Optimizing Magnetic Resonance Imaging for Image-Guided Radiotherapy

by

Lianli Liu

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Electrical Engineering: Systems)
in The University of Michigan
2018

Doctoral Committee:

Professor Jeffrey A. Fessler, Co-Chair
Professor James M. Balter, Co-Chair
Professor Alfred O. Hero III
Professor Douglas C. Noll

Lianli Liu

llliu@umich.edu

ORCID ID: 0000-0002-1095-0953

To my mom and dad

ACKNOWLEDGEMENTS

First, I would like express my sincere gratitude to my advisors, Professor Jeffrey Fessler and Professor James Balter. When I came to Michigan as a master student, Jeff and James granted me the opportunity to work with them, and to continue as a PhD student. Jeff always gives me deep insights into my research, and provides me strong guidance in analyzing problems mathematically. Deriving equations with Jeff has always been an intellectual enjoyment to me. James is very knowledgeable about both imaging physics and radiotherapy, and guides me not only in my research projects, but in various aspects of being a researcher, such as presentation skills and writing styles. I still remembered him edited my first manuscript more than fifteen times, for which I am really thankful to. Both Jeff and James are super friendly, liberal, and very supportive of career decisions I made. I can feel how much they care about their students and I really could not thank them enough.

I would also like to thank my committee members, Professor Douglas Noll and Professor Alfred Hero for the insightful suggestions they provided to my research. I benefited a lot from the depth and broadness of their knowledge. I would also like to thank Professor Yue Cao, for advising me on my research, sharing career stories with me, and suggesting the diffusion-weighted MRI project which becomes the second half of my PhD study.

My next big "Thank You" goes to Dr. Adam Johansson, my best friend and closest lab mate. Adam supports me in every way, from research to daily life. I thank him

for helping me with my research, keeping me accompany for hiking, traveling and cooking, and always being patient and listening to me. This thesis could not be finished without his support. I would also like to thank Priyanka Pramanik, a sweet girl who used to sit back-to-back with me at Argus. She brought me out of my shell, making me feel confident in talking to native speakers, and inspired me a lot with her considerate ways of taking care of others.

I would also like to thank everyone from both radiation oncology group and EECS group. I learned a lot from their work, as well as suggestions they gave to me. I thank Hao Sun, Hung Nien, Shu-Hui Hsu, Hesheng Wang, Madison McGaffin and David Hong for the useful discussions I had with them. Hao Sun and Hung Nien also helped a lot in making my career decisions, for which I am very grateful to. I would also like to thank Yuan Yao and Huan-Tsien Zeng for driving with me when I first got my car, and Sunan Cui, Lise Wei and Yuan Li for sharing girl's talks with me and making my life in Ann Arbor wonderful.

Finally, I would like to thank my mom, Sufen Li, and my dad, Ershuan Liu. Their unconditional love and support are always the greatest inspiration to me. I am particularly grateful for their vision and courage, of sending their only daughter aboard to pursue her career. This thesis is dedicated to them, for everything they give me.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

ABSTRACT

Magnetic resonance imaging (MRI) is playing an increasingly important role in image-guided radiotherapy. MRI provides excellent soft tissue contrast, and is flexible in characterizing various tissue properties including relaxation, diffusion and perfusion. This thesis aims at developing new image analysis and reconstruction algorithms to optimize MRI in support of treatment planning, target delineation and treatment response assessment for radiotherapy.

First, unlike Computed Tomography (CT) images, MRI cannot provide electron density information necessary for radiation dose calculation. To address this, we developed a synthetic CT generation algorithm that generates pseudo CT images from MRI, based on tissue classification results on MRI for female pelvic patients. To improve tissue classification accuracy, we learnt a pelvic bone shape model from a training dataset, and integrated the shape model into an intensity-based fuzzy c-menas classification scheme. The shape-regularized tissue classification algorithm is capable of differentiating tissues that have significant overlap in MRI intensity distributions. Treatment planning dose calculations using synthetic CT image volumes generated from the tissue classification results show acceptably small variations as compared to CT volumes. As MRI artifacts, such as B1 filed inhomogeneity (bias field) may negatively impact the tissue classification accuracy, we also developed an algorithm that integrates the correction of bias field into the tissue classification scheme. We modified the fuzzy c-means classification by modeling the image intensity as the true

intensity corrupted by the multiplicative bias field. A regularization term further ensures the smoothness of the bias field. We solved the optimization problem using a linearized alternating direction method of multipliers (ADMM) method, which is more computational efficient over existing methods.

The second part of this thesis looks at a special MR imaging technique, diffusion-weighted MRI (DWI). By acquiring a series of DWI images with a wide range of b-values, high order diffusion analysis can be performed using the DWI image series and new biomarkers for tumor grading, delineation and treatment response evaluation may be extracted. However, DWI suffers from low signal-to-noise ratio at high b-values, and the multi-b-value acquisition makes the total scan time impractical for clinical use. In this thesis, we proposed an accelerated DWI scheme, that sparsely samples k-space and reconstructs images using a model-based algorithm. Specifically, we built a 3D block-Hankel tensor from k-space samples, and modeled both local and global correlations of the high dimensional k-space data as a low-rank property of the tensor. We also added a phase constraint to account for large phase variations across different b-values, and to allow reconstruction from partial Fourier acquisition, which further accelerates the image acquisition. We proposed an ADMM algorithm to solve the constrained image reconstruction problem. Image reconstructions using both simulated and patient data show improved signal-to-noise ratio. As compared to clinically used parallel imaging scheme which achieves a 4-fold acceleration, our method achieves an 8-fold acceleration. Reconstructed images show reduced reconstruction errors as proved on simulated data and similar diffusion parameter mapping results on patient data.

CHAPTER 1

Introduction

1.1 Image-Guided Radiotherapy (IGRT)

Radiotherapy is the use of radiation, usually X-rays from linear accelerators to treat illness. About 40% patients with cancer have radiotherapy as part of their treatment [1].

Medical imaging provides important guidance to radiotherapy. For example, imaging scans of patients before treatment show the size, shape and position of the cancer as well as the surrounding tissues and bones. Doctors then plan the radiation treatment to give a high dose to the cancer and try to give as low a dose as possible to the surrounding healthy tissue to reduce the risk of side effects [2]. Besides treatment planning, medical imaging has also been used for tumor evaluation, tumor tracking during treatment, verification of radiation therapy delivery as well as treatment response evaluation. To summarize, image-guided radiotherapy (IGRT) refers to the use of imaging during a course of radiation therapy, with decisions made on the basis of imaging [3]. Improving medical imaging techniques has a strong potential in improving the efficacy of radiotherapy, and thus the survival rate and life quality of cancer patients. This thesis develops new image processing techniques to extract richer and more accurate information from medical images and reduce total image

acquisition time, as an attempt to optimize medical imaging for radiotherapy.

1.2   Magnetic Resonance Imaging (MRI) in IGRT

Imaging for radiation therapy treatment planning has been largely based on Computed Tomography (CT) imaging. Recently, Magnetic Resonance Imaging (MRI) is gaining popularity. Comparing to CT scans, MRI scans provide superior soft tissue and tumor contrast for more accurate delineation. Also, MRI is a very flexible imaging modality. By varying imaging parameters, MRI can provide various information, such as spin intensity, relaxation property and diffusitivity of tissues. Such information helps better visualization of abnormalities and serves as potential biomarkers for individual customization of radiation dose to patients as treatment progresses.

However, certain challenges exist in using MRI for IGRT. Firstly, MRI cannot provide radiation attenuation information directly that is needed for radiation dose calculation. As a result, CT scans might have to be acquired in addition to MRI scans. Acquiring two scans would come at significant cost, in both scanning resources and the need to reconcile geometric variations between two models of the patient taken at different times. Such reconciling is usually challenging due to the change of patient anatomy at different times, thus can introduce significant systematic errors in radiotherapy treatment planning and reduce the efficacy of radiotherapy [4]. Besides, certain MRI imaging techniques require long scan time, limiting their usefulness for radiation therapy simulations of immobilized patients. Other challenges include geometric distortion [5], intensity inhomogeneity [6] and other imaging artifacts that may affect image quality and accurate information extraction from images.

1.3   Contributions and Outline

In this work, we focus on improving MRI in two aspects of IGRT: treatment planning and treatment evaluation response. To use MRI only for treatment planning without CT imaging, radiation attenuation maps, commonly referred as "synthetic CT", have to be estimated from MRI images. This appears challenging due to both imaging artifacts in MRI that corrupt true image intensity, and the lack of one to one correspondence between MRI intensities and attenuation properties. To address this challenge, we first developed an MRI intensity non-uniformity ("bias field") correction algorithm that jointly corrects bias field and classifies different tissue types [7]. Next, we developed a bone shape model for female pelvic radiotherapy patients [8][9], to address the intensity overlap between bone and air in MRI images, a major challenge for synthetic CT generation. Then we presented a complete pelvic synthetic CT generation algorithm [10][11] that incorporates the shape model and utilizes both intensity and shape information from MRI data.

Diffusion-weighted MRI (DWI) is a promising imaging technique in treatment response evaluation [12] and other aspects of cancer treatment. It has been shown useful in assessing the tumor cellularity, grading tumors and differentiating recurrent or progressive tumor growth from treatment-induced damage to brain parenchyma in high-grade gliomas after radiation therapy [13]. Recent studies show diffusion-weighted MRI with high b-value has potential use for the early detection of response to therapy [14] and more precise tumor identification for full coverage of tumors in radiotherapy [15]. However, high b-value diffusion-weighted imaging is time consuming due to the extended b-value range and the repeated acquisitions needed to improve image quality. Furthermore, non-mono exponential decay behaviors of tis-

sues have been demonstrated under high b-values [16], thus a new model other than the conventional mono-exponential decay model may be important for extracting more useful metric values. The second part of this work aims at reducing the acquisition time for diffusion-weighted MRI by reconstructing DWI image from sparsely sampled data, using model-based reconstruction algorithms. We built a low-rank tensor model that allows us to exploit the local dependencies and global dependencies of k-space samples simultaneously, and proposed an efficient algorithm to solve the constrained image reconstruction problem [17][18].

