Structured Low-rank Methods for Robust 3D Multi-shot EPI

Xi Chen

St Hilda's College

University of Oxford

Thesis submitted for the degree of

Doctor of Philosophy

Hilary 2023

Abstract

Magnetic resonance imaging (MRI) has inherently slow acquisition speed, and Echo-Planar Imaging (EPI), as an efficient acquisition scheme, has been widely used in functional magnetic resonance imaging (fMRI) where an image series with high temporal resolution is needed to measure neuronal activity. Recently, 3D multi-shot EPI which samples data from an entire 3D volume with repeated shots has been drawing growing interest for fMRI with its high isotropic spatial resolution, particularly at ultra-high fields. However, compared to single-shot EPI, multi-shot EPI is sensitive to any inter-shot instabilities, e.g., subject movement and even physiologically induced field fluctuations. These inter-shot inconsistencies can greatly negate the theoretical benefits of 3D multi-shot EPI over conventional 2D multi-slice acquisitions.

Structured low-rank image reconstruction which regularises under-sampled image reconstruction by exploiting the linear dependencies in MRI data has been successfully demonstrated in a variety of applications. In this thesis, a structured low-rank reconstruction method is optimised for 3D multi-shot EPI imaging together with a dedicated sampling pattern termed seg-CAIPI, in order to enhance the robustness to physiological fluctuations and improve the temporal stability of 3D multi-shot EPI for fMRI at 7T. Moreover, a motion compensated structured low-rank reconstruction framework is also presented for robust 3D multi-shot EPI which further takes into account inter-shot instabilities due to bulk motion. Lastly, this thesis also investigates into the improvement of structured low-rank reconstruction from an algorithmic perspective and presents the locally structured low-rank reconstruction scheme.

Acknowledgement

Firstly, I want to thank my supervisors, Dr. Mark Chiew and Dr. Wenchuan Wu. It is you guys that make my PhD a happy journey. I am very well supported, and you make me feel comfortable to bring it up when I need your help or when I disagree with you. If I am going to be a supervisor in the future, I hope I could do as well as you.

Secondly, I want to thank my parents, who brought me to the world and never ask anything from me. Although we do not talk about work, you have influenced me in every way, and I know I am deeply loved. I would also like to thank all my families, especially my grandma, Yingchun Hao, who is an amazing woman.

Next, I want to thank my colleagues who have helped me in any ways, giving me company, answering my questions, or just being my volunteers for scan. Thank you all for making me feel a sense of belonging in a foreign country. A special thanks to those in my cohort. Wishing you all the best in your future.

I also want to thank myself, although it seems a bit cheeky to do so. I want to take this opportunity to encourage myself. Ten years ago, you did not even know how to run a MATLAB script and you had no idea what you are able to do now, so now looking into the future, I want you to believe in yourself and do not be scared by things you have not yet known how to do.

Lastly, I want to thank whoever is reading this thesis now. You make this work more meaningful, and I hope I have shared something useful with you.

Table of Contents

List of Figures

List of Tables

List of abbreviations

1D, 2D, 3D	One-, two-, three dimensional
ACS	Auto-calibrated signals
ADMM	Alternating direction method of multipliers
BOLD	Blood Oxygen Level Dependent
CAIPI	CAIPIRINHA
CS	Compressed sensing
CT	Computed tomography
dHb	Deoxygenated haemoglobin
DOF	Degrees of freedom
DWI	Diffusion weighted imaging
EPI	Echo-Planar Imaging
ETL	Echo Train Length
FFT	Fast Fourier transform
FISTA	Fast Iterative Shrinkage-Thresholding Algorithm
fMRI	Functional magnetic resonance imaging
FOV	Field of view
GLM	General linear model
GRE	Gradient echo
HB	Oxygenated haemoglobin
IRLS	Iterative reweighted least squares
IRPF	Incremented rank power factorization
GIRAF	Generic iterative reweighted annihilation filter
ISTA	Iterative Shrinkage-Thresholding Algorithm
L+S	Low-rank plus sparsity
LM	Levenberg–Marquardt
LSLR	Locally structured low-rank
mcSLR	Motion compensated structured low-rank

MRI	Magnetic resonance imaging
NRMSE	Normalised root mean square error
PCA	Principle component analysis
RF	Radio frequency
ROI	Region of interest
std	Standard deviation
s.t.	Subject to
SAR	Specific absorption rate
SE	Spin echo
SLR	Structured low-rank
SMS	Simultaneous multi-slice
SNR	Signal-to-noise ratio
SOS	Sum-of-square
SVD	singular value decomposition
TE	Echo time
TR	Repetition time

1 Introduction

1.1 Motivation

Magnetic resonance imaging (MRI) is a medical imaging technique which enables us to see the inside structure of our body and its functionality noninvasively. Along with other imaging modalities like ultrasound and computed tomography (CT), it has been one of the most important breakthroughs that have a profound impact on modern medicine by transforming the whole stream of diagnosis, treatment and prognosis. Compared to other imaging techniques, MRI has comprehensive advantages regarding safety, contrast richness, and the ability to achieve decent spatial and temporal resolutions at the same time. However, MRI also has drawbacks and a prominent one is the inherent slow acquisition speed. The slow speed makes it susceptible to any intra-scan instabilities coming from the subject or the scanner itself. Some conspicuous factors include subject movement like respiration or cardiac pulsation which make cardiac imaging challenging. These physiological movements can even cause dynamic perturbations to the magnetic

field which lead to signal variations in brain imaging. Also, the low imaging speed limits the temporal resolution achievable if an image series is needed to measure the dynamic of physiological states like in functional magnetic resonance imaging (fMRI). Echo-Planar Imaging (EPI), as a fast imaging technique with high efficiency in data acquisition, has been widely used in fMRI, and 3D multi-shot EPI is gaining growing interest especially in high resolution fMRI recently. Compared to conventional 2D single shot EPI, 3D multi-shot EPI has higher signal-to-noise ratio (SNR) (Poser et al., 2010), which is necessary for higher temporal SNR (tSNR) and thus higher sensitivity to neuronal activities. However, multi-shot EPI is more vulnerable to intra-scan instabilities. Specifically, the aforementioned dynamic field perturbations can lead to shot-to-shot image phase variations in multi-shot EPI, which are scaled up at ultra-high fields like 7T. These inter-shot phase inconsistencies result in detrimental temporal variance of the time course, which greatly negates the SNR benefits of 3D multi-shot EPI and reduces its tSNR. Moreover, multi-shot EPI is also vulnerable to inter-shot motion, which leads to motion artefacts and blurring, and also reduces the temporal stability of 3D multi-shot EPI. Thus, the development of robust 3D multi-shot EPI which is insensitive to inter-shot phase variations and motion for a higher tSNR and sensitivity for fMRI has been the focus of this thesis. The robust 3D multi-shot EPI could be particularly useful in brainstem and spinal cord imaging which suffer from significant phase variations, and also potentially advantageous in high resolution imaging like layer-specific fMRI for improved motion robustness.

The removal of image artefacts due to under-sampling for faster acquisition requires careful design of advanced MRI image reconstruction methods, and most of the time prior knowledge about the image is needed as additional regularisations to compensate for the lack of sufficient data. Structured low-rank (SLR) methods, as an emerging field of MRI

image reconstruction, exploit the linear predictability in MRI data (Haldar and Setsompop, 2020; Jacob et al., 2020), that each data sample we acquire is linearly dependent on others, providing redundant information. The SLR reconstruction method MUSSELS (Mani et al., 2017) has been developed to deal with similar phase variations in 2D multi-shot EPI for diffusion weighted imaging (DWI), which circumvents the shot-to-shot inconsistencies by an under-sampled reconstruction for each shot regularised by the linear dependencies between them. Inspired by this, this thesis has investigated into the application of SLR reconstruction in 3D multi-shot EPI for fMRI at 7T to improve its robustness to physiologically induced phase variations. Besides, as discussed earlier, motion can be another source of instabilities which causes problems for 3D multi-shot EPI, and it also impairs the linear dependencies between MRI data of different shots which is the key to SLR reconstruction. Thus, a motion compensated structured low-rank framework has been further proposed for 3D multi-shot EPI which considers inter-shot motion and phase variations as an integrated problem. Moreover, apart from the specific scenario of 3D multi-shot EPI, this thesis has also explored the improvement of SLR method from an algorithmic perspective and proposed the locally structure low-rank approach which has a better reconstruction performance.

1.2 Outline

This thesis consists of 7 chapters. The first introduction chapter briefly depicts the motivation and outline of this thesis.

Chapter 2 gives an introduction on some basic knowledge of MRI that lays the foundation of this thesis, including the physical principles of MRI, two major image reconstruction methods, parallel imaging and compressed sensing, and functional MRI.

Chapter 3 reviews two image reconstruction methods building on the low-rank

property of MRI data: low-rank and structured low-rank reconstruction methods. The application and improvement of structured low-rank image reconstruction are the focuses of this thesis, which are discussed in the following three original research chapters.

Chapter 4 demonstrates how structured low-rank reconstruction can be used to improve the robustness of 3D multi-shot EPI to inter-shot phase variations for fMRI at 7T. The work in this chapter has been published as a full paper in NeuroImage (Chen et al., 2023) and two abstracts in International Society for Magnetic Resonance in Medicine (ISMRM) annual meeting 2020&2021 (Chen et al., 2020, 2021).

Chapter 5 presents the motion compensated structured low-rank reconstruction framework, which further improves the robustness of 3D multi-shot EPI to both inter-shot motion and phase variations. Some of the work in this chapter has been presented at the ISMRM Workshop on Motion Detection & Correction 2022 and ISMRM Workshop on Data Sampling & Image Reconstruction 2023.

Chapter 6 introduces the locally structured low-rank reconstruction approach, which achieves a robust improvement compared to the conventional structured low-rank reconstruction. Some of the work in this chapter has been published as a conference paper in IEEE 19th International Symposium on Biomedical Imaging 2022 (Chen et al., 2022) and an abstract in ISMRM annual meeting 2022 (Chen et al., 2022).

Chapter 7 concludes this thesis with a summary of contributions and discussion of future work.

2 Background

2.1 Introduction

In this chapter, some basic concepts of MRI that lay the foundation of this thesis are presented. Section 2.2 introduces some fundamental physical principles of MRI focusing on three aspects, from how MRI signals are generated, to how versatile contrasts are achieved, and then how spatial encoding is performed. Section 2.3 discusses two major MRI image reconstruction methods, parallel imaging and compressed sensing. Section 2.4 gives a brief review of fMRI and Blood Oxygen Level Dependent (BOLD) contrast. Section 2.5 concludes this chapter.

2.2 MRI Physical Principles

2.2.1 Signal Generation

The Generation of Net Magnetization

MRI is a medical imaging technique which can image the anatomical structure and function of the subject non-invasively. One major difference between MRI and some other medical imaging techniques like CT and ultrasound is the source of signals being detected. CT uses x-ray to produce exogenous signals, in which a beam of x-ray is shot through the body resulting in different amounts of x-ray attenuation for different tissues. Similarly, ultrasound imaging is based on the differentiated reflection of exogenous ultrasound signals by different tissues. In contrast, MRI uses the magnetization generated by certain nuclei, typically the hydrogen protons of water molecules which are naturally abundant in organisms as the endogenous signal sources. However, the magnetization is not generated from hydrogen protons naturally, but relies on the interactions of the nuclei with very strong external magnetic fields, typically (but not necessarily) on the scale of a few tesla (T) which is more than ten thousand times stronger than the magnetic field of the earth. Thus, one of the major hardware components of a MRI scanner is the magnet, and typically a superconducting electromagnet is used to provide 1.5T or higher magnetic fields for clinical scanners. The external magnetic field is called the main magnetic field or static magnetic field, and it is also commonly referred to as the B_0 field. The B_0 field is conventionally along the positive z direction of the laboratory frame, which is along the bore or long axis of a cylindrical MRI scanner.

Atomic and subatomic particles like protons, neutrons and electrons have an intrinsic physical property – spin. Spin is a vector quantity that possesses magnitude and direction.

The net spin of a nucleus is the complex sum of the spins of its constituent protons and neutrons. In contrast to the macroscale world, the magnitude of spin is quantized, and the spin of a nucleus only takes a limited set of discrete values $I = 0, \frac{1}{2}, 1, \frac{3}{2} \ldots$ in units of the reduced Planck constant. For example, hydrogen nuclei have $I = 1/2$. Each nucleus with non-zero spin has a magnetic moment proportional to the spin along the same or opposite direction. Each nuclear spin has $2I + 1$ spin states, so the hydrogen proton only has two spin states. Without an external magnetic field, the two spin states of hydrogen are degenerate, and all the spins and corresponding magnetic moments are pointing at random directions, so the net magnetization is zero. However, when the nuclei are placed in the main magnetic field, they will be aligned by the external magnetic field and reveal two states, "parallel" and "antiparallel". These two states have an energy difference that is proportional to the field strength according to Zeeman Splitting:

$$\Delta E = \gamma h B_0 \qquad \text{Eq. 2.1}$$

Where γ is the gyromagnetic ratio and it is an intrinsic constant of certain nuclei. h is the Planck constant. The parallel nuclei have a lower energy while the anti-parallel nuclei have a higher energy. The numbers of protons of these two energy states follows the Boltzmann distribution:

$$\frac{N^-}{N^+} = e^{-\frac{\Delta E}{kT}} \qquad \text{Eq. 2.2}$$

Where N^+ and N^- are the numbers of nuclei in parallel and anti-parallel states. k is Boltzmann constant, and T is the absolute temperature in Kelvin. The number of parallel nuclei is slightly higher than the number of anti-parallel nuclei, which leads to a net magnetization along the field direction in the macroscale. This net magnetization is termed M_0. Fig. 2.1 shows a schematic diagram illustrating the formation of the net magnetization M_0. Although this difference between N^- and N^+ is very subtle, thanks

to the high abundancy of hydrogen nuclei in human tissues, a non-trivial M_0 is still available as the source of MRI signals. Since ΔE is proportional to the field strength B_0, a stronger magnetic field is desired to achieve a bigger difference between N^- and N^+. More specifically, the net magnetization M_0 is proved to be proportional to B_0, and thus a higher B_0 can provide a stronger signal for imaging. Note although the above classic explanation of the macroscale magnetization is not exactly accurate, the behavior of a macroscale magnetization can be analyzed according to the laws of classical electromagnetism.

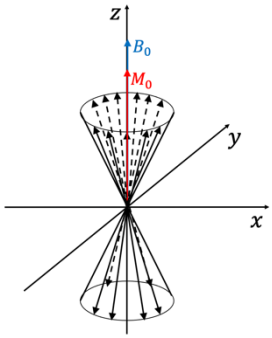

Fig. 2.1 Schematic generation of the net magnetization M_0. The magnetic field is along z direction. Each vector represents a magnetic moment corresponding to a nuclear spin. The projection of each magnetic moment in x-y plane has the equal probability to be at each random direction, so the net magnetization in $x-y$ plane is zero. The projection of each magnetic moment along z direction is more likely to be in the positive direction than the negative direction, creating a non-zero net magnetization along positive z direction. This macroscale magnetization M_0 pointing along the B_0 field direction is the source of MRI signal.

All the magnetic moments not in line with the main magnetic field precess about the field direction like a gyroscope as a consequence of the interaction between the spin and the main magnetic field, which is called Larmor precession. The Larmor precession frequency ω is proportional to the field strength:

$$\omega = \gamma B_0 \qquad \text{Eq. 2.3}$$

Eq. 2.3 is called the Larmor equation. Accordingly, any macroscale magnetizations not in line with the main magnetization field also precess about B_0 at the Larmor

frequency.

The Generation of Transverse Magnetization

The magnetization M_0 at equilibrium is along the z direction parallel to the main magnetic field, which makes it difficult to be detected. To deal with this, M_0 is tilted to be out of alignment with the B_0 field to generate a projection in the $x-y$ plane. As mentioned previously, the hydrogen nuclei have two spin states. The parallel spins have lower energy, and the anti-parallel ones have higher energy. The nuclei in the lower energy state can experience energy transition and end up in the higher energy state if they absorb a photon whose energy is equal to the energy difference between the two states. This phenomenon is called magnetic resonance. The energy of a photon is proportional to its frequency:

$$\Delta E = h\upsilon \qquad \text{Eq. 2.4}$$

Where υ is the frequency of the photon. Compared Eq. 2.4 with Eq. 2.1, we can get:

$$\upsilon = \gamma B_0 \qquad \text{Eq. 2.5}$$

Which shows that energy transition and magnetic resonance happen if the frequency of the external electromagnetic radiation is the same as the Larmor frequency. This frequency is typically in the range of radio waves. To excite the nuclei, an external electromagnetic wave is applied for a very short period, which is typically referred to as the RF pulse. The RF pulse is generated by RF transmitter coils, which are another necessary hardware component of an MRI scanner. The magnetic field produced by the RF pulse is called the B_1 field, and its direction is perpendicular to the B_0 field. The B_1 field generated by RF pulse rotates the net magnetization M_0 from its equilibrium position (the positive z direction) about the B_1 field direction in the macroscale. This

process is called RF excitation. The rotation angle of M_0 is called the flip angle in MRI[1] and it is determined by the strength of the RF pulse and its duration. The 90° and 180° RF pulses are often used, which rotate M_0 by 90° or 180° so it ends up in the $x-y$ plane or along the negative z direction. The projection of the net magnetization in the $x-y$ plane is called the transverse magnetization, denoted as $M_{x,y}$, and the projection in the z direction is called the longitudinal magnetization, denoted as M_z. The transverse magnetization $M_{x,y}$ can be detected by a set of RF receiver coils placed close to the subject, which is the measurable MRI signal. Details about RF receiver coils are discussed in Section 2.3.1.

2.2.2 Contrast

The intensity of the MRI signal must be associated with certain characteristics of the tissue to provide meaningful contrasts. The contrast in CT comes from the different attenuation coefficients between different tissues. In comparison, MRI can provide versatile contrasts not only between different tissues, but also within the same tissue, revealing different characteristics that are related to the physiological or pathological status of the tissue. A natural contrast of MRI is the proton density. Since MRI signals are contributed by the magnetic moments of hydrogen protons, the signal intensity is proportional to the proton density which varies in different tissues. Besides proton density, the relaxation properties of the tissue, like the intrinsic T_1 relaxation, T_2 or T_2^* relaxation, are also important sources of MRI contrast. Relaxation describes the process of excited nuclei releasing energy and returning to their equilibrium. As soon as excitation starts,

[1] Note from the perspective of quantum mechanics, each proton is a weighted superposition of two states simultaneously, and the proton does not only jump between two energy states. The classic explanation in MRI uses these inaccurate concepts, which are generally not very harmful and help the understanding of the complex mechanism. It does have some problems, e.g., explaining RF pulse with flexible flip angles.

relaxation starts too.

The T_1 relaxation is also called spin-lattice relaxation, which describes the recovery of the longitudinal magnetization M_z. It is due to the interaction between the spin and its surrounding lattice that the nucleus emits energy and returns to equilibrium. Assuming a 90° RF pulse is used to flip M_0 from its equilibrium to the transverse plane, the T_1 relaxation process can be described as:

$$M_z(t) = M_0(1 - e^{-t/T_1}) \quad \text{Eq. 2.6}$$

Where T_1 is an intrinsic time constant of the tissue that describes how fast the longitudinal magnetization recovers. When $t = T_1$, M_z has recovered ~63% of M_0. When $t = 5T_1$, M_z has recovered ~99% of M_0. Fig. 2.2a shows the process of T_1 relaxation. Different tissues have different T_1 values, and therefore their longitudinal magnetizations recover at different rates. If the next RF pulse is applied before the longitudinal magnetizations have fully recovered, a contrast in the transverse magnetizations can be established between tissues with different T_1 values.

The T_2 relaxation is also called spin-spin relaxation, which describes the decay of the transverse magnetization $M_{x,y}$. It is due to the natural interactions of different spins. The magnetic moments will gradually dephase due to local field perturbances, which results in a loss of the net transverse magnetization. T_2 relaxation also accompanies T_1 relaxation which changes the transverse component of each magnetic moment. Assuming a 90° RF pulse is used to flip M_0 from its equilibrium to the transverse plane, the T_2 relaxation process can be described by:

$$M_{x,y}(t) = M_0 e^{-t/T_2} \quad \text{Eq. 2.7}$$

Where T_2 is an intrinsic time constant of the tissue that describes how fast the transverse magnetization decays. When $t = T_2$, $M_{x,y}$ is only ~37% of M_0. Fig. 2.2b

shows the process of T_2 relaxation. Different tissues have different T_2 values, and their transverse magnetizations decay at different rates.

In practice, the transverse magnetization decays at a faster speed than the prediction of T_2 relaxation. It is due to the inhomogeneities in the main magnetic field resulting from its intrinsic imperfections and the tissue susceptibility induced field distortions, which speeds up the dephasing process and leads to a more rapid decay of the transverse magnetization. This effect is called T_2^* relaxation, which can be considered as the effective T_2 relaxation in practice.

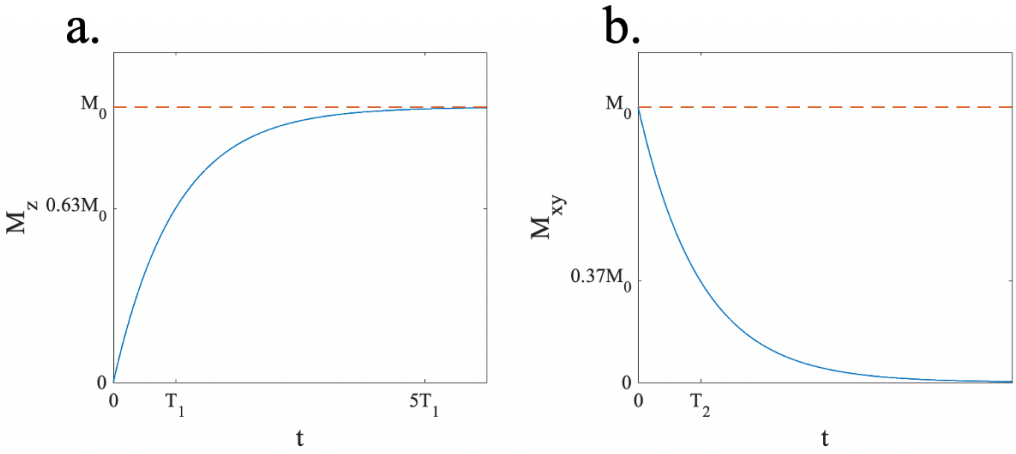

Fig. 2.2 The process of a) T_1 relaxation and b) T_2 relaxation after a 90° RF pulse.

The MRI signal intensity of tissue is determined by its proton density, T_1 and T_2 or T_2^* relaxations simultaneously. Different contrasts can be achieved by changing the weights on these tissue characteristics. An image with the contrast dominated by T_1 relaxation is called a T_1 weighted image. Similarly, an image with the contrast dominated by T_2 or T_2^* relaxation is called a T_2 weighted or T_2^* weighted image. Different weightings can be achieved by manipulating the timings of RF pulses and signal acquisitions. The time between the centre of the RF pulse and the acquisition window is termed echo time (TE). The period of a series of repeating RF pulses is termed repetition time (TR). For example, a T_1 weighted image can be achieved with a short TE and short

TR. The short TE is intended to minimise the contrast from T_2 or T_2^* decay. The short TR leads to incomplete T_1 recovery so that the maximum longitudinal magnetizations available for the transverse components following a RF excitation are different for different tissues. We can also feature the T_2 or T_2^* contrast with a relatively long TE and long TR. The long TR makes all the longitudinal magnetizations mostly recovered to suppress the T_1 contrast. The long TE enhances the signal intensity differences due to the varying T_2 or T_2^* decay rates. Note since T_2^* is sensitive to local field inhomogeneities, T_2^* weighted imaging can be used to detect changes in physiological states which lead to fluctuations in local magnetic field. It is the basis of BOLD contrast typically used in fMRI studies, and this is discussed in detail in Section 2.3.

Despite the same requirements of long TE and long TR, there are some other considerations in order to differentiate T_2 and T_2^* weighted imaging. T_2 weighted imaging requires the signal loss due to local field inhomogeneities to be compensated. This is achieved by a spin echo (SE), one of the most important developments in MRI. Fig. 2.3 shows a schematic illustration of how a spin echo is formed. The key to it is the use of two RF pulses. The first RF pulse rotates the longitudinal magnetization by 90° into the $x - y$ plane, and then the transverse magnetizations within a voxel dephase due to local field inhomogeneities, which speeds up the signal decay on top of the T_2 relaxation. The second RF pulse rotates all the magnetizations 180° about the B_1 field. The consequence is that the magnetizations which have higher Larmor frequencies are now placed behind those with lower Larmor frequencies. Thus, they "catch up" to the slower magnetizations and the net MRI signal increases due to rephase. The maximum signal intensity is achieved when phase coherence is re-established after the rephasing process experiences the same amount of time as the dephaing process, which is the interval between the first 90° RF pulse and the second 180° RF pulse. The increasing

signal evolution forms the first half of the spin echo, corresponding to the rephasing process. After reaching the peak amplitude, the MRI signal will diminish as the magnetizations start dephasing again. This decreasing signal evolution forms the second half of the spin echo. Note only the signal loss due to field inhomogeneities corresponding to T_2^* decay can be restored by spin echo, while the maximum signal at the echo peak is still subject to the intrinsic T_2 decay.

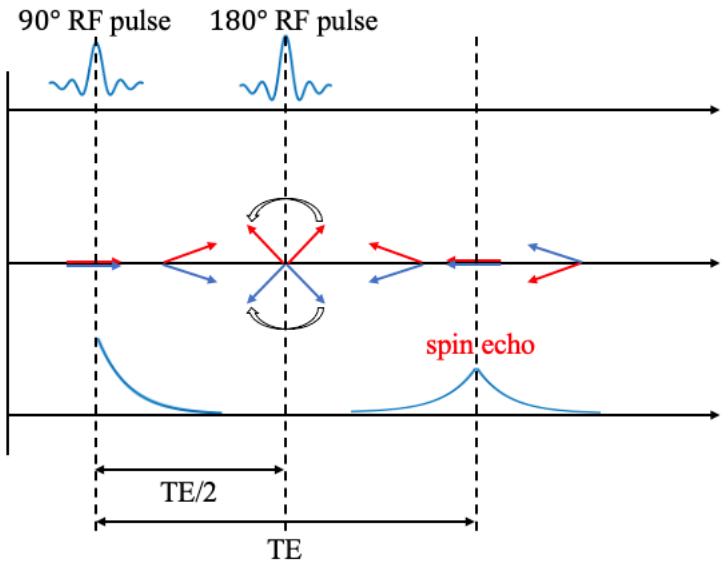

Fig. 2.3 The formation of a spin echo. The red and blue arrows indicate two representative magnetizations in the $x-y$ plane with different precession frequencies. They are in phase by the time the 90° RF pulse ends, and then gradually dephase which leads to a signal decay. The second RF pulse rotates each magnetization by 180° about the B_1 field so it ends up in a symmetrical position with reference to the B_1 field direction. After that, the rephasing stage starts which corresponds to the increasing half of the spin echo. The maximum signal intensity is at the echo centre where these two magnetizations achieve phase coherence again. From there on, the dephasing stage starts again which corresponds to the decreasing half of the spin echo.

In contrast, the counterpart of the spin echo is the gradient echo (GRE) which is used in T_2^* weighted imaging. The only difference between the gradient echo and spin echo is that a single RF pulse is used to form a gradient echo, and can also be defined through its lack of spin echo. The gradient echo has similar signal evolution in the form of an echo

with the peak value in the middle. The corresponding rephasing and dephasing stages of gradient echo are caused by the use of an external magnetic field with opposite polarities, that the phase evolution caused by the external field is reversed by changing its polarity at the beginning of a gradient echo.

2.2.3 Spatial Encoding

Slice Selection

Since magnetic resonance and the generation of transverse magnetization only happens when the frequency of a RF pulse matches the Larmor frequency, which is determined by the external field strength, it is possible to only selectively excite nuclei in our region of interest by manipulating the fulfilment of this condition. The frequency band of a RF pulse can be modulated easily. Also, it is not difficult to achieve a spatially varying magnetic field. The main magnetic field which is responsible for the generation of the net magnetization M_0 is static and homogeneous, but we can change the net magnetic field by having control of another spatially varying magnetic field. This is achieved by three sets of gradient coils, which can be switched on and off independently to create a magnetic field varying linearly along the x, y and z directions respectively. The spatially varying field is called a gradient field featuring a constant gradient across space. The strength of gradient field is much lower than the main magnetic field, and its gradient is on the scale of mT/m. Note the strength of gradient field varies spatially, but its direction is the same as the static field.

The gradient field is responsible for slice-selective excitation, i.e., the excitation is spatially restricted to a certain thin slice for 2D imaging or a thick slab for 3D imaging by applying a gradient field and a RF pulse with matched frequency band simultaneously, as shown in Fig. 2.4. We can select a transverse/sagittal/coronal slice by using a gradient

field along either $x/y/z$ direction, and an oblique slice in any direction by designing a composite gradient field.

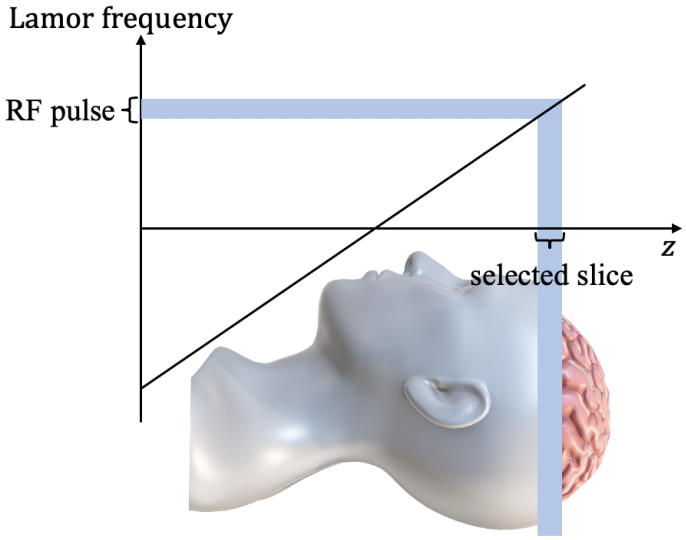

Fig. 2.4 Illustration of slice-selective excitation.

Frequency and Phase Encoding

Although slice-selective excitation has restricted the generation of transverse magnetizations to be only within the slice we are interested in, the RF receiver coils which measure the transverse magnetization are not able to encode its spatial location, and can only measure the integral of all transverse magnetizations across the entire excited volume. To obtain a useful image depicting their spatial distribution correctly, an additional step is needed to encode the spatial information of the transverse magnetization. Spatial encoding in MRI is performed by modifying the transverse magnetization with a location-dependent phase term, which is also achieved by the spatially varying gradient field. As magnetizations at different locations have different precession frequencies in the presence of a gradient field, a spatially varying phase modulation can be established after a period of time. The phase modulation of a magnetization can be represented as a function of the magnetic field deviation through Larmor equation:

$$\Delta\theta(r,t) = \int \Delta w(r,t)dt = \int \gamma \Delta B_0(r,t)\, dt \qquad \text{Eq. 2.8}$$

Where $\Delta\theta(r,t)$ is the phase modulation of the magnetization at location r and time point t. $\Delta B_0(r,t)$ is the spatially and temporally varying net magnetic field, which is generated by a gradient field and can be represented as:

$$\Delta B_0(r,t) \ = G(t)r \qquad \text{Eq. 2.9}$$

Where $G(t)$ denotes the field gradient. The phase modulation can then be represented as a function of the gradient:

$$\Delta\theta(r,t) = \int \gamma G(t)r\, dt \qquad \text{Eq. 2.10}$$

The phase modulation in the complex exponential form is shown as:

$$\varphi(r,t) = e^{-j\Delta\theta(r,t)} = e^{-j\int \gamma G(t)r\, dt} \qquad \text{Eq. 2.11}$$

If we define:

$$k = \frac{\gamma}{2\pi}\int G(t)\, dt \qquad \text{Eq. 2.12}$$

then $\varphi(r,t)$ can be rewritten as:

$$\varphi(r,t) = e^{-j2\pi kr} \qquad \text{Eq. 2.13}$$

It is clear that $e^{-j2\pi kr}$ are the basis functions of the Fourier transform with k denoting the spatial frequency, and then the integral of transverse magnetizations with phase modulations can be expressed as a function of k as:

$$F(k) = \int M_{x,y}(r)e^{-j2\pi kr}\, dr \qquad \text{Eq. 2.14}$$

Where $F(k)$ is now encoded as the Fourier transform coefficient of $M_{x,y}(r)$. The space where the Fourier coefficients $F(k)$ reside in is called k-space, and the coordinate of k-space is the spatial frequency k. It is shown that with the help of gradient field encoding, each integral signal we measure can be mapped to a sample in k-space, and its coordinate is determined by the temporal integral of the gradient $G(t)$.

In order to reconstruct the continuous magnetization distribution from its k-space measurements, technically we need to acquire an infinitely large k-space which is impractical. In practice, we reconstruct a discrete image and the maximum k-space coordinate k_{max} is determined according to the desired spatial resolution Δr following the rule:

$$2k_{max} \propto \frac{1}{\Delta r} \quad \text{Eq. 2.15}$$

The k-space is sampled discretely, and without loss of generality, we consider k-space samples on the Cartesian grid for simplicity. Note the k-space can also be sampled on non-Cartesian grid. According to the Nyquist criterion, the interval of the discrete k-space sampling grid Δk needs to satisfy the following relationship in order to not cause aliasing in the reconstructed imaging:

$$\Delta k \propto \frac{1}{FOV} \quad \text{Eq. 2.16}$$

Since only a single k-space sample can be acquired at a time, the sampling trajectory in k-space needs to be specified. The most common Cartesian sampling trajectory is shown in Fig. 2.5a and the corresponding gradient field is shown in Fig. 2.5b. Following each excitation, a readout samples a series of k-space measurements at the same k_y locations from $-k_{x\,max}$ to $+k_{x\,max}$ sequentially. Note since the readout is fast, the signal decay during this process is typically ignored. Different k_y locations are sampled by multiple readouts. As discussed above, the target k_y coordinate is achieved by applying the G_y gradient field for a certain amount of time, which is switched off before the readout starts so that k_y is fixed during the readout. Typically, different k_y lines are acquired by changing the amplitude of the G_y gradient while keeping the duration the same. This is called phase encoding as the magnetizations at different y locations have been assigned to different phase modulations. Similarly, the target k_x coordinate is

achieved by manipulating the G_x gradient field. Before each readout starts, a negative G_x gradient field is also switched on for a period of time to achieve the desired starting location $-k_{x\,max}$, and then it is kept on during the entire readout so that different k_x coordinates from $-k_{x\,max}$ to $k_{x\,max}$ are traversed as the duration of G_x gradient field increases. This is called frequency encoding as the magnetizations at different x locations are precessing at different frequencies during the readout. Note frequency encoding and phase encoding are intrinsically the same approach as phase modulation is achieved by modifying the precession frequency, and the biggest difference is that frequency encoding direction is sampled continuously in each readout as the inner loop of the k-space acquisition, whereas phase encoding direction is sampled discontinuously in this case by multiple readouts as the outer loop.

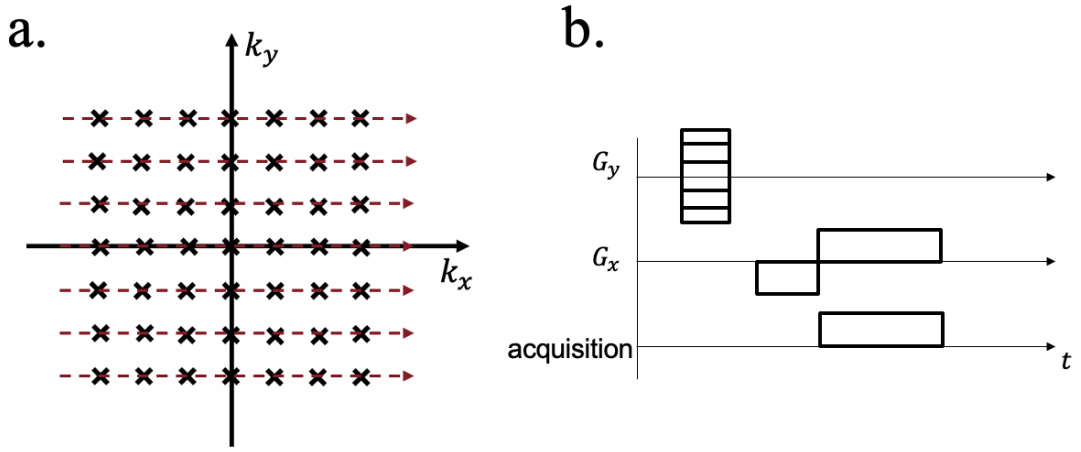

Fig. 2.5 A 2D Cartesian k-space trajectory a) and its corresponding gradient field b). In subfigure a), each cross represents a k-space sample and the samples connected by a dashed line are acquired sequentially in a single excitation.

The above-mentioned trajectory acquires one readout line in k-space per shot. However, we can also acquire multiple k_y lines per shot so that the whole k-space is acquired in a single shot. One of such efficient acquisition methods is called the Echo-Planar Imaging (EPI). A typical EPI readout samples k-space in a "zig zag" way, which is achieved by alternating the polarity of the frequency encoding gradient and inserting a

phase encoding gradient blip between reverse frequency encoding gradients, as shown in Fig. 2.6. The number of k-space lines acquired per shot is called Echo Train Length (ETL). Despite the high sampling efficiency, the long echo train can also lead to various image artefacts. The signal evolution due to T_2^* decay along the phase encoding direction, which is sampled more slowly, can act as a k-space filter and lead to blurring artefacts. Besides, the phase accumulation along the phase encoding direction due to the local field inhomogeneities can result in image distortion artefacts (Jezzard and Clare, 1999). As a consequence, multi-shot EPI which acquires the k-space by a few shots with a reduced ETL can be used to balance sampling efficiency and the above-mentioned artefacts. Due to its fast acquisition speed and robustness against motion, EPI acquisition is widely used in neuroimaging, such as fMRI and DWI.

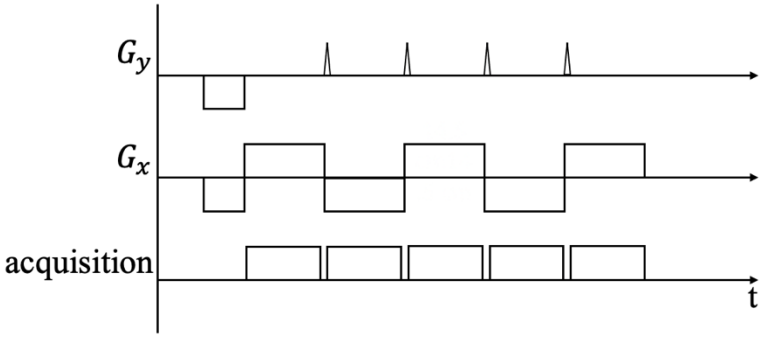

Fig. 2.6 Bipolar EPI acquisition.

After we obtain the fully sampled k-space, the image can be easily reconstructed by performing an inverse Fourier Transform. The fast Fourier Transform (FFT) (Cooley and Tukey, 1965) with high computational speed can be used on the discrete k-space on the Cartesian grid.

2.3 Advanced Image Reconstruction

Although MRI image reconstruction by a one-step inverse FFT operation is simple and fast when the Nyquist criterion is met, the acquisition of a fully sampled k-space is usually time-consuming. Imaging speed has always been a bottleneck of MRI. One of the most

important reasons why MRI is inherently slow is that only one k-space sample can be acquired at a time and thus acquisition time increases as spatial resolution increases. The imaging speed can be improved by collecting the k-space samples faster, e.g., by using a higher bandwidth, but this is limited by the hardware capability as well as SNR. Another method that can accelerate MRI acquisition significantly is under-sampling, i.e., acquiring fewer k-space samples than typically required by the Nyquist sampling criterion. More specifically, imaging acceleration can be achieved by skipping over some phase encoding steps (under-sampling along the frequency encoding direction is not typically helpful), and the reduction of acquisition time is proportional to the number of phase encoding steps skipped assuming each shot acquires a single phase encoding line. The acceleration factor R is defined as the ratio between the total number of samples required by Nyquist criterion and the number of acquired k-space samples. However, according to the relationship in Eq. 2.16, the violation of the Nyquist criterion leads to aliasing artefacts as the effective field of view (FOV) is reduced, as shown in Fig. 2.7. Thus, image reconstruction methods with increased complexity are needed to compensate for the deficiency in k-space data. Parallel imaging (Griswold et al., 2002; Pruessmann et al., 1999) and compressed sensing (Lustig et al., 2007) are two of the most popular under-sampled reconstruction methods. They have greatly promoted the application of MRI technique and are also clinically available nowadays. The following two sections give a brief review of parallel imaging and compressed sensing image reconstruction methods.

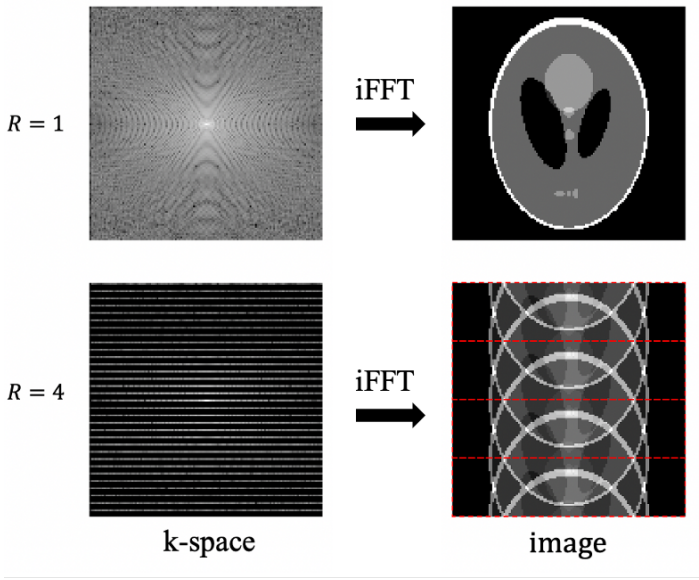

Fig. 2.7 Aliasing artefacts due to k-space under-sampling. A $R = 4$ acceleration factor results in an $1/4$ effective FOV.

2.3.1 Parallel Imaging

Phased array coils are widely used on modern MRI scanners, which consist of an array of multiple receiver coils (Roemer et al., 1990). Each element of phased array coils is a small loop that has high sensitivity to RF signals over a limited spatial coverage. Thus, a large multi-channel array is able to provide a high SNR as well as a large FOV, whereas the single-channel body coil has a large measurement field but low SNR. Although initially designed for the SNR benefit, the use of phased array coils has also inspired and enabled parallel imaging reconstruction methods, which exploits the data redundancy provided by multiple channels. Each individual coil has a spatially varying sensitivity due to its proximity to a local area. Different coils are distributed across space, which leads to differences in their coil sensitivities. Thus, an additional spatial encoding of MRI signals complementary to the gradient field encoding can be achieved by the use of phased array coils, which makes it possible to reconstruct an image from under-sampled k-space data. Depending on how the coil sensitivity encoding is explored, parallel imaging methods can be divided into two categories. One category is image based

methods represented by SENSE (Pruessmann et al., 1999). Another is k-space based methods with GRAPPA (Griswold et al., 2002) being one of the most popular choices. The method SENSE and GRAPPA are discussed in the following sections.

SENSE

The spatial encoding of phased array coils is decoded by SENSE based on the explicit knowledge of coil sensitivity maps. Note the coil sensitivities are subject-dependent, so they need to be acquired for each scan. Since coil sensitivities are supposed to be spatially smooth and contrast-independent, a region covering the centre of k-space is typically acquired by a quick pre-scan or as part of the imaging data to estimate coil sensitivity maps. After we get the image of each coil by inverse Fourier transform, the sensitivity maps can be calculated by:

$$S_i(x,y) = \frac{x_i(x,y)}{\sqrt{\sum_{i=1}^{N_c} x_i(x,y)x_i^*(x,y)}} \qquad \text{Eq. 2.17}$$

Where $x_i(x,y)$ is the low resolution image of the i_{th} coil, and N_c is the total number of coils. The sum-of-square (SOS) combination of image across coils, which is obtained by adding up the absolute square of each complex image and then getting its squared root is used to approximate the image with uniform spatial sensitivity. The coil sensitivity map S_i can be smoothed to further improve its accuracy. Another approach which is widely used nowadays is ESPIRiT (Uecker et al., 2014), which is a parallel imaging reconstruction method similar to GRAPPA. The ESPIRiT method estimates coil sensitivity maps by exploiting the linear dependences in k-space data of different coils.

The SENSE reconstruction is illustrated with a simple example where 2D Cartesian sampling with $R = 2$ along phase encoding direction is used. Since the sampling interval along the phase encoding direction k_y is 2 times larger, the FOV along y direction is halved so that the imaging object is replicated every $\mathrm{FOV}/2$. This leads to aliasing

artefacts if the size of the object is larger than FOV/2. As shown in Fig. 2.8, the pixel y in the aliased image is the superposition of two pixels x_1 and x_2 which are FOV/2 apart in the original unaliased image.

full FOV 1/2 FOV

Fig. 2.8 The origin of aliasing artefacts

What SENSE does is essentially unfolding the aliasing by differentiating x_1 and x_2 from their sum y. It is an underdetermined problem to solve for two unknowns with only one equation. However, the use of phased array coils allows us to solve for these two variables with a set of equations, since the aliased voxel in each coil is a superposition of unaliased voxels weighted by the coil sensitivity of that coil. Assuming there are two coils, this can be formulated as:

$$\begin{aligned} s_{1,1}x_1 + s_{1,2}x_2 &= y_1 \\ s_{2,1}x_1 + s_{2,2}x_2 &= y_2 \end{aligned} \qquad \text{Eq. 2.18}$$

Where $s_{m,n}$ denotes the coil sensitivity of the n_{th} pixel in the m_{th} coil and y_m is the aliased voxel of the m_{th} coil. Thus, the aliased voxels can be unfolded by solving the linear equations above. This unfolding process is repeated for each set of aliased voxels to get the unaliased full-FOV image. Note the number of linearly independent coils sets an upper limit for the acceleration factor.

The SENSE reconstruction can be implemented using a more general formula, cg-SENSE (Pruessmann et al., 2001). It is particularly useful for non-cartesian sampling where the aliasing pattern is much more complex so that it is not practical to form an

explicit matrix representation characterizing the aliasing. The method cg-SENSE considers sensitivity encoding and gradient field encoding together in the signal forward model, and the reconstruction is formulated as the following optimisation which enforces data consistency with the acquired data under the composite encoding model:

$$min\,||AFSX - Y||_2^2 \qquad \text{Eq. 2.19}$$

Where X of size $N^2 \times 1$ corresponds to the image of size $N \times N$ to be reconstructed. S is the sensitivity encoding matrix of size $N_c N^2 \times N^2$, and N_c is the number of coils. F is the Fourier encoding matrix of size $N_c N^2 \times N_c N^2$. A is the sampling matrix of size $N_c N_k \times N_c N^2$, and N_k is the number of acquired k-space samples in each coil. It is clear that with sensitivity encoding, the size of the overall encoding matrix AFS is now N_c times larger. The solution of Eq. 2.19 can be obtained by solving the following equation:

$$E^H EX - E^H Y = 0 \qquad \text{Eq. 2.20}$$

Where $E = AFS$. To achieve high efficiency, Eq. 2.20 can be solved by the conjugate gradient algorithm (Hestenes and Stiefel, 1952) with a relatively small number of iterations. Also, instead of large-scale matrix-vector multiplication, the product with Fourier encoding matrix F can be implemented by FFT operator to improve efficiency in the case of Cartesian sampling.

GRAPPA

Unlike SENSE, the GRAPPA method works primarily in k-space. While SENSE reconstruction aims to reconstruct a single unaliased image, GRAPPA aims to recover the missing k-space data of each coil. More specifically, a missing k-space sample in one coil is recovered as a weighted sum of acquired samples in its neighbourhood from all coils. Note different sets of weights are needed for reconstructing missing samples in different coils. Fig. 2.9 shows the GRAPPA method schematically. In this case, the

acquired samples within a 3 × 3 kernel around the missing target sample from all coils are linearly combined to generate an estimation of the target sample. In practice, a relatively small kernel size is often used to prevent overfitting and improve the noise robustness.

