}{\mathrm{SNR}_{\mathrm{acc}}(\boldsymbol{x}) \cdot \sqrt{R_{\mathrm{eff}}}} = \frac{\sigma_{\mathrm{acc}}(\boldsymbol{x})}{\sigma_{\mathrm{full}}(\boldsymbol{x}) \cdot \sqrt{R_{\mathrm{eff}}}}, \tag{2.19}$$

where $\mathrm{SNR}_{\mathrm{full,acc}}$ and $\sigma_{\mathrm{full,acc}}$ denote the signal-to-noise ratio and the standard deviation in the fully sampled image and in the sub-sampled image after reconstruction, respectively, and R_{eff} denotes the effective acceleration (ratio between the number of lines in the reconstructed $\boldsymbol{k}$–space and the number of acquired lines).

GRAPPA is widely used and has proved to be a robust alternative to SENSE by avoiding the need to estimate the coil profiles. However, the computation of the noise maps in the reconstructed image has shown to be a challenging problem due to the size of the matrices involved in the process. An exact computation will require to follow the noise propagation through all the steps, and has been so far considered unfeasible. To overcome this problem, different strategies have been pursued. First, Monte–Carlo simulations can be carried out to obtain multiple reconstructions corrupting the original data with synthetic noise, and then obtaining the sample standard deviation. Nevertheless, this is a time consuming approach, which led into looking for an analytical alternative. Breuer et. al. proposed to reformulate GRAPPA reconstruction in image–space as a pixel–wise multiplication, which notably simplifies the computation of the noise maps. This equivalence holds exact when the sub-sampling pattern is uniform across $\boldsymbol{k}$–space, which is not the case when the ACS is included within the sampled data or when a variably–density strategy is followed, leading to errors in the estimated noise maps. For this reason, developing an exact analytical method which is computationally feasible remains a challenge.

2.2.2 Echo–Planar Imaging

Data Sampling and Image Reconstruction

Another way to reduce the acquisition time is by collecting multiple lines of $\boldsymbol{k}$–space data after a single RF excitation. The extreme case is when we use a train of echoes after a single excitation (a single shot) with a short phase–encoding blip in order to collect all the data to reconstruct the image. This is known as EPI and in Fig.2.9 we show the sequence diagram together with the $\boldsymbol{k}$–space trajectory.

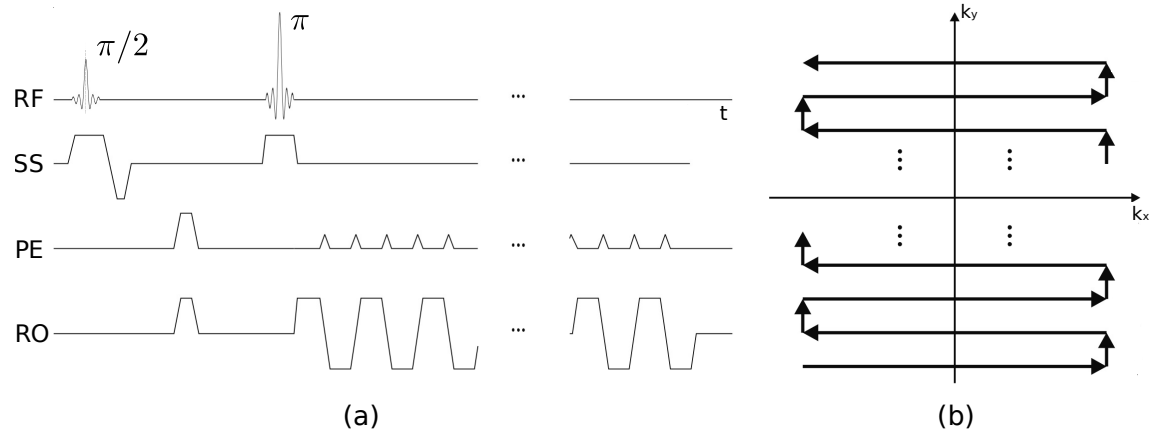

Figure 2.9: Example of EPI imaging sequence for spin–echo (a) together with the trajectory followed in $\boldsymbol{k}$–space (b). *Figure obtained from (Brown et al., 2014).*

EPI, first described by (Mansfield, 2003), is on of the fastest MRI pulse sequences and it can be used both for 2D and 3D imaging, although it is mostly employed in 2D acquisitions. However, compared to conventional spin–echo or gradient–echo imaging, EPI is more prone to different types of artifacts. Reducing the overall scanning time comes at the expense of increasing the echo train length. As a consequence, certain artifacts introduce more severe distortions in the image. First, one of the most common artifacts in EPI is Nyquist ghostings due to mistmatches between odd and even lines, which arise from the delay in the position of the formed echo. Second, we can observe chemical shift artifacts. Unlike in spin–echo or gradient–echo, the displacement between fat and water will appear mainly along the PE direction since the bandwith across this dimension is much smaller. Third, the image will be subject to geometric distortion related to B0–field inhomogeneities, magnetic susceptibility variations or eddy currents, resulting in deformations such as shifts, shearings or rotations. Last, since we acquire the different $\boldsymbol{k}$–space lines at different times, each one will carry a different T_2/T_2^*–weighting, depending on if we use a spin–echo or gradient–echo based sequence. Therefore we might observe blurring along the phase–encoding dimension.

Since many of the artifacts that arise from EPI acquisitions are related to the long echo trains, an intermediate solution consists in segmenting the acquisition into multiple shots, each one after its own RF excitation. This strategy is termed ms–EPI and it provides a trade–off between overall scanning time and artifact distortion. By dividing the acquisition into multiple shots, we can either reduce the artifact corruption or increase the spatial resolution, although the longer scanning time makes this sequence more sensitive to patient motion. In Fig.2.10 we can observe different sampling schemes for the same number of shots.

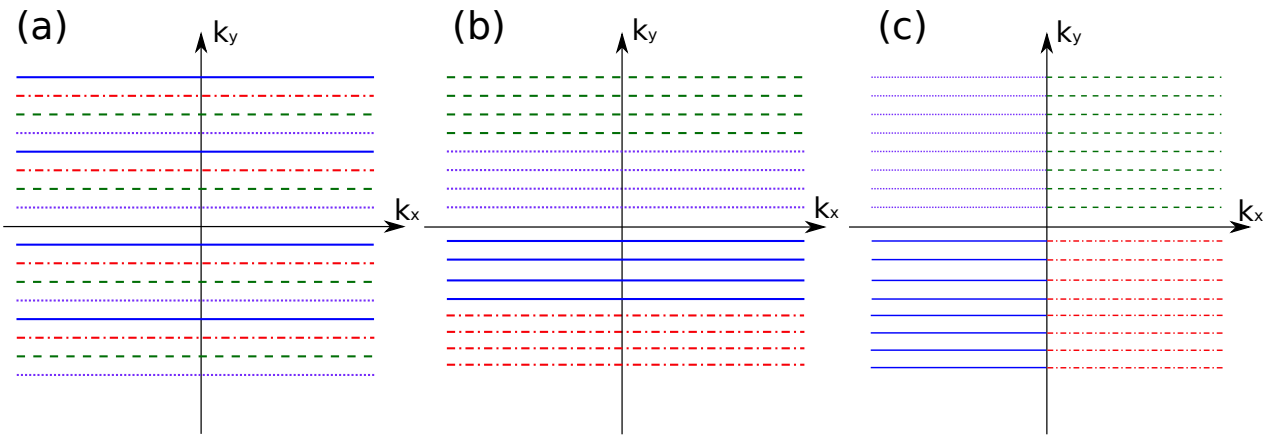

Figure 2.10: Different multishot sampling schemes using interleaved (a), sequential (b) and mosaic (c) strategies.

Application to dMRI

One particular case where ms–EPI has proven to be challenging is dMRI. Diffusion can be described as a mass transport process that results in molecular or particule mixing without bulk motion. Robert Brown observed that microscopic particles suspended in a fluid experiment a random motion. Later on, Einstein introduced the concept of "displacement distribution" to quantify the likelihood that a single particle traverses a certain displacement in a given timeframe. For instance, in the case of free diffusion, the displacement distribution is a Gaussian function whose standard deviation is determined by the diffusion coefficient.

Diffusion MRI aims at noninvasively quantifying the diffusion process of water molecules in living biological tissues such as the human brain white matter. By capturing the diffusion of water molecules it allows to infer the structure of the underlying biological tissues. In Fig.2.11 we show the typical pulsed gradient spin echo MR sequence proposed by Stejskal and Tanner (Stejskal and Tanner, 1965).

In anatomical structures such as the brain, the diffusion of water molecules is not free anymore, but restricted. As a consequence, more complex models are used in order to characterize the 3D displacement distribution. Likewise, this implies that we need to acquire multiple volumes with different diffusion-weighting with the aim of estimating the parameters that fully describe the distribution, the number of these volumes depending on the complexity of the chosen model. Most importantly, this notably increases the acquisition time for dMRI experiments. For this reason, using fast MRI sequences becomes vital and led into widely employing Single–Shot (ss) EPI as the standard choice for dMRI.

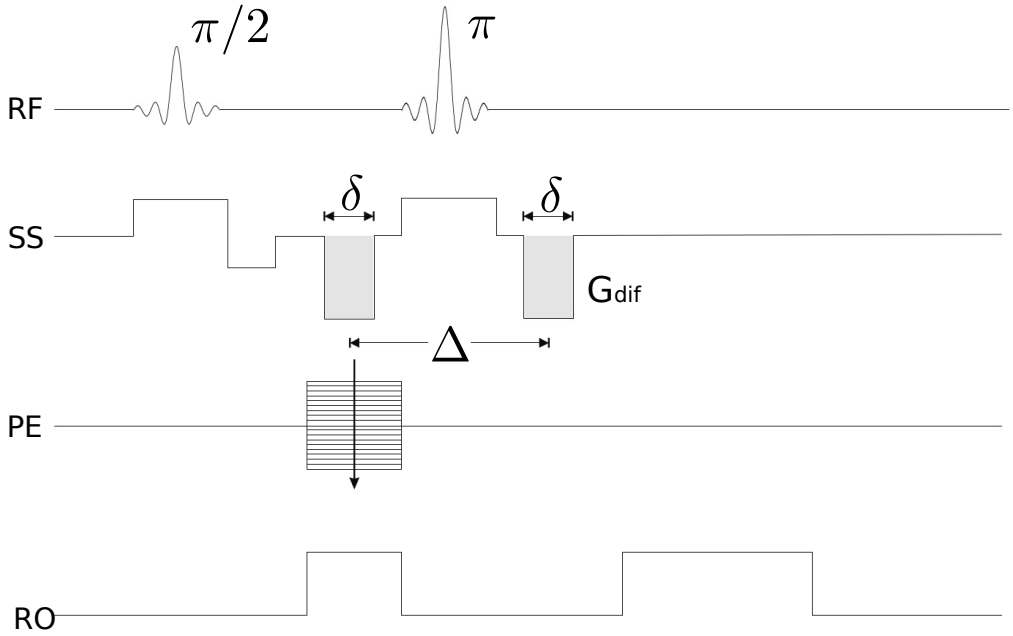

Figure 2.11: Conventional Stejskal–Tanner spin echo sequence with diffusion gradients (shaded) applied along the slice–slection dimension. *Figure obtained from (Brown et al., 2014).*

Motion–induced ghosting

As we have already explained, ss–EPI suffers from different artifacts that limit the spatial resolution we can achieve at a tolerable level of image degradation. In order to increase the spatial resolution, ms–EPI has become a sequence of interest, although its adoption has been challenged by a motion related artifact that is not present in ss–EPI. Since the diffusion experiment is designed to be sensitive to the motion of water molecules in the microscopic scale, even tiny movements of the patient will be captured as well. As we will prove later on, patient motion during the application of the diffusion gradients corrupts the phase of the image, or equivalently, modifies the trajectory of the sampled shot, as shown in Fig.2.12. In dMRI we are mostly interested in the magnitude of the image, so this has not been a limiting factor for ss–EPI. However, in ms–EPI the patient motion might vary from shot to shot, resulting in phase discrepancies that will introduce ghosting in the reconstructed image.

In order to successfully use a multi–shot strategy for dMRI, the phase discrepancies between different shots must be estimated and corrected for. For this purpose, different techniques have been developed; most of them are based on acquiring an extra portion of the center of the $\boldsymbol{k}$–space together with the data corresponding from this shot. From this low–frequency data, sampled at the same $\boldsymbol{k}$–space location for all the shots, the phase differences can be estimated. The main drawback with this approach is that it is time inefficient due to the data acquired only for calibration purposes, which motivated the development of self–calibrated techniques. Unfortunately, even though they use only imaging data for the estimation of the

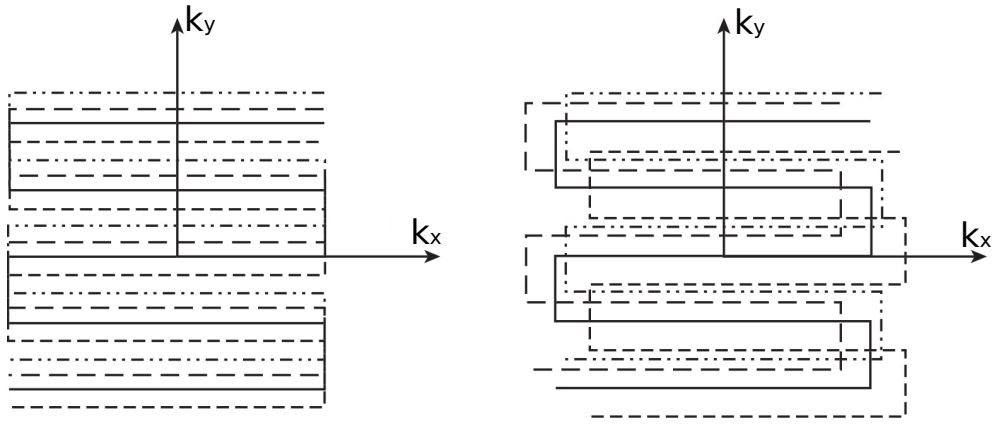

Figure 2.12: Example of a four–shot EPI sampling. (a) No motion occurs during the application of the gradients. (b) Different rigid–motion happens during the application of the diffusion gradients for each shot. The random phase errors introduced in the image–space can be modelled as linear and result in shifts in $\boldsymbol{k}$–space, so the interleaves are no longer equidistant from each other. As a consequence, a non–Uniform Fourier Transform would need to be applied once the information relative to the $\boldsymbol{k}$–space shifts is available. *Figure obtained from (K. Jones, 2011).*

phase maps, they require some degree of redundancy which does not completely remove the inefficiencies. Therefore, free–navigated alternatives have been studied, by exploiting some type of prior knowledge about these phase maps, such as smoothness, limited–support or low–rankness. Following this path, in this work we propose a free–navigated approach that assumes linear phase map corruption arising from bulk motion. Further, we extend this methodology to deal with scenarios of non–rigid motion where we model the phase–maps by means of B–splines.

CHAPTER

THREE

EXACT COMPUTATION OF G–FACTOR MAPS IN 2D GRAPPA

Abstract

Purpose: To exactly characterize the noise distribution and g-factor in 2D GRAPPA acquisitions with non-uniform undersampling patterns by a noise propagation analysis that operates directly in $\boldsymbol{k}$–space.

Theory and Methods: In order to provide a computationally efficient algorithm, we exploit the extensive symmetries and separability into independent blocks in the reconstruction steps to take into account the correlation between all the acquired $\boldsymbol{k}$–space samples. Monte Carlo simulations, phantom experiments and in-vivo tests were conducted to validate both the accuracy and feasibility of the proposed method, as well as to compare it to previous techniques.

Results: Non–uniform undersampling patterns (e.g.: containing a calibration region or based on a variable density sampling approach) cause noise characterization methods based on an image–space analysis to introduce errors. A direct $\boldsymbol{k}$–space based analysis is able to avoid these errors without increasing the computation time with respect to image–space analysis.

Conclusions: By operating directly in the $\boldsymbol{k}$–space, the proposed method is able to provide an exact characterization of noise for uniform and non–uniform undersampling patterns. By exploiting the symmetries and separability into independent blocks through the image reconstruction, the proposed method is able to overcome the computational challenges related to the very large size of the covariance matrices involved.

3.1 Introduction

Noise is an unavoidable source of degradation in MR signals. The principal source of noise is the subject itself, followed by electronic noise during the acquisition of the signal in the receiver chain (Aja Fernández and Vegas Sánchez-Ferrero, 2016). Noise degrades the visual quality of the reconstructed images, and complicates further post–processing techniques, such as segmentation, registration, fMRI analysis or numerical estimation of parameters. Accurate characterization of noise statistics is essential for many different tasks such as quality assurance (Krissian and Aja-Fernández, 2011; Aja-Fernández et al., 2013), protocol optimization (Thünberg and Zetterberg, 2007; Saritas et al., 2011), and tailoring of subsequent post–processing steps (Ghugre et al., 2006; Rabanillo et al., 2016a; Veraart et al., 2011).

As we have already mentioned, noise in the $\boldsymbol{k}$–space is usually assumed to be a zero-mean, spatially uncorrelated Independent and Identically Distributed (IID) complex Gaussian process for each coil, with equal variance in both the real and imaginary parts. If the data is acquired by several receiving coils, the multi-coil noise can be characterized by its covariance matrix. Since the noise distribution of the studied scenarios is well–known, by noise characterization we will refer to the estimation of the parameters that define this distribution (e.g.: the standard deviation of the zero-mean Gaussian case).

For linear image reconstructions, the propagation of noise in $\boldsymbol{k}$–space into the image–space (also known as $\boldsymbol{x}$–space) can be described by matrix operations. In the case of fully-sampled acquisitions, the IID noise behaviour is preserved when the data are transformed into the $\boldsymbol{x}$–space, due to the orthogonality of the iFT. However, in the presence of multiple coils and PI acceleration, reconstructed noise in the $\boldsymbol{x}$–space may show spatial heterogeneities. Pre–calibrated image–space PI methods, such as SENSE, proposed by (Pruessmann et al., 1999), allow an exact characterization of the noise model. Thus, in those methods where the reconstruction is performed in the image–space, a direct noise propagation analysis is computationally efficient, since only channel correlations need to be considered.

Alternatively, in PI methods where the reconstruction takes place in the $\boldsymbol{k}$–space, e.g.: GRAPPA acquisitions (Griswold et al., 2002), the reconstruction is done coil by coil in $\boldsymbol{k}$–space, using data from all the coils to reconstruct the signal from each individual coil. As a consequence, noise characterization becomes challenging, due to the introduction of correlations in the $\boldsymbol{k}$–space that propagate into the $\boldsymbol{x}$–space. A direct noise propagation analysis has been considered computationally infeasible due to the need to operate with very large covariance matrices (Robson et al., 2008).

An early approach to overcome this computational challenge is based on Monte–Carlo simulations, as proposed by (Robson et al., 2008). In Monte–Carlo based methods, noise maps are obtained by repeatedly corrupting the acquired data with synthetic noise properly

scaled and correlated. Unfortunately, Monte–Carlo methods are very time consuming and only provide a noisy estimation of the noise parameters. For this reason, there is significant interest in the development of analytical noise characterization in GRAPPA.

In order to avoid the need for large covariance matrices, several approximated methods have been proposed. In these methods, the problem is simplified by approximating the $\boldsymbol{k}$–space image reconstruction process (which can be viewed as a convolution) as a pixel-wise multiplication, as in $\boldsymbol{x}$–space (Breuer et al., 2009; Aja-Fernández et al., 2011). This reformulation provides computationally efficient noise characterization by avoiding the extensive $\boldsymbol{k}$–space correlations. However, as (Brau et al., 2008; Beatty and Brau, 2016) show, the GRAPPA reconstruction is only equivalent to a convolution when a uniform undersampling pattern is used. Importantly, reconstructions from non-uniform $\boldsymbol{k}$–space undersampling trajectories (e.g.: trajectories including a calibration region in the center of $\boldsymbol{k}$–space, the so-called ACS lines) can only be approximated as a region–by–region convolution, which introduces errors in the boundaries between the different regions.

Although these approximated methods have shown to work properly in multiple scenarios, they will introduce errors in the presence of non-uniformly undersampled $\boldsymbol{k}$–space trajectories. Therefore, development of an exact analytical method is highly desirable, and previous works have focused on this problem (Beatty and Brau, 2016). However, to the best of the authors' knowledge, an accurate and efficient analytical method has not been fully described in the literature.

To overcome the limitations of image–space methods, in this chapter we propose an efficient $\boldsymbol{k}$–space noise propagation analysis. In the proposed method, the noise is propagated through the reconstruction by accounting for all the correlations in the $\boldsymbol{k}$–space. The need to operate with very large covariance matrices is avoided by exploiting extensive symmetries and separability in each of the reconstruction steps. Thus, an exact and computationally efficient solution to the noise distribution in GRAPPA for every reconstructed channel is provided. This analysis is exact under the assumptions of stationarity and uncorrelation in the original undersampled $\boldsymbol{k}$–space acquisition. Furthermore, the analysis is accurate also for more complex $\boldsymbol{k}$–space undersampling, including variable density methods as in (Heidemann et al., 2007).

The resulting coil-by-coil images are usually combined into a magnitude composite image using either complex (Walsh et al., 2000) or magnitude based (Roemer et al., 1990) coil combination approaches. The coil-by-coil noise estimates here provided can subsequently be combined according to the selected coil combination approach, thus characterizing noise in the final coil-combined image. In addition, this method can also be used to evaluate the noise amplification (g-factor) in GRAPPA reconstructions as in (Breuer et al., 2009).

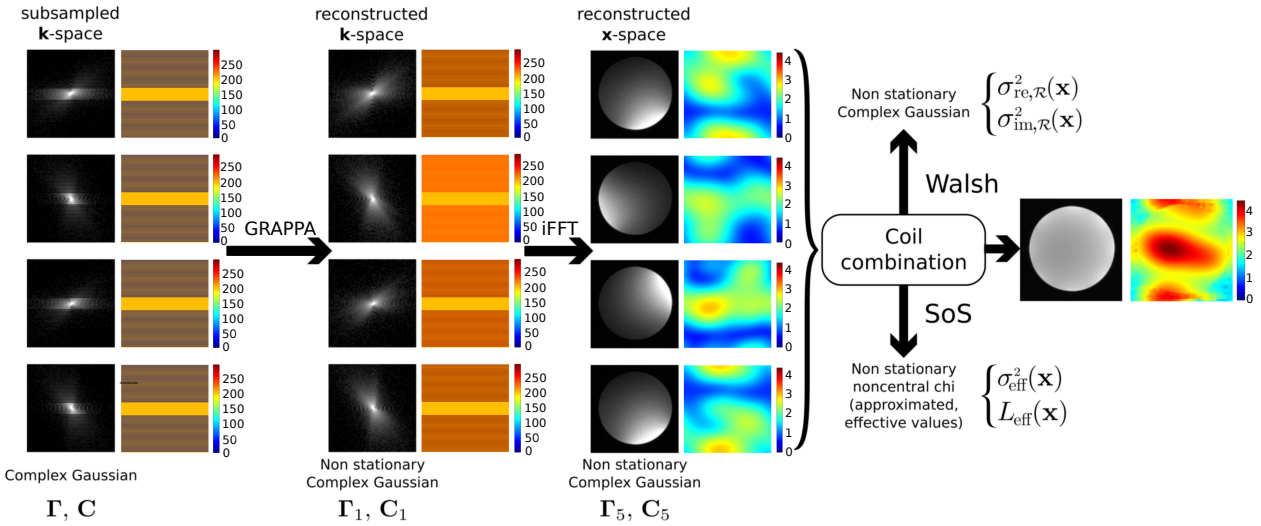

Figure 3.1: Noise models including the main noise parameters for each of the steps within the GRAPPA reconstruction pipeline. The noise distribution of the final composite image depends on the coil combination procedure used.

3.2 Theory

In order to characterize the noise in the final image, each of the steps of the GRAPPA processing must be properly characterized, see Fig.3.1. Initially, noise in the sampled $\boldsymbol{k}$-space is assumed to be a stationary complex Gaussian process. Due to the GRAPPA interpolation, noise in the reconstructed signal can also be modeled as Gaussian, but it becomes position–dependent, i.e., non-stationary. This noise is modeled by the covariance matrices containing the correlations between all the samples in the $\boldsymbol{k}$–space. The size of these matrices is given by $[N_p \cdot N_f \cdot L] \times [N_p \cdot N_f \cdot L]$, where N_p is the number of points in the phase-encoding direction and N_f is the number of points in the FE direction, i.e., the size of the $\boldsymbol{k}$–space to be reconstructed is given by $[N_p, N_f]$. Although theoretically feasible, the direct noise analysis becomes rapidly prohibitive due to the need to perform operations on very large matrices. To overcome this limitation, previously proposed methods have used models based on approximations of the reconstructed signal that simplify the problem. These approximated methods are described in the following section.

3.2.1 Image–space methods for noise characterization in GRAPPA

Both (Breuer et al., 2009) and (Aja-Fernández et al., 2011) have proposed a noise analysis in the image–space based on rewriting Eq.3.1 as a convolution

$$s_l^R(\boldsymbol{k}) = \sum_{m=1}^{L} s_m^S(\boldsymbol{k}) \circledast \mathsf{w}_m(l, \boldsymbol{k}), \tag{3.1}$$

where $\mathsf{w}_m(l, \boldsymbol{k})$ is a convolution kernel built from the reconstruction weights $\omega_m(l, \boldsymbol{c})$ in Eq.2.16, as explained in (Breuer et al., 2009; Brau et al., 2008). Following this definition and based on the Fourier properties, the GRAPPA reconstruction can be expressed as a pixel-wise product in the $\boldsymbol{x}$–space:

$$S_l^R(\boldsymbol{x}) = N_p \cdot N_f \cdot \sum_{m=1}^{L} S_m^S(\boldsymbol{x}) \cdot W_m(l, \boldsymbol{x}), \tag{3.2}$$

where $S_l^R(\boldsymbol{x})$ and $S_l^S(\boldsymbol{x})$ are the reconstructed and sampled signals at coil l-th in the $\boldsymbol{x}$–space, and $W_m(l, \boldsymbol{x})$ are the reconstruction kernels transformed into the $\boldsymbol{x}$–space by the 2D-iFT. We define the matrix $\mathbf{W}(\boldsymbol{x})$ containing all the weights to reconstruct a pixel across all the channels:

$$\mathbf{W}(\boldsymbol{x}) = \begin{pmatrix} W_{11}(\boldsymbol{x}) & \cdots & W_{1L}(\boldsymbol{x}) \\ \vdots & \ddots & \vdots \\ W_{L1}(\boldsymbol{x}) & \cdots & W_{LL}(\boldsymbol{x}) \end{pmatrix}. \tag{3.3}$$

In (Aja-Fernández et al., 2011), the authors consider that an undersampled GRAPPA acquisition is equivalent to zero-padding the missing lines. Furthermore, they assume identical and uncorrelated distributions for the imaginary and real components, characterized by $\boldsymbol{\Sigma}_k$. When the 2D-iFT is computed, we will obtain an aliased image for each channel along the phase–encoding dimension. Under the previous assumptions, this FT preserves spatial stationarity, but scales the noise level in the subsampled images with respect to the fully-sampled case. The noise covariance $\boldsymbol{\Sigma}_x^{\mathcal{S}}$ in the aliased images will be given by:

$$\boldsymbol{\Sigma}_x^{\mathcal{S}} = \frac{N_{\text{acq}}/N_p}{N_p \cdot N_f} \cdot \boldsymbol{\Sigma}_k = \frac{1}{R_{\text{eff}}} \cdot \boldsymbol{\Sigma}_x^{\text{full}}, \tag{3.4}$$

where N_{acq} is the number of acquired lines in the phase-encoding direction and $R_{\text{eff}} = \frac{N_p}{N_{\text{acq}}}$ is the effective acceleration (ratio between the number of acquired lines and the number of lines in the reconstructed $\boldsymbol{k}$–space). $\boldsymbol{\Sigma}_x^{\text{full}} = \frac{1}{N_p \cdot N_f}\boldsymbol{\Sigma}_k$ denotes the covariance matrix in the image–space for a fully-sampled acquisition, which preserves stationarity and spatial uncorrelation due to the orthogonality of the iFT, as described by (Aja-Fernández et al., 2011).

The product in Eq.3.2 will introduce non-stationarity since the transformed kernels $W_m(l, \boldsymbol{x})$ vary across the image. The noise variance $\sigma^2_{\mathcal{R},l}(\boldsymbol{x})$ for coil l in the $\boldsymbol{x}$–space becomes:

$$\sigma^2_{\mathcal{R},l}(\boldsymbol{x}) = \frac{1}{R_{\text{eff}}} \left[\mathbf{W}(\boldsymbol{x}) \cdot \boldsymbol{\Sigma}_x^{\text{full}} \cdot \mathbf{W}^H(\boldsymbol{x})\right]_{ll}. \tag{3.5}$$

In order to quantify the noise amplification due to PI, we will use the definition of the g-factor given in (Breuer et al., 2009):

$$g_l(\boldsymbol{x}) = \frac{\text{SNR}^l_{\text{full}}}{\text{SNR}^l_{\text{acc}}} = \frac{\sqrt{[\mathbf{W}(\boldsymbol{x}) \cdot \boldsymbol{\Sigma}_x^{\text{full}} \cdot \mathbf{W}^H(\boldsymbol{x})]_{ll}}}{\sqrt{R_{\text{eff}}} \cdot \sqrt{[\boldsymbol{\Sigma}_x^{\text{full}}]_{ll}}}, \tag{3.6}$$

where $\text{SNR}^l_{\text{full}}$ and $\text{SNR}^l_{\text{acc}}$ denote the signal-to-noise ratio in the fully sampled image and in the subsampled image after reconstruction for each coil l, respectively.

Finally, coil combination can be performed using the SoS as in (Aja-Fernández et al., 2011), producing a signal that can be approximated using a non–central χ distribution (see Fig.3.1). Alternatively, a complex linear coil combination can be used, as proposed by (Walsh et al., 2000). In this case, the information from each coil is combined using a vector $\vec{\mathbf{m}}(\boldsymbol{x}) = [m_1(\boldsymbol{x}) \cdots m_L(\boldsymbol{x})]^T$:

$$S_T(\boldsymbol{x}) = \vec{\mathbf{m}}(\boldsymbol{x}) \cdot \vec{\mathbf{S}^R}(\boldsymbol{x}) = \sum_{l=1}^{L} m_l(\boldsymbol{x}) S_l^R(\boldsymbol{x}). \tag{3.7}$$

For complex linear coil combination, the noise in the composite image is complex Gaussian , with variance and g-factor as follows:

$$\sigma_T^2(\boldsymbol{x}) = \frac{1}{R_{\text{eff}}}\ \vec{\mathbf{m}}(\boldsymbol{x}) \cdot \mathbf{W}(\boldsymbol{x}) \cdot \boldsymbol{\Sigma}_x^{\text{full}} \cdot \mathbf{W}^H(\boldsymbol{x}) \cdot \vec{\mathbf{m}}^H(\boldsymbol{x}), \tag{3.8}$$

$$g_T(\boldsymbol{x}) = \frac{\sqrt{\vec{\mathbf{m}}(\boldsymbol{x}) \cdot \mathbf{W}(\boldsymbol{x}) \cdot \boldsymbol{\Sigma}_x^{\text{full}} \cdot \mathbf{W}^H(\boldsymbol{x}) \cdot \vec{\mathbf{m}}^H(\boldsymbol{x})}}{\sqrt{R_{\text{eff}}} \cdot \sqrt{\vec{\mathbf{m}}(\boldsymbol{x}) \cdot \boldsymbol{\Sigma}_x^{\text{full}} \cdot \vec{\mathbf{m}}^H(\boldsymbol{x})}}. \tag{3.9}$$

This equivalence between the reconstruction in both spaces holds true when the undersampling pattern is uniform across the $\boldsymbol{k}$–space, as in the case when the weights are obtained from a separate pre-scan. However, if the ACS region is acquired within the image, or a Variable Density GRAPPA (VD-GRAPPA) acquisition scheme, as proposed in (Heidemann et al., 2007), is used, it is not possible to reconstruct the whole $\boldsymbol{k}$–space as a convolution with a single kernel Eq.3.1 (Brau et al., 2008; Beatty and Brau, 2016).

Authors in (Breuer et al., 2009) propose a model based on splitting the $\boldsymbol{k}$–space into M statistacally uncorrelated regions, where each region covers a fraction f_m of the $\boldsymbol{k}$–space and is uniformly undersampled at rate R_m. Each $\boldsymbol{k}$–space region is then reconstructed with its own kernel, resulting in the following expressions for the variance and g-factor in the final image:

$$\sigma_T^2(\boldsymbol{x}) = \sum_{m=1}^{M} \frac{f_m}{R_m} \vec{\mathbf{m}}(\boldsymbol{x}) \cdot \mathbf{W}_m(\boldsymbol{x}) \cdot \boldsymbol{\Sigma}_x^{\text{full}} \cdot \mathbf{W}_m^H(\boldsymbol{x}) \cdot \vec{\mathbf{m}}^H(\boldsymbol{x}), \tag{3.10}$$

$$g_T(\boldsymbol{x}) = \sum_{m=1}^{M} \frac{f_m}{R_m} \frac{\sqrt{\vec{\mathbf{m}}(\boldsymbol{x}) \cdot \mathbf{W}_m(\boldsymbol{x}) \cdot \boldsymbol{\Sigma}_x^{\text{full}} \cdot \mathbf{W}_m^H(\boldsymbol{x}) \cdot \vec{\mathbf{m}}^H(\boldsymbol{x})}}{\sqrt{R_{\text{eff}}} \cdot \sqrt{\vec{\mathbf{m}}(\boldsymbol{x}) \cdot \boldsymbol{\Sigma}_x^{\text{full}} \cdot \vec{\mathbf{m}}^H(\boldsymbol{x})}} \tag{3.11}$$

where $\mathbf{W}_m(\boldsymbol{x})$ is the kernel used in the region m.

Limitations of image–space methods

Although image–space methods have shown to be useful in practical situations, they have fundamental limitations, particularly in the presence of a calibration region or non-uniform

undersampling. One of the advantages of GRAPPA is precisely its ability to accomodate variable density undersampling schemes, where regions with different acceleration factors need to be reconstructed using different kernels. In this scenario assuming that the different $\boldsymbol{k}$–space regions are statistically uncorrelated, image–space methods reconstruct the coil images in several steps. First, the $\boldsymbol{k}$–space is split into the differently undersampled regions. Then, each region is transformed into the image–space and subsequently reconstructed as a product with the 2D-iFT of its kernel, generating an image per region. Finally, these region-specific images are added to obtain the final coil image.

This methodology produces two sources of error. First, the reconstruction in the $\boldsymbol{x}$–space is not exactly equivalent to the reconstruction in the $\boldsymbol{k}$–space. When a kernel is applied to a filtered region, the support of the output overlaps with the adjacent regions, introducing a residual error. Second, when different acceleration factors are used, the acquired lines in the boundaries that separate every pair of regions are used to reconstruct both regions. Consequently, the two regions will be correlated, contradicting the initial assumption. For these reasons, the image–space analysis (including Eq.3.10 and Eq,3.11) is not fully correct. In order to mathematically prove these errors, we provide the Appendix.A.

3.2.2 k–space method for noise characterization in GRAPPA

For the sake of simplicity, our analysis considers: (1) GRAPPA weights are non-stochastic, i.e., they are independent of the noise realization. This is strictly true only if they are estimated from an independent acquisition; however we expect this assumption to be a good approximation even in self-calibrated acquisitions, given the typical overdetermination of GRAPPA weight estimation; (2) acquired points in the $\boldsymbol{k}$–space are IID as in (Henkelman, 1985; McVeigh et al., 1985; Kellman and McVeigh, 2005).*

In order to characterize the noise in the final composite image, each of the steps of the GRAPPA processing must be properly characterized, see Fig.3.1:

k–space Interpolation

The reconstruction for every point $\vec{\mathbf{s}}^R(\boldsymbol{k})$ can be expressed as:

$$\vec{\mathbf{s}}^R(\boldsymbol{k}) = \sum_{\mathbf{c}\in\eta(\boldsymbol{k})} \boldsymbol{\mathcal{W}}^{\boldsymbol{k}}_{\mathbf{c}} \cdot \vec{\mathbf{s}}^S(\mathbf{c}), \tag{3.12}$$

*If the noise shows temporal correlation along the FE direction, our analysis can still be carried out. It would just increase the width of the neighborhood a reference point correlates with, as described in subsection "$\boldsymbol{k}$–space Interpolation".

where the vector $\vec{\mathbf{s}}^S(\mathbf{c}) = \left[s_1^S(\mathbf{c}) \cdots s_L^S(\mathbf{c})\right]$ contains the $\boldsymbol{k}$–space acquired data within the neighborhood $\eta(\boldsymbol{k})$. Also, $\boldsymbol{\mathcal{W}}_{\mathbf{c}}^{\boldsymbol{k}}$ is an $L \times L$ matrix in which the l-th row contains the GRAPPA weights $\vec{\boldsymbol{\omega}}^{\boldsymbol{k}}(l, \mathbf{c}) = [\omega_1^{\boldsymbol{k}}(l, \mathbf{c}) \ldots \omega_L^{\boldsymbol{k}}(l, \mathbf{c})]$ from Eq.2.16 associated to location $\mathbf{c}$ to reconstruct the l-th coil element of the $\boldsymbol{k}$–space vector $\vec{\mathbf{s}}^R(\boldsymbol{k})$.

Notice that although both matrices $\boldsymbol{\mathcal{W}}_{\mathbf{c}}^{\boldsymbol{k}}$ is an $L \times L$ and $\mathbf{W}(\boldsymbol{x})$ are related to the GRAPPA kernels, they are computed differently. $\mathbf{W}(\boldsymbol{x})$ is computed as the iFT of the kernel, as in (Aja-Fernández et al., 2011). It is applied in the image–space and it is position dependent. On the other hand, $\boldsymbol{\mathcal{W}}_{\mathbf{c}}^{\boldsymbol{k}}$ is an $L \times L$ operates in $\boldsymbol{k}$–space and denotes directly the GRAPPA weights without any transformation. Since we consider the general case of using multiple kernels, this matrix can change through $\boldsymbol{k}$–space. Consequently, $\boldsymbol{\mathcal{W}}_{\mathbf{c}}^{\boldsymbol{k}}$ depends on both the point reconstructed $\boldsymbol{k}$ and the point considered for the reconstruction $\mathbf{c}$.

Based on this reconstruction, the acquisition noise $\vec{\mathbf{n}}^S(\boldsymbol{k})$ will propagate into the reconstructed $\boldsymbol{k}$–space as follows:

$$\vec{\mathbf{n}}^R(\boldsymbol{k}) = \sum_{\mathbf{c} \in \eta(\boldsymbol{k})} \boldsymbol{\mathcal{W}}_{\mathbf{c}}^{\boldsymbol{k}} \cdot \vec{\mathbf{n}}^S(\mathbf{c}), \tag{3.13}$$

where the vectors $\vec{\mathbf{n}}^S(\mathbf{c})$ and $\vec{\mathbf{n}}^R(\mathbf{c})$ are the noise in the acquired and reconstructed $\boldsymbol{k}$–space, respectively.