The rest of this thesis organized as follows. Chapter 2 briefly reviews imaging principles of CT and MR imaging. Chapter 3 presents the joint intensity non-uniformity correction and tissue classification algorithm for whole brain MRI [7]. Chapter 4 describes the bone shape model for female pelvic radiotherapy patients [8][9] and Chapter 5 describes the complete female pelvic synthetic CT generation algorithm that integrates the bone shape model in Chapter 4 [10][11]. Chapter 6 presents the low-rank tensor model with phase constraints for accelerated diffusion-weighted MRI reconstruction [17][18]. Finally, Chapter 7 summarizes future work.

CHAPTER 2

Background

2.1 CT Imaging Principles

As a stream of x-ray photons passes through an object, some of them will interact with the object material, either get absorbed or scattered, while the remaining will pass through. The change in the number of photons before and after passing the object is determined by the properties of the object, represented by the attenuation coefficient. Consider a x-ray beam travels through an object along a line at a distance of x, the number of photons will change from $N(0)$ to $N(x)$:

$$N(x) = N(0) \exp(-\mu x), \tag{2.1}$$

where μ is the linear attenuation coefficient and has units $(\text{distance})^{-1}$.

Computed Tomography (CT) produces images that are characterized by the attenuation properties of tissue. During the imaging process, the object is placed between a source array and a detector array, where the source produces x-ray flux and the detector records the photons that pass through the object [19]. Attenuation coefficients are directly related with the electron density of tissues, therefore CT images provide a precise mapping to Relative Electon Density (RED) for dose calculation. Low density material such as air appears dark on the final image while dense material such as bone appears bright. However, CT images are limited in the
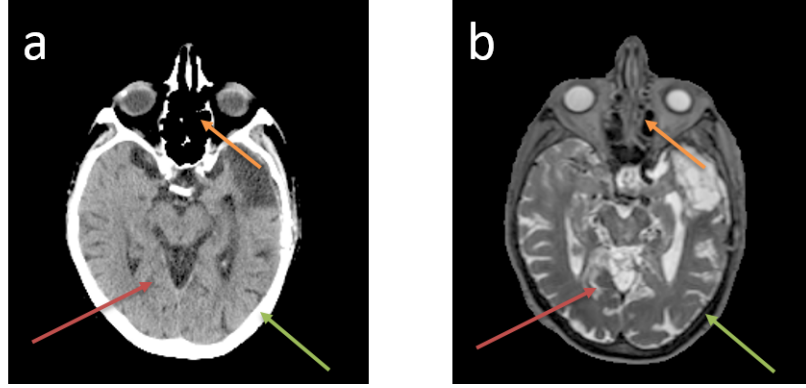
Figure 2.1: (a) Example CT and (b) MR images of one patient with brain tumor. Bone (green arrow) is bright in CT and dark in MR images. Air (yellow arrow) is dark in both CT and MR images. MRI shows better soft tissue contrast (red arrow). The window widths/window levels for CT and MRI are 80/40 and 4099/2049 respectively.

ability of differentiating soft tissues (such as white matter and gray matter in the brain) and tumors. Figure 2.1(a) shows a sample CT images of human brain with tumor.

2.2 MR Imaging Principles

Magnetic Resonance Imaging (MRI) is another imaging methodology that characterizes tissue properties different from CT images. A physical object can be broken down into its constituent molecules, then to atoms, and then to nuclei and their orbiting electrons [20]. Nucleus with odd atomic weights or numbers possess an angular momentum and are referred as spins. In general, MRI images the magnetic moments that result from the quantum mechanical property of nuclear spin. Among all the spins in the body, Hydrogen (^{1}H) is the most abundant and is imaged in most of studies in human MRI [20]. To produce an MRI image, three magnetic fields are needed: the main field ($\vec{B}_0$) for signal preparation, the radio-frequency (RF) field ($\vec{B}_1$) for signal excitation and the gradient field for spatial encoding ($\vec{G}$).

2.2.1 Signal Preparation, Excitation and Relaxation

Without an external magnetic field, the spins in the body are randomly oriented resulting in a net magnetic moment of zero. After applying the main field $\vec{B}_0$, the object develops a distribution of magnetization, $m_0(x, y, z)$, that is part of what we wish to image. This magnetization is aligned with $\vec{B}_0$ (by convention, the direction of $\vec{B}_0$ is referred as the z direction).

The RF field $\vec{B}_1$ is then applied to tip the magnetization away from $\vec{B}_0$. When the magnetization is tipped away from $\vec{B}_0$, by the *Larmor relationship*, it will precess around the z-direction at the *Larmor frequency*

$$\omega = \gamma B_0, \tag{2.2}$$

where γ is the gyromagnetic ratio. For Hydrogen, $\gamma/2\pi = 42.58$ MHz/T [20]. This precessing magnetization induces a current in a receive coil, which is acquired for subsequent processing.

Following excitation (assuming the magnetization is tipped 90 degree away from the main field), the magnetization in the transverse plane decays away with time constant T2 :

$$m_{xy}(t) = m_0 e^{-t/T2}. \tag{2.3}$$

This process is called "T2 relaxation" or "spin-spin relaxation", and the z component recovers with time constant T1:

$$m_z(t) = m_0(1 - e^{-t/T1}). \tag{2.4}$$

This process is called "T1 relaxation" or "spin-lattice relaxation". The above excitation-relaxation process is then repeated.

Since different biological tissues have different T1 and T2 values, by selecting the repetition time TR (time between RF pulses) and echo time TE (observation time

following the RF pulse), we will get "T1-weighted images", where tissues with long T1s appear dark and tissues with short T1s are bright, or "T2-weighted" image, where tissues with long T2s appear bright while tissues with short T2s are dark. Though soft tissues usually have similar electron density, their T1 and T2 values can differ. MR images characterize such differences and provide better soft tissue contrast than CT images. As shown in Figure 2.1(b), tumors are more easily identified from MR images.

2.2.2 Spatial Encoding and k-space

The RF pulse tips the magnetization away from the main field and creates voltage signals in receive coils. Assuming the magnetization is tipped to a plane perpendicular to the $\vec{B}_0$, thus $m(z) = 0$, and considering the precessing process in a reference frame that rotates around $\vec{B}_0$ at the *Larmor frequency* (usually referred as the rotating frame), the received signal is related with the precessing magnetization as

$$S(t) = \int_V m_{xy,\text{rot}}(\vec{r}, t)d\vec{r}. \tag{2.5}$$

Since the main magnetic field is uniform, all spins in the object precess at the same frequency. To establish one-to-one correspondence between spatial location and frequency, gradient fields are applied in the direction of the main field.

The strength of gradient fields vary in space. For example, after applying two gradient fields that vary along x and y direction with linear gradient G_x and G_y respectively, the magnetic field along the main field direction becomes

$$B_z(x, y, z, t) = B_0 + G_x(t)x + G_y(t)y, \tag{2.6}$$

by the *Larmor relationship* in Eq(2.2), the precessing frequency at a spatial location (x, y) is

$$\omega(x, y, t) = \gamma(B_0 + G_x(t)x + G_y(t)y), \tag{2.7}$$

which in the rotating frame is

$$\Delta\omega(x, y, t) = \gamma(G_x(t)x + G_y(t)y), \tag{2.8}$$

the magnetization $m_{xy,rot}(x, y, t)$ develops a spatially varying phase

$$m_{xy,rot}(x, y, t) = m(x, y) \exp(-i\phi(x, y, t)), \tag{2.9}$$

where

$$\phi(x, y, t) = \int_0^t \gamma(G_x(\tau)x + G_y(\tau)y)d\tau. \tag{2.10}$$

Substituting this new magnetization into the signal equation (2.5), it can be shown the received signal is

$$
\begin{aligned}
S(t) &= \iint m_{xy,rot}(x, y, t)dxdy \\
&= \iint m(x, y) \exp(-i\phi(x, y, t))dxdy \\
&= \iint m(x, y) \exp\left(-i\gamma\left(\int_0^t G_x(\tau)d\tau \cdot x + \int_0^t G_y(\tau)d\tau \cdot y\right)\right)dxdy. \quad (2.11)
\end{aligned}
$$

Denoting

$$k_x(t) = \frac{\gamma}{2\pi} \int_0^t G_x(\tau)d\tau \tag{2.12a}$$

$$k_y(t) = \frac{\gamma}{2\pi} \int_0^t G_y(\tau)d\tau, \tag{2.12b}$$

we rewrite Eq(2.11) as

$$
\begin{aligned}
S(t) &= \iint m(x, y) \exp\left(-i2\pi\left(xk_x(t) + yk_y(t)\right)\right)dxdy \\
&= \mathcal{F}\{m(x, y)\}|_{k_x(t), k_y(t)} = M(k_x(t), k_y(t)). \quad (2.13)
\end{aligned}
$$

From Eq(2.13), the signal received by the coil is the Fourier transform of the magnetization of the object. The Fourier transformed space, $M(k_x(t), k_y(t))$, is commonly referred as the k-space.
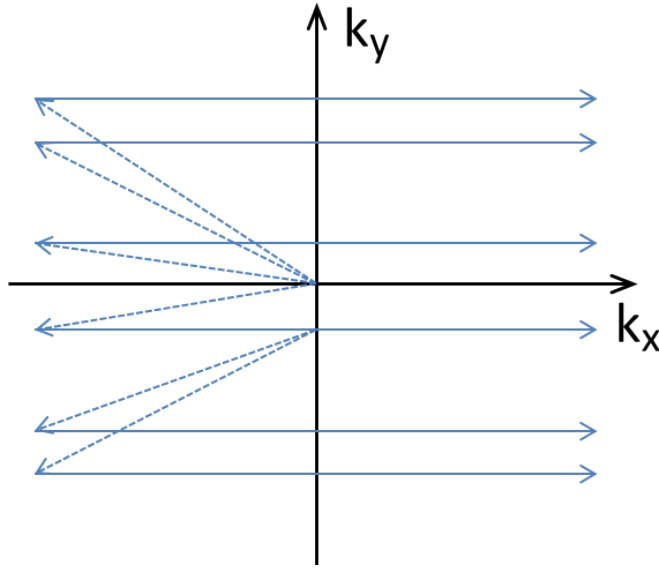
Figure 2.2: The most commonly used sampling pattern for 2D MRI.