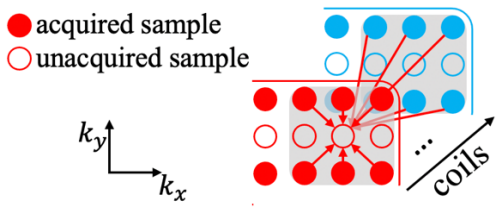

Fig. 2.9 Schematic GRAPPA reconstruction. Solid circles denote acquired samples and hollow circles denote unacquired samples. Coils are distinguished by different colours. The red lines with arrow indicate a set of weights corresponding to each acquired samples in the neighbourhood from all coils that are used to generate the missing sample.

Although GRAPPA does not rely on the explicit knowledge of coil sensitivities, it still needs a fully sampled subset of the central k-space data to calculate the GRAPPA weights, which is called auto-calibrated signals (ACS). The weights are calibrated by enforcing the consistency between the weighted sum of selected neighbouring points across all coils and the acquired sample in the ACS region. After recovering the missing k-space samples with weights learned from ACS data, the unaliased images of all coils are combined using approaches like SOS or adaptive combination (Walsh et al., 2000).

The SPIRiT method (Lustig and Pauly, 2010) which came out later exploits a similar idea with a more general reconstruction framework. There are two major differences between SPIRiT and GRAPPA. Firstly, SPIRiT enforces the consistency between every k-space point and all its neighbours across coils including both acquired and estimated data regardless of the sampling pattern, and this calibration consistency serves as a constraint to regularise the reconstruction. Secondly, the reconstruction is formulated as an optimisation consisting of both the data consistency constraint and calibration consistency constraint.

SNR penalty

The increase in imaging speed achieved by data under-sampling comes with a $\sqrt{R}$ SNR penalty due to reduced acquisition time. Besides this universal $\sqrt{R}$ penalty, parallel imaging reconstruction can also lead to noise amplification as the reconstruction is typically ill-conditioned. Note this noise amplification is non-uniform across the image. Intuitively, if a voxel does not suffer from aliasing (as shown in some of the middle rows in Fig. 2.8), there will be no additional noise amplification. This spatially varying noise amplification is characterised by the g-factor which is defined as:

$$g = \frac{SNR_{full}}{\sqrt{R}SNR_{under}} = \frac{\sigma_{under}}{\sqrt{R}\sigma_{full}} \quad \text{Eq. 2.21}$$

Where SNR_{full} is SNR of the image by fully sampled reconstruction and SNR_{under} is SNR of the image by under-sampled reconstruction. σ_{under} is the standard deviation of the image noise by under-sampled reconstruction and σ_{full} is the standard deviation of the image noise by fully-sampled reconstruction. The g-factor is a useful metric to evaluate and compare the performance of different parallel imaging reconstruction methods. It can be calculated analytically if the image is reconstructed by an analytical solution as in SENSE or by Monte Carlo simulation experiment (Pruessmann et al., 1999; Robson et al., 2008) which is more computationally efficient when an iterative solution is needed as in cg-SENSE.

2.3.2 Compressed Sensing

Compressed sensing (CS) is a general signal processing technique which offers a mathematical framework to reconstruct signals from fewer measurements than conventional sampling theory requires. The application of CS in MRI image reconstruction is one of the most impactful milestones that has greatly accelerated MRI.

In parallel imaging, the data redundancy offered by multi-channel receiver coils is exploited, while CS exploits the intrinsic data redundancy in MRI images. CS can be applied on each coil independently, or used in conjunction with SENSE to reconstruct a coil-combined image as defined in Eq. 2.19 with known coil sensitivities. Without loss of generality, here we consider the reconstruction of a single channel image leaving out the sensitivity encoding for simplicity. In order to perform successful image reconstruction from only sparsely sampled k-space, CS has three key requirements: sparsity, incoherence and non-linear reconstruction. Sparsity defines a necessary property of the image to be reconstructed and incoherence puts a restriction on the sampling. If both these two requirements are satisfied, a non-linear reconstruction can be used to reconstruct an image from fewer k-space samples than Nyquist criterion requires. These three aspects are discussed in detail in the following paragraphs.

Sparsity indicates that the image of interest is required to be sparse in a known transform domain, which means only part of the transform coefficients are non-zero. The sparsifying transform can just be identity transform if the image is already sparse in the original image space, and an example of this is MRI angiography, where the contrast is chosen to only feature vessels which are sparsely located across the image. Allowing sparsity in the transform domain makes the class of potentially applicable images much wider. In fact, most natural images satisfy this requirement, and this is indeed the core of many image compression tools, e.g., JPEG which is based on discrete cosine transform. From the perspective of image compression, transform sparsity indicates this image is compressible that we can only store few significant coefficients and restore a full-size image from them with little or no visual loss of information. In CS, spatial finite differences and the wavelet transform are two of the most widely used sparse transforms. Sparsity clearly indicates that the image has redundant information, and its degrees of

freedom (DOF) is not as high as its number of voxels, which makes it possible to reconstruct an aliasing-free image from only a subset of k-space samples.

Incoherence refers to the fact that the aliasing artefacts in the transform domain due to k-space under-sampling are required to be incoherent. The importance of incoherent sampling can be demonstrated intuitively by a simple example of 1D signal recovery as shown in Fig. 2.10. The original sparse signals and its k-space representations are shown in Fig. 2.10a and Fig. 2.10b. Two k-space under-sampling patterns with the same acceleration factor $R = 4$ are compared. The uniform under-sampling (Fig. 2.10d) leads to coherent aliasing artefacts (Fig. 2.10c), making it impossible to differentiate the true signals from their replicas. In contrast, the random under-sampling (Fig. 2.10f) results in incoherent, noise-like aliasing artefacts (Fig. 2.10e). The two strong signal components stand out from the low strength aliasing artefacts, which can be easily detected by a simple thresholding. Considering physical constraints on the gradients and the non-uniform energy distribution in k-space, variable density pseudo random under-sampling along the phase encoding directions is a sensible way to introduce incoherent aliasing in Cartesian sampling. Moreover, non-Cartesian sampling like radial or spiral sampling are also good choices for CS and can be even more advantageous than Cartesian under-sampling.

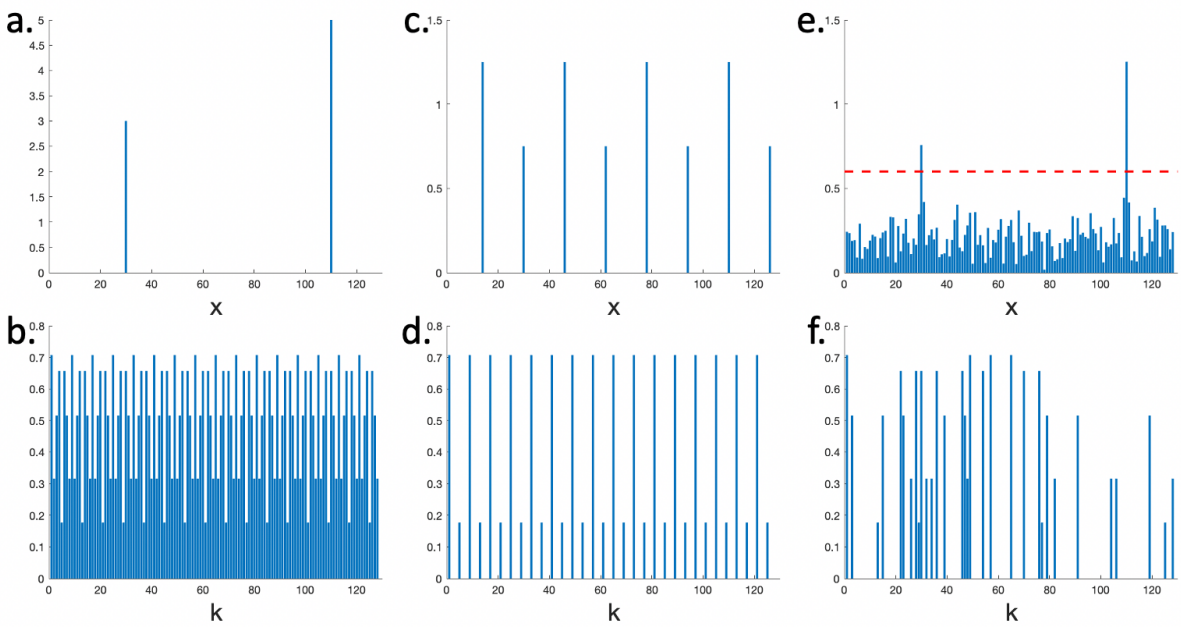

Fig. 2.10 The comparison of incoherent and coherent under-sampling. a) A sparse signal. b) The fully sampled k-space data of the sparse signal. c) The coherent aliasing artefacts resulted from a 4-fold uniform under-sampling in d). e) The incoherent, noise-like aliasing artefacts resulted from a 4-fold random under-sampling in f). The two strong sparse components in e) can be detected by thresholding.

The last ingredient to complete CS is nonlinear reconstruction, which is needed to enforce the prior knowledge of transform domain sparsity as well as taking account of data consistency with the acquired k-space samples. The reconstruction can be formulated as:

$$\min ||FX - Y||_2^2 + \lambda||\Psi X||_1 \qquad \text{Eq. 2.22}$$

Where F is the Fourier transform operator and Y is the acquired k-space data. Ψ denotes the sparsifying transform and ΨX returns the coefficients of X in the transform domain. $||\Psi X||_1$ calculates the l_1 norm of ΨX, which is defined as the sum of absolute values of all the coefficients. Although l_0 norm which is defined as the number of non-zero entries of an input is the ideal choice to enforce sparsity constraints, it is typically not used due to the complexity of solving the non-convex l_0 norm minimisation. Instead, l_1 norm is often used as a convex approximation of l_0 norm which also tends to promote sparse solutions. λ is the regularisation parameter which trades off between the data

consistency constraint and the sparsity constraint.

Some representative methods that can be used to solve the sparsity constrained optimisation are discussed below. One optimisation method widely used in image reconstruction is the alternating direction method of multipliers (ADMM) (Boyd, 2010). Firstly, Eq. 2.22 is reformulated as follows by introducing a dummy variable Z:

$$\begin{aligned} &\underset{X,Z}{\operatorname{argmin}} ||FX - Y||_2^2 + \lambda||Z||_1 \\ &s.t. Z = \Psi X \end{aligned} \qquad \text{Eq. 2.23}$$

It is then rewritten into an unconstrained optimisation by the augmented Lagrangian method (Hestenes, 1969) with a dual variable U:

$$\underset{X,Z,U}{\operatorname{argmin}} ||FX - Y||_2^2 + \lambda||Z||_1 + \mu||Z - \Psi X + U||_2^2 \qquad \text{Eq. 2.24}$$

Where μ is a scalar parameter. The primal variables X, Z and the dual variable U are updated in an iterative fashion. Each iteration consists of three steps and at the $n + 1_{th}$ iteration:

$$\begin{aligned} X_{n+1} &= \underset{X}{\operatorname{argmin}} ||FX - Y||_2^2 + \mu||Z_n - \Psi X + U_n||_2^2 \\ Z_{n+1} &= \underset{Z}{\operatorname{argmin}} \lambda||\mathrm{Z}||_1 + \mu||Z - \Psi X_{n+1} + U_n||_2^2 \\ U_{n+1} &= U_n + Z_{n+1} - \Psi X_{n+1} \end{aligned} \qquad \text{Eq. 2.25}$$

The update step of U ensures that the constraint in Eq. 2.25 is satisfied as the algorithm converges. The subproblem about X with two l_2 terms can be easily solved by methods like the conjugate gradient algorithm. The subproblem about Z with a l_1 term and a l_2 term now has the analytical solution using the soft thresholding function:

$$\hat{Z} = S_{\frac{\lambda}{2\mu}}(\Psi X - U) = sgn(\Psi X - U)max(|\Psi X - U| - \frac{\lambda}{2\mu}, 0) \qquad \text{Eq. 2.26}$$

A plot of the soft threshold shrinkage operator $y = S_\lambda(x)$ is shown in Fig. 2.11a.

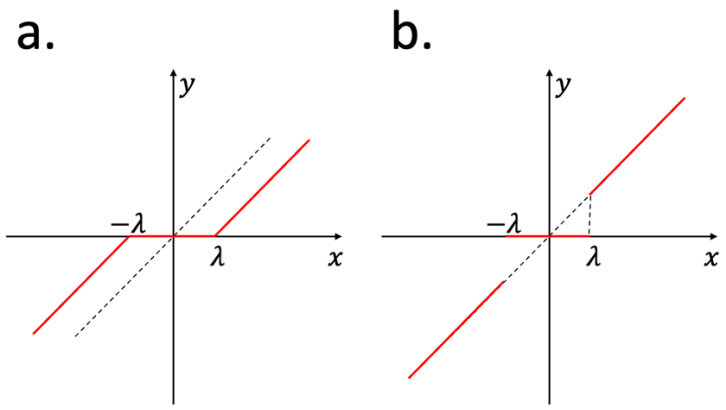

Fig. 2.11 The soft threshold shrinkage operator a) and hard threshold truncation operator b).

Eq. 2.22 is also often solved by the Fast Iterative Shrinkage-Thresholding Algorithm (FISTA) (Beck and Teboulle, 2009), which is a variant of the Iterative Shrinkage-Thresholding Algorithm (ISTA) with faster convergence speed. ISTA is an iterative optimisation algorithm which alternates between a gradient descent step for the l_2 term and a soft thresholding step for the l_1 term. The next estimate in each iteration can be represented as:

$$X_{inter} = X_n - \alpha F^H(FX_n - Y)$$
$$X_{n+1} = \Psi^H S_\rho(\Psi X_{inter}) \qquad \text{Eq. 2.27}$$

Where α is the step size of the gradient descent method. S_ρ is the soft thresholding operator with a threshold ρ as defined previously. The soft thresholding operator shrinks the transform coefficients towards zero to promote sparsity. FISTA achieves a faster convergence by making a modification that the new estimate in the current iteration also depends on the "gradient" of the previous iteration:

$$W_{n+1} = \Psi^H S_\rho\left(\Psi\left(X_n - F^H(FX_n - Y)\right)\right)$$
$$t_{n+1} = \left(1 + \sqrt{1 + 4t_n^2}\right)/2 \qquad \text{Eq. 2.28}$$
$$X_{n+1} = W_{n+1} + (t_n - 1)/t_{n+1}(W_{n+1} - W_n)$$

It is worth mentioning that the soft thresholding operator used in l_1 can often be replaced by the hard threshold truncation (Fig. 2.11b) corresponding to a non-convex rank

constraint, which has different theoretical guarantees (Blumensath and Davies, 2009) but can lead to better reconstruction in certain cases.

Eq. 2.22 can also be solved by the convex conjugate gradient algorithm using an approximate gradient function. The gradient function of l_1 norm is not well-defined since it involves the non-smooth absolute value function. However, the absolute value function can be approximated with a smooth surrogate, such as:

$$|x| = \sqrt{x^* x + \mu} \qquad \text{Eq. 2.29}$$

Where μ is a positive smoothing parameter. Thus, the gradient of $|x|$ can be approximated as $\frac{x}{\sqrt{x^* x+\mu}}$ and the gradient of $||\Psi X||_1$ can also be defined.

2.4 BOLD fMRI

In addition to structural imaging, MRI can also image various physiological functions and brain function is one of them. fMRI based on BOLD contrast (Ogawa et al., 1992, 1990) detects neuronal activity in the brain or spinal cord, which is one of the most important neuroimaging techniques underpinning clinical and basic neuroscience research. In this section, the mechanism of BOLD imaging and some of the considerations from the perspectives of MRI physics and experiment design are discussed.

BOLD imaging is an indirect measurement of neuronal activity (Logothetis et al., 2001) by means of changes in blood flow and oxygenation. The haemodynamic response describes the coupling between neuronal activity and blood flow, with the net effect that the local blood flow and volume is increased to deliver nutrients such as oxygen and glucose in response to the neuronal activity. The increased supply and consumption of oxygen is not balanced, which leads to a change in local blood oxygen level. Oxygen is carried by haemoglobin in red blood cells, which is an iron-containing molecule. The oxygenated haemoglobin (Hb) is diamagnetic while the deoxygenated haemoglobin

(dHb) is paramagnetic which distorts the local magnetic field, so the change in the ratio between Hb and dHb leads to a change in local field inhomogeneities. As discussed before, T_2^* is sensitive to local field inhomogeneities, so T_2^* weighted imaging is often used to detect BOLD signal changes. T_2 weighted imaging can be used too, and it is more sensitive to small vessels and capillaries, whereas gradient echo BOLD is more sensitive to larger vessels and draining veins.

In order to achieve T_2^* contrast, a relatively long TE close to the T_2^* of brain tissues is favored. The optimal TE for maximum BOLD sensitivity is around 25ms at 7T, which is approximately the T_2^* of gray matter. Also, relatively high temporal resolution is needed to capture the dynamic haemodynamic response, typically not longer than a few seconds. Based on these considerations, most fMRI studies use gradient echo EPI acquisition, and 2D simultaneous multi-slice (SMS) (Moeller et al., 2010a; Setsompop et al., 2012a) EPI has been the most popular choice due to its capability to achieve a short TR for high temporal resolution. The simultaneous multi-slice technique allows multiple 2D slices to be excited simultaneously using a multi-band RF pulse, which are then separated using an advanced image reconstruction method and coil sensitivity encoding. In recent years, 3D EPI has also been drawing growing interest in fMRI studies (Huber et al., 2018; Lawrence et al., 2019; Poser et al., 2010) to take advantage of its higher SNR which enables high resolution fMRI. The 3D EPI imaging excites the entire volume every shot and typically samples a partition of the 3D k-space per shot using an EPI readout. A more detailed comparison of 2D and 3D EPI acquisitions for fMRI is discussed in section 4.1.

Task-based fMRI is the most classic type of BOLD fMRI experiment (Buckner, 1998). In task-based fMRI, the subject is asked to perform some cognitive, motor or sensory tasks during the scan to engage brain activation in areas of interest. For example,

one of the most common tasks for studying the motor system is finger tapping, which causes brain activation in the corresponding motor cortex that can be detected using fMRI techniques. These tasks are often cued using experimental stimuli, and the stimuli can be presented using different designs. Two major types of experimental designs are block design and event-related design (Dale and Buckner, 1997; Friston et al., 1999). Block design is statistically powerful and straightforward to analyse. In fMRI studies with block design, the scan consists of alternating time blocks and within each block the subject either performs some tasks or is at resting-state. In contrast, an event-related design presents stimuli in a randomized way which offers greater flexibility. Note the change in MRI signal due to neuronal activity is very subtle, so statistical analysis of the image series is needed to obtain the activation map, which depicts the strength of activation across the brain. Typically the time course of each voxel is fitted to a general linear model (GLM) (Friston et al., 1994) which is the linear combination of each stimulus at its corresponding time point being convolved with the haemodynamic response. Task-based fMRI has enabled functional brain mapping, i.e., identification of brain regions that are involved in specific cognitive processes and behaviors, and visual cortex is one of the very first brain regions that has been studied using fMRI (Belliveau et al., 1991; Brady 1991; Engel et al., 1994). It has also be used in clinical settings to help diagnose and treat neurological and psychiatric disorders, such as epilepsy, stroke, traumatic brain injury, and depression (Bullmore, 2012).

In contrast to task-based fMRI, resting-state fMRI is another type of fMRI study where the subject is at rest during the scan and is not subject to any external stimuli (Barkhof et al., 2014; Biswal et al., 1995; Lv et al., 2018). Resting-state fMRI detects the spontaneous low-frequency MRI signal fluctuations which are correlated across different brain regions to investigate the functional connectivity between different regions of the

brain (Smith et al., 2013). It has enabled the discovery of various resting-state networks (Fox et al., 2005), e.g., the default mode network which has been associated with a variety of cognitive processes (Buckner, 2012). It is a network of brain regions that is active when a person is not focused on the outside world and the brain is at wakeful rest.

2.5 Summary

This chapter gives a brief introduction on some of the MRI basics. Firstly, the physical principles of MRI are introduced from the perspective of signal generation, contrast, and spatial encoding. Next, two categories of advanced image reconstruction methods, parallel imaging and compressed sensing are discussed. Lastly, fMRI based on BOLD contrast which is widely used to study brain functions is presented.

3 Low-rank and Structured Low-rank Image Reconstruction Methods

3.1 Introduction

Low-rank and structured low-rank (SLR) methods are similar but different MRI image reconstruction techniques. Both of them rely on the construction of low-rank matrices from MRI data, which are used as additional low-rank constraints to regularise the reconstruction. However, they build low-rank matrices in different ways. In this chapter, a brief review of low-rank and structured low-rank image reconstruction methods is presented, the latter of which is the focus of this thesis.

3.2 Foundations of Low-rank Methods

In this section, some basic concepts about low-rank model are reviewed which are the key aspects in low-rank and structured low-rank image reconstruction methods. The first part of this section explains the low-rank property of a matrix in theoretical and practical terms, and talks about how this low-rank property can be used for matrix completion, i.e., the estimation of missing matrix entries. The second part provides some solutions to the optimisation problem of low-rank matrix completion.

3.2.1 Low-rank Matrix Completion

The rank of a matrix Φ of size $m \times n$ (without loss of generality assuming $m \geq n$), denoted as $rank(\Phi)$, is defined as the number of independent columns (or rows) of Φ. It equals the dimension of its column space, which is the vector space spanned by its column vectors. This corresponds to the maximal number of linearly independent columns of Φ. Assuming $rank(\Phi) = r$, Φ is low-rank if r is considerably smaller than n. The fact that the dimension of the column space is smaller than the number of the column vectors indicates these column vectors are linearly dependent, and there exists a non-trivial null

space. Thus, the column vectors reside in a low dimensional space and can be represented by linear combinations of the basis vectors. This property can be used to infer the fact that the matrix Φ provides redundant information, and we do not need $m \times n$ independent DOF to characterise this matrix. Note the above statements regarding column vectors and column space apply to row vectors and row space equivalently.

One of the most useful tools for exploring low-rank matrices is the singular value decomposition (SVD). The SVD factorizes a matrix Φ in the form of $\Phi = U\Sigma V^H$, where U is a $m \times m$ unitary matrix, Σ is a $m \times n$ diagonal matrix, and V is a $n \times n$ unitary matrix as shown in Fig. 3.1. The diagonal entries of Σ are the singular values of Φ, typically in a descending order. The rank of a matrix equals the number of its non-zero singular values, so for a rank-r matrix Φ, Σ only has r non-zero diagonal entries. The column vectors of U are called left singular vectors, and they form an orthonormal basis for the column space of Φ. Similarly, the column vectors of V are called right singular vectors and they form an orthonormal basis for the row space of Φ.

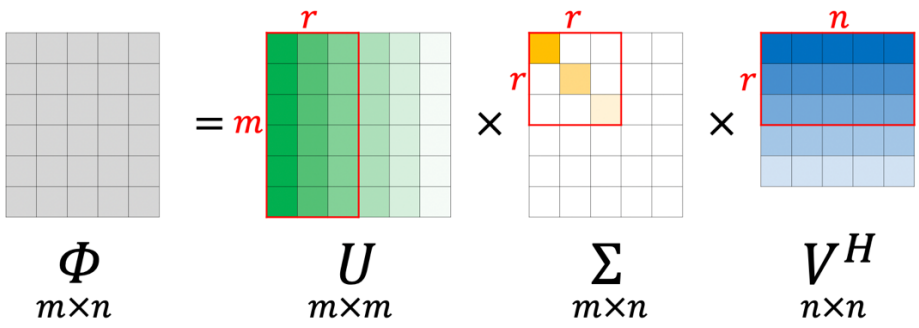

Fig. 3.1 Schematic singular value decomposition.

A low-rank matrix conveys redundant information and fewer DOF than its number of entries, which facilitates matrix completion, i.e., recovering the missing entries of a partially observed matrix. Given the prior knowledge that the matrix is low-rank, matrix completion becomes a well-posed problem with this additional constraint. The problem of low-rank matrix completion can be formulated as the following optimisation:

$$\min rank(X)$$
$$s.t. AX = Y \qquad \text{Eq. 3.1}$$

Where A is a sampling mask selecting all the observed entries of X, and Y contains all the known entries.

Since the rank minimisation problem is NP-hard, a convex relaxation of the problem by minimising the nuclear norm of the matrix is often used:

$$\min ||X||_*$$
$$s.t. AX = Y \qquad \text{Eq. 3.2}$$

Where $||X||_*$ denotes the nuclear norm of X. The nuclear norm of a matrix is defined as the sum of its singular values. This corresponds to the l_1 norm minimisation of singular values instead of the l_0 norm minimisation which is used in $rank(X)$. This replacement of l_0 norm with l_1 norm is analogous to the use of l_1 constraints in compressed sensing.

If the target rank r is known, the low-rank matrix completion can also be formulated using a hard rank constraint as:

$$\min ||AX - Y||_2^2$$
$$s.t. rank(X) = r \qquad \text{Eq. 3.3}$$

Or using an unconstrained, rank penalized formulation as:

$$\min ||AX - Y||_2^2 + \lambda ||X||_* \qquad \text{Eq. 3.4}$$

Where λ is the penalty parameter. Note the optimisation in Eq. 3.3 is a non-convex problem, whereas Eq. 3.4 is convex.

Low-rank matrix completion has a wide range of applications as truly independent random data is rather rare in reality. One example (Candes and Plan, 2010) is the prediction of the interests of users in commercial settings, like the customised recommendation of potentially interested products on Amazon or interested movies on Netflix. If we set up a matrix with each of its rows corresponding to the rating of a user

to all the products, this problem can be formulated as the recovery of the missing entries where users have not rated yet based on the observed entries where the users have already rated. Typically, this matrix is assumed to be low-rank as the preference of customers are usually predictable from a limited number of factors like age, area, profession etc.

3.2.2 Solutions to Low-rank Matrix Completion

The optimisation with a low-rank constraint in Eq. 3.3 and Eq. 3.4 can be solved by different methods. In this section, a few solutions which are typically used in MRI image reconstruction are presented, and they can be broadly divided in SVD-free and SVD-based methods.

One type of SVD-free method is based on matrix factorization which replaces the target low-rank matrix of a bigger size with the multiplication of two matrices of a smaller size (Haldar and Hernando, 2009). The cost function is described as:

$$\min_{\substack{U\in\mathbb{C}^{m\times r}\\ V\in\mathbb{C}^{n\times r}}} ||A(UV^H) - Y||_2^2 \qquad \text{Eq. 3.5}$$

Where U and V are at most rank-r by construction, and their product is also at most rank-r. Eq. 3.5 can be solved by the alternating minimisation approach. Also, instead of starting from the desired rank, the incremented rank power factorization (IRPF) algorithm (Haldar and Liang, 2010) uses an incremented rank strategy which starts solving Eq. 3.5 with $r = 1$ and increases r through iterations until the final desired rank is achieved. Each time the previous result of a smaller rank constraint is used as initialisation in the next iteration. This strategy is claimed to have improved performance and avoid local minima.

Another SVD-free method is the iterative reweighted least squares (IRLS) approach (Mohan and Fazel, 2012) which replaces the nuclear norm minimisation in Eq. 3.4 as a

weighted l_2 norm minimisation based on $||X||_* \leq ||XW^{1/2}||_F^2$, where $W = (X^H X)^{-1/4}$. Thus, Eq. 3.4 can be rewritten as:

$$\min ||AX - Y||_2^2 + \lambda ||XW^{1/2}||_F^2 \quad \text{Eq. 3.6}$$

The LRLS approach updates X and W iteratively so that in each iteration we get the new weight W first and then solve the reweighted version of Eq. 3.6 to update X.

The SVD-based methods for the rank-constrained optimisation are in general similar to the solutions to the sparsity constrained optimisation as discussed in the compressed sensing section in Chapter 2.3.2 (e.g., ADMM and FISTA), as the low-rank constraint minimises the number of singular values and the sparsity constraint minimises the number of transform coefficients (this is discussed in detail in section 3.3.4). These iterative SVD-based methods generally need a step in the form of either soft threshold shrinkage or hard threshold truncation of the singular values to promote the low-rank property of the matrix of interest.

3.3 Low-rank Reconstruction Methods

3.3.1 The Low-rank Model of MRI Data

Due to the inherently slow nature of MRI sampling, k-space is commonly under-sampled to shorten the scan time. MRI image reconstruction is in this case a problem of under-sampled reconstruction where we want to recover a fully sampled k space from only a subset of its measurements. The recovery of missing samples requires us to have prior knowledge about the data to impose regularisation. In compressed sensing, this prior knowledge is that the image is sparse in a transform domain. As discussed above, the low-rank property of a matrix can also be used as the prior knowledge to regularise matrix completion. To employ low-rank matrix completion in MRI image reconstruction, we

need to identify a low-rank matrix from MRI data and interpret the recovery of missing k-space samples as the completion of this low-rank matrix. In fact, MRI data itself in a variety of applications employing spatiotemporal imaging where a space by time image series is expected, can be modelled as a low-rank matrix (Haldar and Liang, 2010; Liang, 2007), and some typical examples include fMRI (Chiew et al., 2015; Nguyen and Glover, 2014), spectroscopic imaging (Lam and Liang, 2014), cardiac imaging (Zhao et al., 2010), and etc. As shown in Fig. 3.2, simply by vectorising the image at each time point (or k-space equivalently) as a column vector and concatenating all these vectors corresponding to different time points along the column dimension, the 3D space by time image series can be reshaped into a 2D matrix, which is called a Casorati matrix. The low-rank property of this space by time matrix is demonstrated empirically. From the perspective of row vectors, the intuition behind this is that some voxels may experience the same dynamic process and thus their time courses are correlated, which means that there might only be a small number of temporal basis functions corresponding to the right singular vectors of the low-rank matrix. For example, in task fMRI with visual cortex stimulus, many voxels in the visual cortex are expected to have very similar signal dynamics correlated with the experimental task. From the perspective of column vectors, it is also not hard to imagine that the images at different time points are linearly correlated, which means there are only a few spatial basis functions corresponding to the left singular values of the low-rank matrix. Also, the low-rank property of this spatiotemporal matrix can be justified empirically by performing SVD decomposition. For example, the work (Chiew et al., 2015) shows a plot of the singular values of fMRI datasets from multiple subjects, where >80% of the signal variance is captured in the top 25% of components.

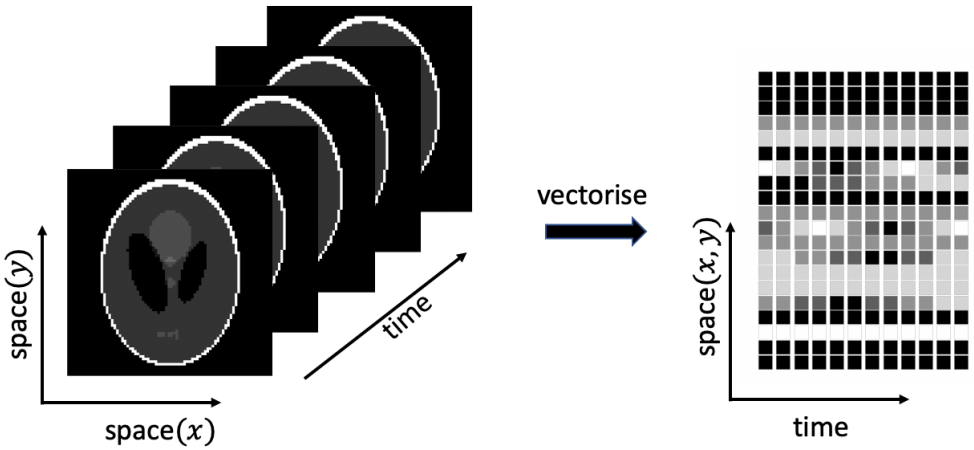

Fig. 3.2 The construction of Casorati matrix from a space by time image series.

3.3.2 Examples of Low-rank Reconstruction Methods

The spatiotemporal imaging using low-rank image reconstruction methods can be broadly divided into two categories. One of them involves the explicit use of the spatial and temporal basis functions, accompanied with specially designed data sampling scheme. These methods typically use a 2-step reconstruction where the basis functions are estimated as the first step. Another type is based on low-rank matrix completion without explicit use of the subspace vectors as well as special data sampling requirements. Both these two types of low-rank reconstruction methods are discussed in this section.

Low-rank reconstruction methods were first introduced to MRI in the form of partially separable functions (Liang, 2007). The spatiotemporal function $f(k,t)$ defined over the k-space and time dimension is modelled as L_{th}-order partially separable:

$$f(k,t) = \sum_{l=1}^{L} c_l(k)\varphi_l(t) \qquad \text{Eq. 3.7}$$

Where $c_l(k)$ is the spatial variation and $\varphi_l(t)$ is the temporal variation. Eq. 3.7 is true if $f(k,t)$ is rank-L, in which case $c_l(k)$ and $\varphi_l(t)$ correspond to the left and right singular vectors of $f(k,t)$ weighted by the square root of singular values. This original work proposed a 2-step reconstruction for computational simplicity. As the first step,

$c_l(k)$ and $\varphi_l(t)$ are estimated separately. Accordingly, a tailored sampling strategy was used which consists of two subsets of k-space, and Fig. 3.3 shows an example which is suitable for spectroscopic imaging. One subset $f_1(k,t)$ has high spatial resolution but limited temporal coverage, and the other subset $f_2(k,t)$ has high temporal resolution but limited k-space coverage. Firstly, $\varphi_l(t)$ is determined from the high temporal resolution subset $f_2(k,t)$ by performing a SVD decomposition. The choice of rank L considers the noise level of the measured data and the L most significant singular vectors are chosen as $\varphi_l(t)$. Secondly, $c_l(k)$ are calculated by fitting Eq. 3.7 to $f_1(k,t)$. Finally, the full matrix $f(k,t)$ is generated as in Eq. 3.7.

The method SPICE (Lam and Liang, 2014) uses a similar two-step reconstruction to significantly accelerate spectroscopic imaging. The method fLORA (Nguyen and Glover, 2014) exploits the partially separable model on fMRI data together with sparsity constraints based on features of BOLD signals for image reconstruction of sparsely sampled fMRI data. The k-t PCA method (Pedersen et al., 2009), although it does not exploit the low-rank model for image reconstruction directly as previous methods, also learns temporal basis functions using principle component analysis (PCA) on the training data which has a limited coverage in k-space but full coverage of the entire time frame, i.e., low spatial resolution and full temporal resolution. In some applications like quantitative imaging where the time dimension corresponds to signal evolution due to T_2/T_2^* relaxation, the basis functions can be estimated from simulation data. The low-rank model with predetermined temporal subspace obtained from Bloch simulations has also been used to improve the accuracy and speed of MRI fingerprinting (Zhao et al., 2018).

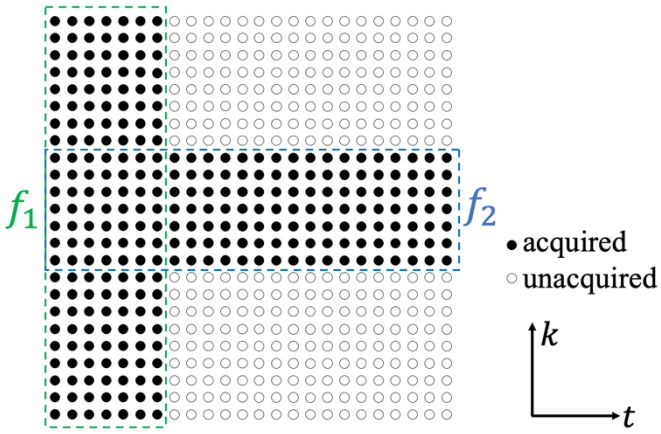

Fig. 3.3 A tailored sampling pattern for 2-step low-rank reconstruction.

As mentioned above, there are also some other low-rank reconstruction methods using a direct low-rank matrix completion without explicit estimation of the spatial and temporal basis functions based on a tailored sampling strategy. In later work (Haldar and Liang, 2010), the partially separable functions model is exploited as a low-rank matrix completion problem and the low-rank constrained reconstruction is formulated similar to Eq. 3.3 above when the rank r is known as a prior, or similar to Eq. 3.4 to enforce the low-rank constraint in a softer way to avoid explicit determination of the target rank r. The reconstruction on randomly sampled data has been validated in this work, although the optimal sampling for low-rank reconstruction is still an open problem. More theoretical discussions about the low-rank reconstruction framework for dynamic imaging was provided in the work (Haldar and Liang, 2011).

The spatiotemporal matrix of cardiac imaging has also been claimed to be partially separable on the order of 5-16, and this low-rank matrix recovery framework has been used in cardiac MRI without (Zhao et al., 2010) and with a joint sparsity constraint (Lingala et al., 2011; Zhao et al., 2012). There are also a variety of image reconstruction methods which exploit the low-rank spatiotemporal structure in fMRI data (Chiew et al., 2016, 2015; Lam et al., 2013; Mason et al., 2021). It is worth mentioning that PCA dimensionality reduction is often involved in fMRI analysis, which is also based on the

low dimensional structure of fMRI data. The method PCLA (Gao et al., 2014) has developed a phase constrained low-rank model for DWI to exploit the correlations between different diffusion directions while accounting for the drastic phase changes between them.

3.3.3 Other Forms of Low-rank Image Reconstruction

In addition to the low-rank reconstruction framework discussed above, there are also some other variants of low-rank reconstruction methods, e.g., low-rank tensor reconstruction, low rank plus sparsity model, and locally low-rank reconstruction, which are discussed in this section.

Low-rank Tensor Reconstruction

The low-rank reconstruction methods for spatiotemporal imaging have been extended to low-rank tensor reconstruction (Christodoulou et al., 2018; He et al., 2016) for MRI data with higher dimensionality. The low-rank model can be particularly effective on higher dimensional dataset in the sense that the number of entries of the low-rank tensor grow exponentially with the number of dimensions (curse of dimensionality), while its DOF only increases approximately linearly with the number of dimensions. Specifically, the multitasking method (Christodoulou et al., 2018) which uses a low-rank tensor model has proven successful for motion-resolved quantitative cardiovascular imaging, where the dataset with space by multiple time dimensions corresponding to T_1/T_2 relaxation, cardiac motion, respiratory motion and etc. is modelled as a low-rank tensor. Similar to the 2-step low-rank reconstruction, the data acquisition also consists of sparsely sampled imaging data which is interleaved with frequently sampled auxiliary training data for subspace estimation corresponding to physiological time dimensions e.g., cardiac and respiratory motion. Also, the temporal basis functions of the non-physiological time

dimensions, e.g., T_1 relaxation can be predetermined from a training dictionary generated by Bloch equations. The spatial basis functions are then reconstructed by fitting the measured data with the estimated temporal basis functions.

Low-rank + Sparsity

In some of the methods mentioned above, the low-rank model has been combined with sparsity constraints. In these methods, the low-rank and sparsity constraints are imposed independently, and the images to be reconstructed are assumed to have both low-rank and transform sparsity properties at the same time. In contrast, there are also some methods utilizing a low-rank plus sparsity (L+S) model (Chiew et al., 2018; Otazo et al., 2015; Petrov et al., 2017; Weizman et al., 2017), where the image time series is decomposed into a low-rank component and a sparse component. The low-rank component models the temporally coherent background and the sparse component models the dynamic signal changes. Incoherence between the L and S components are required for robust decomposition of background and dynamic components, i.e., the low-rank component is not sparse and vice versa. The L+S model has been demonstrated to have superior compressibility of dynamic MRI data compared to a low-rank only model in applications like cardiac imaging.

Locally Low-rank Reconstruction

In all previous methods, the low-rank matrix is constructed from the entire image or k-space matrix. In contrast, there are also some locally low-rank reconstruction methods which build a series of low-rank matrices from local patches of the image. The work (Trzasko and Manduca, 2011) is one of the first which proposed to spatially partition the image at each time point into overlapping blocks, and promote low-rank structure on the image series corresponding to each block independently. It was demonstrated that the locally low-rank approach can potentially outperform a global low rank approach when

the row and column dimensions of the global matrix differ significantly. Also, the locally low-rank model can be useful in some cases where the global low-rank model does not apply. For example, the CLEAR method (Trzasko and Manduca, 2012) uses the locally low-rank constraints to perform calibration-free parallel imaging reconstruction, where the coil sensitivities or the calibration data is not available. This method is based on the property that the coil sensitivities are spatially smooth. Thus, the matrices constructed from the coil sensitivity-modulated local patches of different coils are supposed to be low-rank. Similarly, the method proposed in (Hu et al., 2019) assumes the phase variations between different shots in a multi-shot acquisition are spatially smooth, so that each matrix constructed from phase-modulated local patches of different shots is low-rank. The locally low-rank constraint has also been used to accelerate parameter mapping (Zhang et al., 2015).

3.3.4 Comparison to Compressed Sensing

Fundamentally, low-rank constraints and compressed sensing both exploit data redundancy for the recovery of under-sampled signals, which actually have fewer DOF than the number of entries. Only in this case it is possible to recover the missing entries from a subset of observed entries. This data redundancy needs to be formulated explicitly as a prior to constrain the reconstruction. In compressed sensing, the prior knowledge is that the image is sparse in some transform domains. In low-rank methods, the prior knowledge is that the matrix constructed from MRI data is low-rank. In fact, they can be connected by means of SVD decomposition. Based on $\Phi = U\Sigma V^H$, we can get $U^H \Phi V = \Sigma$, so the multiplication of left and right singular vectors can be viewed in some sense as the sparsifying transform, and Σ containing the transform coefficients is sparse with only a limited number of non-zero values. Similarly, incoherent sampling is required in both

compressed sensing and low-rank matrix completion. In compressed sensing, incoherent sampling ensures that under-sampling artefacts manifest as the noise-like distributions in the transform domain, so it is possible to distinguish true signal components from artefacts by promoting sparsity. In low-rank matrix completion, the coherent under-sampling artefacts can also have low-rank structure and end up being mixed with the main signal components, which makes it difficult to isolate them by cutting off the tail of the singular value distribution.

The main difference between compressed sensing and low-rank reconstructions is that in compressed sensing, the sparsifying transform needs to be predetermined and is data-invariant. However, the "sparsifying transform" in low-rank methods – the left and right singular vectors, are data-dependent and no knowledge of their actual forms is needed ahead of time.

3.4 Structured Low-rank Reconstruction Methods

3.4.1 The Structured Low-rank Model of MRI Data

Besides the conventional low-rank image reconstruction discussed in Section 3.3, structured low-rank image reconstruction, as another subspace method, has been increasingly popular and demonstrated in a variety of MRI applications. In contrast to conventional low-rank reconstruction where the low-rank matrix is the multi-dimensional dataset itself, structured low-rank reconstruction uses a transformation of the k-space, which typically has a block-Hankel (or Toeplitz) structure as shown in Fig. 3.4, as the low-rank matrix whose low-rankness can be theoretically justified. This change makes structured low-rank reconstruction a much more flexible framework which is no longer limited to multi-dimensional imaging like spatiotemporal imaging. The mechanism of structured low-rank reconstruction is illustrated next via a specific example, the

LORAKS method (Haldar, 2014), which is one of the earliest work of structured low-rank reconstruction in MRI.

The LORAKS method demonstrates how a structured low-rank model can be built on a single image without any information from additional dimensions corresponding to time points, coils, etc., and with the only prior knowledge that it has limited image support, which is a very common condition. It is worth mentioning that limited support means the image is sparse in the original space domain. Without loss of generality, we consider a 2D image of limited support, which is denoted by $\rho(x, y)$. Assuming $f(x, y)$ is any smooth function which has a complementary image support, i.e., it only has non-zero values outside of the support of $\rho(x, y)$, we can get the following equation:

$$\rho(x, y) f(x, y) = 0 \qquad \text{Eq. 3.8}$$

By applying Fourier transform on both sides, this relation can also be represented in k-space:

$$\tilde{\rho}(kx, ky) * \tilde{f}(kx, ky) = 0 \qquad \text{Eq. 3.9}$$

Where $\tilde{\rho}(kx, ky)$ and $\tilde{f}(kx, ky)$ are the Fourier transforms of $\rho(x, y)$ and $f(x, y)$ respectively. $*$ denotes matrix convolution. This kind of equations with 0 on one side are termed annihilation relations, which are a key concept in structured low-rank reconstruction. Since $f(x, y)$ is spatially smooth by definition, its Fourier transform $\tilde{f}(kx, ky)$ has a limited support, which is assumed to be $\{-d \leq kx \leq d, -d \leq ky \leq d\}$. The matrix convolution output in Eq. 3.9 can implemented as:

$$\sum_{u=-d, v=-d}^{u=d, v=d} \tilde{\rho}(kx - u, ky - v) \tilde{f}(u, v) = 0 \qquad \text{Eq. 3.10}$$

This indicates that the entries of each $(2d + 1) \times (2d + 1)$ patch of $\tilde{\rho}(kx, ky)$ are linearly dependent as there exists a set of weights corresponding to $\tilde{f}(kx, ky)$ with which

their weighted sum is zero. As shown in Fig. 3.4, we can construct a matrix by vectorising each patch of $\tilde{\rho}(kx, ky)$ selected by a $(2d+1) \times (2d+1)$ sliding window into a row vector and concatenating all the row vectors along the row dimension as this window slides through k-space. This matrix is a block-Hankel structured matrix, and the operator H is defined as the block-Hankel transformation. As shown in Fig. 3.5, if the sliding window traverses k-space in the same direction as how different entries in each patch are vectorised (first vertical then horizontal in this case), this block-Hankel structured matrix will have a constant antidiagonal structure within and between blocks. As defined in LORAKS (Haldar, 2014), this block-Hankel structured matrix is called the C-matrix. According to Eq. 3.10, the C-matrix constructed from $\tilde{\rho}(kx, ky)$ has a null-space vector corresponding to the entries of $\tilde{f}(kx, ky)$. Since we can often construct a large number of linearly independent f that satisfy the annihilation relations described in Eq. 3.8 and Eq. 3.9, the C-matrix has a non-trivial null-space and its column space lies in a low dimensional signal subspace, which indicates it is low-rank. It is shown that the dimension of its column space and thus its rank depends on the specific support of $\rho(x, y)$.

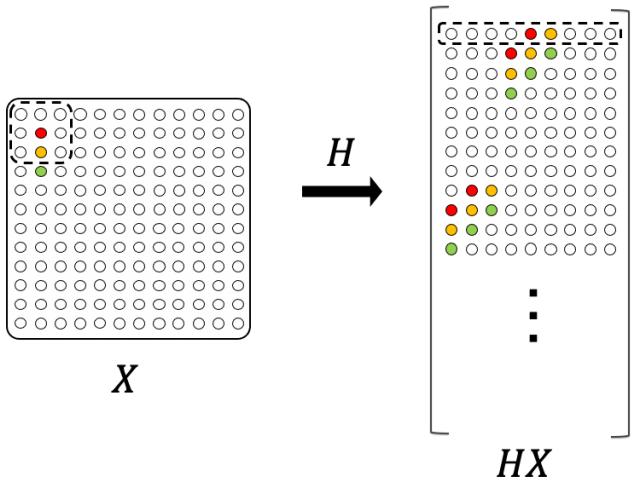

Fig. 3.4 The construction of block-Hankel structured matrix from a single k-space.

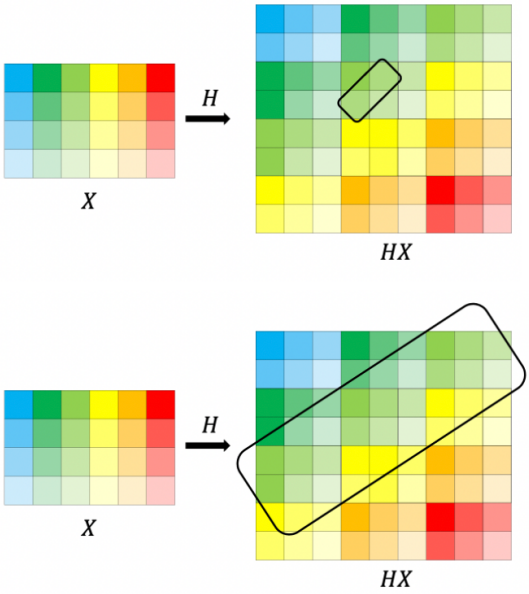

Fig. 3.5 The constant antidiagonal structure within and between blocks of the block-Hankel structured matrix.

Unlike conventional low-rank matrices where there is a direct and intuitive mapping between the signal correlations in multi-dimensional dataset and the low-rank model, in structured low-rank reconstruction, two additional steps are needed to build up this mapping. Firstly, we need to formulate the prior knowledge about the image, e.g., limited image support, as an annihilation relation. Secondly, we need to identify a low-rank matrix based on this annihilation relation. In the example above, the annihilation relations in image space and k-space are shown in Eq. 3.8 and Eq. 3.9 respectively. Accordingly, a low-rank matrix is identified as the C-matrix, in the format of a block-Hankel structured matrix constructed from a single k-space. Note different annihilation relations can be formulated exploiting various signal properties, from which different low-rank matrices, typically having a block-Hankel structure, can be built.