Further, GRAPPA interpolation introduces correlation between two arbitrary points in the reconstructed $\boldsymbol{k}$-space. After this first stage of the reconstruction, the resulting noise correlation matrices $\boldsymbol{\Gamma}_1$ and $\mathbf{C}_1$ (where the index 1 simply indicates the order of appearance of these correlation matrices in the present derivation) can be expressed as (see Fig.4.2.a)

$$\begin{aligned}
\boldsymbol{\Gamma}_1^{ij} &= E\{\vec{\mathbf{n}}^R(\boldsymbol{k}_i) \cdot \vec{\mathbf{n}}^R(\boldsymbol{k}_j)^H\} \\
&= E\left\{ \left(\sum_{\mathbf{c}_i \in \eta_i} \boldsymbol{\mathcal{W}}_{\mathbf{c}_i}^{\boldsymbol{k}_i} \cdot \vec{\mathbf{s}}^S(\mathbf{c}_i) \right) \left(\sum_{\mathbf{c}_j \in \eta_j} \boldsymbol{\mathcal{W}}_{\mathbf{c}_j}^{\boldsymbol{k}_j} \cdot \vec{\mathbf{s}}^S(\mathbf{c}_j) \right)^H \right\} \\
&= \sum_{\mathbf{c}_i \in \eta_i} \sum_{\mathbf{c}_j \in \eta_j} \boldsymbol{\mathcal{W}}_{\mathbf{c}_i}^{\boldsymbol{k}_i} \cdot E\left\{ \vec{\mathbf{s}}^S(\mathbf{c}_i) \cdot \vec{\mathbf{s}}^S(\mathbf{c}_j)^H \right\} \cdot {\boldsymbol{\mathcal{W}}_{\mathbf{c}_j}^{\boldsymbol{k}_j}}^H \\
&= \sum_{\mathbf{c} \in (\eta_i \cap \eta_j)} \boldsymbol{\mathcal{W}}_{\mathbf{c}_i}^{\boldsymbol{k}_i} \cdot \boldsymbol{\Gamma}_a \cdot {\boldsymbol{\mathcal{W}}_{\mathbf{c}_j}^{\boldsymbol{k}_i}}^H
\end{aligned} \tag{3.14}$$

and equivalently

$$\mathbf{C}_1^{ij} = E\{\vec{\mathbf{n}}^R(\boldsymbol{k}_i) \cdot \vec{\mathbf{n}}^R(\boldsymbol{k}_j)^T\} = \sum_{\mathbf{c} \in (\eta_i \cap \eta_j)} \boldsymbol{\mathcal{W}}_{\mathbf{c}_i}^{\boldsymbol{k}_i} \cdot \mathbf{C}_a \cdot {\boldsymbol{\mathcal{W}}_{\mathbf{c}_j}^{\boldsymbol{k}_i}}^T. \tag{3.15}$$

This implies that a point in the reconstructed $\boldsymbol{k}$–space correlates with all the points with overlapping neighborhoods. The points in the acquired lines are left untouched, and thus their kernel only contains one point (themselves), so $\boldsymbol{\mathcal{W}}_{\mathbf{c}}^{\boldsymbol{k}}$ is the identity matrix for $\boldsymbol{k} = \mathbf{c}$ and the null matrix otherwise. Fig.3.2 shows an example of the correlations introduced by GRAPPA reconstruction.

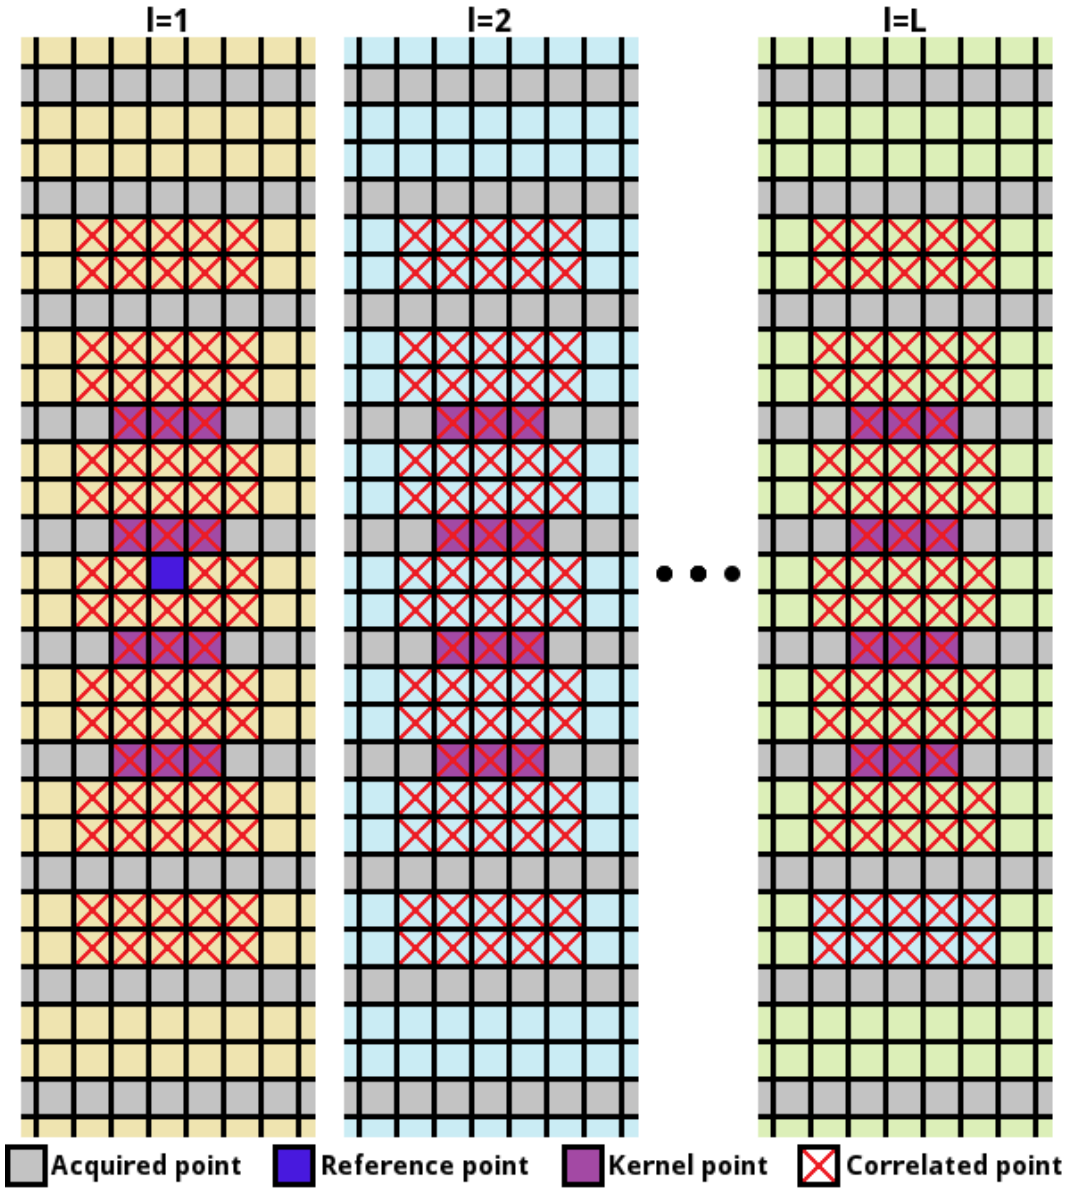

Figure 3.2: Example of correlations between a point and its neighborhood for the case of 2D-GRAPPA with acceleration factor 3 and kernel size 4×3. A missing point in $\boldsymbol{k}$–space is estimated as an interpolation of its neighborhood. Consequently, it will only be correlated with a reduced number of points in $\boldsymbol{k}$–space consisting of the points in its neighborhod and those points whose neighborhood overlaps with its own neighborhood.

We define the size of the $\boldsymbol{k}$–space to be reconstructed as $[N_p, N_f]$, where p refers to the phase-encoding direction and f to the FE direction. Similarly, if we define the GRAPPA kernel size as $[K_p, K_f]$, a reconstructed point correlates in the FE direction with points in $2K_f - 1$ columns, where a column refers to the data for a particular FE value and a row refers to the data for a particular phase-encoding value. In the phase-encoding direction, the number of points with which a reconstructed point correlates depends on the row since the kernel may vary from row to row. Let us now stack a set of consecutive $2K_f - 1$ columns into a vector; the column in position K_f in the stack will be referred to as reference column. If we now stack L such vectors (one from each coil), we construct a $N_p \cdot (2K_f - 1) \cdot L$-component vector, with correlation matrix defined as follows (for simplicity we only keep track of the evolution of $\boldsymbol{\Gamma}$):

$$\boldsymbol{\Gamma}_2 = \begin{pmatrix} \mathbf{B}_{11} & \mathbf{B}_{12} & \cdots & \mathbf{B}_{1L} \\ \mathbf{B}_{21} & \mathbf{B}_{22} & \cdots & \mathbf{B}_{2L} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{B}_{L1} & \mathbf{B}_{L2} & \cdots & \mathbf{B}_{LL} \end{pmatrix}, \tag{3.16}$$

where $\mathbf{B}_{ij}$ contains the correlations of the reference column and its $2K_f - 2$ neighboring columns in the i-th coil and this equivalent vector in the j-th coil. This matrix presents Hermitian symmetry, $\mathbf{B}_{ij} = \mathbf{B}_{ji}^H$; thus, we only need to compute $L \cdot (L+1)/2$ blocks.

Since GRAPPA uses the same kernel for all the reconstructed points in a row, under the stationarity assumption, the correlation between a reference column and its surrounding columns is the same, independently of the column picked for reference. This gives rise to a block-Toeplitz structure for each of the blocks $\mathbf{B}_{ij}$ in Eq.3.16. Specifically, if the reference column has index K_f, its left hand side columns are indexed within the interval $1 \leq i \leq K_f - 1$ and its right hand side columns are indexed within $K_f + 1 \leq i \leq 2K_f - 1$, then the first column of $\mathbf{B}_{ij}$ is defined as

$$\begin{bmatrix} \mathbf{b}_{ij}^0 & \mathbf{b}_{ij}^1 & \cdots & \mathbf{b}_{ij}^{K_f-1} & \overline{\mathbf{0}} & \cdots & \overline{\mathbf{0}} \end{bmatrix}^T, \tag{3.17}$$

where $\mathbf{b}_{ij}^m$ is the $N_p \times N_p$ covariance matrix of two columns, the subtraction of the indices of which equals m, and $\overline{\mathbf{0}}$ denotes a null matrix of dimensions $N_p \times N_p$. These covariances are obtained by selecting the appropriate components of $\boldsymbol{\Gamma}_1$ in Eq.3.14. As for the first row,

$$\begin{bmatrix} \mathbf{b}_{ij}^0 & \mathbf{b}_{ij}^{-1} & \cdots & \mathbf{b}_{ij}^{1-K_f} & \overline{\mathbf{0}} & \cdots & \overline{\mathbf{0}} \end{bmatrix}. \tag{3.18}$$

Notice that we only need to compute $2K_f - 1$ sub-blocks due to the block-Toeplitz structure as well as the presence of null correlations known beforehand. Unfortunately no conjugate symmetry generally holds for these sub-blocks, but only for $i = j$. Therefore, the overall number of sub-blocks needed to build Eq.3.16 is at most $(2K_f - 1) \cdot L \cdot (L+1)/2$.

Column iFT

After reconstructing the $\boldsymbol{k}$–space, data are transformed into the image–space by a 2D-iFT, which can be decomposed into two 1D-iFT. Computing the 1D-iFT along the phase-encoding direction preserves the number of points with which a reference point correlates across the FE direction.

The 1D-iFT can be expressed as a matrix operation with matrix $\mathbf{F}_I$. For our vector stack, the 1D-iFT of each column for each coil can be obtained by the product with:

$$\boldsymbol{\mathcal{F}}_I = \begin{pmatrix} \mathbf{F}_I & \overline{\mathbf{0}} & \cdots & \overline{\mathbf{0}} \\ \overline{\mathbf{0}} & \mathbf{F}_I & \cdots & \overline{\mathbf{0}} \\ \vdots & \vdots & \ddots & \vdots \\ \overline{\mathbf{0}} & \overline{\mathbf{0}} & \cdots & \mathbf{F}_I \end{pmatrix}. \tag{3.19}$$

The correlation matrix of the transformed vector is straightforwardly obtained as

$$\boldsymbol{\Gamma}_3 = \boldsymbol{\mathcal{F}}_I \cdot \boldsymbol{\Gamma}_2 \cdot \boldsymbol{\mathcal{F}}_I^H = \begin{pmatrix} \mathbf{D}_{11} & \mathbf{D}_{12} & \cdots & \mathbf{D}_{1L} \\ \mathbf{D}_{21} & \mathbf{D}_{22} & \cdots & \mathbf{D}_{2L} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{D}_{L1} & \mathbf{D}_{L2} & \cdots & \mathbf{D}_{LL} \end{pmatrix}, \tag{3.20}$$

where each $\mathbf{D}_{ij}$ shows the same block-Toeplitz structure as $\mathbf{B}_{ij}$, i.e., just a replacement of $\mathbf{b}$ by $\mathbf{d}$ is needed in Eq.3.17 and Eq.3.18 to build $\mathbf{D}_{ij}$, due to the fact that

$$\mathbf{d}_{ij}^m = \mathbf{F}_I \cdot \mathbf{b}_{ij}^m \cdot \mathbf{F}_I^H. \tag{3.21}$$

Once again, the number of sub-blocks $\mathbf{d}_{ij}$ needed to build $\mathbf{D}_{ij}$ in Eq.3.20 is given by $(2K_f - 1) \cdot L \cdot (L + 1)/2$, i.e., it coincides with the number needed to build Eq.3.16 because of equality in Eq.3.21.

Row iFT

To complete the 2D-iFT a second 1D-iFT is done along the row dimension. The correlation within each row suffices to obtain the noise maps. It is worth mentioning that this analysis generalizes to calculate the correlations between any pair of rows directly.

First, we create a column vector by stacking a selected row (set up as a column) in $\boldsymbol{k}$–space with the same rows in the other coils; this vector will have $N_f \cdot L$ components. The correlation

matrix of this vector can be defined as follows, where the dependance with y indicates that this process is done for every row.

$$\mathbf{\Gamma}_4(y) = \begin{pmatrix} \mathbf{G}_{11} & \mathbf{G}_{12} & \cdots & \mathbf{G}_{1L} \\ \mathbf{G}_{21} & \mathbf{G}_{22} & \cdots & \mathbf{G}_{2L} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{G}_{L1} & \mathbf{G}_{L2} & \cdots & \mathbf{G}_{LL} \end{pmatrix}, \tag{3.22}$$

where every block $\mathbf{G}_{lm}$ has dimensions $N_f \times N_f$. and contains the cross-correlations of the selected row in coils l and m within the hybrid space (k_f, y). These correlations are obtained by selecting the appropriate components in Eq.3.20. Assuming GRAPPA performs a circular interpolation, this will result in a cyclical structure along the FE direction. Specifically, the first column of block $\mathbf{G}_{ij}$ is defined as

$$\vec{\mathbf{g}_{ij}} = \left[g_{ij}^0,\ g_{ij}^1,\ \cdots,\ g_{ij}^{K_f-1},\ 0,\ \cdots,\ 0,\ g_{ij}^{1-K_f},\ \cdots,\ g_{ij}^{-1}\right]^T, \tag{3.23}$$

with $N_f - (2K_f - 1)$ zeroes and values g_{ij}^m are taken from the components in sub-block $\mathbf{d}_{ij}^m$ in Eq.3.21 that correspond to the selected row. Then, the j-th row of $\mathbf{G}_{ij}$, $2 \leq j \leq N_f - 1$, is obtained as a rightward circular shift of the row $j-1$, which results in the blocks $\mathbf{G}_{lm}$ being circulant matrices.

Computing the row 1D-iFT provides a correlation matrix given by:

$$\mathbf{\Gamma}_5(y) = \boldsymbol{\mathcal{F}}_I \cdot \mathbf{\Gamma}_4(y) \cdot \boldsymbol{\mathcal{F}}_I^H = \begin{pmatrix} \mathbf{H}_{11} & \mathbf{H}_{12} & \cdots & \mathbf{H}_{1L} \\ \mathbf{H}_{21} & \mathbf{H}_{22} & \cdots & \mathbf{H}_{2L} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{H}_{L1} & \mathbf{H}_{L2} & \cdots & \mathbf{H}_{LL} \end{pmatrix}, \tag{3.24}$$

where the blocks $\mathbf{H}_{ij}$ are:

$$\mathbf{H}_{ij} = \mathbf{F_I} \cdot \mathbf{G}_{ij} \cdot \mathbf{F_I}^H. \tag{3.25}$$

Taking into account that circulant matrices are diagonalized by the FT (see (Davis, 1979)), Eq.3.25 is simplified to:

$$\mathbf{H}_{ij} = \mathbf{F}_I \cdot \mathbf{G}_{ij} \cdot \mathbf{F}_I^H = \mathrm{diag}(\mathbf{F}_I \cdot \vec{\mathbf{g}_{ij}}). \tag{3.26}$$

Again, we have Hermitian symmetry, so we only need to compute $[N_f \cdot L \cdot (L+1)/2]$ 1D-iFT for this step.

Coil combination

A final composite image is obtained by merging the data from every coil into a single image. This can be done using the SoS as in (Aja-Fernández et al., 2011), which produces a signal

that can be approximated by a non–central χ distribution (see Fig.3.1), or, alternatively, in a more general way as a properly weighted linear combination as described by (Roemer et al., 1990). In this chapter we will use the linear combination proposed in (Walsh et al., 2000), since it allows the reconstruction to be written as a linear matrix operation. It combines the information from each coil using a vector $\vec{\mathbf{m}}(\boldsymbol{x}) = [m_1(\boldsymbol{x}) \cdots m_L(\boldsymbol{x})]^T$, where the $\boldsymbol{x}$ dependance indicates that the operation is pixel–wise:

$$S_T(\boldsymbol{x}) = \vec{\mathbf{m}}(\boldsymbol{x}) \cdot \vec{\mathbf{S}^R}(\boldsymbol{x}) = \sum_{l=1}^{L} m_l(\boldsymbol{x}) S_l^R(\boldsymbol{x}). \tag{3.27}$$

In order to characterize the distribution in every final reconstructed pixel, we first need to obtain the correlation matrices for each pixel $\vec{\mathbf{p}}$ along the coil dimension. The matrix $\mathbf{\Gamma}_5$ contains the correlations between each point in a row and all the points in that row in all the coils. The 1D-iFT in the FE direction introduces non-stationarity along this dimension, so we need to proceed pixel by pixel and extract the correlation matrices $\mathbf{\Gamma}_6$ and $\mathbf{C}_6$ for a selected position, which consists of picking the corresponding entry in the diagonal of $\mathbf{\Gamma}_5$ or $\mathbf{C}_5$.

If the complex images are linearly combined as in Eq.3.27, the final image preserves the Gaussian behavior (see Fig.3.1), although it presents spatial correlations and non-stationarity. For every pixel in the composite image, we can define

$$\Gamma_7(\boldsymbol{x}) = \vec{\mathbf{m}}_l(\boldsymbol{x}) \cdot \mathbf{\Gamma}_6(\boldsymbol{x}) \cdot \vec{\mathbf{m}}_l^H(\boldsymbol{x}). \tag{3.28}$$

The variance of noise for the real and imaginary components can be calculated by:

$$\begin{aligned} \sigma^2_{\text{re},\mathcal{R}}(\boldsymbol{x}) &= \frac{1}{2}\mathbb{R}e\left\{\Gamma_7 + C_7\right\}, \\ \sigma^2_{\text{im},\mathcal{R}}(\boldsymbol{x}) &= \frac{1}{2}\mathbb{R}e\left\{\Gamma_7 - C_7\right\}, \end{aligned} \tag{3.29}$$

where $\mathcal{R}$ indicates that this is the noise in the reconstructed subsampled image. Note that correlation between real and imaginary components for a pixel can exist, which are computed as:

$$\begin{aligned} \sigma^2_{\text{re-im},\mathcal{R}}(\boldsymbol{x}) &= \frac{1}{2}\mathbb{I}m\left\{C_7 - \Gamma_7\right\}, \\ \sigma^2_{\text{im-re},\mathcal{R}}(\boldsymbol{x}) &= \frac{1}{2}\mathbb{I}m\left\{\Gamma_7 + C_7\right\}. \end{aligned} \tag{3.30}$$

Finally, the g-factor map is derived from the previous equation defining an *average* variance $\sigma^2_{\mathcal{R}}(\boldsymbol{x})$

$$\sigma^2_{\mathcal{R}}(\boldsymbol{x}) = \frac{\sigma^2_{\text{re},\mathcal{R}}(\boldsymbol{x}) + \sigma^2_{\text{im},\mathcal{R}}(\boldsymbol{x})}{2}, \tag{3.31}$$

$$g_{\mathcal{R}}(\boldsymbol{x}) = \frac{\sqrt{\sigma^2_{\mathrm{re},\mathcal{R}}(\boldsymbol{x}) + \sigma^2_{\mathrm{im},\mathcal{R}}(\boldsymbol{x})}}{\sqrt{R_{\mathrm{eff}}} \cdot \sqrt{\vec{\mathbf{m}}(\boldsymbol{x}) \cdot \left(\boldsymbol{\Sigma}^{\mathrm{full}}_{x,\mathrm{re}} + \boldsymbol{\Sigma}^{\mathrm{full}}_{x,\mathrm{im}}\right) \cdot \vec{\mathbf{m}}^H(\boldsymbol{x})}}, \tag{3.32}$$

where $\boldsymbol{\Sigma}^{\mathrm{full}}_{x,\mathrm{re}}$ and $\boldsymbol{\Sigma}^{\mathrm{full}}_{x,\mathrm{im}}$ refer to the covariance matrices for the real and imaginary parts in the coil images when the $\boldsymbol{k}$–space is fully sampled.

Although only results for the linear combination in Eq.3.27 have been studied here, the application to other coil combination methods (e.g.: SoS) is straightforward once the covariance matrices $\boldsymbol{\Gamma}_6$ and $\mathbf{C}_6$ are calculated.

Summary of the procedure

The procedure to obtain $\boldsymbol{\Gamma}$ matrices is graphically depicted in Fig.4.2 and it can be summarized as follows ($\mathbf{C}$ matrices are obtained similarly):

1. Calculate $\mathbf{b}^m_{ij}$; the number of such sub-blocks is $(2K_f - 1) \cdot L \cdot (L+1)/2$, and their dimension is $N_p \times N_p$. Each entry in these sub-blocks is obtained from Eq.3.14, by choosing the appropriate components.

2. Calculate $\mathbf{d}^m_{ij}$ using Eq.3.21.

3. Calculate $\vec{\mathbf{g}_{ij}}$ in Eq.3.23 by choosing the appropriate components of $\mathbf{d}^m_{ij}$.

4. Calculate $\mathbf{H}_{ij}$ using Eq.3.26.

5. Create matrix $\boldsymbol{\Gamma}_6$ by selecting a pixel index in the row, say l, then $[\boldsymbol{\Gamma}_6]_{(i,j)} = \mathbf{H}_{i,j}(l,l)$.

6. Apply Eq.4.13, Eq.3.31 and Eq.3.32 to obtain pixel-wise noise characterization.

Fig.4.2 shows an outline of the noise propagation throughout the entire GRAPPA reconstruction. The image is acquired through different channels and noise can be considered stationary across the sampled $\boldsymbol{k}$–space for each channel, although correlations between channels may exist. GRAPPA then reconstructs the $\boldsymbol{k}$–space for each channel interpolating the missing $\boldsymbol{k}$–space areas, an operation that introduces spatial non-stationarity both in the $\boldsymbol{k}$–space and in the image–space, as well as correlations between adjacent $\boldsymbol{k}$-space locations. Last, the different channels need to be combined in order to provide the final image, which will also show spatial non-stationarity.

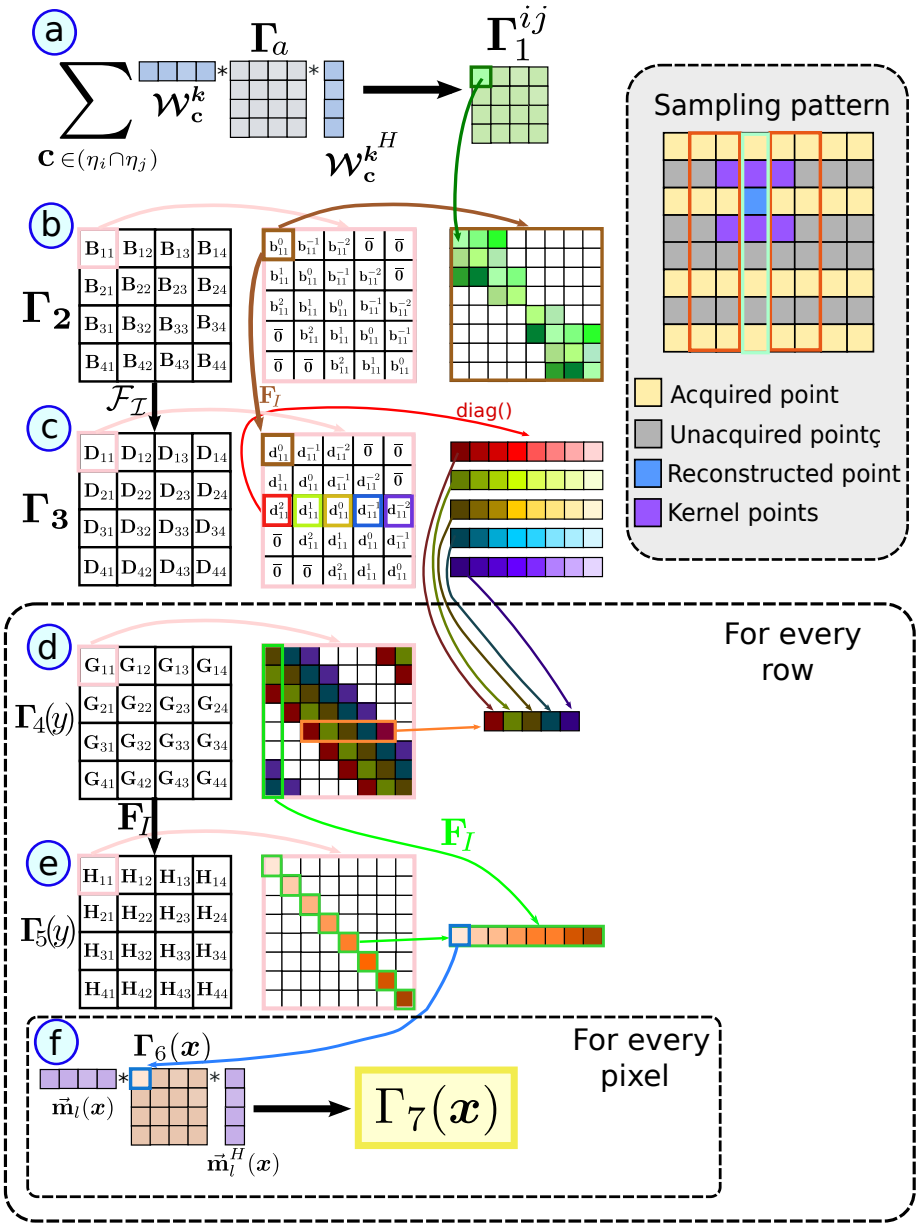

Figure 3.3: Graphical description of the algorithm. For this example we consider a matrix size 8×8 with the subsampling pattern shown in the upper right corner and L=4 coils. (a) Each reconstructed point is correlated with all the points in its kernel as well as with all the reconstructed points whose kernel correlates with the kernel of the reference one. On the other hand, the unacquired points correlate only with the reconstructed points that had them on its kernel. The point to point correlations can be computed using Eq.3.14. (b) A reference column (green column on the top right) is correlated with $2K_f - 1$ columns, marked in orange. These correlations are stored in Γ_2, which shows a block structure. (c) Computing the iFT across the phase-encoding direction results in Γ_3, where each of the blocks can be computed separately using Eq.3.21, and we can keep only the diagonal of the blocks since we only consider the correlations within a row. (d) For every row, we rearrange properly the correlation of a point and its $2K_f - 1$ neighboors. These correlations are stationary across the column, resulting in a Toeplitz structure for every block. (e) The iFT across the FE direction results in diagonal blocks due to this Toeplitz structure and can be computed very fast using Eq.3.26. (f) Finally, for every pixel in the row we simply apply Eq.4.13 using the Walsh coil-combination vector.

3.2.3 Computational complexity

A naive $\boldsymbol{k}$–space approach would require large amounts of memory to store complex matrices ($\boldsymbol{\Gamma}$ and $\mathbf{C}$) of size $[N_p \cdot N_f \cdot L] \times [N_p \cdot N_f \cdot L]$ containing the correlations between all the points across all the channels. Following this approach, $O\left((N_p \cdot N_f \cdot L)^3\right)$ complex multiplications would be performed to compute the noise maps. However, it is worth noticing that an equivalent GRAPPA reconstruction can be performed in the hybrid space (k_f, y) after carrying out the iFT along the FE direction as in (Brau et al., 2008). This would require to store N_f complex matrices of size $[N_p \cdot L] \times [N_p \cdot L]$ containing the correlations between all the points in every column across all the channels. Following this approach, $O\left(N_f \cdot (N_p \cdot L)^3\right)$ complex multiplications would be performed to compute the noise maps.

By taking advantage of the separability into independent blocks, our method achieves the storage of all the correlations in $(2K_f - 1) \cdot L \cdot (L+1)/2$ complex matrices of size $N_p \times N_p$. Since these blocks can be processed independently, it is possible to operate with them sequentially and to avoid the need to store all of their data simultaneously. Regarding the number of operations, it is necessary to compute $[(2K_f - 1) \cdot L \cdot (L+1)/2]$ 2D-iFT for the phase-encoding direction and $[N_f \cdot L \cdot (L+1)/2]$ 1D-iFT for the FE direction. Overall, this results in a total number of $(4K_f - 1) \cdot [L \cdot (L+1)/2] \cdot O\left(N_f^2 \cdot \log(N_f)\right)$ complex multiplications. We would like to remark that our methodology could be used with the equivalent reconstruction in the aforementioned hybrid space (k_f, y). Following this approach, each column is reconstructed independently, thus removing the need to keep track of the correlations between adjacent columns. However, the weights to reconstruct each column are no longer the same for all the colums. Consequently, it is necessary to repeat the column noise analysis separately for each column. Operating sequentially on the columns and exploiting the Hermitian symmetry would allow to store just $L \cdot (L+1)/2$ matrices of size $N_p \times N_p$. Separate processing of the independent blocks would require the computation of $[N_f \cdot L \cdot (L+1)/2]$ 2D-iFT, which would result in $[L \cdot (L+1)] \cdot O\left(N_f^3 \cdot \log(N_f)\right)$ complex multiplications. Compared to the direct $\boldsymbol{k}$–space analysis, the alternative hybrid approach would reduce the memory requirements, but at the expense of increasing the number of operations.

For comparison with our proposed method, the direct analysis in the image–space (Breuer et al., 2009; Aja-Fernández et al., 2011) can highly reduce the memory requirements. Since the channel correlations in the acquired sub-sampled images are spatially stationary across the image, it suffices to store the $L \times L$ complex correlation matrices for a single pixel. Thus, the main memory requirement arises from the need to store the $N_f \times N_p$ reconstruction kernels $W_m(l, \boldsymbol{x})$ transformed into the image domain. Further, $M \cdot L^2 \cdot O\left(N_f^2 \cdot \log(N_f)\right)$ complex multiplications are performed to compute the 2D-iFT of the kernels for the M differently undersampled regions, and $M \cdot L^2 \cdot N_p \cdot N_f$ complex multiplications are performed to compute the noise maps for all the channels.

3.3 Methods

Two data sets are considered for the experiments (see Fig.4.3):

- *Simulated brain data set:* a reference axial brain MR image was obtained from the BrainWeb database ((D.L. et al., 1998)). This is a T1-weighted image, with intensity non-uniformity set to INU=0%, slice thickness=1 mm and intensity range normalized to [0,255]. An 8-coil acquisition was simulated by modulating the image using artificial sensitivity maps coded for each coil as in (Aja Fernández and Vegas Sánchez-Ferrero, 2016; Aja-Fernández et al., 2014). The noise-free coil images were transformed into the $\boldsymbol{k}$–space and corrupted with synthetic Gaussian noise characterized by the matrices $\mathbf{\Gamma}_k$ and $\mathbf{C}_k$ with SNR=30 for each coil, and the correlation coefficient between coils was set to $\rho = 0.1$. For statistical purposes, 4000 realizations of each image were used.

- *Water phantom acquisition:* 100 realizations of the same fully-encoded slice of a water phantom doped with 3.3685 g/L of nickel chloride hydrate ($NiCl_2$-$6H_2O$) and 2.4g/L of sodium chloride (NaCl). They were scanned in an 8-channel head coil on a 3.0T scanner (MR750, GE Healthcare, Waukesha, WI). Acquisition parameters included: TE/TR=2.0/11.8ms, flip angle=3°, FOV=220×220mm^2, matrix size=128×128, slice thickness= 3mm, bandwidth= ±62.5KHz, total scan time=641.2 seconds. In order to ensure steady-state, we acquired 200 realizations and discarded the first 100. Also, we corrected for B_0 field drift related phase variations by a pre-processing step that estimated the phase-shift between realizations from the center of the $\boldsymbol{k}$–space as a cubic function of time and removed it afterwards.

- *In vivo acquisition:* in order to prove the feasibility of the proposed method, a fully encoded 2D axial slice head experiment was perfermod on a volunteer and informed consent was obtained prior to the acquisition. A FSE (Fast Spin Echo) sequence was used with a 32-channel head coil on a 3.0T scanner (MR750, GE Healthcare, Waukesha, WI). Acquisition parameters included: TE/TR=30.8/2000ms, FOV=260×260mm^2, matrix size=256×256, slice thickness= 3mm, bandwidth= ±22.5KHz, total scan time=132 seconds.

From the first two data sets, three experiments were performed:

1. First, we studied the case of uniform undersampling patterns. The $\boldsymbol{k}$–space data corresponding to the first realization in each data set was subsampled with acceleration factor $R = 2$ and 32 ACS lines. The GRAPPA kernel and the combination vector given by (Walsh et al., 2000) are estimated from the ACS lines from the first realization.

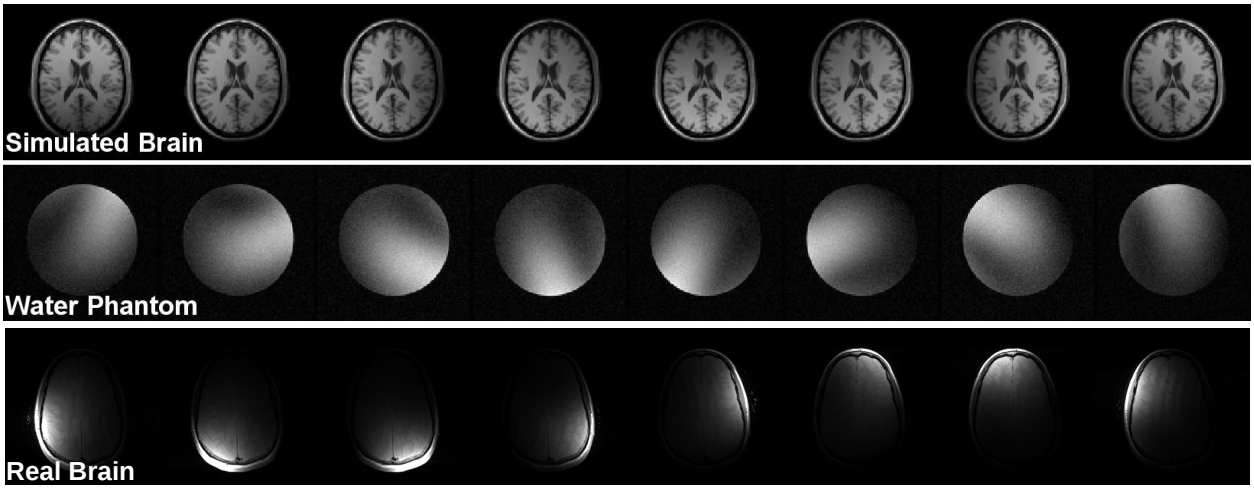

Figure 3.4: Datasets (coil-by-coil images) used for the Monte–Carlo simulation, the phantom and the *in–vivo* experiments. Only 8 of the 32 channels are shown for the case of the *in–vivo* dataset.

For each realization, the $\boldsymbol{k}$–space data were undersampled uniformly without including ACS lines. Next, each realization was reconstructed using the same kernel (from the first realization) and linearly combined with the previously computed coil combination vector. Note that, in this case, the assumptions underlying image–space based methods are correct, and therefore image–space methods should be equivalent to the proposed $\boldsymbol{k}$–space method for this first case.

2. Second, we considered the case of non-uniform subsampling patterns due to the presence of the ACS lines. For this purpose, the $\boldsymbol{k}$–space data for every realization in each data set was subsampled with an ACS region containing 32 lines. In order to study the effect of the kernel size and the acceleration factor, three combinations of parameters were used: $\{R = 3, \text{kernel} = [2, 3]\}$, $\{R = 4, \text{kernel} = [2, 3]\}$ and $\{R = 3, \text{kernel} = [4, 3]\}$. For each scenario, the first realization of each data set was used to compute the GRAPPA kernel and the coil combination vector, which then were used to reconstruct the final image.

3. Third, we considered the case of non-uniform subsampling patterns caused by a Variable Density reconstruction. For this purpose, the first image in each data set was subsampled with different acceleration factors in different $\boldsymbol{k}$–space regions (R=2, 3 or 4) and an unaccelerated region at the $\boldsymbol{k}$–space center containing 32 ACS lines. The ACS were used to obtain the different GRAPPA kernels, from which the coil images were reconstructed and used to estimate the coil combination vector. Then, for each realization, the $\boldsymbol{k}$–space coils were subsampled non-uniformly with the same pattern, reconstructed using GRAPPA with the previously computed kernel and linearly combined as in (Walsh et al., 2000) with the previously computed combination vector.

For these experiments, g-factor maps were obtained in three different ways: First, we followed a Monte–Carlo strategy to compute the sample standard deviation for each pixel from all the realizations for both the unaccelerated and the accelerated images. The g-factor maps are obtained by simply applying its definition; Second, we directly computed the g-factor maps using the image–space method described in Eq.3.9; Third, we directly computed the g-factor maps using the proposed $\boldsymbol{k}$–space method as in Eq.3.32. The last two methods require as inputs the $\boldsymbol{k}$–space noise matrices $\boldsymbol{\Gamma}$ and $\mathbf{C}$ in the subsampled images, the GRAPPA kernel and the coil-combination vector. For the simulated data set, all the parameters were known, whereas for the acquired phantom, these matrices needed to be estimated. Since multiple realizations of the same acquisition were available, we estimated them as the sample covariance matrices obtained across all the realizations. Furthermore, since these matrices are assumed to be stationary across the $\boldsymbol{k}$–space, we computed them for every $\boldsymbol{k}$–space sample and averaged them afterwards. However, in order to minimize errors caused by the residual B_0 drift after its removal, only the points where no signal was present were considered for the averaging. For an *in–vivo* acquisition, the estimation of these matrices should be done from a single image as in (Aja Fernández and Vegas Sánchez-Ferrero, 2016). However, the objective of this study is to show that the g-factor maps can be computed exactly when these matrices are known.

Last, in order to prove the applicability to clinical scenarios, we studied the impact of the acceleration rate and the size of the kernel using the *in–vivo* scan. The acquired $\boldsymbol{k}$–space was subsampled with an ACS region containing 32 lines using all possible combinations with $R = \{2, 3, 4\}$ and kernel $= \{[2, 3], [4, 3]\}$.

All image reconstruction and g-factor maps estimation were performed using Matlab and run on a standard PC with an Intel®Core™ i5-4210M @2.6 GHz Processor and 7.5 GB of RAM. In the spirit of reproducible research, we provide a software package including both the data sets and the code that we used, allowing to reproduce all the results included in this manuscript. It can be downloaded from `http://lpi.tel.uva.es/grappa_kspace`.