The gradients, $G_x(t)$ and $G_y(t)$ control the trajectory on which the sample locations fall. It is also possible to apply gradients that vary along different trajectories instead of x and y direction. Different gradient fields result in different sampling patterns of the k-space. Figure 2.2 shows the most commonly used sampling pattern for 2D imaging. In this case, the entire k-space is sampled line-by-line along x direction. $G_x(t)$ remains constant during the acquisition to sample each line, while $G_y(t)$ is only turned on briefly before the acquisition of each line to move the sampling location one line above/below. The x direction is known as the "gradient encoding direction" while the y direction is known as the "phase encoding direction".

2.2.3 Spatial Range and Resolution

By Nyquist's sampling theorem, to recover the image of the object, we must sample the k-space densely enough, otherwise aliasing will occur. Also, to recover details of the image, we must acquire enough k-space samples to cover sufficiently high frequency components.

For 2D imaging, denote the sample spacing in k-space as Δk_x and Δk_y, and define

$$W_{kx} = N_x \Delta k_x \tag{2.14a}$$

$$W_{ky} = N_y \Delta k_y, \tag{2.14b}$$

where N_x and N_y are the number of k-space samples along x and y direction respectively. The field of view of an acquisition is

$$\text{FOV}_x = \frac{1}{\Delta k_x} \tag{2.15a}$$

$$\text{FOV}_y = \frac{1}{\Delta k_y}. \tag{2.15b}$$

The field of view gives the limits of the spatial range of the object to be imaged. If the object is larger than the field of view, aliasing will occur.

The spatial resolution is determined by

$$\Delta x = \frac{1}{W_{kx}} \tag{2.16a}$$

$$\Delta y = \frac{1}{W_{ky}}. \tag{2.16b}$$

Therefore, to image the object without aliasing and achieve high-resolution image, we must sample k-space densely enough with sufficient number of samples. This makes the scan time of MRI long in some applications.

2.3 Special MR Imaging Sequences Relevant to This Study

From section 2.2, MRI is a very flexible imaging modality with various imaging parameters that can be adjusted, such as RF field, gradient field (sampling pattern), echo time (TE) and repetition time (TR). Combining these imaging parameters results in an imaging sequence and different imaging sequences can characterize different biological characteristics. In this section, we briefly review two special imaging sequences that are relevant to our study.
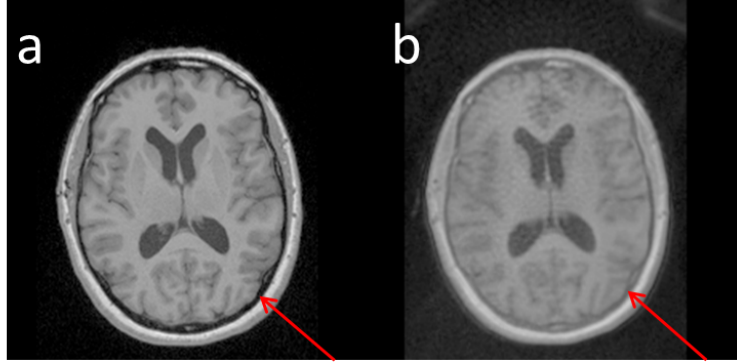
Figure 2.3: Example (a) T1-weighted and (b) UTE image of one patient. Bony tissues (red arrow) show higher signal on UTE than on T1-weighted image.

2.3.1 Ultra-short TE Imaging

Although MRI presents superior soft tissue contrast over CT imaging, visualizing bony tissue in MRI appears challenging. Bony tissues have very short T2 values. For example, cortical bone has a mean T2 of 0.42 ms to 0.50 ms, while most clinically used MRI systems use TE values ranging from 5 ms to 200 ms [21]. By Eq(2.3), signals from tissues with such short T2 will decay away before detection. As a result, bony tissues usually appear dark on MRI images, making them inseparable from tissues with low spin density, such as air cavities.

Ultra-short TE imaging (UTE) is developed to visualize tissues with short T2, such as bone. The echo time for a UTE sequence ranges from 0.05 ms to 0.50 ms [21], so that signals from tissues with short T2s can be detected before decaying away. Figure 2.3 shows an example of UTE imaging, where bony tissues show higher signals than conventional imaging.

Some additional care needs to be taken to design a sequence with such short TE. For example, center-out radial sampling is used instead of Cartesian sampling shown in Figure 2.2, with the aim of eliminating phase encoding and reducing TE. Besides, the RF pulse is truncated. Slice selection is achieved in two halves with reversed

gradients and data acquisition can begin as soon as each half of the slice selection is completed [21].

Various UTE sequences have been proposed. Besides the UTE sequence with radial sampling as mentioned above (referred as UTE in the rest of this work), another UTE sequence that combines Cartesian sampling at the center of the k-space with radial sampling at the outer, named pointwise encoding time reduction with radial acquisition (PETRA) [22] was also used in this work.

2.3.2 Diffusion-Weighted Imaging (DWI)

Besides imaging the T1 and T2 properties of tissues, it is also possible to use MRI to image the random movement of water molecules and learn about tissue properties from water mobility.

The random movement of particles suspended in a fluid or gas is referred as Brownian motion. The statistical relationship between the average distance that particles move over an interval of time is established as [23]

$$< \Delta r^2 >= 6D\Delta t, \tag{2.17}$$

where $< \Delta r^2 >$ represents the average squared displacement of a particle, allowed to diffuse freely over the time interval Δt, and D is the diffusion coefficient that is affected by both temperature and media properties. The diffusion coefficient D of pure water at body temperature (37°C) is approximately 3×10^{-3}mm^2/s [24].

When water molecules move randomly in tissues, they will encounter many cellular and subcellular impediments, as well as impediment presented by the tortuosity of cell packing in dense tissues [25]. As a result, water mobility is reduced by non-water tissue constituents. By probing water mobility using MR imaging, we can learn important information about such non-water constituents tissues, including physic-
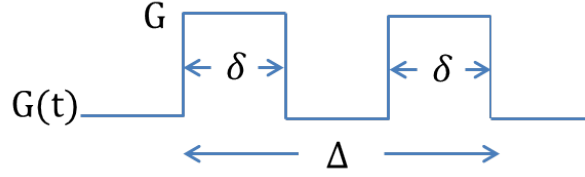
Figure 2.4: Waveform of the diffusion sensitizing gradient

ochemical properties of the tissue as well as its structural components [26]. Such information is useful in detecting and analyzing abnormalities, for example, tumor characterization. Water mobility in tumors is usually even more reduced than in normal tissues. Therefore tumors will present contrast different from normal tissues on diffusion-sensitive MR imaging and tumor-related information can be learned from such contrast.

Diffusion-weighted MRI incorporates diffusion-sensitization gradients into an imaging sequence. The default diffusion sensitive sequence is the two-pulse gradient waveform shown in Figure 2.4. Define b-value [27] as

$$b = (\gamma G \delta)^2 [\Delta - \frac{\delta}{3}], \qquad (2.18)$$

signal attenuation due to random spin diffusion is given by

$$S(b) = S_0 \exp(-Db), \qquad (2.19)$$

where S_0 is the signal without diffusion attenuation ($b = 0$). We can vary b-value by changing the gradient amplitude G or gradient duration δ in Eq(2.18) and get a series of diffusion signals as a function of b-value. Fitting the signal series to Eq(2.19) gives the diffusion coefficient.

Eq (2.19) assumes isotropic diffusion, which dose not always hold for tissues [28]. For anisotropic diffusion, the single valued diffusion coefficient in Eq (2.19) is gener-

14

alized to a diffusion tensor

$$D = \begin{bmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{yz} & D_{yy} & D_{yz} \\ D_{zx} & D_{zy} & D_{zz} \end{bmatrix}, \tag{2.20}$$

and the b-value in Eq (2.18) is generalized to

$$b_{ij} = \gamma^2 \int_0^{TE} \left(\int_0^t G_i(t')dt' \right) \left(\int_0^t G_j(t')dt' \right) dt, \tag{2.21}$$

where G_i and G_j denote gradients along different directions. The diffusion signal equation is

$$S(b) = S_0 \exp\left(- \sum_{i=1}^{3} \sum_{j=1}^{3} b_{ij} D_{ij} \right). \tag{2.22}$$

One can improve the accuracy of diffusion tensor calculation by acquiring a large number of non-colinear diffusion gradient directions, this is referred as diffusion tensor imaging (DTI) [29]. For diffusion-weighted imaging (DWI), the diffusion tensor is simplified to the mean diffusivity [24], D_{ave} as

$$D_{\text{ave}} = \frac{D_{xx} + D_{yy} + D_{zz}}{3}. \tag{2.23}$$

D_{ave} is important as it accurately summarizes water mobility and is rotationally invariant [30]. Only three orthogonal gradient directions are needed to calculate D_{ave}. The b-values associated with the three gradient directions have the same amplitude, thus we use the amplitude only to describe the b-value in DWI. The images acquired using the three gradients, denoted as $\mathbf{I}_x, \mathbf{I}_y, \mathbf{I}_z$ are reconstructed separately first. The final DWI image volume $\mathbf{I}$ is obtained by taking the geometric mean of the three images

$$\mathbf{I} = \sqrt[3]{\mathbf{I}_x \mathbf{I}_y \mathbf{I}_z}. \tag{2.24}$$

After obtaining a series of DWI image volumes under different b-values, a linear fitting will be performed to calculate the mean diffusivity using the mono-exponential decay model in Eq(2.19). The calculated mean diffusivity is qualified as the "apparent diffusion coefficient" (ADC).