3.4.2 Examples of Structured Low-rank Reconstruction Methods

Based on the applications, structured low-rank image reconstruction methods can be

broadly grouped into two categories. One of them works on the MRI dataset with single channel/contrast/time point etc. As shown above, the LORAKS method demonstrates a structured low-rank model from the k-space of a single image with limited image support. Besides, LORAKS also presents another structured low-rank model on the image with slowly varying phase, in which case an annihilation relation can be formulated as:

$$\rho(x,y)h(x,y) = \rho^*(x,y)h^*(x,y) \qquad \text{Eq. 3.11}$$

Where $\rho^*(x,y)$ is the complex conjugate of $\rho(x,y)$. $h(x,y)$ is any smooth function that satisfies:

$$h(x,y) = \exp\big(-1j\angle\rho(x,y)\big)\ \forall\rho(x,y) \neq 0 \qquad \text{Eq. 3.12}$$

According to Eq. 3.11, an annihilation relation in k-space can be derived as:

$$\tilde{\rho}(kx,ky) * \tilde{h}(kx,ky) - \tilde{\rho}^*(-kx,-ky) * \tilde{h}^*(-kx,-ky) = 0 \qquad \text{Eq. 3.13}$$

Since the phase of $\rho(x,y)$ and $h(x,y)$ are spatially smooth, $\tilde{h}(kx,ky)$ has a limited support in k-space, assuming to be $\{-d \le kx \le d, -d \le ky \le d\}$. Eq. 3.13 can then be implemented as:

$$\sum_{u=-d,v=-d}^{u=d,v=d} \tilde{\rho}(kx-u,ky-v)\tilde{h}(u,v) - \sum_{u=-d,v=-d}^{u=d,v=d} \tilde{\rho}^*(-kx-u,-ky-v)\tilde{h}^*(u,v) = 0 \qquad \text{Eq. 3.14}$$

Considering the real and imaginary components of complex numbers respectively (i.e., $\rho = \rho_{real} + 1j\rho_{imag}$), Eq. 3.14 can be rewritten into:

For real component:

$$\sum_{u=-d,v=-d}^{u=d,v=d} \tilde{\rho}_{real}(kx-u,ky-v)\tilde{h}_{real}(u,v) - \sum_{u=-d,v=-d}^{u=d,v=d} \tilde{\rho}_{imag}(kx-u,ky-v)\tilde{h}_{imag}(u,v) - \sum_{u=-d,v=-d}^{u=d,v=d} \tilde{\rho}_{real}(-kx-u,-ky-v)\tilde{h}_{real}(u,v) + \sum_{u=-d,v=-d}^{u=d,v=d} \tilde{\rho}_{imag}(-kx-u,-ky-v)\tilde{h}_{imag}(u,v) = 0 \quad \text{Eq. 3.15}$$

For imaginary component:

$$\sum_{i=-d,j=-d}^{i=d,j=d} \tilde{\rho}_{real}(kx-i,ky-j)\tilde{h}_{imag}(i,j) + \sum_{i=-d,j=-d}^{i=d,j=d} \tilde{\rho}_{imag}(kx-i,ky-j)\tilde{h}_{real}(i,j) + \sum_{i=-d,j=-d}^{i=d,j=d} \tilde{\rho}_{real}(-kx-i,-ky-j)\tilde{h}_{imag}(i,j) + \sum_{i=-d,j=-d}^{i=d,j=d} \tilde{\rho}_{imag}(-kx-i,-ky-j)\tilde{h}_{real}(i,j) = 0 \quad \text{Eq. 3.16}$$

As defined in LORAKS (Haldar, 2014), the S-matrix can be generated as:

$$\begin{bmatrix} H\left(\tilde{\rho}_{real}(kx,ky)\right) - H\left(\tilde{\rho}_{real}(-kx,-ky)\right) & -H\left(\tilde{\rho}_{imag}(kx,ky)\right) + H\left(\tilde{\rho}_{imag}(-kx,-ky)\right) \\ H\left(\tilde{\rho}_{imag}(kx,ky)\right) + H\left(\tilde{\rho}_{imag}(-kx,-ky)\right) & H\left(\tilde{\rho}_{real}(kx,ky)\right) + H\left(\tilde{\rho}_{real}(-kx,-ky)\right) \end{bmatrix} \quad \text{Eq. 3.17}$$

Where H is the operator which performs block-Hankel transformation as defined before. According to Eq. 3.15 and Eq. 3.16, this S-matrix has a null space vector

corresponding to $\tilde{h}(kx, ky)$. Since we can easily modify Eq. 3.11 by multiplying another smooth phase term on each side of the equation, this S-matrix potentially has a large null-space and is low-rank.

Another category of structured low-rank reconstruction methods works in high dimensional imaging scenarios, as in the conventional low-rank methods in Section 3.3. An important application of structured low-rank reconstruction methods is to deal with the image phase inconsistencies between different shots, which can result in a variety of image artefacts depending on the imaging context. For example, EPI Nyquist ghost artefacts which originate from the mismatch between the positive and negative readouts due to gradient delay and eddy currents can be described as the phase inconsistencies between old and even echoes, since a shift in k-space is equivalent to a phase modulation in image space. Structured low-rank reconstruction methods (J. Lee et al., 2016; Lobos et al., 2021, 2018) have been used successfully to correct for the EPI Nyquist ghost artefacts. Similarly, structured low-rank reconstruction has also been used to deal with the trajectory errors for radial sampling (Mani et al., 2018), where the shifted readouts due to gradient delay and eddy currents can be described as the image phase inconsistencies between different shots. For high resolution DWI using multi-shot EPI acquisition, the shot-to-shot phase variations due to subtle subject movement being magnified by the strong diffusion gradient is a significant issue, which has been addressed by the structured low-rank method MUSSELS (Mani et al., 2017). fMRI imaging with 3D multi-shot EPI acquisition at ultra-high fields like 7T also suffers from inter-shot phase variations arising from physiologically induced field changes, and Chapter 4 discusses how structured low-rank reconstruction can be used to deal with this issue in detail (Chen et al., 2023).

Typically, methods dealing with these issues rely on measurement of the phase variations by navigators. In contrast, structured low-rank reconstruction methods

formulate the problem as the reconstruction of a higher dimensional dataset with an additional shot dimension accounting for the phase fluctuations, i.e., the under-sampled k-space of each shot in a given phase state is recovered, and then the images corresponding to each individual shot are SOS combined which is insensitive to their phase differences. To accommodate the higher under-sampling factor for each shot, all shots are jointly reconstructed to exploit the linear correlation between them, as they only vary in spatially smooth phase while having the same magnitude signal. This linear correlation leads to the low-rank property of the block-Hankel structured matrix constructed from the multi-shot k-space. This process is illustrated below starting with the existence of the annihilation relation:

$$M_i(x,y)\Phi_j(x,y) - M_j(x,y)\Phi_i(x,y) = 0 \qquad \text{Eq. 3.18}$$

Where $M_i(x,y) = M(x,y)\Phi_i(x,y)$ is the image for the i_{th} shot, which is expressed as the uncorrupted complex image $M(x,y)$ being modulated by the phase modulation $\Phi_i(x,y)$. The same applies for $M_j(x,y)$. This annihilation relation can also be expressed in k-space as:

$$\widetilde{M}_i(kx,ky) * \widetilde{\Phi}_j(kx,ky) - \widetilde{M}_j(kx,ky) * \widetilde{\Phi}_i(kx,ky) = 0 \qquad \text{Eq. 3.19}$$

Where $\widetilde{M}_i(kx,ky)$ and $\widetilde{M}_j(kx,ky)$ are the Fourier transforms of $M_i(x,y)$ and $M_j(x,y)$. $\widetilde{\Phi}_i(kx,ky)$ and $\widetilde{\Phi}_j(kx,ky)$ are the Fourier transforms of $\Phi_i(x,y)$ and $\Phi_j(x,y)$. Since the phase variations are assumed to be spatially smooth, $\widetilde{\Phi}(kx,ky)$ has a limited support in k-space which is assumed to be $\{-d \le kx \le d, -d \le ky \le d\}$, and Eq. 3.19 can be implemented as:

$$\sum_{u=-d,v=-d}^{u=d,v=d} \widetilde{M}_i(kx-u,ky-v)\widetilde{\Phi}_j(u,v) - \sum_{u=-d,v=-d}^{u=d,v=d} \widetilde{M}_j(kx-u,ky-v)\widetilde{\Phi}_i(u,v) = 0 \quad \text{Eq. 3.20}$$

By using the block-Hankel structured matrix representations of k-space data, the convolution operation in Eq. 3.19 and Eq. 3.20 can be implemented as matrix multiplication:

$$H(\widetilde{M}_i(kx,ky))vec(\widetilde{\Phi}_j(kx,ky) - H(\widetilde{M}_j(kx,ky))vec(\widetilde{\Phi}_i(kx,ky)) = 0 \quad \text{Eq. 3.21}$$

Where vec is the vectorisation operator. We can further rewrite Eq. 3.21 in a matrix form as:

$$\left[H(\widetilde{M}_i(kx,ky))\ \ H(\widetilde{M}_j(kx,ky))\right]\begin{bmatrix} vec(\widetilde{\Phi}_j(kx,ky)) \\ vec(\widetilde{\Phi}_i(kx,ky)) \end{bmatrix} = 0 \quad \text{Eq. 3.22}$$

As this annihilation relation holds true for any two shots, we can infer that the matrix generated by concatenating the block-Hankel structured matrices of all N shots as $\left[H(\widetilde{M}_1)\ H(\widetilde{M}_2), \ldots\ H(\widetilde{M}_N)\right]$, has a non-trivial null space and is low-rank. Fig. 3.6 shows how to construct this enlarged block-Hankel structured matrix from multi-shot k-space, which is denoted by the operator H_s. By enforcing its low-rank structure, we can recover the missing k-space data of each shot using a low-rank matrix completion method. Note as mentioned above in Section 3.3.2, the phase constrained low-rank method PCLA has to incorporate a low-resolution phase estimation explicitly to ensure the low-rank structure of the matrix consisting of data from different diffusion weighted directions, whereas by using the block-Hankel structured matrix representation of k-space, no phase estimation is needed in the structured low-rank method.

A similar format of block-Hankel structured matrix construction from multi-shot k-

space can also be adapted to a single k-space. If we consider an image with smooth phase, its conjugate symmetrical k-space can be viewed as from a virtual shot which has the conjugate image phase of the original image (Blaimer et al., 2009). Thus, a low-rank matrix can be generated as the concatenation of two block-Hankel structured matrices constructed from the original shot and the virtual shot respectively (Bilgic et al., 2019; Kim and Haldar, 2018), which is called the VC-matrix. Note the VC-matrix and the S-matrix both exploit the smoothness of the image phase, but they are constructed in different ways.

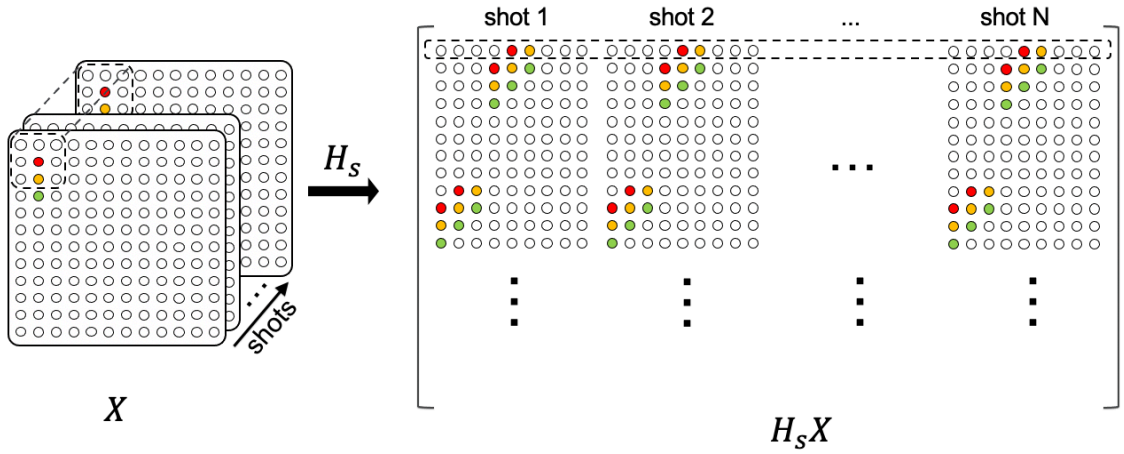

Fig. 3.6 The construction of block-Hankel structured matrix from multi-shot k-space.

The methods SAKE (Shin et al., 2014a) and P-LORAKS (Haldar and Zhuo, 2016) perform similar calibration-less parallel imaging reconstruction based on structured low-rank matrix completion. The structured matrix in this case is the concatenation of the block-Hankel structured matrices constructed from each coil, like how it is constructed from the multi-shot k-space where the smooth coil sensitivities play the same role as the phase variations. Note in this case the block-Hankel structured matrix is the same as the calibration matrix in k-space based parallel imaging reconstruction methods like GRAPPA (Griswold et al., 2002) and ESPIRiT (Uecker et al., 2014). By enforcing the low-rankness on this structured matrix, we are exploiting the prior knowledge that there

is data redundancy in the multi-channel k-space due to sensitivity encoding.

Structured low-rank methods have also been used for calibration-free B_0 inhomogeneity correction of EPI data (Balachandrasekaran et al., 2019; Siemonsma et al., 2020). The idea resembles the methods dealing with inter-shot phase variations. The problem is reformulated as the reconstruction of image time series with an additional time dimension using time segmentation, so that each segment has a reduced temporal footprint and is less susceptible to the field inhomogeneity. The exponential signal evolution is exploited to formulate an annihilation relation in k-t space, from which a structured low-rank matrix is generated.

The structured low-rank matrix in methods above has been extended to a structured low-rank tensor for a higher dimensional dataset with more than two fundamental dimensions. The method (Yi et al., 2021) proposed to perform a calibration-less, multi-channel, multi-contrast joint reconstruction through a block-Hankel structured low-rank tensor completion framework. Considering a block-Hankel structured low-rank matrix constructed from multi-channel k-space, a tensor can be generated by stacking the structured matrix corresponding to each contrast along a third "contrast" dimension. Different signal properties are exploited by enforcing the low-rank property on different modes of unfolding matrices of this structured tensor. By merging the channel and contrast dimensions, different contrasts are considered as virtual channels. In the meantime, by combining the kernel and contrast dimensions, different contrasts are treated as from virtual k-space kernels corresponding to images with common spatial support and shared coil sensitivity. A calibration-less, multi-channel, multi-slice reconstruction method using structured low-rank tensor completion is also presented in a similar way (Liu et al., 2021). The method (Hess et al., 2021) uses a calibration-less joint transmit and receive low-rank tensor completion algorithm for accelerated parallel

transmit mapping.

3.4.3 Comparison to Low-rank Reconstruction

The main difference between structured low-rank and conventional low-rank reconstruction is that the low-rank matrix is constructed in different ways. In structured low-rank reconstruction, the low-rank matrix is the block-Hankel structured matrix, whereas in low-rank reconstruction it is simply the reshaped version of multi-dimensional data itself. The conventional low-rank reconstruction can be considered a special case of structured low-rank reconstruction with a 1×1 kernel of the block-Hankel transformation. The discussion above has shown that the size of the kernel corresponds to the size of the local neighbourhood in which the linear dependencies in k-space is restricted. As the linear dependencies in k-space result from the convolution operation due to factors like sensitivity encoding, the intrinsic range of "leakage" is determined by the size of the convolution kernel which corresponds to e.g., the k-space support of the coil sensitivities. Thus, typically a larger kernel of the block-Hankel transformation leads to better accuracy of the structured low-rank model. This formulation also makes structured low-rank reconstruction capable of exploring the linear dependencies in a single k-space, but the conventional low-rank reconstruction, with a 1×1 kernel, does not work in this case.

3.5 Summary

In this chapter, MRI image reconstruction methods based on the low-rank structure of the matrix of interest are discussed, which are divided into two categories: conventional low-rank reconstruction methods and structured low-rank reconstruction methods. Although these methods vary in the exact formulation of low-rank constraint in different applications, they all share the explicit use of linear correlations in MRI data. No man is

an island. No k-space point is an island. Thus, it is possible to infer the missing data from the acquired data.

4 Structured Low-rank Reconstruction for 3D Multi-shot EPI fMRI

4.1 Introduction

As briefly mentioned before, 3D multi-shot EPI suffers from shot-to-shot physiological fluctuations, and SLR reconstruction can be exploited to address this issue for fMRI. In this Chapter, a detailed discussion of this work is presented.

Although 2D SMS-EPI acquisition has been the workhorse of fMRI studies (Moeller et al., 2010b), 3D encoding methods are drawing growing interest and increasingly being explored as alternatives (Poser et al., 2010). Compared to 2D multi-slice imaging, 3D imaging can provide a SNR advantage proportional to the square root of the number of shots since the whole volume is excited repeatedly by every shot, despite the lower steady state magnetization due to the much shorter time interval between subsequent excitations of the same tissue. Also, 3D imaging is less likely to run into specific absorption rate (SAR) issues due to the much lower optimal flip angles compared to 2D SMS-EPI. In addition, 3D imaging can achieve high isotropic resolution without being limited by RF slice profiles along the slice encoding direction. All these advantages contribute to the prevalence of 3D encoding in high isotropic resolution fMRI, which has received strong interests for applications such as layer-specific fMRI (Lawrence et al., 2019).

Unlike 2D EPI, multi-shot acquisition is typically required for 3D EPI in practice due to the limit on echo train length. The main disadvantage of 3D multi-shot EPI imaging, compared to 2D acquisitions which use a single-shot readout like 2D SMS-EPI, is that it suffers from increased vulnerability to shot-to-shot inconsistencies, e.g., due to physiological fluctuations. The tSNR of a time series, defined as its temporal mean divided by the temporal standard deviation (std), reflects the temporal stability of a given time course, and relates to the sensitivity to subtle activations of the fMRI measurement. In the absence of any temporal fluctuations like physiological noise or scanner instabilities, tSNR is the same as the image SNR of each volume. In practice, physiological noise, which scales with the signal intensity, creates an asymptotic limit on the achievable tSNR (Krüger and Glover, 2001; Triantafyllou et al., 2006, 2005). Previous work has shown that this tSNR limit depends on the number of shots used for a conventional 3D multi-shot EPI acquisition, and that a larger number of shots is usually

associated with a lower tSNR at the same SNR level (van der Zwaag et al., 2012), since the shot-to-shot k-space inconsistencies can result in temporally varying ghost artefacts and thus a reduction in tSNR. Hence, the SNR benefits enabled by 3D encoding are not fully realized for 3D multi-shot EPI acquisition, and it might only be able to offer a higher tSNR than 2D single shot EPI in low SNR, thermal noise dominated regimes (Jorge et al., 2013; Lutti et al., 2013). Thus, a typical choice is to use 3D EPI at 1 mm isotropic resolution or smaller voxel sizes, and 2D EPI for bigger voxel sizes. Previous work (Huber et al., 2018) has showed that the tSNR curves for 3D and 2D EPI crossed at around 1 mm isotropic resolution for BOLD fMRI at 7T.

The physiological fluctuations which trouble 3D multi-shot EPI acquisition can manifest themselves as rather localized fluctuations of blood and cerebrospinal fluid related to cardiac pulsation, and spatially varying image phase modulations resulting from the B_0 field fluctuations mainly caused by the movement of the chest during respiration (Tijssen et al., 2011; Zahneisen et al., 2014), the latter of which is particularly troublesome for multi-shot imaging. Also, as the off-resonance effects scale with field strength (Triantafyllou et al., 2006), the physiologically induced inter-shot phase variations can be more detrimental at ultra-high fields like 7T, which plays an important role in high resolution fMRI due to its ability to boost SNR in small voxel regimes.

While various methods have been proposed to reduce physiological noise for fMRI, most of them are post-processing methods, either “model-based” (Agrawal et al., 2020; Glover et al., 2000; Hutton et al., 2011; Jorge et al., 2013; Kasper et al., 2017; Lutti et al., 2013; Tijssen et al., 2014) or “data-driven” (Behzadi et al., 2007; Jorge et al., 2013; Kasper et al., 2017; Salimi-Khorshidi et al., 2014) that work on reconstructed image time series. Unlike these post-processing methods, reconstruction methods which take into account the characteristics of multi-shot acquisitions generally rely on navigators to

estimate the shot-to-shot phase variations (Barry et al., 2008; Barry and Menon, 2005). However, navigator techniques estimate dynamic phase information at the cost of prolonged acquisition time and reduced temporal resolution, and it can be particularly challenging to acquire 3D navigators. In the context of DWI, subtle motion induced inter-shot phase variations due to the use of strong diffusion encoding gradients present a similar problem for multi-shot EPI acquisitions, and a variety of methods have been proposed to deal with this issue. These methods in general are also based on explicit phase estimates (Atkinson et al., 2006; Chen et al., 2013; Porter and Heidemann, 2009). Recently, structured low-rank (SLR) image reconstruction methods based on block-Hankel structured low-rank matrix completion have been proposed to resolve the odd/even echo phase differences (J. Lee et al., 2016; Lobos et al., 2018, 2021) as well as the shot-to-shot phase variations for EPI acquisitions (Mani et al., 2017). Specifically, the MUSSELS method (Mani et al., 2017) originally proposed for 2D multi-shot DWI data has demonstrated its ability to account for inter-shot phase variations without requiring explicit knowledge of them, and has shown superior performance compared to explicit phase estimation methods. While few applications of structured low-rank reconstructions have focused on fMRI, the ALOHA method (J. Lee et al., 2016) and the RAC-LORAKS method (Lobos et al., 2021) evaluated Nyquist ghost corrections in fMRI or fMRI-like experiments.

In this chapter, a structured low-rank image reconstruction method is proposed for 3D multi-shot EPI for fMRI at 7T to reduce its vulnerability to physiologically induced inter-shot phase variations and thus improve its temporal stability, at different isotropic spatial resolutions between 1.0 mm and 1.8 mm. To adapt to the proposed reconstruction, a 3D CAIPIRINHA (CAIPI) sampling trajectory with interleaved ordering along k_z, termed “seg-CAIPI” is also presented. Both simulation and in vivo experiments have

demonstrated that the combination of the seg-CAIPI sampling trajectory and structured low-rank reconstruction method provides a robust way to improve the tSNR of 3D multi-shot EPI imaging for fMRI at 7T.

4.2 Methods

An ideal and straightforward solution to reduce vulnerability to shot-to-shot inconsistencies is to reconstruct an individual image for each shot and perform a phase-insensitive shot combination in image space, which is the strategy of the proposed SLR reconstruction. However, as 3D EPI acquisitions for whole brain imaging typically consist of a large number of shots, the proposed method bins multiple consecutive shots together into shot groups, and performs joint reconstruction for each shot group instead to improve the conditioning of the reconstruction. In accordance with the proposed reconstruction, the seg-CAIPI trajectory aims at optimising the sampling pattern for each shot group. The seg-CAIPI sampling and the joint reconstruction of all shot groups by SLR method are demonstrated in detail in the following sections.

4.2.1 Seg-CAIPI Sampling Trajectory

CAIPI sampling patterns (Breuer et al., 2005; Setsompop et al., 2012b) have reduced g-factors compared to conventional non-CAIPI sampling, and 2D CAIPIRINHA have been widely used in 3D EPI acquisitions (Breuer et al., 2006). A conventional 3D blipped-CAIPI multi-shot EPI sampling pattern (Narsude et al., 2016) is shown in Fig. 4.1a, which uses a sequential ordering along the second phase encoding direction, and it is referred to as the "standard" trajectory in this work. The proposed seg-CAIPI sampling which uses interleaved ordering along k_z combined with a k_z-blipped CAIPI pattern is shown in Fig. 4.1b-e. Note unlike some previous works (Hendriks et al., 2020; Stirnberg and Stöcker, 2021), here the segmentation and interleaved ordering are introduced along the k_z

direction instead of along the k_y direction for each k_z plane.

The parameter "width" is used to describe the offset along k_z between two consecutive shots with identical trajectories. The width of the standard trajectory is the same as R_z, the under-sampling factor along the shot dimension. For seg-CAIPI, its width is now $R_z \times N$, since interleaved ordering along k_z with *N* interleaves is used. Thus, k-space is traversed *N*-fold faster along k_z. The span of each shot along k_z for seg-CAIPI is also defined by the width parameter, indicating an *N*-fold wider span in accordance with the *N*-fold larger offset between two consecutive shots. Note this allows for a Δk_z blip size $\geq R_z$. The next interleave traverses the k-space again to acquire missing samples with a complementary trajectory, which has the same Δk_z blip size and k_z-span but is offset by R_z along k_z. Each specific standard blipped-CAIPI and seg-CAIPI trajectory is marked by $(\text{width}, \Delta k_z)$. Although different Δk_z blip sizes might be chosen for the standard and the seg-CAIPI schemes, the overall sampling masks are designed to be exactly the same. Typically, all shots corresponding to an interleave are binned as a shot group. In the example shown in Fig. 4.1, each shot group consists of three consecutive shots in Fig. 4.1b&d, and two shots in Fig. 4.1c&e respectively. In general, each shot group has a more uniform under-sampling pattern for the seg-CAIPI scheme, while for the standard trajectory, each shot group if defined in the same way has a more localized sampling pattern.

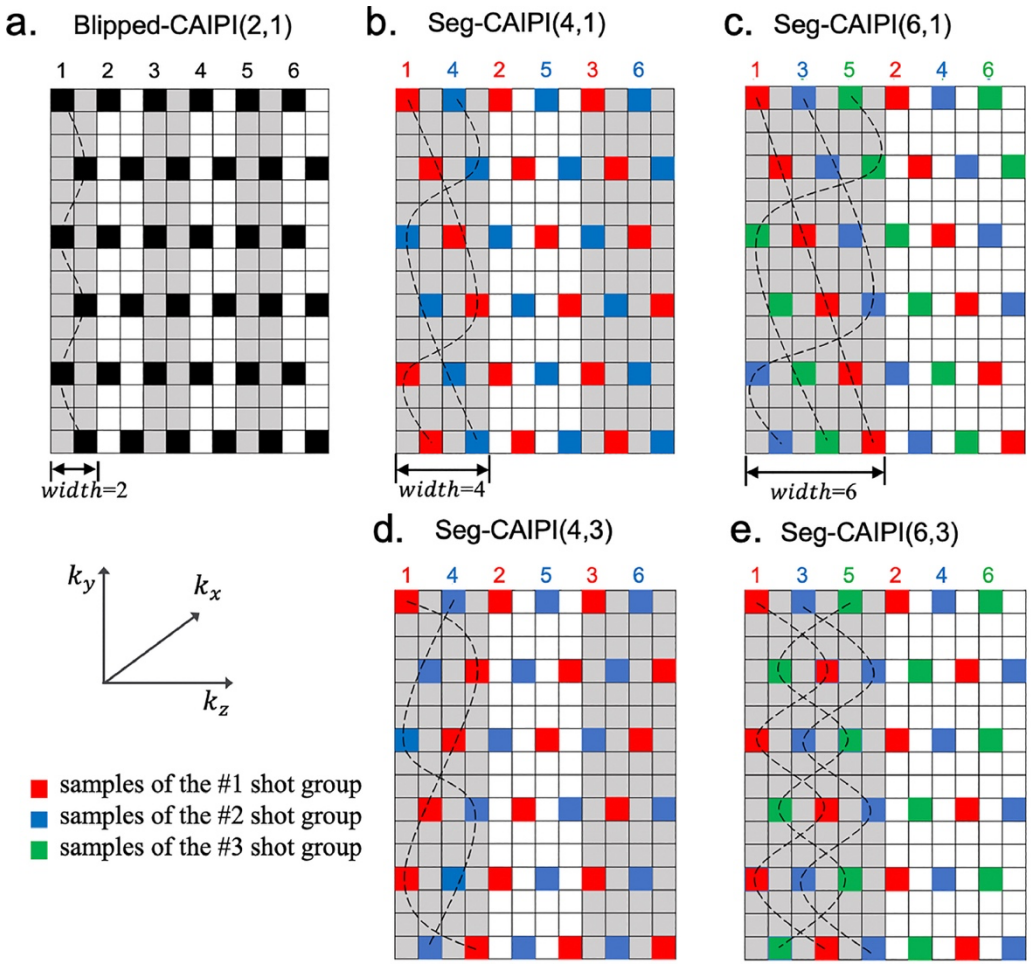

Fig. 4.1 Schematic 3D multi-shot EPI sampling patterns. a) Standard blipped-CAIPI. b)-e) The proposed seg-CAIPI scheme with varying widths and Δk_z blips. The sampling patterns are shown in k_y-k_z plane and each sample represents a readout line. All the samples connected by a dashed line correspond to a single shot. Different shot groups are marked in different colours. The under-sampling factor $R_y \times R_z$ is 3×2 in this case. For seg-CAIPI sampling patterns with width $= 4$ in b) and d), 2 shot groups are used. For seg-CAIPI sampling patterns with width $= 6$ in c) and e), 3 shot groups are used.

4.2.2 Structured Low-rank Constrained Reconstruction

Conventional reconstruction methods combine all acquired shots in k-space directly regardless of the shot-to-shot inconsistencies. To alleviate the intra-volume inconsistencies, an individual image for each shot group which consists of fewer shots is reconstructed. The images of all shot groups are then SOS combined to avoid phase cancellation effects. Similar to the method MUSSELS (Mani et al., 2017), the missing k-

space data of all shot groups are jointly recovered by exploiting the linear dependencies among them, assuming the images for each shot group have the same magnitude and differ only by the physiologically induced phase modulations. To enable this, the intra-shot group variations are ignored, as respiratory phase effects (~3-5s period) are expected to be relatively temporally coherent with respect to the shot TR time scale. This redundancy in multi-shot group k-space can be expressed as the low rank property of its block-Hankel structured matrix representation, as shown in Section 3.4.2. Note the structured low-rank matrix formulation also leverages all other linear inter-dependencies such as those arising from limited image support (Haldar, 2014; Haldar and Zhuo, 2016) simultaneously. Fig. 4.2 shows the process of shot binning and one way to construct a block-Hankel structured matrix from the multi-shot group k-space. The image reconstruction is formulated as a structured low-rank constrained parallel imaging optimisation problem:

$$\min 1/2\|EX - Y\|_2^2 + \lambda||H_s X||_* \qquad \text{Eq. 4.1}$$

Where X is the k-space of all shot groups to be reconstructed and Y is the measured multi-channel, multi-shot group k-space data. The operator E performs the composition of $AFSF^H$, where F and F^H are the Fourier and inverse Fourier transform respectively. S denotes the multi-coil sensitivity encoding operator, and A selects the sampled k-space locations for each shot group. The operator H_s applied to X generates the block-Hankel structured matrix of X. $||H_s X||_*$ denotes the nuclear norm of $H_s X$, and this convex constraint is used to promote the low-rank structure in this work. λ is the regularisation parameter which weights the low-rank constraint. Eq. 4.1 is solved using the ADMM (Boyd, 2010) algorithm. As discussed in Section 2.3.2, one step of the ADMM algorithm performs soft thresholding on the singular values. In this work, a hard thresholding is used instead, which has shown better performance in this problem especially when the

tSNR is low.

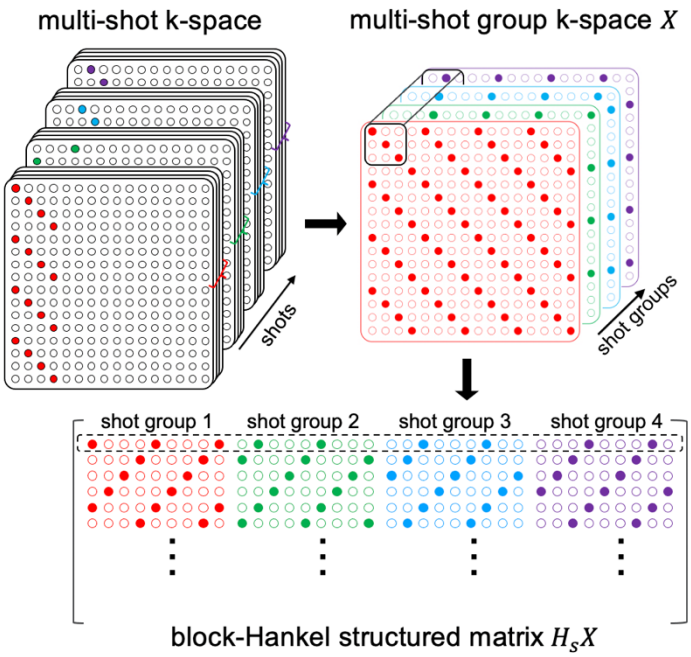

Fig. 4.2 Illustration of the shot-binning on multi-shot k-space and the block-Hankel structured matrix construction on multi-shot group k-space. A small number of shots acquired consecutively are binned together as a shot group. X denotes the multi-shot group k-space data and H_sX corresponds to the block-Hankel structured matrix representation of X. A small patch of X selected by a sliding window in each shot group is vectorised to generate a row of the block-Hankel structured matrix.

4.3 Experiments

Simulation and in vivo experiments at 7T were performed to validate the performance of the proposed method. All in vivo data were collected on a Siemens Magnetom 7T scanner (Siemens Healthineers, Erlangen, Germany) equipped with a 32-channel head-only receive coil with single-channel transmit (Nova Medical, Wilmington, MA, USA). All subjects were scanned with informed consent under a technical development protocol approved by the local ethics committee. Under-sampling was applied along two dimensions with the total acceleration factor $R = R_y \times R_z$, where R_y and R_z denote the acceleration along the primary phase encoding dimension and the shot dimension respectively. The reconstruction was implemented in MATLAB R2020a (MathWorks,

Inc.). Coil sensitivity maps were calculated using ESPIRiT implemented in the BART toolbox (Uecker et al., 2014). The Nyquist ghosting was corrected using the 3-line reference scan averaged cross all the shots to perform a conventional linear phase correction. The number of shot groups used for the proposed reconstruction was set to be the same as the number of interleaves (width/R_z), to achieve roughly uniform sampling for each shot group. The kernel size for the block-Hankel transformation was chosen empirically to be 6×6. Reconstruction hyperparameters were chosen empirically to maximize the resulting tSNR of a single sagittal slice and then used for all slices. $\lambda = 3E - 3$ was used for the simulation datasets and $\lambda = 3E - 4$ or $6E - 4$ were used for the in vivo datasets. Since image artefacts arising from shot-to-shot inconsistencies do not occur in the readout direction, an inverse Fourier transform was performed along the fully sampled readout direction, followed by separate reconstruction of each 2D k_y-k_z slice, which reduced the computational burden.

The performance of the proposed method which uses SLR reconstruction on the seg-CAIPI sampling data was evaluated by comparing to the conventional method which uses SENSE (Pruessmann et al., 2001) reconstruction on the shot-combined data acquired by the standard blipped-CAIPI sampling, as well as a 2D SMS-EPI using the CMRR multi-band EPI sequence and online reconstruction (Moeller et al., 2010). Note the online reconstruction of the CMRR SMS-EPI sequence may yield better tSNR or z-statistic values due to the optimised image reconstruction and data processing pipeline.

4.3.1 Simulation Experiments

Multi-shot k-space datasets consisting of a single k_y-k_z plane were synthesized by modulating a ground truth image with a series of physiologically plausible phase variation maps, which were measured from a 2D EPI time series acquired in sagittal orientation at

7T with 1.5mm isotropic resolution and TE/TR=20/40ms. The measured phase variation maps were then fit to a second order spherical harmonic basis set for denoising. The k-space data for each shot was generated by resampling the 2D k-space time-series for both the standard trajectory and the seg-CAIPI trajectory. The matrix size of the simulation data is 140×96, and a fully sampled acquisition consists of 96 shots.

The impact of thermal noise and acceleration factor, both having an influence on the thermal to physiological noise ratio and thus the performance of the proposed method was examined. Three different levels of complex, zero-mean Gaussian noise (low/medium/high) were added to the above k-space data respectively. The physiological noise-free datasets with the same amount of Gaussian noise as the phase variations corrupted datasets were also generated as the reference datasets, to indicate the tSNR upper bounds of the datasets with phase variations. The mean SNRs of these reference datasets are approximately 97/65/43 for the low/medium/high thermal noise levels respectively. The performance of the proposed method using seg-CAIPI(6,1) sampling was validated across different thermal noise levels in the fully sampled regime. The performance of the proposed method at varying acceleration factors were also investigated at the medium thermal noise level, where seg-CAIPI(6,2) was used for the proposed method at $R = 1 \times 3$ and 3×1, and blipped-CAIPI(3,2) was used for the conventional method at $R = 1 \times 3$. The results are shown in Fig. 4.3. Fig. 4.3a compares the mean tSNR of the conventional and proposed methods at low/medium/high thermal noise levels in a fully sampled regime. The proposed method achieves a significantly higher tSNR than the conventional method in all 3 cases, which approaches the tSNR of the reference data with thermal noise only. As expected, the relative tSNR differences between different methods and the reference data reduce when thermal noise increases. Fig. 4.3b compares the mean tSNR of the conventional and proposed methods at different

acceleration factors at the same medium thermal noise level. Similar to Fig. 4.3a, the proposed method achieves a higher tSNR than the conventional method in all 3 cases, while the tSNR differences between different methods get smaller at higher acceleration factors, which lead to higher thermal noise amplification.

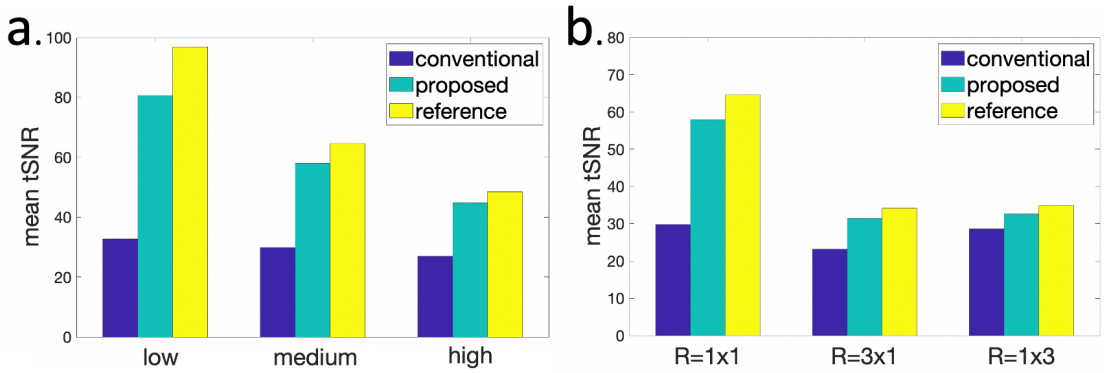

Fig. 4.3 The reconstruction results of the simulation data at different thermal noise levels and different acceleration factors. The dataset without inter-shot phase variations is used as the reference. a) Mean tSNR of different methods at varying thermal noise levels in the fully sampled regime. b) Mean tSNR of different methods at varying acceleration factors at the medium thermal noise level.

Fig. 4.4 shows the comparison between the proposed method with seg-CAIPI(8,2) sampling and conventional method with blipped-CAIPI(4,2) sampling at $R = 2 \times 4$ in a thermal noise-free regime, where the temporal fluctuations of the time course are induced by phase variations only. As shown in Fig. 4.4a, the temporal mean images of these two methods are comparable, and the proposed method achieves a tSNR improvement by reducing the temporal standard deviation. Fig. 4.4b shows the mean power spectrum of the reconstructed time series across the masked brain. Compared to the physiological noise-free reference data, the conventional 3D multi-shot EPI method has higher noise level across a broad frequency range, and the proposed method reduces the apparent broadband physiologically induced noise.

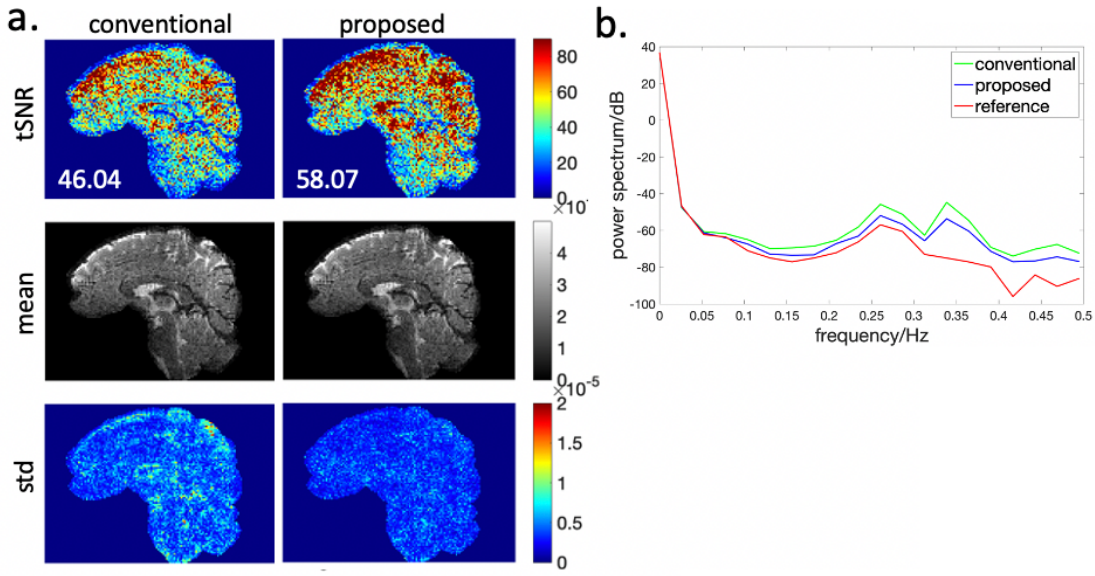

Fig. 4.4 The reconstruction results of the simulation data at $R = 2 \times 4$ in a thermal noise-free regime. The dataset without inter-shot phase variations is used as the reference. a) Maps of tSNR (mean tSNR shown on the bottom left), temporal mean magnitude and temporal standard deviation of the conventional method (left) and the proposed method (right). b) The mean power spectrum across the masked brain of different methods.

Fig. 4.5 shows a simulation experiment where the seg-CAIPI sampling was compared with both conventional blipped-CAIPI using a sequential ordering along k_z (standard blipped-CAIPI, referred to as "conventional-sequential" in this case) and using an interleaved ordering along k_z (referred to as "conventional-interleaved"). Despite the same interleaved ordering, the conventional-interleaved sampling and seg-CAIPI sampling has different spans of each shot along k_z. These datasets were simulated at $R = 2 \times 2$ and medium thermal noise level. The conventional-sequential and conventional-interleaved samplings both used a blip size $\Delta k_z = 1$ and a k_z-span of 2. The seg-CAIPI sampling used $\Delta k_z = 3$ and a k_z-span of 8. The proposed reconstruction used 4 shot groups with 12 shots per shot group for all these three trajectories. Fig. 4.5 shows the sampling trajectory for the first shot and the first shot group of these three sampling patterns. Fig. 4.5b shows the temporal mean images of SENSE and the proposed reconstruction for these three sampling patterns respectively. Both the conventional-interleaved and seg-CAIPI have a lower tSNR than the conventional-sequential sampling

with a shot-combined SENSE reconstruction. This is probably because the phase inconsistencies are mapped coherently with a sequential sampling along k_z, whereas interleaved sampling leads to incoherent mapping of phase variations, which can result in more artefacts if a simple shot-combined reconstruction is used. In contrast, the proposed method which reconstructs an individual image for each shot group benefits from the interleaved ordering and a more uniform under-sampling pattern for each shot group, which is demonstrated by the highest tSNR achieved by the seg-CAIPI sampling compared to other two sampling patterns.

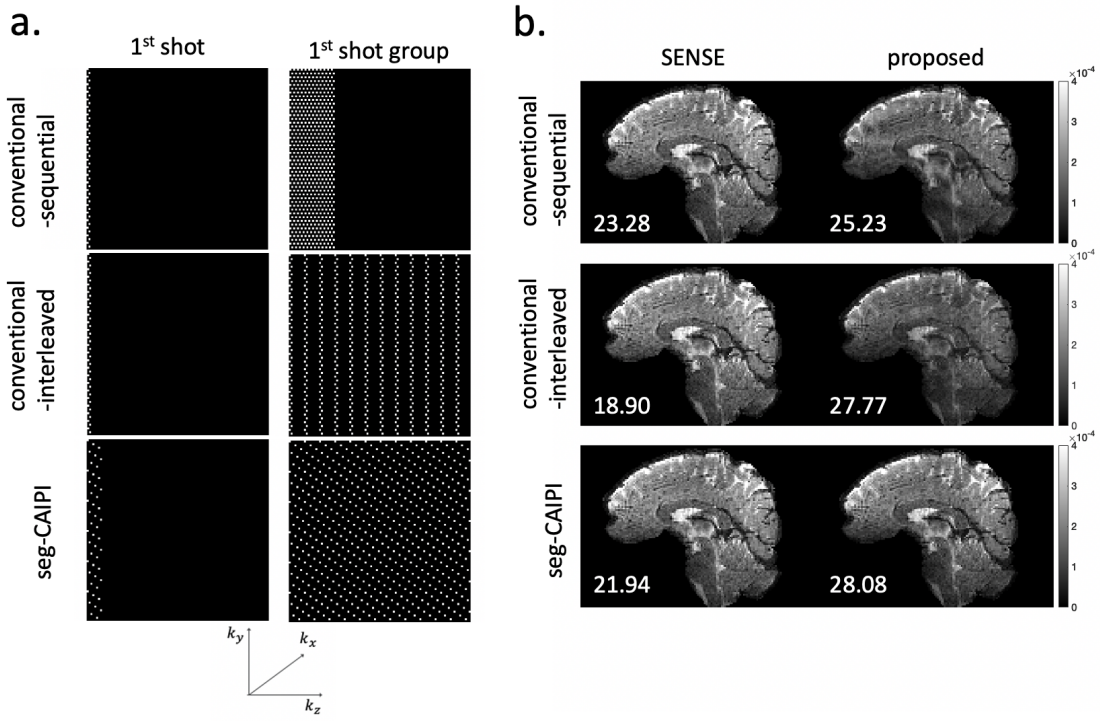

Fig. 4.5 The comparison of three different sampling patterns and their reconstruction performance on the simulation data. The three sampling patterns are: 1) seg-CAIPI, 2) the conventional blipped-CAIPI using a sequential ordering along k_z ("conventional-sequential"), and 3) the conventional blipped-CAIPI using an interleaved ordering along k_z ("conventional-interleaved"). a) The trajectory of the first shot and first shot group. b) The temporal mean magnitude images of SENSE and the proposed reconstruction. The mean tSNR is shown on the bottom left for each temporal magnitude image.

Another simulation experiment is performed to validate the ability of the proposed method to recovery phase variations. Two types of datasets were generated. One of them

only has inter-shot group phase variations, i.e., the image phase is consistent within each shot group, so the ground truth phase for each shot group is available. Another dataset is subject to shot-to-shot phase variations as in previous simulations, so the ground truth phase of each shot group is not available. These datasets were simulated at $R = 2 \times 2$ and medium thermal noise level. The seg-CAIPI(8,3) sampling pattern was used for the proposed method. Fig. 4.6 shows the reconstruction results of this experiment, where the phase variation maps of different shot groups calculated as the relative phase differences with reference to the first shot group (in radian) are shown. Fig. 4.6a shows the results of the dataset with inter-shot group phase variations only, highlighting good correspondence between the recovered phase variations and the ground truth phase modulations. Fig. 4.6b shows the results of the dataset with shot-to-shot phase variations. In this case, distinct "average" phase modulations per shot-group are still recovered.