3.4 Results

Fig.4.5 shows the g-factor for each pixel in the synthetic phantom image reconstructed using GRAPPA with different strategies: uniform undersampling, non-uniform undersampling with ACS and non-uniform undersampling with VD-GRAPPA. The last two columns show the differences between the two methods against the Monte Carlo estimation. One of the strengths of our method is that it provides the accurate noise maps under the aforementioned assumptions of deterministic GRAPPA weights and I.I.D. $\boldsymbol{k}$–space samples. The theoretical derivation suffices to prove it since we have simply kept track of the evolution of the noise

variance throughout the linear reconstruction. No extra assumptions or approximations have been introduced in the procedure, yielding the exact analytical derivation.

Regarding the acquired phantom, since only 100 repetitions were available, the Monte Carlo estimates are much noisier that in the synthetic scenario. Since our theoretical derivation proves that our analysis is exact under the aforementioned assumptions, and given the empirical support provided by the synthetic experiments, we compare directly the image–space results against our results. This allows us to show the improvement introduced by the $\boldsymbol{k}$–space analysis that otherwise would be hidden under the randomness inherent to the Monte Carlo estimation, due to the small number of available repetitions.

When the undersampling pattern is uniform across the phase-encoding direction, both methods are completely equivalent, and thus there is no error in the estimation of the g-factor maps. However, when the undersampling pattern is not uniform, both methods provide a slightly different estimation as expected, with errors ranging from 5% in the areas of high g–factor up to 30% in areas of lower g–factor. Since the errors introduced by the image–space method arise from the reconstruction of the boundaries between differently-sampled regions, their magnitude will depend on the undersampling pattern.

As previously reported, the image–space method assumes the $\boldsymbol{k}$–space can be split into independent regions. Importantly, this assumption introduces residual errors due to neglected overlapping between regions, with the number of overlapping lines depending on the size of the kernel. Further, since the acquired lines in the boundaries that separate every two regions are used to reconstruct both regions, the missing lines in the two regions that are reconstructed from the shared acquired line will be correlated, with the number of correlated lines depending on the acceleration factor in the regions. Fig.4.5 and Fig.4.6 show how the error grows when the size of the kernel, the acceleration factor or the number of regions are increased.

In order to show that the $\boldsymbol{k}$–space analysis can be applied to *in–vivo* datasets, the g-factor maps obtained with different acceleration rates and kernel sizes are shown in Fig.4.7. It can be seen that increasing the acceleration rate results in an increase of the g-factor effect on the noise amplification, whereas increasing the kernel size results in an improved performance of the reconstruction regarding noise amplification. It can also be observed that by reducing the acceleration rate, or increasing the kernel size, the noise amplification smooths spatially.

Finally, as for computational load, we performed a test with a 256×256 image using both the image–space and the $\boldsymbol{k}$–space methods. The image–space method required 4.66 seconds, whereas our proposed $\boldsymbol{k}$–space method required 4.06 seconds. Importantly, the Matlab code used in this chapter has not been optimized for speed, so significant time reduction may be achieved if the high degree of paralelism of some reconstruction steps in our method is accounted for in the programming.

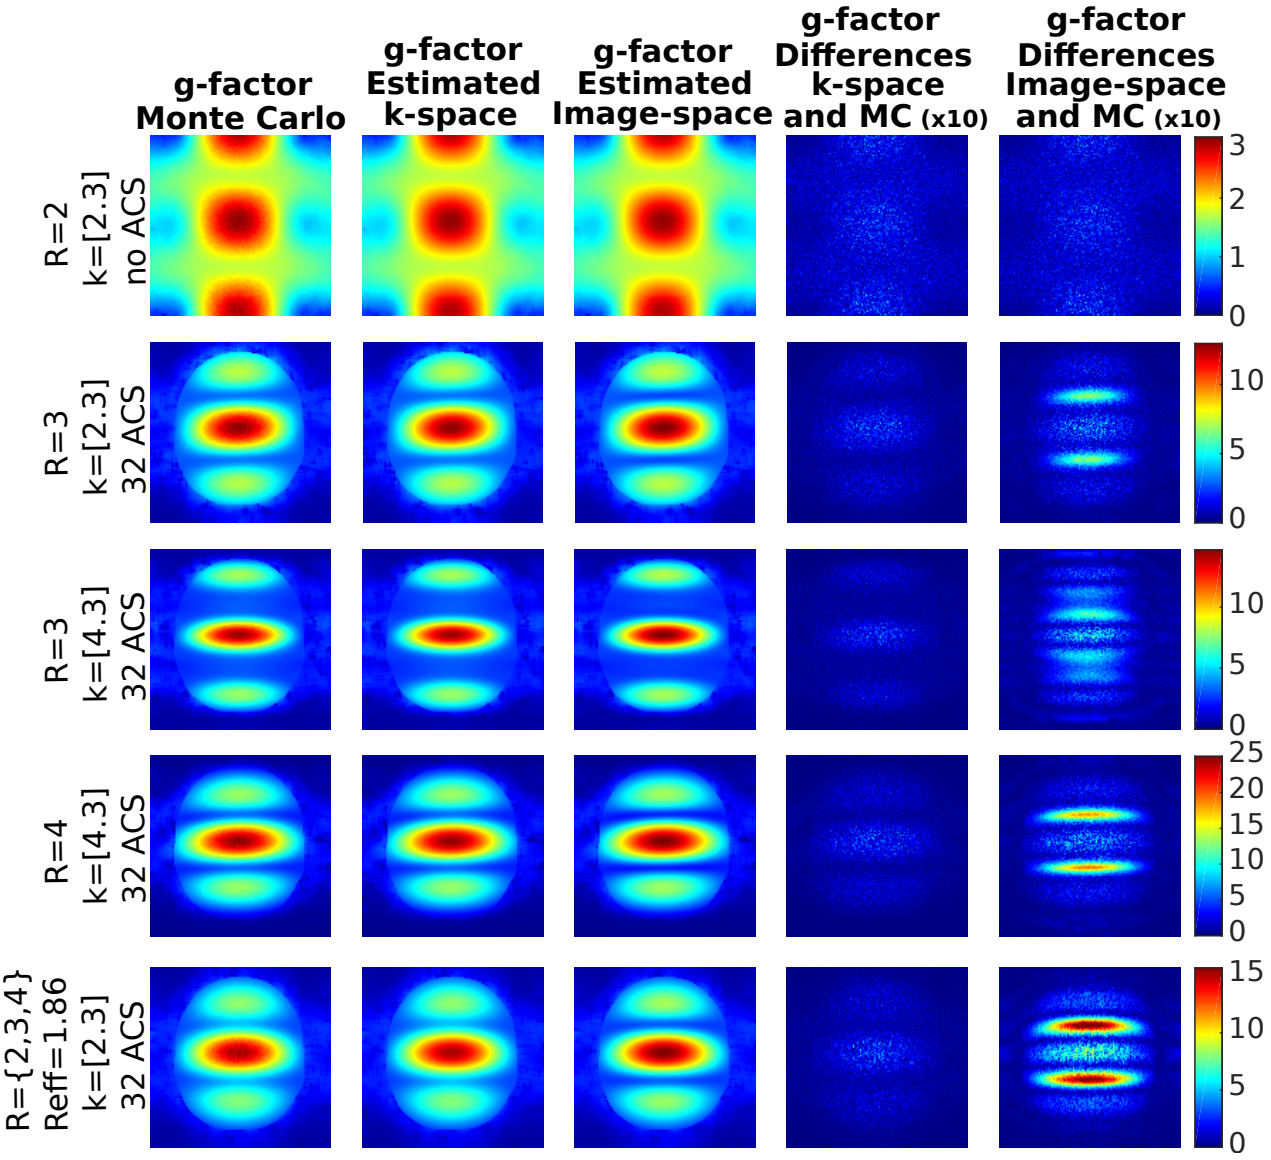

Figure 3.5: g-factor maps for the synthetic phantom obtained through a Monte–Carlo strategy (first column) and estimated using the proposed $\boldsymbol{k}$–space method (second column) and the image–space method (third column). The last two columns show the absolute differences between the two methods against the Monte Carlo estimation. Several escenarios are studied: uniform undersampling with R=2 and kernel=[2,3] (first row), non-uniform undersampling including 32 ACS with R=3 and kernel=[2,3] (second row), R=3 and kernel=[4,3] (third row), R=4 and kernel=[2,3] (fourth row), and non-uniform undersampling with a Variable Density approach (fifth row).

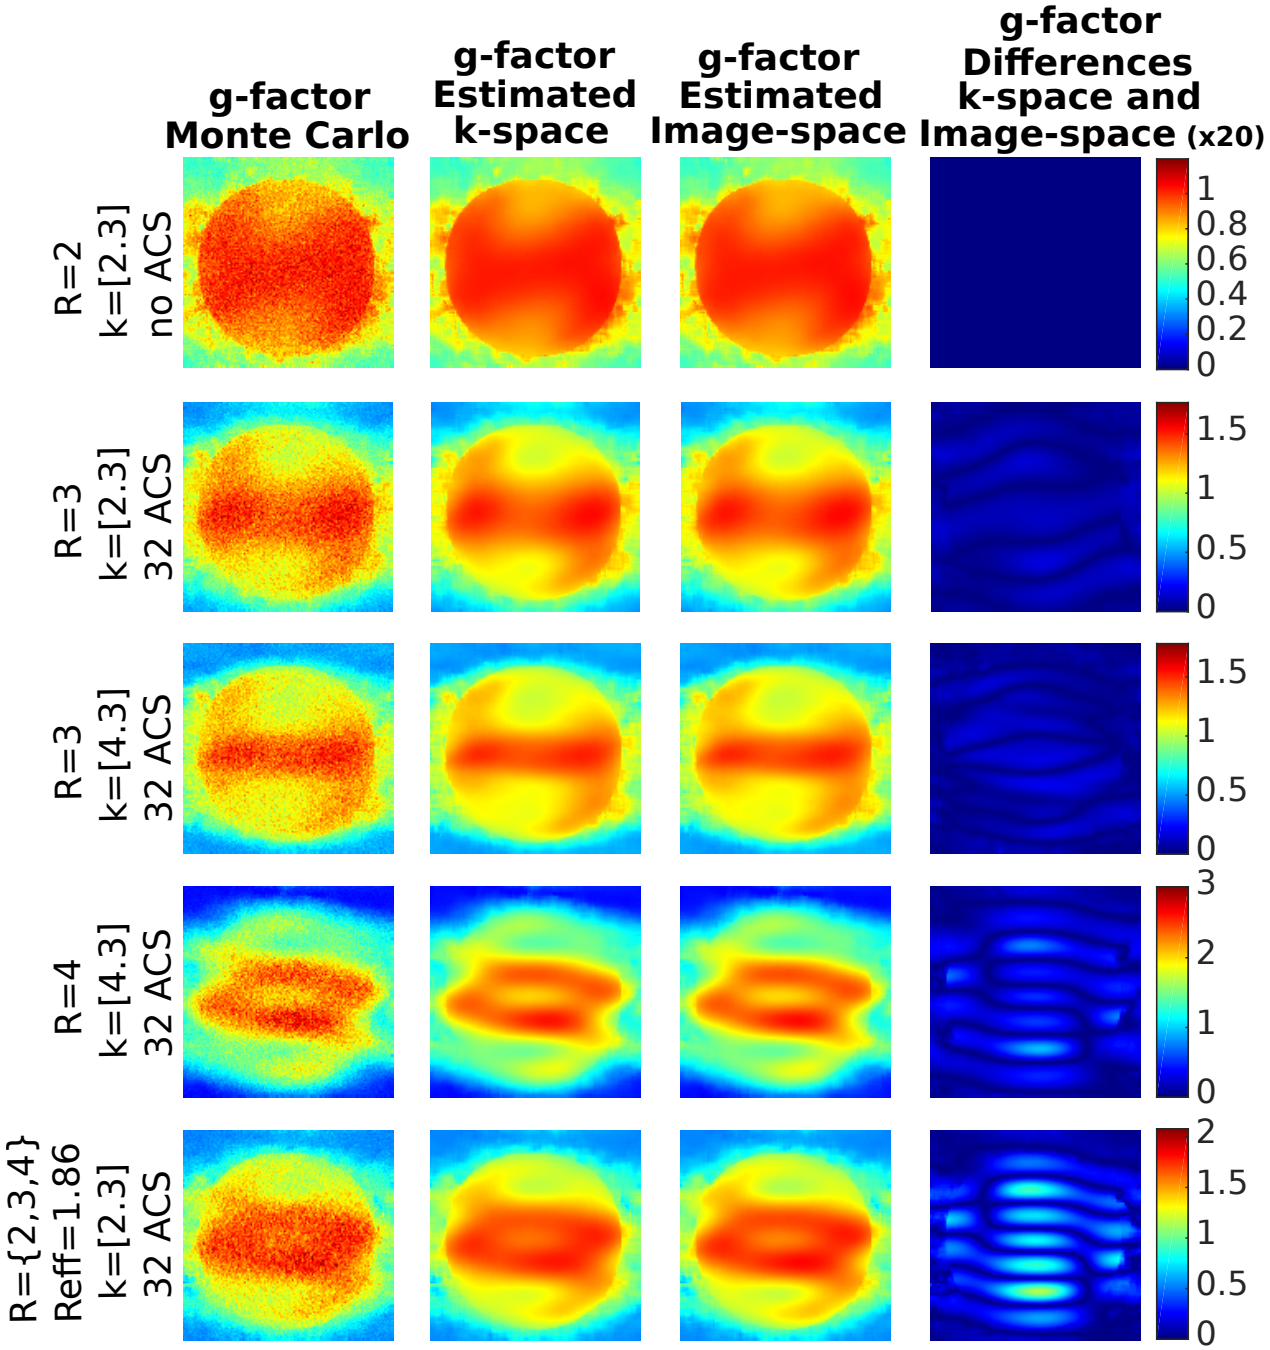

Figure 3.6: g-factor maps for the scanned water phantom obtained through a Monte–Carlo strategy (first column) and estimated using the proposed $\boldsymbol{k}$–space method (second column) and the image–space method (third column). The last column shows the absolute differences between the exact g-factor maps ($\boldsymbol{k}$–space) and the approximated (image–space). Same scenarios as for the synthetic phantom are shown.

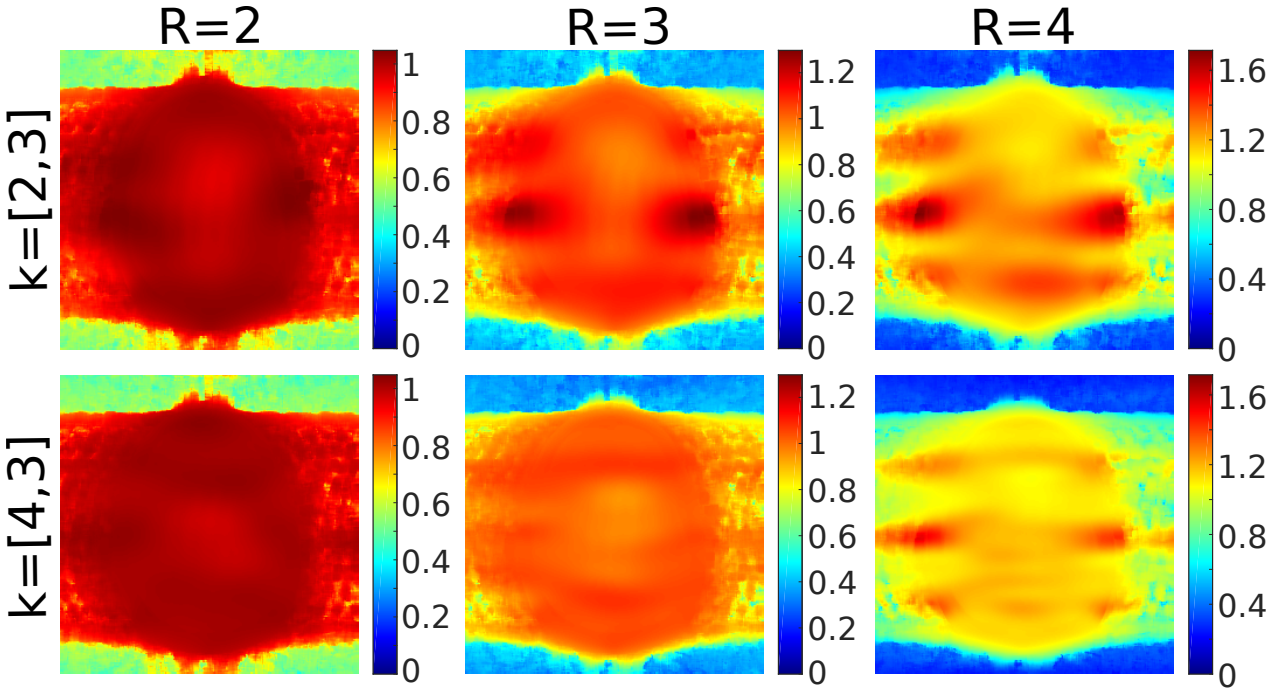

Figure 3.7: g-factor maps for the *in–vivo* brain experiment obtained through the proposed $\boldsymbol{k}$–space method. Several uniform undersampling escenarios are considered with acceleration rate R=2 (first column), R=3 (second column) and R=4 (third column). The kernel size impact is also studied using kernel=[2,3] (first row), non-uniform undersampling including 32 ACS with R=3 and kernel=[2,3] (second row), R=3 and kernel=[4,3] (third row), R=4 and kernel=[2,3] (fourth row), and non-uniform undersampling with a Variable Density approach (fifth row).

3.5 Discussion

An analytical method for noise characterization in GRAPPA reconstructions has been proposed. The method allows for an exact characterization of noise under the assumptions of stationarity and uncorrelation in the original $\boldsymbol{k}$–space undersampled acquisition. The strength of the method lies in two cornerstones: (1) by operating directly in the $\boldsymbol{k}$–space, we succeed in providing the exact characterization of noise by accounting for all the $\boldsymbol{k}$–space correlations; (2) by exploiting the extensive symmetries and the separability in the reconstruction steps, we overcome the computational challenges related to the very large size of the covariance matrices. As a result, the proposed method provides an accurate characterization of noise under the assumptions of stationarity and uncorrelation in the original $\boldsymbol{k}$–space undersampled acquisition.

The accurate modeling of noise provided by the theoretical method is also confirmed by the experiments, where it outperforms previously proposed methods. The Monte–Carlo method proposed in (Robson et al., 2008) is a versatile approach for noise characterization. However, this method is very time consuming, due to the need to reconstruct multiple images, as well as subject to errors due to the limited number of replicas for the Monte–Carlo estimation. Compared with this approach, our proposal provides an exact characterization of noise and avoids the need to generate and reconstruct multiple replicas of the image.

The image–space method proposed in (Breuer et al., 2009; Aja-Fernández et al., 2011), succeeds in providing a close approximation to the actual noise maps, which holds exact for the case of uniform undersampling patterns, but they are subject to errors when a non-uniform undersampling pattern is used. By operating directly in the $\boldsymbol{k}$–space, our proposed method is able to avoid these errors and provide an accurate estimation in a feasible computation time. However, for simple patterns the image–space method provides a good approximation of the noise and g–factor maps with reduced memory requirements. Importantly, our method does not provide a better approximation compared to image–space method (i.e.: an incremental improvement), but rather provides the exact solution to the problem of calculating noise maps (subject to the same assumptions as image–space methods), and without increasing the computation time. Furthermore, although we have focused on estimating the noise at every location in the image, our method allows to obtain the correlation between any two points in the image, something that cannot be accomplished through image–space methods. This could be useful for applications such as the optimal filtering proposed by (Sprenger et al., 2016).

This study has some limitations: First, the analysis here presented is focused on GRAPPA reconstruction followed by a linear coil combination. For other coil combination strategies, the final step of the analysis must be redone. If SoS is considered, for instance, the noise

distribution can be approximated as a non-Central χ distribution, as shown in (Aja-Fernández et al., 2011), whose effective values must be calculated. However, this step is a straightforward extension of this manuscript once the final covariance matrices are built. It would just require to replace their matrix $\mathbf{C_x}$ (obtained using the image–space analysis) with the covariance matrices $\mathbf{\Gamma}_6(x)$ and $\mathbf{C}_6(x)$; Second, this manuscript focuses on 2D-GRAPPA reconstructions accelerated in one dimension. Nevertheless, we believe the proposed approach can be extended to acquisitions with acceleration in two dimensions, or 3D acquisitions. It is worth noticing that the error introduced by the image–space method arises from the boundaries between differently sampled regions and it is expected to grow with the number and size of these boundaries. A 3D acquisition with two undersampling directions offers great flexibility to use highly non-uniform sampling patterns that would result in more boundaries, so in such a scenario our accurate k–space analysis may provide an increased advantage with respect to the image–space analysis. On the other hand, this flexibility may complicate the analysis due to the possible loss of the cyclic structure of the GRAPPA reconstruction along the phase encoding directions. We leave the details of this generalization to 3D scenarios to future research. Third, our analysis is focused on Cartesian acquisitions. If non-Cartesian acquisitions (eg: radial or spiral trajectories), are followed by a linear interpolation into a Cartesian grid, our procedure can still be applied, although this process could increase the correlations in $\boldsymbol{k}$–space, as well as introducing non-stationarity in the $\boldsymbol{k}$–space since each location may be regridded using different weights. A potential extention to our method could be non-Cartesian acquisitions regridded with the GROG operator, as in (Seiberlich et al., 2008c), since no extra-correlations would be introduced and only the non-stationarity should be taken into account. Fourth, the proposed method assumes that the kernel used for reconstruction is independent of the noise in the acquired image, which is not exact when the kernel is autocalibrated from the data instead of from a separate pre-scan. However, this assumption is expected to be a good approximation due to the typical overdetermination of GRAPPA weights estimation from the ACS region, and for this reason it has become a common assumption in the literature.

3.6 Conclusion

We have proposed a novel method for noise characterization in GRAPPA reconstructions. Through a careful direct $\boldsymbol{k}$–space analysis that exploits both the symmetry and separability in the reconstruction steps, our method provides an exact noise characterization under the assumptions of stationarity and uncorrelation in the acquired $\boldsymbol{k}$–space.

CHAPTER

FOUR

EXACT COMPUTATION OF G–FACTOR MAPS IN 3D GRAPPA

Abstract

Purpose: To characterize the noise distributions in 3D-MRI accelerated acquisitions reconstructed with GRAPPA using an exact noise propagation analysis that operates directly in $\boldsymbol{k}$–space.

Theory and Methods: We exploit the extensive symmetries and separability in the reconstruction steps to account for the correlation between all the acquired $\boldsymbol{k}$–space samples. Monte Carlo simulations and multi-repetition phantom experiments were conducted to test both the accuracy and feasibility of the proposed method; an *in–vivo* experiment was performed to assess the applicability of our method to clinical scenarios.

Results: Our theoretical derivation shows that the direct $\boldsymbol{k}$–space analysis renders an exact noise characterization under the assumptions of stationarity and uncorrelation in the original $\boldsymbol{k}$–space. Simulations and phantom experiments provide empirical support to the theoretical proof. Finally, the *in–vivo* experiment demonstrates the ability of the proposed method to assess the impact of the sub-sampling pattern on the overall noise behavior.

Conclusions: By operating directly in the $\boldsymbol{k}$–space, the proposed method is able to provide an exact characterization of noise for any Cartesian pattern sub-sampled along the two phase–encoding directions. Exploitation of the symmetries and separability into independent blocks through the image reconstruction procedure allows us to overcome the computational challenges related to the very large size of the covariance matrices involved.

4.1 Introduction

As we explained in the previous chapter, characterization of noise statistics in MR images is essential for multiple applications. In particular, commonly used PI methods where the reconstruction takes place in the $\boldsymbol{k}$–space, such as GRAPPA, introduce correlations in the $\boldsymbol{k}$–space that propagate into the $\boldsymbol{x}$–space. In (Robson et al., 2008) it has been shown that a direct noise propagation analysis requires handling large covariance matrices, which renders this analysis challenging. In the previous chapter we have shown as well that, by exploiting the extensive symmetries and separability in each step of the reconstruction, a computationally efficient and exact noise analysis can be obtained for 2D Cartesian acquisitions, where acceleration is applied along a single (phase-encoding) dimension. However, no feasible and exact noise analysis has been reported for 3D Cartesian acquisitions.

Indeed, 3D acquisitions are commonplace in clinical and research MRI exams, both in brain and body imaging applications. Compared to 2D imaging, 3D acquisitions enable high (eg: isotropic) spatial resolution with high SNR efficiency. Further, the high SNR efficiency and multi-dimensional sampling in the presence of dense receive coil arrays are ideally suited for PI acceleration along multiple dimensions, as exploited in (Stevens et al., 2008; Meisamy et al., 2011; Kalia et al., 2017). Importantly, (Weiger et al., 2002; Blaimer et al., 2006) have shown how PI techniques can be combined with 3D imaging by sub-sampling the $\boldsymbol{k}$–space along both phase–encoding directions, resulting in highly flexible sub-sampling schemes, with multiple different kernels applied over different regions of $\boldsymbol{k}$–space, as done by (Wang et al., 2011). However, GRAPPA reconstruction in 3D imaging introduces correlations in $\boldsymbol{k}$–space across the three dimensions, greatly increasing the size of the matrices involved in the noise propagation analysis. This can easily become prohibitive regarding both computation time and memory requirements, so we have shown how to extend the philosophy in (Rabanillo et al., 2017) and provide the implementation for doing so. Furthermore, for 3D-GRAPPA the reconstruction is in general non-stationary along the two PE directions available for subsampling, forcing to considerate individually each entry along this dimensions. We explain how to overcome this challenge not present in 2D acquisitions by a previous step that identifies all the columns with an equivalent pattern, notably reducing the computation load.

In the previous chapter, we described two approaches to approximate the noise maps alleviating the computational burden in 2D GRAPPA: the Monte–Carlo (MC) philosophy described by (Robson et al., 2008) and the $\boldsymbol{x}$–space methods suggested by (Breuer et al., 2009). Although these techniques can be easily extended to 3D GRAPPA, their practical application is more problematic compared to the 2D case. In particular, the MC method requires to reconstruct multiple replicas of the sampled data, increasing the time computation since now we must perform 3D GRAPPA reconstructions. Alternatively, we showed in the previous chapter that the $\boldsymbol{x}$–space method introduces systematic errors in the estimated

noise maps. They are generally associated with the boundaries between $\boldsymbol{k}$–space regions with different different sampling patterns, which is specially troublesome in 3D acquisitions obtained with highly non–uniform sub-sampling patterns.

In this chapter, we propose a method for exact noise characterization of 3D GRAPPA reconstructions. This method generalizes the 2D approach described in the previous chapter by exploiting the extensive symmetries and separability in each step of the 3D GRAPPA reconstruction. In particular, the proposed method can account for multiple kernels within the same image, such as in the variable density sampling scenarios shown in (Wang et al., 2011), which is a shortcoming of previous methods.

4.2 Theory

4.2.1 k–space method for noise characterization in GRAPPA

Similarly to the previously described 2D imaging case, we rely on the following assumptions: (1) GRAPPA weights are non-stochastic, i.e., they are independent of the noise realization. This holds strictly true for the case when they are computed from a separate calibration scan as in (Breuer et al., 2005). Nevertheless, due to the overdetermination of GRAPPA weights estimation, this assumption is common even in self-calibrated acquisitions (2); noise at different acquired points in the $\boldsymbol{k}$–space is IID* as in (Henkelman, 1985; McVeigh et al., 1985; Kellman and McVeigh, 2005).

In order to characterize the noise in the final composite image, we proceed to characterize the noise propagation through each reconstruction step (see Fig.4.2):

k–Space interpolation

We can build the correlation between a point reconstructed $\boldsymbol{k}$ and the point considered for its reconstruction $\mathbf{c}$ as in the previous chapter, thus obtaining matrices $\mathbf{\Gamma}_1$ and $\mathbf{C}_1$, as depicted in Fig.4.2.a. We will reconstruct a $\boldsymbol{k}$–space of size $[N_f, N_{p_1}, N_{p_2}]$, where once again f denotes the FE direction and we distinguish between two different PE directions $\{p_1, p_2\}$. We will hereafter refer to a column as the data along the first dimension (fixing the other two dimensions), a row the data along the second dimension and a layer the data along the third dimension. Once again, we hereafter focus on the calculation of matrix Γ since the computation of the final matrix C is analogous.

*If the noise shows temporal correlation along the FE direction, our analysis can still be applied. This potential correlation would simply increase the width of the neighborhood a reference point correlates with, as described in subsection "$\boldsymbol{k}$–space Interpolation".

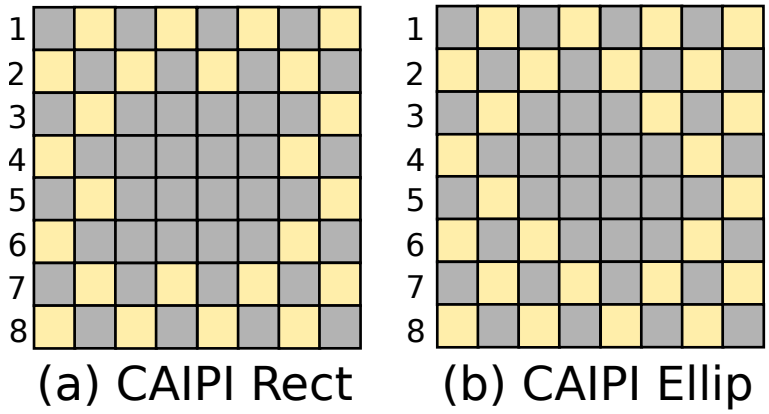

Figure 4.1: Sampling patterns for CAIPI Rect and CAIPI Ellip with a matrix size 8×8 to illustrate the correlations within a row chosen as a reference. (a) For the CAIPI Rect case a there are two unique rows. The first one (row number 1) is in the area that does not include the ACS region. Row 7 presents the same pattern and rows 2 and 8 present a shifted version of it (with shift equal to 1). The second row (row number 3) is within the area that includes the ACS region, and it is equivalent to row number 5, with rows 4 and 6 presenting a shifted version of its pattern (shift equal to 1 again). (b) For the CAIPI Ellip case, we have once again column 1 as a unique column outside the ACS region (with rows 2, 7 and 8 replicating it), but now we have two different unique rows within the ACS region: rows 3 and 4, with rows 6 and 5 being a shifted version of them (shift equal to 1) respectively. This shows that in general for the ellipsoidal region more rows with a unique pattern can be found, and this number will grow as the matrix size does.

The size of the 3D reconstruction kernel is given by $[K_f, K_{p_1}, K_{p_2}]$ and therefore a reconstructed point will correlate with $2K_f - 1$ points in the FE direction. Importantly, its correlations along the PE directions may not be stationary across the $\boldsymbol{k}$-space since we may use different reconstruction kernels for different regions. We will hereafter use the term reconstruction pattern of a row to refer to the set containing the position of the points that are used in its reconstruction, as well as the kernel used for each of them. Taking into account the stationarity of the reconstruction along the FE dimension, and in order to operate efficiently, our algorithm starts by identifying the rows that present a unique reconstruction pattern with respect to its neighboring rows, i.e., all the rows with the same reconstruction pattern will have the same correlation matrix (see Fig.4.1). We will hence avoid the unnecessary burden of computing the same correlation matrices multiple times.

Let N_{uq} be the number of rows with unique reconstruction patterns. We can use Eq.3.16 to compute the correlation matrix $\mathbf{\Gamma}_2^{m_{uq}}$ between a row with pattern m_{uq} (where $\{m_{uq} = 1, \ldots, N_{uq}\}$) and its neighboring rows in a similar fashion to (see Fig.4.2.b):

$$\mathbf{\Gamma}_2^{m_{uq}} = \begin{pmatrix} \mathbf{B}_{11}^{m_{uq}} & \mathbf{B}_{12}^{m_{uq}} & \cdots & \mathbf{B}_{1L}^{m_{uq}} \\ \mathbf{B}_{21}^{m_{uq}} & \mathbf{B}_{22}^{m_{uq}} & \cdots & \mathbf{B}_{2L}^{m_{uq}} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{B}_{L1}^{m_{uq}} & \mathbf{B}_{L2}^{m_{uq}} & \cdots & \mathbf{B}_{LL}^{m_{uq}} \end{pmatrix}, \tag{4.1}$$

where Hermitian symmetry guarantees that $\mathbf{B}_{ij}^{m_{uq}} = (\mathbf{B}_{ji}^{m_{uq}})^H$ and thus we only need to

compute $L \cdot (L+1)/2$ blocks.

The same division into sub–blocks of size $N_{p_1} \times N_{p_1}$ as in Eq.3.17 can be done, each of them containing the correlation between each pair of the rows contained in the neighborhood. Likewise, these sub-blocks will be transformed independently. It is important to notice that in the 3D case, this computation needs to be repeated for each row with a unique pattern. Exploiting the stationarity along the first dimension and the identification of unique columns, we only need to compute the central row of the block matrix $\mathbf{B}_{ij}^{m_{uq}}$ containing the correlation sub-blocks $\{\mathbf{b}_{ij,m_{uq}^p}^{m_{uq}}, m_{uq}^p = 1, \ldots, M_{uq}\}$ between the reference row and its neighboring rows.

The total number of rows considered for a reference row is given by $M_{uq} = (2K_f - 1) \cdot M_{uq}^{p_2}$, where $M_{uq}^{p_2}$ is the number of row that the reference row correlates with in the third dimension, which depends on the pattern indexed by m_{uq}. Therefore, the overall number of sub-blocks we need to keep track of is $M_{Tot} = \sum_{m_{uq}=1}^{M_{uq}} M_{uq}^{p_2} \cdot (2K_f - 1) \cdot L \cdot (L+1)/2$.

Row iFT

After filling the missing lines in the $\boldsymbol{k}$–space with GRAPPA interpolation, the data are transformed into the image–space through a 3D iFT, which is equivalent to successively computing a 1D-iFT along each of the dimensions. If we start by applying the iFT along the first PE dimension, the number of rows that a reference row correlates with is preserved. Moreover, the 1D-iFT can be expressed as a matrix product using the Fourier associated matrix $\mathbf{F}_I$ defined in Eq.3.19. However, the number of blocks is now increased to $M_{uq} \cdot L$, which in the hybrid–space $[k_f, y, k_{p_2}]$ results in the correlation matrix (see Fig.4.2.c)

$$\boldsymbol{\Gamma}_3^{m_{uq}} = \boldsymbol{\mathcal{F}}_I \cdot \boldsymbol{\Gamma}_2^{m_{uq}} \cdot \boldsymbol{\mathcal{F}}_I^H = \begin{pmatrix} \mathbf{D}_{11}^{m_{uq}} & \mathbf{D}_{12}^{m_{uq}} & \cdots & \mathbf{D}_{1L}^{m_{uq}} \\ \mathbf{D}_{21}^{m_{uq}} & \mathbf{D}_{22}^{m_{uq}} & \cdots & \mathbf{D}_{2L}^{m_{uq}} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{D}_{L1}^{m_{uq}} & \mathbf{D}_{L2}^{m_{uq}} & \cdots & \mathbf{D}_{LL}^{m_{uq}} \end{pmatrix}, \tag{4.2}$$

where each $\mathbf{D}_{ij}^{m_{uq}}$ block is equally composed of sub–blocks $\{\mathbf{d}_{ij,m_{uq}^p}^{m_{uq}}, m_{uq}^p = 1, \ldots, M_{uq}\}$ that can be computed independently by:

$$\mathbf{d}_{ij,m_{uq}^p}^{m_{uq}} = \mathbf{F}_I \cdot \mathbf{b}_{ij,m_{uq}^p}^{m_{uq}} \cdot \mathbf{F}_I^H. \tag{4.3}$$

Importantly, rows whose reconstruction pattern is a shifted version (along the second dimension) of another row are considered to have an equivalent (non-unique) pattern, thereby reducing the computational requirements for the proposed algorithm. This is due to the fact that the diagonal of $\boldsymbol{\Gamma}_3^{m_{uq}}$ for two rows with equivalent pattern (except for a shift) is equal (see Appendix.B) and, as shown in the the previous chapter, the proposed algorithm only relies on the the diagonal of these matrices at this step.

Layer iFT

The next step in the 3D-iFT is to perform the 1D-iFT along the second PE dimension (layer). At this step, it suffices to keep track of the correlations between the points with the same coordinate in the second dimension, since in the following stages points that do not share this second coordinate will not be combined.

Additionally, all the layers along the first dimension (column) behave equivalently due to the stationarity of the GRAPPA reconstruction along the column dimension. For each layer along the second dimension (dependence on k_f is removed due to the aforementioned stationarity along the columns), as shown with the y-dependence, we create the correlation matrix (see Fig.4.2.d):

$$\mathbf{\Gamma}_4(y) = \begin{pmatrix} \mathbf{G}_{11} & \mathbf{G}_{12} & \cdots & \mathbf{G}_{1L} \\ \mathbf{G}_{21} & \mathbf{G}_{22} & \cdots & \mathbf{G}_{2L} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{G}_{L1} & \mathbf{G}_{L2} & \cdots & \mathbf{G}_{LL} \end{pmatrix}, \tag{4.4}$$

where each block $\mathbf{G}_{lm}$ contains the correlations between the reference layer and the $2K_f - 1$ adjacent layers along the FE dimension in coils l and m within the hybrid–space (k_f, y, k_{p_2}). As in the 2D case, due to the stationarity of the GRAPPA kernel along the FE dimension and the assumption of noise stationarity, the correlation between a reference layer and its surrounding layers coincides, independently of the layer picked for reference, giving rise to a block-Toeplitz structure. If we index the $2K_f - 1$ layers placing the reference as the central one (index K_f), then the first column of the block $\mathbf{G}_{ij}$ (with length $2K_f - 1$ blocks) is defined as

$$\begin{bmatrix} \mathbf{g}_{ij}^0 & \mathbf{g}_{ij}^1 & \cdots & \mathbf{g}_{ij}^{K_f-1} & \overline{\mathbf{0}} & \cdots & \overline{\mathbf{0}} \end{bmatrix}^T, \tag{4.5}$$

where $\mathbf{g}_{ij}^m$ is the $N_{p_2} \times N_{p_2}$ covariance matrix of two layers, the subtraction of the indices of which equals m, and $\overline{\mathbf{0}}$ denotes a null matrix of dimensions $N_{p_2} \times N_{p_2}$. These correlations are obtained by simply reallocating the elements of the diagonal —at position y— from the previously computed blocks $\mathbf{d}_{ij,m_{uq}^p}^{m_{uq}}$ for every unique row in Eq.3.21. As for the first row, we get

$$\begin{bmatrix} \mathbf{g}_{ij}^0 & \mathbf{g}_{ij}^{-1} & \cdots & \mathbf{g}_{ij}^{1-K_f} & \overline{\mathbf{0}} & \cdots & \overline{\mathbf{0}} \end{bmatrix}. \tag{4.6}$$

Due to the block-Toeplitz structure and the presence of null correlations known beforehand, we only need to compute $2K_f - 1$ sub-blocks. Therefore, the overall number of sub-blocks needed to build Eq.3.22 is at most $(2K_f - 1) \cdot L \cdot (L + 1)/2$. Finally, computing the row 1D-iFT provides a correlation matrix given by (see Fig.4.2.e):

$$\mathbf{\Gamma}_5(y) = \boldsymbol{\mathcal{F}}_I \cdot \mathbf{\Gamma}_4(y) \cdot \boldsymbol{\mathcal{F}}_I^H = \begin{pmatrix} \mathbf{H}_{11} & \mathbf{H}_{12} & \cdots & \mathbf{H}_{1L} \\ \mathbf{H}_{21} & \mathbf{H}_{22} & \cdots & \mathbf{H}_{2L} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{H}_{L1} & \mathbf{H}_{L2} & \cdots & \mathbf{H}_{LL} \end{pmatrix}, \tag{4.7}$$

where the blocks $\mathbf{H}_{ij}$ are in turn composed of sub-blocks $\mathbf{h}_{ij}^{m}$ that are obtained by:

$$\mathbf{h}_{ij}^{m} = \mathbf{F_I} \cdot \mathbf{g}_{ij}^{m} \cdot \mathbf{F_I}^{H}. \tag{4.8}$$

Column iFT

The last step to complete the 3D-iFT is the 1D-iFT along the first dimension (FE), for which we only need the correlation within each column. For each column, we can build the matrix containing the correlation within that column across all the coil channels (see Fig.4.2.f):

$$\mathbf{\Gamma}_6(y,z) = \begin{pmatrix} \mathbf{U}_{11} & \mathbf{U}_{12} & \cdots & \mathbf{U}_{1L} \\ \mathbf{U}_{21} & \mathbf{U}_{22} & \cdots & \mathbf{U}_{2L} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{U}_{L1} & \mathbf{U}_{L2} & \cdots & \mathbf{U}_{LL} \end{pmatrix}, \tag{4.9}$$

where every block $\mathbf{U}_{lm}$ has dimensions $N_f \times N_f$ and contains the cross–correlations of the selected column in coils l and m within the hybrid space (k_f, y, z). These correlations, as in the 2D case, are computed by simply choosing the appropriate components in Eq.4.7. As explained in the previous chapter, GRAPPA performs a circular interpolation that introduces a cyclical structure along the FE direction. We define the first column of block $\mathbf{U}_{ij}$ as

$$\vec{\mathbf{u}_{ij}} = \left[u_{ij}^{0},\ u_{ij}^{1},\ \cdots,\ u_{ij}^{K_f-1},\ 0,\ \cdots,\ 0,\ u_{ij}^{1-K_f},\ \cdots,\ u_{ij}^{-1}\right]^{T}, \tag{4.10}$$

with $N_f - (2K_f - 1)$ zeroes and values u_{ij}^{m} taken from the components in sub-block $\mathbf{h}_{ij}^{m}$ in Eq.3.21 that correspond to the selected column. As a result of the stationarity along the FE dimension, the j-th row of $\mathbf{U}_{ij}$, $2 \leq j \leq N_f - 1$, is obtained as a simple rightward circular shift of the row $j-1$, which results in the blocks $\mathbf{U}_{lm}$ being circulant matrices. The correlation matrix after the 1D-iFT is given by (see Fig.4.2.g):

$$\mathbf{\Gamma}_7(y,z) = \boldsymbol{\mathcal{F}}_I \cdot \mathbf{\Gamma}_6(y,z) \cdot \boldsymbol{\mathcal{F}}_I^{H} = \begin{pmatrix} \mathbf{V}_{11} & \mathbf{V}_{12} & \cdots & \mathbf{V}_{1L} \\ \mathbf{V}_{21} & \mathbf{V}_{22} & \cdots & \mathbf{V}_{2L} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{V}_{L1} & \mathbf{V}_{L2} & \cdots & \mathbf{V}_{LL} \end{pmatrix}, \tag{4.11}$$

where the blocks $\mathbf{V}_{ij}$ are simply computed as:

$$\mathbf{V}_{ij} = \mathbf{F}_I \cdot \mathbf{U}_{ij} \cdot \mathbf{F}_I^{H} = \mathrm{diag}(\mathbf{F}_I \cdot \vec{\mathbf{u}_{ij}}). \tag{4.12}$$

Note that we only need to compute $[N_f \cdot L \cdot (L+1)/2]$ 1D-iFT due to Hermitian symmetry in this step.