Eq(2.19) and Eq(2.22) describe molecular mobility by simple diffusion coefficients, while water movement through cellular tissues is very complicated and can be affected by various factors, such as perfusion induced by blood flow and microcirculation [24]. The signal attenuation due to perfusion is most apparent at relatively low b-values (approximately $b \leq 100 \mathrm{s/mm}^2$) [31]. Therefore in practice, besides the DWI image at $b = 0$, at least 2 other DWI images with one at modest b-values ($\sim 50 - 200 \mathrm{\ s/mm}^2$) and another at a higher b-values ($\sim 800 - 1000 \mathrm{\ s/mm}^2$) are needed to reduce the effect of perfusion in calculating ADC. Another factor that will influence the diffusion measurement is motion. Even the modest bulk tissue motion can result in large phase shifts. To address this, all spatial encoding is performed in a single shot (known as echo planar imaging (EPI) [32]). The phase values of the reconstructed images can then be discarded when analyzing DWI data.

Joint MRI Bias Field Correction and Tissue Classification in Support of Whole Brain Synthetic CT Generation for Radiotherapy

3.1 Introduction

[1] CT and MR images are typically used in radiotherapy today, as CT provides direct radiation attenuation information needed for radiation dose calculation while MR images provide superior soft tissue contrast for tumor and organ delineation. MR images are aligned to CT images to transfer soft tissue contrast information seen on MRI into the CT-based patient representation for treatment planning. In addition to using significant resources, such a transfer could potentially introduce systematic errors due to registration uncertainties [4].

Interest is emerging to use MRI without CT to support radiation dose calculation, attenuation correction of positron emission tomography (PET) in a PET-MRI system and some aspects of image guidance in radiation therapy. To support these roles, radiation attenuation information, or synthetic CT (MRCT) volumes, need to be generated from MRI scans.

Various techniques have been proposed for MRCT generation. These include atlas based methods, where electron density maps are generated by aligning an atlas

---

[1]This chapter is based my work presented at the 2_{nd} MR in RT symposium, St Louis, 2014 [7]

derived from reference CT images with target MR images [33] [34] [35] [36]. However, such methods are somewhat limited in their ability to adapt to patient anatomical variations. Other algorithms apply a relation between the attenuation properties and image intensities of one or more MRI scans to generate MRCT images [4] [37] [38]. The major challenge of such methods is the ambiguity in the correspondence between attenuation properties and image intensities. For example, bony tissues have low signal intensities in MR images due to their short T2 and air has low signal due to extremely low electron density, yet the attenuation properties of bone and air are significantly different. Imaging artifacts, partial volume effects and noise in MRI further complicate the intensity distribution of different tissues, resulting in misclassifications of tissue types, and thus wrong attenuation assignments in MRCT images.

Ultrashort echo time scanning techniques, as described in section 2.3.1, are able to yield signals from tissues with short T2 and have been used in MRCT generation for the head to improve the separation of bone from air. For example, Hsu *et al.* described an intensity-based classification scheme for generating synthetic CT image volumes using multiple MR images, including UTE image, as input [4]. However, there are a number of factors that influence the image intensity in MRI, including the B1 field inhomogeneity that is impacted by RF transmit and receive variations across the patient [6]. The image intensities corrupted by the B1 field inhomogeneity (also referred as the bias field) may negatively impact intensity-driven quantitative image analysis. It has been noted that bias field correction may be needed for accurate synthetic CT generation [39].

Various bias field correction methods have been proposed, including prospective ones that attempt to correct the bias field during the imaging procedure [6] [40] [41],

by using specially designed imaging sequences or multiple coils. Retrospective methods have also been used that involve image processing after image reconstruction. As compared to prospective methods, retrospective techniques are more general as they do not make assumptions about the image acquisition process and are able to further resolve patient-induced inhomogeneity [42]. One commonly used retrospective method estimates the bias field by sharpening the intensity histogram of bias field-corrupted image volumes, without any prior knowledge of the underlying patient-generated image signals [43] [44]. Other investigations have integrated bias field estimation directly into classification solutions for segmentation of MR image data. In these methods, a regularization scheme is usually needed to ensure that the bias field is slow-varying. Example regularization schemes used include Gaussian smoothing [45] [46], B-spline fitting [47] and a neighborhood constraint [48]. These methods focus mainly on classification of brain tissue (white matter, gray matter, fluid) only.

This chapter investigates a joint bias field correction algorithm for whole brain (including bone and air cavities) data, with special focus on the potential for bias field correction to improve the accuracy of synthetic CT generation while maintaining computational efficiency.

3.2 Method and Materials

3.2.1 Image acquisition

6 patients were scanned on a 3T MRI scanner (Skyra, Siemens, Erlangen, Germany) under an institution review board-approved protocol. T2-weighted image volumes were acquired using a 3D Spin Echo (SPACE) sequence with TE/TR = 409/3200 ms, flip angle = 120°, and voxel size =1×1×1 mm^3. T1-weighted as well as fat and water images were generated using a 3D ultrafast gradient echo se-

quence (VIBE Dixon, initially designed for breath-held liver imaging) with TE/TR =2.46/4.1 ms, flip angle 9°, and voxel size≈1.27×1.27×1 mm³. Fat and water images were computed from the in-phase (T1-weighted) and out-of-phase VIBE images. To limit mis-classification of flowing blood as bone, a 2D time-of-flight sequence was acquired with TE = 4 ms, TR = 8.6 ms and flip angle = 20°. This sequence creates brighter intensities on blood vessels than the brain tissue due to the time-of-flight (TOF) effect. To limit misclassification of air as bone, PETRA images were acquired with TE/TR = 0.07/4.69 ms, flip angle 6°, and voxel size ≈ 1.71×1.71×1.71 mm³. Patients were scanned immobilized in custom thermoplastic masks. An 18-channel surface coil was suspended above the patient's head using a plastic bridge, and used in combination with a posterior spine coil for imaging. The vendor's B1 field normalization was applied to all image acquisitions and the configuration of the coils and subjects was not modified between acquisitions, with the expectation of a common residual bias field across all images. All patients also had CT scans acquired on a wide bore CT simulation system (Brilliance, Philips, Cleveland OH) with slice thickness of 1-3 mm and in-plane pixel size ≈ 0.5×0.5 mm². CT image volumes were registered to T2-weighted image volumes using rigid alignment. All images were reformatted to a common voxel size of 1×1×1mm³. Example MR and CT images used from one subject are shown in Figure 3.1 and Figure 3.2 respectively.

3.2.2 Modified fuzzy c-means for joint bias field correction and classification

In the synthetic CT generation scheme investigated by [4], voxel classification is achieved via fuzzy c-means (FCM) clustering with a spatial constraint. Given an image dataset $\{x_{ks}\}_{k=1,s=1}^{N,d}$ from d MRI volumes of the same object, each containing N voxels, as well as the total number of tissue classes presented c, standard fuzzy c-means classification with a spatial constraint was performed to get the probabilistic
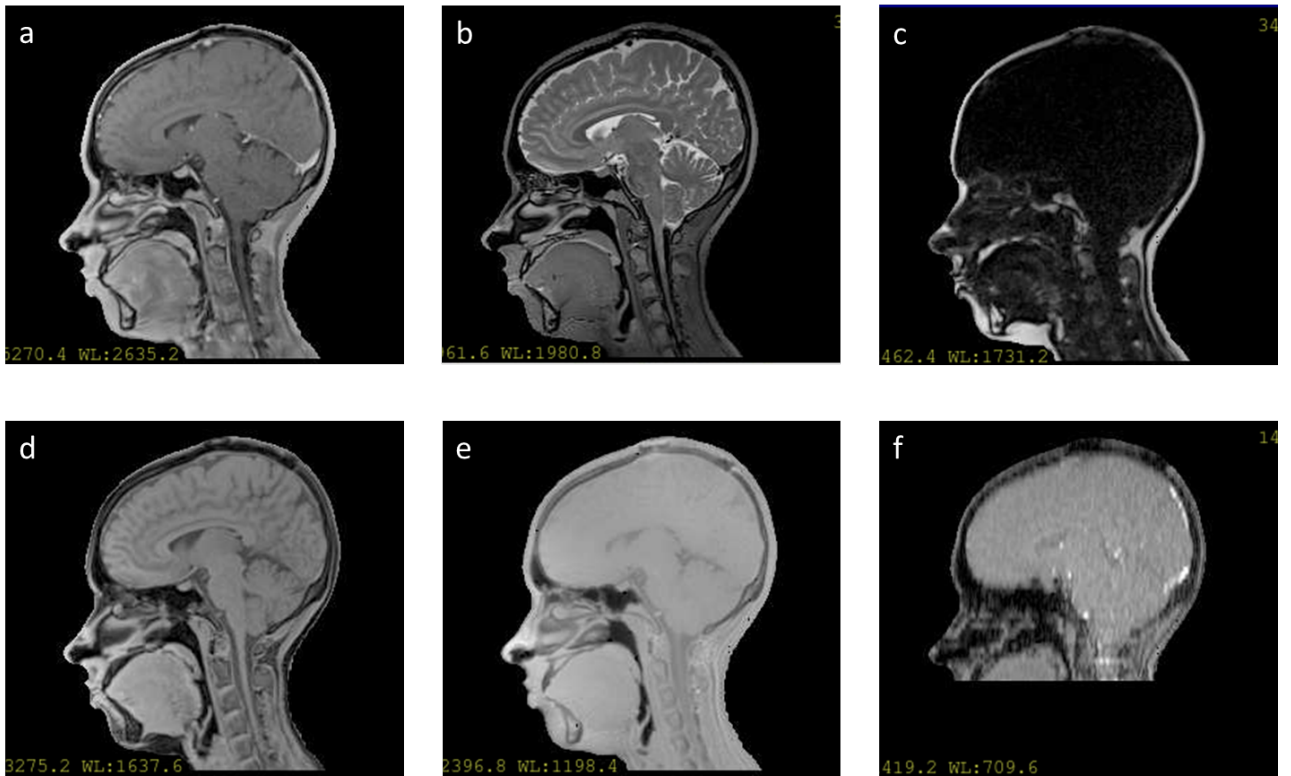
Figure 3.1: Example MR images from one study subject. (a) T1-weighted, (b) T2-weighted, (c) Fat, (d) Water, (e) PETRA and (f) TOF image.