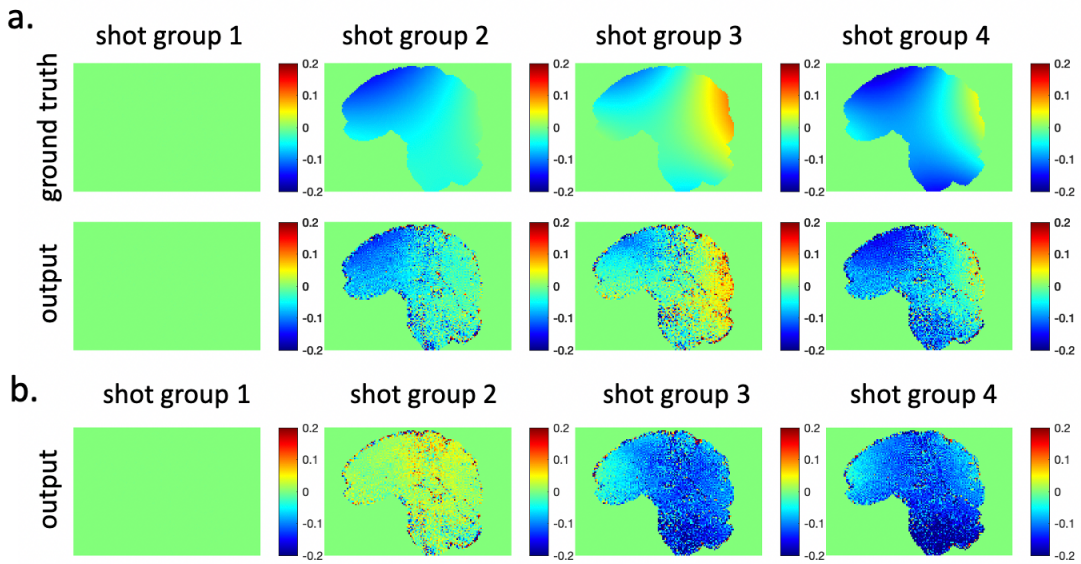

Fig. 4.6 The phase variation maps of different shot groups reconstructed by the proposed method on two different simulation datasets. a) a simulation dataset with inter-shot group phase variations only. b) a simulation dataset with shot-to-shot phase variations. The phase variation maps are shown as the phase differences with reference to the first shot group. In subfigure a), the phase variation maps of the ground truth (first row) and the reconstruction output (second row) are shown. In subfigure b), the phase variation maps of the reconstruction output are shown (ground truth not available).

4.3.2 In Vivo Experiments

Resting-state and task-based fMRI datasets were acquired on seven healthy volunteers to validate the performance of the proposed method. The performance of different methods were compared at varying spatial resolutions (1/1.2/1.5/1.8 mm isotropic resolution) and acceleration factors where the ratio between thermal noise and physiological noise changes. The impact of width used for the proposed method which trades off between intra-shot group consistency and the data available for each shot group was also investigated. Each resting-state dataset consists of 40 or 60 volumes. A task-based fMRI experiment was performed on six of the subjects to assess the impact of the proposed method on BOLD activation. Each subject was scanned twice with both the proposed 3D seg-CAIPI EPI and 2D SMS-EPI acquisitions in a counterbalanced order, and each scan was approximately 4 minutes with 170 volumes. The experiment used a 30/30s off/on block design, and the subjects were instructed to perform a finger tapping task with both hands when a 10 Hz black and white flashing checkerboard was shown every 30 seconds. The functional images were processed using FEAT (Jenkinson et al., 2012; Woolrich et al., 2001), and minimal data pre-processing was used which includes high pass filtering and motion correction by MCFLIRT (Jenkinson et al., 2002).

All in vivo scanning protocols are shown in Table 4.1. Slice orientation was transverse with phase encoding along the anterior-posterior direction. A slab-selective sinc excitation was used without slice oversampling. All standard, seg-CAIPI trajectory data and SMS-EPI data used for comparison were acquired with matched TE/TR/bandwidth. Ernst angles for 3D EPI and SMS-EPI were used respectively. For a given width, the Δk_z blip size was chosen following the principle of minimising the largest distance between two samples.

iso. res. (mm^3)	matrix size	BW (Hz/pixel)	R_y	R_z/ MB factor	TE (ms)	volume TR (ms)	shot TR (ms)	flip angle	FOV shift	echo spacing (ms)	sampling	# shots/shot group
1.8	116 ×116 ×96	2052	2/4	2	23	2640	55	15	-	0.57/0.59	blipped-CAIPI(2,1)	48
										0.65	seg-CAIPI(8,3)	12
		2156	2	4	23	1320	55	15	-	0.59	blipped-CAIPI(4,2)	24
										0.63	seg-CAIPI(8,2)	12
							-	64	FOV/4	0.61	SMS-EPI	-
1.5	140 ×140 ×96	1880	3	2	24	2592	54	15	-	0.64	blipped-CAIPI(2,1)	48
										0.72	seg-CAIPI(8,3)	12
							-	78	FOV/6	0.68	SMS-EPI	-
		1880	3	3	24	1728	54	15	-	0.66	blipped-CAIPI(3,2)	32
										0.7	seg-CAIPI(6,2)	16
							-	70	FOV/12	0.68	SMS-EPI	-
1.2	174 ×174 ×120	1596	3	2	29	3900	65	16	-	0.73	blipped-CAIPI(2,1)	60
										0.75-0.81	seg-CAIPI(4,1), (6,3), (8,3)	30/20/15
							-	85	FOV/6	0.77	SMS-EPI	-
1	210 ×210 ×120	1254	4	2	32	4080	68	15	-	1	blipped-CAIPI(2,1)	60
										1	seg-CAIPI(6,3)	20

Table 4.1 The scanning protocols of the in vivo experiments.

Fig. 4.7 shows the results of an experiment where a) the standard blipped-CAIPI(2,1) data and b) the seg-CAIPI(8,3) data were both reconstructed by the same SENSE and the proposed reconstruction. These datasets were acquired at 1.8mm isotropic resolution and resting-state with $R = 4 \times 2$. As shown in Fig. 4.7a, the proposed reconstruction achieves a higher tSNR than SENSE reconstruction on the same standard sampling data, but its fidelity of the temporal mean image is much lower. In contrast, as shown in Fig. 4.7b, the proposed reconstruction on the seg-CAIPI data achieves both a higher tSNR and high-fidelity temporal mean image compared to SENSE reconstruction. It is demonstrated that data acquired with the standard trajectory are not compatible with the proposed reconstruction, due to the clustered nature of sequentially acquired shots, and thus the proposed method was compared to the conventional method which uses SENSE reconstruction on the standard sampling data in most of this work.

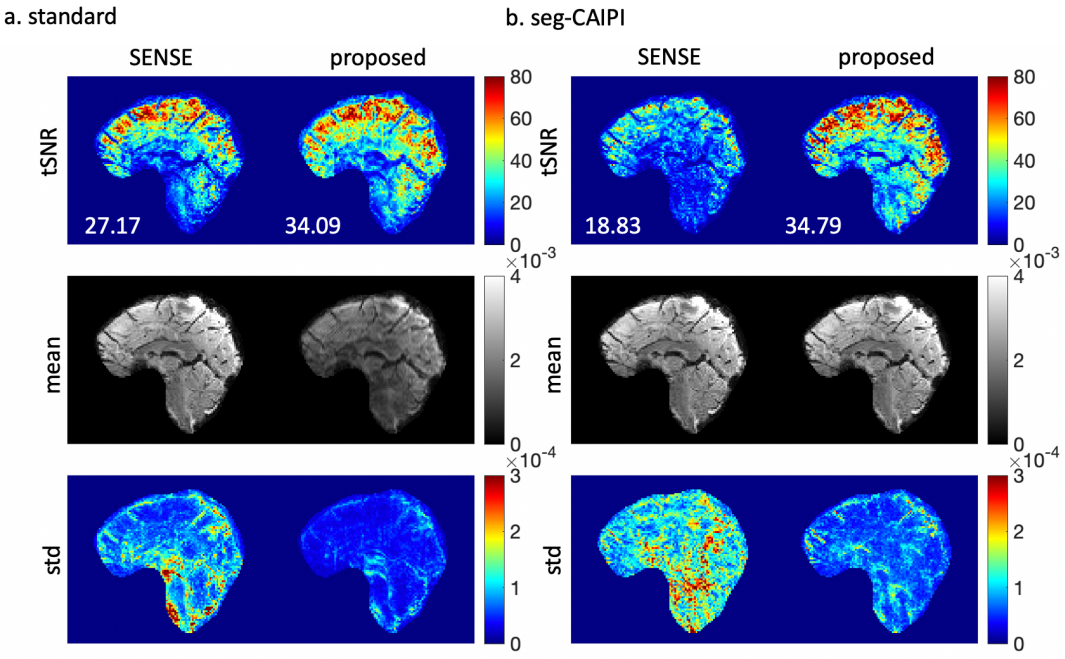

Fig. 4.7 The comparison of the same SENSE and SLR reconstructions on in vivo resting-state data acquired with different sampling trajectories. a) the standard blipped-CAIPI(2,1) trajectory. b) the seg-CAIPI(8,3) trajectory. The mean tSNR calculated for this 2D slice is shown on the bottom left for each tSNR map.

Fig. 4.8 shows the results of an experiment investigating the impact of the regularisation parameter λ on the proposed reconstruction at different acceleration factors, $R = 4 \times 2$ and $R = 2 \times 2$. The seg-CAIPI(8,3) datasets were acquired at 1.8mm isotropic resolution and resting-state. $\lambda = 3E - 4$ is considered a good choice in this case, whereas $\lambda = 3E - 3$ leads to over-regularisation and a lower tSNR. At the higher acceleration factor $R = 4 \times 2$, the lower tSNR of the reconstruction using $\lambda = 3E - 3$ is mostly due its reduced image magnitude. At the lower acceleration factor $R = 2 \times 2$, the lower tSNR of the reconstruction using $\lambda = 3E - 3$ is due to its higher temporal standard deviation.

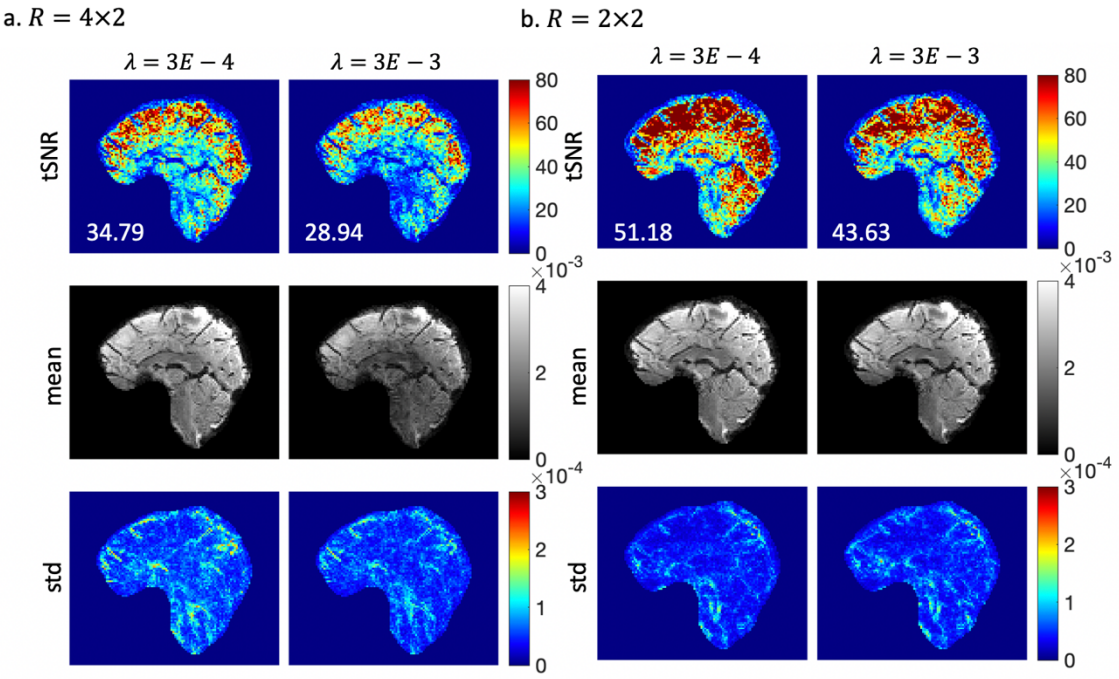

Fig. 4.8 The impact of regularisation parameter λ on the proposed reconstruction at different acceleration factors. a) $R = 4 \times 2$ and b) $R = 2 \times 2$. The mean tSNR calculated for this 2D slice is shown on the bottom left for each tSNR map.

Fig. 4.9 and Fig. 4.10 show the impact of width on the performance of the proposed method on resting-state in vivo datasets across different spatial resolutions with the same acceleration factor, and across different acceleration factors at the same resolution respectively. Fig. 4.9 shows the reconstruction results of a) a 1.2mm and b) a 1.5mm isotropic resolution in vivo datasets acquired at $R = 3 \times 2$. At 1.2mm isotropic resolution, width $= 6$ (3 shot groups) and width $= 8$ (4 shot groups) have comparable performance, whereas width $= 8$ achieves a significantly higher tSNR than width $= 6$ at 1.5mm isotropic resolution. This is probably because the reconstruction is better conditioned at the lower spatial resolution, and thus having more shot groups at the cost of less data per shot group could have more benefit. Fig. 4.10 shows the reconstruction results of two sets of in vivo data at 1.8mm isotropic resolution acquired at a) $R = 2 \times 2$ and b) $R = 4 \times 2$. At $R = 4 \times 2$, width $= 8$ (4 shot groups) achieves a slightly higher tSNR than width $= 4$ (2 shot groups), whereas width $= 8$ achieves a significantly

higher tSNR than width $= 4$ at $R = 2 \times 2$. Similarly, this is probably because the reconstruction is better conditioned at the lower acceleration factor, and thus having more shot groups at the cost of less data per shot group could have more benefit.

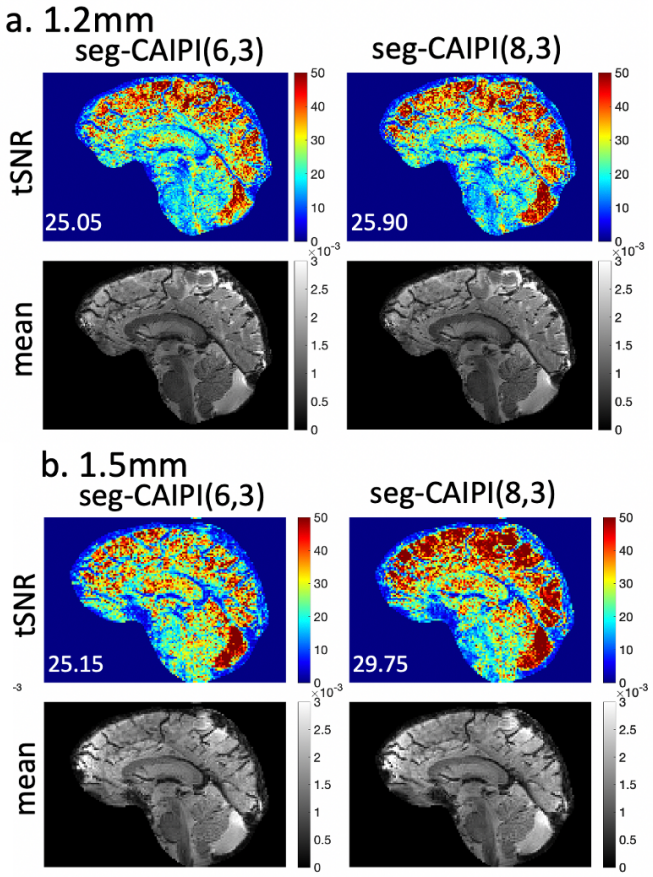

Fig. 4.9 The impact of width on the proposed method at different spatial resolutions. a)1.2mm isotropic resolution. b) 1.5mm isotropic resolution. All datasets were acquire at $R = 3 \times 2$. The mean tSNR calculated for this 2D slice is shown on the bottom left for each tSNR map.

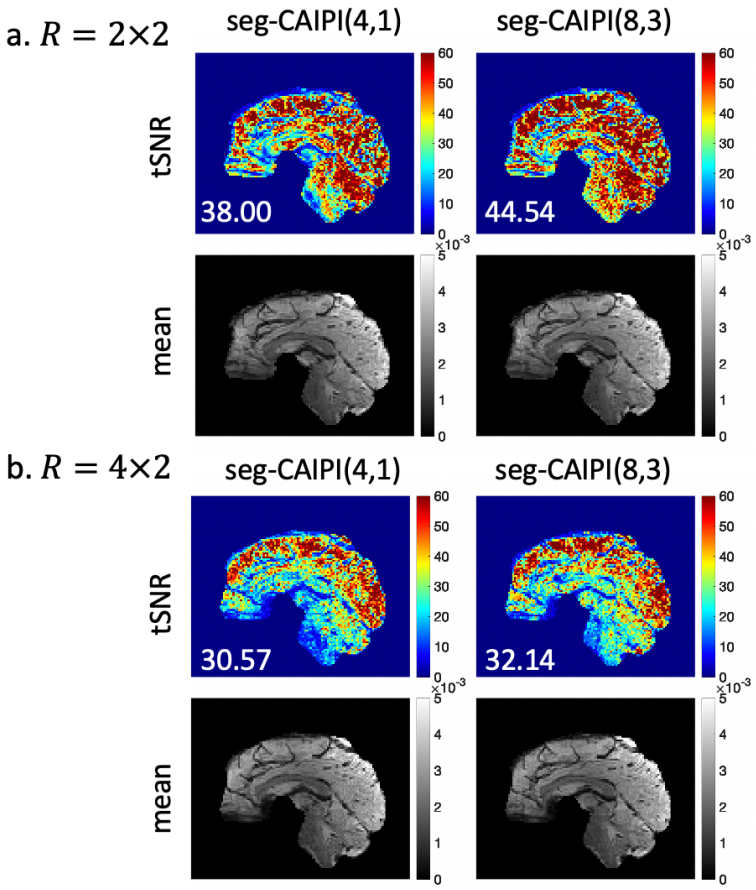

Fig. 4.10 The impact of width on the proposed method at different acceleration factors. a) $= 2 \times 2$. b) $= 4 \times 2$. All datasets were acquired at 1.8mm isotropic resolution. The mean tSNR calculated for this 2D slice is shown on the bottom left for each tSNR map.

Fig. 4.11 shows the comparison between the conventional and the proposed 3D methods in the in vivo resting-state experiment at 1.8mm isotropic resolution with acceleration factors $R = 2 \times 2$ and $R = 4 \times 2$. In both cases, the proposed method achieves considerably higher tSNR than the conventional 3D method. Note the tSNR of conventional method does not improve significantly as R_y decreases, which indicates that the higher SNR at $R = 2 \times 2$ was negated by the higher physiological noise compared to $R = 4 \times 2$. Accordingly, the proposed 3D method achieves a larger tSNR gain at $R = 2 \times 2$, resulting in considerably higher tSNR at the lower R_y. It is also worth noting that the proposed 3D method achieves even higher tSNR at $R = 4 \times 2$ than the conventional method at $R = 2 \times 2$.

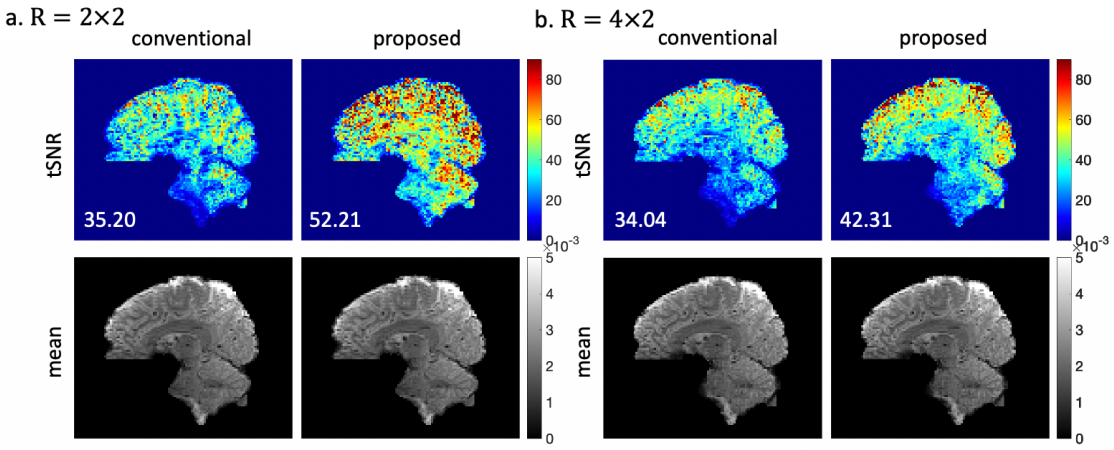

Fig. 4.11 The reconstruction results of the 1.8mm isotropic resolution resting-state datasets acquired at a) $R = 2 \times 2$ and b) $R = 4 \times 2$. The proposed method used seg-CAIPI(8,3) sampling trajectory. Mean tSNR across the whole brain is shown on the bottom left for each tSNR map.

Table 4.2 shows the comparison of mean tSNR between the conventional, proposed 3D methods as well as 2D SMS-EPI imaging in the in vivo resting-state experiment at 1.8mm isotropic resolution with $R_y = 2$ and R_z/MB factor $= 4$ for 5 subjects. In this case, the difference in tSNR between the conventional 3D method and SMS-EPI is small and the proposed 3D method only achieves a small improvement over the conventional approach. This tSNR improvement is smaller compared to the tSNR improvement at $R = 2 \times 2$ shown above, which is likely because fewer shots per volume can result in lower physiological noise and higher thermal noise. As shown in Fig. 4.11 and Table 4.2, the tSNR gain of the proposed method varies depending on the acceleration factor R_y and R_z, despite the same spatial resolution. This is likely because the acceleration factor alters physiological noise by influencing the sensitivity to physiological variations as well as thermal noise by changing the g-factor.

subject #	conventional 3D	proposed 3D	SMS-EPI
1	32.29	35.48	**38.36**
2	34.51	37.37	**38.62**
3	30.16	**34.93**	33.11
4	**31.19**	28.56	29.71
5	**36.19**	35.64	35.02
average	32.87	34.40	**34.96**

Table 4.2 The mean tSNR of different methods in the 1.8mm isotropic resolution resting-state experiment of 5 subjects. The proposed method used seg-CAIPI(8,2) sampling trajectory. The acceleration factor is $R = 2 \times 4$.

Table 4.3 compares the max and mean z-statistics between the proposed 3D and SMS-EPI methods in the task fMRI experiment at 1.8 mm isotropic resolution with $R = 2 \times 4$ of 6 subjects. The ROI was independently defined over motor and visual cortices for each subject. Fig. 4.12 shows the activation z-statistic maps of subject 3. The proposed 3D method and SMS-EPI show comparable max and mean activations in this case.

subject #	max z stats		mean z stats	
	proposed	SMS-EPI	proposed	SMS-EPI
1	**23.09**	21.70	**6.56**	6.01
2	**18.52**	14.86	3.84	**4.02**
3	20.19	**20.70**	5.66	**6.68**
4	17.39	**19.32**	3.43	**5.77**
5	17.46	**21.63**	4.59	**7.49**
6	**22.33**	22.09	7.71	**8.27**

Table 4.3 The max z-statistic and mean z-statistic within the selected ROI of different methods in the 1.8mm isotropic resolution task fMRI experiment of 6 subjects. The proposed method used seg-CAIPI(8,2) sampling trajectory. The acceleration factor is $R = 2 \times 4$.

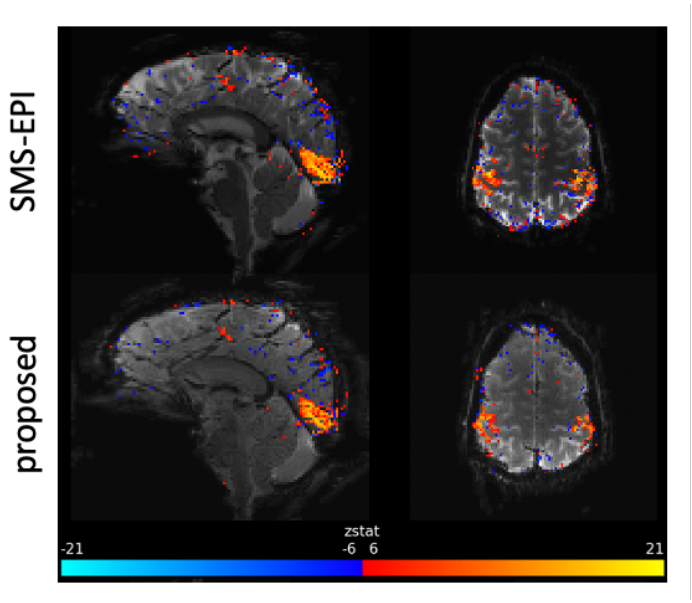

Fig. 4.12 Representative activation maps of the 1.8 mm isotropic resolution task fMRI datasets of subject 3 in Table 4.3.

Fig. 4.13 shows the comparison between different methods in the in vivo resting-state experiment at 1.5mm isotropic resolution with acceleration factors $R = 3 \times 2$ and $R = 3 \times 3$. When $R_z = 2$, the proposed 3D method achieves a significant tSNR improvement compared to the conventional 3D method, which enables the 3D imaging method to have higher tSNR than SMS-EPI. When $R_z = 3$, the proposed method achieves a moderate tSNR improvement compared to the conventional 3D method, which is comparable to SMS-EPI in this regime. Similar to the results shown in Fig. 4.11, the tSNR of conventional method does not improve as R_z decreases, which is likely again due to the limiting impact of physiological noise. In contrast, the proposed 3D method achieves a significantly higher tSNR when R_z is reduced, demonstrating that the impact of physiological noise has been greatly reduced at $R = 3 \times 2$.

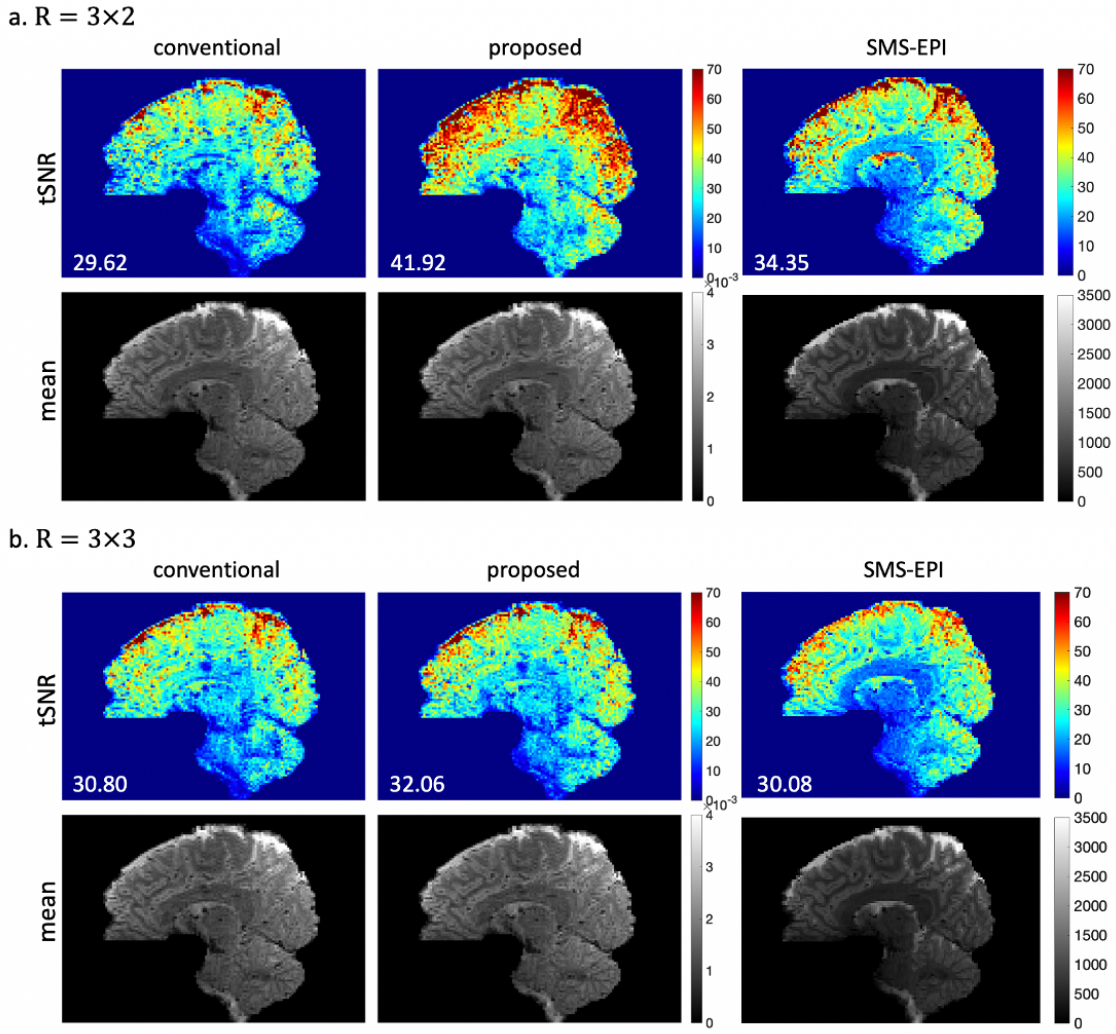

Fig. 4.13 The reconstruction results of the 1.5mm isotropic resolution resting-state datasets acquired at a) $R = 3 \times 2$ and b) $R = 3 \times 3$. The proposed method used seg-CAIPI(8,3) at $R = 3 \times 2$ and seg-CAIPI(6,2) at $R = 3 \times 3$. Mean tSNR across the whole brain is shown on the bottom left for each tSNR map.

Fig. 4.14 shows the comparison between different methods in the in vivo resting-state experiment at 1.2mm isotropic resolution with acceleration factor $R = 3 \times 2$. The proposed method achieves a mean tSNR comparable to SMS-EPI, which is notably higher than the conventional method.

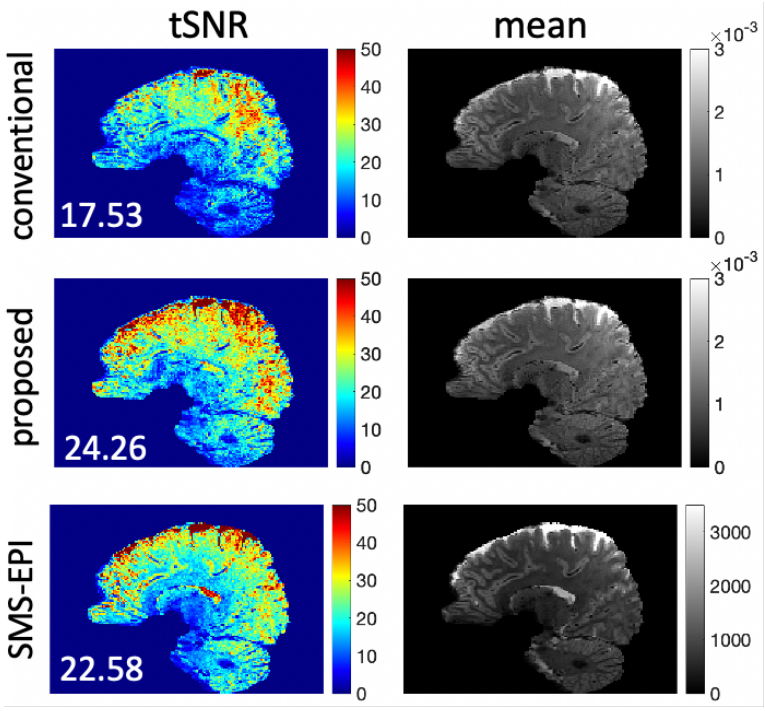

Fig. 4.14 The reconstruction results of the 1.2mm isotropic resolution resting-state datasets. The proposed method used seg-CAIPI(4,1) sampling trajectory. The acceleration factor is $R = 3 \times 2$. Mean tSNR across the whole brain is shown on the bottom left for each tSNR map.

Fig. 4.15 shows the reconstruction results of the 1mm isotropic resolution in vivo resting-state datasets. The proposed method achieves a moderate tSNR improvement compared to the conventional method. At this high resolution, tSNR values are relatively low, and the improvement over the conventional 3D method is limited by the thermal noise dominance. However, there does appear to be a localized improvement in tSNR in the cerebellum, for example, consistent with greater impact of respiratory phase fluctuations nearer to the lung cavity. Fig. 4.16 shows the comparison of the proposed reconstruction using soft thresholding and hard thresholding on the 1mm isotropic resolution dataset. Despite comparable tSNR, hard thresholding approach achieves a notably better image quality in this case.

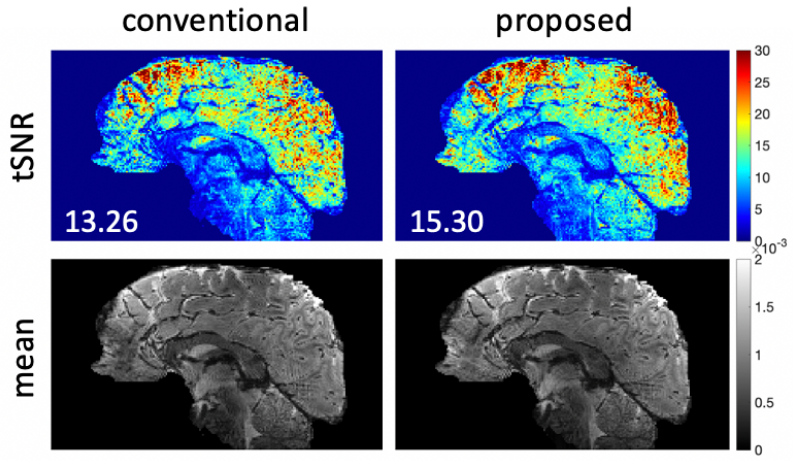

Fig. 4.15 The reconstruction results of the 1mm isotropic resolution resting-state datasets. The proposed method used seg-CAIPI(6,3) sampling trajectory. The acceleration factor is $R = 4 \times 2$. Mean tSNR across the whole brain is shown on the bottom left for each tSNR map.

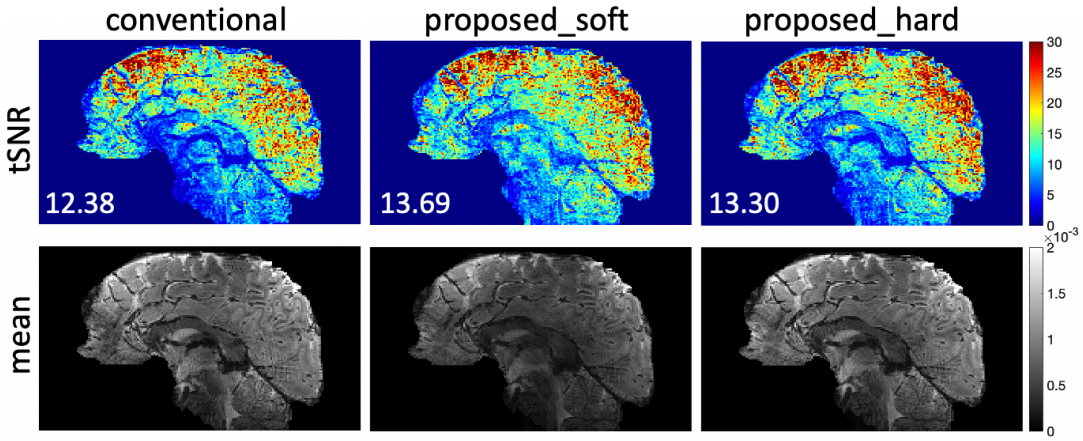

Fig. 4.16 The comparison of the proposed reconstruction using soft thresholding and hard thresholding on the 1mm isotropic resolution resting-state dataset. The mean tSNR calculated for this 2D slice is shown on the bottom left for each tSNR map.

4.4 Discussion

In this chapter, a method to improve the robustness of 3D multi-shot EPI acquisition to physiological fluctuations is presented for fMRI at 7T, which incorporates the use of the 3D seg-CAIPI sampling scheme and a reconstruction method based on structured low-rank matrix completion. In vivo experiments have demonstrated that the proposed 3D method outperforms the conventional 3D method across a wide range of isotropic resolutions from 1mm to 1.8mm and varying acceleration factors. The tSNR

improvement gained by the proposed 3D method is particularly significant where the conventional 3D method suffers from a great tSNR loss compared to 2D SMS-EPI imaging, suggesting the high impact of physiological noise. Now with the proposed method, 3D EPI imaging can achieve a tSNR comparable or higher than 2D SMS-EPI. Thus, the proposed method could increase the utility of 3D EPI in regimes currently dominated by 2D multi-slice methods.

It is important to note that the proposed reconstruction operates on each image volume independently, and only jointly reconstructs shot groups corresponding to a single volume. Hence, it reduces the vulnerability to physiological variations within the acquisition window of each volume, while retaining full temporal DOF across different volumes and leaving the BOLD-related signal fluctuations and temporal characteristics intact. Besides, it deals with the inter-shot phase variations without requiring any information or external measurements of the physiological traces. Moreover, although images of all shot groups are SOS combined in the proposed method, other image combination approaches which preserve the phase information in the final output image can also be considered for applications where complex data are necessary.

The performance of the proposed method in terms of its ability to boost tSNR is related to thermal to physiological noise ratio, which is difficult to characterise and can be affected by many factors jointly including spatial resolution and acceleration factor. In general, the proposed method is more beneficial in higher SNR regimes e.g., either low spatial resolution or low acceleration factor given that other conditions are the same, where the SNR benefit obtained by higher signal strength or more data available could be negated by higher physiological noise. Although the tSNR gain varies in different sampling protocols, the proposed method is beneficial in most cases where the inter-shot phase variations play a non-trivial role in limiting the tSNR. The proposed method could

also be less effective as the conditioning gets worse at very high acceleration factor.

The seg-CAIPI trajectory is a combination of interleaved sampling along the shot dimension with CAIPI blipping. Similar to other 3D EPI blipped-CAIPI schemes (Narsude et al., 2016; Stirnberg and Stöcker, 2021; Wang et al., 2019), the seg-CAIPI sampling acquires 3D k-space using a series of blipped, band-limited (in k_z) readouts, except that it returns to acquire different samples in the same "band" in different interleaves. Also, each shot of the seg-CAIPI sampling covers a wider band along k_z compared to its conventional alternatives, as this sampling pattern aims to optimise the sampling trajectory of each shot group in addition to the overall shot-combined sampling trajectory. This seg-CAIPI sampling pattern is used to facilitate the joint reconstruction of images for each shot group, instead of a single shot-combined image that suffers from phase cancellation, which is the key to the proposed reconstruction. While significantly better performance of the proposed reconstruction is observed on the seg-CAIPI data compared to the conventional blipped-CAIPI data with sequential or interleaved ordering, the optimal sampling for structured low-rank matrix recovery is still a topic of study (Haldar, 2014). Therefore, alternative 3D CAIPI sampling with different shot widths and shot interleaving schemes could be explored as well. It is worth mentioning that although the seg-CAIPI sampling pattern benefits the proposed reconstruction, it may not be optimal for conventional shot-combined SENSE reconstructions. The seg-CAIPI sampling achieves lower tSNR than conventional sampling by the same SENSE reconstruction in most cases, which is probably because the interleaved sampling scheme leads to greater signal variations between adjacent k-space samples that the signal variations are mapped incoherently across k-space.

The temporal scale of physiological variations is reduced from the acquisition window of each volume to each shot group by reconstructing an individual image for each shot

group. Thus, there is a trade-off between the number of shot groups and the amount of data available for each group. Increasing the number of shot groups can enhance the intra-shot group consistency, but may make the reconstruction more challenging due to the reduced number of samples for each shot group. In comparison, decreasing the number of shot groups might result in strong phase variations within the shot group that impairs the efficacy of the proposed reconstruction. Note the minimal TE achievable is limited by the largest k_z rewind blip of the seg-CAIPI acquisition, which could be another consideration limiting high width parameters and thus the number of shot groups, while $\text{width} \leq 8$ used in this work does not prolong the minimal TE compared to the standard blipped-CAIPI sampling. In this work, 2-4 shot groups have been demonstrated to be effective across the 3D whole brain imaging protocols tested, which are accomplished by $\text{width} = 4, 6 \text{ or } 8$. In general, using more shot groups is more beneficial in low spatial resolution, low acceleration factor regimes. At 1.8mm isotropic resolution, 4 shot groups ($\text{width} = 8$) are recommended. At 1.2mm isotropic resolution, this choice does not make a significant difference. Extending the current framework to hybrid radial-Cartesian sampling like 3D TURBINE (Chiew et al., 2016) may allow for more flexibility to choose the optimal binning window for each shot group retrospectively, as the golden angle scheme can provide near-uniform coverage at any window size. Radial sampling also has intrinsic robustness to temporal fluctuations, which may also benefit the robustness of 3D sampling further.

The proposed approach is based on the SLR reconstruction method MUSSELS (Mani et al., 2017) for DWI using 2D multi-shot EPI acquisition. The idea of structured low-rank matrix completion has also been successfully employed in some other applications such as calibration-less parallel imaging reconstruction (Haldar and Zhuo, 2016; Liu et al., 2021; Shin et al., 2014b; Yi et al., 2021), EPI Nyquist ghost correction (J. Lee et al.,

2016; Lobos et al., 2018, 2021), EPI B_0 inhomogeneity correction (Balachandrasekaran et al., 2019; Siemonsma et al., 2020), and trajectory error correction (Mani et al., 2018). The inherent linear dependency of the phase-corrupted multi-shot data has enabled us to leverage the low-rank constraint on its block-Hankel structured matrix representation, assuming image phase fluctuations driven by respiration are relatively smooth in the spatial domain, an assumption which has been employed in other work (Wallace et al., 2020). A future extension of this work could be the joint reconstruction across multiple volumes to improve the conditioning of the reconstruction in highly under-sampled regimes, but special care should be taken not to alter the temporal DOF of the time series. Initialising the current intra-volume reconstruction with data across multiple volumes could be a promising option for improving convergence.

The reconstruction strategy employed here empirically estimates the hyperparameter λ which trades off between the data consistency term and the low-rank constraint term. The parameter was hand tuned to optimise for the output tSNR, which provides a useful heuristic for hyperparameter tuning in this problem without requiring any additional training data or prior knowledge. Note the proposed reconstruction works on each volume independently, while the tSNR is a measurement over the entire time course which reflects both the temporal variations and the fidelity of the image magnitude strength. In this work, under-regularisation results in high temporal standard deviation and over-regularisation could lead to either high temporal standard deviation or a reduced image magnitude which prevents the tSNR from increasing despite the low temporal standard deviation, depending on the acceleration factor. The acquired data is normalised before reconstruction and the optimal λ is very robust (3E-4 or 6E-4) across all in vivo datasets.

The proposed reconstruction incorporates the structured low-rank constraint in the SENSE-based parallel imaging formulation which uses coil sensitivities explicitly,

derived from a separate multi-shot calibration dataset in this work. Compared to calibration-free formulations which construct the block-Hankel structured matrix from the multi-coil k-space, the calibration-based formulation which reconstructs a single-coil k-space is more computationally efficient as the block-Hankel structured matrix could have a much-reduced size. Also, the coil sensitivity maps are easy to obtain by reference scan, and the moderate phase variations in fMRI typically do not lead to visible artefacts in the calculated coil sensitivities. However, in the worst case scenario when high fidelity coil sensitivity maps cannot be obtained from the calibration data, alternative approaches can be used, such as employing the low-rank tensor representation (Hess et al., 2021; Liu et al., 2021; Yi et al., 2021) of multi-channel k-space for a calibration-less reconstruction without explicit use of coil sensitivity maps, or using a calibration consistency constraint that jointly identifies a coil-null space from the imaging and calibration data, without trusting either dataset completely (Lobos et al., 2021). In addition, partial Fourier sampling which is typically used to reduce TE was not employed in this work, but it is also compatible with the seg-CAIPI trajectory and the proposed reconstruction. When partial Fourier sampling is used, the VC-matrix (Bilgic et al., 2019; Kim and Haldar, 2018), which incorporates the conjugate symmetrical k-space into the block-Hankel structured matrix construction as virtual shots might be beneficial. This approach was evaluated on the retrospectively under-sampled partial Fourier data. However, the reconstruction results were not further improved compared to the results achieved without using virtual shots, which is probably because the baseline image phase is not sufficiently smooth to effectively leverage the virtual conjugate shot constraints.

The convex nuclear norm minimisation is used to enforce the low-rank property of the block-Hankel structured matrix in this work. However, instead of convex soft thresholding, a non-convex hard thresholding is used in the nuclear norm minimisation

subproblem in ADMM, as it has shown its ability to achieve better image quality than soft thresholding in this work, particularly in low tSNR regimes at 1mm isotropic resolution, at the cost of slightly lower output tSNR.

One limitation of the proposed reconstruction is the computation time. The reconstruction takes ~120 seconds (3.1 GHz Intel Core i7 and 16 GB RAM) for a 2D 116×96 sagittal matrix at $R = 2 \times 2$. Since all sagittal slices can be reconstructed independently, the reconstruction does parallelize well with additional computational resources. Also, the reconstruction time might be further reduced by using some other optimisation algorithms with higher computational efficiency like the SVD-free methods discussed in Section 3.2.3 (D. Lee et al., 2016; Mohan and Fazel, 2012). Another limitation of this work is that no consideration of motion artefacts is taken into account. As the low-rank property of the block-Hankel structured matrix representation relies on the assumption that different shot groups share the same image magnitude, motion induced magnitude mismatch between shot groups could violate this assumption. Thus, a future extension of this work could be to incorporate motion estimates in the forward model to preserve the low-rankness of the block-Hankel structured matrix representation and ultimately improve the robustness of 3D multi-shot EPI further. One particularly useful application of the proposed method could be brainstem and spinal cord imaging, where multi-shot 3D imaging can provide the high spatial resolution needed to resolve important structures, but respiration induced phase variations are expected to be more significant, and motion artefacts could also be more severe.

4.5 Summary

In this chapter, a structured low-rank image reconstruction method is proposed for 3D multi-shot EPI, which in coordination with the 3D seg-CAIPI sampling pattern can reduce

the vulnerability to physiologically induced inter-shot phase variations, and thus improves the temporal stability of 3D multi-shot EPI for fMRI at 7T.

5 Motion Compensated Structured Low-rank Reconstruction for 3D Multi-shot EPI

5.1 Introduction

As discussed previously, multi-shot EPI imaging can suffer from shot-to-shot phase variations due to various sources of instability. In DWI, phase variations can emerge even with subtle subject motion (e.g., cardiac pulsations) in the presence of strong diffusion

gradients, whereas in fMRI, they are contributed by physiologically induced field fluctuations, mostly due to respiration. Although the inter-shot phase variations may not produce noticeable ghost artefacts in fMRI, they can result in unwanted temporal variance and reduce the temporal stability of the time course. To address this issue in 3D multi-shot EPI for fMRI, a structured low-rank reconstruction method has been proposed in Chapter 4 which partitions all shots into several shot groups with each having a reduced temporal footprint and therefore reduced phase inconsistencies. An individual image for each shot group is reconstructed and the images for different shot groups are then SOS combined taking account of phase inconsistencies between them. Since each shot group is highly under-sampled, a joint reconstruction of all shot groups is performed to explore the linear dependencies among them by enforcing a low-rank constraint on the block-Hankel structured matrix constructed from the multi-shot group k-space.

In addition to physiologically induced phase variations which are always present in vivo, multi-shot acquisitions are also vulnerable to bulk motion (i.e., head rotation and translation), which causes not only inter-volume inconsistencies shown as misregistration between different volumes, but also image artefacts arising from intra-volume, shot-to-shot inconsistencies. These motion induced inter-shot inconsistencies are also problematic for the aforementioned SLR reconstruction for multi-shot EPI imaging, as it breaks the basic assumption that different shots only differ in phase but have the same image magnitude. Thus, any motion induced signal magnitude variations could reduce the linear inter-dependencies of different shots. Accordingly, the joint reconstruction leveraging the structured low-rank constraint does not have a significant benefit as expected. To illustrate this, Fig. 5.1 compares the normalised singular values of the block-Hankel structured matrices constructed from two simulated multi-shot datasets, with and without inter-shot motion respectively. It is shown that the low-rank property of the

block-Hankel structured matrix is largely invalid due to inter-shot motion.

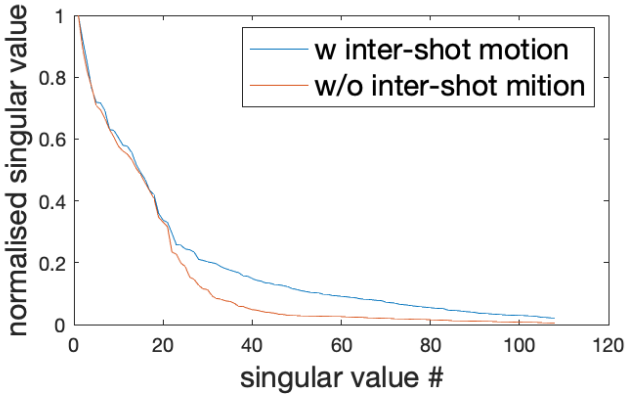

Fig. 5.1 The normalised singular values of two block-Hankel structured matrices constructed from simulated multi-shot datasets with and without inter-shot motion. Each dataset has 3 shots in total and the one with inter-shot motion is subject to a 3° rotation between every two consecutive shots.