Coil combination

As we did in the 2D case, we create the correlation matrices $\mathbf{\Gamma}_8$ and $\mathbf{C}_8$ for each pixel $\vec{\mathbf{p}}$ along the coil dimensions. This simply consists of picking the corresponding entry in the diagonal of $\mathbf{\Gamma}_7$ or $\mathbf{C}_7$. We proceed as in the previous chapter and for every pixel in the final image we define (see Fig.4.2.i)

$$\Gamma_9(\boldsymbol{x}) = \vec{\mathbf{m}}_l(\boldsymbol{x}) \cdot \mathbf{\Gamma}_8(\boldsymbol{x}) \cdot \vec{\mathbf{m}}_l^H(\boldsymbol{x}), \tag{4.13}$$

then the variance of noise for the real and imaginary components can be calculated by:

$$\begin{aligned} \sigma^2_{\mathrm{re},\mathcal{R}}(\boldsymbol{x}) &= \frac{1}{2} 2\mathbb{R}e\left\{\Gamma_9 + C_9\right\}, \\ \sigma^2_{\mathrm{im},\mathcal{R}}(\boldsymbol{x}) &= \frac{1}{2} 2\mathbb{R}e\left\{\Gamma_9 - C_9\right\}, \end{aligned} \tag{4.14}$$

Finally, the noise and g-factor maps can be computed as in Eq.3.31 and Eq.3.32, respectively.

Summary of the procedure

The procedure to obtain $\mathbf{\Gamma}$ matrices is graphically depicted in Fig.4.2 and can be summarized as follows ($\mathbf{C}$ matrices are obtained similarly):

(i). Identify the number of rows with a unique sampling pattern. Rows whose pattern matches a shifted version of another row are not explicitly considered.

(ii). For each unique row, calculate the sub-blocks $\mathbf{b}^{m_{uq}}_{ij,m^p_{uq}}$ using Eq.3.14. The number of such sub-blocks is $M_{Tot} = \sum_{m_{uq}=1}^{M_{uq}} M_{uq}^{p_2} \cdot (2K_f - 1) \cdot L \cdot (L+1)/2$, and their dimension is $N_{p_1} \times N_{p_1}$. Each entry in these sub-blocks is obtained from Eq.3.14, by choosing the appropriate components.

(iii). Calculate the sub-blocks $\mathbf{d}^{m_{uq}}_{ij,m^{p_2}_{uq}}$ resulting from the row 1D-iFT using Eq.4.3.

(iv). For each layer, build the sub-blocks $\mathbf{g}^m_{ij}$ that build up the matrix Eq.4.4 by properly reallocating the elements of the diagonal from $\mathbf{d}^{m_{uq}}_{ij,m^{p_2}_{uq}}$.

(v). Compute the sub-blocks $\mathbf{h}^m_{ij}$ resulting from the layer 1D-iFT using Eq.4.8.

(vi). For each layer and each column, calculate $\vec{\mathbf{u}_{ij}}$ in Eq.4.10 by choosing the appropriate components of $\mathbf{h}^m_{ij}$.

(vii). Calculate $\mathbf{V}_{ij}$ resulting from the column 1D-iFT using Eq.4.12.

(viii). For every pixel, create its matrix $\mathbf{\Gamma}_8$ by selecting the right element in the diagonal of $\mathbf{V}_{ij}$.

(ix). Apply equations Eq.4.13, Eq.3.31 and Eq.3.32 to obtain pixel-wise noise characterization.

4.2.2 Computational complexity

In terms of memory requirements, the most demanding step in our method is the computation of the correlation matrices between each unique row and its neighbors in the original $\boldsymbol{k}$–space. Since we operate sequentially along the third dimension (i.e., in each iteration of a loop we consider only the rows whose distance to the reference row coincide), we need to store $[(2K_f - 1) \cdot L \cdot (L+1)/2 \cdot M_{uq}]$ complex matrices of size $N_{p_1} \times N_{p_1}$, and after the iFT we only keep the diagonal of each matrix. However, if memory constraints demand so, it would be possible to reduce the number of matrices stored at the same time by simply operating sequentially along the first dimension or the coil dimension as well.

In terms of the number of operations, the most demanding step is the computation of the iFT. It is necessary to compute $\left[(2K_f - 1) \cdot L \cdot (L+1)/2 \cdot M_{uq} \cdot \bar{M}_{uq}^{p_2}\right]$ 2D-iFT for the first phase–encoding direction, where $\bar{M}_{uq}^{p_2}$ denotes the average number of columns that a unique row correlates with along the layer dimension; $((2K_f - 1) \cdot L \cdot (L+1)/2 \cdot N_{p_1}]$ 2D-iFT for the second phase–encoding direction and $[N_{p_1} \cdot N_{p_2} \cdot L \cdot (L+1)/2]$ for the FE direction, resulting in a total number of $(4K_f - 1) \cdot [L \cdot (L+1)/2] \cdot N_{p_1} \cdot N_{p_2} \cdot O\left(\overline{N_i \cdot \log(N_i)}\right)$ complex multiplications approximately, where $\overline{N_i \cdot \log(N_i)}$ refers to $(N_{p_1} \cdot \log(N_{p_1}) + N_{p_2} \cdot \log(N_{p_2}) + N_f \cdot \log(N_f))$.

4.3 Methods

Three datasets are considered for the experiments (see Fig.4.3)

- *Simulated abdomen dataset:* we have synthetized a 3D volume using the simulation environment XCAT provided by (Segars et al., 2010) based on the extended cardio-torso phantom. We simulated a T1-weighted acquisition using the following acquisition parameters: TE/TR=1.5/3ms, flip angle=60°, acquisition matrix size=60×60×32. A 32-coil acquisition was simulated by modulating the image using artificial sensitivity maps coded for each coil as in (Aja Fernández and Vegas Sánchez-Ferrero, 2016; Aja-Fernández et al., 2014). The noise-free coil images were transformed into the $\boldsymbol{k}$–space and corrupted with synthetic Gaussian noise characterized by the matrices $\mathbf{\Gamma}_k$ and $\mathbf{C}_k$ with SNR=25 for each coil, and the correlation coefficient between coils was set to $\rho = 0.1$. For statistical purposes, 4000 realizations of each image were used.

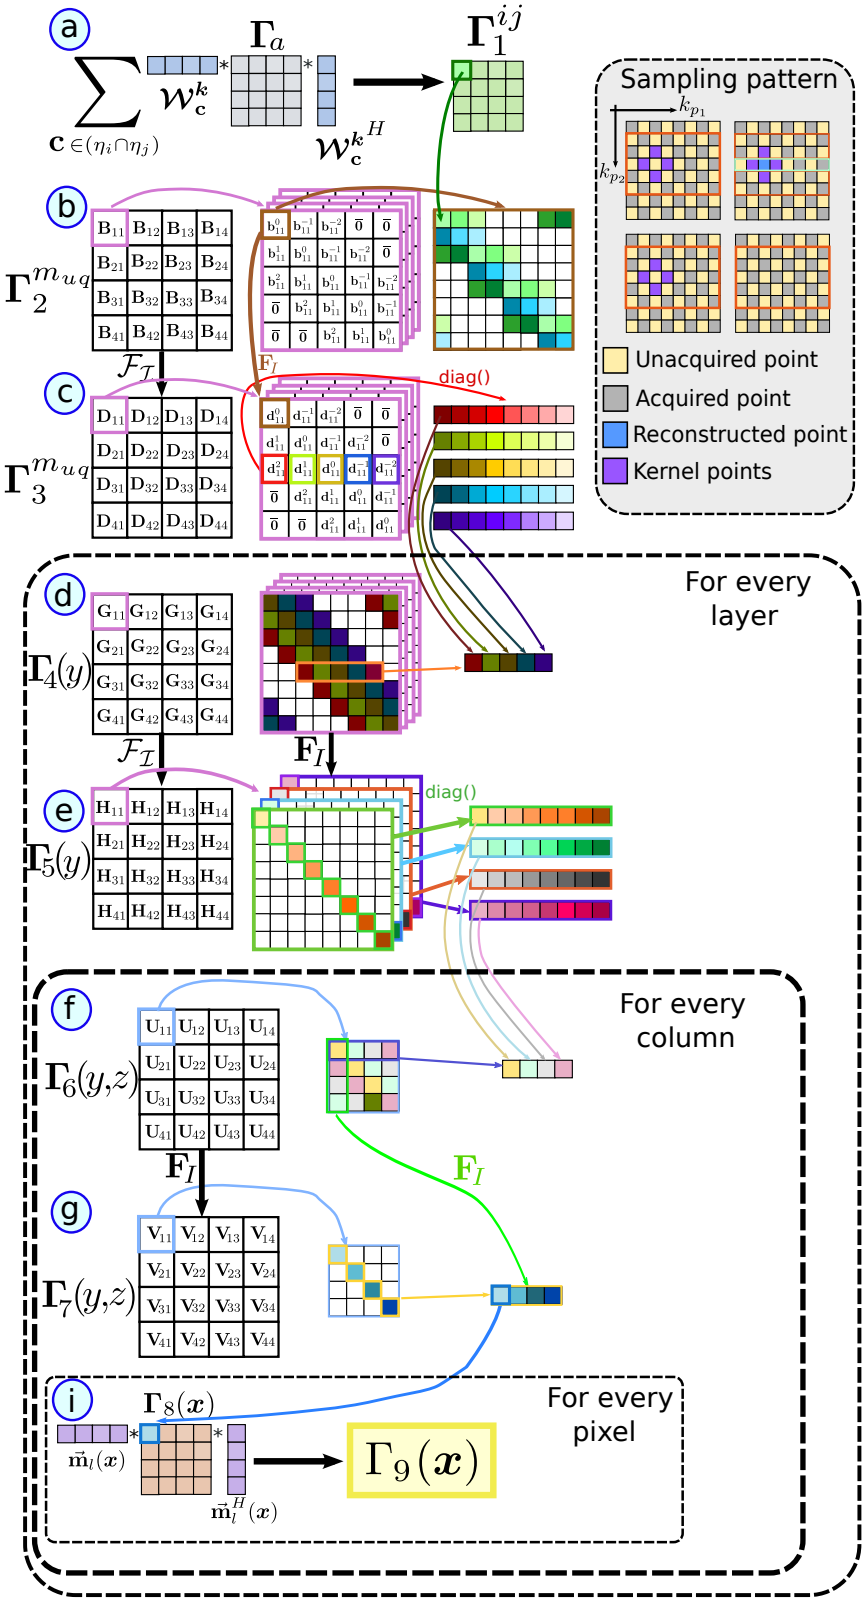

Figure 4.2: Graphical description of the algorithm with a matrix size $4 \times 8 \times 8$ (notice that the $\boldsymbol{k}$–space is rotated so the subsampling pattern along the two PE dimensions can be observed easily) with the subsampling pattern shown in the upper right corner and L=4 coils. (a) Correlation between $\boldsymbol{k}$–space points computed using Eq.3.14. (b) Correlation between a reference row (green row) with $2K_{p_2} - 1$ rows along the second PE dimension and with all the rows along the FE dimension (marked in orange). (c) Correlation after the iFT along the first PE direction. (d) Correlation between a point and its $2K_f - 1$ neighboors along all the columns. (e) Correlation after the iFT along the second PE direction. (f) These correlations result in a Toeplitz structure due to the stationary across the column. (g) Correlation after the iFT along the FE direction . (h) Correlation in the final image between different coils

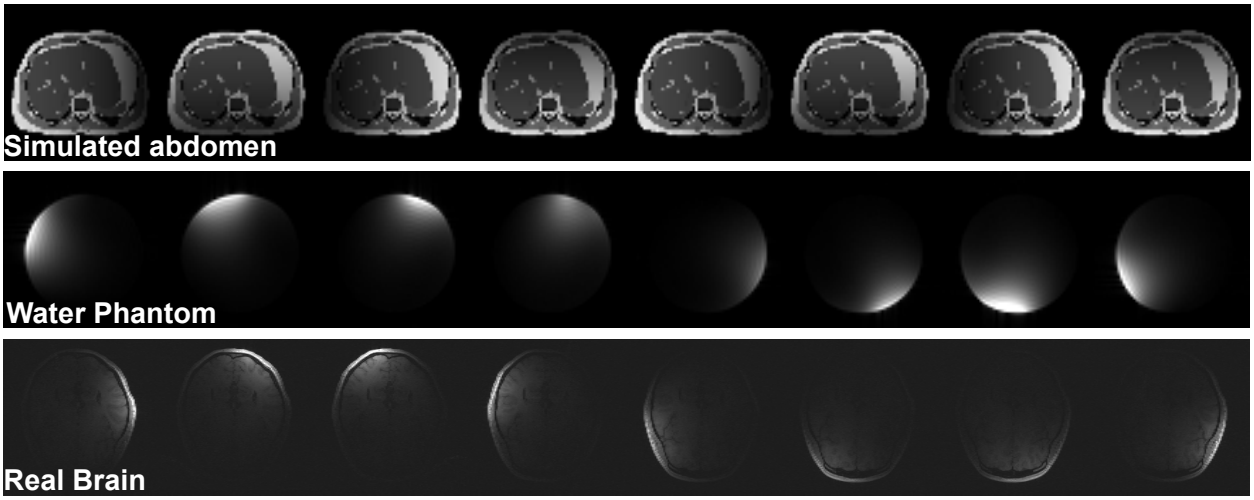

Figure 4.3: The first three rows show the datasets (coil-by-coil images) used for the Monte–Carlo simulation, the phantom and the *in–vivo* experiments. Due to space constraints, we only show 8 of the 32 channels of the central slice for each dataset.

- *Water phantom acquisition:* A MR phantom sphere with solution (GE Medical Systems, Milwaukee, WI) was scanned in a 32-channel head coil on a 3.0T scanner (MR750, GE Healthcare, Waukesha, WI). A spoiled gradient-echo acquisition with 100 realizations of the same fully-encoded $\boldsymbol{k}$–space sampling was used. Acquisition parameters included: coronal view, TE/TR=0.96/3.69ms, flip angle=12°, FOV=22×22×30.7cm^3, acquisition matrix size=60×60×32, bandwidth= ±62.5KHz. We corrected for B_0 field drift related phase variations and magnitude decay (Vos et al., 2016) by a pre-processing step. First we estimated the phase-shift between realizations from the center of the $\boldsymbol{k}$–space as a cubic function of time and removed it afterwards. And, second, we estimated the magnitude-decay in the $\boldsymbol{k}$–space as a linear function and substracted it in order not to affect the noise.

- *In vivo acquisition:* in order to assess the feasibility of the proposed method, after obtaining the approval fo the local institutional review board (IRB), a volunteer was scanned in a 32-channel head coil on a 3.0T scanner (MR750, GE Healthcare, Waukesha, WI). A spoiled gradient-echo acquisition of a fully-encoded $\boldsymbol{k}$–space sampling was used. Acquisition parameters included: coronal view, TE/TR=2.2/5.7ms, flip angle=12°,FOV=22×22×22cm^3, matrix size=220×220×220, bandwidth= ±62.5KHz.

From each of the first two datasets, four sub-sampling scenarios were considered (see Fig.4.4).

(i). A CAIPIRINHA type sub-sampling pattern along the two PE dimensions with an acceleration rate R=2. A fully–sampled rectangular ACS region in the center of the

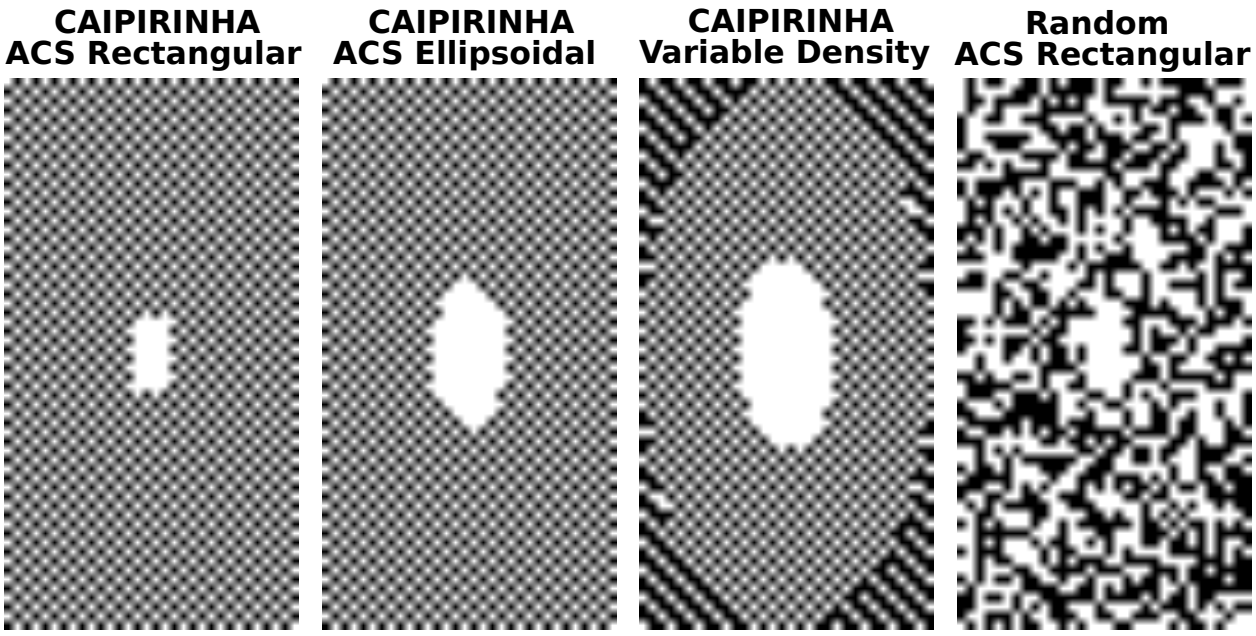

Figure 4.4: Sub-sampling patterns used for the phantom experiments, all containing a central fully–sampled ACS region.

$\boldsymbol{k}$–space was acquired containing 8×4 fully sampled lines along the FE direction (CAIPI-Rect). We used a GRAPPA kernel of size $[3, 3, 3]$ to reconstruct the missing data.

(ii). Usage of an ellipsoidal ACS region over the same CAIPIRINHA type sub-sampling pattern (CAIPI-Ellip). The axis of the ellipsoidal ACS were of size 8×4 and the reconstruction kernel size was kept at $[3, 3, 3]$.

(iii). A variable density sampling (CAIPI-VD) as in (Heidemann et al., 2007). We divided the $\boldsymbol{k}$–space into three ellipsoidal regions. The center was a fully–sampled ellipsoid (ACS region) with axis 8×4. Then, a middle region with axis 60×32 was sub-sampled at a rate $R = 2$ and finally the outer region covering the corners was sub-sampled at a rate $R = 3$. The inner region was reconstructed with a $[3, 3, 3]$ kernel, whereas for the outer region we used a $[5, 5, 3]$ kernel.

(iv). A random sub-sampling pattern along the two PE directions with acceleration rate $R = 2$, preserving an ACS region of size 8×4. Again, a $[3, 3, 3]$ kernel was used for reconstruction. The method will be hereafter referred to as Random.

For these experiments, g-factor maps were obtained in two different ways. First, we followed a MC approach to calculate the sample standard deviation for each pixel from all the realizations for both the unaccelerated and the accelerated images. The g-factor maps are computed using Eq.2.19. Second, we directly computed the g-factor maps using the proposed $\boldsymbol{k}$–space method as in Eq.3.32. Our method requires as inputs the $\boldsymbol{k}$–space noise matrices $\boldsymbol{\Gamma}$

and $\mathbf{C}$ in the sub-sampled images, the GRAPPA kernel and the coil-combination vector. For the simulated scenario, these parameters were known. For the phantom acquisition, these matrices were estimated as the sample covariance matrices obtained across realizations and further averaged across locations (due to stationarity).

Finally, in order to prove the applicability to clinical scenarios, we studied the impact of the sub-sampling pattern using the high–resolution *in–vivo* scan. Two different sub-sampling patterns were considered at an acceleration rate of $R = 6$: Rectangular type ($R = 3 \times 2$) and CAIPIRINHA-type ($R = 3 \times 2$), as shown in Fig.4.7. The acquired $\boldsymbol{k}$–space was sub-sampled preserving a 16×16 fully sampled ACS region. In this case we needed to estimate the noise from a single repetition, which we did using the Mean Absolute Deviation estimator after a wavelet filtering the preserves the high-frequency band, which is supposed to contain only the noise, as in (Aja Fernández and Vegas Sánchez-Ferrero, 2016).

All image reconstruction and g-factor maps estimation were performed using Matlab and run on a shared computer with an Intel®Xeon®CPU E5-2695 v3 @2.30GHz Processor and 110 GB of RAM. In the spirit of reproducible research, we provide a software package including both the datasets and the code that we used, allowing to reproduce all the results included in this manuscript. This package can be downloaded from `http://www.lpi.tel.uva.es/3dgrappa`.

4.4 Results

Fig.4.5 shows the g-factor for each pixel at the intermediate slice of the 3D simulated dataset reconstructed using 3D-GRAPPA with the four sub-sampling patterns described above. The first two columns show the MC empirical g-factor maps and the analytical $\boldsymbol{k}$–space g-factor maps obtained with our method, respectively. The third colum shows the differences between the analytical estimation against the MC reference. Finally, the last column shows a boxplot diagram of absolute error distribution along the third dimension (along the slices). Fig.4.6 shows the same results obtained for the water phantom.

Fig.4.7 shows both the g-factor maps obtained with different sub-sampling patterns as well as the coil-combined images together with the subsampling patterns used along the two PE directions. This experiment makes used of the *in-vivo* data described above.

Finally, as for computational load, we include the computation time to obtain the g-factor maps with the synthetic phantom for each of the sub-sampling patterns previously described in Table 4.1. As expected, higher numbers of rows with unique patterns result in higher computational time requirements.

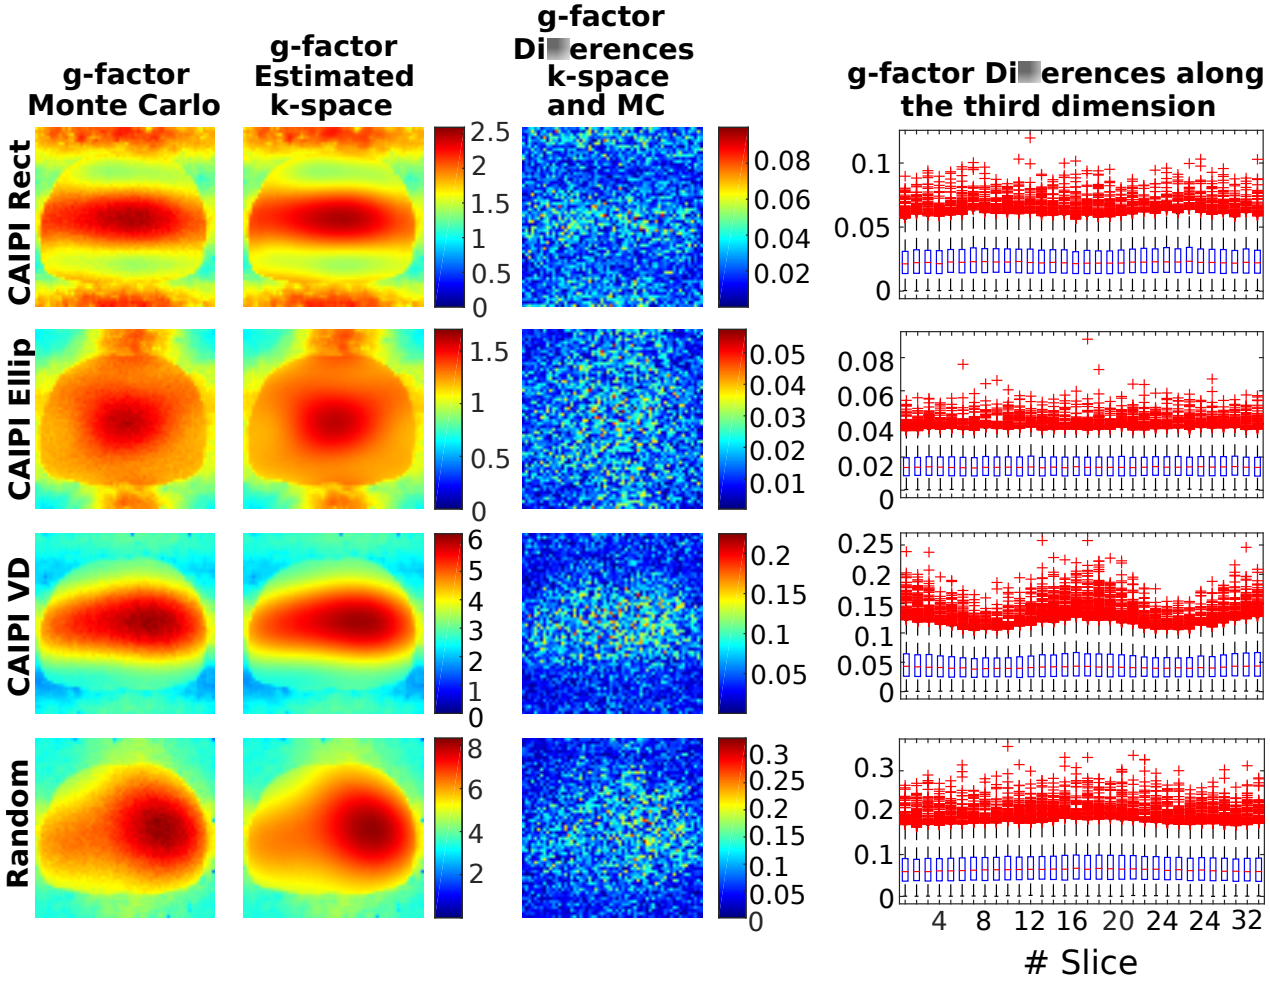

Figure 4.5: g-factor maps for the synthetic phantom obtained through a MC strategy (first column) and estimated using the proposed $\boldsymbol{k}$–space method (second column) for the intermediate slice of the 3D volume. The third column shows the absolute differences between the $\boldsymbol{k}$–space method against the MC estimation. The last column shows the absolute error distribution along the slice dimension. The central mark indicates the median within the slice, and the bottom and top edges of the box indicate the 25th and 75th percentiles, respectively. The whiskers extend to the most extreme data points not considered outliers, and the outliers are plotted individually using the '+' red symbol. Several scenarios are studied: CAIPIRINHA–type with rectangular ACS region (first row), CAIPIRINHA–type with ellipsoidal ACS region (second row), CAIPIRINHA–type with Variable Density sampling (third row) and random sub-sampling with rectangular ACS region (fourth row).

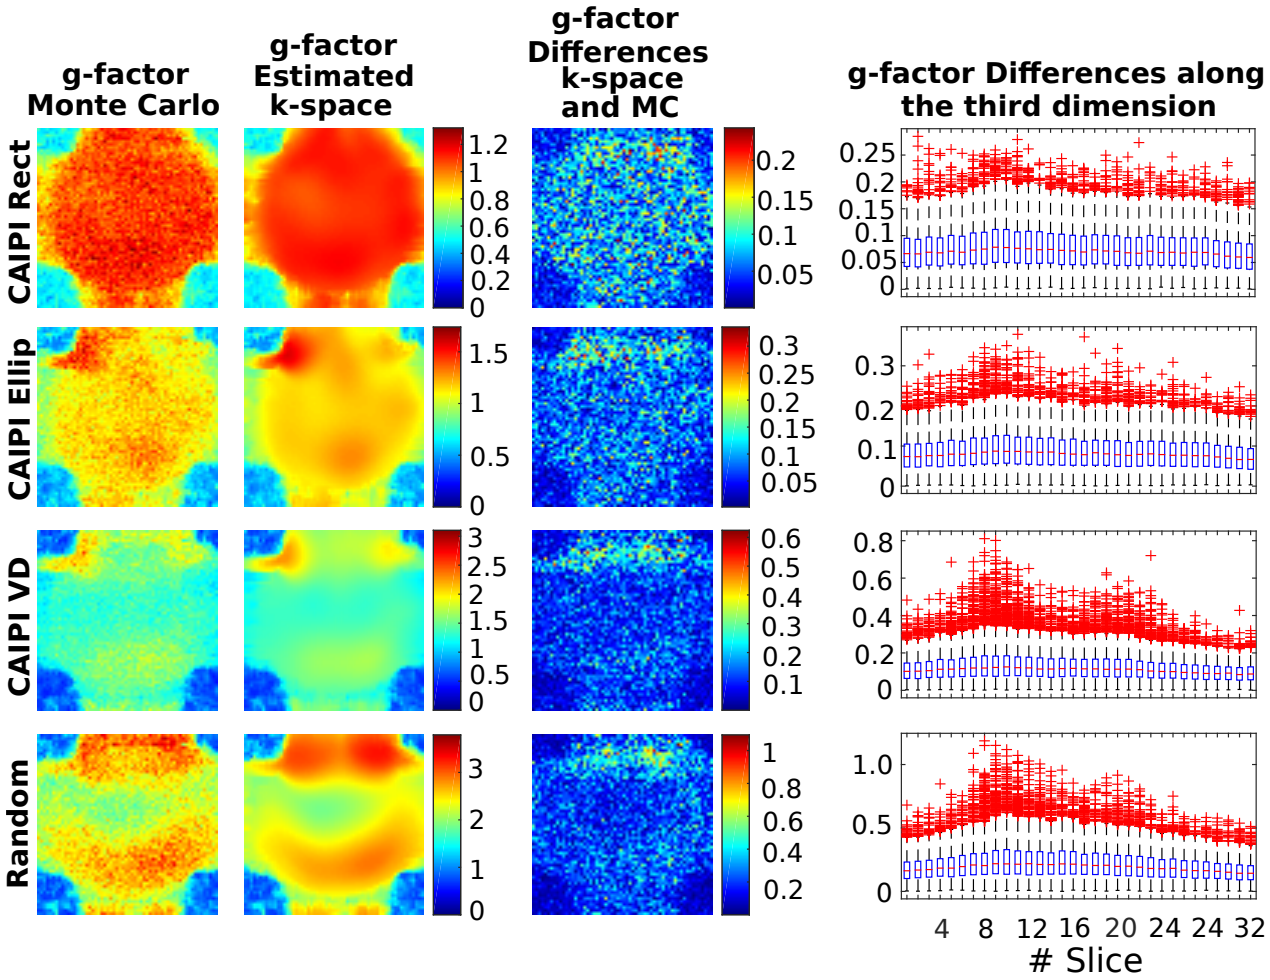

Figure 4.6: g-factor maps for the scanned water phantom obtained through a MC strategy (first column) and estimated using the proposed $\boldsymbol{k}$–space method (second column). The third column shows the absolute differences between the $\boldsymbol{k}$–space method against the MC estimation. The last column shows the absolute error distribution along the slice dimension. Same scenarios as for the synthetic phantom are shown.

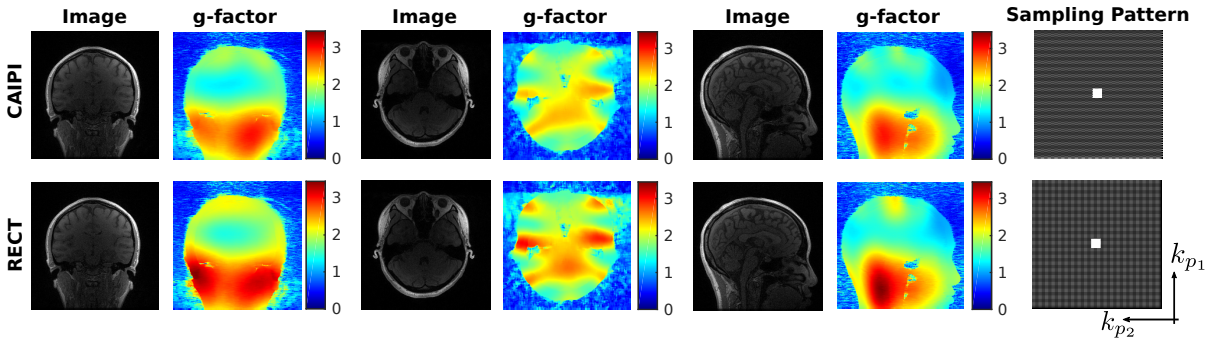

Figure 4.7: Calculated g-factor maps for the *in–vivo* brain experiment obtained through the proposed $\boldsymbol{k}$–space method. Two different subsampling patterns are studied at the same acceleration rate $R = 6$: Rectangular type ($R = 3 \times 2$) and CAIPIRINHA-type ($R = 3 \times 2$). Reconstructed images after coil–combination and g-factor maps are shown for both scenarios including a view from each orthogonal plane in the 3D in-vivo dataset. Last column shows the subsampling pattern used along the two PE dimensions.

Table 4.1: Computation times (seconds)

CAIPI Rect	CAIPI Ellip	CAIPI VD	Random
121.32	215.28	688.71	629.08

4.5 Discussion

In this work, we have proposed an analytical method for noise characterization in 3D-GRAPPA reconstructions that provides an exact characterization of noise under two common assumptions: stationarity and uncorrelation of noise in the acquired $\boldsymbol{k}$–space. This is a challenging problem due to the very large size of the covariance matrices involved in the noise propagation, which is overcome by exploiting the extensive symmetries and the separability in the reconstruction steps.

The proposed analytical method is able to provide an accurate estimation of the g-factor maps under all the sub-sampling scenarios tested, both for the synthetic and real phantom. The differences observed between the g-factor maps corresponding to the analytical approach and the MC approach in the synthetic experiments can be atributed to the approximate behaviour inherent to MC methods, since they need an infinite number of realizations to be exact. With the aim of showing the applicability of our method to clinical scenarios, a high–resolution *in–vivo* experiment was carried out comparing CAIPI and Rectangular type sampling patterns. As expected, it can be observed that the CAIPI acquisition has lower noise amplification across the axial slice, especially in the posterior region. Importantly, the proposed method may have application for the optimization of acquisition protocols (including the choice of acceleration and sub-sampling scheme in 3D imaging), as well as for the development of optimized post-processing methods that rely on the noise statistics at each voxel.

This study presents some limitations. First, the analysis focuses on GRAPPA reconstruction followed by a linear coil combination such as (Walsh et al., 2000). If SoS is used instead, (Aja-Fernández et al., 2011) have shown that the noise distribution can be approximated as a non-Central χ distribution, whose effective parameters need to be calculated. However, noise characterization of SoS coil combination is straightforward from the final covariance matrices $\mathbf{\Gamma}_6(x)$ and $\mathbf{C}_6(x)$ by simply applying them to Eq.(17-18) from (Aja-Fernández et al., 2011). Second, as in (Rabanillo et al., 2017), we have assumed that the kernel used for reconstruction is independent of noise in the acquired image, which is not strictly true when the kernel is autocalibrated from the data instead of pre-calibrated from a separate pre-scan. However, this assumption is expected to be a good approximation due to the typical overdetermination of GRAPPA weights estimation from the ACS region, and for this reason it has become a common assumption in the literature. Similarly, we have also assumed that the vector used

for coil combination is noise–independent. Third, our analysis is focused on providing the final voxel–wise g-factor maps. However, for certain applications such as optimal filtering in (Sprenger et al., 2016) it may be useful to additionally characterize the noise correlation between different locations in the reconstructed image. Importantly, our proposed method can be extended to provide the correlation between any two points in the final image, thus allowing to obtain cross-correlation maps.

Although this manuscript does not directly consider non-Cartesian 3D acquisitions, the proposed techniques can also be extended to non-Cartesian acquisitions (eg: radial or spiral trajectories) if preceded by a linear interpolation into a Cartesian grid. Such an interpolation would introduce correlations in the sub-sampled $\boldsymbol{k}$–space and may introduce non-stationarity if the interpolation weights are not uniform across the $\boldsymbol{k}$–space. In such scenarios, it would suffice to take into account this non-stationarity when building the correlation matrices of the sub-sampled Cartesian $\boldsymbol{k}$–space at the first step of our analysis.

4.6 Conclusion

We have extended the $\boldsymbol{k}$–space noise analysis proposed in (Rabanillo et al., 2017) to characterize noise in 3D-GRAPPA reconstructions of volumes sub-sampled along two PE directions. By directly operating in $\boldsymbol{k}$–space our analysis is able to exactly characterize the noise under the assumptions of stationarity and uncorrelation in the acquired $\boldsymbol{k}$–space. By exploiting both the symmetry and separability in the reconstruction steps, the proposed method enables efficient and exact noise characterization in 3D-GRAPPA .

CHAPTER

FIVE

JOINT IMAGE RECONSTRUCTION AND PHASE CORRUPTION MAPS ESTIMATION IN MULTI-SHOT ECHO PLANAR IMAGING

Abstract

Purpose: To develop a navigator-free method for Multi–Shot acquisitions in dMRI EPI under microscopic intra–shot subject motion that both reconstructs the image and estimates the induced shot–to–shot phase discrepancies.