Figure 3.2: Example reference CT image from a study subject

membership u_{ik} of the k^{th} voxel belonging to the i^{th} tissue class, whose intensity centroid on the s^{th} MR image is denoted by v_{is}:

$$\hat{\mathbf{u}}, \hat{\mathbf{v}} = \arg\min \sum_{s=1}^{d}\sum_{i=1}^{c}\sum_{k=1}^{N} u_{ik}^{m}\|x_{ks} - v_{is}\|^{2} + \alpha \sum_{s=1}^{d}\sum_{i=1}^{c}\sum_{k=1}^{N} u_{ik}^{m}\|\bar{x}_{ks} - v_{is}\|^{2}$$

$$s.t. \sum_{i=1}^{c} u_{ik} = 1, \quad \forall k = 1, \cdots, N, \tag{3.1}$$

where m is the fuzzy degree that takes the partial volume effect into consideration. The first term is a simple fuzzy c-means clustering. The second term is a spatial constraint to improve connectivity and suppress noise, where $\bar{x}_{ks}$ denotes the median of the neighbors within a kernel and α controls the weight for the spatial constraint.

3.2.2.1 Modeling the bias field

The observed voxel intensity x_k is not the true intensity value but the true intensity y_k corrupted by a bias field that takes value b_k at voxel location k. The bias field is usually modeled as a slowly-varying multiplicative signal, and thus the relationship between true and observed voxel intensities can be expressed as

$$x_k = y_k b_k + w_k, \tag{3.2}$$

where w_k is additive random noise.

One way to deal with the effect of bias fields on tissue classification is to perform image processing using algorithms such as N4itk [44]. Fuzzy c-means classification can then be performed on bias-corrected images. Previous studies showed that more than one MRI volume (or contrast) may be required to achieve accurate classification of the major tissue types in the head [4] [38]. Image processing methods such as N4itk [44] typically estimate a bias field from each individual MR image volume, and could yield inconsistent results across images of different contrasts. Another

possible approach involves estimating the bias field jointly from images of multiple contrasts. This methodology can be further integrated with synthetic CT generation by modifying fuzzy c-means classification to incorporate bias field estimation and correction. Considering that the bias field should be consistent across multiple image volumes, the use of multiple (e.g., T1-weighted, T2-weighted) contrasts covering the same scanned volume may improve the accuracy of estimation, as the impact of true anatomic features will be less likely to influence erroneous estimations of the bias field. The simultaneous estimation of a single bias field from multiple images during classification may also be more computationally efficient than individual estimation on each separate image followed by classification.

To decouple the bias field from true image voxel intensities, a log transform is usually applied to convert the multiplicative interaction with the image to an additive one [45] [47] [48]. However, log transforms are nonlinear and may degrade the original contrast, resulting in inaccurate classification especially in poor contrast areas. On the other hand, direct division to recover true signal is unstable due to the presence of noise. Therefore, we choose to adjust the centroid of each prototype class for each voxel k according to the estimated bias field in a method similar to a previously published technique [46]. The resulting modified fuzzy c-means formulation for joint classification and bias field correction can be expressed as

$$
\hat{\mathbf{u}}, \hat{\mathbf{v}}, \hat{\mathbf{b}} = \arg\min \sum_{s=1}^{d} \sum_{i=1}^{c} \sum_{k=1}^{N} u_{ik}^{m} \|x_{ks} - b_k v_{is}\|^2 + \alpha \sum_{s=1}^{d} \sum_{i=1}^{c} \sum_{k=1}^{N} u_{ik}^{m} \|\bar{x}_{ks} - b_k v_{is}\|^2
$$
$$
s.t. \sum_{i=1}^{c} u_{ik} = 1, \quad \forall k = 1, \cdots, N. \tag{3.3}
$$

3.2.2.2 Regularization scheme

The problem formulated above is an ill-posed one. The bias field computed directly from Eq(3.3) will be a residual image that compensates for classification error, which violates the prior knowledge that the bias field is slowly varying. To encourage smoothness, we introduce a quadratic penalty term into the objective function:

$$\hat{\mathbf{u}}, \hat{\mathbf{v}}, \hat{\mathbf{b}} = \arg\min \sum_{s=1}^{d} \sum_{i=1}^{c} \sum_{k=1}^{N} u_{ik}^m \|x_{ks} - b_k v_{is}\|^2 + \alpha \sum_{s=1}^{d} \sum_{i=1}^{c} \sum_{k=1}^{N} u_{ik}^m \|\bar{x}_{ks} - b_k v_{is}\|^2 + \lambda\|\mathbf{R}\mathbf{b}\|_2^2$$

$$s.t. \sum_{i=1}^{c} u_{ik} = 1, \quad \forall k = 1, \cdots, N, \tag{3.4}$$

where $\mathbf{R}$ is the finite difference matrix that calculates the difference between neighboring voxels along the axial, sagittal and coronal directions, $\mathbf{b} = [b_1, b_2, \cdots, b_N]^T$ is the column vector that organizes voxel values of the bias field lexicographically.

3.2.2.3 Minimization with respect to u and b

The objective function in Eq(3.4) can be minimized by updating $\mathbf{u}$, $\mathbf{b}$ and $\mathbf{v}$ iteratively. In our investigation, the intensities of all MR images are normalized before classification and the centroid of each class is fixed with a set of pre-determined values, thus we only need to minimize with respect to $\mathbf{u}$ and $\mathbf{b}$.

The update equation for u_{ik} is

$$u_{ik} = \frac{\sum_{s=1}^{d} \left(\|x_{ks} - b_k v_{is}\|^2 + \alpha\|\bar{x}_{ks} - b_k v_{is}\|^2\right)^{-1/(m-1)}}{\sum_{s=1}^{d} \sum_{j=1}^{c} \left(\|x_{ks} - b_k v_{js}\|^2 + \alpha\|\bar{x}_{ks} - b_k v_{js}\|^2\right)^{-1/(m-1)}}, \tag{3.5}$$

to minimize the objective function with respect to $\mathbf{b}$, we rearrange terms in Eq(3.4) and write the regularized optimization problem in the matrix form:

$$\hat{\mathbf{b}} = \arg\min_{\mathbf{b}} \|\mathbf{z} - \mathbf{D}\mathbf{b}\|_{\mathbf{W}}^2 + \lambda\|\mathbf{R}\mathbf{b}\|_2^2, \tag{3.6}$$

where $\mathbf{D} = \text{diag}\{d_k\}$ with $d_k = \sqrt{(1+\alpha)\sum_{s=1}^{d}\sum_{i=1}^{c} u_{ik}^m v_{is}^2}$. $\mathbf{z} = [z_1, z_2, \cdots, z_N]^T$, with $z_k = \frac{\sum_{s=1}^{d}\sum_{i=1}^{c} u_{ik}^m v_{is}(x_{ks}+\alpha\bar{x}_{ks})}{d_k}$, $\mathbf{W} = \text{diag}\{w_k\}$ with $w_k \in \{0, 1\}$ to exclude background voxels.

Eq(3.6) has a closed-form solution $\hat{\mathbf{b}} = [\mathbf{D}^T\mathbf{W}\mathbf{D}+\lambda\mathbf{R}^T\mathbf{R}]^{-1}\mathbf{D}^T\mathbf{W}\mathbf{z}$. However, due to the large size of medical data, it is not practical to compute this solution. Allison *et al* [49] proposed to solve this problem using Augmented Lagrangian methods. Firstly, the finite difference matrix $\mathbf{R}$ was factorized into a circulant matrix $\mathbf{C}$ and a binary diagonal matrix $\mathbf{B}$ such that $\mathbf{R} = \mathbf{B}\mathbf{C}$. Since $\mathbf{C}$ is a finite difference matrix with periodic boundary conditions, it will penalize the difference between voxels on opposing boundaries. A binary mask $\mathbf{B}$ is needed to eliminate this effect. For an illustrative picture of $\mathbf{R},\mathbf{C}$ and $\mathbf{B}$ for 1D case, one can refer to the paper of the augmented Lagrangian method [49]. A variable splitting scheme is next introduced to convert the unconstrained problem into an equivalent constrained problem

$$\hat{\mathbf{b}}, \hat{\mathbf{u}} = \arg\min_{\mathbf{b},\mathbf{u}} \|\mathbf{h} - \mathbf{A}\mathbf{u}\|_2^2 \quad s.t. \mathbf{u} = \mathbf{T}\mathbf{b}, \tag{3.7}$$

where

$$\mathbf{u} = \begin{bmatrix} \mathbf{u}_1 \\ \mathbf{u}_0 \end{bmatrix}, \quad \mathbf{T} = \begin{bmatrix} \mathbf{I} \\ \mathbf{C} \end{bmatrix}, \quad \mathbf{h} = \begin{bmatrix} \mathbf{W}^{1/2}\mathbf{z} \\ \mathbf{0} \end{bmatrix} \quad \text{and} \quad \mathbf{A} = \begin{bmatrix} \mathbf{W}^{1/2}\mathbf{D} & \mathbf{0} \\ \mathbf{0} & \sqrt{\lambda}\mathbf{B} \end{bmatrix}. \tag{3.8}$$

An alternating direction method of multipliers (ADMM) algorithm was designed to solve the constrained optimization problem in Eq(3.7) as shown in algorithm 1.

Both $\mathbf{D}_2$ and $\mathbf{B}_2$ in algorithm 1 are diagonal matrices, and thus the most computationally expensive part is computing $\mathbf{C}\mathbf{b}^{(j+1)}$ and $\mathbf{Q}^H\mathbf{\Phi}_2^{-1}\mathbf{Q}$, which has a computational complexity of $O(N \log N)$ due to the circular convolution and fast Fourier transform (FFT).