One previous method which we refer to as aligned reconstruction (Cordero-Grande et al., 2016), deals with multi-shot structural imaging in the presence of rigid motion. This approach jointly estimates the structural image and rigid motion parameters in a fully data driven way by exploiting the sensitivity encoding redundancy. The method TAMER (Haskell et al., 2018) performs similar joint image-motion optimisation and features a high computation efficiency. The method DISORDER (Cordero-Grande et al., 2020) further improves the performance and efficiency of aligned reconstruction using distributed and incoherent sampling orders. Recently, this method has been further improved at 7T by incorporating the modelling of pose-dependent B_0 field changes (Brackenier et al., 2022). However, none of these motion correction methods developed for structural imaging considers the physiologically induced phase variations of multi-shot EPI acquisition and exploit additional structured low-rank constraints, and they have also not yet been examined on 3D multi-shot EPI data.

In this chapter, a motion compensated SLR (mcSLR) reconstruction framework for robust 3D multi-shot EPI imaging which takes into account both inter-shot motion and

inter-shot phase variations is proposed. The proposed method can correct for the motion induced inconsistencies between acquired shots and thus facilitate improved SLR reconstruction in the presence of rigid motion. The performance of mcSLR has been validated by simulation and in vivo experiments at 7T.

5.2 Methods

5.2.1 Formulation of the Optimisation Problem

The proposed mcSLR reconstruction incorporates rigid-body motion transforms in the forward model to compensate for motion induced inconsistencies. In line with the SLR formulation, intra-volume motion is decomposed into *intra*-shot group motion between different temporal subdivisions within each shot group, and *inter*-shot group motion between the images of different shot groups. Fig. 5.2 shows the diagram of the forward model of mcSLR. The images of different shot groups to be reconstructed are assumed to be aligned, so we can apply a low-rank constraint on the block-Hankel structured matrix constructed from their k-space. The proposed reconstruction is formulated as the following optimisation problem:

$$\min||AFST_{intra}T_{inter}X - Y||_2^2 + \lambda||H_sFX||_* \qquad \text{Eq. 5.1}$$

Where X: $[x_1, x_2, \dots x_n]$ consists of 3D images of n shot groups. T_{inter} is the *inter*-shot group motion transform and T_{intra} is the *intra*-shot group motion transform. The operator S applies coil sensitivity encoding and F performs Fourier transform. A denotes the sampling operator. Y is the acquire k-space data. The operator H_s constructs the block-Hankel structured matrix from the k-space data of all shot groups FX.

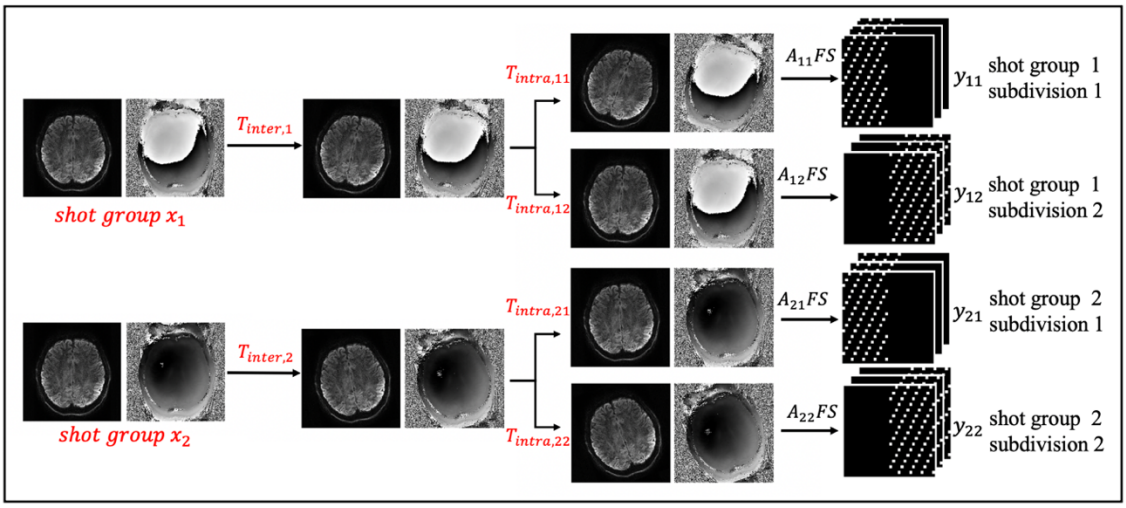

Fig. 5.2 The diagram of the forward model of mcSLR reconstruction. The images of different shot groups (e.g., x_1, x_2) to be reconstructed are assumed to be aligned. The inter-shot group motion operator T_{inter} transforms them to their different positions respectively. The intra-shot group motion operator T_{intra} then maps the image of each shot group to the actual positions corresponding to different temporal subdivisions of this shot group.

For the implementation of the rigid motion transforms T_{intra} and T_{inter}, an interpolation-free representation is used as in aligned reconstruction (Cordero-Grande et al., 2016; Unser et al., 1995). This formulation decomposes a rotation into three consecutive shears to perform high quality rotations without the need of regridding. Thus, a 3D rigid motion transform with translation parameters $[\delta_1, \delta_2, \delta_3]$ and rotation parameters $[\theta_1, \theta_3, \theta_3]$ can be represented as a series of linear phase modulations in the hybrid $r-k$ space as:

$$
\begin{aligned}
T = {} & F^H \Phi_{tra} F \\
& \times F_2^H \Phi_{rot,\theta_1}^{tan} F_2 F_3^H \Phi_{rot,\theta_1}^{sin} F_3 F_2^H \Phi_{rot,\theta_1}^{tan} F_2 \\
& \times F_3^H \Phi_{rot,\theta_2}^{tan} F_3 F_1^H \Phi_{rot,\theta_2}^{sin} F_1 F_3^H \Phi_{rot,\theta_2}^{tan} F_3 \\
& \times F_1^H \Phi_{rot,\theta_3}^{tan} F_1 F_2^H \Phi_{rot,\theta_3}^{sin} F_2 F_1^H \Phi_{rot,\theta_3}^{tan} F_1
\end{aligned}
\qquad \text{Eq. 5.2}
$$

Where F_i represents Fourier transform along the i_{th} dimension. Φ_{tra} denotes the phase modulation in k-space corresponding to translation. $\Phi_{rot,*}^{tan}$ and $\Phi_{rot,*}^{sin}$ denote the

phase modulations in the hybrid $r-k$ space that are used to implement the decomposed rotation. The entries of Φ_{tra}, $\Phi_{rot,*}^{tan}$ and $\Phi_{rot,*}^{sin}$ are given as:

$$\phi_{tra}(k_1,k_2,k_3) = e^{-1j(\delta_1 k_1+\delta_2 k_2+\delta_3 k_3)}$$

$$\begin{aligned}
&\phi_{rot,\theta_1}^{tan}(k_2,r_3) = e^{1j(tan(\theta_1/2)k_2 r_3)} \quad \phi_{rot,\theta_1}^{sin}(k_3,r_2) = e^{1j(sin(\theta_1/2)k_3 r_2)} \\
&\phi_{rot,\theta_2}^{tan}(k_3,r_1) = e^{1j(tan(\theta_2/2)k_3 r_1)} \quad \phi_{rot,\theta_2}^{sin}(k_1,r_3) = e^{1j(sin(\theta_2/2)k_1 r_3)} \\
&\phi_{rot,\theta_3}^{tan}(k_1,r_2) = e^{1j(tan(\theta_3/2)k_1 r_2)} \quad \phi_{rot,\theta_3}^{sin}(k_2,r_1) = e^{1j(sin(\theta_3/2)k_2 r_1)}
\end{aligned} \qquad \text{Eq. 5.3}$$

Where r_* is the coordinate in image space and k_* is the coordinate in k-space.

5.2.2 Solution to the Optimisation Problem

All three variables X, T_{intra}, T_{inter} are unknown and need to be estimated. Since the proposed method does not rely on external measurements or navigators to acquire the motion traces, these three variables are jointly estimated in a fully data driven way. The problem is solved by alternating between three subproblems in an iterative fashion. For each subproblem, one variable is updated while fixing the other two as estimated from the last iteration.

Subproblem 1- Estimate *X* using a Pseudo 3D SLR Constraint

The images of all shot groups are reconstructed by solving the following optimisation with motion estimates $\hat{T}_{intra}$ and $\hat{T}_{inter}$ fixed:

$$\underset{X}{\text{argmin}} ||AFS\hat{T}_{intra}\hat{T}_{inter}X - Y|_2^2 + \lambda||H_s FX||_* \qquad \text{Eq. 5.4}$$

In Chapter 4, the 3D reconstruction of the whole volume was decomposed into 2D reconstructions of each $y-z$ slice after performing 1D inverse Fourier transform along the fully sampled k_x direction, which can alleviate the computational challenge of 3D

reconstruction and enable the use of parallel computing resources. However, due to the 3D motion transform involved in this problem, the optimisation needs to be solved simultaneously for the 3D volume as a whole, which presents a challenge for 3D structured low-rank reconstruction as the block-Hankel structured matrix will be of extremely large size making the computation infeasible. This challenge is circumvented by using a pseudo 3D structured low-rank constraint which applies a low-rank constraint on the block-Hankel structured matrix constructed from each 2D $r-k$ slice. The 2D image slice can be chosen along any direction, and Fig. 5.3 shows an example where a block-Hankel structured matrix is constructed from each 2D k-space of a $y-z$ image slice. The optimisation with this pseudo 3D structured low-rank constraint is reformulated as:

$$\underset{X}{\operatorname{argmin}} ||AFS\hat{T}_{intra}\hat{T}_{inter}X - Y||_2^2 + \lambda \sum_{l=1}^{N_l} ||H_s F X_l||_* \qquad \text{Eq. 5.5}$$

Where X_l denotes a 2D image slice of X along either x, y or z direction. l is the slice index and N_l denotes the number of slices along the chosen direction. F is the 2D Fourier transform operator. Note FX_l corresponds to the k-space of a 2D image slice, not a 2D slice picked from a 3D k-space. H_s builds the block-Hankel structured matrix from a 2D k-space using a 2D kernel. This optimisation subproblem is solved by ADMM as in conventional SLR reconstruction.

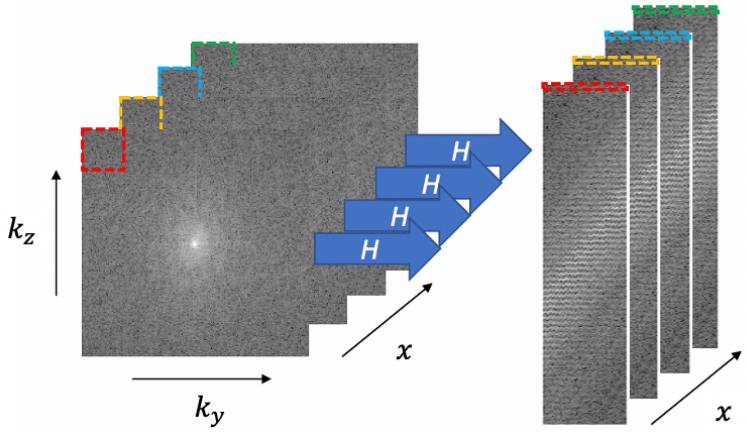

Fig. 5.3 Pseudo 3D structured low-rank constraint. Instead of constructing a block-Hankel structured matrix from the 3D k-space with a 3D kernel, a block-Hankel structured matrix is constructed from the 2D k-space of each $y - z$ slice.

Subproblem 2 - Estimate T_{inter} using Image-based Registration Method

Since the images of different shot groups are available after solving subproblem 1, the inter-shot group motion parameters between them can be estimated by image-based registration methods. In this work, the image registration method FLIRT (Jenkinson et al., 2012) from the FSL toolbox is used. The first shot group is chosen as the reference volume, and all other shot groups are registered to it to get the residual motion between them, which are then concatenated with the motion estimates from last iteration to get the new update. Fig. 5.4 shows the diagram of this pipeline.

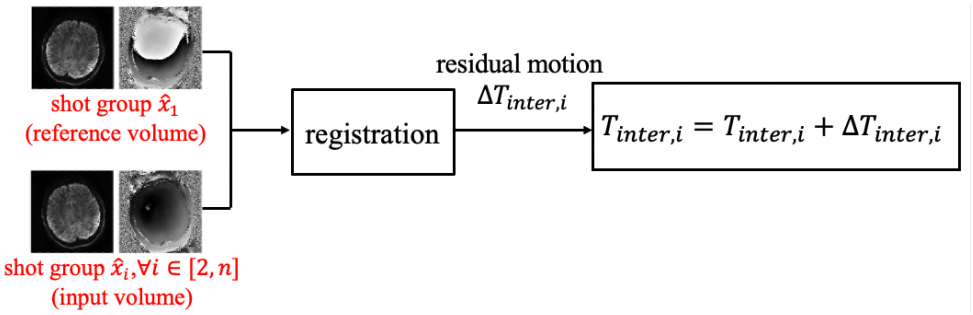

Fig. 5.4 The pipeline of inter-shot group motion estimation. The images of different shot groups are reconstructed by solving subproblem 1. The first shot group is chosen as the reference, and every other shot group is registered to the reference to get the residual motion estimates, which are then concatenated with the motion estimates from last iteration to obtain the new estimates.

Subproblem 3 - Estimate T_{intra} by Solving the Optimisation of Data Consistency Constraint

Since now the image for each shot group is available, the intra-shot group motion parameters can be estimated by solving the optimisation which enforces data consistency constraint between the motion artefacts-free image and the motion corrupted multi-channel k-space data of each temporal subdivision, as in the aligned reconstruction approach. The optimisation is formulated as:

$$\underset{T^d_{intra,i}}{\operatorname{argmin}} ||A_i^d F S T^d_{intra,i} \hat{T}_{inter,i} \hat{x}_i - Y_i^d||_2^2$$
$$\forall d = 1, \dots m;\ i = 1, \dots n \qquad \text{Eq. 5.6}$$

Where i is the shot group index. $\hat{x}_i$ denotes the image of shot group i. $\hat{T}_{inter,i}$ is the inter-shot group motion transform corresponding to shot group i. d is the temporal subdivision index within each shot group. A_i^d and Y_i^d are the sampling mask and the motion-corrupted k-space data of the d_{th} subdivision of shot group i. It is solved by a modified Levenberg–Marquardt (LM) algorithm (Levenberg, 1944; Marquardt, 1963) based on the MATLAB implementation provided in aligned reconstruction (Cordero-Grande et al., 2016). The LM algorithm is an iterative method which can be viewed as a combination of gradient descent and Gauss-Newton methods. In each iteration, an increment of the motion parameters ΔT is calculated as:

$$\Delta T = -(wI_d + J^H J)^{-1} G \qquad \text{Eq. 5.7}$$

Where w is the damping parameter which balances between gradient descent and Gauss-Newton methods. I_d is the unit matrix. J is the Jacobian matrix and $J^H J$ is an approximation of the Hessian matrix. G is the gradient of the cost function. When w is very big, ΔT is determined by G as in gradient descent method. When w is small, ΔT is

determined by $(J^H J)^{-1} G$ as in Gauss-Newton method. The cost function of a potential new estimate $\hat{T} + \Delta T$ is compared to the cost function of the current estimate $\hat{T}$, and $\hat{T} + \Delta T$ will be accepted as the new estimate only if it has a lower cost function value. w is updated automatically as follows: when the cost function value of $\hat{T} + \Delta T$ gets smaller, w will be reduced in next iteration, otherwise it will be increased. A modification on the LM algorithm is also presented in this work such that instead of a scalar damping parameter, a vector w in accordance with the dimensionality of the motion parameter (6 for rigid motion) is used. It is observed that this modification leads to a significant improvement in simulations, and a validation of this modification is shown in the Experiments section.

5.2.3 Additional Implementation Details

The Non-uniqueness of Inter and Intra-shot Group Motion

In this work, intra-volume motion is decomposed into inter-shot group and intra-shot group motion in accordance with the formulation of SLR reconstruction, and by doing so they can be more easily estimated by different methods. Since we have the images of different shot groups available, it is straightforward to estimate inter-shot group motion by image registration-based methods. For each shot group, intra-shot group motion can be readily estimated by solving the optimisation of data consistency constraint as in aligned reconstruction. In the forward model, there is indeed no need to use this two-step motion transform, and the two motion terms $\hat{T}_{intra}\hat{T}_{inter}$ can be combined first to improve the computational efficiency when solving subproblem 1.

It is worth noting that there are no unique solutions for the inter and intra-shot group motion without further constraints. For example, given $\hat{X}$, the intermediate image $T_{inter}\hat{X}$ can be at random positions and the cost function does not change as long as the

concatenation of inter and intra-shot group motion $T_{intra}T_{inter}$ remains the same (i.e. T_{intra} undoes any additional misalignment introduced by T_{inter}). Thus, to avoid potential drifting effect, the position of the intermediate image $T_{inter,i}x_i$ is fixed to be the same as the actual position of its first subdivision, which is achieved by subtracting the motion parameters of the first subdivision $T^1_{intra,i}$ from the motion parameters of all other subdivisions $T^d_{intra,i}, d \in [2, m]$ as the last step of subproblem 3. Note the subtracted motion parameters of the first subdivision $T^1_{intra,i}$ are then added back to the inter-shot group motion $T_{inter,i}$ directly (except the first shot group). This corresponds to an additional constraint that $T^1_{intra,i} = 0$. Apart from the non-unique decomposition of T_{intra} and T_{inter}, there is also ambiguity in the position of the aligned shot groups X without further constraints, i.e., we can easily replace $\hat{X}$ by $T_0\hat{X}$ with a random motion term T_0, and get the same cost function value with a corresponding motion term $T_{intra}T_{inter}T_0^H$. However, since the image of the first shot group is assigned to be the reference volume when estimating inter-shot group motion, this corresponds to another constraint that $T_{inter,1} = 0$. These two additional constraints make sure that the position of the aligned shot groups X is fixed to be the same as the actual position of the first subdivision of the first shot group, and all three variables have a unique solution.

Initialisation of Motion Estimates

To speed up convergence of the whole problem, aligned reconstruction is performed first on a single combined shot group consisting of every shot regardless of the phase variations to initialise the motion parameters. More specifically, the motion estimates of aligned reconstruction are used to initialise T_{intra} directly without being decomposed into T_{intra} and T_{inter} first, since after the first iteration of mcSLR, the intra-shot group motion

of the first subdivision will be subtracted from T_{intra} and added to T_{inter} as mentioned earlier.

Input Image of Subproblem 3

As shown previously in Eq. 5.6, the image of each shot group at its respective position, $\hat{T}_{inter,i}\hat{x}_i$, is used as the input image for subproblem 3. In practice, the input image $\hat{x}_i{}'$ is a combination of all shot groups defined as:

$$\hat{x}_i{}' = \frac{1}{n}\left(\sum_{i=1}^{n}|\hat{x}_i|\right)\exp\left(1j\angle(\hat{x}_i)\right) \qquad \text{Eq. 5.8}$$

The magnitude of $\hat{x}_i{}'$ is the average of the magnitude images of all shot groups. The phase of $\hat{x}_i{}'$ is the same as the original phase of $\hat{x}_i$ considering the phase variations between different shot groups. It has been observed that using $\hat{x}_i{}'$ could speed up the convergence of the problem compared to using the original image $\hat{x}_i$ in some simulations, likely due to the higher SNR of $\hat{x}_i{}'$.

5.3 Experiments

Simulation experiments based on 2D DWI and 3D fMRI data and in vivo experiments at 7T were performed to validate the performance of mcSLR. The kernel size of the block-Hankel transformation was chosen empirically to be 6×6. $\lambda = 3E-4$ was used for the DWI simulation. $\lambda = 1E-1$ or $3E-1$ was used for the fMRI simulation. $\lambda = 6E-2$ was used for the in vivo data. Coil compression (Buehrer et al., 2007) was performed to reduce reconstruction time. The performance of the proposed method was evaluated by comparing to SENSE reconstruction, aligned reconstruction as well as conventional structured low-rank reconstruction.

5.3.1 Simulation Experiments

2D DWI Simulation

A 3-shot 2D EPI dataset was simulated with inter-shot phase variations obtained from a multi-shot DWI dataset and a 5° in-plane rotation between every two consecutive shots. The matrix size of the DWI data is 210×210. Cartesian sampling with a $R_y = 2$ uniform under-sampling pattern and 9 compressed receive channels were used. The Normalised Root Mean Square Error (NRMSE) calculated between the reconstructed magnitude image and the ground truth magnitude image was used to evaluate the performance of different reconstruction methods. Due to the small number of shots in this 2D case, each shot group consists of a single shot and only inter-shot group motion was considered.

The results of the DWI dataset are shown in Fig. 5.5. The SLR reconstruction followed by a direct SOS combination suffers from substantial motion artefacts. Although a post hoc registration before SOS combination can largely remove the motion artefacts in this simplified scenario, as shown in the bottom row of Fig. 5.5a, it does not improve the efficacy of SLR constraint. Thus, it still has a worse NRMSE than the SOS combined image of mcSLR. The improved efficacy of the SLR constraint in mcSLR can also be demonstrated by the lower NRMSE of the first shot group compared to SLR as shown in Fig. 5.5b.

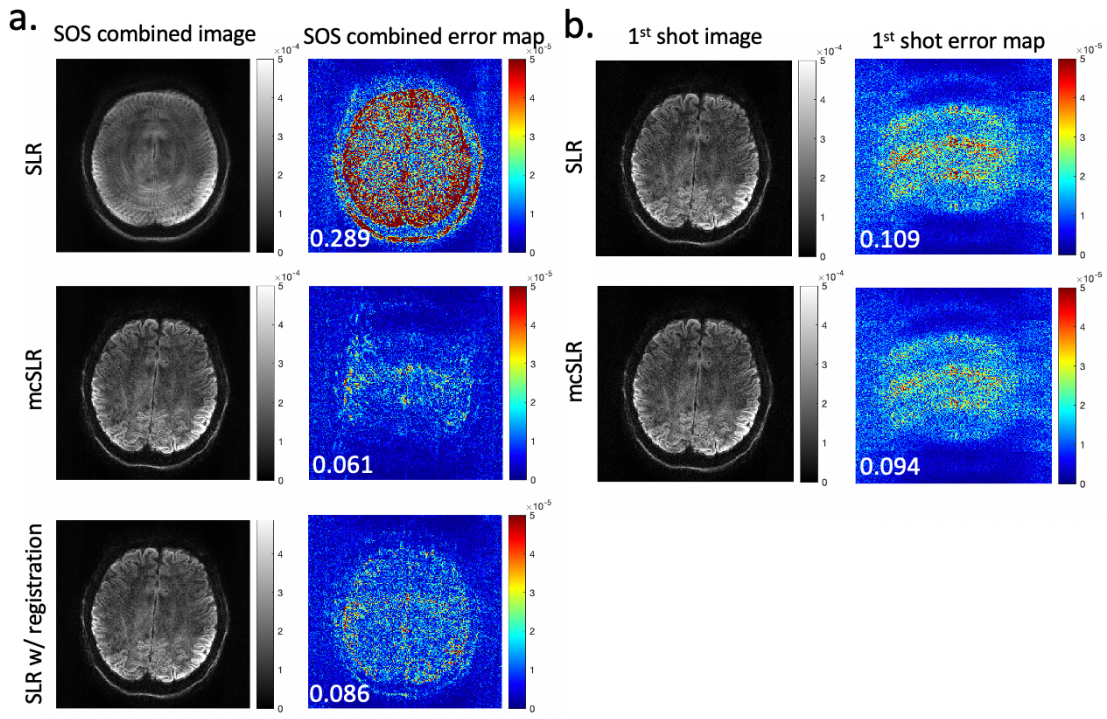

Fig. 5.5 The simulation results of the DWI dataset. a) The SOS combined image and its error map. b) The image of the first shot and its error map. The NRMSE is shown in each error map respectively. The bottom row of a) shows the results of SLR with multi-shot images being registered before SOS combination. Although post hoc registration can remove the motion artefacts of SLR in this simplified scenario, it does not improve the efficacy of SLR constraint, thus the NRMSE is still higher than mcSLR for both the SOS combined image and the image of the first shot group.

3D fMRI Simulation

A 2D sagittal slice of a 3D EPI dataset with inter-shot motion and inter-shot phase variations was simulated. The magnitude image and phase variations were obtained from a 2D EPI time series acquired at 7T with 1.5mm isotropic resolution and TE/TR=20/40ms. The matrix size of the 2D slice is 140×96. An acceleration factor $R_y \times R_z = 2 \times 2$ (48 shots per volume) and 11 compressed receiver channels were used for simulation. The seg-CAIPI(8,3) sampling trajectory as described in Chapter 4 was used. Four shot groups were used for SLR and mcSLR reconstructions. Continuous motion was generated by linear interpolation from N discrete motion samples which are randomly chosen from a uniform distribution within $[-\frac{\theta}{2}, +\frac{\theta}{2}]$ degrees for in-plane rotation and $[-\frac{\delta}{2}, +\frac{\delta}{2}]$ mm for

1D translation along the k_y direction. A time course consisting of 16 volumes was generated. The temporal resolution of motion estimates was 40ms, which is the same as the shot TR.

Fig. 5.6 shows the results of a simulation experiment where motion estimates were obtained by the LM algorithm using a) a scalar damping parameter and b) a vector damping parameter. The ground truth image was used as the input for motion estimation. The motion estimates for one subdivision versus number of iterations of the LM algorithm are shown. Each subdivision consists of 6 shots ($R = 32$) in this experiment. The LM algorithm using a vector damping parameter leads to a significantly faster convergence at the cost of slightly worse estimation for the translation parameter along y direction (middle row).

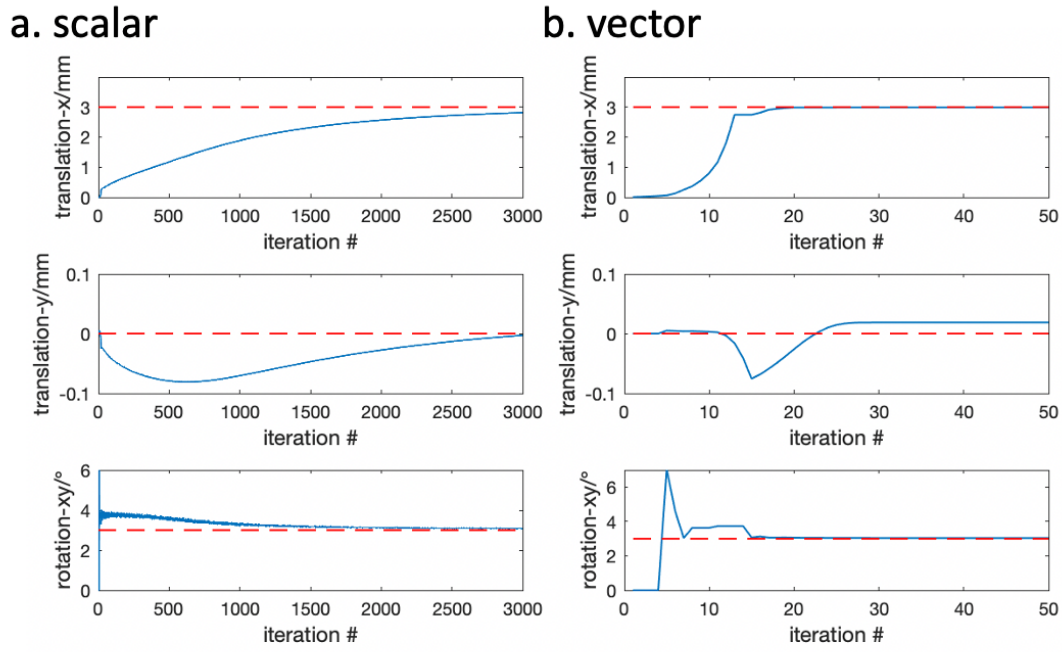

Fig. 5.6 The convergence speed of the Levenberg–Marquardt algorithm with a) a scalar damping parameter and b) a vector damping parameter. The ground truth motion parameter is indicated by the red dashed line and the motion estimates through iteration are represented by the blue solid line. The ground truth image is used as the input to the Levenberg–Marquardt algorithm.

The proposed method was validated at different ranges and frequencies of motion and the results are shown in Fig. 5.7. The range of motion tested is $\theta = 6°, 9°, 12°, 15°$, for

which the motion was randomly sampled within $[-3°/2, +3°/2]$ and then amplified by 2/3/4/5 times. Different frequencies of motion were achieved by interpolating $N = 8/16/24$ discrete samples to generate the continuous 48-shot motion traces. It is shown that both aligned reconstruction and mcSLR have better image quality than SLR, and mcSLR has an overall more robust performance than aligned reconstruction.

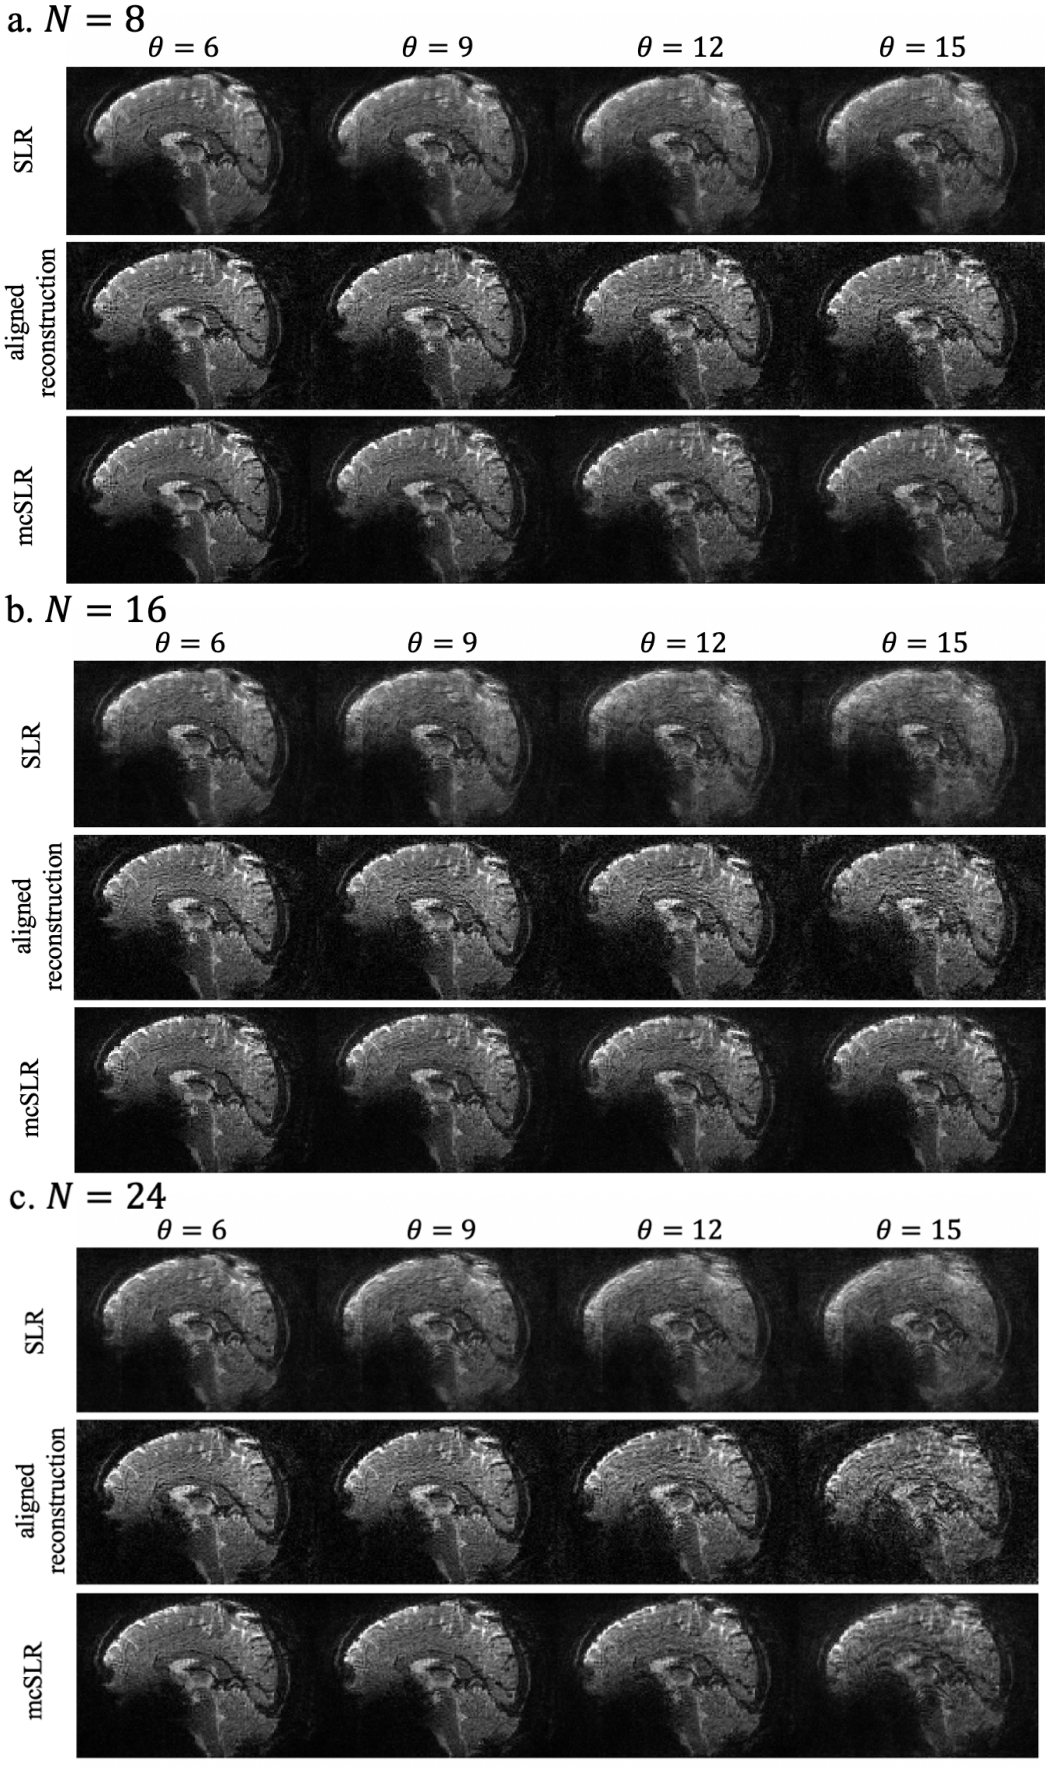

Fig. 5.7 The comparison of different reconstruction methods in simulation experiments with different ranges and frequencies of motion.

The performance of different reconstruction methods was also evaluated over the time course, where the motion traces were generated with $\theta = 3°$, $\delta = 3mm$ and $N = 8$. The reconstructed images of one reference volume are shown in Fig. 5.8. The result of SENSE reconstruction shows substantial artefacts. The results of the SLR reconstruction suffer from blurring. In contrast, both aligned reconstruction and mcSLR are able to reconstruct clean and sharp images, and the red ellipse highlights the improved sharpness in small structures. While the temporally varying image artefacts due to phase variations are not visible in a single volume, they reduce the temporal stability of a time course, which is shown in the tSNR comparison in Fig. 5.9. The time course was registered before tSNR calculation, so that residual inter-volume motion effects do not contribute to the observed temporal variance. The temporal mean images of SENSE and SLR are both blurry and they also have considerably lower tSNR, whereas aligned reconstruction and mcSLR can largely restore the image resolution and improve tSNR. Despite comparable image quality, mcSLR achieves a 47% higher tSNR than aligned reconstruction.

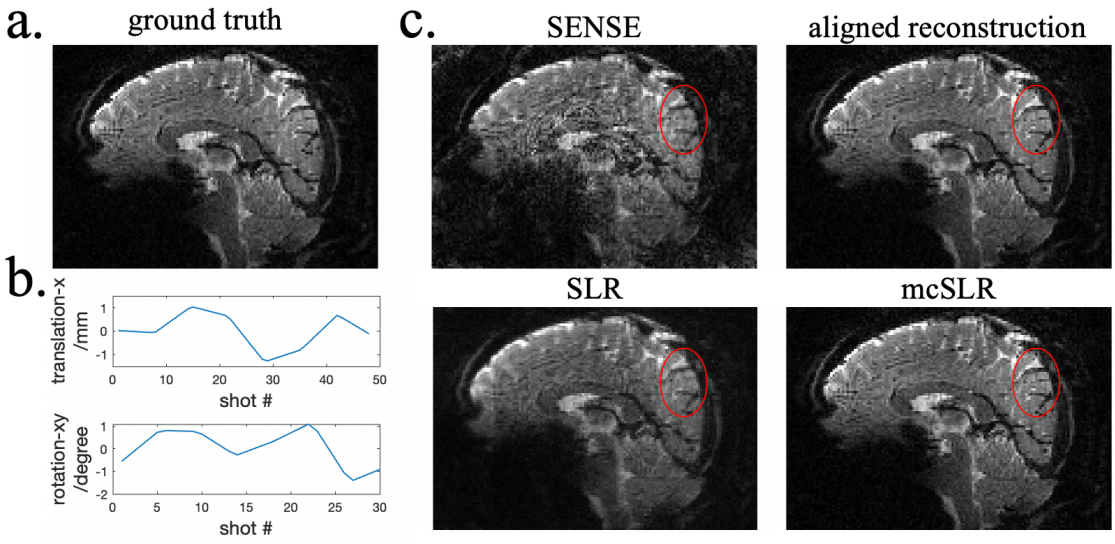

Fig. 5.8 The comparison of different reconstruction methods on a single simulated volume. a) The ground truth magnitude image. b) The simulated motion traces. c) The reconstructed images. The motion traces were generated with $\theta = 3°$, $\delta = 3mm$ and $N = 8$.

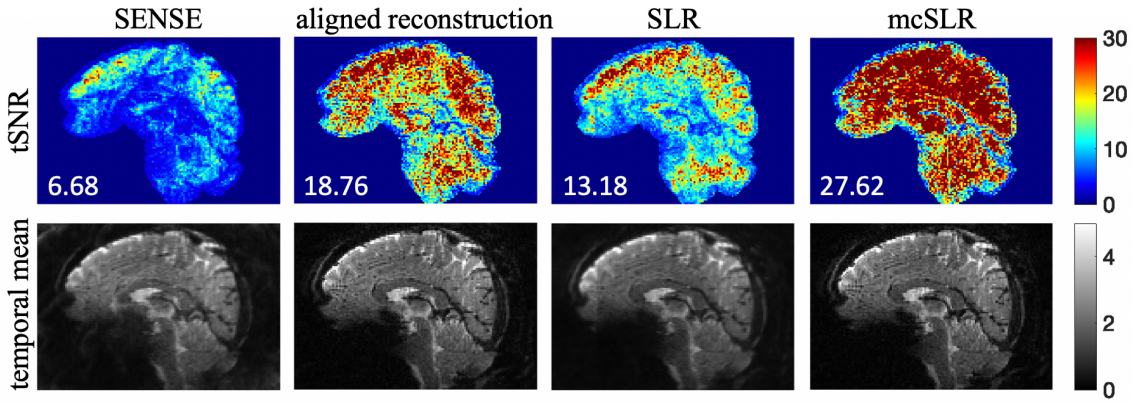

Fig. 5.9 The comparison of different reconstructions on the simulated time course. The tSNR maps and temporal mean magnitude images are shown. The mean tSNR value is shown in each tSNR map respectively. The time course was registered prior to tSNR calculation.

5.3.2 In Vivo Experiments

In vivo experiments at 7T were performed on three healthy subjects to validate the performance of the proposed mcSLR reconstruction. One subject was instructed to keep still during the scan, and this is the same motion-free dataset shown in Fig. 4.10. Two subjects were instructed to perform deliberate motion continuously during the scan. The motion patterns are divided into two levels depending on the maximum amplitude of rotation: 2.5° for medium motion and 5° for large motion. The imaging parameters were: 1.8mm isotropic resolution, matrix size $= 116 \times 116 \times 96$, TE/TR $= 22/45$ms, $R_y \times R_z = 2 \times 2$. The seg-CAIPI(8,3) sampling trajectory was used. Four shot groups were used for SLR and mcSLR reconstructions. The temporal resolution of motion estimates was 270ms for both aligned reconstruction and mcSLR (6 shots per motion state).

The performance of the pseudo 3D structured low-rank constraint with 2D slices along different directions were compared on one volume which suffers from very subtle motion. The results are shown in Fig. 5.10, where the reconstructed images and the error maps with reference to the aligned reconstruction are compared. It is shown that the three choices of slice direction do not lead to significant differences.

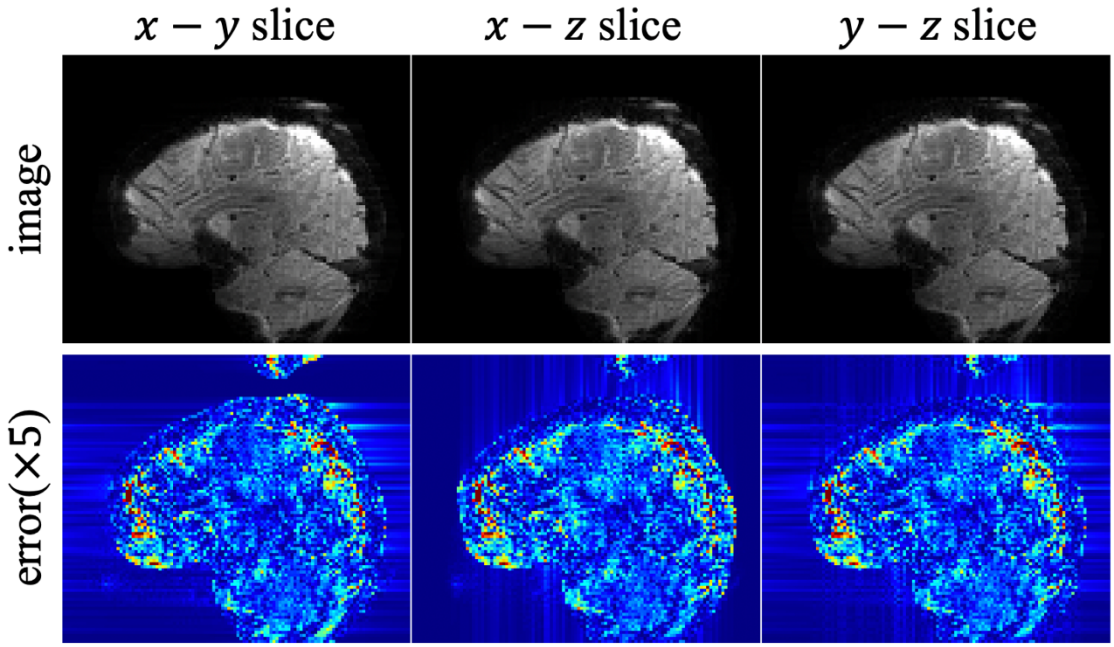

Fig. 5.10 The comparison of SLR reconstruction with pseudo 3D structured low-rank constraint on 2D slices in different directions. The error map was calculated by comparing to the result of aligned reconstruction.

Fig. 5.12 shows the reconstruction results of the conventional SLR method and the proposed mcSLR method on the motion-free dataset. These two methods show comparable results without any noticeable differences.

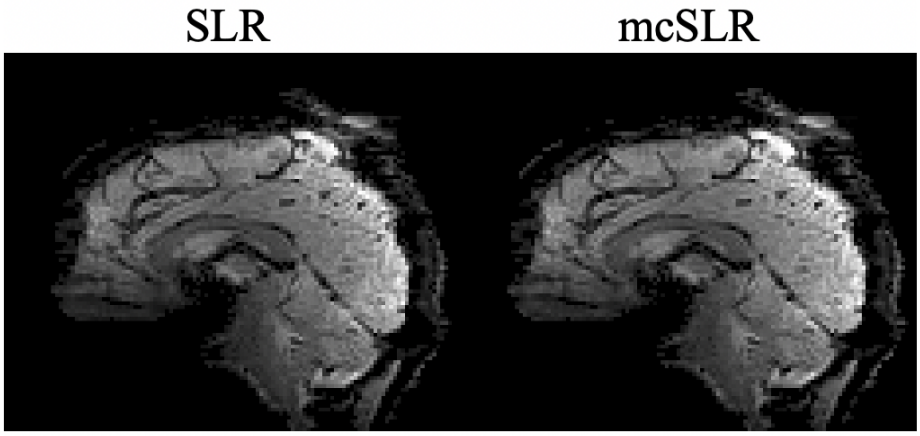

Fig. 5.11 The comparison of different reconstruction methods on the motion-free dataset.

Fig. 5.12 shows the results of different reconstruction methods on one volume with medium motion. The red ellipse in the sagittal view indicates the blurring of SLR compared to both aligned reconstruction and mcSLR. The green ellipse in axial view

indicates the residual artefacts of aligned reconstruction. The proposed mcSLR method demonstrates the overall best performance. The motion estimates of aligned reconstruction and mcSLR for each subdivision are shown in Fig. 5.13 and they are in good agreement in the primary motion dimensions: rotation in the $y-z$ plane and translation along the y direction. Note the motion estimates of mcSLR are shown as the combination of T_{inter} and T_{intra}.

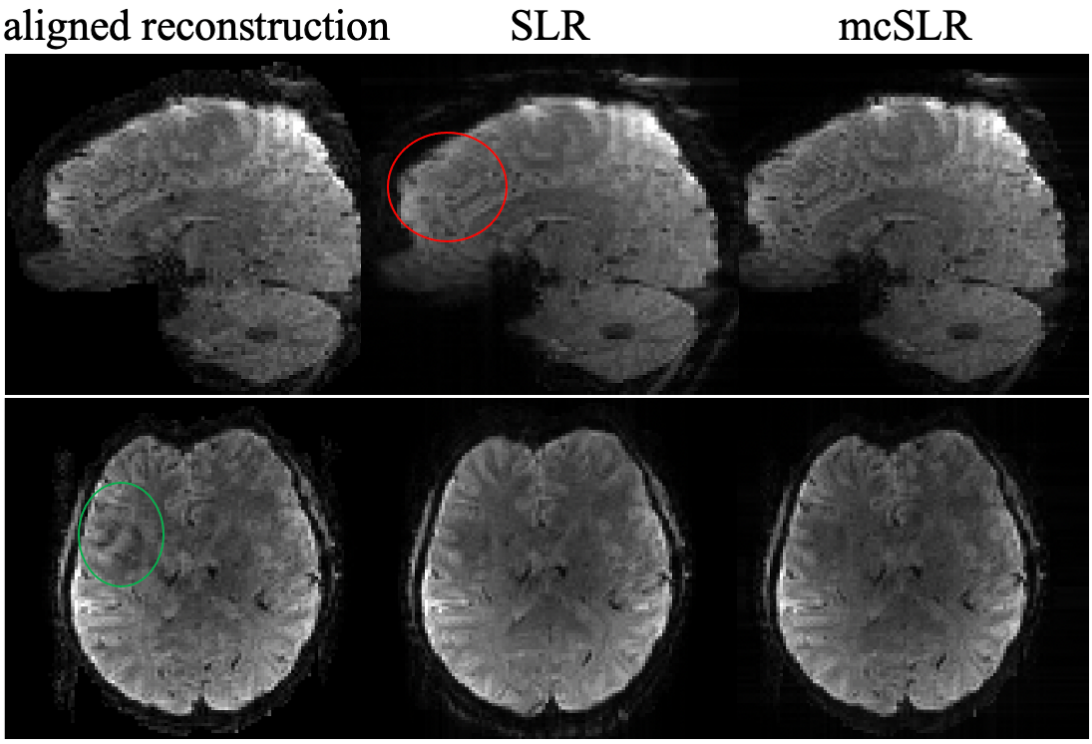

Fig. 5.12 The comparison of different reconstruction methods on one volume with medium motion. The red ellipse in the sagittal view indicates the blurring of SLR compared to both aligned reconstruction and mcSLR. The green ellipse in axial view indicates the residual artefacts of aligned reconstruction.