Theory and Methods: We incorporate the well-known fact that under rigid motion, the phase corruption maps behave linearly and therefore they can be fully characterized by three parameters per shot. By incorporating this prior knowledge into a maximum likelihood formulation, we can alternate between image reconstruction and phase estimation following a greedy fashion. Finally, we extend this methodology to the more complex scenario of non–rigid motion by modelling the resulting phase maps by means of B–splines.

Results: We use synthetically generated data to validate our algorithm under different levels of subject motion, noise corruption and number of shots, both under the assumption of linear phase maps and smooth phase maps. We use as well a real phantom and realistic phases from in–vivo ss-EPI data to validate our proposals. We observe that the proposed method is able to reconstruct the images for up to 8 shots under the assumption of linear or smooth phase maps.

Conclusions: By incorporating the prior knowledge of rigid–motion into the theoretical maximum likelihood formulation, we are able to reconstruct diffusion multi–shot

images up to 8 shots under subject motion–induced phase corruption. Reaching the global optimum of the resulting non–convex function is achieved by means of a proper initialization based on a shot–by–shot SENSE reconstruction and a further greedy algorithm that iterates between the image reconstruction and phase estimation problems.

5.1 Introduction

dMRI is a non-invasive technique that allows for quantification of water molecules diffusion in biological tissues. Due to its speed and insensitivity to motion, ss-EPI has become the most commonly used sequence in dMRI. However, its long echo-trains result in significant image distortions due to B0-field inhomogeneities, eddy currents, T2*-blurring or chemical shift (Jaermann et al., 2006). In order to achieve higher resolution and less distortion using ss-EPI, different alternatives have been studied. One option would be to combine it with PI, but the acceleration rate is limited by the noise amplification. Another possibility would be to use rFOV techniques (Saritas et al., 2008), where only a portion of the object is excited (Pauly et al., 1989), leading to shorter echo-trains since a lower resolution $\boldsymbol{k}$–space suffices to reconstruct the image. The negative aspect is that this technique is uniquely useful when we are only interested in a specific part of the object (Saritas et al., 2008).

A different approach to reduce the EPI echo train is ms–EPI, as in (Anderson and Gore, 1994; Butts et al., 1996; Miller and Pauly, 2003). This consists in segmenting the readout into multiple shots, which can be done in an interleaved (Atkinson et al., 2006; Jeong et al., 2013) or sequential (Holdsworth et al., 2008; Porter and Heidemann, 2009) fashion. This reduces the length of the echo-trains and diminishes, as a result, the distortion caused by the aforementioned effects. However, physiological motion during the application of the diffusion sensitizing gradients might cause phase inconsistencies between different shots that cause additional ghosting artifacts. For this reason, ms–EPI sequences need to incorporate a phase-correction technique that accounts for these phase mismatches between shots.

Subject-induced phase maps in ms–EPI behave similarly to sensitivity maps in PI in the sense that they can be modelled as a pixel-wise multiplication between the original image and corresponding phase map. This analogy has motivated the translation of PI techniques to reconstruct the multi-shot data under shot-to-shot phase discrepancies, where the shots would play the role of the coils. The main difference, however, is that in PI the data is sampled at the exact same exact location for all the coils, whereas in ms–EPI, each shot samples a different trajectory. Consequently, although many well-known ms–EPI techniques follow different PI methodologies, this slight discrepancy in the trajectories needs to be incorporated into the formulation.

State-of-the-art methods can be classified into three different groups: extra–navigated methods, self–navigated methods and navigator–free methods. The first group relies on the assumption that the image and the callibration (navigator) data share the same phase information. Consequently, the phase maps can be measured from navigator data and incorporated into a SENSE-type reconstruction (Van et al., 2011). Nevertheless, it should be noticed that the image–echo and the navigator–echo might be acquired with a different bandwith along the phase–encoding direction, which results in different geometric distortions between the image and the navigator.

These mismatches can be compensated by means of a registration step, but it will increase the computation complexity and might still introduce errors. As an alternative, self–navigated methods acquire the calibration data within the imaging sequence, and further reconstruct the image based on different PI techniques such as GRAPPA (Liu et al., 2016), SPIRiT (Dong et al., 2018), LORAKS (Lobos et al., 2018), PROPELLER (Pipe et al., 2002) or SNAILS (Liu et al., 2004). These sequences are thus inmune to image–navigator mismatches, but the need to oversample the center of the $\boldsymbol{k}$–space renders them time-inefficient.

Finally, navigator-free techniques have been developed with the aim of retrospectively correcting for ghosting artifacts from the images themselves. Furthermore, the absence of navigator echoes not only increases time efficiency but also avoids the geometrical mismatches between the estimated phase maps and the shot images. In order to estimate the phase maps, different priors have been incorporated, such as phase smoothness (Chen et al., 2013), rigid–motion (Van et al., 2011) or low-rankness (Mani et al., 2017). However, phase estimation becomes challenging due to noise amplification as we increase the number of shots if we rely on PI techniques since the number of shots is equivalent to the subsampling rate of each shot.

In this paper, we present a navigator-free method to reconstruct ms–EPI images under the assumption of intra-shot rigid motion. Anderson et. al. proved that the resulting phase errors are linear in the image space, which turn into shifts and constant offsets in $\boldsymbol{k}$–space (Anderson and Gore, 1994). Incorporating this knowledge into the model, we present an image reconstruction procedure based on the optimization of a functional used to jointly estimate the $\boldsymbol{k}$–space shifts and offsets and the diffusion image by exploiting the sensitivity encoding (Pruessmann et al., 1999) redundancy provided by the coil array. This philosophy has already proven succesful for the correction between even and odd lines, where corruption is modelled only along the fully–sampled readout direction (Ianni J.D., 2018). Furthermore, we extend this methodology to the case of non–rigid motion, assuming the derived non–linear phase maps vary smoothly and can be modelled by means of B–splines.

5.2 Theory

5.2.1 Overview: Diffusion MRI and phase accumulation due to rigid motion

dMRI is aimed at probing the microstructure of biological tissues by capturing the local diffusion of water molecules, which is restricted or hindered by anatomical structures. Initially, diffusion was considered a source of degration since it resulted in a signal decay due to the different dephasing observed by the moving particles in the presence of $\vec{\mathbf{B}}_{\mathbf{0}}$-field inhomogeneities. However, (Stejskal and Tanner, 1965) were able to use the observed decay to their advantage in order to measure the motion of water molecules inside the tissues. In their seminar paper (Stejskal and Tanner, 1965), they presented a novel pulsed-gradient spin echo MR sequence. Named after them, it still remains the standard scheme for dMRI (see Fig.2.11).

The main novelty is the introduction of the diffusion sensitized gradients $\vec{\mathbf{G}}$. During their application, the position $\vec{\mathbf{R}}$ of the water molecules is encoded in its phase $\phi(\vec{\mathbf{R}})$:

$$\phi(\vec{\mathbf{R}}) = \gamma \int_0^{\delta} \vec{\mathbf{G}}^T(t)\vec{\mathbf{R}}(t)\, dt, \tag{5.1}$$

where refers to γ the gyromagnetic ratio of the water proton. When the gradients are constant as in Fig.2.11, what we encode in the phase of the spins is its average position $E\{\vec{\mathbf{R}}\}$ during the application of the gradients:

$$\phi(\vec{\mathbf{R}}) = \gamma\delta\vec{\mathbf{G}}^T E\{\vec{\mathbf{R}}\}. \tag{5.2}$$

Furthermore, the diffusion gradients are applied twice with an 180° RF–pulse in between, which is equivalent to applying them with inverse polarity. Therefore, the overall phase of the water molecules will be the difference of the accumulated phase during the application of both gradients:

$$\Delta\phi(\vec{\mathbf{R}}) = \gamma\delta\vec{\mathbf{G}}^T\Delta\vec{\mathbf{R}}, \tag{5.3}$$

where $\Delta\vec{\mathbf{R}} = E\{\vec{\mathbf{R}}_2\} - E\{\vec{\mathbf{R}}_1\}$ contains the vector difference between the average position of a spin during the application of the two diffusion pulses.

To understand the impact of motion during the application of the diffusion gradients, let us define the following coordinate systems, as in (Anderson and Gore, 1994):

- $[\hat{\mathbf{X}}, \hat{\mathbf{Y}}, \hat{\mathbf{Z}}]$: this system is defined by the imaging gradients, where X refers to the readout or FE direction, Y refers to the phase–encoding direction and Z refers to the slice selection direction. The position of any spin can be specified by its coordinates in this system as:

$$\vec{\mathbf{R}} = X\hat{\mathbf{X}} + Y\hat{\mathbf{Y}} + Z\hat{\mathbf{Z}}. \tag{5.4}$$

- $[\hat{\mathbf{x}}, \hat{\mathbf{y}}, \hat{\mathbf{z}}]$: this is the body–centered system that moves together with the patient, where x refers to the anterior/posterior direction, y refers to the right/left direction and z refers to the inferior/superior direction. The origin of this system has coordinates $\vec{\mathbf{R}}_0 = [X_0, Y_0, Z_0]$ in the system $[\hat{\mathbf{X}}, \hat{\mathbf{Y}}, \hat{\mathbf{Z}}]$. Equivalently, the position of a spin can be defined as:

$$\vec{\mathbf{r}} = x\hat{\mathbf{x}} + y\hat{\mathbf{y}} + z\hat{\mathbf{z}}. \tag{5.5}$$

  and we can relate these two coordinate systems by:

$$\vec{\mathbf{R}} = X_0\hat{\mathbf{X}} + Y_0\hat{\mathbf{Y}} + Z_0\hat{\mathbf{Z}} + \vec{\mathbf{r}}. \tag{5.6}$$

dMRI is designed to be sensitive to the motion of water molecules, which occurs in the scale of micrometers. As a consequence, any bulk motion that takes place in this scale will result in an undesired phase accumulation. Importantly, due to the short interval during which we apply the diffusion gradients, we assume that the bulk motion scale is several orders of magnitude smaller than the resolution of the image. Or equivalently, we can work under the assumption of intra–voxel motion.

Hence, for sufficiently small displacements, we can describe any change in the position of a spin in the gradient frame by:

$$\vec{\mathbf{R}}(t) = \vec{\mathbf{R}}_0(t) + \vec{\boldsymbol{\theta}}(t) \times \vec{\mathbf{r}}, \tag{5.7}$$

where $\vec{\mathbf{R}}_0(t)$ describes the change in the body frame origin (translation) and $\vec{\boldsymbol{\theta}}(t)$ defines the change in the body frame orientation (rotation). Here we have assumed that we can describe a rotation as a cross product under the assumption of small rotations (Goldstein, 1980).

The MR signal $\hat{s}(k_x, k_y)$ measured in $\boldsymbol{k}$–space under these conditions is given by:

$$\hat{s}(k_x, k_y) = \int p(x, y) e^{i\varphi(x,y,z)} e^{i(k_x x + k_y y)} dx dy, \tag{5.8}$$

where $p(x, y)$ is the MR signal weighted by diffusion and relaxation decay and $\varphi(x, y, z)$ is the accumulated phase at each location during the application of the diffusion gradients. This accumulated phase is given by

$$\varphi(x, y, z) = \gamma \int \vec{\mathbf{G}}^T(t)\vec{\mathbf{R}}(t)\, dt, \tag{5.9}$$

or equivalently:

$$\begin{aligned} \varphi(x, y, z) &= \gamma \int \vec{\mathbf{G}}^T(t)\vec{\mathbf{R}}_0(t)\, dt + \left[\gamma \int \vec{\mathbf{G}}^T(t) \times \vec{\boldsymbol{\theta}}(t)\, dt\right] \vec{\mathbf{r}} \\ &= \Phi_x + \Phi_y + \Phi_z + \vec{\mathbf{d}}\vec{\mathbf{k}}\vec{\mathbf{r}}, \end{aligned} \tag{5.10}$$

where we can write the phase errors due to translation as:

$$\begin{aligned}\Phi_x &= \gamma \int G_x(t) X_0(t)\, dt \\ \Phi_y &= \gamma \int G_y(t) Y_0(t)\, dt \\ \Phi_z &= \gamma \int G_z(t) Z_0(t)\, dt\end{aligned} \tag{5.11}$$

and the phase errors due to rotation as:

$$\begin{aligned}dk_x &= \gamma \int \left[G_y(t)\theta_z(t) - G_z(t)\theta_y(t)\right]\, dt \\ dk_y &= \gamma \int \left[G_z(t)\theta_x(t) - G_x(t)\theta_z(t)\right]\, dt \\ dk_z &= \gamma \int \left[G_x(t)\theta_y(t) - G_y(t)\theta_x(t)\right]\, dt.\end{aligned} \tag{5.12}$$

Finally, we can write the measured signal as:

$$\hat{s}(k_x, k_y) = e^{i(\Phi_x+\Phi_y+\Phi_z)} \cdot \left[\int e^{i\, dk_z z} dz\right] \int p(x,y) e^{i[(k_x+dk_x)x+(k_y+dk_y)y]} dxdy, \tag{5.13}$$

which we can relate to the ideal motion–free signal:

$$s(k_x, k_y) = \int p(x,y) e^{i(k_x x + k_y y)} dxdy. \tag{5.14}$$

By using the FT properties, we can write:

$$\hat{s}(k_x, k_y) = e^{i(\Phi_x+\Phi_y+\Phi_z)} \cdot \left[\int e^{i\, dk_z \cdot z} dz\right] \cdot s(k_x + dk_x, k_y + dk_y). \tag{5.15}$$

Using this derivation, and under the assumption of small displacements, we have shown that rigid motion during the application of the diffusion sensitized gradients results in a shift in $\boldsymbol{k}$–space accompanied by a multiplication by a complex constant. This relation can be translated into the image–space, where we will observe a linear phase corruption of the desired image:

$$\hat{p}(x,y) = \Phi_0 \cdot p(x,y) \cdot e^{i(dk_x x + dk_y y)}, \tag{5.16}$$

which is well characterized by three parameters:

$$\hat{p}(x,y) = p(x,y) \cdot e^{i(\theta_0 + \theta_1 x + \theta_2 y)}, \tag{5.17}$$

given by:

$$\begin{aligned}\theta_0 &= \Phi_x + \Phi_y + \Phi_z - i\ \log\left(\int e^{i\, dk_z z} dz\right) \\ \theta_1 &= dk_x \\ \theta_2 &= dk_y.\end{aligned}$$

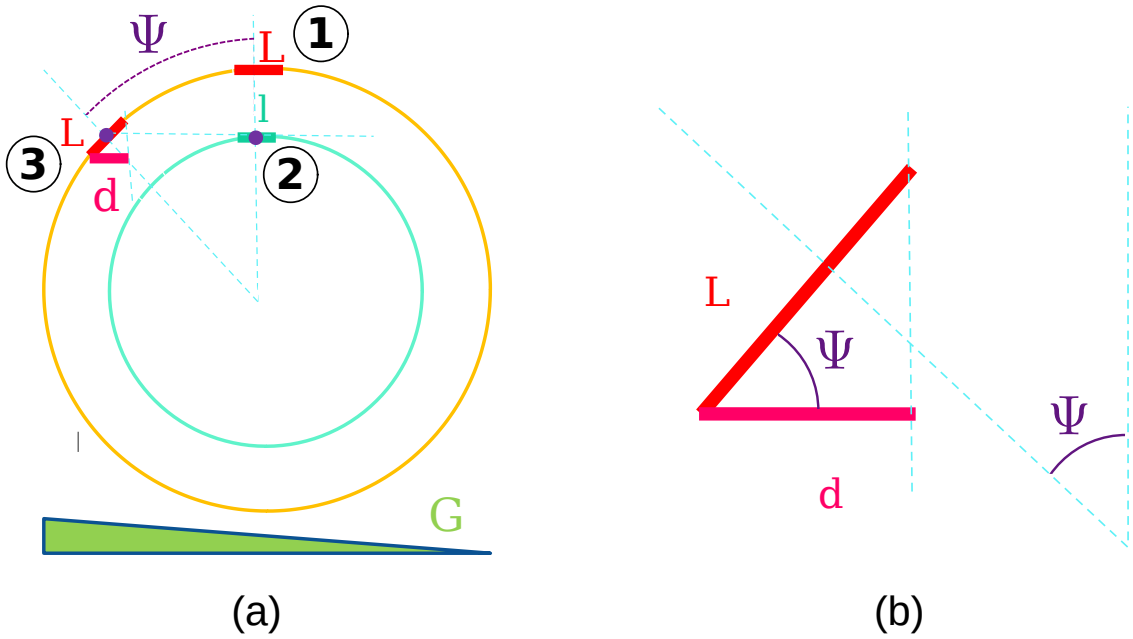

Figure 5.1: Geometric description of the effects of rigid motion in the phase of the acquired image. (a) Assuming an horizontal diffusion gradient, points 1 and 2, in a vertical line, acquire a phase that is directly proportional to their distance to the center of rotation. (b) On the other hand, points 2 and 3, in an horizontal line, acquire exactly the same phase even though they are at a different distance from the center of rotation. This is because what matters in diffusion is the movement along the diffusion direction, and as we can see using trigonometry relationships and triangle similarity, both points traverse the same horizontal distance.

Alternatively, we can explain the effects of rigid motion using geometry, as depicted in Fig.5.1. For this purpose, it is important to notice that the diffusion sequence is sensitive to the motion of the spins along the direction defined by the diffusion gradient. Furthermore, what matters is not the trajectory followed by the spins, but its initial and final position with respect to this direction, as observed in Eq.5.2.

Consequently, when the whole object experiences a translation, all the spins undergo exactly the same displacement along the diffusion direction, and therefore they accumulate the exact same phase error, given by:

$$\theta_0 = \gamma\delta\vec{\mathbf{G}}^T(t)\Delta\vec{\mathbf{r}}, \tag{5.18}$$

For the case of a rotation, we use once again the assumption of intra–voxel motion. In order to better understand the effect of a rotation, let us assume we apply a gradient along the horizontal direction as in Fig.5.1.(b). In a rotation, the displacement experienced by the spins is proportional to their distance to the center of rotation. However, as we have already mentioned, what matters is the displacement relative to the diffusion gradient direction, not the absolute displacement.

In Fig.5.1 we first study the relative displacement of two points, 1 and 2, located on a line perpendicular to the direction of the diffusion gradient. Since their distance to the center of rotation is different, their displacement L and l will differ as well, leading to the following phase eror for each of them:

$$\begin{aligned} \varphi_1 &= \gamma\delta|\vec{\mathbf{G}}|L \\ \varphi_2 &= \gamma\delta|\vec{\mathbf{G}}|l. \end{aligned} \tag{5.19}$$

Once again, under the assumption of small rotations, we can consider that the arc of the circumference is a straight line. Using similar triangle properties, we can relate the arc L and l of the circunferences and their radius r and R:

$$\frac{l}{L} = \frac{r}{R}, \tag{5.20}$$

which shows that the phase accumulated error grows linearly in the direction perpendicular to the applied diffusion gradient.

Moreover, we can study as well the relative displacement of two points, 2 and 3, placed in a line parallel to the gradient, that therefore are at a different distance from the center of rotation. In this example we suppose the point 3 is at the same distance from the rotation center than point 1, forming and angle Ψ. From trigonometry we know that its cosine can be computed from the radius of the two circumferences:

$$\cos(\Psi) = \frac{r}{R}. \tag{5.21}$$

Furthermore, we observe, as shown in Fig.5.1.(b), that this triangle is similar to the one formed with the displacement experienced by point 3 and its projection along the diffusion gradient direcion. Consequently, we have:

$$\cos(\Psi) = \frac{r}{R} = \frac{d}{L}, \tag{5.22}$$

which combined with Eq.5.20 shows that the relative displacements d and L along the gradient direction are equal, and therefore the accumulated phase errors $\varphi_2 = \varphi_3$. Hence, the phase accumulated error does not vary along the direction parallel to the gradient. In conclusion, we observe geometrically that rigid motion results in a linear phase map as previously proven mathematically.

5.2.2 Multi–shot EPI for Diffusion MRI

dMRI requires the acquisition of multiple volumes with different diffusion weighting in order to estimate the parameters characterizing the undergoing random process that describes the motion of water molecules. Further, the number of volumes needed depends on the assumptions about the nature of the diffusion process (Basser et al., 1994; Ozarslan and Mareci, 2003; Jensen et al., 2005; Wedeen et al., 2008; Tuch, 2004; Assaf and Basser, 2005; Alexander et al., 2010; Zhang et al., 2012; Ozarslan et al., 2013). As a consequence, the objective in dMRI is to minimize the scan time so that a higher resolution in the space of measurements, the so–called $\vec{\mathbf{q}}$–space, can be achieved for the same overal acquisition duration *.

ss–EPI has established itself as the standard modality for dMRI since it allows to acquire the whole $\boldsymbol{k}$–space in a single excitation (Mansfield, 2003). However, the long echo trains give rise to different types or artifacts such as geometrical distortions related to $\vec{\mathbf{B}}_0$–field inhomogeneities and eddy currents, chemical shift or T2*–blurring. In order to mitigate the distortion and increase the spatial resolution, a common strategy consists in segmenting the acquisition into multiple shots (ms–EPI), which reduces the duration of the echoes (see Fig.2.10).

Nevertheless, as we have previously explained, dMRI is very sensitive to bulk motion due to the presence of the diffusion gradients, which results in undesired phase accumulation. This is not a problem in ss–EPI since we only aply once the diffusion sensitized gradients, but it poses a major challenge in ms–EPI, where even the slightest differences in patient motion during the application of the gradients corresponding to each shot result in phase discrepancies. If not corrected for, they will introduce ghosting in the reconstructed image, notably degrading its quality.

*The $\vec{\mathbf{q}}$–space is a Transformed domain related to the $\vec{\mathbf{R}}$–space where the probability distribution of the displacement of water molecules can be computed. These two spaces are related through the FT.

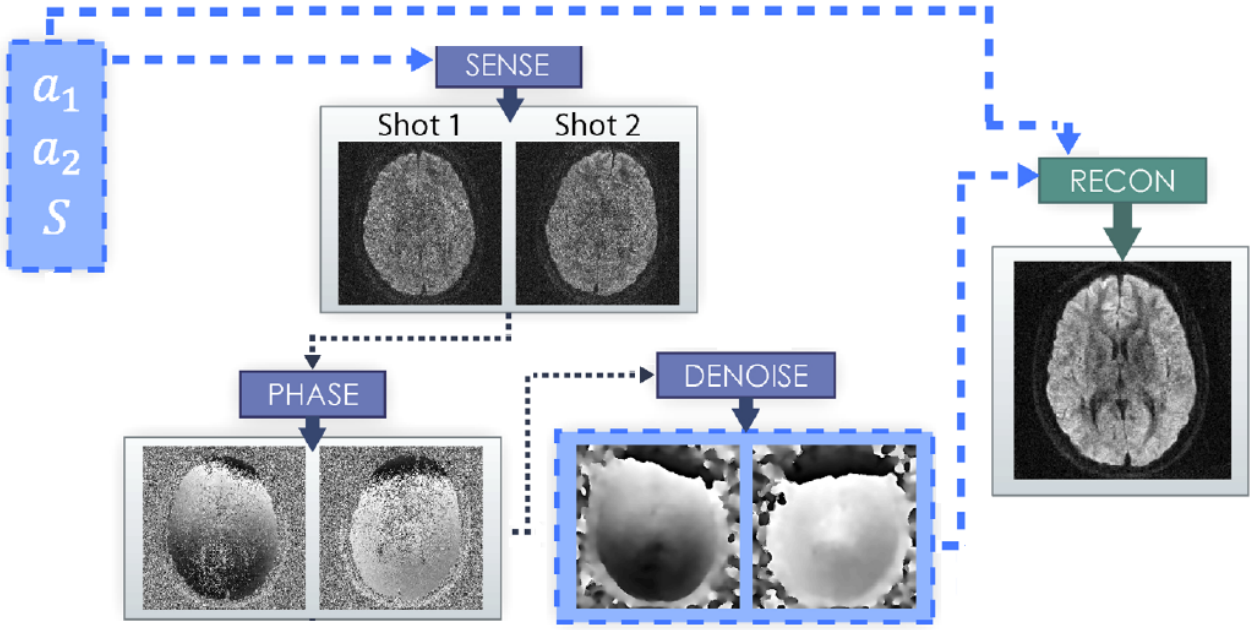

Figure 5.2: Scheme of MUSE reconstruction. First each shot is reconstructed from multi–channel data using SENSE. Then, the phase of each of them is retrieved and filtered using Total Variation to denoise it. Finally, the phase maps are incorporated into a SENSE–type reconstruction that combines the information from all the shots and coils.

Ideally, the objective would be to estimate the phase discrepancy maps from the very data itself, avoiding the need to acquire navigator data and the resulting time inefficiency. The cost of not acquiring callibration data is having to impose some prior knowledge about the phase maps corrupting these shots. Different models have been studied in the literature, such as phase linearity (Van et al., 2011), phase smoothness (Chen et al., 2013; Lobos et al., 2018) and limited image support (Mani et al., 2017). In this work we will present a methodology that deals both with linear phase maps derived from rigid motion and smooth non–linear phase maps arising from more complex motion such as pulsatile motion.

We will compare our method with the MUltiplexed Sensitivity Encoding (MUSE) technique (Chen et al., 2013), which imposes the prior of smooth phase maps. In Fig.5.2 we depict the pipeline followed by this alternative scheme. Importantly, data is acquired with a multi–channel antena, sampling the imaging data simultaneously at different coils. First of all, all the shots are reconstructed by means of SENSE (Pruessmann et al., 1999) to obtain a fully–sampled image for each of them. Then, a Total–Variation (TV) (Rudin et al., 1992) filter is applied to recover a denoised estimate of the phase maps. Finally, these phase maps are incorporated into a SENSE type reconstruction (Pruessmann et al., 2001) using the data from all coils and shots, obtaining the reconstructed final image.

5.2.3 Problem formulation

Rigid motion

Under the assumption of rigid motion during the application of diffusion-sensitized gradients, the phase corruption for each shot becomes linear and can be characterized by three parameters in 2D imaging. Thus, the reconstruction for parallel multishot imaging can be performed by means of the Maximum Likelihood formulation in matrix form as:

$$(\vec{\mathbf{x}}^*, \vec{\boldsymbol{\theta}}^*) = \underset{\vec{\mathbf{x}}, \vec{\boldsymbol{\theta}}^*}{\operatorname{argmin}} \left\| \mathbf{A}\boldsymbol{\mathcal{F}}\mathbf{S}\boldsymbol{\mathcal{P}}(\vec{\boldsymbol{\theta}}^*)\vec{\mathbf{x}} - \vec{\mathbf{y}} \right\|_2^2. \tag{5.23}$$

The aim is to reconstruct a 2D image $\vec{\mathbf{x}}$ of size $N = N_1 N_2$, where N_l refers to the number of voxels along the dimension l using an array containing C coils from $M = ESC$ samples of a discretized $\boldsymbol{k}$–space grid of size $K = K_1 K_2$. E refers to the number of sampled points per shot and S is the number of shots. Extension to 3D is straightforward for this model. The terms in Eq.5.23 are represented by the following matrices:

- $\vec{\mathbf{y}}$ is a vector of size $M \times 1$ containing measured multi-shot $\boldsymbol{k}$–space data.
- $\mathbf{A}$ is a sampling matrix of size $M \times KSC$ given by

$$\mathbf{A} = \begin{pmatrix} \mathbf{A}_{11} & \cdots & \mathbf{0} & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots & \vdots & \vdots \\ \mathbf{0} & \cdots & \mathbf{A}_{1C} & \cdots & \mathbf{0} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \cdots & \mathbf{0} & \cdots & \mathbf{A}_{SC} \end{pmatrix}, \tag{5.24}$$

  where A_{sc} is a matrix of size $E \times K$ whose entries are equal to 1 if the sample e of the shot s matches the $\boldsymbol{k}$–space location indexed by k and 0 otherwise.

- $\boldsymbol{\mathcal{F}}$ is a matrix of size $KSC \times NSC$ that performs the Discrete Fourier Transform (DFT) given by the diagonal block structure

$$\boldsymbol{\mathcal{F}} = \begin{pmatrix} \mathbf{F}_{11} & \cdots & \mathbf{0} & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots & \vdots & \vdots \\ \mathbf{0} & \cdots & \mathbf{F}_{1C} & \cdots & \mathbf{0} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \cdots & \mathbf{0} & \cdots & \mathbf{F}_{SC} \end{pmatrix}, \tag{5.25}$$

  where $\mathbf{F}_{sc}$ is a $K \times N$ matrix performing the 2D DFT.

- $\mathbf{S}$ is a matrix of size $NSC \times NS$ containing properly arranged coil sensitivity maps:

$$\mathbf{S} = \begin{pmatrix} \mathbf{S}_{11} & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{S}_{1C} & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \cdots & \mathbf{S}_{SC} \end{pmatrix}, \tag{5.26}$$

  where $\mathbf{S}_{sc}$ is a diagonal matrix of size $N \times N$ containing the sensitivity profile of coil c.

- $\boldsymbol{\mathcal{P}}(\vec{\boldsymbol{\theta}})$ is a matrix of size $NS \times N$ containing the phase maps given by

$$\boldsymbol{\mathcal{P}}(\vec{\boldsymbol{\theta}}) = \begin{pmatrix} \mathbf{P}(\vec{\boldsymbol{\theta}}_{\mathbf{1}}) \\ \vdots \\ \mathbf{P}(\vec{\boldsymbol{\theta}}_{\boldsymbol{S}}) \end{pmatrix}, \tag{5.27}$$

  where $\mathbf{P}(\vec{\boldsymbol{\theta}}_{\boldsymbol{s}})$ is a diagonal matrix of size $N \times N$ containing the linear phase map for shot s characterized by three parameters $\vec{\boldsymbol{\theta}}_{\boldsymbol{s}} = [\theta_{s0}, \theta_{s1}, \theta_{s2}]$. Its diagonal is given by the vector $\vec{\mathbf{p}}_s(\vec{\boldsymbol{\theta}}_{\boldsymbol{s}})$:

$$\vec{\mathbf{p}}_s(\vec{\boldsymbol{\theta}}_s) = e^{j(\theta_{s0} + \theta_{s1} \cdot r_1 + \theta_{s2} \cdot r_2)}, \tag{5.28}$$

  where r_l denotes the spatial coordinate vector along dimension l.

- $\vec{\mathbf{x}}$ is a vector of size $N \times 1$ containing the reconstructed image.

This is graphically depicted in Fig.5.3, where we can observe the forward process that transforms the original image into the multi–shot and multi–channel sampled data.

In order to solve the problem posed in Eq.5.23, it is important to notice that our minimization problem is a non-linear least squares optimization problem, which does not have a closed-form solution in general. Directly solving the joint optimization problem, however, is practically intractable. For this reason, we proceed in a greedy iterative fashion, alternating between solving the following two subproblems (Ying and Sheng, 2007):

- Image reconstruction problem

$$\vec{\mathbf{x}}^* = \underset{\vec{\mathbf{x}}}{\operatorname{argmin}} \left\| \mathbf{A}\boldsymbol{\mathcal{F}}\mathbf{S}\boldsymbol{\mathcal{P}}(\vec{\boldsymbol{\theta}}^*)\vec{\mathbf{x}} - \vec{\mathbf{y}} \right\|_2^2. \tag{5.29}$$

  This subproblem can be solved by means of the conjugate gradient method described in (Pruessmann et al., 2001) applied to the system:

$$\boldsymbol{\mathcal{P}}^H(\vec{\boldsymbol{\theta}}^*)\mathbf{S}^H\boldsymbol{\mathcal{F}}^H\mathbf{A}^H\mathbf{A}\boldsymbol{\mathcal{F}}\mathbf{S}\boldsymbol{\mathcal{P}}(\vec{\boldsymbol{\theta}}^*)\vec{\mathbf{x}} = \boldsymbol{\mathcal{P}}^H(\boldsymbol{\theta}^*)\mathbf{S}^H\boldsymbol{\mathcal{F}}^H\mathbf{A}^H\vec{\mathbf{y}} \tag{5.30}$$

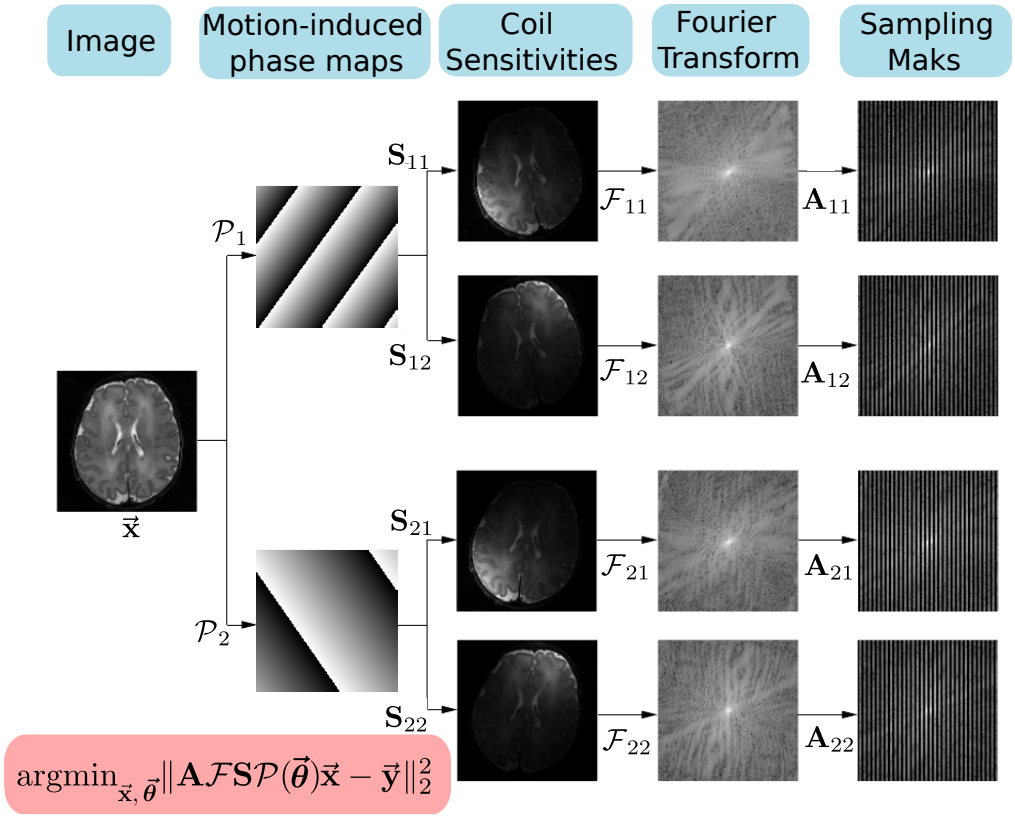

Figure 5.3: Forward model in 2D of the acquisition process for multi–shot data in the presence of rigid motion. The original image to be reconstructed $\boldsymbol{x}$ might be corrupted by different phase maps $\boldsymbol{P}(\boldsymbol{\theta_s})$ derived from different states of rigid motion during the application of the diffusion sensitized gradients. Here $s \in [1, \dots, S]$ denotes the shot ($S = 2$ in this example). The availability of multiples receive channels (2 in this example) allows to encode the measured information by means of their sensitivity profiles S_{sc}, where $c \in [1, \dots, C]$ denotes the channel ($C = 2$ in this example). The data is measured in the $\boldsymbol{k}$–space, which involves the FT given by F_{sc}. Finally, each shot subsamples a different trajectory of the $\boldsymbol{k}$–space, encoded in the matrix A_{sc}.

- Phase maps estimation

$$\vec{\boldsymbol{\theta}}^* = \underset{\vec{\theta}}{\operatorname{argmin}} \left\| \mathbf{A}\boldsymbol{\mathcal{F}}\mathbf{S}\boldsymbol{\mathcal{P}}(\vec{\boldsymbol{\theta}})\vec{\mathbf{x}}^* - \vec{\mathbf{y}} \right\|_2^2. \tag{5.31}$$

We look for the solution that nulls the gradient of the objective function:

$$\nabla_{\vec{\boldsymbol{\theta}}} \left(\left\| \mathbf{A}\boldsymbol{\mathcal{F}}\mathbf{S}\boldsymbol{\mathcal{P}}(\vec{\boldsymbol{\theta}})\vec{\mathbf{x}}^* - \vec{\mathbf{y}} \right\|_2^2 \right) = 0, \tag{5.32}$$

and for that purpose we have chosen to use the Newton's algorithm, which uses information of the Hessian. Details are provided in the Appendix.C.

Since the objective function is nonconvex, it will in general contain multiple local minima. For this reason, a proper initialization becomes vital for our algorithm to find the global minimum. In section 5.3 we will explain the different strategies we have considered for this initialization and we will carry out an experiment to determine our choice for the final algorithm.

Non–rigid motion

Unluckily, there might be a motion component that is non–rigid as a consequence of more complex patterns mainly related to tissue deformation arising from the blood and cerebrospinal fluid pulsation (Miller and Pauly, 2003). In the case of the brain, for example, this non–rigid component is stronger as we approach the cervicals area and the distortion it introduces becomes more problematic as we increase the b–value (Greitz et al., 1992). Importantly, the presence of non–rigid motion results in non–linear phase corruption (Atkinson et al., 2006). Hereby we will assume that these non–linear phase maps present a smooth spatial variation, which is a common assumption in ms–EPI (Uecker et al., 2009; Chen et al., 2013; Mani et al., 2017; Lobos et al., 2018).

In order to model these non linear phase maps we will make use of B–splines. We define them over a set of knots, whose coordinates are given by $\vec{\mathbf{p}}^{\boldsymbol{u}} = (p_1^{\boldsymbol{u}}, \ldots, p_L^{\boldsymbol{u}})$. The superscript $\boldsymbol{u}$ refers to the vector $\vec{\mathbf{u}} = (u_1, \ldots, u_L)$, which contains the indices of the knot, taking an integer value over the range $-C_l \leq u_l \leq D_l$, where L refers to the number of dimensions ($L = 2$ in our case). We can now define the B–spline to approximate the phase maps $\varphi_s^{nr}(\boldsymbol{r})$ resulting from non–rigid motion for each shot s, characterized by the parameters $\vec{\boldsymbol{\xi}}_s = \{\xi_s^{\boldsymbol{u}}\}$:

$$\varphi_s^{nr}(\vec{\mathbf{r}}, \vec{\boldsymbol{\xi}}_s) = \sum_{u_1=-C_1}^{D_1} \cdots \sum_{u_L=-C_L}^{D_L} \left(\prod_{l=1}^{L} B_E \left(\frac{r_l - p_l^{\boldsymbol{u}}}{\Delta_l} \right) \right) \xi_s^{\boldsymbol{u}}, \tag{5.33}$$

where Δ_1 contains the distance between knots along dimension l and $B_E(\boldsymbol{r})$ represents the uniform B–spline of degree E, which can be constructed by means of the Cox-de Boor recursion

formula (Boor, 2003). Further details are provided in Appendix.D. As an example, the B–spline of order 2 is given by:

$$B_2(r) = \begin{cases} \frac{4r^2+12r+9}{8} & \text{if } -1.5 \leq r \leq -0.5 \\ \frac{-4r^2+3}{4} & \text{if } -0.5 \leq r \leq 0.5 \\ \frac{4r^2-12r+9}{8} & \text{if } 0.5 \leq r \leq 1.5 \\ 0 & \text{otherwise} \end{cases} \tag{5.34}$$

Given this phase map, we can add an extra operator $\boldsymbol{\mathcal{Q}}(\vec{\boldsymbol{\xi}})$ in the forward transform, which is a matrix of size $NS \times N$ containing the phase maps given by

$$\boldsymbol{\mathcal{Q}}(\vec{\boldsymbol{\xi}}) = \begin{pmatrix} \boldsymbol{\mathcal{Q}}(\vec{\boldsymbol{\xi}}_1) \\ \vdots \\ \boldsymbol{\mathcal{Q}}(\vec{\boldsymbol{\xi}}_S) \end{pmatrix}, \tag{5.35}$$

where $\boldsymbol{\mathcal{Q}}(\vec{\boldsymbol{\xi}}_s)$ is a diagonal matrix of size $N \times N$ containing the non–linear phase map for shot s characterized by the parameters $\vec{\boldsymbol{\xi}}_s$ and whose diagonal is given by the vector $\vec{\mathbf{q}}_s(\vec{\boldsymbol{\xi}}_s)$:

$$\vec{\mathbf{q}}_s(\vec{\boldsymbol{\xi}}_s) = e^{j\varphi_s^{nr}(\boldsymbol{r},\vec{\boldsymbol{\xi}}_s)}. \tag{5.36}$$

Therefore, our updated Maximum Likelihood formulation looks like:

$$(\vec{\mathbf{x}}^*, \vec{\boldsymbol{\theta}}^*, \vec{\boldsymbol{\xi}}^*) = \underset{\vec{\mathbf{x}},\vec{\boldsymbol{\theta}},\vec{\boldsymbol{\xi}}}{\operatorname{argmin}} \left\| \mathbf{A}\boldsymbol{\mathcal{F}}\mathbf{S}\boldsymbol{\mathcal{Q}}(\vec{\boldsymbol{\xi}})\boldsymbol{\mathcal{P}}(\vec{\boldsymbol{\theta}})\vec{\mathbf{x}} - \vec{\mathbf{y}} \right\|_2^2. \tag{5.37}$$

The pseudocode for the optimization algorithm is detailed in Alg.1.