To further accelerate the algorithm, we used a linearized Augmented Largrangian method as proposed by Nien *et al* [50]. To update $\mathbf{b}$, the algorithm majorizes the quadratic penalty term by its separable quadratic surrogate. Instead of updating $\mathbf{b}$ as

$$\mathbf{b}^{(j+1)} = \arg\min_{\mathbf{b}} \frac{1}{2}\|\mathbf{u}^{(j)} - \mathbf{T}\mathbf{b} - \boldsymbol{\eta}^{(j)}\|^2, \tag{3.9}$$

Algorithm 1 ADMM

Initialize $\mathbf{u}_1^{(0)} = \mathbf{b}^{(0)}, \mathbf{u}_0^{(0)} = \mathbf{Cb}^{(0)}, \boldsymbol{\eta}_0^{(0)} = \mathbf{0}, \boldsymbol{\eta}_1^0 = \mathbf{0}$ and $j = 0$

Set $\mathbf{D}_2^{-1} = [\mathbf{D^H WD + I}]^{-1}$, $\mathbf{z}_2 = \mathbf{D^H Wz}$, $\mathbf{B}_2^{-1} = [\lambda \mathbf{B}^H \mathbf{B} + \mathbf{I}]^{-1}$
Set $\boldsymbol{\Phi}_2^{-1} = [\mathbf{I} + \boldsymbol{\Phi}]^{-1}$, where $\boldsymbol{\Phi}$ is the diagonal matrix containing the spectrum of $\mathbf{C^H C}$

Repeat until stop criterion is achieved:
  1) $\mathbf{b}^{(j+1)} = \mathbf{Q}^H \boldsymbol{\Phi}_2^{-1} \mathbf{Q} \times [\mathbf{C^H}(\mathbf{u}_0^{(j)} - \boldsymbol{\eta}_0^{(j)}) + (\mathbf{u}_1^j - \boldsymbol{\eta}_1^{(j)})]$
    where $\mathbf{Q}$ is the multidimensional discrete Fourier transform such that $\mathbf{C^H C} = \mathbf{Q}^H \boldsymbol{\Phi} \mathbf{Q}$

  2) $\mathbf{u}_1^{(j+1)} = \mathbf{D}_2^{-1}[\mathbf{D}^H \mathbf{Wz} + (\mathbf{b}^{(j+1)} + \boldsymbol{\eta}_1^{(j)})]$
    $\mathbf{u}_0^{(j+1)} = \mathbf{B}_2^{-1}(\mathbf{Cb}^{(j+1)} + \boldsymbol{\eta}_0^{(j)})$

  3) $\boldsymbol{\eta}_1^{(j+1)} = \boldsymbol{\eta}_1^{(j)} - (\mathbf{u}_1^{(j+1)} - \mathbf{b}^{(j+1)})$
    $\boldsymbol{\eta}_0^{(j+1)} = \boldsymbol{\eta}_0^{(j)} - (\mathbf{u}_0^{(j+1)} - \mathbf{Cb}^{(j+1)})$

  4) $j = j + 1$

---

we calculate

$$\mathbf{b}^{(j+1)} = \arg\min_{\mathbf{b}} \frac{1}{2}\|\mathbf{u}^{(j)} - \mathbf{Tb} - \boldsymbol{\eta}^{(j)}\|^2 + \langle \mathbf{C}^H(\mathbf{Cb} - (\mathbf{u}_0^{(j)} - \boldsymbol{\eta}_0^{(j)})), \mathbf{b} - \mathbf{b}^{(j)}\rangle + \frac{L}{2}\|\mathbf{b} - \mathbf{b}^{(j)}\|^2,$$

(3.10)

where L is the maximum eigenvalue of $\mathbf{C^H C}$. The problem in Eq(3.10) has a simple

closed-form solution

$$\mathbf{b}^{(j+1)} = \frac{1}{1+L}(\mathbf{u}_1^{(j)} - \boldsymbol{\eta}_1^{(j)}) + \frac{L}{1+L}\left(\mathbf{b}^{(j)} - (1/L)\mathbf{C}^H\left(\mathbf{Cb}^{(j)} - (\mathbf{u}_0^{(j)} - \boldsymbol{\eta}_0^{(j)})\right)\right). \quad (3.11)$$

Furthermore, since we do not need to calculate the inverse of $\mathbf{C^H C}$ using fast Fourier

transform, $\mathbf{C^H C}$ does not have to be block circulant. Thus $\mathbf{C}$ can be a simple linear

convolution matrix with a small convolution kernel ($[-1, 2, -1]$) and without periodic

boundary conditions. Therefore we can calculate $\mathbf{Cb}^{(j)}$ with a linear computational

complexity. Besides, for further computational efficiency gains, one can estimate $\mathbf{b}$

from down sampled image data and extrapolate to full scale, while still effectively

removing bias field effects, a technique similar to that used in the N3 and N4itk

methods [43] [44].

3.2.2.4 Evaluation and Validation

We evaluated the proposed algorithm on the 6 multicontrast MRI data sets. First, we preprocessd the MR images using methods similar to those previously reported [4], the skin surface was determined as the union of contours extracted from T1-weighted, T2-weighted, fat, water and PETRA images with intensity thresholds of $20, 10, 10, 30$ and 100 respectively. The extracted external skin surface was further processed using morphologic operations (dilation and erosion) to clean up gaps and holes. Vessel masks were generated from TOF images with an intensity threshold of 60. A normalization factor was calculated as the averaged mean intensity of T1-weighted and T2-weighted images within the skin surface. T1-weighted, T2-weighted, fat and water images were normalized by multiplying voxel values by 1000 and then dividing by this normalization factor. PETRA images were normalized independently to have a mean voxel intensity of 1000. After normalization, air masks within the skin surface were generated using the normalized PETRA and water images with intensity thresholds of 200 and 400. The voxels outside the skin surface, as well as those encompassed by the vessel and air masks, were excluded from classification.

After the above preprocessing steps, the probabilities of remaining voxels containing each of five classes of tissue/material (fat, fluid, white matter, grey matter and bone) were determined using the proposed method, with $m = 1.5, \alpha = 3.8$ and $\lambda = 2^5$. The pre-determined centroids of each class are shown in Table 3.1. By using the pre-determined centroid as the solution for $\mathbf{v}$, the initial estimates of membership $\mathbf{u}$ are reasonably close to the optimum. Therefore instead of updating the membership and the estimated bias field iteratively, the bias field was estimated in a single shot, and the final classification result was calculated using Eq(3.5). For estimation of the bias field, image data ($\mathbf{z}$ and $\mathbf{D}$) were first down sampled from $256 \times 256 \times 176$ to

Table 3.1: Centroid of each tissue class and image contrast used for classification

	T1-weighted	T2-weighted	Fat	Water
White matter	1488	783	143	1339
Grey matter	1161	1418	118	1051
Fluid	714	2784	113	631
Fat	2385	936	2132	733
Bone	589	260	301	424

$64 \times 64 \times 32$. The 5 classes (fat, fluid, white matter, grey matter and bone) were assigned CT attenuation values of $-50, 0, 40, 40$ and 800 HU respectively. Synthetic CT image volumes were generated by multiplying the probability of each voxel being represented in each class with the attenuation value of that class and summing the result across all classes. The synthetic CT images generated using the proposed method were compared to synthetic CT images generated using the original fuzzy c-means method under the same parameter settings. The mean absolute error between the synthetic CT and the reference CT image volumes within brain region was calculated. The percentage of soft tissue voxels misclassified as bone within the brain region was also calculated, where soft tissue voxels were determined as voxels with intensity between 0 and 50 HU from the reference CT images, and misclassification was defined as any brain tissue assigned a composite intensity above 50 HU. Finally, the normalized bone probability, defined as the sum total of bone class membership probability of all voxels within the brain divided by the total number of brain voxels, was calculated.

3.3   Results

Figure 3.3 shows example synthetic CT images generated using the proposed method and the original fuzzy c-means classification. The synthetic CT volumes generated from the MR volumes without bias field correction show artifacts due to the misclassification of some soft tissue as bone. Figure 3.4 shows example probability
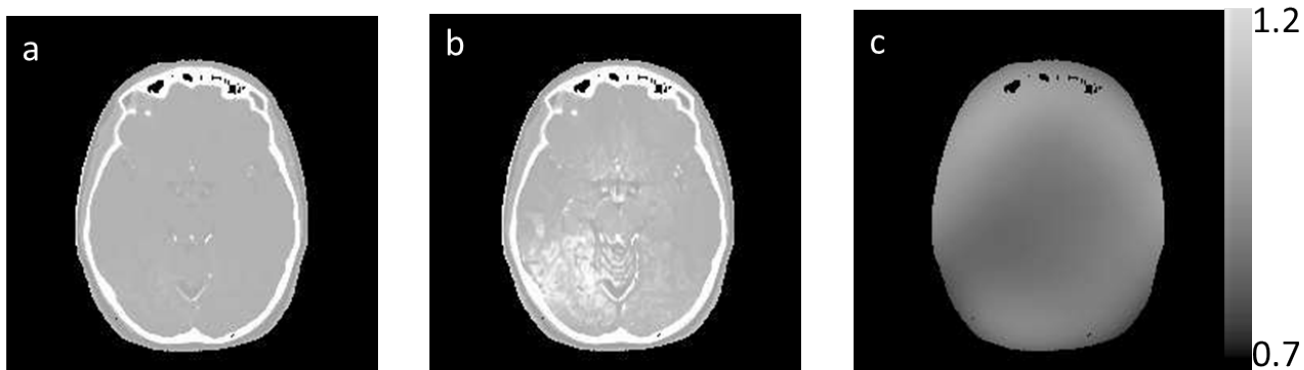
Figure 3.3: (a)Synthetic CT generated using the proposed modified fuzzy c-means classification method with integrated bias field correction. (b) Synthetic CT generated using the original fuzzy c-means classification method. (c) Bias field estimated using the proposed method (holes are due to air masked out prior to processing).
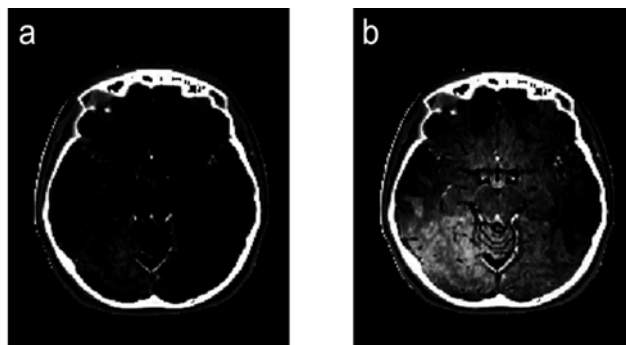
Figure 3.4: (a) Bone probability map generated using the proposed modified fuzz c-means classification method with integrated bias field correction, (b) Bone probability map generated using the original fuzzy c-means classification method.

distribution of bone voxels using the two methods, where the misclassification using the original fuzzy c-means method was obvious. Due to the bias field, intensities of soft tissue voxels were decreased in some regions to the point where they overlapped with bone voxel intensities. These artifacts were significantly reduced using the proposed method.