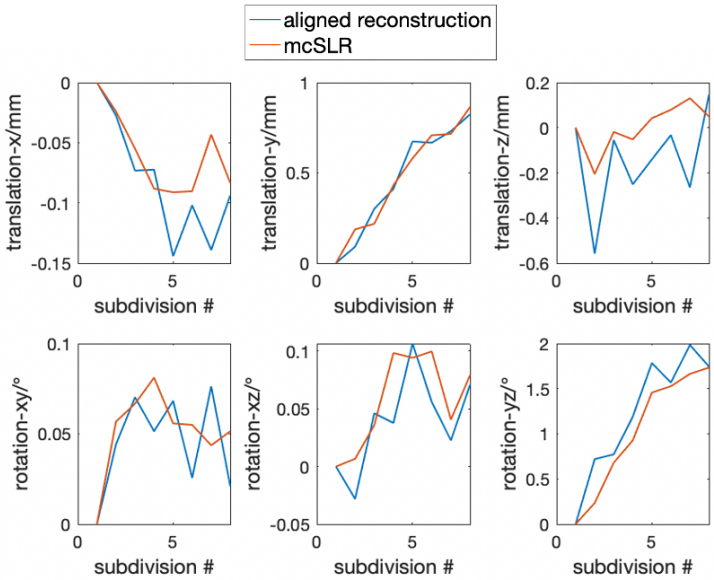

Fig. 5.13 The motion estimates of aligned reconstruction and mcSLR of the results shown in Fig. 5.12.

Fig. 5.14 shows the results of another volume with medium motion artefacts. In this case, aligned reconstruction largely removes the aliasing and blurring compared to SENSE reconstruction, but it still retains some artefacts. The results of SLR reconstruction do not have such severe aliasing artefacts as SENSE or aligned reconstruction, but suffer from motion induced blurring. This is probably because SLR does the SOS combination of the images of different shot groups, and since the k-space of each shot group is sampled sequentially from $-k_z$ to $+k_z$ by seg-CAIPI sampling, motion is still mapped coherently across the k-space of each shot group, which leads to relatively minor blurring. However, similar to the incoherent sampling of DISORDER (Cordero-Grande et al., 2020), the interleaved ordering of seg-CAIPI leads to incoherent mapping of motion in shot-combined reconstruction like SENSE, in which case the motion artefacts are more severe. Besides, the phase variations between different shot groups might also contribute to the severe artefacts of SENSE and aligned reconstruction. The proposed mcSLR method further improves the sharpness compared to SLR which suffers from blurring and shows superior image quality compared to aligned

reconstruction. The motion estimates of aligned reconstruction and mcSLR for each subdivision are shown in Fig. 5.15, and they do not show a significant difference.

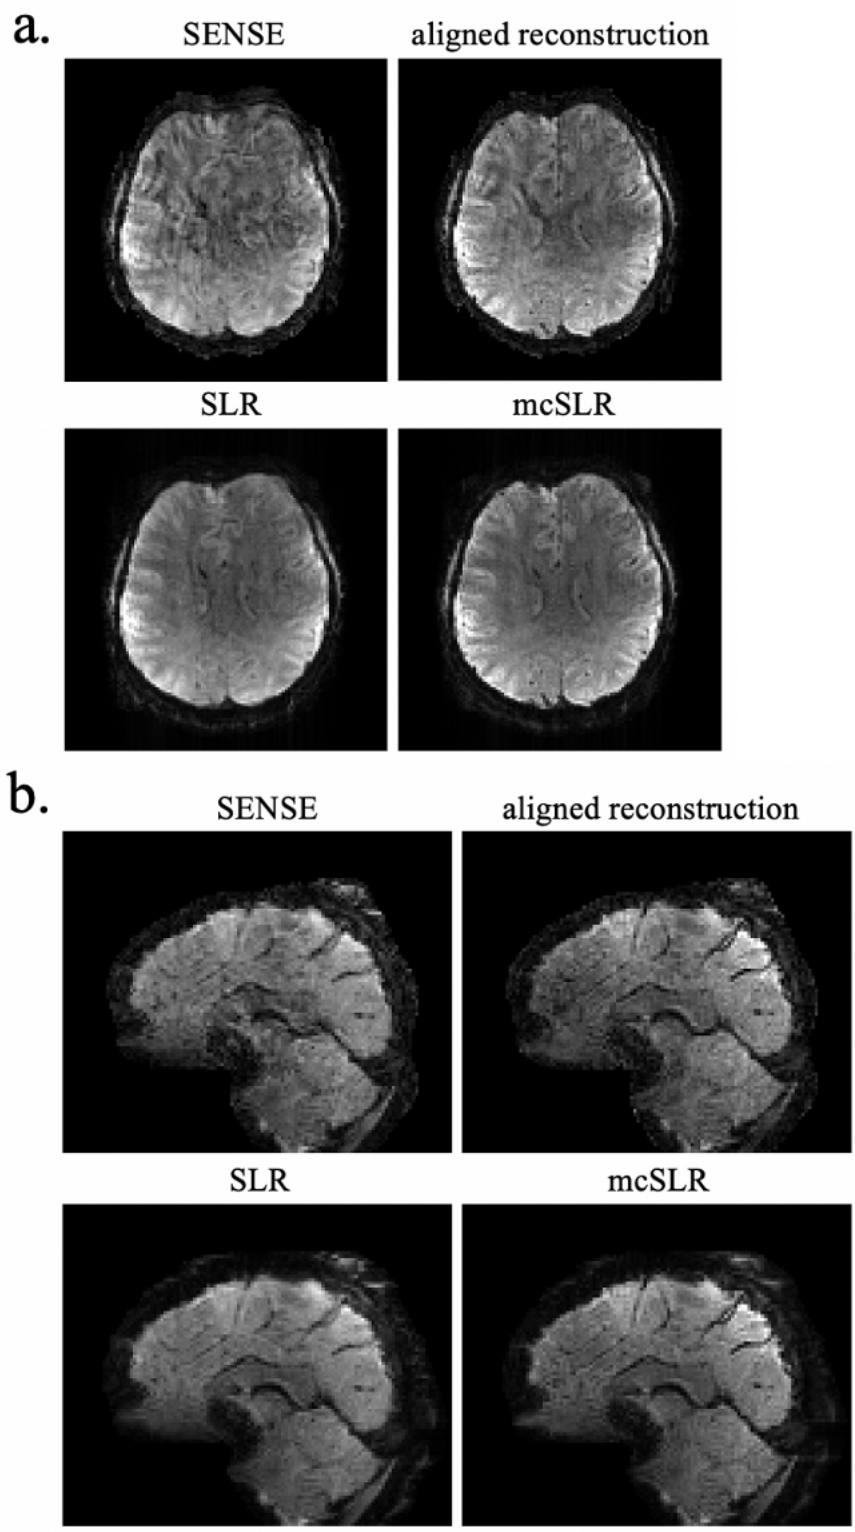

Fig. 5.14 The comparison of different reconstruction methods on another volume with medium motion. a) axial view and b) sagittal view.

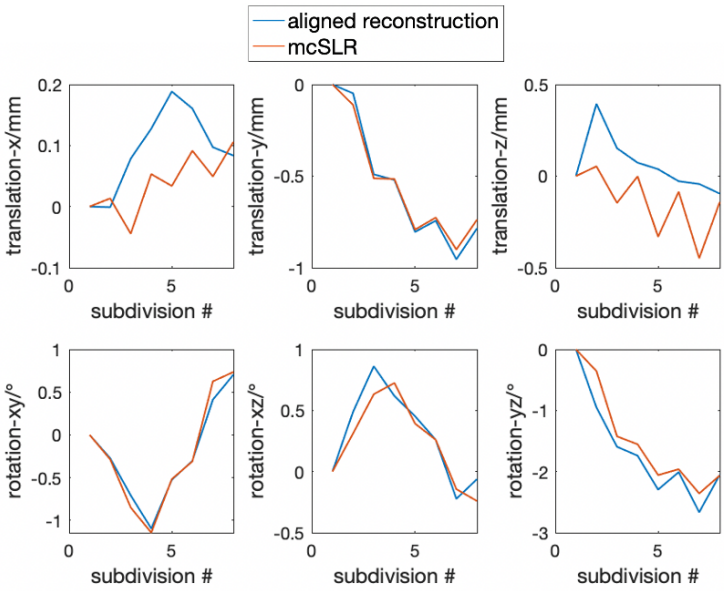

Fig. 5.15 The motion estimates of aligned reconstruction and mcSLR of the results shown in Fig. 5.14.

Fig. 5.16 shows SLR reconstruction of the same volume as Fig. 5.14 with different choices of the regularisation parameter λ. When $\lambda = 0$, it corresponds to a SENSE reconstruction for each shot group which are then SOS combined. Note it still shows fewer aliasing artefacts than the shot-combined SENSE likely due to the coherent mapping of motion and phase variations as discussed previously. Moreover, an overweighted SLR constraint ($\lambda = 3E - 1$) in the presence of motion can also introduce noticeable motion artefacts similar to SENSE as shown in Fig. 5.14. The results obtained with $\lambda = 6E - 2$ are the choice of this work.

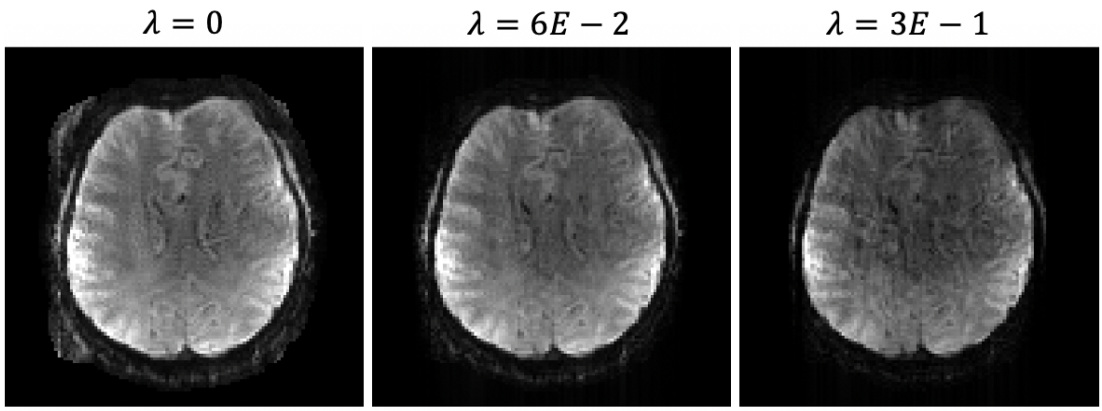

Fig. 5.16 The comparison of SLR reconstruction with different choices of the regularisation parameter λ.

Fig. 5.17 shows the results of one volume with large motion where both the SOS combined image and the image of the first shot group of SLR and mcSLR are shown. The result of aligned reconstruction shows improved sharpness compared to the SOS combined image as well as the first shot group image of SLR, but it suffers from more image artefacts especially in the cerebellum. Some fine structures are better delineated in the image of the first shot group by mcSLR compared to SLR, which indicates successful intra-shot group motion correction. The SOS combined image of SLR is blurrier than its first shot group image due to inter-shot group motion. However, the first shot group image and SOS combined image of mcSLR are in good alignment and show comparable sharpness, which indicates successful inter-shot group motion correction.

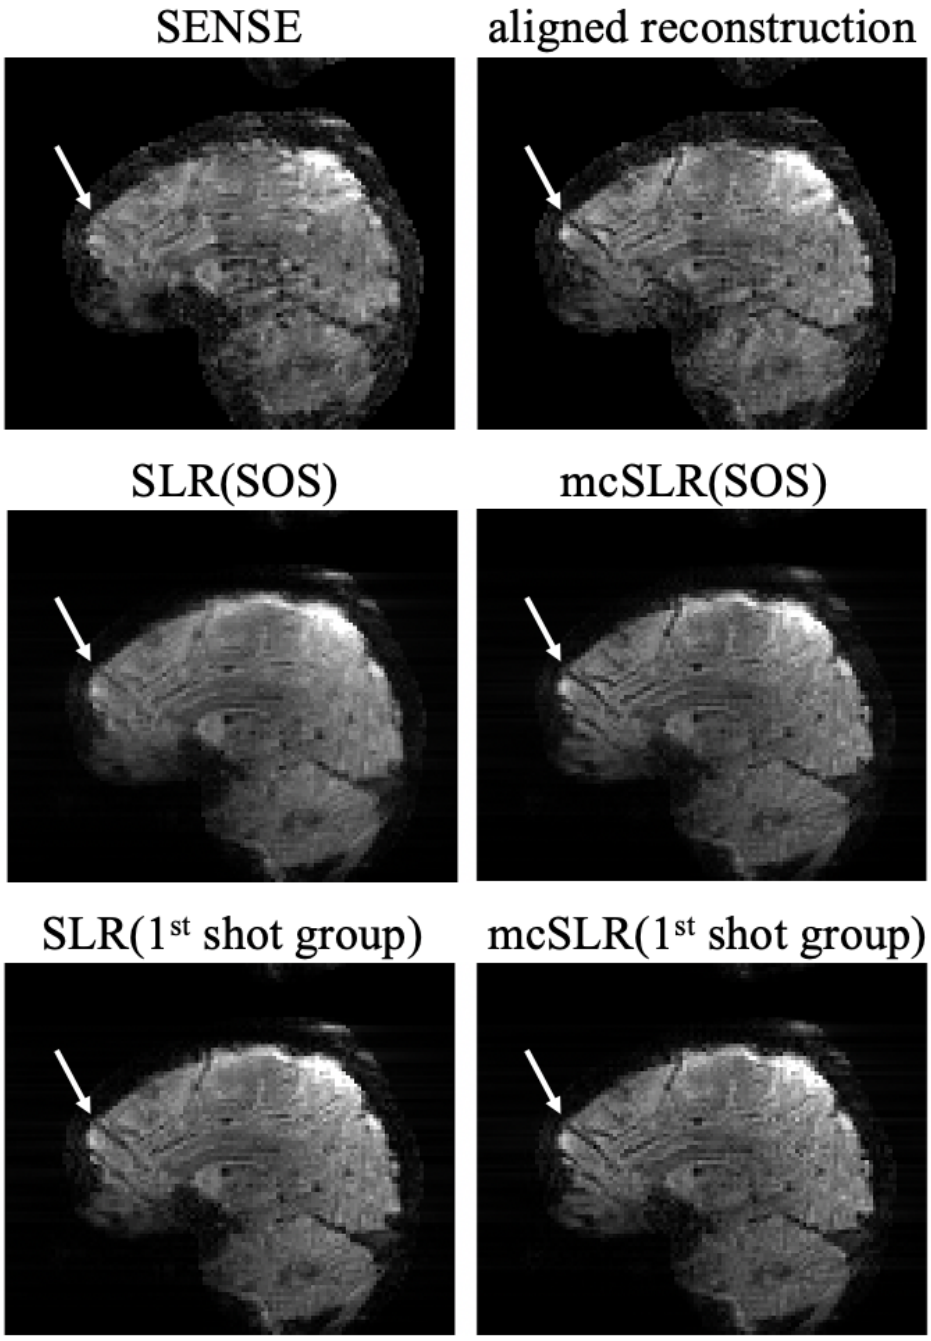

Fig. 5.17 The comparison of different reconstruction methods on one volume with large motion. The white arrow highlights a small structure with varying sharpness by different reconstruction methods.

Fig. 5.18 shows the results of another volume with large motion. In this case, aligned reconstruction failed to achieve an improvement compared to SENSE. The results of SLR reconstruction have overall better image quality but strong blurring and poor contrast between grey matter and white matter. In contrast, the results of mcSLR reconstruction

still show improved sharpness and better contrast that are necessary to disentangle small structures.

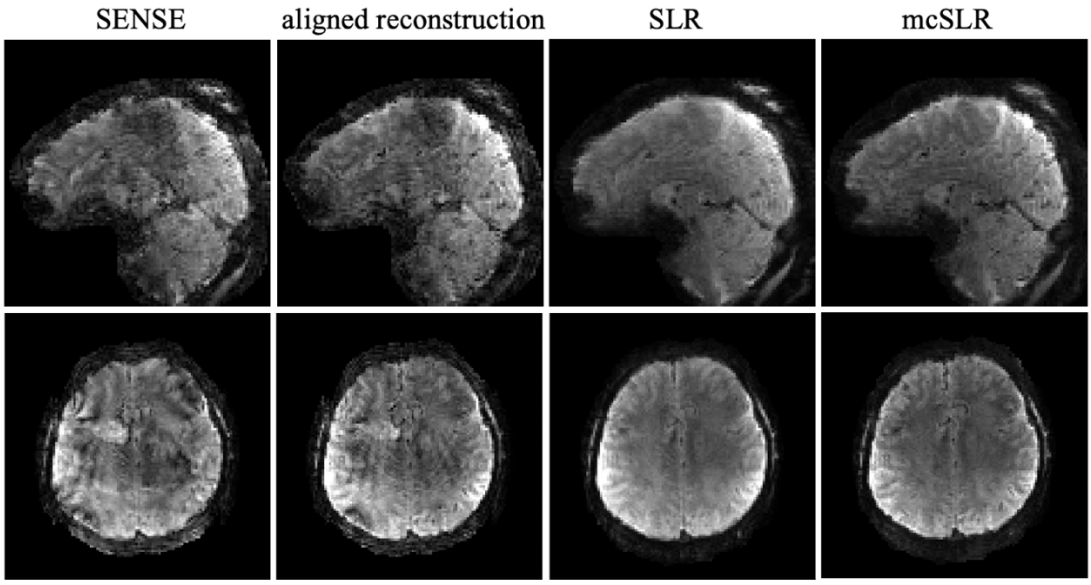

Fig. 5.18 The comparison of different reconstruction methods on another volume with large motion.

The motion estimates of aligned reconstruction and mcSLR of the results shown in Fig. 5.18 are compared in Fig. 5.19. They are in good agreement in the primary motion dimensions: rotation in the $y-z$ plane and translation along the y and z directions. To investigate the failure of aligned reconstruction, a shot-combined SENSE reconstruction with the motion estimates from mcSLR was also performed. As shown in the middle column of Fig. 5.20, it shows similar artefacts as aligned reconstruction, which also confirms the agreement of the motion estimates between aligned reconstruction and mcSLR. Besides, a SENSE reconstruction on each shot group (SOS combined later) with motion estimates from mcSLR (mcSLR with $\lambda = 0$) is also performed. As shown in the last column of Fig. 5.20, the SENSE reconstruction on each shot group, despite using the same motion estimates from mcSLR as the middle column, shows much reduced motion artefacts. The phase differences of each shot group relative to the first shot group by mcSLR reconstruction are shown in Fig. 5.21, where the phase variations between

different shot groups can be observed. These experiments show that the artefacts in aligned reconstruction are likely due to incoherent motion mapping and phase variations.

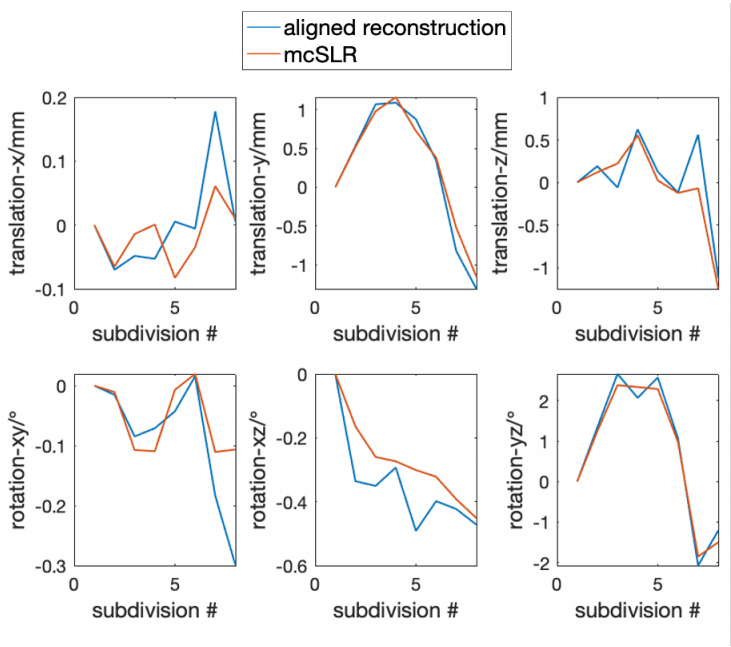

Fig. 5.19 The motion estimates of aligned reconstruction and mcSLR of the results shown in Fig. 5.18.

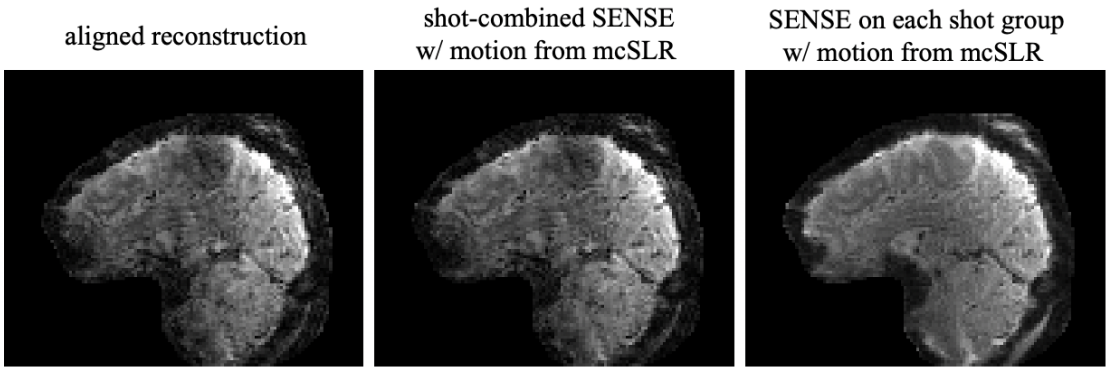

Fig. 5.20 The comparison of aligned reconstruction, shot-combined SENSE reconstruction, and SENSE reconstruction on each shot group. The latter two both used the same motion estimates from mcSLR.

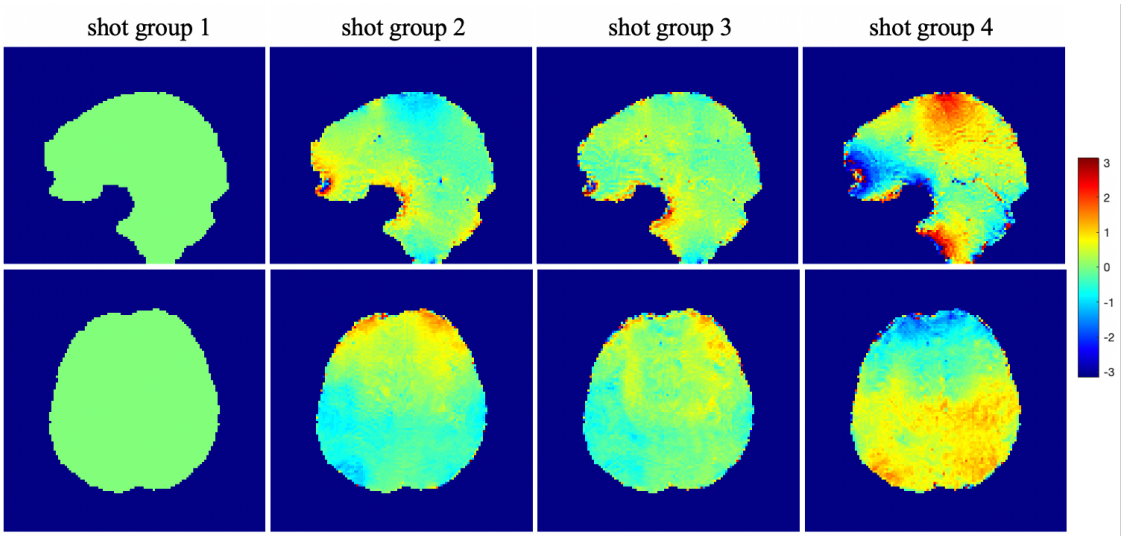

Fig. 5.21 The phase differences of each shot group relative to the first shot group by mcSLR reconstruction of the results shown in Fig. 5.18.

Fig. 5.22 shows the reconstruction results of a different subject with medium motion. In this case, aligned reconstruction and mcSLR show improved sharpness compared to SENSE and SLR.

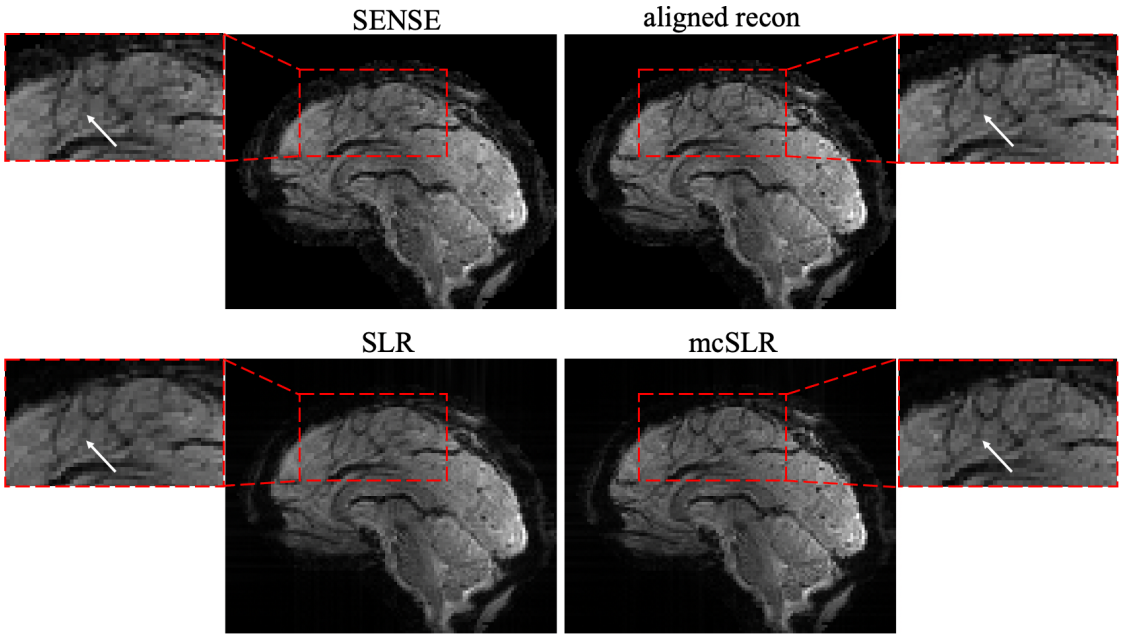

Fig. 5.22 The comparison of different reconstruction methods on a different subject. The white arrow highlights a small structure with varying sharpness by different reconstruction methods.

5.4 Discussion

In this work, a motion compensated structured low-rank reconstruction framework for 3D multi-shot EPI is proposed, which improves its robustness to both inter-shot motion and inter-shot phase variations. Simulation and in vivo experiments have shown that the proposed method can reduce image artefacts and improve the temporal stability of 3D multi-shot EPI in the presence of rigid head motion and physiologically induced phase variations.

The proposed solution to inter-shot inconsistencies is built on the SLR framework which deals with inter-shot phase variations as discussed in Chapter 4. The basic idea of SLR reconstruction for 3D multi-shot EPI is to reconstruct an individual image for each shot group which has a much-reduced time window of signal integration compared to the entire volume. To deal with the higher under-sampling factor of each shot group, SLR method jointly reconstructs all shot groups to exploit their linear dependencies which is formulated as the low-rank property of the block-Hankel structured matrix constructed from the multi-shot group k-space. This low-rank property is based on the assumption that different shot groups only differ in image phase while having consistent image magnitude. This assumption is clearly no longer true in the presence of inter-shot motion. Thus, SLR reconstruction with motion compensation can not only reduce motion artefacts, but also improve the validity of the structured low-rank constraint for the correction of phase variations.

The proposed mcSLR framework considers motion and phase variations at different time scales. The head motion is considered at the shot-to-shot (intra-shot group) level, as motion can be less temporally coherent compared to the physiologically induced phase variations. Also, the characterisation of each motion state only needs six rigid motion parameters instead of an individual image. The physiologically induced phase variations

are assumed to be temporally coherent, so they are considered at the inter-shot group level to improve the conditioning of the reconstruction, as the SLR framework uses a model-free characterisation of phase variations such that each state of phase variations needs to be specified by an individual image. It is worth mentioning that bulk motion can induce B_0 field fluctuations and image phase variations as well. The SLR formulation automatically accounts for any spatially smooth phase variations between shot groups, so any motion induced phase variations are also compensated to some degree. However, the phase variations at the intra-shot group level are not considered currently, which could be a direction of future work. The work proposed recently (Brackenier et al., 2022) has considered motion induced B_0 variations to improve the performance of aligned reconstruction for anatomical imaging at 7T. This method uses a physics-inspired B_0 model which describes the B_0 variation map as a linear combination of the pitch (around left–right axis) and roll (around posterior-anterior axis) rotation angles, and thus only two additional linear coefficient maps instead of a B_0 variation map for each motion state need to be estimated. Similar approaches might be considered to improve the proposed method further by modeling motion induced intra-shot group phase variations.

In accordance with the formulation of SLR reconstruction, inter-shot motion is modeled using a hierarchical strategy, which is decomposed into inter-shot group motion and intra-shot group motion. This motion decomposition makes it easier to estimate motion at different levels by different methods. The inter-shot group motion can be estimated robustly by image registration-based methods, and the method FLIRT (Jenkinson et al., 2002) which is dedicated to brain imaging was used in this work. Within each shot group which has a reduced time window and thus less integration of signal variations compared to the whole volume, the intra-shot group motion is estimated by solving the optimisation with data consistency constraint that exploits sensitivity

encoding redundancy as in aligned reconstruction.

The proposed mcSLR reconstruction was also used in conjunction with the seg-CAIPI sampling pattern as in conventional SLR reconstruction. Previous work has shown that the seg-CAIPI data with the interleaved ordering along k_z direction can have lower tSNR than the sequentially sampled data if conventional shot-combined SENSE reconstruction is used, whereas the seg-CAIPI data with SLR reconstruction is able to achieve a higher tSNR than the sequentially sampled data with SENSE reconstruction. Similarly, the seg-CAIPI sampling can also lead to more severe motion artefacts compared to sequential ordering if a simple shot-combined SENSE reconstruction is used, as motion is mapped incoherently across k-space. However, this is not responsible for the bad performance of aligned reconstruction when motion is strong, as the higher sensitivity of incoherent sampling to motion can benefit motion estimation as demonstrated in previous work (Cordero-Grande et al., 2016; Cordero-Grande et al., 2020).

The long computation time of the 3D reconstruction problem is one limitation of the proposed method. Currently the reconstruction of a 3D volume can take more than one day with multiple CPUs on our compute cluster. The conventional 3D SLR reconstruction takes about 45 minutes per iteration, and the mcSLR reconstruction takes about 60 minutes per iteration, which corresponds to 1/3 longer reconstruction time given the same prescribed number of iterations is used. Besides, the initialisation of motion parameters by aligned reconstruction also takes typically more than 4 hours on motion-corrupted data and about 3 hours on motion-free data. Thus, the optimisation of computational efficiency is a major task of future work. Specifically, the generic iterative reweighted annihilation filter (GIRAF) algorithm (Ongie and Jacob, 2017) which is SVD-free and matrix lifting-free is worthy of investigation (Mani et al., 2020). In addition, some automatic motion evaluation and classification procedures in the upstream pipeline could also be helpful to

reduce the overall computational time of the time course by determining on which volumes mcSLR reconstruction is needed.

5.5 Summary

In this chapter, a motion compensated structured low-rank (mcSLR) reconstruction method for robust 3D multi-shot EPI is proposed, which jointly models inter-shot motion and inter-shot phase variations. The proposed method improves the motion robustness of SLR framework which accounts for phase variations, and thus reduce motion artefacts and improve the temporal stability of 3D multi-shot EPI.

6 Locally Structured Low-rank Reconstruction

6.1 Introduction

As discussed in previous chapters, structured low-rank image reconstruction methods (Haldar and Setsompop, 2020; Jacob et al., 2020) have been successfully employed in a variety of applications such as phase-constrained imaging (Haldar, 2014), calibration-less parallel imaging reconstruction (Haldar and Zhuo, 2016; Shin et al., 2014a), multi-shot EPI reconstruction (Chen et al., 2023; Mani et al., 2017), EPI Nyquist ghost correction (J. Lee et al., 2016; Lobos et al., 2021, 2018), etc. SLR image reconstruction methods are typically based on block-Hankel structured low-rank matrix completion and provide a

powerful framework that offers great flexibility to exploit different forms of linear dependencies in MRI data to constrain image reconstruction. More specifically, the low-rank matrix in SLR methods is typically a structured matrix which consists of a single or multiple concatenated block-Hankel structured matrix representations of k-space data. The low-rank property of this block-Hankel structured matrix is validated through the existence of a non-trivial null-space and the low dimensional column space (assuming each kernel is vectorised into a row vector), arising from annihilation relations that describe signal properties of MRI data. Accordingly, the low-rank matrix in SLR methods can be constructed in different ways to exploit different signal properties, e.g., the C-matrix, S-matrix (Haldar, 2014), or VC-matrix (Bilgic et al., 2019; Kim and Haldar, 2018). A detailed discussion on these different structured low-rank matrix constructions can be found in Section 3.4.1 and 3.4.2.

What most structured low-rank methods share in common is the global enforcement of the low-rank constraint on the entire block-Hankel structured matrix. This implementation promotes shift-invariant annihilation relations in k-space, such that all row vectors across the block-Hankel structured matrix share the same null space vectors and all local patches across k-space share the same annihilation filters. Note as the kernel size of block-Hankel transformation is typically chosen to be smaller than 10×10, the number of row vectors of the block-Hankel structured matrix is typically more than 100 times larger than the number of column vectors, assuming the number of rows counts the number of possible patch positions, and the number of columns is the vectorised kernel dimension.

The method STEP (Zhou et al., 2016) which is an alternative to the SLR method SAKE (Shin et al., 2014b) for calibration-less parallel imaging, is an exception to the global structured low-rank method as discussed above. It divides k-space into 5×5

overlapping blocks and enforces structured low-rank constraints on each of them respectively to adapt for spatially varying signal characteristics. Similarly, as a type of conventional low-rank reconstruction discussed in Section 3.3.3, locally low-rank reconstruction methods build a series of low-rank matrices from local patches of the image, which can be useful in some cases where the global low-rank model does not apply, or potentially outperform global low rank methods when the row and column dimensions of the Casorati matrix differ significantly (Trzasko and Manduca, 2012).

In this chapter, a new SLR matrix recovery method, termed locally structured low-rank (LSLR) is proposed, which relaxes the single low-rank constraint on the full block-Hankel structured matrix by enforcing low-rank constraints on its submatrices with each consisting of a subset of row vectors. This LSLR method provides a sufficiently flexible modification which works with virtually any conventional SLR methods. In this work, simulation experiments based on numerical phantoms, experimental T_1 weighted and GRE data have been performed to validate the superior performance of the proposed LSLR method compared to conventional SLR method across a wide range of imaging settings.

6.2 Methods

Without loss of generality, the LSLR method is demonstrated for a simple 2D calibration-less parallel imaging reconstruction as a representative example with straightforward extensions to other applications. In this problem, the cost function of SLR reconstruction with a non-convex low-rank constraint is formulated as:

$$\begin{aligned} &\min\|AX - Y\|^2 \\ &s.t. rank(H_c X) = r \end{aligned} \qquad \text{Eq. 6.1}$$

Where $X \in \mathbb{C}^{n^2 N_c}$ corresponds to the k-space data of a $n \times n$ image with N_c channels,

and the forward measurement model A: $\mathbb{C}^{n^2 N_c} \rightarrow \mathbb{C}^{N_k N_c}$ represents the sampling operator which selects N_k k-space samples from each channel. $Y \in \mathbb{C}^{N_k N_c}$ denotes the acquired k-space data, so the under-sampling factor is defined as N_k/n^2. Let H: $\mathbb{C}^{n^2} \rightarrow \mathbb{C}^{l \times d^2}$ represent a basic block-Hankel transformation operator with a $d \times d$ kernel, so that HX is constructed by vertically concatenating all $l = (n - d + 1)^2$ row vectors corresponding to $d \times d$ patches of X. Note here H is defined in the form of a C-matrix, and it can also be defined as the S-matrix or VC-matrix which are of different sizes. The augmented operator H_c: $\mathbb{C}^{n^2 N_c} \rightarrow \mathbb{C}^{l \times d^2 N_c}$ further concatenates the block-Hankel transformation of each channel horizontally along the column dimension. Note as mentioned before, typically the number of rows l of $H_c X$ is far larger than its number of columns $d^2 N_c$. A non-convex strict rank constraint with a rank parameter $r < d^2 N_c$ is used to enforce the low-rank property on $H_c X$.

The conventional SLR approach enforces the low-rank property on the full block-Hankel structured matrix, which corresponds to a single set of null space vectors (or equivalently annihilation filters) for the block-Hankel structured matrix. In contrast, the proposed LSLR method formulates the reconstruction as:

$$\begin{gathered} \min \|AX - Y\|^2 \\ s.t. rank(\Gamma_i H_c X) = r \\ \forall \Gamma_i \in \Omega = \{\Gamma_1, \Gamma_2, \dots, \Gamma_m\} \end{gathered} \qquad \text{Eq. 6.2}$$

Where Γ_i: $\mathbb{C}^{l \times d^2 N_c} \rightarrow \mathbb{C}^{s \times d^2 N_c}$ is the operator that selects $s = l/m$ consecutive rows from $H_c X$ to form the i_{th} submatrix, as shown in Fig. 6.1. Accordingly, the adjoint operator Γ_i^*: $\mathbb{C}^{s \times d^2 N_c} \rightarrow \mathbb{C}^{l \times d^2 N_c}$ places the submatrix rows back to their original positions with all other row vectors being 0. m is the number of non-overlapping submatrices. When l/m is not an integer, s is typically rounded to be the smallest integer

close to it for simplicity. Ω is the set which consists of m operators selecting all non-overlapping submatrices of H_cX, e.g., Γ_1 selects the rows of index $1{:}\,s$ and Γ_2 selects the rows of index $s+1{:}\,2s$, etc. Thus, when $m=1$ and $s=l$, LSLR reconstruction is the same as SLR reconstruction. As the submatrices select a subset of row vectors of the block-Hankel structured matrix, the rank of each submatrix is bounded by the rank of the global block-Hankel structured matrix, and therefore is still low-rank given a sufficient number of rows are selected. This formulation relaxes the single low-rank constraint over H_cX to a set of low-rank constraints on m disjoint submatrices.

The optimisation problem of LSLR is solved by iterative algorithms as in global SLR reconstruction. To promote shift-invariance and avoid potential boundary artefacts, a cycle-spinning procedure is employed which randomly shifts the positions of the submatrices between $[1, l]$ at every iteration of the optimisation algorithm, as shown in Fig. 6.2. The proposed approach can also be viewed as a stochastic algorithm that iterates over randomly chosen subsets of all possible $s \times d^2 N_c$ submatrices.

In this work, the optimisation of LSLR is solved by ADMM algorithm (Boyd, 2010), which reformulates Eq. 6.2 as:

$$\begin{gathered}\underset{X,Z,U}{\operatorname{argmin}} \|AX - Y\|_2^2 + \lambda \|H_cX - Z + U\|_2^2 \\ s.t.\, rank(\Gamma_i Z) = r \\ \forall \Gamma_i \in \Omega = \{\Gamma_1, \Gamma_2, \ldots, \Gamma_m\}\end{gathered} \qquad \text{Eq. 6.3}$$

Note a different set Ω is regenerated at each iteration. Details of the algorithm are shown in Table 6.1. The optimisation of SLR is solved in the same way with $m=1$, except that a fixed Ω is used at every iteration.

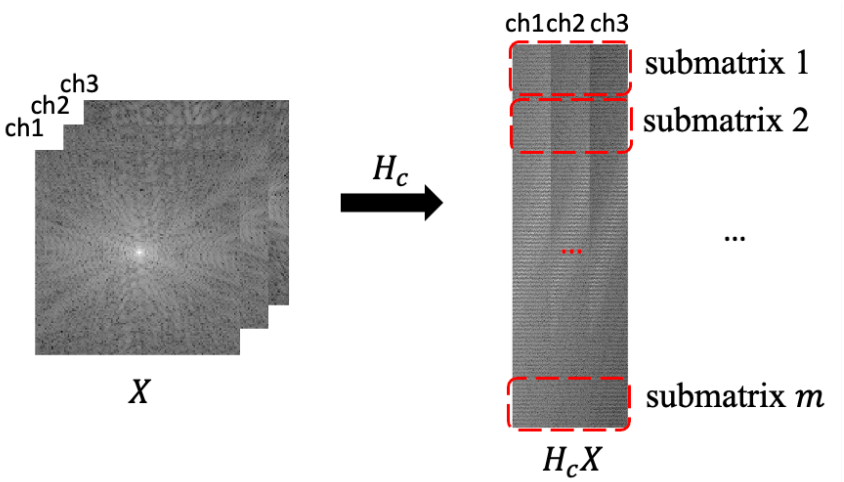

Fig. 6.1 Schematic submatrix selection on the block-Hankel structured matrix. Each submatrix consists of a subset of consecutive rows.

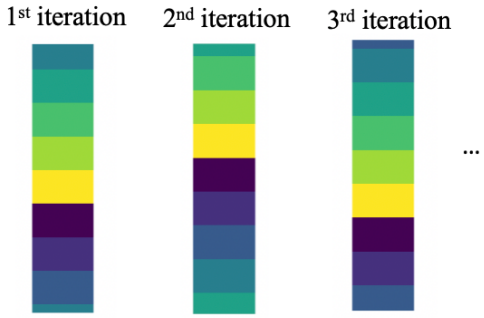

Fig. 6.2 Schematic cycle-spinning procedure. Each colour represents the row vectors belonging to the same submatrix. The positions of submatrices are randomly shifted at each iteration. When the continuous row index of a submatrix exceeds the number of rows, it wraps around and restarts from the first row.

Step 1. initialise $X_0 = A^H Y,\ Z_0 = H_c X_0, U_0 = 0, j = 0$
Step 2. solve subproblem a, b, c alternately at the j_{th} iteration. a. Update X_{j+1}: $X_{j+1} = \left(A^H A + \lambda {H_c}^H H_c\right)^{-1}(A^H Y + \lambda {H_c}^H (Z_j - U_j))$ b. Update Z_{j+1}: 1) Generate a new Ω_j, and then for each $\Gamma_i \in \Omega_j$: Calculate SVD: $\Gamma_i\left(H_c X_{j+1} + U_j\right) = U_i \Sigma_i V_i^H$ Perform hard threshold truncation on Σ_i so that $\sigma_k = 0$ if $k > r$, where σ_k is the k_{th} largest singular value along the diagonal of Σ_i 2) $Z_{j+1} = \sum_{\Gamma_i \in \Omega_j} \Gamma_i^* (U_i \Sigma_i V_i^H)$ c. Update U_{j+1}: $U_{j+1} = H_c X_{j+1} - Z_{j+1} + U_j$
Step 3. Set $j = j + 1$, repeat Step 2 until a prescribed number of iterations or a stopping criterion is reached

Table 6.1 ADMM algorithm for solving LSLR.

The LSLR method can also be adapted to structured low-rank tensor reconstruction for higher dimensional datasets. Here it is demonstrated for a calibration-less, multi-echo joint reconstruction problem. The cost function of the locally structured low-rank tensor reconstruction using a non-convex formulation is shown as:

$$\begin{aligned} &min\|AX - Y\|^2 \\ &s.t. rank(\Gamma_i H_1 X) = r1,\ \forall \Gamma_i \in \Omega_1 \\ &rank\left(\Gamma_j H_2 X\right) = r2,\ \forall \Gamma_j \in \Omega_2 \end{aligned} \qquad \text{Eq. 6.4}$$

Where $X \in \mathbb{C}^{n^2 N_c N_e}$ corresponds to the k-space data of a $n \times n$ image with N_c channels and N_e echoes. $Y \in \mathbb{C}^{N_k N_c N_e}$ denotes the acquired k-space data and A: $\mathbb{C}^{n^2 N_c N_e} \rightarrow \mathbb{C}^{N_k N_c N_e}$ represents the sampling operator. As defined previously, H_c: $\mathbb{C}^{n^2 N_c} \rightarrow \mathbb{C}^{l \times d^2 N_c}$ generates the block-Hankel structured matrix from multi-coil k-space, and then a low-

rank tensor $\mathbb{C}^{l\times d^2N_c\times N_e}$ can be constructed by concatenating N_e block-Hankel structured matrices in a third dimension. This tensor can be unfolded into a low-rank matrix in two ways, as shown in Fig. 6.3. The operator H_1: $\mathbb{C}^{n^2N_cN_e} \rightarrow \mathbb{C}^{l\times d^2N_cN_e}$ generates the one-mode unfolding matrix by concatenating the block-Hankel structured matrix representation of each echo along the column (channel) dimension, and H_2: $\mathbb{C}^{n^2N_cN_e}$ $\rightarrow \mathbb{C}^{lN_e\times d^2N_c}$ generates the two-mode unfolding matrix by concatenating them along the row (patch) dimension. $r1$ and $r2$ are the rank parameters, and Γ_i and Γ_j are the submatrix selection operators for each unfolding matrix respectively. As shown in Fig. 6.3, Γ_i: $\mathbb{C}^{l\times d^2N_cN_e} \rightarrow \mathbb{C}^{s\times d^2N_cN_e}$ selects a submatrix of s consecutive rows from H_1X. Ω_1 is the partition consisting of all possible non-overlapping Γ_i-selecting operators. Γ_j: $\mathbb{C}^{lN_e\times d^2N_c} \rightarrow \mathbb{C}^{s\times d^2N_c}$ selects a submatrix of of s/N_e consecutive rows from the block-Hankel structured matrix representation of each echo (corresponding to the same k-space patches of each echo) and vertically concatenates them together. Ω_2 is the partition consisting of all possible non-overlapping Γ_j-selecting operators. As H_1X and H_2X are assumed to be low-rank and have a low dimensional column space, each submatrix should also be low-rank. The cycle-spinning procedure is also used in the LSLR tensor reconstruction that the positions of the submatrices are randomly shifted at every iteration. The optimisation is solved by the ADMM algorithm in similar ways with adaption to the two low-rank constraints.

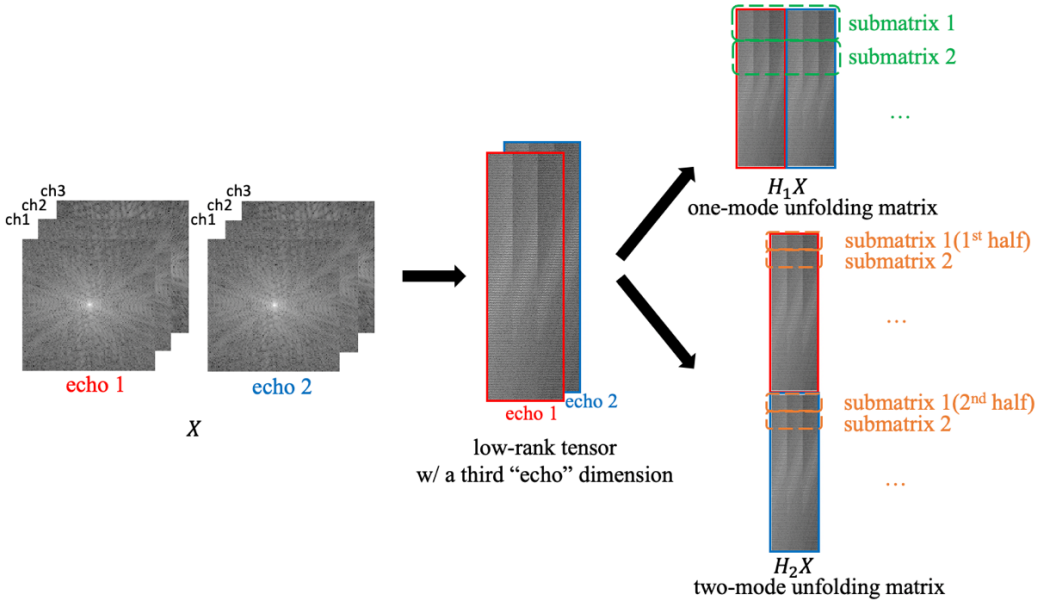

Fig. 6.3 Schematic submatrix selection on different unfolding matrices of the block-Hankel structured tensor. For the one-mode unfolding matrix which treats different echoes as virtual channels, each submatrix consists of a subset of consecutive rows. For the two-mode unfolding matrix which treats different echoes as virtual patches, each submatrix consists of two subsets of consecutive rows corresponding to k-space patches at the same locations from two echoes.