Details for solving Eq.5.40 are provided in the Appendix.E. As for the initialization indicated in the first step of Alg.1, $\vec{\mathbf{x}}^{(0)}$ and $\vec{\boldsymbol{\xi}}^{(0)}$ were initialized to 0, whereas $\vec{\boldsymbol{\theta}}^{(0)}$ was initialized following a three-step procedure as depicted in Fig.5.4:

1. SENSE Reconstruction: we reconstruct shot–by–shot using SENSE with Tikhonov regularization, which can be done using the classical formulation as in (Pruessmann et al., 1999):
$$\vec{\mathbf{x}}_s^* = \underset{\vec{\mathbf{x}}_s}{\operatorname{argmin}} \|\mathbf{A}\boldsymbol{\mathcal{F}}\mathbf{S}\vec{\mathbf{x}}_s - \vec{\mathbf{x}}_s\|_2^2 + \lambda\|\vec{\mathbf{x}}_s\|_2^2. \tag{5.41}$$
2. Peak detection: according to FT properties, constant offsets are preserved in $\boldsymbol{k}$–space and linear phases introduce $\boldsymbol{k}$–space shifts. Under no phase corruption and assuming ideal spin–echo acquisitions, our image would be real and therefore the peak of its FT would be placed in the centre of the spectrum. This allows us to obtain an initial guess of the phase parameters $\boldsymbol{\theta_s} = [\theta_{s0}, \theta_{s1}, \theta_{s2}]$ by detecting the position and phase of that peak.

Algorithm 1 Optimization algorithm

1) Initialize $\vec{\mathbf{x}}^{(0)} = 0, \vec{\boldsymbol{\theta}}^{(0)} = \vec{\boldsymbol{\theta}}_{\text{init}}, \vec{\boldsymbol{\xi}}^{(0)} = 0$

2) for $i = 1, \ldots,$ until convergence

a) Use Conjugate–Gradient to solve

$$\vec{\mathbf{x}}^{(i)} = \underset{\vec{\mathbf{x}}}{\text{argmin}} \left\| \mathbf{A}\boldsymbol{\mathcal{F}}\mathbf{S}\boldsymbol{\mathcal{Q}}(\vec{\boldsymbol{\xi}}^{(i-1)})\boldsymbol{\mathcal{P}}(\vec{\boldsymbol{\theta}}^{(i-1)})\vec{\mathbf{x}}^{(i-1)} - \vec{\mathbf{y}} \right\|_2^2. \tag{5.38}$$

b) Use Newton's method to solve

$$\vec{\boldsymbol{\theta}}^{(i)} = \underset{\vec{\boldsymbol{\theta}}}{\text{argmin}} \left\| \mathbf{A}\boldsymbol{\mathcal{F}}\mathbf{S}\boldsymbol{\mathcal{Q}}(\vec{\boldsymbol{\xi}}^{(i-1)})\boldsymbol{\mathcal{P}}(\vec{\boldsymbol{\theta}}^{(i-1)})\vec{\mathbf{x}}^{(i)} - \vec{\mathbf{y}} \right\|_2^2. \tag{5.39}$$

c) Use Levenberg–Marquard's method to solve

$$\vec{\boldsymbol{\xi}}^{(i)} = \underset{\vec{\boldsymbol{\xi}}}{\text{argmin}} \left\| \mathbf{A}\boldsymbol{\mathcal{F}}\mathbf{S}\boldsymbol{\mathcal{Q}}(\vec{\boldsymbol{\xi}}^{(i-1)})\boldsymbol{\mathcal{P}}(\vec{\boldsymbol{\theta}}^{(i)})\vec{\mathbf{x}}^{(i)} - \vec{\mathbf{y}} \right\|_2^2. \tag{5.40}$$

3) Obtain the final reconstructed image by solving again Eq.5.38

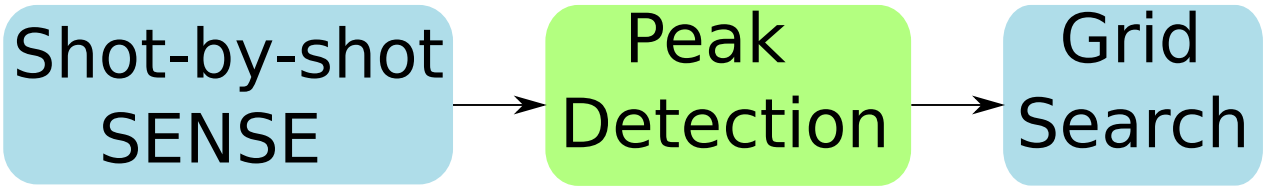

Figure 5.4: Three–step procedure for the initialization of linear phase parameters. First each shot was reconstructed using a regularized version of SENSE. Second, an equalization filter was applied in order to take into account noise–amplification due to g–factor effects. Third, the peak of the spectrum was detected from which the linear phase parameters can be obtained. Fourth, a grid search was carried out in order to scape from wrong peak detection for a particular shot due to noise amplification.

3. Grid search: peak detection may be subject to noise amplification and it can be particularly problematic in the case where the peak was originally shifted to the intermediate location in between two acquired lines. In order to be able to escape from this situation where the peak for a certain shot might not be well–identified, we add gradient-free search on a discretized grid in the space of phases. In order to alleviate the computational burden, it operates in an Iterative Conditional Modes (ICM) fashion. Therefore, while fixing the phases of the rest of the shots, we search the three–dimensional space of phases for a certain shot trying to minimize Eq.5.39.

The updated forward model for the case of non–rigid motion is graphically depicted in Fig.5.5.

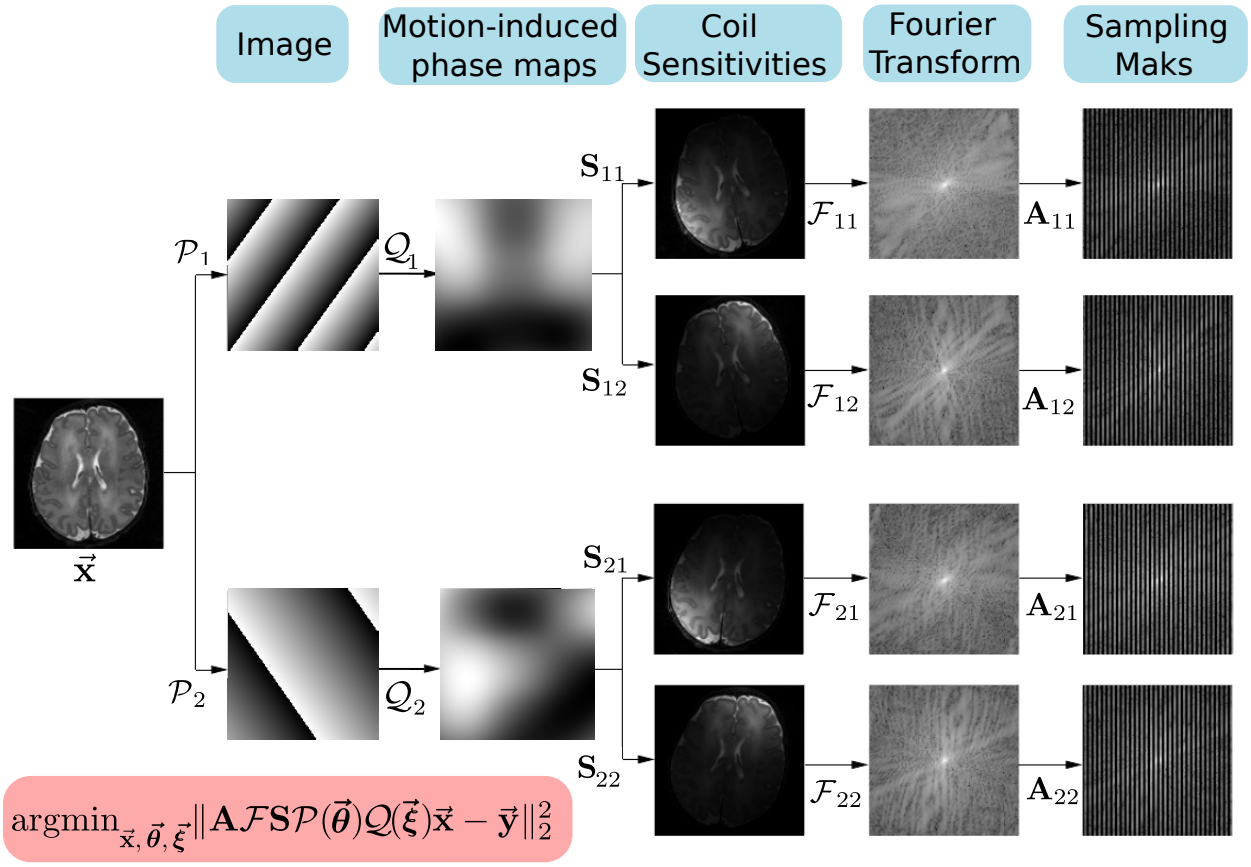

Figure 5.5: Forward model in 2D of the acquisition process for multi–shot data in the presence of non–rigid motion. The original image to be reconstructed $\boldsymbol{x}$ might be corrupted by different linear phase maps $\boldsymbol{P}(\boldsymbol{\theta_s})$ derived from different states of rigid motion during the application of the diffusion sensitized gradients, as well as by different non–linear phase maps $\boldsymbol{Q}(\boldsymbol{\xi_s})$ resulting from different states of non–rigid motion. Here $s \in [1, \ldots, S]$ denotes the shot ($S = 2$ in this example). The availability of multiples receive channels (2 in this example) allows to encode the measured information by means of their sensitivity profiles S_{sc}, where $c \in [1, \ldots, C]$ denotes the channel ($C = 2$ in this example). The data is measured in the $\boldsymbol{k}$–space, which involves the FT given by F_{sc}. Finally, each shot subsamples a different trajectory of the $\boldsymbol{k}$–space, encoded in the matrix A_{sc}.

5.3 Methods

5.3.1 Algorithms

We have studied both our rigid–motion method (termed ItLin), which solves Eq.5.23 and the non–rigid motion one (ItNL), which solves Eq5.37. We have compared them with three differet alternatives:

1. MUSE: following (Chen et al., 2013), we reconstruct each shot using SENSE and from it we get an estimate of its phase map after applying a Total-Variation (TV) filter. We then reconstruct the image solving Eq.5.29.

2. Linearly parameterized MUSE (LinMUSE): since the original MUSE implementation was developed for a more general scenario where the only prior knowledge about the phase is that it is smooth, we implemented a MUSE version that incorporates the prior knowledge of linear phase corruption maps. The difference with the previous algorithm is that, instead of taking the phase of each shot-reconstructed image after the TV-filter, we estimate the three parameters caracterising the linear phase map from the peak and position of the $\boldsymbol{k}$–space peak.

3. LinMUSE with $L2-$Tikhonov regularization (LinMUSE+Tik): under noise corruption, peak detection in $\boldsymbol{k}$-space can be problematic, especially when the peak is shifted midway between two acquired lines. For this reason, a certain degree of regularization in the SENSE reconstruction, as in Eq.5.41, can help avoid misidentifying the peak as amplified noise. It is important to notice that regularization may induce errors as well due to incorrect removal of ghosting artifacts.

5.3.2 Datasets

We carried out different experiments using the following datasets:

- Simulated multi–shot brain dataset: we have used a T_2 neonatal brain axial image acquired with a fast spin echo sequence on a $3T$ PHILIPS ACHIEVA TX with a head coil array of 32 channels and using the following parameters: resolution= $0.8 \times 0.8\text{mm}^2$, slice thickness= 1.6mm, TE=145ms, TR=12s, flip angle $\alpha = 90^\circ$. Coil sensitivity profiles were estimated from a separate pre-scan using (Allison et al., 2013). The image was reconstructed without zero filling in order to preserve the resolution and then cropped to a 128×128 matrix so the brain adjustes to the FOV. To generate our synthetic ms–EPI data, we applied the forward model described in Eq.5.23 and Eq.5.37, depending on the experiment.

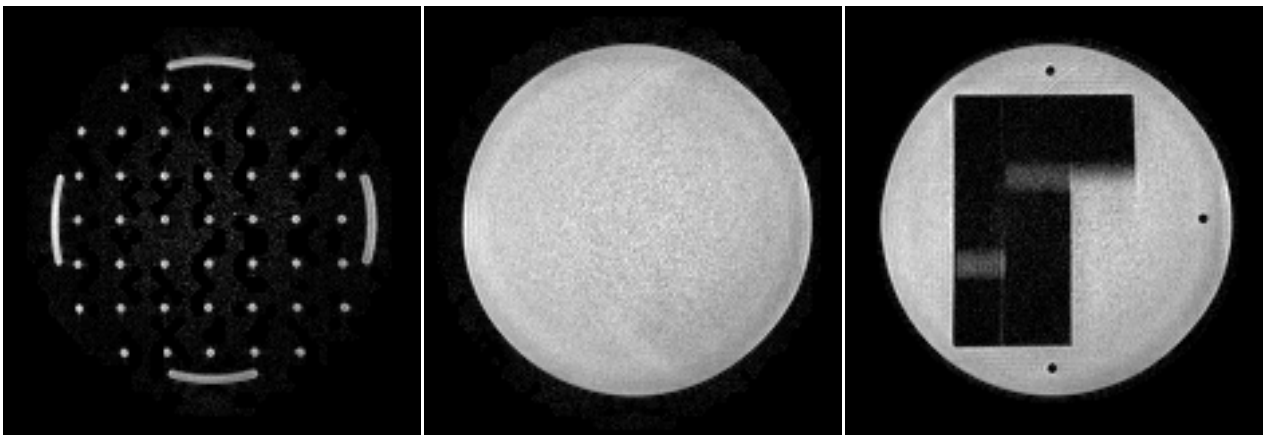

Figure 5.6: Water phantom acquired with ss–EPI and SENSE acceleration factor 2 and 8 coils, considered our silver ground-truth. We have acquired three different slices of the phantom, shown from left to right.

- Water phantom acquisition: we have scanned a water phantom acquired with a diffusion turbo spin echo sequence on a $3T$ PHILIPS ACHIEVA TX with a head coil array of 8 channels. Coil sensitivity profiles were estimated from a separate pre-scan using (Allison et al., 2013) and the following parameters were fixed for all acquisitions: FOV= $24 \times 24 \times 4.8$ cm^3, slice thickness= 3mm, flip angle $\alpha = 90°$, b–value= 1000 s/mm^2. We have acquired 3 different slices varying the number of shots, which resulted in the following acquisition parameters:

Table 5.1: Acquisition parameters

# shots	Acc	TE (ms)	TR (s)	Resolution (mm^2)	Matrix size
1	2	103	0.885	1.5×1.51	160×160
2	1	132	1.12	1.5×1.52	160×158
4	1	102	0.857	1.5×1.54	160×156
8	1	92	0.752	1.5×1.58	160×152
10	1	89	0.725	1.5×1.6	160×150

  We have considered the 1–shot acquisition with acceleration factor 2 (denoted Acc in Table.5.1) as the reference, which can be seen in Fig.5.6.

- Synthetic ms-EPI acquisition corrupted with realistic phase–maps from *in–vivo* ss-EPI brain data: We have used a dataset of diffusion images consisting of three different slices acquired at different b-values: $\{0, 100, 400, 900, 1600, 2500, 3600, 4900, 6400, 8100, 10000\}$ s/mm^2. They were collected using a diffusion turbo spin echo sequence on a $3T$ PHILIPS ACHIEVA TX with a head coil array of 32 channels. Once again, coil sensitivity maps were estimated from a separate pre-scan using (Allison et al., 2013). The acquisition parameters we used were: resolution= 2.5×2.5mm^2, FOV= $24 \times 24 \times 5$ cm^3, slice

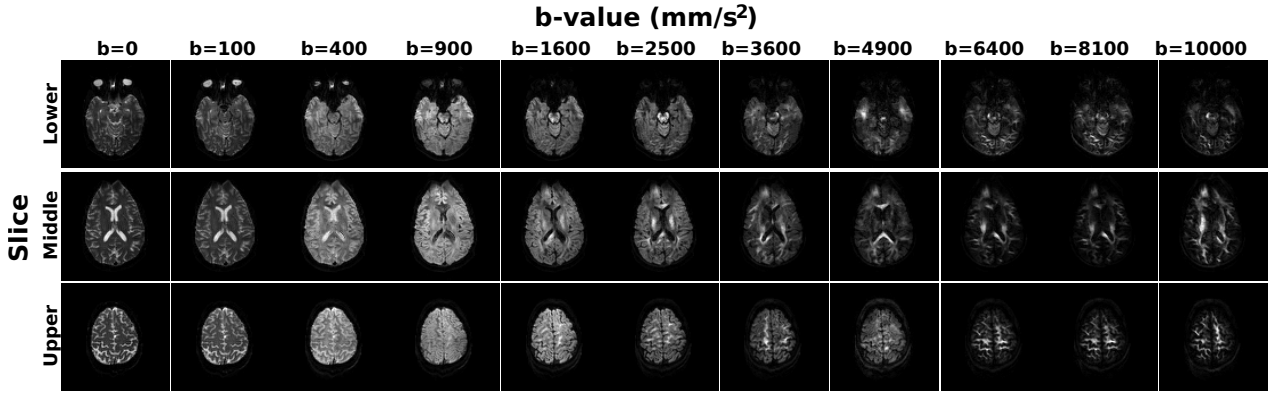

Figure 5.7: In–vivo ss–EPI data used to synthetize real phases in order to simulate a multi–shot acquisition. Each column corresponds to a different b–value from $\{0, 100, 400, 900, 1600, 2500, 3600, 4900, 6400, 8100, 10000\}$ s/mm^2 and each row corresponds to a different slice in the brain.

thickness= 2.5mm, flip angle $\alpha = 90°$, TE=99ms, TR=6.9s, matrix size=96×96. These images were reconstructed in the complex domain after being filtered by means of the Marchenko–Pastur denoising technique proposed in (Veraart et al., 2016). The magnitude images are shown in Fig.5.7.

5.3.3 Experiments

First, we have carried out two experiments on our data using the synthetic phantom previously described:

- In order to visually assess the ability of the aforementioned methods to reconstruct the image, we compared the reconstructions at a low SNR scenario (SNR=1) for 4, 6, 8, 10 and 12 shots. The phase offset was randomly generated between $[-\pi, \pi]$, whereas the linear phase ramp in Eq.5.28 was generated to shift the peak in the region covering the central S lines, with S denoting the number of shots. We focused on this region of linear shifts since the worst scenario arises as previously stated, when the peak is shifted midway between two acquired lines. With respect to those acquired lines themselves, their absolute position in $\boldsymbol{k}$–space does not seem to play a role as for peak determination.

- We carried out 100 reconstructions using $S = 8$ shots for varying levels of SNR, where the SNR was computed as the mean absolute signal divided by the standard deviation of the synthetic uniformly gaussian added noise. The phase parameters were generated as in the previous experiment. The experiment compares the absolute estimation error for each phase parameter in Eq.5.28 for our method and for LinMUSE.

With the aim of studying the ability of the B–splines method (termed ItNL) to reconstruct the images, we carried out 100 reconstruction using $S = 8$ shots for varying levels of SNR, where the SNR was computed as previously explained. The linear phase parameters were generated as in the previous experiment, whereas the non–linear phase maps were synthetically generated by means of a B–spline using a grid of 5×5 knots, whose corresponding control points were randomly generated in the range $[-\pi, \pi]$. The experiment compares the mean squared error for the reconstructed images using our linear method (ItLin), our non–linear method (ItNL) and standard MUSE with respect to the ground–truth.

Second, we have reconstructed the phantom data acquired using different number of shots. Given the expected linear phase corruption, we have only compared our linear algorithm ItLin against LinMuse with Tikhonov regularization and standard Muse. Due to the different resolution for the different acquisition, we have assessed our reconstructions by visually comparing them to the images obtained from the alternative methods.

Third, we have simulated an 8–shot acquisition from the ss–EPI *in–vivo* brain acquisition. For each of the sampled b–values, multiple diffusion directions were available, which has allowed us to synthetize realistic phase corruption maps. To do so, we randomly chose 8 of the diffusion images for each b-value and used their phase maps to synthetically generate our $\boldsymbol{k}$–space, therefore obtaining a realistic model-free phase corruption acquisition. In particular, we have reconstructed the data using both our linear (LinIt) and non–linear (ItNL) algorithms, as well as the standard version of MUSE. In order to assess the quality of the reconstructions, we have computed the relative MSE against the ground-truth.

5.4 Results

In Fig.5.8 we show the reconstructed images for a low SNR (SNR=1) scenario varying the number of shots. The first column shows the ground-truth, and the next columns show the reconstruction for the described method, LinMUSE and standard MUSE.

We observe that for $S = 4$, all reconstructions seem to perform similarly, which is consistent with the literature (Chen et al., 2013). However, when we increase the number of shots up to 8, we observe that while our method is able to reconstruct properly the image, MUSE-type reconstructions introduce ghosting artifacts. Increasing the number of shots augments the g-factor noise amplification associated to SENSE, which in turn increments the uncertainity in the detection of the peak of the spectrum. Specifically, when the phase corruption shifts the peak to the midway position between two sampled lines for a particular shot, it becomes more likely to identify it at a noisy location, and this can result in a high-frequency ghosting artifact.

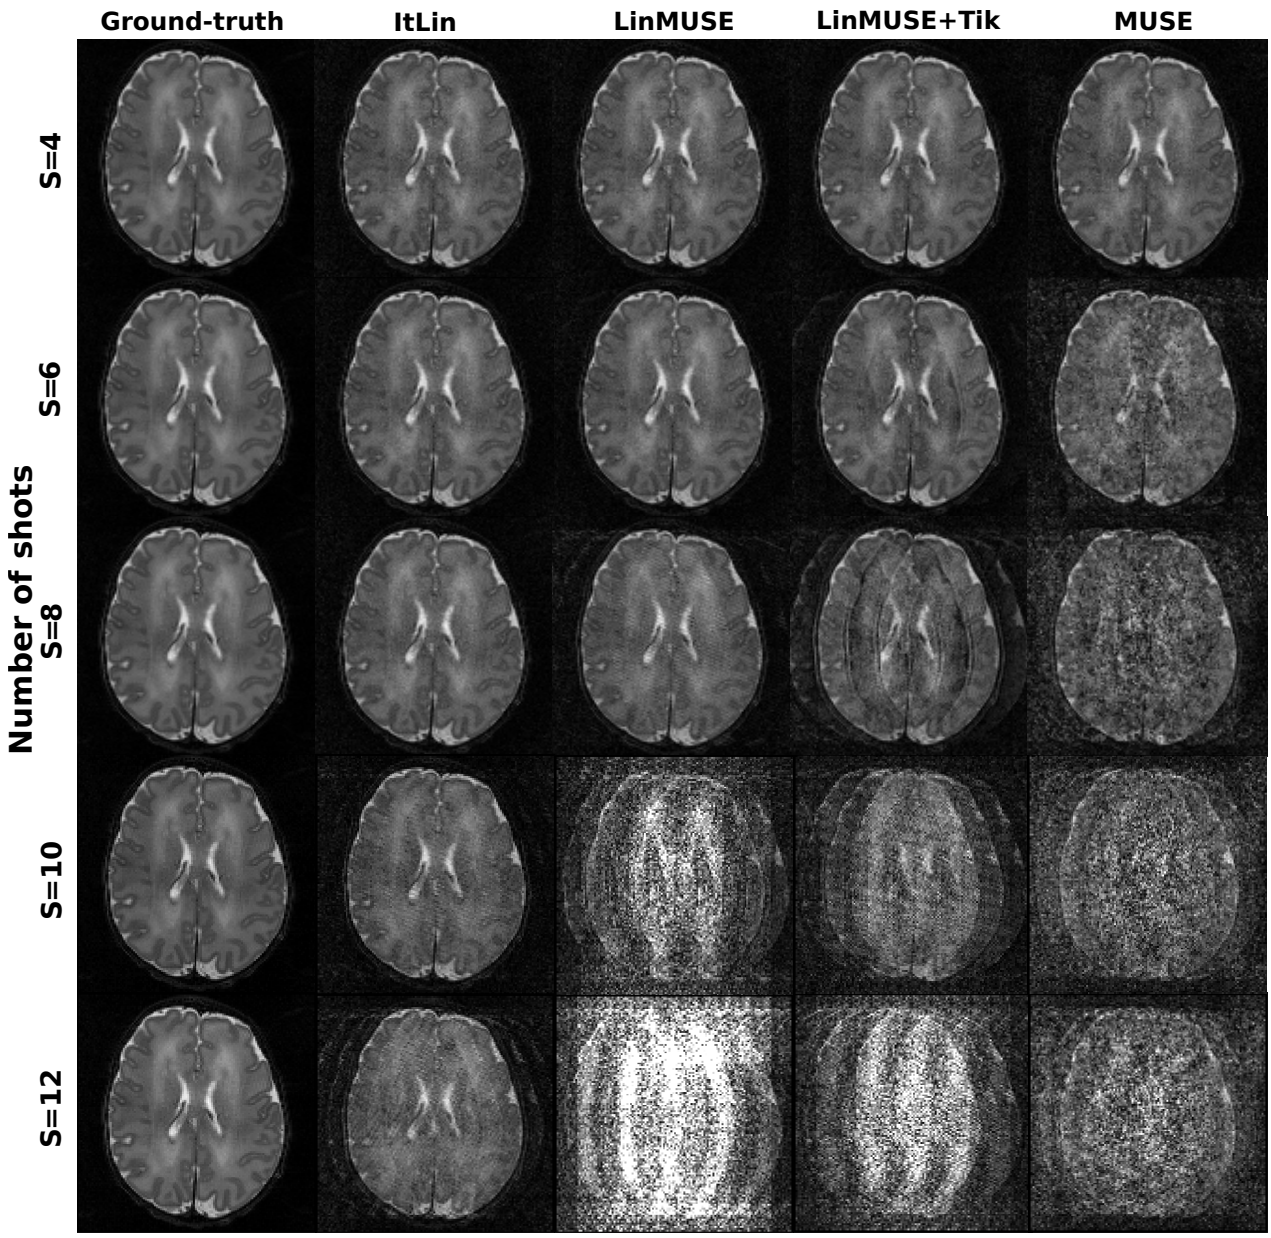

Figure 5.8: Reconstructed images in a low SNR scenario for different number of shots: 4 (first row), 6 (second row), 8 (third row), 10 (fourth row) and 12 (fifth row). The first column shows the ground-truth and the following ones the reconstructed images with ItLin (second column), LinMUSE (third column), LinMUSE with Tikhonov regularization (fourth column) and standard MUSE (fifth column).

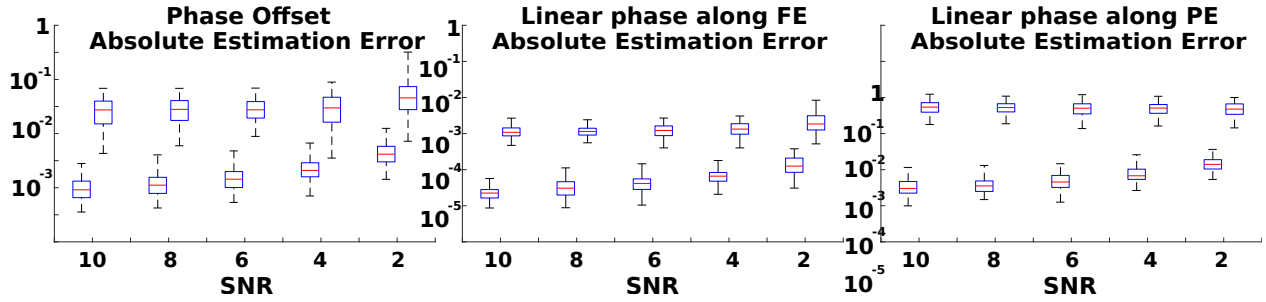

Figure 5.9: Boxplots of the absolute errors for the phase parameters estimation: offset (first column), FE slope (second column) and PE slope (third column). In each figure and for each value of the SNR, the leftmost boxplot shows the error associated to ItLin, whereas the rightmost boxplot shows the error of LinMUSE. We must point out that multiple outliers were present for the lowest SNR scenario in all of the studied reconstructions, although we removed them from the figure for a more clear visualization.

Tikhonov regularization can increase the accuracy in the peak detection at the cost of reducing its precision, resulting in low-frequency ghosting artifacts. This increased accuracy increases the likelihood of escaping local minima by initializing our algorithm closer to the global minimum. On the other hand, by the extra steps included in our method, i.e. search on a fine grid around the initialization and Newton's iterative optimization, we seem to be able to obtain a ghost-free image even for the case of 8 shots. However, for 10 shots, we observe that our method starts to introduce ghosting artifacts, related to the inability to escape local minima due to the increased uncertainity in the peak detection from which we obtain initialize our algorithm. Further increasing the number of shots to 12 results in a very distorted image.

In Fig.5.9 we can see that the absolute errors for the estimation of the three different parameters: the constant offset, the slope along the FE direction and the slope along the phase–encoding (PE) direction. We compared the errors between our method and LinMUSE for varying levels of SNR. We observe that our method is able to get a closer estimate to the actual phase parameters. This is consistent with the fact that our method is initialized with the phase estimates provided by LinMUSE and from there a finer estimation is done based on both the initial grid search and the subsequent Newton's method based descent. It is also remarkable that the estimation error is higher along the PE direction compared to the FE direction. This is related to the fact that the undersampling of each shot occurs along the PE direction, increasing the difficulty to detect the peak along this direction.

Regarding the non–rigid motion scenario, we include both Fig.5.10 showing the numerical MSE for different reconstruction algorithms and Fig.5.11 depicting the final images for a particular experiment at both high and low SNR. We can observe that for high SNR scenarios, the non–linear algorithm is able to retrieve both the linear phase parameters as well as the non–linear control point values characterizing the B–splines for the different shots, allowing to reconstruct a ghosting–free image. When the SNR is reduced, in most of the cases it is able

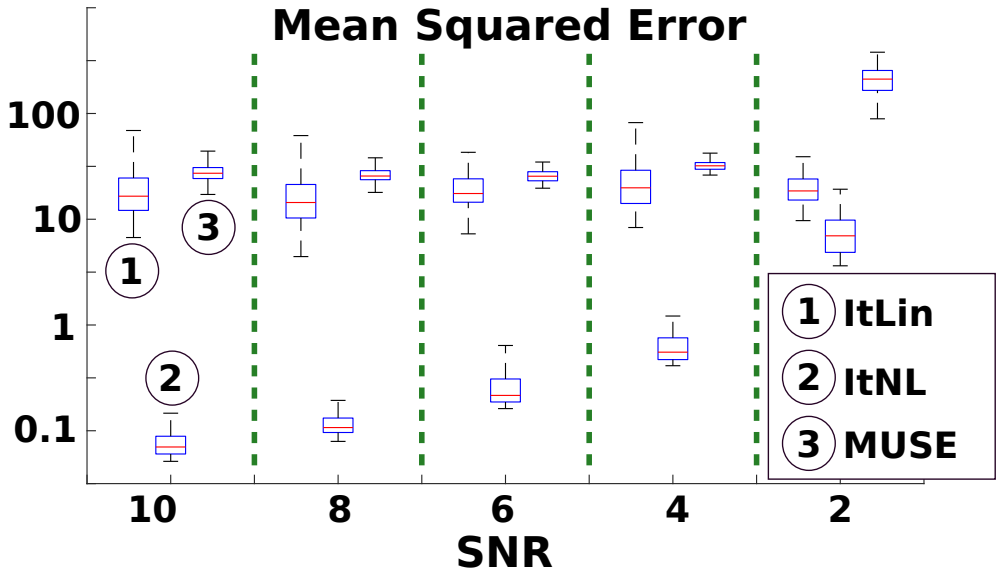

Figure 5.10: Boxplots of the mean squared errors for the reconstructed images at different levels of SNR for our linear iterative algorithm (ItLin), our non–linear iterative algorithm (ItNL) and standard MUSE.

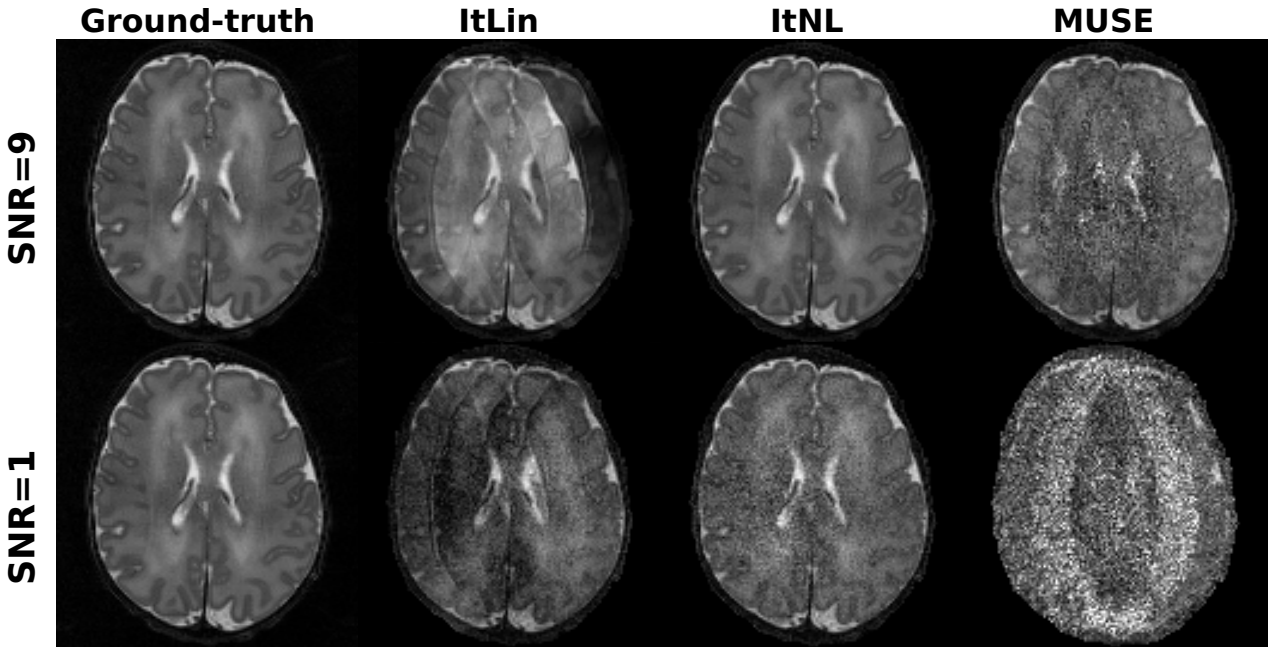

Figure 5.11: Reconstructed images in both a high (first row) and a low SNR scenario (second row) for a randomly picked experiment comparing the different methods. The first column shows the ground-truth and the following ones the reconstructed images with ItLin (second column), ItNL (third column) and standard MUSE (fourth column). We must point out that we show a case where our method succeeds at recovering the image although, for the lowest SNR scenario, it is not completely robust since some of the experiments showed some residual ghosting.

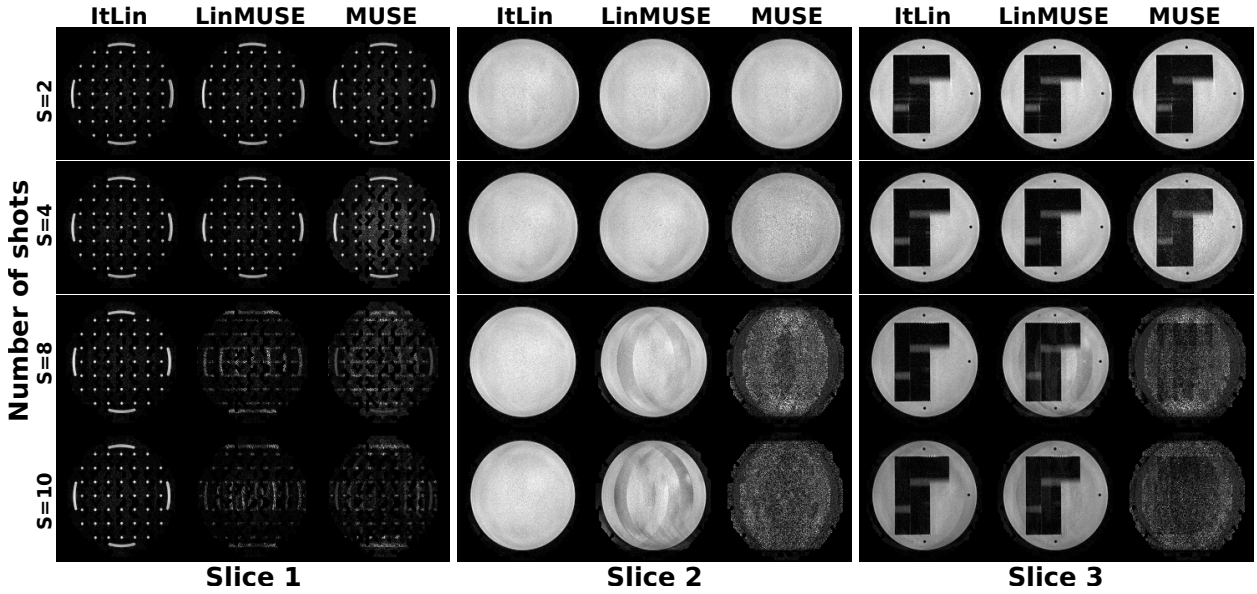

Figure 5.12: Reconstructed images for the three different slices comparing the different methods. The first column shows our reconstruction using the linear algorithm (ItLin), the second column show the results for LinMUSE with Tikhonov regularization and the last column shows the results for standard MUSE. Different rows refer to different number of shots.

to recover the image, but it is not completely robust since some of the images still showed residual ghosting artifacts. The non–linear phase maps have the effect of a convolution in $\boldsymbol{k}$–space, which smooths the peak of the spectrum and thus reduces its level with respect to the noise floor, rendering it harder to detect. On the other hand, when ItLin is applied, the final image is notably distorted by ghosting artifacts, which can be explained because of the discrepancy between the model and reality. Finally, as expected from the previous experiments, standard MUSE is noticeably affected by g–factor noise amplificaMean Squared Errortion associated to the high undersampling factor for each shot, which notably hinders the estimation of the shot phase–maps after SENSE reconstruction.