The mean absolute error (MAE) of intensity in the brain between synthetic CT images generated with and without bias field correction and the reference CT images for each patient is shown in Table 3.2. The percentage of soft tissue voxels misclassi-

29

Table 3.2: Quantitative analysis on MRCT images generated with and without bias field correction

Patient ID	MAE without bias correction	MAE with bias correction	Misclassification rate without bias correction	Misclassification rate with bias correction	Normalized summed bone probability without bias correction	Normalized summed bone probability with bias correction
1	32	22	24%	9%	3%	2%
2	80	25	74%	12%	9%	2%
3	90	48	59%	13%	14%	5%
4	51	16	38%	7%	6%	1%
5	80	25	63%	14%	15%	2%
6	89	25	53%	14%	9%	2%

fied as bone is also shown. With the proposed method, the averaged mean absolute error among 6 patients was reduced by 43.5 HU and the percentage of misclassified brain voxels was reduced by 40%. Table 3.2 also shows the summed bone probability map with in the soft tissue region, normalized by the total number of soft tissue voxels. The proposed method decreased the normalized bone probability within soft tissue by 7%.

3.4   Discussion

In this work, we investigated an integrated method for joint tissue classification and bias field correction of multiple contrast MR images to support generation of synthetic CT. The proposed method improves the quality of synthetic CT volumes generated by reducing artifacts resulting from misclassification. Evaluation in terms of classification accuracy and quality of the synthetic CT images validates the efficacy of the proposed method. It should be noted that the misclassification constraint was deliberately set to be very strict, hence the large number of misclassified voxels both with and without bias correction. In addition, some vessels were still not covered by the mask from the TOF images, leading to a constant number of voxels being misclassified regardless of whether bias field correction was applied or not.

Joint estimation takes advantage of multispectral images of different contrasts to reduce error due to low signal regions present in one or more but not all images studied. With intensity normalization and use of pre-determined cluster centroids, the proposed method is able to remove bias field effects in a single shot, while most existing methods require an iterative process. In addition, the quadratic penalty term enforces a natural constraint on smoothness of the bias field, without making underlying assumptions to model the bias field shape, such as a B-spline basis or Gaussian kernel [43] [44] [47].

The linearized Augmented Largrangian method used to calculate the bias field was run for 500 iterations to give satisfactory results. Each iteration requires linear convolution and simple addition. Thus the computational complexity for each iteration is simply $O(N)$. A commonly used bias field estimation algorithm N4itk [44], which is publicly available in a medical image analysis software environment (SLICER, Surgical Planning Laboratory, Brigham and Womens Hospital, Boston, MA), has default values for iterations at each of three resolution levels of 500, 400 and 300. The B-spline interpolation method used to smooth the estimated bias field in N4itk has a computational complexity of $O(N + n)$ [51], where n is the total number of control points used for B-spline interpolation. Our proposed method does not need to calculate and sharpen image histograms by deconvolution and produces classification results simultaneously, thus reducing the computational effort in generating synthetic CT volumes.

For future work, more thorough evaluation of the proposed method and its impact on contouring, treatment planning, and dose calculation will be performed, including a more extensive study on a larger patient dataset.

3.5   Conclusion

A joint bias field correction and tissue classification method has been developed and evaluated. Correcting bias field of MRI data reduces the error in tissue classification thus improving the accuracy of synthetic CT generation. By performing intensity normalization and applying a linearized Augmented Largrangian method, the computational efficiency was improved over existing bias field correction algorithms.

CHAPTER 4

Pelvic Bone Shape Model in Support of Pelvic Synthetic CT Generation

4.1 Introduction

[1]While intensity thresholding of UTE image data to mask air out in support of synthetic CT generation has been evaluated and validated for attenuation mapping in the head [4], this process is challenging in the pelvis due to the mobility of air over short time periods. Intra-session changes may lead to air regions estimated from UTE imaging sequence not spatially matching the locations of air in other MR volumes from the same scanning session, as shown in Figure 4.1. This potential spatial mismatch could lead to misclassification of air as bone in the pelvis. Several methods [52] [53] [54] [55] manually contoured the bony part in pelvis before synthetic CT generation, which can be time-consuming and non-repeatable.

One approach to address this issue is to add bone shape information to assist the air/bone separation process. As air is rarely in close contact with bone in the pelvis, defining the space in which bone exists would allow for identifying air in remaining low MR intensity regions outside of bone for exclusion from tissue classification, and thus also remove the need for UTE imaging. Achieving this goal requires a model that describes possible shape variance of the pelvic bone. Various methods

---

[1]This chapter is based on publications [8] [9]
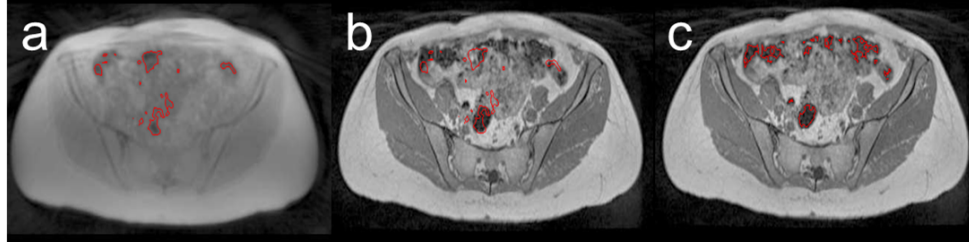
Figure 4.1: (a) The air regions (contoured in red) estimated from PETRA image overlaid on (b) T1-weighted image of the same patient from the same scan session. The air regions estimated from PETRA do not agree with the (c) air regions (contoured in red) presented on T1-weighted image.

have been proposed for shape model construction, including 1) active shape models where shape variability is represented by the variance of coordinates of landmarks placed on the object [56] [57]; 2) active contour models where shape variability is represented by the variance of object contours [58] [59]; 3) statistical deformation models where shape variability is represented by a series of deformation fields that deform a reference image to match a group of training images [60] [61] [62]. As compared to active shape models and active contour models, statistical deformation models do not require explicit contouring of objects, and are thus less expensive in terms of effort for model construction. In this study, a pelvic bone shape model is developed and investigated based upon deformable alignment of pelvic CT image volumes across a female population to assist bone identification in MRI as the first step in the synthetic CT generation. As actual attenuation values will be determined in a subsequent probabilistic tissue classification step that permits voxels to have combinations of tissues, the goal here is to find a bulk space where air is unlikely to exist instead of contouring bone tissues accurately. This reduced demand for specificity in bone location labeling allows a trade-off between model accuracy versus model complexity. As compared to deformable registration based methods [36], the shape model has fewer parameters to be optimized and the search space appears to be free of local minima, thus is more computationally efficient and robust.

34
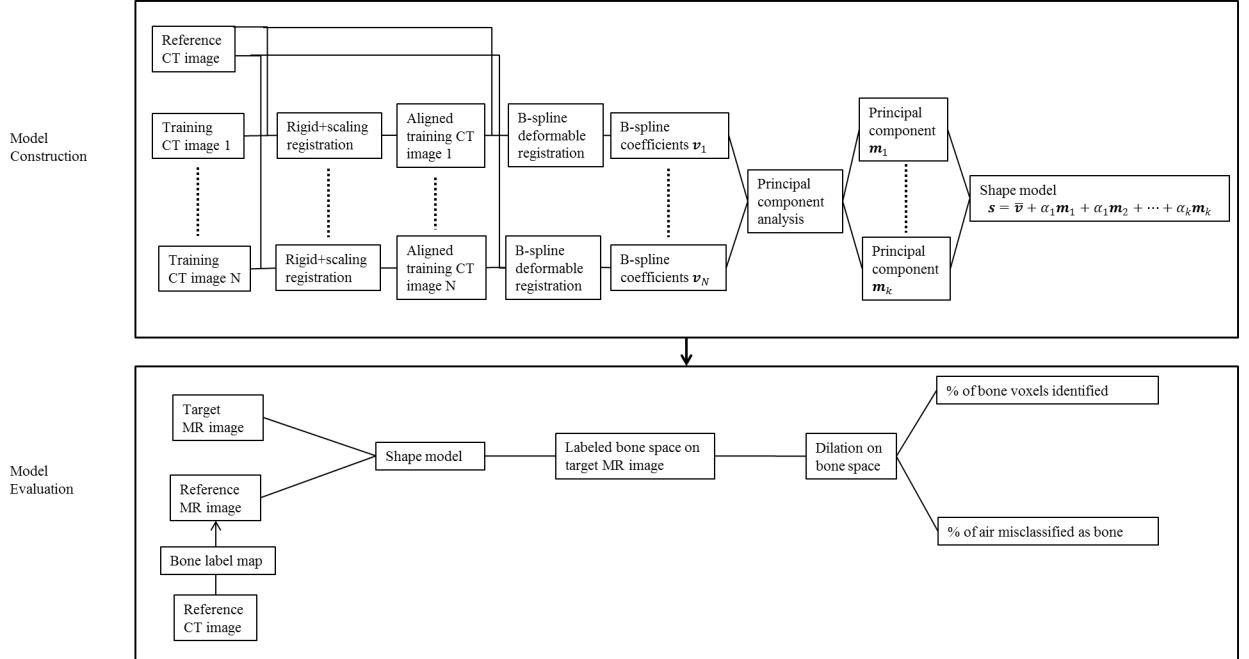
Figure 4.2: Flow chart showing the model construction and evaluation processes.