6.3 Experiments

Simulation experiments based on numerical phantoms and experimental T_1 weighted and GRE data were performed to validate the performance of LSLR. For each numerical phantom, the single coil k-space datasets with added complex, zero-mean Gaussian noise were generated. The maximum k-space energy is normalised to 1, and the noise level is specified by its standard deviation. The matrix size of the numerical phantom is 80×80. For the experimental T_1 weighted and GRE data, 8 synthetic channels which were compressed from 32-channel source data were used, and the matrix size was down-sampled to be 128×128 for higher computational efficiency. The GRE data have 6 echoes in total and the first two echoes were used for reconstruction. The reconstruction performance was evaluated using NRMSE. For the experimental data, the fully sampled reconstruction was used as a proxy for the ground truth. The rank parameter

r was hand tuned to produce the lowest NRMSE for SLR and the same rank was used for most LSLR reconstructions for simplicity, except when a different optimal rank can produce a notably lower NRMSE. The number of submatrices m was also optimised for the lowest NRMSE for LSLR, which is denoted in parentheses as LSLR(m). The ADMM penalty parameters λ were 10^{-10} and 10^{-4}, and the kernel sizes of the block-Hankel transformation operator were 9×9 and 5×5 for the numerical phantom and experimental data respectively.

6.3.1 Numerical Phantom

A numerical simulation experiment was performed to compare different implementation choices of LSLR, including the submatrix selection modes and the used of cycle-spinning procedure. As shown in Fig. 6.4, a Shepp-Logan phantom with smooth phase and a 2D random under-sampling mask (50.5% under-sampling) were used to generate the simulation data. The standard deviation of the complex Gaussian noise is 10^{-3}. The C-matrix style block-Hankel structured matrix was used in SLR and LSLR reconstruction.

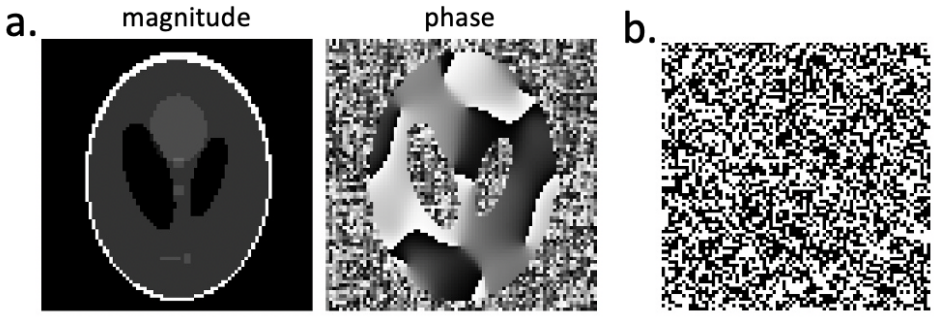

Fig. 6.4. Simulation setup. a) The ground truth image. b) The 2D random sampling mask.

Implementation Choice - Submatrix Selection Mode

Fig. 6.5 shows three different ways to select submatrices from the block-Hankel structured matrix. As introduced above, block mode selects each submatrix as a

continuous block. In contrast, stripe mode selects discontinuous and equally distributed row vectors as a submatrix, and random mode selects a random subset of row vectors (non-overlapping) as a submatrix. The positions of submatrices are fixed through iteration for stripe mode whereas random mode uses the same cycle spinning procedure as block mode. Fig. 6.6 compares the reconstruction performance of the three submatrix selection modes. The block mode shows a significantly better performance than the other two and is used as the choice of LSLR. This is likely because the adjacent row vectors are also close in the original k-space, which share more similar signal characteristics than those k-space data which are further apart.

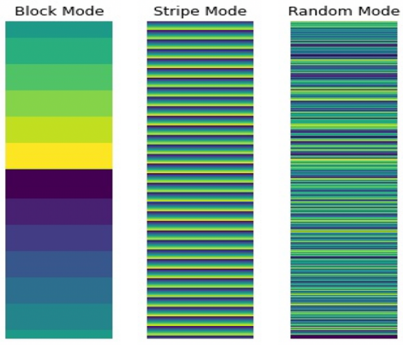

Fig. 6.5 Three submatrix selection modes. Each colour represents the row vectors belonging to the same submatrix.

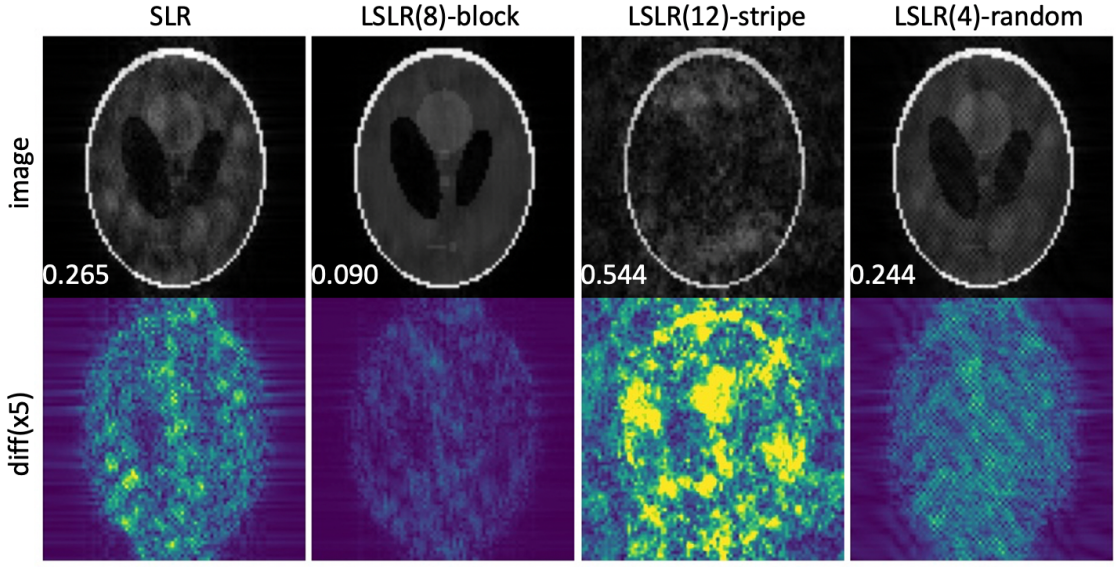

Fig. 6.6. The comparison of the reconstruction performance of different submatrix selection modes. The reconstructed images (top) and their error maps compared to the ground truth (bottom) are shown. The NRMSE values are shown in each image respectively. $r = 70$ was used.

Implementation Choice - Cycle-spinning

Fig. 6.7 compares the reconstruction performance of LSLR using the cycle-spinning procedure and using fixed submatrix locations. It is shown that the cycle-spinning procedure leads to a substantial improvement.

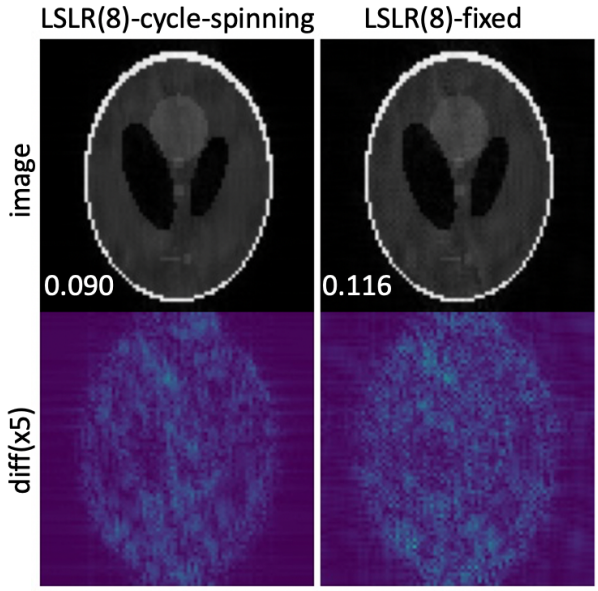

Fig. 6.7. The comparison of the reconstruction performance of LSLR using cycle-spinning procedure and using fixed submatrix locations. The reconstructed images (top) and their error maps compared to the ground truth (bottom) are shown. The NRMSE values are shown in each image respectively. $r = 70$ was used.

The performance of LSLR was validated across a variety of conditions, including different styles of block-Hankel structured matrices, different thermal noise levels, different levels of rank-deficiency of the block-Hankel structured matrix and on a white noise-like phantom. The details of each experiment are shown below.

Performance Validation - Different Block-Hankel Structured Matrices

The performance of LSLR reconstruction was validated for different styles of block-Hankel structured matrices, including C-matrix, S-matrix and VC-matrix. A $3/4$ partial Fourier under-sampling was used in this experiment and the results are shown in Fig. 6.8. The reconstructions using the S-matrix and VC-matrix have comparable performance, which is consistent with the fact that they are different implementations exploiting the

same image property of smooth phase and limited support. The reconstruction using C-matrix has the worst performance as only the limited support prior is utilized with this formula. In all three cases, LSLR reconstruction achieves a much lower NRMSE than SLR, as clearly indicated in the error map which shows the difference compared to the ground truth.

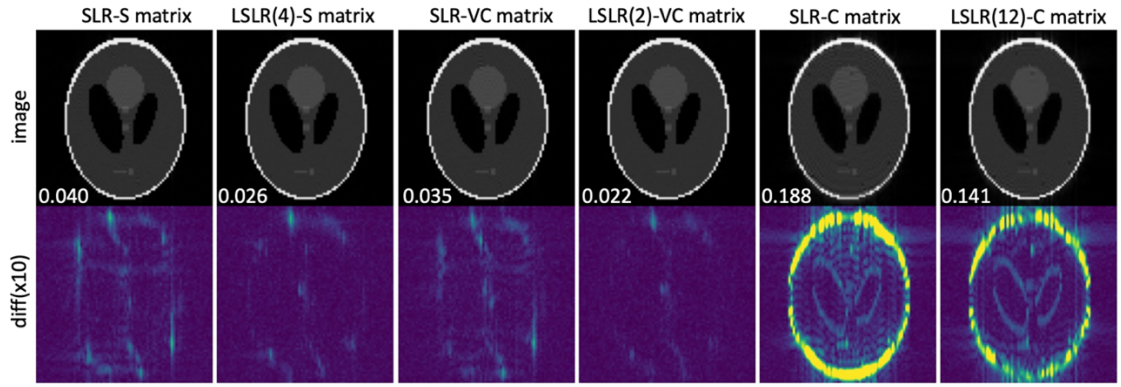

Fig. 6.8. The comparison of SLR and LSLR reconstructions with different low-rank matrix constructions. The reconstructed images (top) and their error maps compared to the ground truth (bottom) are shown. The NRMSE values are shown in each image respectively. $r = 130$ was used for the S-matrix. $r = 135$ was used for the VC-matrix. $r = 45$ was used for the C-matrix.

Performance Validation - Different Noise Levels

The performance of LSLR reconstruction was also validated across different noise levels, corresponding to $std = \{0,1,3,6,9\} \times 10^{-3}$. The S-matrix and partial Fourier sampling were used for SLR and LSLR reconstructions in this experiment. The NRMSE of SLR and LSLR reconstructions at different Gaussian noise levels are compared in Table 6.2. The results of LSLR reconstruction have lower NRMSE than SLR reconstruction across different noise levels with $std = \{0,1,3,6\} \times 10^{-3}$, and the differences between two approaches get smaller as noise increases until they achieve comparable NRMSE at $std = 9 \times 10^{-3}$.

std of noise	0	10^{-3}	3×10^{-3}	6×10^{-3}	9×10^{-3}
NRMSE of SLR	0.020	0.040	0.062	0.098	0.137
NRMSE of LSLR	**0.013**	**0.026**	**0.050**	**0.094**	**0.136**

Table 6.2. The comparison of SLR and LSLR reconstructions under different noise levels. LSLR with 2 submatrices was used at $std = 9 \times 10^{-3}$ and LSLR with 4 submatrices was used for all other datasets. $r = 130$ was used for $std = \{0,1\} \times 10^{-3}$, $r = 120$ was used for $std = \{3,6\} \times 10^{-3}$ and $r = 110$ was used for $std = 9 \times 10^{-3}$.

Performance Validation - Different Levels of Rank-deficiency Due to Varying Image Support

Four different numerical phantoms with varying image support were generated to simulate different levels of rank deficiency of the block-Hankel structured matrix. The C-matrix and the same random under-sampling in Fig. 6.4b were used in this experiment. As shown in Fig. 6.9, from left to right, the image support increases and the rank of the C-matrix also increases. The reconstruction results are shown in Fig. 6.10. LSLR does not achieve an improvement compared to SLR on the first phantom, where the rank of the C-matrix is the lowest. As the rank of the C-matrix increases from left to right on different phantoms, the differences between the NRMSE of SLR and LSLR increase. This suggests that the improvement of LSLR is likely due to the increased rank of the global block-Hankel structured matrix, whereas the relaxed low-rank constraints on submatrices of LSLR are able to fit the data better when the global low-rank model is not strictly true, as shown on the last phantom in this experiment.

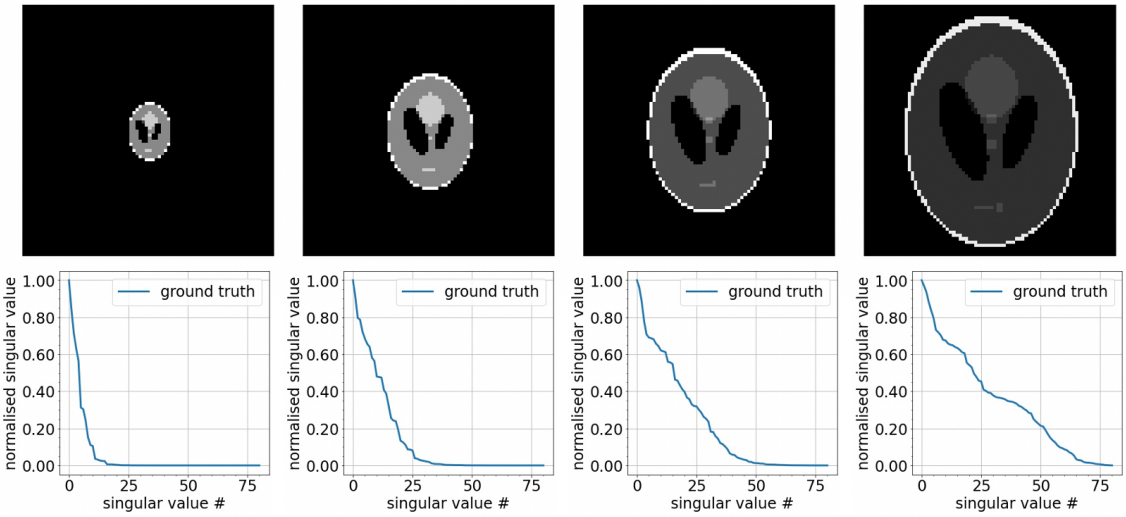

Fig. 6.9. The ground truth images and normalised singular value distributions of the C-matrix of four different simulation datasets. From left to right, the image support of the numerical phantom increases and the rank of the C-matrix also increases.

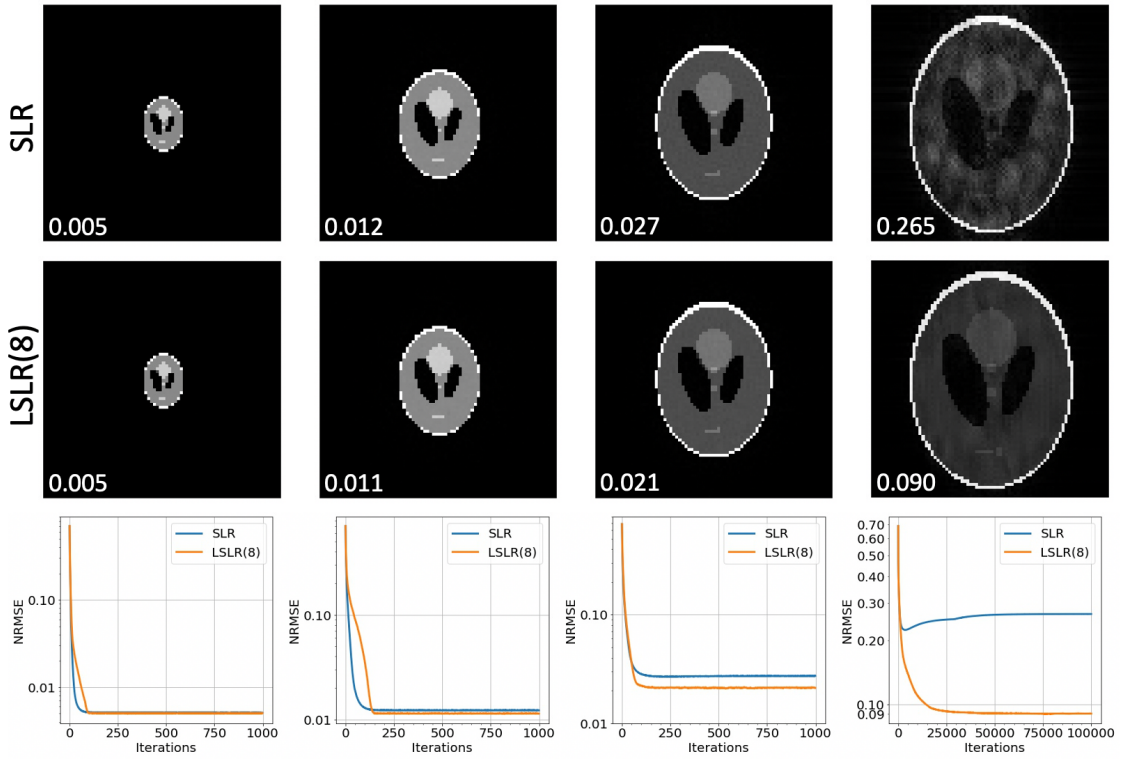

Fig. 6.10. The comparison of SLR and LSLR reconstructions on the four simulation datasets shown in Fig. 6.9. The reconstructed images and the plots of NRMSE through iterations are shown. The NRMSE values are shown in each image respectively. $r = 35$ was used for the first phantom (from left to right). $r = 40$ was used for the second phantom. $r = 45$ was used for the third phantom. $r = 70$ was used for the fourth phantom.

Performance Validation - Different Levels of Rank-deficiency Due to Varying Kernel Sizes

The level of rank deficiency of the block-Hankel structured matrix was also modulated by changing the kernel size. The C-matrix and the third phantom in Fig. 6.9 was used in this experiment. As shown in Fig. 6.11, the C-matrix is more rank-deficient, i.e., the singular values decrease faster, as the kernel size increases. The reconstruction results are shown in Fig. 6.12. LSLR achieves a lower NRMSE than SLR in all three cases. The performance of SLR gets better as kernel size increases, and its lowest NRMSE is achieved with the biggest 21×21 kernel. However, LSLR basically has already achieved its best performance with a 15×15 kernel, with a marginal further improvement by a larger 21×21 kernel, and it is better than the best of SLR. Note the lowest NRMSE of SLR with an even larger 27×27 kernel is 0.0187, which is still worse than LSLR with just a 15×15 kernel.

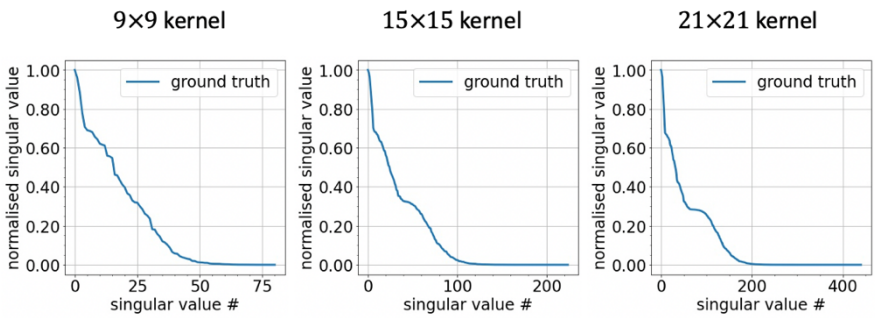

Fig. 6.11. The normalised singular value distributions of the C-matrix with different kernel sizes of the third phantom shown in Fig. 6.8.

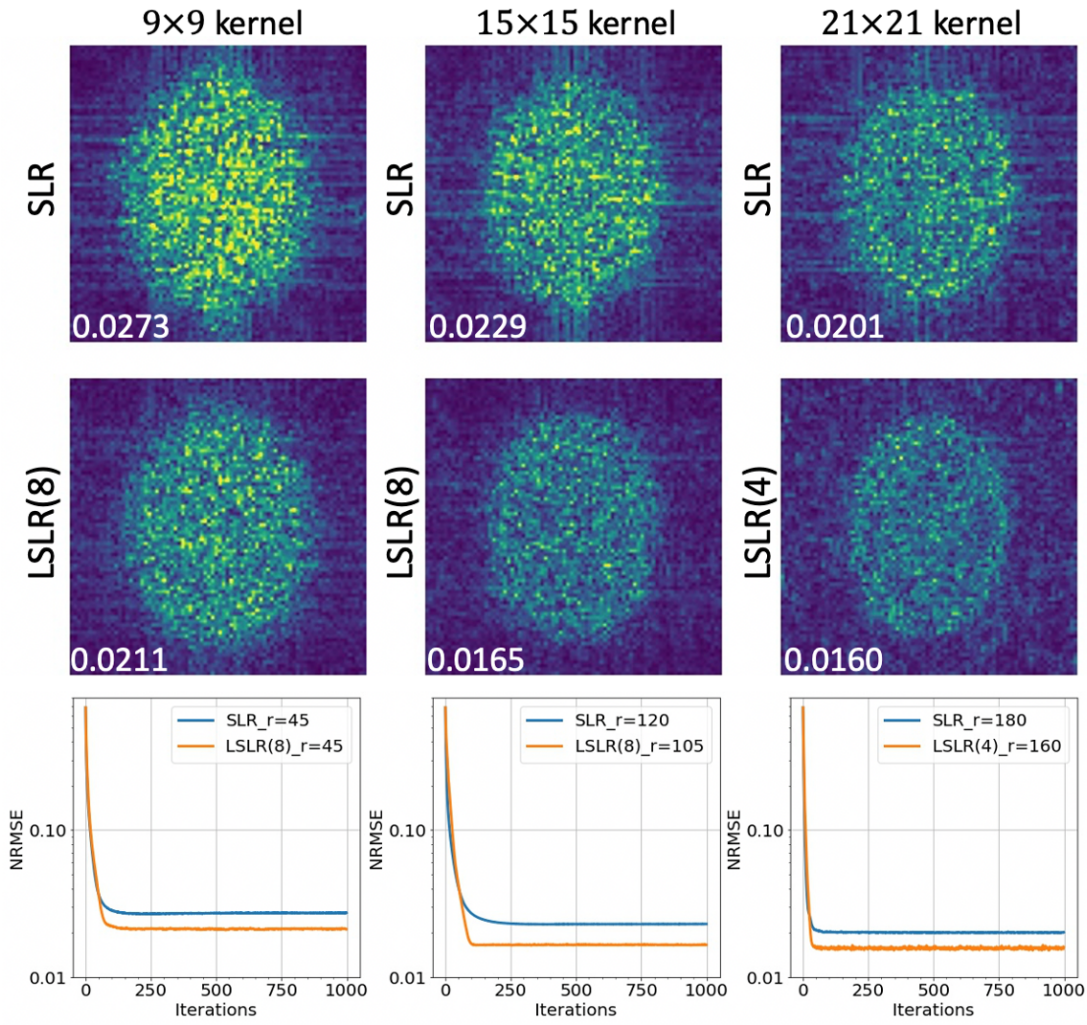

Fig. 6.12. The comparison of SLR and LSLR reconstructions with different kernel sizes of the C-matrix. The error maps of the reconstruction results compared to the ground truth and the plots of NRMSE through iterations are shown. The NRMSE values are shown in each error map respectively.

Performance Validation - White Noise-like Phantom

A hypothesis regarding the improvement of LSLR over SLR is that LSLR lightens the inherent weighting effect on different frequency components in k-space. Since k-space centre has a much higher energy, it might play a dominant role in k-space based methods, i.e., the low-rank constraint is dominated by row vectors of the block-Hankel matrix corresponding to k-space centre. This hypothesis was proved wrong with a simulation experiment using a white noise-like phantom, as shown in Fig. 6.13. The phantom was generated from an image of white noise which has almost uniform energy distribution in

k-space, and the image support is limited to ensure the low-rank structure of its C-matrix.

The reconstruction results are shown in Fig. 6.14. Despite the great reconstruction of SLR, LSLR still achieves a 2 × reduction in NRMSE in this case.

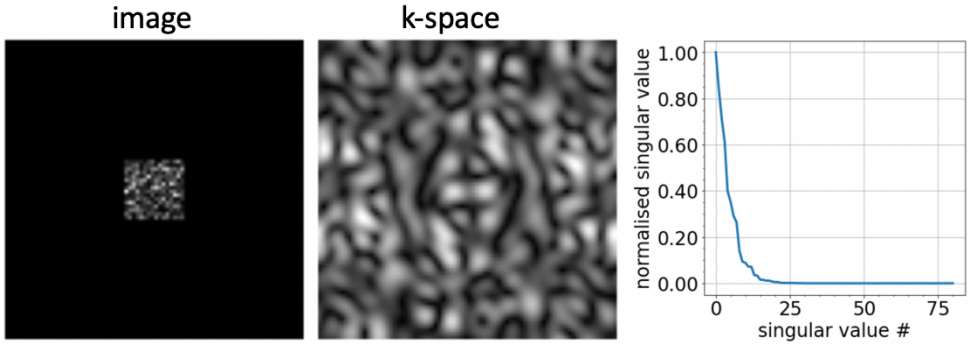

Fig. 6.13. The ground truth image, k-space and normalised singular value distribution of the C-matrix of a white noise-like phantom.

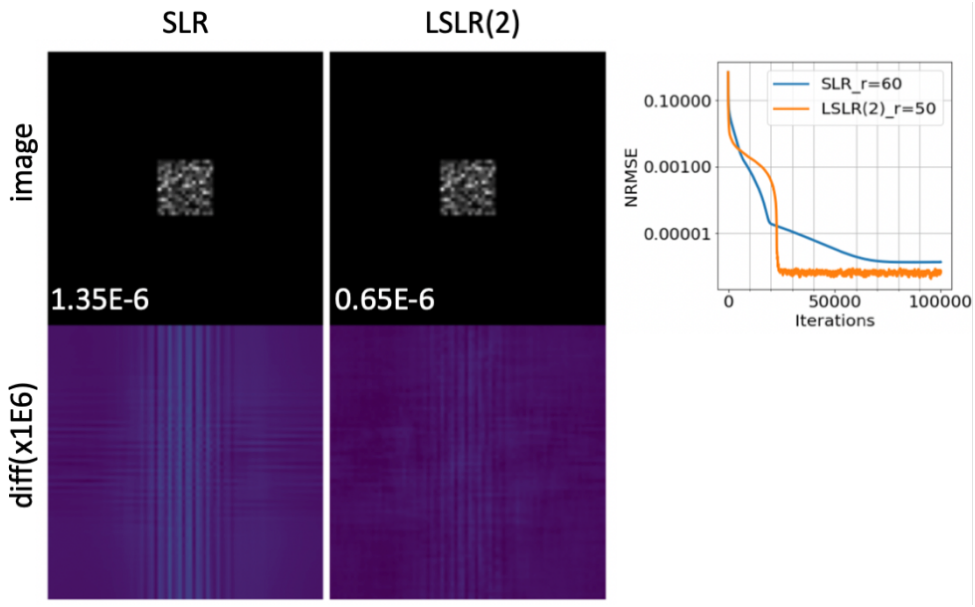

Fig. 6.14. The comparison of SLR and LSLR reconstructions on the white noise-like phantom shown in Fig. 6.13. The reconstructed images and their corresponding error maps as well as the plot of NRMSE through iterations are shown. The NRMSE values are shown in each image respectively. $r = 60$ was used for SLR and $r = 50$ was used for LSLR.

6.3.2 Experimental Data

T_1 Weighted Data

The performance of LSLR was validated on a T_1 weighted dataset, where 1D random under-sampling (45%) was performed retrospectively. The S-matrix was used for SLR and LSLR reconstructions. As shown in Fig. 6.15, LSLR reconstruction has a better performance than SLR reconstruction in this case.

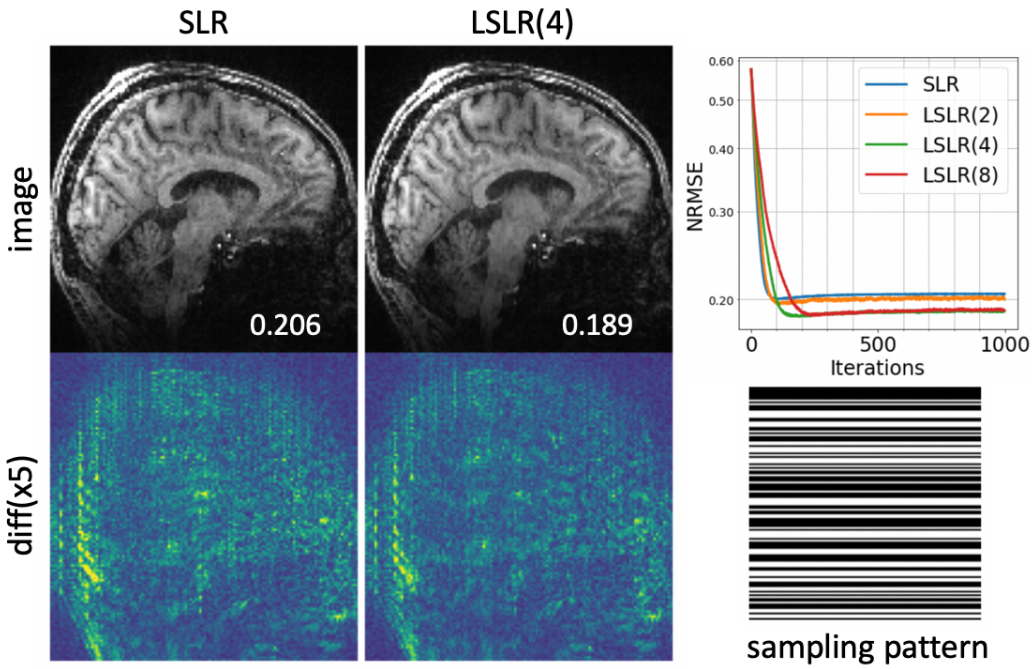

Fig. 6.15. The reconstruction results of the T_1 weighted data with 1D random under-sampling. The reconstructed images and their error maps compared to the fully sampled reconstruction are shown. The plot of NRMSE through iterations and the sampling pattern are also shown. The NRMSE values are shown in each image respectively. The S-matrix and $r = 140$ was used.

GRE Data

The performance of locally structured low-rank matrix reconstruction of a single echo and locally structured low-rank tensor reconstruction of two echoes were both validated on the GRE data. Fig. 6.16 shows the results of a single echo reconstruction, where 2D random under-sampling (50.4%) was performed retrospectively. The C-matrix was used for SLR and LSLR reconstructions. Fig. 6.17 shows the results of another single echo reconstruction where a uniform 2 × under-sampling with additional 3/4 partial Fourier (total under-sampling 38.3%) was performed. The VC-matrix was used in this experiment.

LSLR reconstruction achieves a lower NRMSE than SLR reconstruction in both two cases.

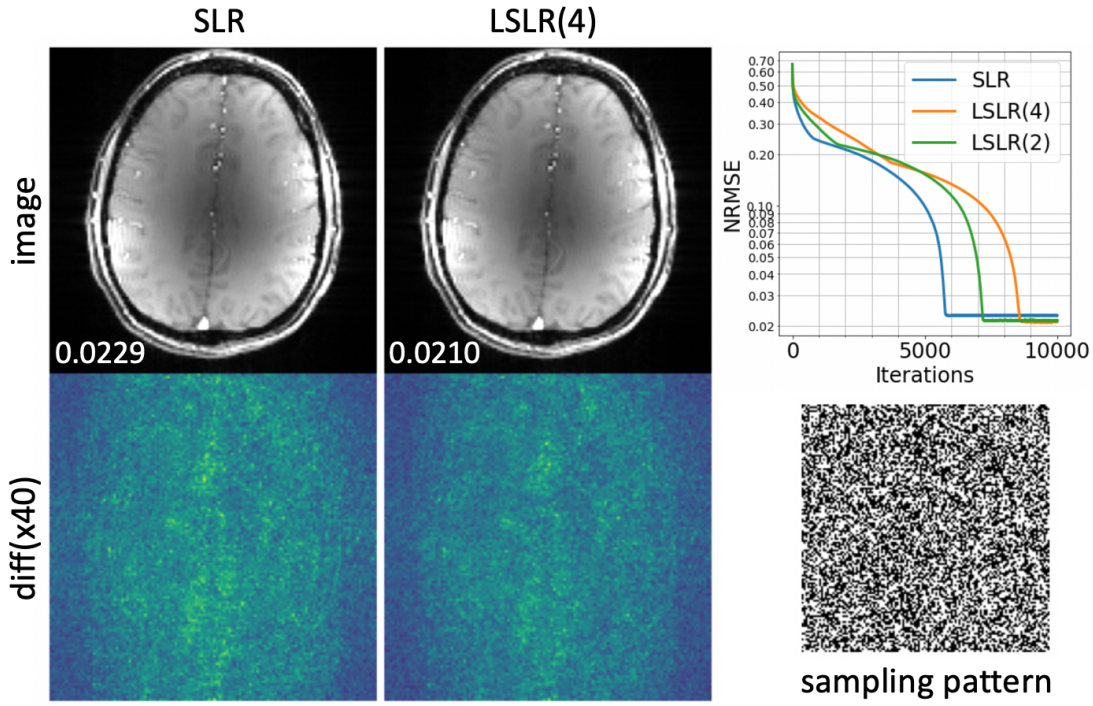

Fig. 6.16. The single echo reconstruction results of the GRE data with 2D random under-sampling. The reconstructed images and their error maps compared to the fully sampled reconstruction are shown. The plot of NRMSE through iterations and the sampling pattern are also shown. The NRMSE values are shown in each image respectively. The C-matrix and $r = 120$ were used.

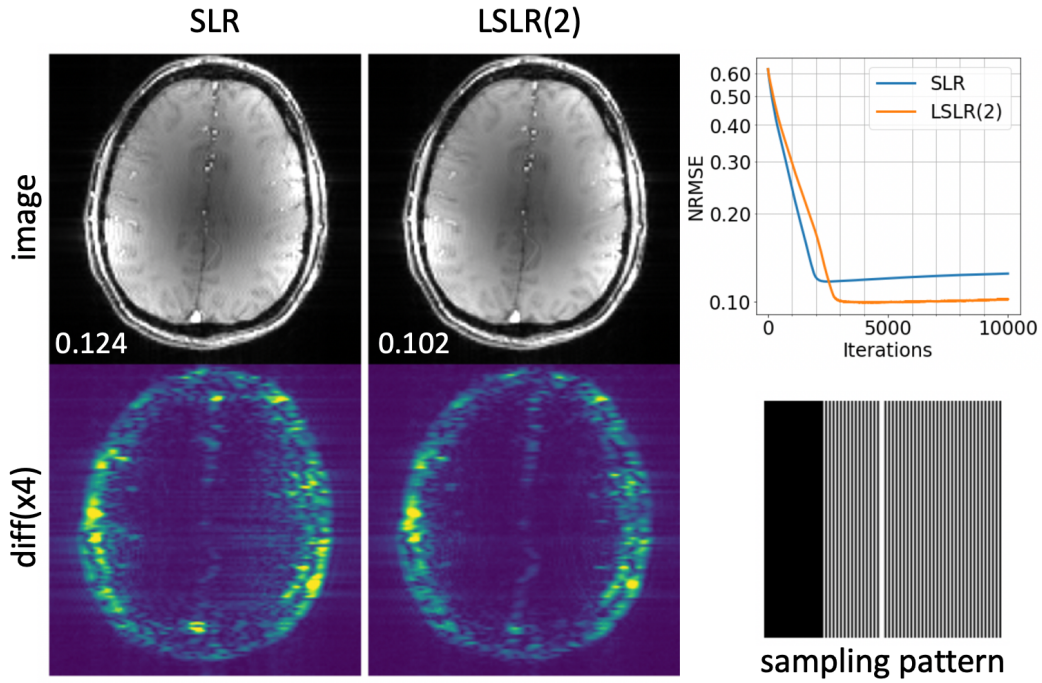

Fig. 6.17. The single echo reconstruction results of the GRE data with uniform and partial Fourier under-sampling. The reconstructed images and their error maps compared to the fully sampled reconstruction are shown. The plot of NRMSE through iterations and the sampling pattern are also shown. The NRMSE values are shown in each image respectively. The VC-matrix and $r = 180$ were used.

Fig. 6.18 shows the results of a two-echo joint tensor reconstruction where 4 × uniform under-sampling with additional 3/4 partial Fourier was performed. Fig. 6.19 shows the results of another two-echo joint structured low-rank tensor reconstruction where complementary 1D random under-sampling was used for each echo. The VC-matrix was used in these two experiments. LSLR has a consistently better reconstruction performance in both cases.

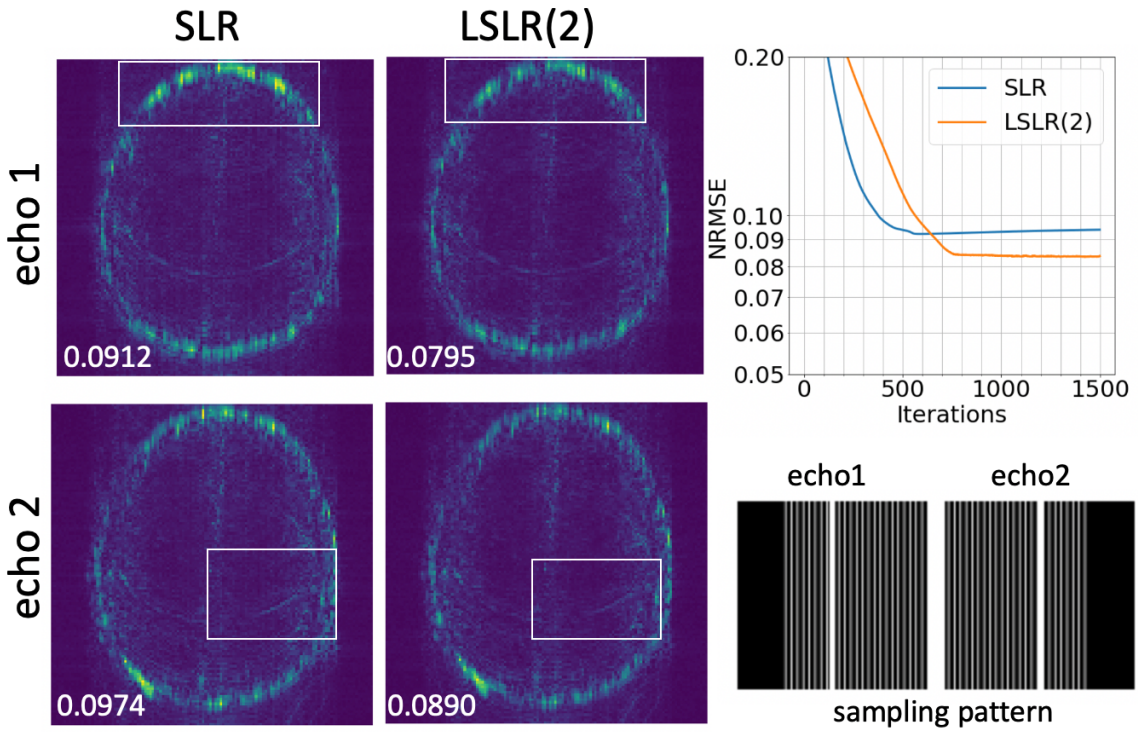

Fig. 6.18. The two-echo joint tensor reconstruction results of the GRE data with uniform and partial Fourier under-sampling. The error maps of the reconstructed images compared to the fully sampled construction are shown for the first echo and second echo respectively. The NRMSE for each echo is shown in the corresponding error map. The plot of NRMSE through iterations corresponds to the results of two echoes. The sampling patterns of the two echoes are also shown. The VC-matrix and $r_1 = 215, r_2 = 165$ were used. The white boxes highlight the differences between SLR and LSLR.

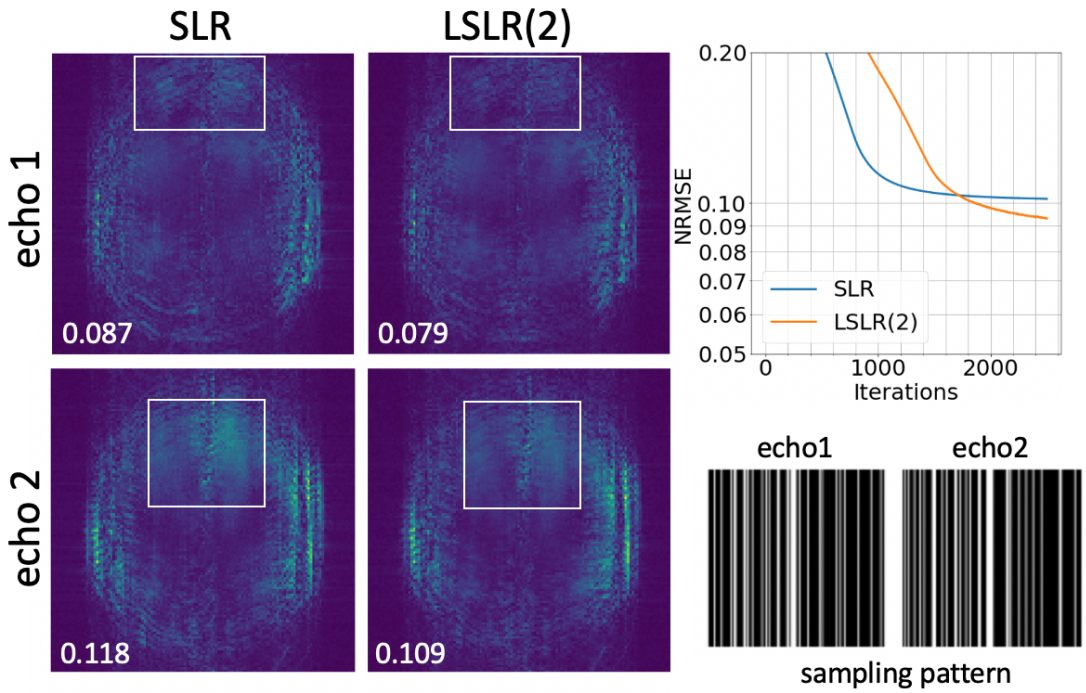

Fig. 6.19. The two-echo joint tensor reconstruction results of the GRE data with 1D random under-sampling. The error maps of the reconstructed images compared to the fully sampled construction are shown for the first echo and second echo respectively. The NRMSE for each echo is shown in the corresponding error map. The plot of NRMSE through iterations corresponds to the results of two echoes. The sampling patterns of the two echoes are also shown. The VC-matrix and $r_1 = 230, r_2 = 150$ were used. The white boxes highlight the differences between SLR and LSLR.

6.4 Discussion

In this chapter, the locally structured low-rank image reconstruction method is proposed. Unlike SLR reconstruction which enforces a low-rank constraint on the entire block-Hankel structured matrix, LSLR reconstruction imposes low-rank constraints on randomly chosen submatrices of the block-Hankel structured matrix, which essentially relaxes the global low-rank constraint. Simulation experiments based on numerical phantom and experimental data have verified that LSLR outperforms SLR in a variety of scenarios including different low-rank matrix constructions (C-matrix, S-matrix and VC-matrix), noise levels, sampling patterns, data types, etc. While LSLR is demonstrated here using a non-convex formulation which is solved by ADMM optimisation, it would be

straightforward to implement this method with other forms of low-rank constraints and optimisation schemes.

It is worth noting that although LSLR enforces low-rank constraints on submatrices of the block-Hankel structured matrix, the cycle-spinning procedure can still encourage common null space vectors/annihilation filters across the entire block-Hankel structured matrix and promote the global low-rank structure. Thus, it is still consistent with the theoretically justified shift-invariant linear dependencies in MRI data and only relaxes the global low-rank constraint. As shown in some numerical simulations, this relaxation can be particularly helpful when the global matrix is not sufficiently low-rank.

The relaxed low-rank constraints on submatrices can achieve more flexibility by potentially allowing for variations in the null-space vectors to some degree. The spatially varying signal characteristics considered in some methods (Park et al., 2005; Park and Park, 2012; Xu et al., 2015; Zhou et al., 2016) might be contributing to the better performance of LSLR. In practice, the variations of annihilation filters across k-space might be advantageous when time-dependent sampling considerations are significant, such as in long, single-shot acquisition schemes like EPI imaging. However, these potentially shift-variant characteristics must not be the only contributing factor as the benefit of LSLR was also validated in numerical simulations besides experimental data. Besides, the simulation experiment based on white noise-like phantom has also excluded spatially varying energy distribution of k-space as a potential reason that leads to the improvement of LSLR. In short, further theoretical validation is needed to fully understand the mechanism of LSLR.

Although the optimal number of submatrices of LSLR is data-dependent, typically a small number, i.e., 2, 4 or 8 is sufficient for good reconstruction performance based on the empirical assessments performed in this work, which in some sense confirms the

global low-rank structure and that LSLR only relaxes this in small scale. The optimal rank of LSLR is typically not bigger than the optimal rank of SLR, as the rank of a submatrix can only be smaller or equal to the rank of the full matrix.

One drawback of LSLR is that its computation time can be slightly longer than SLR sometimes, given the same prescribed number of iterations is used. The implication of using submatrices in terms of reconstruction time is twofold. Each submatrix is of a smaller size which reduces the scale of the time-consuming SVD, while the same procedure needs to be repeated for each submatrix. The computation time for SLR and LSLR(4) of the T_1 weighted data were 36 and 54 minutes respectively for 3000 iterations. The computation time for SLR and LSLR(4) of the single echo GRE data with 2D random sampling were 67 and 66 minutes respectively for 10000 iterations. In terms of convergence speed, SLR is often slightly faster than LSLR on the experimental data. Besides, another drawback of LSLR is that both the rank and the number of submatrices need to be tuned for its best performance, whereas SLR only has the rank parameter that needs to be optimised.

6.5 Summary

In this chapter, locally structured low-rank image reconstruction method is presented which relaxes the global low-rank constraint on the full block-Hankel structured matrix by imposing low-rank constraints on its submatrices. LSLR achieves a robust improvement compared to SLR in a wide range of settings.

7 Summary

7.1 Questions Addressed in This Thesis

This thesis focuses on the development of structured low-rank reconstruction methods for 3D multi-shot EPI imaging, in order to achieve improved robustness against physiological and motion artefacts. Although low-rank reconstruction has been applied in MRI for many years as one of the most established reconstruction methods in addition to parallel imaging and compressed sensing, structured low-rank reconstruction is a relatively new and emerging research field. Despite using the same mathematical frameworks to express the low-rank constraint, conventional low-rank reconstruction and structured low-rank reconstruction construct their low-rank matrices in very different ways. Conventional low-rank reconstruction is mostly used in spatiotemporal imaging or in applications acquiring any series of images. It uses the multi-dimensional dataset itself as the low-rank matrix, whose low-rank property is validated empirically. In contrast, the low-rank matrix in structured low-rank reconstruction is typically in the format of block-

Hankel structured matrices constructed from either single or multi-dimensional k-space data, whose low-rankness originates from the signal correlations within local neighbourhoods in single or among multi-dimensional k-space data with more comprehensive theoretical justifications. Note the SLR reconstruction framework can express conventional low-rank reconstructions if the local neighbourhood shrinks to be a single k-space sample by using a 1×1 kernel for the block-Hankel transformation. This makes structured low-rank reconstruction a more generalised and comprehensive framework which can exploit various linear dependencies in MRI data for a wide range of applications. One of the most important applications is the reconstruction of multi-shot data corrupted by shot-to-shot phase variations due to various sources, which is an appealing feature of SLR as it cannot be easily handled by conventional methods like compressed sensing or parallel imaging. The method MUSSELS (Mani et al., 2017), for example, has been used successfully in the reconstruction of 2D multi-shot DWI data to deal with shot-to-shot phase inconsistencies.