With respect to the real phantom experiments, we include Fig.5.12 depicting the reconstructed images for the varying number of shots for the three different slices. We can observe that our linear algorithm is able to reconstruct the images up to 10 shots, although a certain degree of distortion can be observed for the highest number of shots in the last slice. On the other hand, LinMuse and standard MUSE are able to reconstruct high–quality images up to 4 shots, but introduce different levels of distortions when we move to 8 or 10 shots.

Finally, Fig.5.13 shows the relative MSE in dB computed as the ratio between the mean squared intensity of the ground–truth and the mean squared error. It is shown for the three considered slices as a function of the b–value. We can observe that the error increases as the b–value does, which is consistent with the SNR decay in the acquired data, as well as with the increased sensitivity to non–rigid motion at high b–values that hinders the image

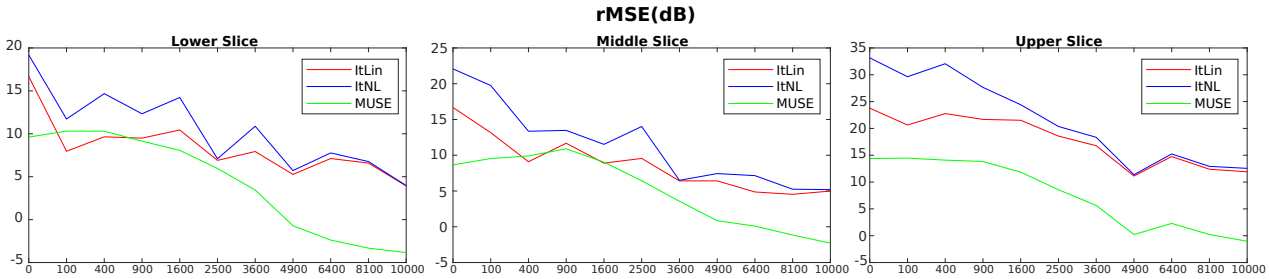

Figure 5.13: Reconstructed images for the three different slices comparing the different methods at b=2500 s/mm^2. The first column shows the reconstruction and the followinc columns compare our linear algorithm (ItLin), our non–linear algorithm (ItNL) and standard MUSE. Different rows correspond to different slices.

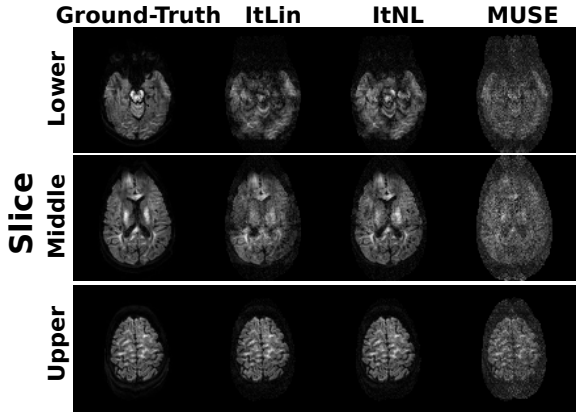

Figure 5.14: Relative MSE for the three different slices at different b-values comparing the reconstructions using ItLin, ItNL and standard MUSE.

reconstruction. Likewise, it can be seen that the error grows as we descend along the z–axis. This is consistent with the observation that non–rigid pulsatile motion is stronger in the brain stem and within the basal ganglia (Greitz et al., 1992). We include an example of the reconstructed images for b=2500 s/mm^2 in Fig.5.14. We can observe that for the upper slice, where the motion is mostly rigid, both ItLin and ItLN perform similarly. When we descend to the intermediate slice, we observe that ItLin introduces ghosting artifacts, which are partially removed by ItNL. Last, at the lower slice, both of them result in a distorted reconstruction.

5.5 Discussion

In this work we have proposed a joint procedure to both reconstruct the images and to estimate the phase error maps for ms–EPI suffering from phase discrepancies between shots as a consequence of patient motion during the application of the diffusion-sensitized gradients in dMRI. Initially we have considered the case of rigid motion, which results in linear phase corruption, as proven by (Anderson and Gore, 1994). Under this assumption, phase maps are completely characterized by three parameters: one related to the constant offset and two

defining the slopes along the PE and FE directions. With the aim of reconstructing the diffusion image, we have used a SENSE–based Maximum Likelihood formulation, which has already been proposed for navigated ms–EPI, as in (Van et al., 2011). However, our proposal has the advantage of increasing time efficiency since it is able to estimate the phase–maps from the imaging data itself. Moreover, our philosophy has already been used successfully by (Ianni J.D., 2018), who were able to correct for even–odd mismatches, although their model only considers $\boldsymbol{k}$–space shifts along the fully–sampled readout direction.

Importantly, we have extended our formulation to deal with the more complex case of non–linear phase corruption maps, which can arise from non–rigid pulsatile motion. In order to include this scenario in our formulation, we have assumed that the resulting phase–maps can be assumed to vary smoothly and therefore we have used B–splines as a means of modelling the maps. Consequently, our final algorithm preserves the linear phase component, which is assumed to be responsible of rapid changes, and adds the non–linear component in order to capture slow variations resulting from non–rigid motion. This idea was inspired from the JSENSE formulation proposed by (Ying and Sheng, 2007), where they jointly reconstruct the image and estimate the sensitivity profiles. In our case, the sensitivity maps are known beforehand, and the phase–maps play the role of the coil profiles.

Key to the success of our method is the initialization of the algorithm. Our cost function is not convex and the algorithm could easily get trapped in local minima. In order to avoid it, our current implementation takes advantage of the ideas behind methods such as MUSE. Accordingly, we initialize the phase maps using a linearly parametrized MUSE reconstruction that identifies the position and phase of the peak of the spectrum for each shot. In order to test the ability of our method to reconstruct the diffusion images, we have carried out different experiments on synthetic data varying the level of noise degradation and number of shots. Moreover, we have tested our algorithms using a real phantom as well as synthetized data corrupted with real motion–derived phase–maps. Importantly, our method seems able to obtain accurate reconstructions up to 8 shots for linear and smooth phase corruption.

This study presents some limitations. First, we not tested the methods in an *in–vivo* ms–EPI acquisition, but instead we have used a simulated data and a real phantom. However, from the experiments using real phase–coruptions maps, we expect the linear assumption to be reasonable for certain areas of the brain, or if cardiac triggering is used during the acquisition to limit the non-linear phase effects of pulsatile motion (Miller and Pauly, 2003; O'Halloran et al., 2013). For this reason, we believe the results are promising regarding its applicability to clinical scenarios. Second, we have only compared our method to our own implementation of MUSE (Chen et al., 2013) and a linearly parametrized version of it. These methods, however, contain different parameters that need to be tuned, and although we have tried our best to be fair, room for improvement may exist. Furthermore, other algorithms such as LORAKS (Lobos et al., 2018) or MUSSELS (Mani et al., 2017) could be considered as

well. Third, our method seems sensitive to the initialization, which could motivate the study of different alternatives that may be more robust to g-factor noise amplification, such as using the centre of mass or the MLE for estimating the linear phase parameters. Furthermore, we have separated the global optimization problem and iterate between solving each of the three subproblems, but a global optimization could enhance the performance. Fourth, we have only considered tha case of Cartesian sampling, but our formulation is compatible with different trajectories such as spirals, so it would be interesting to test it under different sampling patterns. Finally, we have focused on correcting for microscopic intra–shot motion, which causes phase discrepancies between different shots due to the diffusion–sensitized gradients. However, it implicitly assumes there is no need to correct for shot–to–shot misalignments. We believe our method could be extended to deal with both microscopic intra–shot motion and macroscopic inter–shot motion, similarly to what has been done with MUSE in (Guhaniyogi et al., 2016). In order to do so, one could incorporate an extra term in the Maximum Likelihood formulation that accounts for shot–to–shot motion, which can be done following (Cordero-Grande et al., 2016).

5.6 Conclusion

We have developed a method that builds upon state-of-the-art techniques and is able to better estimate the phase maps corruption on ms–EPI acquisitions, resulting in an increased ability to remove ghosting artifacts from the reconstructed images. Under the assumption of linear or smooth non–linear phase corruption, we pose a joint formulation that is able both to reconstruct the images and to estimate the phase maps in a greedy iterative fashion.

CHAPTER

SIX

CONCLUSIONS, CONTRIBUTIONS AND FUTURE WORK

Like any other type of diagnostic imaging modality, MRI is also susceptible to different kinds of artifacts. Multiple sources can cause them, such as hardware imperfections in the MR scanner, physical/temporal constraints, patient motion or improper modelling of the image formation process. Moreover, they present a high variability depending on the particular type of MR scanner, imaging modality or reconstruction pipeline. Nevertheless, the overall result is that they can introduce undesired distortion in the final image, which can lead into an incorrect diagnose. For this reason, understanding their sources and effects is key in order to be able to prevent and/or remove them.

MRI is a powerful modility due to its flexibility and sensitivity to different tissue properties, but the majority of MR sequences require a lengthy scan time in order to acquire sufficient data to form an image. Importantly, long scan times can result in uncomfort for patients and increase the likelihood of motion during the acquisition, which in turn would introduce artifacts in the images. Furthermore, prolonged acquisition times hinder the obtention of images of moving structures such as the heart, or in which contrast evolves over time, as is the case with flowing blood in an MR angiography. Consequently, there is an interest in reducing the scanning time, but this might come at the expense of introducing new types or artifacts.

In this dissertation, we have focused on two different ways of accelerating MRI acquisitions and the related artifacts we can observe in the reconstructed images. On the one hand, PI exploits the availability of multiple coils to reduce the amount of $\boldsymbol{k}$–space data it is necessary to collect to create the image. However, the sub–sampling of the $\boldsymbol{k}$–space renders more complex the reconstruction process, which no longer consists of a simple 2D/3D-iFT. Moreover, the more complicated reconstructions result in the appearance of spatially varying noise maps in

the final image, violating the common assumption of homogeneous noise corruption. Accurate characterization of the noise statistics has proven essential for many different tasks, but it has been considered very challenging for one of the most commonly used PI techniques: GRAPPA. It operates directly in the $\boldsymbol{k}$–space estimating the missing data as a lineal combination of local neighborhoods. This process introduces local correlations in $\boldsymbol{k}$–space that need to be taken into account when data is transformed into the image domain. As a result, one needs to work with very large covariance matrices, which rendered the direct noise propagation analysis computationally infeasible. In this thesis, we have shown how to provide an exact noise characterization both for 2D and 3D GRAPPA acquisitions by exploiting the presence of extensive symmetries and the block separability in the reconstruction steps.

On the other hand, in contrast to Spin–Warp Imaging, Echo–Planar Imaging segments the acquisition into multiple shots by sampling multiple lines of the $\boldsymbol{k}$–space within a single excitation, thus reducing the scan time. However, they are particularly problematic in dMRI due to shot–to–shot discrepancies. This modality is designed to be sensitive to the microscopic movements of water molecules, which is encoded in the phase of the spins. Consequently, it is sensitive as well to any kind of bulk motion from the patient, which translates into phase corruption of the images. If the motion differs from shot to shot, the resulting phase discrepancies will lead into ghosting artifacts in the reconstructed images. In this dissertation, we have proposed a navigator–free scheme based on a Maximum Likelihood formulation to iteratively reconstruct the image and estimate the phase–maps under the assumption of linearity or smoothness.

Knowledge of the origin of MRI artifacts allows for better modelling of their effects, and it is the first step in order to elliminate them. In the case of noise corruption, we have focused on providing an exact estimation of noise amplification maps due to GRAPPA reconstructions under the assumption of stationarity and uncorrelation in the original $\boldsymbol{k}$–space undersampled acquisition. Furthermore, the availability of accurate characterization of noise maps has proven crucial for many different tasks such as quality assurance (Krissian and Aja-Fernández, 2011; Aja-Fernández et al., 2013), protocol optimization (Thünberg and Zetterberg, 2007; Saritas et al., 2011), and tailoring of subsequent post–processing steps (Ghugre et al., 2006; Rabanillo et al., 2016a; Veraart et al., 2011). Alternatively, in the case of ms–EPI dMRI acquisitions suffering from shot–to–shot phase discrepancies derived from patient motion, we have directly modelled subject motion and incorporated it into the reconstruction problem. In this way, we have been able to recover ghosting–free images up to 8 shots under both linear and smooth phase corruption.

6.1 Directions for future research

We have focused on Cartesian sampling for the study of noise amplification in GRAPPA reconstructions. However, non–Cartesian sampling schemes have proven to be an interesting alternative since they offer an efficient $\boldsymbol{k}$–space coverage and can supress off-resonance effects. Nevertheless, they usually require a regridding operation into a Cartesian scheme before the FT is applied. Amongst the many approaches proposed to do so we can include the GRAPPA philosophy, which has been successfully applied both to fully–sampled (see (Seiberlich et al., 2007, 2008b)) and undersampled data (see (Seiberlich et al., 2008a)). The reconstruction consists of a first step where each sampled point is shifted to the closest location in the Cartesian grid by a simple scaling as defined by the GROG operator, providing a non–uniform Cartesian coverage. Second, the missing points are estimated from its adjacents using the GRAPPA weights, which are no longer stationary across the $\boldsymbol{k}$–space due to the non–uniformity. Finally the FT is applied and the different channels are combined to form the final image. Regarding noise, the first step preserves the $\boldsymbol{k}$–space uncorrelation, although it introduces spatial heterogeneity. However, once the data is regridded, the noise analysis should not differ from the case of random sampling, where the $\boldsymbol{k}$–space coverage is likewise non–uniform. It would suffice to incorporate the level of noise and cross–correlation between different channels at the initial points.

Importantly, one of the fields that has been revolutionised by the developments of PI is the field of Multi–Band (MB) Imaging, also known as Simultaneous Multi–Slice (SMS). SMS Imaging dates back to the late 1980s, when researches proposed to increase the scan efficiency by exciting several slides simultaneously and thus removing the need for an RF pulse per slice. This was made possible thanks to the understanding of the physics behind selective RF pulses (Muller, 1988). Nevertheless, since phased–arrays were not developed yet, an RF–pulse was used to modulate the phase (Glover, 1991) of the excited slices and the FOV was increased in order to fit all of them. However, this FOV expansion resulted in the need to acquire the same amount of data than in a conventional single–slice acquisition. To further increase the scan efficiency, researchers proposed to build on the idea of phase cycling modulations to shift the different slices, but without increasing the FOV. Moreover, this is equivalent to the application of a gradient encoding along the slice direction (Setsompop et al., 2012), allowing for a direct acquisition in a 3D $\boldsymbol{k}$–space. Nonetheless, preserving the original FOV results in the expected overlapping of the different slices, which could later be separated during reconstruction by means of PI techniques such as GRAPPA (Setsompop et al., 2012). Hence, our noise characterization algorithm could be directly applied to SMS scenarios. Furthermore, it is important to notice that SMS acquisitions take advantage of 3D acquisitions by means of Fourier averaging. Consequently, SMS reduces the acquisition time without incurring in a SNR penalty with respect to single slice imaging, other than g-factor

effects, which could be quantified by our method.

An application where the availability of the exact noise maps could improve the start–of–the art is the estimation of decay parameters in accelerated multi–contrast acquisitions individually reconstructed using GRAPPA. In such scenarios, the g–factor maps might vary for the different images since the GRAPPA weights are separately computed for each image. This is done assuming they will not differ much since they are intimately connected to the sensitivity profiles of the channels, which are expected to be stable over time in such an acquisition. However, we have observed (see (Rabanillo et al., 2016a)) that they can actually vary across the different images, rendering further computation sub–optimal in scenarios such as T2*–estimation or ADC estimation in multi–echo or multi b–value acquisitions, respectively. A potential path to achieve a more accurate estimation could be to carry out a maximum–likelihood estimation of the parameters of interest. The optimum estimator under Gaussian, independent and identically distributed noise corruption would be a weighted least squares that should incorporate the noise maps variability. Alternatively, one could consider as well a joint procedure that uses all the available data to compute a single set of GRAPPA weights and thus preserves the same noise maps for all of the reconstructed images.

On the other hand, the algorithm we have developed for image reconstruction in ms–EPI reconstruction under shot–to–shot phase map discrepancies seems to present a strong dependancy on the initialization. We have proposed to initialize the phase maps by means of a shot–by–shot regularized SENSE reconstruction followed by a detector to identify the position of the spectrum peak, which is directly related to the slope of the linear phase components. However, this step can become very sensitive to the presence of noise. In particular, it could result in an improper initialization in scenarios of low Signal–to–Noise ratio or high noise amplification due to g–factor effects for a high number of shots. For this reason, it could be interesting to explore stochastic algorithms that can deal with non–convex optimization and which could excape local minima without the need to initialize the iterative procedure close to the global minimum.

Another alternative to approach this problem would be to learn a function for the estimation of the phase maps of an image given its multi–shot sampled $\boldsymbol{k}$–space data, where each shot is subject to a different phase corruption, similar to what has been done by (Balakrishnan et al., 2018) for the image registration problem. This could be done, for example, by means of a convolutional neural network that takes as an input the multi–shot $\boldsymbol{k}$–space and returns the shot phase map images. Once the network is trained, it could be integrated in our iterative algorithm as a module to solve the non–convex phase estimation problem. In this fashion, not only this step could be accelerated (at the cost of the training step), but hopefully one could overcome or, at least, alleviate, the initialization dependance.

Finally, our ms–EPI is focused on correcting for phase distortions derived from microscopic intra–shot motion. However, it would be interesting to extend this philosophy with the

aim of dealing with macroscopic misregistrations between different shots as well, similar to the extension proposed for MUSE, presented by (Guhaniyogi et al., 2016). In order to incorporate this idea into our ML formulation, we could include an extra term accounting for the macroscopic deformation for every shot. Once again, different deformation models could be applied, such as rigid (Cordero-Grande et al., 2016) or non–rigid deformations (Royuela-del Val et al., 2018). Importantly, as it is pointed out in (Guhaniyogi et al., 2016), it would be necessary to obtain a pre-estimate of the diffusion tensor and incorporate it into the recosntruction problem. This is because macroscopic motion would cause the different shots to suffer from different attenuation due to diffusion, since the effective diffusion gradient they observe would be different.

6.2 Contributions

In this dissertation we have focused on better characterizing the distortion introduced by two different types of artifacts in accelerated acquisitions and we have developed different tools to deal with them in a feasible computational time, namely, the estimation of exact g–factor maps in 2D/3D GRAPPA acquisitions and the reconstruction of multishot Diffusion–Weighted Images suffering from intra–shot motion.

The main contributions of this dissertation are the following:

1. *Exact and time–efficient algorithm for the estimation of g–factor maps in 2D GRAPPA acquisitions under the assumptions of stationarity and uncorrelation in the original undersampled $\boldsymbol{k}$–space acquisition.* GRAPPA reconstruction operates directly in $\boldsymbol{k}$–space and therefore an exact noise analysis requires to take into account the propagation of noise correlation through the whole imaging process, including interpolation, FT and coil combination. This procedure can be very time and memory consuming due to the size of the correlation matrices involved. To overcome this limitation, we have proposed an algorithm that exploits the symmetries in the aforementioned matrices as well as the separability into blocks that can be processed independently.

2. *Exact and time–efficient algorithm for the estimation of g–factor maps in 3D GRAPPA acquisitions under the assumptions of stationarity and uncorrelation in the original undersampled $\boldsymbol{k}$–space acquisition.* 3D MRI acquisitions provide higher SNR, are subject to higher acceleration factors and can help reduce distortions derived from through–plane motion. However, memory requirements can be a limiting factor since the 3D GRAPPA reconstruction introduces $\boldsymbol{k}$–space correlations along the three dimensions. Furthermore, they also offer greater versatility regarding $\boldsymbol{k}$–space sampling since both PE directions can be used to reduce acquisition time. We have extended our previous

algorithm to deal with the case of 3D acquisitions by, once again, exploiting symmetries and separability. Moreover, in order to alleviate memmory requirements, our algorithm includes a step to identify the unique sampling patterns present in the undersampling scheme.

3. *Joint image reconstruction and phase map estimation method for multi–shot dMRI under the assumption of intra–shot rigid motion.* Multi–shot dMRI is very sensitive to bulk motion since discrepancies between the different shots results in different phase corruption for each of them. If not corrected for, this would introduce ghosting artifacts in the reconstructed image. With the aim of increasing the efficiency of the acquisition, we have proposed a free–navigated algorithm that exploits the knowledge that rigid motion results in linear phase maps corruption. Parametrizing in this manner the phase maps has allowed us to pose a maximum likelihood formulation that iteratively estimates both the parameters characterizing the linear phase corruption and the reconstructed image.

4. *Joint image reconstruction and phase map estimation method for multi–shot dMRI under the assumption of intra–shot non–rigid motion.* The assumption of rigid–motion may not be realistic for certain regions due to the impact of pulsatile motion. This type of motion introduces non–linear components in the phase maps, which need to be modelled in order to obtain a successful reconstruction. We have proposed to use B–splines to capture this non–linear variation under the assumption of smoothness. Therefore, we have extended our joint algorithm to alternate between the estimation of the linear component parameters, the parameters characterizing the B–spline and the reconstruction of the image.

APPENDIX

ONE

PROOF THAT IMAGE–SPACE EQUIVALENT RECONSTRUCTION DOES NOT HOLD EXACT FOR VD SAMPLING PATTERNS

The method described in (Breuer et al., 2009) performs the image–space analysis for the VD-GRAPPA case by assuming that the $\boldsymbol{k}$–space reconstruction has an equivalence in the image domain. They consider that the $\boldsymbol{k}$–space GRAPPA reconstruction is equivalent to performing a convolution by regions. This way, the $\boldsymbol{k}$–space is separated in different regions (equivalent to filtering the $\boldsymbol{k}$–space), each region is convolved with its own kernel, and then the resulting regions are simply added. In the image–space this would be equivalent to transform each $\boldsymbol{k}$–space region to the image domain via iFT, multiply pixel–wise (according to the convolution theorem) each resulting image by the iFT of its corresponding kernel, and sum the obtained images.

For noise characterization purposes, this would allow to perform the noise analysis in the image–domain for each region, and then add the resulting noise maps. However, for this to be valid, it is necessary that these regions are uncorrelated. Otherwise, the final step is not a simple addition of the noise maps for each region. It would be necessary to take into account these correlations between the regions.

The main problem is that this reconstruction is not exactly equivalent to what GRAPPA reconstruction does. We start by the case of an uniform undersampling pattern with an ACS region (see Fig.A.1). The reconstruction would separate the $\boldsymbol{k}$–space in two regions (the ACS region and the undersampled region), and then apply a different kernel to each region (an identity kernel to the ACS region and the GRAPPA kernel to the undersampled region).

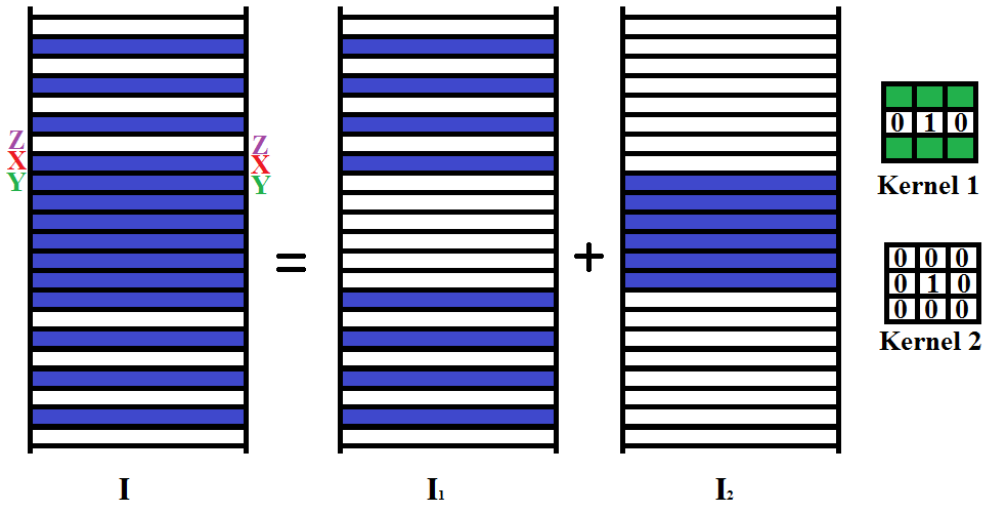

Figure A.1: Example of the equivalent reconstruction proposed by (Breuer et al., 2009) for a uniform undersampling pattern with acceleration rate R=2 and an ACS region for calibration. The original $\boldsymbol{k}$–space I is separated into two regions: the undersampled region (I_1) and the ACS region (I_2). Each region is convolved with its own kernel, i.e., the GRAPPA kernel and an identity kernel respectively, and finally the results are added. However, this reconstruction is not equivalent to GRAPPA reconstruction because of the boundaries between regions. This is the case of line X, that should not change during the reconstruction. If it is placed at I_1, the application of the kernel will introduce a residual error at line Y in image I_1, whereas if it is placed at I_2, then line Z will not be reconstructed properly in I_1 because line X is not present there.

However, the boundaries (for example line X in Fig.A.1) between the two regions give rise to uncertainity.

This acquired line, that should not change during the reconstruction, could be placed both at I_1 or I_2 since both kernels will leave it untouched. If it is considered in I_1 (as in Fig.A.1), the application of the kernel will introduce a residual error at line Y in image I_1. This is because when the GRAPPA kernel is applied to line Y at I_1 it is reconstructed from line X, but line Y should not change since it is in the ACS region, and this residual error is preserved when the two regions are added. On the other hand, if line X is placed at I_2, then line Z will not be reconstructed properly in image I_1 because line X is not present there.

For the more complex case of a variable density undersampling, such as the one shown in Fig.A.2, there is still an ambiguity in the boundaries. If line X is considered in I_1, then lines Y_1, Y_2 and Y_3 will not be properly reconstructed. The reason is twofold. **First** because of the residual introduced at lines Y_i by the application of the first kernel to line X (these lines should be reconstructed with the second kernel, not with the first one). **Second** because line X is missing in I_2, so the application of the second kernel to lines Y_i will not reconstruct them properly since there is not data at line X to reconstruct them. Equivalently, if line X is included in I_2, then lines Z_1, Z_2 and Z_3 will not be reconstructed in the right manner. And last, if line X is included in both I_1 and I_2, it will be present when reconstructing all lines,

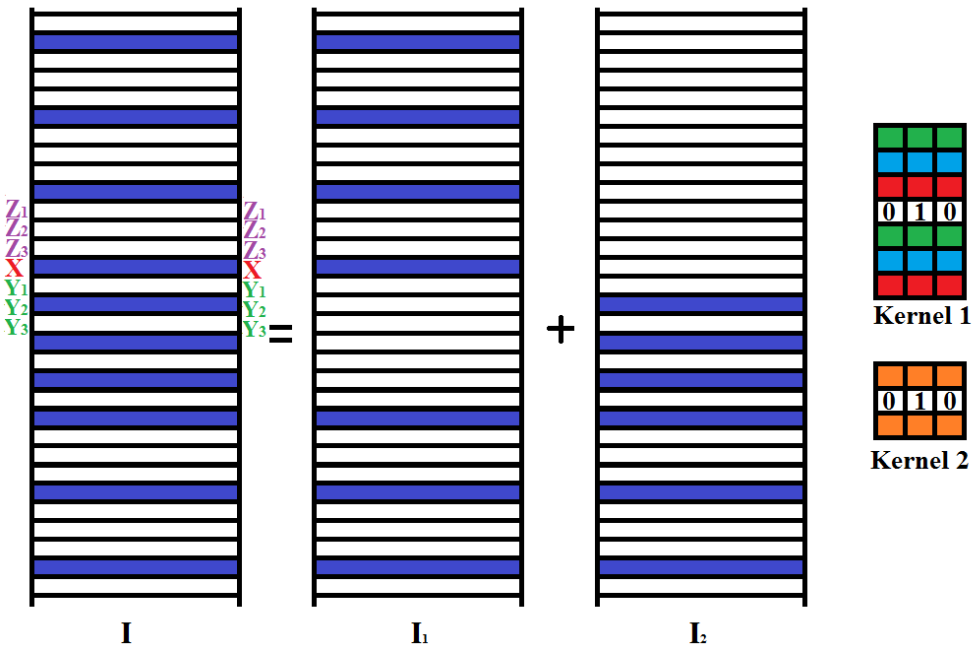

Figure A.2: Example of the equivalent reconstruction proposed by (Breuer et al., 2009) for a variable-density pattern with acceleration rates R=2 and R=4. If the boundary line X is considered in I_1, then lines Y_1, Y_2 and Y_3 will not be properly reconstructed and the reason is twofold: because of the residual introduced at lines Y_i by the application of the first kernel to line X at I_1 and because line X is missing in I_2. Equivalently, if line X is included in I_2, then lines Z_1, Z_2 and Z_3 will not be reconstructed in the right manner. And last, if line X is included in both I_1 and I_2, it will be present when reconstructing all lines, but it would appear twice when the regions are added, it will correlate the regions and it will still introduce the residual errors.

but it would appear twice when the regions are added, it will correlate the regions and it will still introduce the residual errors.

To sum up, the image–space methodology produces two sources of error. First, the reconstruction in the image–space is not exactly equivalent to the reconstruction in $\boldsymbol{k}$–space. When a kernel is applied to a filtered region, the support of the output overlaps with the adjacent regions, introducing a residual error. These residuals could be removed by an extra filtering before combining the regions. However, regarding the noise analysis, this filtering becomes a convolution in the image–space and it would be necessary to know the correlation between points in the image–space (correlations introduced by the initial filtering used to separate the regions). Second, when different acceleration factors are used, the acquired lines in the boundaries that separate every pair of regions are used to reconstruct both regions. Consequently, the two regions will be correlated even if the extra-filtering is included, contradicting the initial assumption.

APPENDIX

TWO

PROOF THAT SHIFTED COLUMNS ARE EQUIVALENT

In this appendix we will present the proof that two rows that present the same pattern except for a shift provide a matrix $\mathbf{\Gamma}_3^{m_{uq}}$ with the same diagonal. We can write the correlation matrix $\mathbf{\Gamma}$ of a reference row with another row as follows:

$$\mathbf{\Gamma}_{a,1} = \begin{pmatrix} \vec{\mathbf{c}_1}[n] & \vec{\mathbf{c}_2}[n] & \cdots & \vec{\mathbf{c}_{N_x}}[n] \end{pmatrix}, \tag{B.1}$$

where the vectors $\vec{\mathbf{c}_i}[n], n = \{0 \dots N_x - 1\}$ denote the columns of this matrix. Let us consider a different reference row that correlates with a surrounding one as a shifted version of the already described reference row. We will initially consider a shift of a single element. Consequently, its correlation matrix can be described as:

$$\mathbf{\Gamma}_{b,1} = \begin{pmatrix} \vec{\mathbf{c}_{N_x}}[n-1] & \vec{\mathbf{c}_1}[n-1] & \cdots & \vec{\mathbf{c}_{N_x-1}}[n-1] \end{pmatrix}. \tag{B.2}$$

It should be noted that $\vec{\mathbf{c}_i}[n-1], n = \{0 \dots N_x - 1\}$ denotes a circular shift of the column vector. Computing the iFT along the row dimension results in operating as follows on the $\mathbf{\Gamma}$ matrix, as stated in Eq.3.20:

$$\mathbf{\Gamma}_2 = \mathbf{F}_I \cdot \mathbf{\Gamma}_1 \cdot \mathbf{F}_I^H = [\mathbf{F}_I \cdot (\mathbf{F}_I \cdot \mathbf{\Gamma}_1)^H]^H = [\mathbf{F}_I \cdot \mathbf{\Psi}^H]^H. \tag{B.3}$$

where we have defined an intermediate matrix $\mathbf{\Psi} = \mathbf{F}_I \cdot \mathbf{\Gamma}_1$. Defining this intermediate matrix for the first reference column as:

$$\mathbf{\Psi}_a = \begin{pmatrix} \vec{\mathbf{d}_1}[n] & \vec{\mathbf{d}_2}[n] & \cdots & \vec{\mathbf{d}_{N_x}}[n] \end{pmatrix}, \tag{B.4}$$

and using the Shift Theorem of the FT, it is inmediate to observe that:

$$\mathbf{\Psi}_b = \begin{pmatrix} \vec{\mathbf{d}_{N_x}}[n] \cdot \vec{\mathbf{w}}[n] & \vec{\mathbf{d}_1}[n] \cdot \vec{\mathbf{w}}[n] & \cdots & \vec{\mathbf{d}_{N_x-1}}[n] \cdot \vec{\mathbf{w}}[n] \end{pmatrix}, \tag{B.5}$$

where $\vec{\mathbf{w}}[n] = exp\left(\iota\frac{2\cdot\pi\cdot n}{N_x}\right), n = \{0\dots N_x - 1\}$. To avoid confusion with the indices, we have denoted the imaginary unit by ι. If we denote the rows of the matrix $\mathbf{\Psi}_a$ as $\vec{\mathbf{r}}_j^a[n], n = \{0\dots N_x - 1\}$, we can see that the rows of the matrix $\mathbf{\Psi}_b$ can be expressed as $\vec{\mathbf{r}}_j^b[n] = \vec{\mathbf{r}}_j^a[n-1]\cdot exp\left(\iota\frac{2\cdot\pi\cdot j}{N_x}\right)$. When we compute the final matrix $\mathbf{\Gamma}_2$, we operate on $\mathbf{\Psi}^H$, that is to say, on the rows of $\mathbf{\Psi}$. Using again the Shift Theorem, we can see that:

$$\mathbf{\Gamma}_{b,2} = \mathbf{\Gamma}_{b,1} \circ exp\left(\iota\frac{2\cdot\pi\cdot(i-j)}{N_x}\right), \tag{B.6}$$

where $\{i, j\} = \{0\dots N_x - 1\}$ denote the row and column of the matrix. After operating along the first dimension, we only care about the correlations of a reference point with other points that are in the same plane (equal position in the first dimension). This means that we only need to keep the diagonal of the previous matrix, as already explained. Comparing the diagonals of these two matrices, we can clearly see that $\text{diag}(\mathbf{\Gamma}_{b,2}) = \text{diag}(\mathbf{\Gamma}_{b,2})$. Finally, we have derived this proof for the case of a one–element shift. For the general case of a shift of any size, since it can be seen as the successive application of many one-element shifts, the previous equality will hold.

APPENDIX

THREE

DETAILS FOR THE IMPLEMENTATION OF LINEAR PHASE MAP'S PARAMETERS OPTIMIZATION

In order to solve Eq.5.32*, we can define:

$$f(\vec{\boldsymbol{\theta}}) = \left\| \mathbf{A}\boldsymbol{\mathcal{F}}\mathbf{S}\boldsymbol{\mathcal{P}}(\vec{\boldsymbol{\theta}})\vec{\mathbf{x}} - \vec{\mathbf{y}} \right\|_2^2, \tag{C.1}$$

which allows us to compute the gradient

$$\nabla_l^f = \frac{\partial f(\vec{\boldsymbol{\theta}})}{\partial \theta_l} = 2\mathbb{R}e\left\{ \vec{\mathbf{w}}^H \vec{\mathbf{w}}_{(l)} \right\}, \tag{C.2}$$

where

$$\begin{cases} \vec{\mathbf{w}} & = \mathbf{A}\boldsymbol{\mathcal{F}}\mathbf{S}\boldsymbol{\mathcal{P}}(\vec{\boldsymbol{\theta}})\vec{\mathbf{x}} - \vec{\mathbf{y}} \\ \vec{\mathbf{w}}_{(l)} & = \mathbf{A}\boldsymbol{\mathcal{F}}\mathbf{S}\frac{\partial \boldsymbol{\mathcal{P}}(\vec{\boldsymbol{\theta}})}{\partial \theta_l}\vec{\mathbf{x}} \end{cases} \tag{C.3}$$

and the Hessian is given by:

$$H_{lm}^f = \frac{\partial^2 f(\vec{\boldsymbol{\theta}})}{\partial \theta_l \theta_m} = 2\mathbb{R}e\left\{ \vec{\mathbf{w}}_{(l)}^H \vec{\mathbf{w}}_{(m)} + \vec{\mathbf{w}}^H \vec{\mathbf{w}}_{(lm)} \right\}, \tag{C.4}$$

where

$$\vec{\mathbf{w}}_{(lm)} = \mathbf{A}\boldsymbol{\mathcal{F}}\mathbf{S}\frac{\partial^2 \boldsymbol{\mathcal{P}}(\vec{\boldsymbol{\theta}})}{\partial \theta_l \theta_m}\vec{\mathbf{x}}. \tag{C.5}$$

Importantly, $\boldsymbol{\mathcal{P}}(\vec{\boldsymbol{\theta}})$ is a diagonal matrix whose is given by Eq.5.28:

$$\vec{\mathbf{p}}(\vec{\boldsymbol{\theta}}) = e^{j(\theta_0 + \theta_1 \cdot r_1 + \theta_2 \cdot r_2)}, \tag{C.6}$$

*We will drop the shot dependance since the problem can be separated and solved shot–by–shot

and therefore its gradient and Hessian will be diagonal matrices as well. In particular, the diagonal of the gradient will be given by:

$$\frac{\partial \vec{\mathbf{p}}(\vec{\boldsymbol{\theta}})}{\partial \theta_l} = \begin{cases} j \cdot \vec{\mathbf{p}} & \text{if } l = 1 \\ j \cdot \vec{\mathbf{r}}_l \circ \vec{\mathbf{p}} & \text{if } l = \{2, 3\} \end{cases}, \tag{C.7}$$

whereas the diagonal of the Hessian can be computed as:

$$\frac{\partial^2 \vec{\mathbf{p}}(\vec{\boldsymbol{\theta}})}{\partial \theta_l \theta_m} = \begin{cases} -\vec{\mathbf{p}} & \text{if } l = 1, m = 1 \\ -\vec{\mathbf{r}}_l \circ \vec{\mathbf{r}}_m \circ \vec{\mathbf{p}} & \text{if } l = \{2, 3\}, m = \{2, 3\} \\ -\vec{\mathbf{r}}_m \circ \vec{\mathbf{p}} & \text{if } l = 1, m = \{2, 3\} \end{cases}. \tag{C.8}$$

We use a damped version of Newton's method in which each iteration is given by the rule:

$$\vec{\boldsymbol{\theta}}_{(n+1)} = \vec{\boldsymbol{\theta}}_{(n)} - \left(w_{(m)} \boldsymbol{Id} + \boldsymbol{H}^f \right)^{-1} \cdot \nabla_l^f, \tag{C.9}$$

where n denotes the phase estimation iteration and the weight $w_{(n)}$ is updated according to

$$w_{(n+1)} = \begin{cases} 2w_{(n)} & \text{if } f(\vec{\boldsymbol{\theta}}_{(n+1)}) > f(\vec{\boldsymbol{\theta}}_{(n)}) \\ w_{(n)}/1.2 & \text{otherwise} \end{cases}. \tag{C.10}$$

APPENDIX

FOUR

DESCRIPTION OF B-SPLINE PHASE–MAPS CHARACTERIZATION

We will model the non linear phase–maps by means of B–splines. For this purpose we define a mesh of knots over the image space as we illustrate in Fig.D.1, which will be given by coordinates $\vec{\mathbf{u}} = (u_1, \ldots, u_L) = (u_1, \ldots, u_L)$, where L refers to the number of dimensions ($L = 2$ in our case) and the knot coordinates take an integer value over the range $-C_l \leq u_l \leq D_l$. We can define the minimum box that covers our image (or our region of interest if we use a mask) by its limits $\left(\left[r_1^C, r_1^D\right], \ldots, \left[r_L^C, r_L^D\right]\right)$. In order to define the dimensions of our knots mesh, we can define as well the influence radius of a knot as:

$$R_l = \frac{(E+1)\cdot\Delta_l}{2}, l = 1 \ldots L, \tag{D.1}$$

where E refers to the order of the considered B-Spline, as we will define later on, and $\vec{\boldsymbol{\Delta}} = (\Delta_1, \ldots, \Delta_L)$ provides the resolution in pixels of the knots mesh along each dimension. Further, if we define the center of the knot mesh at point $\boldsymbol{c} = (c_1, \ldots, c_L)$, we can compute the coordinates of a knot $\vec{\mathbf{u}} = (u_1, \ldots, u_L)$ (given in the knots coordinate system) in the image coordinate system by simply doing:

$$\vec{\mathbf{p}}^{\boldsymbol{u}} = \vec{\mathbf{c}} + \vec{\boldsymbol{\Delta}} \circ \vec{\mathbf{u}}, \tag{D.2}$$

where $\circ$ refers to the Hadamard product.