Inspired by the great success of MUSSELS, I have explored the application of SLR reconstruction in 3D multi-shot EPI, which has been drawing an increasing interest in fMRI as an alternative to 2D single shot EPI, particularly for high-resolution imaging at ultra-high fields. However, multi-shot EPI acquisition suffers from inter-shot phase variations due to physiologically induced B_0 fluctuations in fMRI. These inter-shot inconsistencies can impair the temporal stability of the time course by contributing to spurious temporal variance, and thus the SNR benefit of 3D acquisition is effectively diminished in 3D multi-shot EPI, which then does not necessarily demonstrate a tSNR advantage compared to 2D single-shot EPI. The SLR reconstruction framework deals with this issue by resolving the phase differences across shots, which reconstructs an individual image for each phase state and then performs a phase-insensitive (e.g., SOS)

shot combination. The images for all phase states are jointly reconstructed to exploit the linear dependencies between them, which are formulated as the low-rank property of the block-Hankel structured matrix representation of their k-space data. Since the number of shots is large in 3D multi-shot EPI and the phase variations are assumed to be temporally coherent, a phase state is assigned for each shot group which consists of a small number of consecutive shots rather than a single shot to improve the conditioning of the reconstruction. To facilitate the reconstruction, a CAIPI sampling trajectory with interleaved ordering along the shot dimension was used in order to achieve an approximately uniform under-sampling pattern for each shot group. Simulation and in-vivo experiments have demonstrated that the proposed 3D method is able to achieve a tSNR improvement over conventional 3D method across a wide range of isotropic resolutions from 1mm to 1.8mm at 7T, making the tSNR of 3D multi-shot EPI higher than 2D SMS-EPI in some scenarios where it used to be the opposite.

The linear dependencies between different shot groups are built on the assumption that they only differ in image phase while having the same magnitude. Thus, any motion induced signal changes can violate this assumption and reduce the validity of the structured low-rank constraint. In addition, multi-shot acquisitions are also sensitive to inter-shot motion which can lead to ghost artefacts and blurring. To deal with these issues, a motion compensated structured low-rank reconstruction (mcSLR) has been proposed to further improve the robustness of 3D multi-shot EPI to both inter-shot motion and phase variations. Motion compensation is achieved by incorporating a rigid motion transform in the forward model, which is considered at the shot-to-shot level. The motion parameters are jointly estimated with the multi-shot group images in a fully data driven way, which are decomposed into inter- and intra-shot group motion and estimated by different approaches. In mcSLR framework, the low-rank constraint is enforced on the

block-Hankel structured k-space representation of the multi-shot group images which are free from motion artefacts and in good alignment. Simulation and in vivo experiments have validated the performance of mcSLR reconstruction, which can largely remove the image artefacts as well as improve the temporal stability of 3D multi-shot EPI in the presence of rigid motion.

On top of exploring the application of SLR reconstruction in 3D multi-shot EPI and the development of a motion compensated SLR reconstruction, I was also interested in exploring the optimisation of SLR constraints from an algorithmic perspective. Inspired by the characteristics of typical block-Hankel structured matrices, namely 1) they have a much larger number of rows corresponding to the number of k-space patches than the number of columns, and 2) their column space is demonstrated to have a reduced dimension indicating the low-rankness, a locally structured low-rank reconstruction scheme (LSLR) has been proposed. The LSLR reconstruction relaxes the global low-rank constraint over the entire block-Hankel structured matrix to a set of low-rank constraints over each of its non-overlapping submatrices consisting of a subset of row vectors. The positions of these submatrices are shifted randomly at every iteration of the optimisation algorithm to avoid boundary artefacts and promote shift-invariant signal correlations and global low-rankness. Realistic simulation experiments suggest that LSLR outperforms global SLR reconstruction very robustly in a wide range of imaging settings.

7.2 Unanswered Questions and Future Directions

Throughout the journey of solving the problems mentioned above, I have also come across some interesting questions which have not yet been answered and might be worth exploring in the future as extensions of the current work.

Chapter 3 has shown how structured low-rank reconstruction can be used to deal with

the phase variations in 3D multi-shot EPI data. Unlike conventional phase estimation-based methods, SLR reconstruction circumvents the estimation of phase variations by phase-resolved imaging which reconstructs an image for each phase state corresponding to a shot or a shot group. Although SLR provides a universal solution to this problem, it may not be the most efficient in the sense that each phase map needs to be represented by a full-size image, especially when the number of phase states to be resolved is large. Although the intrinsic DOF of these images is regularised by the low-rank constraint, the number of variables to be estimated still increases significantly as the number of phase states increases. Thus, inspired by the motion correction method aligned reconstruction (Cordero-Grande et al., 2016), a question that I am interested in is whether a data-driven joint reconstruction of a single shot-combined image and a series of phase estimations work? Previous work which uses Free Induction Decay navigators (Wallace et al., 2020) to estimate the B_0 variations assumes that they are spatially smooth (the same assumption as in SLR reconstruction) which can be accurately represented by a second order spherical harmonic model. To improve the conditioning of the joint reconstruction problem, this lower order expansion of phase variations might be considered so that the number of unknowns needed to identify each phase variation map can be greatly reduced from the size of the image to only a few coefficients. Note both this approach and SLR reconstruction are intrinsically similar subspace methods. The SLR reconstruction exploits the existence of a null space of the block-Hankel structured matrix representation of multi-shot k-space, i.e., it lies in a lower dimensional subspace, whereas the above-mentioned joint reconstruction with lower order phase models specifies the spherical harmonic basis functions which spans a lower dimensional subspace explicitly. Moreover, since the phase variations are assumed to be temporally coherent, an extended spatial-temporal model of them might be useful to further restrain the subspace of the

multi-shot data, while the temporal correlations of the phase variations have not been exploited in current SLR reconstruction.

The current reconstruction for 3D multi-shot EPI does not account for distortion artefacts. A possible solution to diminish distortion artefacts is to reduce the echo train length by using segmented readout along k_y as well as k_z, similar to (Stirnberg and Stöcker, 2021). The SLR reconstruction for 3D multi-shot EPI still applies in this case, while the challenge is the increased number of shots. Presumably it is better to distribute different segments along k_y into different shot groups so that the k-space of each shot group is still sampled sequentially along k_y, and then any motion/physiological fluctuations are mapped coherently along k_y for each shot group. Accordingly, different segments along k_y are not acquired sequentially as an inner loop but alternates with the acquisition of different segments along k_z as an outer loop. Other disadvantages including prolonged acquisition time can possibly be counterbalanced by the higher SNR due to shorter TE or lower bandwidth. Thus, it is interesting to see whether the use of segmentation along k_y can lead to significant benefits in general.

As for the motion compensated structured low-rank reconstruction, a possibly useful question to look into is whether the incorporation of a motion-free reference image could improve the reconstruction efficiency, similar to the method SAMER (Polak et al., 2022) which uses a low-resolution scout image to speed up aligned reconstruction (Cordero-Grande et al., 2016). The joint estimation of the image and motion parameters is solved by alternating minimisation currently which could have a low computational efficiency. A separate motion-free image, if available, might be useful to serve as the reference to derive the motion estimates, which is decoupled from the estimation of the image. Thus, the estimation of the image and motion parameters might only need to be performed once

or just for a few iterations. Although the availability of a reference image is an extra precondition of the reconstruction, it is not a demanding requirement to meet in practice. For instance, the low-resolution calibration data required for the calculation of coil sensitivities, or a motion-free volume, or even the reference-free mcSLR reconstruction of one volume can possibly be used as the reference image. However, possible inconsistencies between the reference image and the motion corrupted imaging data including the phase inconsistencies might be an issue.

In the locally structured low-rank reconstruction framework, the low-rank constraint over the block-Hankel structured matrix is relaxed by using a set of low-rank constraints over submatrices of the block-Hankel structured matrix. Although the efficacy of this relaxation has been validated experimentally and we have formulated several possible hypotheses to explain the observed benefit, the theoretical justification which can help us understand the mechanism behind this approach better is still an open question. Besides, since it is a general modification on the global low-rank constraint, it would be interesting to see whether this approach can be applied in some other applications utilising a low-rank model, e.g., image denoising.

Similar to low-rank and structured low-rank reconstruction, manifold learning can be viewed as a generalisation of subspace methods which is based on the assumption that the data in higher dimensional space can be projected onto an embedded lower dimensional manifold. The biggest difference between SLR reconstruction and manifold learning is that SLR uses a linear mapping between the higher dimensional signal space and lower dimensional signal space whereas manifold learning uses non-linear mappings. Manifold learning methods have showed successful applications in the reconstruction of high dimensional MRI data, i.e., in cardiac imaging where respiration and cardiac pulsation correspond to additional data dimensions (Ahmed et al., 2020; Nakarmi et al.,

2017, 2017; Poddar and Jacob, 2016; Zou et al., 2021). Inspired by this, it would be interesting to see whether the leap from linear mapping to non-linear mapping could have significant benefits in our scenarios by exploiting the non-linear signal correlations and pushing the compactness of the lower dimensional signal space further.

7.3 Closing Remarks

Although structured low-rank reconstruction has made further progress standing on the shoulders of low-rank reconstruction, it is not yet widely known and applied. I started my DPhil journey along this road by trying to replicate its success in the application of 3D multi-shot EPI for fMRI. After that I was keen to fix its weak spot by proposing a motion compensated version of structured low-rank reconstruction and challenged the routine implementation by presenting the locally structured low-rank scheme. I am interested to see to the directions in which structured low-rank reconstruction will continue to develop, and whether it will gain prosperity or prove to be a last generation of iPhone following its forerunner low-rank reconstruction in the future.

8 References

Agrawal, U., Brown, E.N., Lewis, L.D., 2020. Model-based physiological noise removal in fast fMRI. NeuroImage 205, 116231. https://doi.org/10.1016/j.neuroimage.2019.116231

Ahmed, A.H., Zhou, R., Yang, Y., Nagpal, P., Salerno, M., Jacob, M., 2020. Free-Breathing and Ungated Dynamic MRI Using Navigator-Less Spiral SToRM. IEEE Transactions on Medical Imaging 39, 3933–3943. https://doi.org/10.1109/TMI.2020.3008329

Atkinson, D., Counsell, S., Hajnal, J.V., Batchelor, P.G., Hill, D.L.G., Larkman, D.J., 2006. Nonlinear phase correction of navigated multi-coil diffusion images. Magn. Reson. Med. 56, 1135–1139. https://doi.org/10.1002/mrm.21046

Balachandrasekaran, A., Mani, M., Jacob, M., 2019. Calibration-Free B0 Correction of EPI Data Using Structured Low Rank Matrix Recovery. IEEE Trans. Med. Imaging 38, 979–990. https://doi.org/10.1109/TMI.2018.2876423

Barkhof, F., Haller, S., Rombouts, S.A.R.B., 2014. Resting-State Functional MR Imaging: A New Window to the Brain. Radiology 272, 29–49. https://doi.org/10.1148/radiol.14132388

Barry, R.L., Martyn Klassen, L., Williams, J.M., Menon, R.S., 2008. Hybrid two-dimensional navigator correction: A new technique to suppress respiratory-induced physiological noise in multi-shot echo-planar functional MRI. NeuroImage 39, 1142–1150. https://doi.org/10.1016/j.neuroimage.2007.09.060

Barry, R.L., Menon, R.S., 2005. Modeling and suppression of respiration-related physiological noise in echo-planar functional magnetic resonance imaging using global and one-dimensional navigator echo correction. Magn. Reson. Med. 54, 411–418. https://doi.org/10.1002/mrm.20591

Behzadi, Y., Restom, K., Liau, J., Liu, T.T., 2007. A Component Based Noise Correction Method (CompCor) for BOLD and Perfusion Based fMRI. Neuroimage 37, 90–101. https://doi.org/10.1016/j.neuroimage.2007.04.042

Belliveau, J.W., Kennedy, D.N., McKinstry, R.C., Buchbinder, B.R., Weisskoff, R.M., Cohen, M.S., Vevea, J.M., Brady, T.J., Rosen, B.R., 1991. Functional mapping

of the human visual cortex by magnetic resonance imaging. Science 254, 716–719.
Bilgic, B., Liao, C., Manhard, M.K., Tian, Q., Chatnuntawech, I., Iyer, S.S., Cauley, S.F., Feiweier, T., Giri, S., Hu, Y., 2019. Robust high-quality multi-shot EPI with low-rank prior and machine learning, in: Proc Int Soc Magn Reson Med 2019.
Biswal, B., Zerrin Yetkin, F., Haughton, V.M., Hyde, J.S., 1995. Functional connectivity in the motor cortex of resting human brain using echo-planar mri. Mag. Reson. Med. 34, 537–541. https://doi.org/10.1002/mrm.1910340409
Blaimer, M., Gutberlet, M., Kellman, P., Breuer, F.A., Köstler, H., Griswold, M.A., 2009. Virtual coil concept for improved parallel MRI employing conjugate symmetric signals. Magn Reson Med 61, 93–102. https://doi.org/10.1002/mrm.21652
Blumensath, T., Davies, M.E., 2009. Iterative hard thresholding for compressed sensing. Applied and Computational Harmonic Analysis 27, 265–274. https://doi.org/10.1016/j.acha.2009.04.002
Boyd, S., 2010. Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. FNT in Machine Learning 3, 1–122. https://doi.org/10.1561/2200000016
Brackenier, Y., Cordero-Grande, L., Tomi-Tricot, R., Wilkinson, T., Bridgen, P., Price, A., Malik, S.J., De Vita, E., Hajnal, J.V., 2022. Data-driven motion-corrected brain MRI incorporating pose-dependent B0 fields. Magn. Reson. Med. 88, 817–831. https://doi.org/10.1002/mrm.29255
Brady, T.J., 1991. Future prospects for MR imaging. in: Proc IntSoc Magn Reson Med 1991.
Breuer, F.A., Blaimer, M., Heidemann, R.M., Mueller, M.F., Griswold, M.A., Jakob, P.M., 2005. Controlled aliasing in parallel imaging results in higher acceleration (CAIPIRINHA) for multi-slice imaging. Magn. Reson. Med. 53, 684–691. https://doi.org/10.1002/mrm.20401
Breuer, F.A., Blaimer, M., Mueller, M.F., Seiberlich, N., Heidemann, R.M., Griswold, M.A., Jakob, P.M., 2006. Controlled aliasing in volumetric parallel imaging (2D CAIPIRINHA). Magn. Reson. Med. 55, 549–556. https://doi.org/10.1002/mrm.20787
Buehrer, M., Pruessmann, K.P., Boesiger, P., Kozerke, S., 2007. Array compression for MRI with large coil arrays. Magn. Reson. Med. 57, 1131–1139. https://doi.org/10.1002/mrm.21237
Buckner, R.L., 1998. Event-related fMRI and the hemodynamic response. Human Brain Mapping 6, 373–377. https://doi.org/10.1002/(SICI)1097-0193(1998)6:5/6<373::AID-HBM8>3.0.CO;2-P
Buckner, R.L., 2012. The serendipitous discovery of the brain's default network. NeuroImage, 20 YEARS OF fMRI 62, 1137–1145. https://doi.org/10.1016/j.neuroimage.2011.10.035
Bullmore, E., 2012. The future of functional MRI in clinical medicine. NeuroImage, 20 YEARS OF fMRI 62, 1267–1271. https://doi.org/10.1016/j.neuroimage.2012.01.026
Candes, E.J., Plan, Y., 2010. Matrix Completion With Noise. Proceedings of the IEEE 98, 925–936. https://doi.org/10.1109/JPROC.2009.2035722
Chen, N.-K., Guidon, A., Chang, H.-C., Song, A.W., 2013. A robust multi-shot scan strategy for high-resolution diffusion weighted MRI enabled by multiplexed sensitivity-encoding (MUSE). Neuroimage 72, 41–47. https://doi.org/10.1016/j.neuroimage.2013.01.038

Chen, X., Wu, W., Chiew, M., 2023. Improving robustness of 3D multi-shot EPI by structured low-rank reconstruction of segmented CAIPI sampling for fMRI at 7T. NeuroImage 267, 119827. https://doi.org/10.1016/j.neuroimage.2022.119827

Chen, X., Wu, W., Chiew, M., 2022. Locally Structured Low-Rank MR Image Reconstruction using Submatrix Constraints, in: 2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI). Presented at the 2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI), pp. 1–4. https://doi.org/10.1109/ISBI52829.2022.9761692

Chen, X., Wu, W., Chiew, M., 2020. Reduced Inter-shot Physiological Variability in 3D Multi-Shot fMRI using Structured Low-Rank Matrix Completion. in: Proc Int Soc Magn Reson Med 2020.

Chen, X., Wu, W., Chiew, M., 2021. Respiratory fluctuations in 3D fMRI from inter-shot phase variations can be reduced by low-rank reconstruction of segmented CAIPI sampling. in: Proc IntSoc Magn Reson Med 2021.

Chen, X., Wu, W., Chiew, M., 2022. A Locally Structured Low-Rank Tensor Method Using Submatrix Constraints for Joint Multi-echo Image Reconstruction. in: Proc IntSoc Magn Reson Med 2022.

Chiew, M., Graedel, N.N., McNab, J.A., Smith, S.M., Miller, K.L., 2016. Accelerating functional MRI using fixed-rank approximations and radial-cartesian sampling. Magn. Reson. Med. 76, 1825–1836. https://doi.org/10.1002/mrm.26079

Chiew, M., Graedel, N.N., Miller, K.L., 2018. Recovering task fMRI signals from highly under-sampled data with low-rank and temporal subspace constraints. NeuroImage 174, 97–110.

Chiew, M., Smith, S.M., Koopmans, P.J., Graedel, N.N., Blumensath, T., Miller, K.L., 2015. k-t FASTER: Acceleration of functional MRI data acquisition using low rank constraints. Magn. Reson. Med. 74, 353–364. https://doi.org/10.1002/mrm.25395

Christodoulou, A.G., Shaw, J.L., Nguyen, C., Yang, Q., Xie, Y., Wang, N., Li, D., 2018. Magnetic resonance multitasking for motion-resolved quantitative cardiovascular imaging. Nat Biomed Eng 2, 215–226. https://doi.org/10.1038/s41551-018-0217-y

Cooley, J.W., Tukey, J.W., 1965. An algorithm for the machine calculation of complex Fourier series. Math. Comp. 19, 297–301. https://doi.org/10.1090/S0025-5718-1965-0178586-1

Cordero-Grande, L., Ferrazzi, G., Teixeira, R.P.A.G., O'Muircheartaigh, J., Price, A.N., Hajnal, J.V., 2020. Motion-corrected MRI with DISORDER: Distributed and incoherent sample orders for reconstruction deblurring using encoding redundancy. Magn. Reson. Med. 84, 713–726. https://doi.org/10.1002/mrm.28157

Cordero-Grande, L., Teixeira, R.P.A.G., Hughes, E.J., Hutter, J., Price, A.N., Hajnal, J.V., 2016. Sensitivity Encoding for Aligned Multishot Magnetic Resonance Reconstruction. IEEE Transactions on Computational Imaging 2, 266–280. https://doi.org/10.1109/TCI.2016.2557069

Dale, A.M., Buckner, R.L., 1997. Selective averaging of rapidly presented individual trials using fMRI. Human Brain Mapping 5, 329–340. https://doi.org/10.1002/(SICI)1097-0193(1997)5:5<329::AID-HBM1>3.0.CO;2-5

Fox, M.D., Snyder, A.Z., Vincent, J.L., Corbetta, M., Van Essen, D.C., Raichle, M.E., 2005. The human brain is intrinsically organized into dynamic, anticorrelated

functional networks. Proc Natl Acad Sci U S A 102, 9673–9678. https://doi.org/10.1073/pnas.0504136102
Friston, K.J., Holmes, A.P., Worsley, K.J., Poline, J.-P., Frith, C.D., Frackowiak, R.S.J., 1994. Statistical parametric maps in functional imaging: A general linear approach. Human Brain Mapping 2, 189–210. https://doi.org/10.1002/hbm.460020402
Friston, K.J., Zarahn, E., Josephs, O., Henson, R.N.A., Dale, A.M., 1999. Stochastic Designs in Event-Related fMRI. NeuroImage 10, 607–619. https://doi.org/10.1006/nimg.1999.0498
Gao, H., Li, L., Zhang, K., Zhou, W., Hu, X., 2014. PCLR: Phase-constrained low-rank model for compressive diffusion-weighted MRI. Magn. Reson. Med. 72, 1330–1341. https://doi.org/10.1002/mrm.25052
Glover, G.H., Li, T.-Q., Ress, D., 2000. Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR. Magn. Reson. Med. 44, 162–167. https://doi.org/10.1002/1522-2594(200007)44:1<162::AID-MRM23>3.0.CO;2-E
Griswold, M.A., Jakob, P.M., Heidemann, R.M., Nittka, M., Jellus, V., Wang, J., Kiefer, B., Haase, A., 2002. Generalized autocalibrating partially parallel acquisitions (GRAPPA). Magn. Reson. Med. 47, 1202–1210. https://doi.org/10.1002/mrm.10171
Haldar, J.P., 2014. Low-rank modeling of local k-space neighborhoods (LORAKS) for constrained MRI. IEEE Trans Med Imaging 33, 668–681. https://doi.org/10.1109/TMI.2013.2293974
Haldar, J.P., Hernando, D., 2009. Rank-Constrained Solutions to Linear Matrix Equations Using PowerFactorization. IEEE Signal Process Lett 16, 584–587. https://doi.org/10.1109/LSP.2009.2018223
Haldar, J.P., Liang, Z.-P., 2011. Low-rank approximations for dynamic imaging, in: 2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro. Presented at the 2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro, pp. 1052–1055. https://doi.org/10.1109/ISBI.2011.5872582
Haldar, J.P., Liang, Z.-P., 2010. Spatiotemporal imaging with partially separable functions: A matrix recovery approach, in: 2010 IEEE International Symposium on Biomedical Imaging: From Nano to Macro. Presented at the 2010 IEEE International Symposium on Biomedical Imaging: From Nano to Macro, pp. 716–719. https://doi.org/10.1109/ISBI.2010.5490076
Haldar, J.P., Setsompop, K., 2020. Linear Predictability in Magnetic Resonance Imaging Reconstruction: Leveraging Shift-Invariant Fourier Structure for Faster and Better Imaging. IEEE Signal Process. Mag. 37, 69–82. https://doi.org/10.1109/MSP.2019.2949570
Haldar, J.P., Zhuo, J., 2016. P-LORAKS: Low-rank modeling of local k-space neighborhoods with parallel imaging data: P-LORAKS: Low Rank-Modeling of Parallel Imaging Data. Magn. Reson. Med. 75, 1499–1514. https://doi.org/10.1002/mrm.25717
Haskell, M.W., Cauley, S.F., Wald, L.L., 2018. TArgeted Motion Estimation and Reduction (TAMER): Data Consistency Based Motion Mitigation for MRI Using a Reduced Model Joint Optimization. IEEE Transactions on Medical Imaging 37, 1253–1265. https://doi.org/10.1109/TMI.2018.2791482
He, J., Liu, Q., Christodoulou, A.G., Ma, C., Lam, F., Liang, Z.-P., 2016. Accelerated High-Dimensional MR Imaging With Sparse Sampling Using Low-Rank Tensors.

IEEE Transactions on Medical Imaging 35, 2119–2129. https://doi.org/10.1109/TMI.2016.2550204

Hendriks, A.D., D'Agata, F., Raimondo, L., Schakel, T., Geerts, L., Luijten, P.R., Klomp, D.W.J., Petridou, N., 2020. Pushing functional MRI spatial and temporal resolution further: High-density receive arrays combined with shot-selective 2D CAIPIRINHA for 3D echo-planar imaging at 7 T. NMR in Biomedicine 33. https://doi.org/10.1002/nbm.4281

Hess, A.T., Dragonu, I., Chiew, M., 2021. Accelerated calibrationless parallel transmit mapping using joint transmit and receive low-rank tensor completion. Magn. Reson. Med. 86, 2454–2467. https://doi.org/10.1002/mrm.28880

Hestenes, M.R., 1969. Multiplier and gradient methods. J Optim Theory Appl 4, 303–320. https://doi.org/10.1007/BF00927673

Hestenes, M.R., Stiefel, E., n.d. Methods of Conjugate Gradients for Solving Linear Systems.

Hu, Y., Levine, E.G., Tian, Q., Moran, C.J., Wang, X., Taviani, V., Vasanawala, S.S., McNab, J.A., Daniel, B.A., Hargreaves, B.L., 2019. Motion-robust reconstruction of multishot diffusion-weighted images without phase estimation through locally low-rank regularization. Magn. Reson. Med. 81, 1181–1190. https://doi.org/10.1002/mrm.27488

Huber, L., Ivanov, D., Handwerker, D.A., Marrett, S., Guidi, M., Uludağ, K., Bandettini, P.A., Poser, B.A., 2018. Techniques for blood volume fMRI with VASO: From low-resolution mapping towards sub-millimeter layer-dependent applications. Neuroimage 164, 131–143. https://doi.org/10.1016/j.neuroimage.2016.11.039

Hutton, C., Josephs, O., Stadler, J., Featherstone, E., Reid, A., Speck, O., Bernarding, J., Weiskopf, N., 2011. The impact of physiological noise correction on fMRI at 7 T. NeuroImage 57, 101–112. https://doi.org/10.1016/j.neuroimage.2011.04.018

Jacob, M., Mani, M.P., Ye, J.C., 2020. Structured Low-Rank Algorithms: Theory, Magnetic Resonance Applications, and Links to Machine Learning. IEEE Signal Process. Mag. 37, 54–68. https://doi.org/10.1109/MSP.2019.2950432

Jenkinson, M., Bannister, P., Brady, M., Smith, S., 2002. Improved optimization for the robust and accurate linear registration and motion correction of brain images. Neuroimage 17, 825–841. https://doi.org/10.1016/s1053-8119(02)91132-8

Jenkinson, M., Beckmann, C.F., Behrens, T.E.J., Woolrich, M.W., Smith, S.M., 2012. FSL. NeuroImage, 20 YEARS OF fMRI 62, 782–790. https://doi.org/10.1016/j.neuroimage.2011.09.015

Jezzard, P., Clare, S., 1999. Sources of distortion in functional MRI data. Human Brain Mapping 8, 80–85. https://doi.org/10.1002/(SICI)1097-0193(1999)8:2/3<80::AID-HBM2>3.0.CO;2-C

Jorge, J., Figueiredo, P., van der Zwaag, W., Marques, J.P., 2013. Signal fluctuations in fMRI data acquired with 2D-EPI and 3D-EPI at 7 Tesla. Magnetic Resonance Imaging 31, 212–220. https://doi.org/10.1016/j.mri.2012.07.001

Kasper, L., Bollmann, S., Diaconescu, A.O., Hutton, C., Heinzle, J., Iglesias, S., Hauser, T.U., Sebold, M., Manjaly, Z.-M., Pruessmann, K.P., Stephan, K.E., 2017. The PhysIO Toolbox for Modeling Physiological Noise in fMRI Data. Journal of Neuroscience Methods 276, 56–72. https://doi.org/10.1016/j.jneumeth.2016.10.019

Kim, T.H., Haldar, J.P., 2018. LORAKS software version 2.0: Faster implementation and enhanced capabilities. University of Southern California, Los Angeles, CA, Tech. Rep. USC-SIPI-443.

Krüger, G., Glover, G.H., 2001. Physiological noise in oxygenation-sensitive magnetic resonance imaging. Magn. Reson. Med. 46, 631–637. https://doi.org/10.1002/mrm.1240

Lam, F., Liang, Z.-P., 2014. A subspace approach to high-resolution spectroscopic imaging. Magn. Reson. Med. 71, 1349–1357. https://doi.org/10.1002/mrm.25168

Lam, F., Zhao, B., Liu, Y., Liang, Z.-P., Weinei, M., Schuff, N., 2013. Accelerated fMRI using low-rank model and sparsity constraints. Proc. Intl. Soc. Mag. Reson. Med. 21.

Lawrence, S.J.D., Formisano, E., Muckli, L., de Lange, F.P., 2019. Laminar fMRI: Applications for cognitive neuroscience. NeuroImage 197, 785–791. https://doi.org/10.1016/j.neuroimage.2017.07.004

Lee, D., Jin, K.H., Kim, E.Y., Park, S.-H., Ye, J.C., 2016. Acceleration of MR parameter mapping using annihilating filter-based low rank hankel matrix (ALOHA): MR Parameter Mapping Using ALOHA. Magn. Reson. Med. 76, 1848–1864. https://doi.org/10.1002/mrm.26081

Lee, J., Jin, K.H., Ye, J.C., 2016. Reference-free single-pass EPI Nyquist ghost correction using annihilating filter-based low rank Hankel matrix (ALOHA): Reference-Free Single-Pass EPI Nyquist Ghost Correction using ALOHA. Magn. Reson. Med. 76, 1775–1789. https://doi.org/10.1002/mrm.26077

Levenberg, K., 1944. A method for the solution of certain non-linear problems in least squares. Quart. Appl. Math. 2, 164–168. https://doi.org/10.1090/qam/10666

Liang, Z., 2007. SPATIOTEMPORAL IMAGINGWITH PARTIALLY SEPARABLE FUNCTIONS, in: 2007 4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro. Presented at the 2007 4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro, pp. 988–991. https://doi.org/10.1109/ISBI.2007.357020

Lingala, S.G., Hu, Y., DiBella, E., Jacob, M., 2011. Accelerated Dynamic MRI Exploiting Sparsity and Low-Rank Structure: k-t SLR. IEEE Transactions on Medical Imaging 30, 1042–1054. https://doi.org/10.1109/TMI.2010.2100850

Liu, Y., Yi, Z., Zhao, Y., Chen, F., Feng, Y., Guo, H., Leong, A.T.L., Wu, E.X., 2021. Calibrationless parallel imaging reconstruction for multislice MR data using low-rank tensor completion. Magn. Reson. Med. 85, 897–911. https://doi.org/10.1002/mrm.28480

Lobos, R.A., Hoge, W.S., Javed, A., Liao, C., Setsompop, K., Nayak, K.S., Haldar, J.P., 2021. Robust autocalibrated structured low-rank EPI ghost correction. Magn. Reson. Med. 85, 3403–3419. https://doi.org/10.1002/mrm.28638

Lobos, R.A., Kim, T.H., Hoge, W.S., Haldar, J.P., 2018. Navigator-Free EPI Ghost Correction With Structured Low-Rank Matrix Models: New Theory and Methods. IEEE Trans. Med. Imaging 37, 2390–2402. https://doi.org/10.1109/TMI.2018.2822053

Logothetis, N.K., Pauls, J., Augath, M., Trinath, T., Oeltermann, A., 2001. Neurophysiological investigation of the basis of the fMRI signal. Nature 412, 150–157. https://doi.org/10.1038/35084005

Lustig, M., Donoho, D., Pauly, J.M., 2007. Sparse MRI: The application of compressed sensing for rapid MR imaging. Magn. Reson. Med. 58, 1182–1195. https://doi.org/10.1002/mrm.21391

Lustig, M., Pauly, J.M., 2010. SPIRiT: Iterative self-consistent parallel imaging reconstruction from arbitrary k-space. Magn. Reson. Med. 64, 457–471. https://doi.org/10.1002/mrm.22428

Lutti, A., Thomas, D.L., Hutton, C., Weiskopf, N., 2013. High-resolution functional MRI at 3 T: 3D/2D echo-planar imaging with optimized physiological noise correction. Magn. Reson. Med. 69, 1657–1664. https://doi.org/10.1002/mrm.24398

Lv, H., Wang, Z., Tong, E., Williams, L.M., Zaharchuk, G., Zeineh, M., Goldstein-Piekarski, A.N., Ball, T.M., Liao, C., Wintermark, M., 2018. Resting-State Functional MRI: Everything That Nonexperts Have Always Wanted to Know. AJNR Am J Neuroradiol 39, 1390–1399. https://doi.org/10.3174/ajnr.A5527

Mani, M., Aggarwal, H.K., Magnotta, V., Jacob, M., 2020. Improved MUSSELS reconstruction for high-resolution multi-shot diffusion weighted imaging. Magn. Reson. Med. 83, 2253–2263. https://doi.org/10.1002/mrm.28090

Mani, M., Jacob, M., Kelley, D., Magnotta, V., 2017. Multi-shot sensitivity-encoded diffusion data recovery using structured low-rank matrix completion (MUSSELS). Magn. Reson. Med. 78, 494–507. https://doi.org/10.1002/mrm.26382

Mani, M., Magnotta, V., Jacob, M., 2018. A general algorithm for compensation of trajectory errors: Application to radial imaging. Magn. Reson. Med. 80, 1605–1613. https://doi.org/10.1002/mrm.27148

Marquardt, D.W., 1963. An Algorithm for Least-Squares Estimation of Nonlinear Parameters. Journal of the Society for Industrial and Applied Mathematics 11, 431–441. https://doi.org/10.1137/0111030

Mason, H.T., Graedel, N.N., Miller, K.L., Chiew, M., 2021. Subspace-constrained approaches to low-rank fMRI acceleration. Neuroimage 238, 118235. https://doi.org/10.1016/j.neuroimage.2021.118235

Moeller, S., Yacoub, E., Olman, C.A., Auerbach, E., Strupp, J., Harel, N., Uğurbil, K., 2010a. Multiband multislice GE-EPI at 7 tesla, with 16-fold acceleration using partial parallel imaging with application to high spatial and temporal whole-brain fMRI. Magn. Reson. Med. 63, 1144–1153.

Moeller, S., Yacoub, E., Olman, C.A., Auerbach, E., Strupp, J., Harel, N., Uğurbil, K., 2010b. Multiband multislice GE-EPI at 7 tesla, with 16-fold acceleration using partial parallel imaging with application to high spatial and temporal whole-brain fMRI. Magn. Reson. Med. 63, 1144–1153. https://doi.org/10.1002/mrm.22361

Mohan, K., Fazel, M., 2012. Iterative reweighted algorithms for matrix rank minimization. The Journal of Machine Learning Research 13, 3441–3473.

Nakarmi, U., Wang, Y., Lyu, J., Liang, D., Ying, L., 2017. A Kernel-Based Low-Rank (KLR) Model for Low-Dimensional Manifold Recovery in Highly Accelerated Dynamic MRI. IEEE Transactions on Medical Imaging 36, 2297–2307. https://doi.org/10.1109/TMI.2017.2723871

Narsude, M., Gallichan, D., van der Zwaag, W., Gruetter, R., Marques, J.P., 2016. Three-dimensional echo planar imaging with controlled aliasing: A sequence for high temporal resolution functional MRI: 3D-EPI-CAIPI: A Sequence for High Temporal Resolution fMRI. Magn. Reson. Med. 75, 2350–2361. https://doi.org/10.1002/mrm.25835

Nguyen, H., Glover, G., 2014. Field-corrected imaging for sparsely-sampled fMRI by exploiting low-rank spatiotemporal structure, in: Proc Int Soc Magn Reson Med 2014.

Ogawa, S., Lee, T.M., Kay, A.R., Tank, D.W., 1990. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. Proc. Natl. Acad. Sci. U.S.A. 87, 9868–9872. https://doi.org/10.1073/pnas.87.24.9868

Ogawa, S., Tank, D.W., Menon, R., Ellermann, J.M., Kim, S.G., Merkle, H., Ugurbil, K., 1992. Intrinsic signal changes accompanying sensory stimulation: functional

brain mapping with magnetic resonance imaging. Proceedings of the National Academy of Sciences 89, 5951–5955. https://doi.org/10.1073/pnas.89.13.5951
Ongie, G., Jacob, M., 2017. A Fast Algorithm for Convolutional Structured Low-Rank Matrix Recovery. IEEE Trans. Comput. Imaging 3, 535–550. https://doi.org/10.1109/TCI.2017.2721819
Otazo, R., Candès, E., Sodickson, D.K., 2015. Low-rank plus sparse matrix decomposition for accelerated dynamic MRI with separation of background and dynamic components. Magn. Reson. Med. 73, 1125–1136. https://doi.org/10.1002/mrm.25240
Park, J., Zhang, Q., Jellus, V., Simonetti, O., Li, D., 2005. Artifact and noise suppression in GRAPPA imaging using improved k-space coil calibration and variable density sampling. Magn. Reson. Med. 53, 186–193. https://doi.org/10.1002/mrm.20328
Park, S., Park, J., 2012. Adaptive self-calibrating iterative GRAPPA reconstruction. Magn. Reson. Med. 67, 1721–1729. https://doi.org/10.1002/mrm.23188
Pedersen, H., Kozerke, S., Ringgaard, S., Nehrke, K., Kim, W.Y., 2009. k-t PCA: Temporally constrained k-t BLAST reconstruction using principal component analysis. Magn. Reson. Med. 62, 706–716. https://doi.org/10.1002/mrm.22052
Petrov, A.Y., Herbst, M., Andrew Stenger, V., 2017. Improving temporal resolution in fMRI using a 3D spiral acquisition and low rank plus sparse (L+S) reconstruction. NeuroImage 157, 660–674. https://doi.org/10.1016/j.neuroimage.2017.06.004
Poddar, S., Jacob, M., 2016. Dynamic MRI Using SmooThness Regularization on Manifolds (SToRM). IEEE Transactions on Medical Imaging 35, 1106–1115. https://doi.org/10.1109/TMI.2015.2509245
Polak, D., Splitthoff, D.N., Clifford, B., Lo, W.-C., Huang, S.Y., Conklin, J., Wald, L.L., Setsompop, K., Cauley, S., 2022. Scout accelerated motion estimation and reduction (SAMER). Magn. Reson. Med. 87, 163–178. https://doi.org/10.1002/mrm.28971
Porter, D.A., Heidemann, R.M., 2009. High resolution diffusion-weighted imaging using readout-segmented echo-planar imaging, parallel imaging and a two-dimensional navigator-based reacquisition. Magn. Reson. Med. 62, 468–475. https://doi.org/10.1002/mrm.22024
Poser, B.A., Koopmans, P.J., Witzel, T., Wald, L.L., Barth, M., 2010. Three dimensional echo-planar imaging at 7 Tesla. NeuroImage 51, 261–266. https://doi.org/10.1016/j.neuroimage.2010.01.108
Pruessmann, K.P., Weiger, M., Börnert, P., Boesiger, P., 2001. Advances in sensitivity encoding with arbitrary k-space trajectories. Magn. Reson. Med. 46, 638–651. https://doi.org/10.1002/mrm.1241
Pruessmann, K.P., Weiger, M., Scheidegger, M.B., Boesiger, P., 1999. SENSE: Sensitivity encoding for fast MRI. Magn. Reson. Med. 42, 952–962. https://doi.org/10.1002/(SICI)1522-2594(199911)42:5<952::AID-MRM16>3.0.CO;2-S
Robson, P.M., Grant, A.K., Madhuranthakam, A.J., Lattanzi, R., Sodickson, D.K., McKenzie, C.A., 2008. Comprehensive quantification of signal-to-noise ratio and g-factor for image-based and k-space-based parallel imaging reconstructions. Magn. Reson. Med. 60, 895–907. https://doi.org/10.1002/mrm.21728
Roemer, P.B., Edelstein, W.A., Hayes, C.E., Souza, S.P., Mueller, O.M., 1990. The NMR phased array. Magn. Reson. Med. 16, 192–225. https://doi.org/10.1002/mrm.1910160203
Salimi-Khorshidi, G., Douaud, G., Beckmann, C.F., Glasser, M.F., Griffanti, L., Smith, S.M., 2014. Automatic denoising of functional MRI data: Combining independent

component analysis and hierarchical fusion of classifiers. NeuroImage 90, 449–468. https://doi.org/10.1016/j.neuroimage.2013.11.046

Setsompop, K., Gagoski, B.A., Polimeni, J.R., Witzel, T., Wedeen, V.J., Wald, L.L., 2012a. Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g-factor penalty. Magn. Reson. Med. 67, 1210–1224. https://doi.org/10.1002/mrm.23097

Setsompop, K., Gagoski, B.A., Polimeni, J.R., Witzel, T., Wedeen, V.J., Wald, L.L., 2012b. Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced *g* -factor penalty: Blipped-CAIPI for Simultaneous Multislice EPI. Magn. Reson. Med. 67, 1210–1224. https://doi.org/10.1002/mrm.23097

Shin, P.J., Larson, P.E.Z., Ohliger, M.A., Elad, M., Pauly, J.M., Vigneron, D.B., Lustig, M., 2014a. Calibrationless Parallel Imaging Reconstruction Based on Structured Low-Rank Matrix Completion. Magn. Reson. Med. 72, 959–970. https://doi.org/10.1002/mrm.24997

Shin, P.J., Larson, P.E.Z., Ohliger, M.A., Elad, M., Pauly, J.M., Vigneron, D.B., Lustig, M., 2014b. Calibrationless parallel imaging reconstruction based on structured low-rank matrix completion. Magn. Reson. Med. 72, 959–970. https://doi.org/10.1002/mrm.24997

Siemonsma, S., Kruger, S., Balachandrasekaran, A., Mani, M., Jacob, M., 2020. Multi-Echo Recovery with Field Inhomogeneity Compensation Using Structured Low-Rank Matrix Completion, in: 2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI). Presented at the 2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI), pp. 1074–1077. https://doi.org/10.1109/ISBI45749.2020.9098418

Smith, S.M., Vidaurre, D., Beckmann, C.F., Glasser, M.F., Jenkinson, M., Miller, K.L., Nichols, T.E., Robinson, E.C., Salimi-Khorshidi, G., Woolrich, M.W., Barch, D.M., Uğurbil, K., Van Essen, D.C., 2013. Functional connectomics from resting-state fMRI. Trends in Cognitive Sciences, Special Issue: The Connectome 17, 666–682. https://doi.org/10.1016/j.tics.2013.09.016

Stirnberg, R., Stöcker, T., 2021. Segmented K‐space blipped‐controlled aliasing in parallel imaging for high spatiotemporal resolution EPI. Magn. Reson. Med. 85, 1540–1551. https://doi.org/10.1002/mrm.28486

Tijssen, R.H.N., Jenkinson, M., Brooks, J.C.W., Jezzard, P., Miller, K.L., 2014. Optimizing RetroICor and RetroKCor corrections for multi-shot 3D FMRI acquisitions. Neuroimage 84, 394–405. https://doi.org/10.1016/j.neuroimage.2013.08.062

Tijssen, R.H.N., Okell, T.W., Miller, K.L., 2011. Real-time cardiac synchronization with fixed volume frame rate for reducing physiological instabilities in 3D FMRI. NeuroImage 57, 1364–1375. https://doi.org/10.1016/j.neuroimage.2011.05.070

Triantafyllou, C., Hoge, R.D., Krueger, G., Wiggins, C.J., Potthast, A., Wiggins, G.C., Wald, L.L., 2005. Comparison of physiological noise at 1.5 T, 3 T and 7 T and optimization of fMRI acquisition parameters. NeuroImage 26, 243–250. https://doi.org/10.1016/j.neuroimage.2005.01.007

Triantafyllou, C., Hoge, R.D., Wald, L.L., 2006. Effect of spatial smoothing on physiological noise in high-resolution fMRI. NeuroImage 32, 551–557. https://doi.org/10.1016/j.neuroimage.2006.04.182

Trzasko, J.D., Manduca, A., Borisch, E., 2011. Local versus Global Low-Rank Promotion in Dynamic MRI Series Reconstruction. in: Proc Int Soc Magn Reson Med 2011.

Trzasko, J.D., Manduca, A., 2012. CLEAR: Calibration-free parallel imaging using locally low-rank encouraging reconstruction, in: Proc. Int. Soc. Magn. Reson. Med. in: Proc Int Soc Magn Reson Med 2012.

Uecker, M., Lai, P., Murphy, M.J., Virtue, P., Elad, M., Pauly, J.M., Vasanawala, S.S., Lustig, M., 2014. ESPIRiT-an eigenvalue approach to autocalibrating parallel MRI: Where SENSE meets GRAPPA. Magn. Reson. Med. 71, 990–1001. https://doi.org/10.1002/mrm.24751

Unser, M., Thevenaz, P., Yaroslavsky, L., 1995. Convolution-based interpolation for fast, high-quality rotation of images. IEEE Trans Image Process 4, 1371–1381. https://doi.org/10.1109/83.465102

van der Zwaag, W., Marques, J.P., Kober, T., Glover, G., Gruetter, R., Krueger, G., 2012. Temporal SNR characteristics in segmented 3D-EPI at 7T. Magn. Reson. Med. 67, 344–352. https://doi.org/10.1002/mrm.23007

Wallace, T.E., Afacan, O., Kober, T., Warfield, S.K., 2020. Rapid measurement and correction of spatiotemporal B0 field changes using FID navigators and a multi-channel reference image. Magn. Reson. Med. 83, 575–589. https://doi.org/10.1002/mrm.27957

Walsh, D.O., Gmitro, A.F., Marcellin, M.W., 2000. Adaptive reconstruction of phased array MR imagery. Magn. Reson. Med. 43, 682–690. https://doi.org/10.1002/(SICI)1522-2594(200005)43:5<682::AID-MRM10>3.0.CO;2-G

Wang, F., Dong, Z., Reese, T.G., Bilgic, B., Katherine Manhard, M., Chen, J., Polimeni, J.R., Wald, L.L., Setsompop, K., 2019. Echo planar time-resolved imaging (EPTI). Magn. Reson. Med. 81, 3599–3615. https://doi.org/10.1002/mrm.27673

Weizman, L., Miller, K.L., Eldar, Y.C., Chiew, M., 2017. PEAR: PEriodic And fixed Rank separation for fast fMRI. Medical Physics 44, 6166–6182. https://doi.org/10.1002/mp.12599

Woolrich, M.W., Ripley, B.D., Brady, M., Smith, S.M., 2001. Temporal Autocorrelation in Univariate Linear Modeling of FMRI Data. NeuroImage 14, 1370–1386. https://doi.org/10.1006/nimg.2001.0931

Xu, L., Guo, L., Liu, X., Kang, L., Chen, W., Feng, Y., 2015. GRAPPA reconstruction with spatially varying calibration of self-constraint. Magn. Reson. Med. 74, 1057–1069. https://doi.org/10.1002/mrm.25496

Yi, Z., Liu, Y., Zhao, Y., Xiao, L., Leong, A.T.L., Feng, Y., Chen, F., Wu, E.X., 2021. Joint calibrationless reconstruction of highly undersampled multicontrast MR datasets using a low‐rank Hankel tensor completion framework. Magn. Reson. Med. 85, 3256–3271. https://doi.org/10.1002/mrm.28674

Zahneisen, B., Assländer, J., LeVan, P., Hugger, T., Reisert, M., Ernst, T., Hennig, J., 2014. Quantification and correction of respiration induced dynamic field map changes in fMRI using 3D single shot techniques: Respiration Induced Field Map Dynamics. Magn. Reson. Med. 71, 1093–1102. https://doi.org/10.1002/mrm.24771

Zhang, T., Pauly, J.M., Levesque, I.R., 2015. Accelerating Parameter Mapping with a Locally Low Rank Constraint. Magn. Reson. Med. 73, 655–661. https://doi.org/10.1002/mrm.25161

Zhao, B., Haldar, J.P., Brinegar, C., Liang, Z.-P., 2010. Low rank matrix recovery for real-time cardiac MRI, in: 2010 IEEE International Symposium on Biomedical Imaging: From Nano to Macro. Presented at the 2010 IEEE International Symposium on Biomedical Imaging: From Nano to Macro, pp. 996–999. https://doi.org/10.1109/ISBI.2010.5490156

Zhao, B., Haldar, J.P., Christodoulou, A.G., Liang, Z.-P., 2012. Image reconstruction from highly undersampled (k, t)-space data with joint partial separability and sparsity constraints. IEEE Transactions on Medical Imaging 31, 1809–1820. https://doi.org/10.1109/TMI.2012.2203921

Zhao, B., Setsompop, K., Adalsteinsson, E., Gagoski, B., Ye, H., Ma, D., Jiang, Y., Ellen Grant, P., Griswold, M.A., Wald, L.L., 2018. Improved magnetic resonance fingerprinting reconstruction with low-rank and subspace modeling. Magn. Reson. Med. 79, 933–942. https://doi.org/10.1002/mrm.26701

Zhou, Z., Wang, J., Balu, N., Li, R., Yuan, C., 2016. STEP: Self-supporting tailored k-space estimation for parallel imaging reconstruction. Magn. Reson. Med. 75, 750–761. https://doi.org/10.1002/mrm.25663

Zou, Q., Ahmed, A.H., Nagpal, P., Kruger, S., Jacob, M., 2021. Dynamic Imaging Using a Deep Generative SToRM (Gen-SToRM) Model. IEEE Transactions on Medical Imaging 40, 3102–3112. https://doi.org/10.1109/TMI.2021.3065948