In order to define the knots, which can cover a region bigger than the image, we notice a knot $\vec{\mathbf{u}}$ will reach the points in the image (or in the ROI) if and only if:

$$\left(p_l^{\boldsymbol{u}} + R_l > r_l^C\right) \wedge \left(p_l^{\boldsymbol{u}} - R_l < r_l^D\right), \tag{D.3}$$

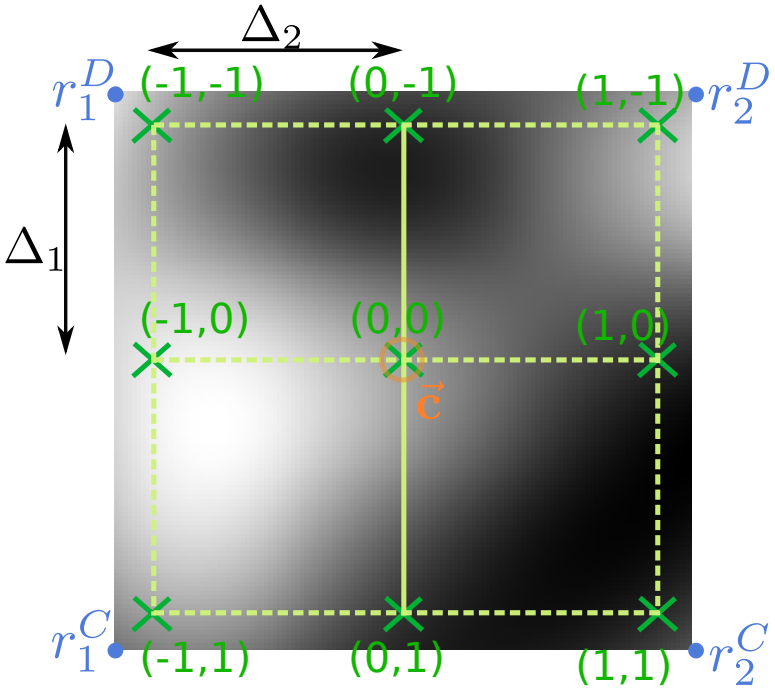

Figure D.1: Phase map defined over the 2D region $\left(\left[r_1^C, r_1^D\right], \left[r_2^C, r_2^D\right]\right)$, indicated in blue. We can define over it a knot mesh (green crosses) of size 3×3 in this example, where we indicate the coordinates of the knots $\vec{\mathbf{u}}$ in green next to them. The resolution of the knot mesh is given in the image coordinate system by (Δ_1, Δ_2).

where $\wedge$ denotes the logical AND operation. Using the previous inequalities, we can define the bound integers C_l and D_l in order to guarantee that it is not possible to add further knots whose influence includes the image (or ROI):

$$\begin{aligned} C_l &= \left\lfloor \frac{c_l - r_l^C + R_l}{\Delta_l} \right\rfloor, \\ D_l &= \left\lfloor \frac{r_l^D - c_l + R_l}{\Delta_l} \right\rfloor, \end{aligned} \tag{D.4}$$

where $\lfloor x \rfloor$ refers to the maximum integer lower or equal to x.

With these definitions, we can now define the B–spline to approximate the phase maps $\varphi_s^{nr}(\boldsymbol{r})$ resulting from non–rigid motion for each shot s:

$$\varphi_s^{nr}(\boldsymbol{r}, \vec{\boldsymbol{\xi}}_s) = \sum_{u_1=-C_1}^{D_1} \cdots \sum_{u_L=-C_L}^{D_L} \left(\prod_{l=1}^{L} B_E\left(v_l\left(r_l - p_l^{\boldsymbol{u}}\right)\right) \right) \xi_s^{\boldsymbol{u}}, \tag{D.5}$$

where we have as many parameters $\vec{\boldsymbol{\xi}}_s = \{\xi_s^{\boldsymbol{u}}\}$ for each shot as knots. Furthermore, we make use of the transform function $\vec{\mathbf{v}}(\vec{\mathbf{r}})$:

$$\vec{\mathbf{v}}(\vec{\mathbf{r}}) = \vec{\mathbf{r}} \circ \vec{\boldsymbol{\Delta}}^{-1}, \tag{D.6}$$

where $\vec{\boldsymbol{\Delta}}^{-1} = (1/\Delta_1, \ldots, 1/\Delta_L)$ refers to the Hadamard inverse of vector $\vec{\boldsymbol{\Delta}}$. Moreover, $B_E(r)$ represents the uniform B–spline of degree E, which can be constructed by means of the Cox-de

Boor recursion formula (Boor, 2003). As an example, the initial B–splines are given by:

$$
\begin{aligned}
B_0(r) &= \begin{cases} 1 & \text{if } -0.5 \le r \le 0.5 \\ 0 & \text{otherwise} \end{cases} \\
B_1(r) &= \begin{cases} r+1 & \text{if } -1 \le r \le 0 \\ 1-r & \text{if } 0 \le r \le 1 \\ 0 & \text{otherwise} \end{cases} \\
B_2(r) &= \begin{cases} \frac{4r^2+12r+9}{8} & \text{if } -1.5 \le r \le -0.5 \\ \frac{-4r^2+3}{4} & \text{if } -0.5 \le r \le 0.5 \\ \frac{4r^2-12r+9}{8} & \text{if } 0.5 \le r \le 1.5 \\ 0 & \text{otherwise} \end{cases} \\
B_3(r) &= \begin{cases} \frac{r^3+6r^2+12r+8}{6} & \text{if } -2 \le r \le -1 \\ \frac{-3r^3-6r^2+4}{6} & \text{if } -1 \le r \le 0 \\ \frac{3r^3-6r^2+4}{6} & \text{if } 0 \le r \le 1 \\ \frac{-r^3+6r^2-12r+8}{6} & \text{if } 1 \le r \le 2 \\ 0 & \text{otherwise} \end{cases} \\
&\dots
\end{aligned}
\tag{D.7}
$$

APPENDIX

FIVE

DETAILS FOR THE IMPLEMENTATION OF B-SPLINE'S PARAMETERS OPTIMIZATION

In this appendix we will provide details on the implementation of the Levenberg–Marquardt algorithm, which was our choice for the estimation of the B–spline's parameters $\vec{\boldsymbol{\xi}} = \{\xi^{\boldsymbol{u}}\}$. In order to solve Eq.5.40, we can define:

$$g(\vec{\boldsymbol{\xi}}) = \left\| \mathbf{A}\boldsymbol{\mathcal{F}}\mathbf{S}\boldsymbol{\mathcal{Q}}(\vec{\boldsymbol{\xi}})\boldsymbol{\mathcal{P}}(\vec{\boldsymbol{\theta}})\vec{\mathbf{x}} - \vec{\mathbf{y}} \right\|_2^2, \tag{E.1}$$

which allows us to compute the gradient

$$\nabla_l^g = \frac{\partial g(\vec{\boldsymbol{\xi}})}{\partial \xi^{\boldsymbol{l}}} = 2\mathbb{R}e\left\{ \vec{\mathbf{s}}^H \vec{\mathbf{s}}_{(l)} \right\}, \tag{E.2}$$

where

$$\begin{cases} \vec{\mathbf{s}} & = \mathbf{A}\boldsymbol{\mathcal{F}}\mathbf{S}\boldsymbol{\mathcal{Q}}(\vec{\boldsymbol{\xi}})\boldsymbol{\mathcal{P}}(\vec{\boldsymbol{\theta}})\vec{\mathbf{x}} - \vec{\mathbf{y}} \\ \vec{\mathbf{s}}_{(l)} & = \mathbf{A}\boldsymbol{\mathcal{F}}\mathbf{S}\frac{\partial \boldsymbol{\mathcal{Q}}(\vec{\boldsymbol{\xi}})}{\partial \xi^{\boldsymbol{l}}}\boldsymbol{\mathcal{P}}(\vec{\boldsymbol{\theta}})\vec{\mathbf{x}} \end{cases} \tag{E.3}$$

and the Hessian is approximated by the first term in Eq.C.5:

$$H_{lm}^g = \frac{\partial^2 g(\vec{\boldsymbol{\xi}})}{\partial \xi^{\boldsymbol{l}} \xi^{\boldsymbol{m}}} \approx 2\mathbb{R}e\left\{ \vec{\mathbf{s}}_{(l)}^H \vec{\mathbf{s}}_{(m)} \right\}. \tag{E.4}$$

Once again, $\boldsymbol{\mathcal{Q}}(\vec{\boldsymbol{\xi}})$ is a diagonal matrix, whose diagonal is given by:

$$\vec{\mathbf{q}}(\vec{\boldsymbol{\xi}}) = e^{j\varphi^{nr}(\vec{\mathbf{r}},\vec{\boldsymbol{\xi}})} \tag{E.5}$$

and therefore we can easily compute the derivative

$$\frac{\partial \vec{\mathbf{q}}(\vec{\boldsymbol{\xi}})}{\partial \xi^{\boldsymbol{m}}} = j\frac{\partial \varphi^{nr}(\vec{\mathbf{r}}, \vec{\boldsymbol{\xi}})}{\partial \xi^{\boldsymbol{m}}} e^{j\varphi^{nr}(\vec{\mathbf{r}},\vec{\boldsymbol{\xi}})}. \tag{E.6}$$

where the derivative of the phase map given in Eq.D.5 can be computed from:

$$\frac{\partial \varphi_s^{nr}(\vec{\mathbf{r}}, \vec{\boldsymbol{\xi}})}{\partial \xi^{\boldsymbol{m}}} = \prod_{l=1}^{L} B_E\left(v_l\left(r_l - p_l^{\boldsymbol{m}}\right)\right). \tag{E.7}$$

In order to implement the B–spline in Eq.D.5, we exploit the condition of compact support of B–splines. We know that the coordinates of a pixel $\boldsymbol{p}$ in the knots frame are given by:

$$\vec{\mathbf{u}}^p = (\vec{\mathbf{p}} - \vec{\mathbf{c}}) \circ \vec{\boldsymbol{\Delta}}^{-1}, \tag{E.8}$$

allowing us to infer the range of knots $\boldsymbol{\mathcal{U}}^{\boldsymbol{p}}$ that have influence on a given pixel, which we can store in a matrix:

$$\mathcal{U}_l^{\boldsymbol{p}} = \begin{cases} \left[\lfloor u_l^{\boldsymbol{p}} \rceil - \frac{E}{2}, \lfloor u_l^{\boldsymbol{p}} \rceil + \frac{E}{2}\right] & \text{if E is even} \\ \left[\lfloor u_l^{\boldsymbol{p}} \rfloor - \frac{E-1}{2}), \lfloor u_l^{\boldsymbol{p}} \rfloor - \frac{E+1}{2})\right] & \text{if E is odd} \end{cases}, \tag{E.9}$$

where $\lfloor x \rfloor$ refers to the maximum integer lower or equal to x (floor) and $\lfloor x \rceil$ refers to the closer integer to x (round).

Consequently, we can compute the contribution of $\prod_{l=1}^{L} B_E\left(v_l\left(r_l - p_l^{\boldsymbol{u}}\right)\right)$ for each pixel in the image. Further, we only need to do this computation for the knots that affect a given pixel, which are given by Eq.E.9. Assuming we have the same number of knots $(E+1)$ along each dimension, we can store these contributions in a $N_1 \times \cdots \times_L \times (E+1)^L$ matrix. Therefore, given the concrete values for the parameters vector $\vec{\boldsymbol{\xi}}$, in order to compute Eq.D.5 we simply need to select for each pixel the ones affecting it, given by Eq.E.9, multiply them by the corresponding contribution of the knot at the pixels and further add the whole contributions.

In order to compute the gradient in Eq.E.7, we proceed in a similar fashion. We know that a given knot $\boldsymbol{u}$ has influence over a region of pixels $\boldsymbol{\mathcal{X}}^{\boldsymbol{u}}$ whose limits are given by:

$$\mathcal{X}_l^{\boldsymbol{u}} = \left[\lceil p_l^{\boldsymbol{u}} - R_l \rceil, \lfloor p_l^{\boldsymbol{u}} + R_l \rfloor\right], \tag{E.10}$$

where R_l is the radius of influence of a knot defined in Eq.D.1 and $\lceil x \rceil$ refers to the minimum integer greater or equal to x (ceil).

Therefore, we can save these indices delimiting the pixels affected for each knot in a $(D_1 + C_1 + 1) \times \ldots \times (D_L + C_L + 1) \times 2 \times L$ matrix and the non–zero contributions of $\prod_{l=1}^{L} B_E\left(v_l\left(r_l - p_l^{\boldsymbol{u}}\right)\right)$ for each pixel at a given knot in another $(D_1 + C_1 + 1) \times \ldots \times (D_L + C_L + 1) \times (2\lfloor R_1 \rfloor + 1) \times \ldots \times (2\lfloor R_L \rfloor + 1)$. This allows us to precompute the derivative images of the B–spline $\frac{\partial \varphi_s^{nr}(\boldsymbol{r}, \boldsymbol{\xi})}{\partial \xi^{\boldsymbol{m}}}$ since it does not depend on the parameters themselves.

BIBLIOGRAPHY

Aja-Fernández, S., Brion, V., and Tristán-Vega, A. (2013). Effective noise estimation and filtering from correlated multiple-coil MR data. *Magnetic Resonance Imaging*, 31:272–285.

Aja-Fernández, S., Tristán-Vega, A., and Hoge, W. (2011). Statistical noise analysis in GRAPPA using a parametrized non-central chi approximation model. *Magnetic Resonance in Medicine*, 65(4):1195–1206.

Aja-Fernández, S., Vegas-Sánchez-Ferrero, G., and Tristán-Vega, A. (2014). Noise estimation in parallel MRI: GRAPPA and SENSE. *Magnetic Resonance Imaging*, 32(3):281–290.

Aja Fernández, S. and Vegas Sánchez-Ferrero, G. (2016). *Statistical Analysis of Noise in MRI*. Springer, Switzerland.

Alexander, D., Hubbard, P., Hall, M., Moore, E., Ptito, M., Parker, G., and T.B., D. (2010). Orientationally invariant indices of axon diameter and density from diffusion MRI. *Neuroimage*, 52(4):1374–1389.

Allison, M., Ramani, S., and Fessler, J. (2013). Accelerated regularized estimation of MR coil sensitivities using augmented Lagrangian methods. *IEEE Transactions on Medical Imaging*, 32(3):556–564.

Anderson, A. and Gore, J. (1994). Analysis and correction of motion artifacts in diffusion weighted imaging. *Magnetic Resonance in Medicine*, 32(3):379–387.

Assaf, Y. and Basser, P. (2005). Composite hindered and restricted model of diffusion (CHARMED) MR imaging of the human brain. *Neuroimage*, 27(1):48–58.

Atkinson, D., Counsell, S., Hajnal, J., Batchelor, P., Hill, D., and Larkman, D. (2006). Nonlinear phase correction of navigated multi-coil diffusion images. *Magnetic Resonance in Medicine*, 56(5):1135–1139.

Balakrishnan, G., Zhao, A., Sabuncu, M., Guttag, J., and Dalca, A. (2018). Voxelmorph: A learning framework fordeformable medical image registration. *IEEE Transactions on Medical Imaging*.

Barth, M., Breuer, F., Koopmans, P., Norris, D., and Poser, B. (2016). Simultaneous multislice (sms) imaging techniques. *Magnetic Resonance in Medicine*, 75(1):63–81.

Basser, P., Mattiello, J., and Le Bihan, D. (1994). MR diffusion tensor spectroscopy and imaging. *Biophysical Journal*, 66(1):259–267.

Beatty, P. and Brau, A. (2016). Analytical computation of g-Factor maps for autocalibrated Parallel Imaging. *Proceedings of the 16th Annual Meeting of the International Society for Magnetic Resonance in Medicine (ISMRM '8), Toronto, Canada,*, page 1297.

Bernstein, M. A., King, K. F., and Zhou, X. J., editors (2004). *Handbook of MRI Pulse Sequences*. Elsevier Academic Press, Burlington, MA, USA.

Blaimer, M., Breuer, F., Mueller, M., Seiberlich, N., Ebel, D., Heidemann, R., Griswold, M., and Jakob, P. (2006). 2D-GRAPPA-Operator for faster 3D parallel MRI. *Magnetic Resonance in Medicine*, 56:1359–1364.

Boor, C. d., editor (2003). *A Practical Guide to Splines*. Springer-Verlag, Berlin, Germany.

Brau, A., Beatty, P., Skare, S., and Bammer, R. (2008). Comparison of reconstruction accuracy and efficiency among autocalibrating data-driven parallel imaging methods. *Magnetic Resonance in Medicine*, 59:382–395.

Breuer, F., Blaimer, M., Heidemann, R., Mueller, M., Griswold, M., and Jakob, P. (2005). Controlled aliasing in parallel imaging results in higher acceleration (CAIPIRINHA) for multi-slice imaging. *Magnetic Resonance in Medicine*, 53:684–691.

Breuer, F., Kannengiesser, S., Blaimer, M., Seiberlich, N., Jakob, P., and Griswold, M. (2009). General formulation for quantitative g-factor calculation in GRAPPA reconstructions. *Magnetic Resonance in Medicine*, 62(3):739–746.

Brown, R. W., Cheng, Y. N., Mark, H., R., T. M., and Ramesh, V. (2014). *Magnetic Resonance Imaging: Physical Principles and Sequence Design*. Wiley–IEEE Press, Piscataway, NJ, USA.

Butts, K., de Crespigny, A., Pauly, J., and Moseley, M. (1996). Diffusion-weighted interleaved echo-planar imaging with a pair of orthogonal navigator echoes. *Magnetic Resonance in Medicine*, 35(5):763–770.

Carlson, J. and Minemura, T. (1993). Imaging time reduction through multiple receiver coil data acquisition and image reconstruction. *Magnetic Resonance in Medicine*, 29(1):681–687.

Chen, N., Guidon, A., Chang, H., and Song, A. (2013). A robust multi-shot scan strategy for high-resolution diffusion weighted MRI enabled by multiplexed sensitivity-encoding (MUSE). *Neuroimage*, 72:41–47.

Collins, C., Liu, W., Schreiber, W., Yang, Q., and Smith, M. (2005). Central brightening due to constructive interference with, without, and despite dielectric resonance. *Journal on Magnetic Resonance Imaging*, 21(2):192–196.

Constantinides, C., Atalar, E., and McVeigh, E. (1997). Signal-to-noise measurements in magnitude images from NMR phased arrays. *Magnetic Resonance in Medicine*, 38(5):852–857.

Cordero-Grande, L., Teixeira, R. P. A. G., Hughes, E. J., Hutter, J., Price, A. N., and Hajnal, J. V. (2016). Sensitivity encoding for aligned multishot magnetic resonance reconstruction. *IEEE Transactions on Computational Imaging*, 2(3):266–280.

Davis, P. (1979). *Circulant matrices*. Wiley.

Deshmane, A., Gulani, V., Griswold, M., and Seiberlich, N. (2012). Parallel MR Imaging. *Journal of Magnetic Resonance Imaging*, 36(1):55–72.

D.L., C., Zijdenbos, A., Kollokian, V., Sled, J., Kabani, N., Holmes, C., and Evans, A. (1998). Design and construction of a realistic digital brain phantom. *IEEE Transactions on Medical Imaging*, 17(3):463–468.

Dong, Z., Wang, F., Ma, X., Zhang, Z., Dai, E., Yuan, C., and Guo, H. (2018). Interleaved EPI diffusion imaging using spirit-based reconstruction with virtual coil compression. *Magnetic Resonance in Medicine*, 79(3):1525–1531.

Eichner, C., Cauley, S., Adad, J., Moller, H., Turner, R., Setsompop, K., and Wald, L. (2015). Real diffusion–weighted MRI enabling true signal averaging and increased diffusion contrast. *Neuroimage*, 122:373–384.

Ghugre, N., Enriquez, C., Coates, T., Nelson, M., and Wood, J. (2006). Improved R2* measurements in myocardial iron overload. *Journal of Magnetic Resonance Imaging*, 23:9–16.

Glover, G. (1991). Phase–Offset Multiplanar (POMP) volume imaging– a new technique. *Journal of Magnetic Resonance Imaging*, 1(1):457–461.

Goldstein, H. (1980). *Classical Mechanics*. Addison–Wesleay, Reading.

Goodman, N. (1963). Statistical analysis based on a certain multivariate complex Gaussian distribution (An introduction). *Annals Math. Statist.*, 34:152–177.

Greitz, D., Wirestam, R., Franck, A., Nordell, B., Thomsen, C., and Stahlberg, F. (1992). Pulsatile brain movement and associated hydrodynamics studied by magnetic resonance phase imaging. The Monro-Kellie doctrine revisited. *Neuroradiology*, 34(5):370–380.

Griswold, M., Jacob, P., Heidermann, R., Nittka, M., Jellus, V., Wang, J., Kiefer, B., and Haase, A. (2002). Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magnetic Resonance Imaging*, 47(6):1202–10.

Guhaniyogi, S., Chu, M., Chang, H., Song, A., and Chen, N. (2016). Motion immune diffusion imaging using augmented MUSE for high-resolution multi-shot EPI. *Magnetic Resonance in Medicine*, 75(2):639–652.

Hanson, L. (2008). Is quantum mechanics necessary for understanding magnetic resonance? *Concepts Magn Reson Part A*, 32(5):329–340.

Heidemann, R., Griswold, M., Seiberlich, N., Nittka, M., Kannengiesser, S., Kiefer, B., and Jakob, P. (2007). Fast method of 1D non-Cartesian parallel imaging using GRAPPA. *Magnetic Resonance in Medicine*, 57(6):1037–46.

Henkelman, R. (1985). Measurement of signal intensities in the presence of noise in MR images. *Medical Physics*, 12:232–233.

Henkelman, R. (1987). Artifacts in magnetic resonance imaging. *Rev. Magn. Reson. Med.*, 2(7).

Hernando, D., Kramer, J., and Reeder, S. (2013). Multipeak Fat-Corrected Complex R2* Relaxometry: Theory, Optimization, and Clinical Validation. *Magnetic Resonance Imaging*, 70:1319–1331.

Holdsworth, S., Skare, S., Newbould, R., Guzmann, R., Blevins, N., and Bammer, R. (2008). Readout-segmented EPI for rapid high resolution diffusion imaging at 3T. *Journal of Radiology*, 65(1):36–46.

Hutchinson, M. and Raff, U. (1988). Fast MRI data acquisition using multiple detectors. *Magnetic Resonance in Medicine*, 6(1):87–91.

Ianni J.D., Welch E.B., G. W. (2018). Ghost reduction in echo-planar imaging by joint reconstruction of images and line-to-line delays and phase errors. *Magnetic Resonance in Medicine*, 79(6):3114–3121.

Jaermann, T., Pruessmann, K., Valavanis, A., Kollias, S., and P., B. (2006). Influence of sense on image properties in high-resolution diffusion tensor imaging. *Magnetic Resonance in Medicine*, 55(2):335–342.

Jensen, J., Helpern, J., Ramani, A., Lu, H., and Kaczynski, K. (2005). Diffusional kurtosis imaging. the quantification of non-Gaussian water diffusion by means of magnetic resonance imaging. *Magnetic Resonance in Medicine*, 53(6):1432–1440.

Jeong, H., Gore, J., and Anderson, A. (2013). High-resolution human diffusion tensor imaging using 2-D navigated multishot SENSE EPI at 7T. *Magnetic Resonance in Medicine*, 69(3):793–802.

Jezzard, P., Barnett, A., and Pierpaoli, C. (1998). Characterization of and correction for eddy current artifacts in echo planar diffusion imaging. *Magnetic Resonance in Medicine*, 39(5):801–812.

K. Jones, D. (2011). *Diffusion MRI. Theory, Methods and Applications*. Wiley–IEEE Press, Oxford, UK.

Kalia, V., Fritz, B., Johnson, R., Gilson, W., Raithel, E., and Fritz, J. (2017). CAIPIRINHA accelerated SPACE enables 10-min isotropic 3D TSE MRI of the ankle for optimized visualization of curved and oblique ligaments and tendons. *European Radiology*, 27:3652–3661.

Kellman, P. and McVeigh, E. (2005). Image Reconstruction in SNR Units: A General Method for SNR Measurement. *Magnetic Resonance in Medicine*, 54:1439–1447.

Krissian, K. and Aja-Fernández, S. (2011). Noise-Driven anisotropic Diffusion filtering on MRI. *IEEE Transactions on Image Processing*, 65:138–145.

Larkman, D., Hajnal, J., Herlihy, A., Coutts, G., Young, I., and Ehnholm, G. (2001). Use of multicoil arrays for separation of signal from multiple slicessimultaneously excited. *Journal of Magnetic Resonance Imaging*, 13:313–317.

Larkman, D. and Nunes, R. (2007). Parallel magnetic resonance imaging. *Physics in Medicine and Biology*, 52:15–55.

Le Bihan, D., Poupon, C., Amadon, A., and Lethimonnier, F. (2006). Artifacts and pitfalls in diffusion mri. *Journal on Magnetic Resonance Imaging*, 24(3):478–488.

Liang, Z.-P. and Lauterbur, P. C. (2000). *Principles of magnetic resonance imaging: a signal processing perspective*. Wiley–IEEE Press, Piscataway, NJ, USA.

Liu, C., Bammer, R., Kim, D., and Moseley, M. (2004). Self-navigated interleaved spiral (SNAILS): application to high-resolution diffusion tensor imaging. *Magnetic Resonance in Medicine*, 52(6):1388–1396.

Liu, W., Zhao, X., Y., M., Tang, X., and Gao, J. (2016). DWI using navigated interleaved multishot EPI with realigned GRAPPA reconstruction. *Magnetic Resonance in Medicine*, 75(1):280–286.

Lobos, R., Kim, T., Hoge, W., and Haldar, J. (2018). Navigator-free EPI ghost correction with structured low-rank matrix models: New theory and methods. *Magnetic Resonance in Medicine*, 37(11):2390 – 2402.

Macovski, A. (1996). Noise in MRI. *Magnetic Resonance in Medicine*, 36(3):494–497.

Mani, M., Jacob, M., Kelley, D., and Magnotta, V. (2017). Multi-shot sensitivity-encoded diffusion data recovery using structured low-rank matrix completion (MUSSELS). *Magnetic Resonance in Medicine*, 78(2):494–507.

Mansfield, P. (2003). Snapshot MRI. *Nobel Lecture 8 Dec 2003*.

McVeigh, E., Henkelman, R., and Bronskill, M. (1985). Noise and filtration in magnetic resonance imaging. *Med. Phys.*, 12:586–591.

Meisamy, S., Hines, C., Hamilton, G., Sirlin, C., McKenzie, C., Yu, H., Brittain, J., and Reeder, S. (2011). Quantification of hepatic steatosis with T1-independent, T2-corrected MR imaging with spectral modeling of fat: blinded comparison with MR spectroscopy. *Radiology*, 258:767–775.

Miller, K. and Pauly, J. (2003). Nonlinear phase correction for navigated diffusion imaging. *Magnetic Resonance in Medicine*, 50(2):343–353.

Muller, S. (1988). Multifrequency selective RF pulses for multislice MR Imaging. *Magnetic Resonance in Medicine*, 6(1):364–371.

O'Halloran, R., Holdsworth, S., Aksoy, M., and R., B. (2013). Model for the correction of motion-induced phase errors in multishot diffusion-weighted-mri of the head: are cardiac-motion-induced phase errors reproducible from beat-to-beat? *Magnetic Resonance in Medicine*, 68(2):430–440.

Ozarslan, E., Koay, C., Shepherd, T., Komlosh, M., Irfanoglu, M., Pierpaoli, C., and Basser, P. (2013). Mean apparent propagator (MAP) MRI. A novel diffusion imaging method for mapping tissue microstructure. *Neuroimage*, 78:16–32.

Ozarslan, E. and Mareci, T. (2003). Generalized diffusion tensor imaging and analytical relationships between diffusion tensor imaging and high angular resolution diffusion imaging. *Magnetic Resonance in Medicine*, 50(5):955–965.

Pauly, J., Nishimura, D., and Macovski, A. (1989). A k-space analysis of small-tip-angle excitation. *Journal of Magnetic Resonance*, 81(1):43–56.

Pieciak, T., Rabanillo, I., and Aja-Fernández, S. (2018). Bias correction for non-stationary noise filtering in MRI. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI)*.

Pipe, J., Farthing, V., and Forbes, K. (2002). Multishot diffusion-weighted fse using PROPELLER MRI. *Magnetic Resonance in Medicine*, 47(1):42–52.

Porter, D. and Heidemann, R. (2009). High resolution diffusion-weighted imaging using readout-segmented echo-planar imaging, parallel imaging and a two-dimensional navigator-based reacquisition. *Magnetic Resonance in Medicine*, 62(2):468–475.

Prah, D., Paulson, E., Nencka, A., and Schmainda, K. (2010). A simple method for rectified noise floor suppression: phase–corrected real data reconstruction with application to Diffusion–Weighted Imaging. *Magnetic Resonance in Medicine*, 64:418–429.

Pruessmann, K. P., Weiger, M., Börnert, P., and Boesiger, P. (2001). Advances in sensitivity encoding with arbitrary k-space trajectories. *Magnetic Resonance in Medicine*, 46(4):638–651.

Pruessmann, K. P., Weiger, M., Scheidegger, M., and Boesiger, P. (1999). SENSE: Sensitivity encoding for fast MRI. *Magnetic Resonance in Medicine*, 42(5):952–62.

Rabanillo, I., Aja-Fernández, S., Alberola-López, C., and Hernando, D. (2017). Exact calculation of noise maps and g-factor in GRAPPA using a k–space analysis. *IEEE Transactions on Medical Imaging*.

Rabanillo, I., Aja-Fernández, S., Alberola-López, C., and Hernando, D. (2018a). Exact calculation of noise maps and g -factor in GRAPPA using a k–space analysis. *IEEE Transactions on Medical Imaging*, 37(2):480–490.

Rabanillo, I., Hernando, D., and Aja-Fernández, S. (2016a). Variation of noise in multi-echo MRI acquisitions using Parallel Imaging. *Proceedings of the 33rd European Society for Magnetic Resonance in Medicine and Biology Scientific Meeting (ESMRMB '16), Vienna, Austria,*, 29:400.

Rabanillo, I., Holmes, J., Guidon, A., Aja-Fernández, S., Alberola-López, C., Reeder, S., and Hernando, D. (2016b). ADC measurement accuracy in quantitative diffusion phantoms using reduced field-of-view and multi-shot acquisitions. In *Proceedings of the 33rd International Society of Magnetic Resonance in Medicine Workshop on Breaking the Barriers of Diffusion MRI*.

Rabanillo, I., Sanz-Estebanez, S., Aja-Fernández, S., Hajnal, J., Alberola-López, C., and Cordero-Grande, L. (2018b). Joint image reconstruction and phase corruption maps estimation in multi-shot echo planar imaging. In *Proceedings of the 21st International Conference on Medical Image Computing & Computer Assisted Intervention. International Workshop on Computational Diffusion MRI (MICCAI/CDMRI 2018)*.

Rabanillo, I., Zhu, A., Aja-Fernández, S., Alberola-López, C., and Hernando, D. (2018c). Computation of exact g-factor maps in 3D GRAPPA reconstructions. *Magnetic Resonance in Medicine*.

Robson, P., Grant, A., Madhuranthakam, A., Lattanzi, R., Sodickson, D., and McKenzie, C. (2008). Comprehensive quantification of signal-to-noise ratio and g-factor for image-based and k-space-based parallel imaging reconstructions. *Magnetic Resonance in Medicine*, 60:895.

Roemer, P., Edelstein, W., Hayes, C., Souza, S., and Mueller, O. (1990). The NMR phased array. *Magnetic Resonance in Medicine*, 16:192–225.

Royuela-del Val, J., Godino-Moya, A., Menchon-Lara, R., Sanz-Estebanez, S., Simmross-Wattenberg, F., Martín-Fernández, M., and Alberola-López, C. (2018). Joint non-rigid motion and image mri sense reconstruction. In *Proceedings of the 26th annual meeting of the International Society of Magnetic Resonance in Medicine*, Paris, France. (Accepted).

Rudin, L., Osher, S., and Fatemi, E. (1992). Nonlinear total variation based noise removal algorithms. *Physica D Nonlinear Phenomena*, 60(1–4):259–268.

Sacolick, L., Wiesinger, F., Hancu, I., and Vogel, M. (2010). B1 mapping by bloch-siegert shift. *Magnetic Resonance in Medicine*, 63(5):1315–1322.

Sanz-Estébanez, S., Rabanillo, I., Royuela-del Val, J., Aja-Fernández, S., and Alberola-López, C. (2018). Joint groupwise registration and ADC estimation in the liver using a b-value weighted metric. *Magnetic Resonance Imaging*, 37:1–9.

Saritas, E., Cunningham, C., Lee, J., Han, E., and Nishimura, D. (2008). DWI of the spinal cord with reduced FOV single-shot EPI. *Magnetic Resonance in Medicine*, 60(2):468–473.

Saritas, E., Lee, J., and Nishimura, D. (2011). SNR- dependence of optimal parameters for apparent diffusion coefficient measurements. *IEEE Transactions on Medical Imaging*, 2:424–437.

Segars, W., Sturgeon, G., Mendonca, S., Grimes, J., and Tsui, B. (2010). 4D XCAT phantom for multimodality imaging research. *Med. Phys.*, 37:4902–4915.

Seiberlich, N., Breuer, F., Blaimer, M., Barkauskas, K., Jakob, P., and Griswold, M. (2007). Non-cartesian data reconstruction using GRAPPA operator gridding (grog). *Magnetic Resonance in Medicine*, 58(6):1259–1265.

Seiberlich, N., Breuer, F., Blaimer, M., Deidenmann, R., Jakob, P., and Griswold, M. (2008a). Reconstruction of undersampled non-cartesian data sets using pseudo-cartesian GRAPPA in conjunction with grog. *Magnetic Resonance in Medicine*, 59(5):1127–1137.

Seiberlich, N., Breuer, F., Blaimer, M., Jakob, P., and Griswold, M. (2008b). Self-calibrating GRAPPA operator gridding for radial and spiral trajectories. *Magnetic Resonance in Medicine*, 59(4):930–935.

Seiberlich, N., Breuer, F., Heidemann, R., Blaimer, M., Griswold, M., and Jakob, P. (2008c). Reconstruction of undersampled non-cartesian data sets using pseudo-cartesian GRAPPA in conjunction with grog. *Magnetic Resonance in Medicine*, 59:1127–1137.

Setsompop, K., Gagoski, B., Polimeni, J., Witzel, T., Wedeen, V., and Wald, L. (2012). Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g-factor penalty. *Magnetic Resonance in Medicine*, 67(5):1210–1224.

Sprenger, T., Sperl, J., Fernandez, B., Haase, A., and Menzel, M. (2016). Real valued diffusion-weighted imaging using decorrelated phase filtering. *Magnetic Resonance Imaging*, 77:559–572.

Stejskal, E. and Tanner, J. (1965). Spin diffusion measurements: Spin echoes in the presence of a time-dependent field gradient. *The Journal of Chemical Physics*, 42(1):288–292.

Stevens, K., Busse, R., Han, E., Brau, A., Beatty, P., Beaulieu, C. F., and Gold, G. (2008). Ankle: Isotropic MR Imaging with 3D-FSE-Cube—Initial experience in healthy volunteers. *Radiology*, 248:1026–1033.

Suetens, P. (2002). *Fundamentals of Medical Imaging*. Cambridge University Press, Cambridge, UK.

Tannus, A. and Garwood, M. (1997). Adiabatic pulses. *NMR in Biomedicine*, 10:423–434.

Thünberg, P. and Zetterberg, P. (2007). Noise distribution in SENSE- and GRAPPA- reconstructed images: a computer simulation study. *Magnetic Resonance Imaging*, 25:1089–94.

Tuch, D. (2004). Q-ball imaging. *Magnetic Resonance in Medicine*, 52(6):1358–1372.

Uecker, M., Karaus, A., and Frahm, J. (2009). Inverse reconstruction method for segmented multishot diffusion-weighted MRI with multiple coils. *Magnetic Resonance in Medicine*, 62(5):1342–1348.

Van, A., Hernando, D., and Sutton, B. (2011). Motion-induced phase error estimation and correction in 3D diffusion tensor imaging. *IEEE Transactions on Medical Imaging*, 30(11):1933–1940.

Veraart, J., Novikov, D., D., C., Ades-Aron, B., Sijbers, J., and Fieremans, E. (2016). Denoising of diffusion mri using random matrix theory. *Neuroimage*, 142:394–406.

Veraart, J., Poot, D., Van Hecke, W., Blockx, I., Van der Linden, A., Verhoye, M., and Sijbers, J. (2011). More accurate estimation of diffusion tensor parametersusing diffusion kurtosis imaging. *Magnetic Resonance in Medicine*, 65:138–145.

Vos, S., Tax, C., Luijten, P., Ourselin, S., Leemans, A., and Froeling, M. (2016). The importance of correcting for signal drift in diffusion MRI. *Magnetic Resonance in Medicine*, 77:285–299.

Walsh, D., Gmitro, A., and Marcelling, M. (2000). Adaptive reconstruction of phased array MR imagery. *Magnetic Resonance in Medicine*, 43(5):682–690.

Wang, K., Busse, R., Holmes, J., Beatty, P., Brittain, J., Francois, C., Reeder, S., Du, J., and Korosec, F. (2011). Interleaved variable density sampling with a constrained parallel imaging reconstruction for dynamic contrast-enhanced mr angiography. *Magnetic Resonance in Medicine*, 66:428–436.

Wedeen, V., Wang, R., Schmahmann, J., Benner, T., Tseng, W., Dai, G., Pandya, D., Hagmann, P., Arceuil, H., and Crespigny, A. (2008). Diffusion spectrum magnetic resonance imaging (DSI) tractography of crossing fibers. *Neuroimage*, 41(4):1267–1277.

Weiger, M., Pruessmann, K., and Boesiger, P. (2002). 2D SENSE for faster 3D MRI. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 14:10–19.

Wiens, C., Artz, N., Jang, H., McMillan, A., Koch, K., and Reeder, S. (2018). Fully phase-encoded mri near metallic implants using ultrashort echo times and broadband excitation. *Magnetic Resonance in Medicine*, 79(4):2156–2163.

Wissmann, L., Santelli, C., Segars, W. P., and Kozerke, S. (2014). MRXCAT: Realistic numerical phantoms for cardiovascular magnetic resonance. *Journal of Cardiovascular Magnetic Resonance*, 16(1):63.

Ying, L. and Sheng, J. (2007). Joint image reconstruction and sensitivity estimation in SENSE (JSENSE). *Magnetic Resonance in Medicine*, 57(6):1196–1202.

Zaitsev, M., Maclaren, J., and Herbst, M. (2015). Motion artifacts in MRI: A complex problem with many partial solutions. *Journal on Magnetic Resonance Imaging*, 42(4):887–901.

Zhang, H., Schneider, T., Wheeler-Kingshott, C., and Alexander, D. (2012). NODDI. Practical in vivo neurite orientation dispersion and density imaging of the human brain. *Neuroimage*, 61(4):1000–1016.

Zhang, Y., Holmes, J., Rabanillo, I., Guidon, A., Wells, S., and Hernando, D. (2018). Quantitative diffusion MRI using reduced field-of-view and multi-shot acquisition techniques: Validation in phantoms and prostate imaging. *Magnetic Resonance Imaging*, 51:173–181.