4.2   Method and Materials

We built the model from CT images where bone structure can be easily defined, and applied it to MR images for model evaluation. To exclude individual differences that are not due to shape variance, all images were first (rigidly plus scale) aligned into the same coordinate system. The reference CT was deformably aligned to the resulting images, and Principal Component Analysis (PCA) of the population of deformations was used to assemble the shape model. The shape model was then used to guide the deformation of the reference MR image with labeled bone regions to match the target MR image. Dilation was further introduced to improve the coverage of bone voxels. Figure 4.2 shows a flow chart of the model construction and application/evaluation processes.

4.2.1 Image data for shape model construction

Under institution review board approval, CT image volumes acquired from 30 female patients who underwent simulation for external beam radiotherapy were investigated. Of these patients, 17 also had corresponding MR scans. CT scans were acquired with slice thickness of 3 mm and in-plane pixel size of 1×1 mm^2. MR scans consisted of multi-contrast MRI images including T1-weighted images generated using a 3D gradient echo sequence (VIBE Dixon, initially designed for breath-hold liver imaging) with TE (in/out-of-phase) equals 2.46/1.23 ms, TR equals 4.1 ms, flip angle 90°, and voxel size of $1.5\times1.5\times1.5$ mm^3. Fat and water images were computed from the in-phase (T1-weighted) and out-of-phase VIBE images. All MR and CT images were acquired using a 3T MR scanner (Skyra, Siemens Medical Systems, Erlangen, Germany) and CT simulator (Philips Healthcare, Cleveland, OH), respectively. All scans were acquired with patients lightly immobilized in vacuum-formed bean bags, and with no other controls on physiological movement. Intensity inhomogeneity correction was applied to MR images using the N4ITK algorithm [44], implemented in a publicly available image analysis software environment (SLICER, surgical processing laboratory, Brigham and Womens Hospital, Boston, MA). All MR images were aligned to their corresponding CT images through rigid registration implemented in SLICER, with the aim of aligning bony structures accurately without considering soft tissues, as soft tissues can deform across scanning sessions. All image volumes were reformatted in axial orientation with voxel sizes interpolated to $1\times1\times1$ mm^3. Example images used from one subject are shown in Figure 4.3. Variations in subject body composition as well as internal motion during scans were present in the study population. Figure 4.4 shows one sample image from the dataset that shows motion artifacts as well as one sample image that has relatively high quality.
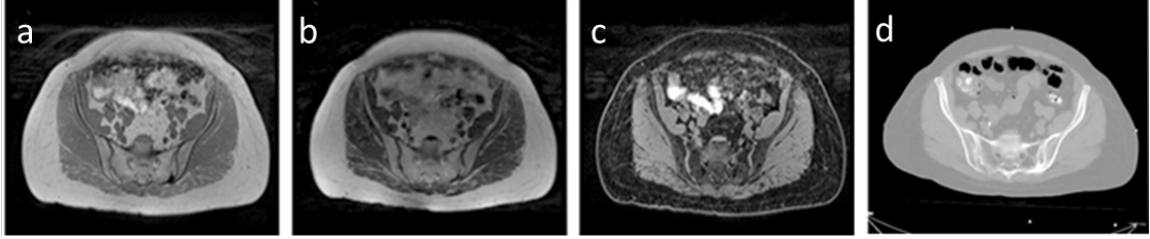
Figure 4.3: Example MR and CT images from one study subject. (a) T1-weighted (in-phase), (b) Fat, (c) Water and (d) CT image.
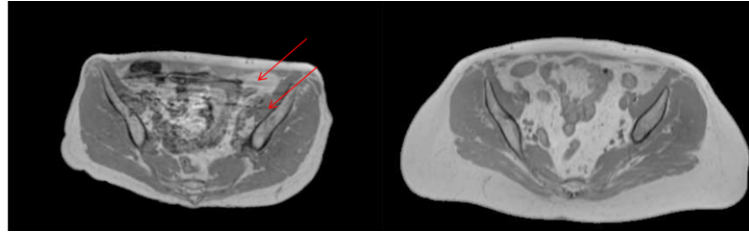
Figure 4.4: Example T1-weighted MR images used for this study with different image qualities. Left: image that shows motion artifacts (lines indicated by red arrows). Right: image of relatively high quality.

4.2.2 Model construction

4.2.2.1 Alignment of training images

Since we were interested in bone identification only, irrelevant structures were first excluded to reduce alignment complexity. Bony surfaces were identified on CT images using an intensity threshold of 150 Hounsfield units (HU). These surfaces were dilated by 3mm to include contrast with adjacent surrounding soft tissues. Using an in-house Functional Imaging Analysis Tool (FIAT), morphological operations (3 mm dilation, filling holes and 3 mm erosion) were subsequently applied to clean up holes and gaps. Non-pelvic bones, such as vertebrae and femurs, were manually removed. Example CT images before and after preprocessing are shown in Figure 4.5.

After the above preprocessing steps, one CT image volume was selected randomly from the dataset as the reference. To exclude individual differences that were not due to shape variance, the remaining training images in the dataset were first aligned to
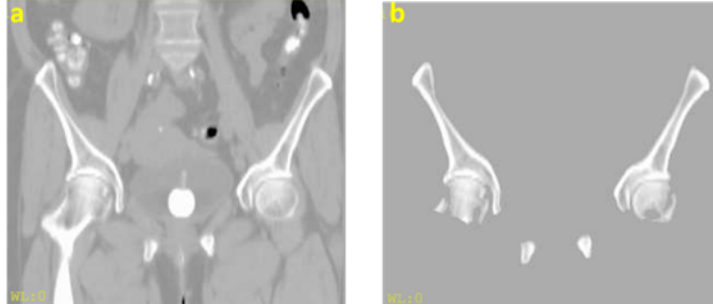
Figure 4.5: CT image (a) before and (b) after preprocessing.
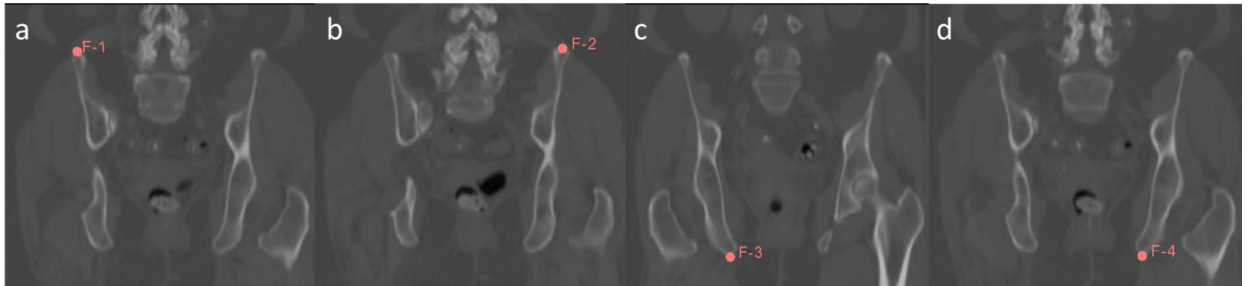
Figure 4.6: Landmarks placed on CT image volumes to assist alignment.

the reference image using rigid registration plus global scaling. Then the reference image was deformed to match each of the training images, using a multilevel B-spline deformable registration with subsampling rates of [10 10 5], [4 4 2] and [2 2 1] voxels and B-spline grid sizes of 50, 30 and 15 mm respectively. The regularization parameter for smoothness was 0.005 and the maximal number of iterations was 50 for each level. The mean square error between the target and reference image was chosen as the similarity metric for both rigid and B-spline deformable registration processes. Four anatomic landmarks were identified at the top and the bottom and the left and the right of the pelvis, as shown in Figure 4.6, to initialize the rigid plus scaling alignment. All alignments were performed using SLICER. Figure 4.7 shows the reference image as well as an example deformation field that warps the reference image to match a training image.
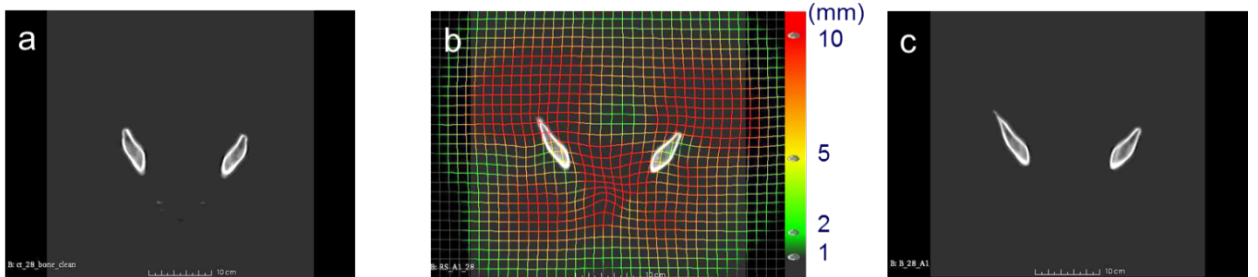
Figure 4.7: (a) Reference image. (b) Target image with deformation vectors. (c) Deformed reference image.

4.2.2.2 Principal component analysis of B-spline coefficients

To extract representative modes of shape variance, PCA is usually performed on a matrix composed of deformation fields obtained from the non-rigid alignment process. The deformation fields in our method were calculated from B-spline coefficients at control points through simple interpolation, and thus the high dimensional deformation fields were mapped from the low dimensional B-spline coefficients through a linear transformation matrix. Therefore the information encoded in deformation vectors and B-spline coefficients is equivalent. PCA was performed on the B-spline coefficients instead of the deformation vectors to yield more compact descriptions of shape variance and save computational effort by reducing the matrix size.

The B-spline coefficients of the i^{th} deformation field were organized into a vector $\mathbf{v}_i$ and PCA was performed on the matrix

$$\mathbf{V} = [\mathbf{v}_1, \mathbf{v}_2, \cdots, \mathbf{v}_N]. \tag{4.1}$$

PCA of $\mathbf{V}$ yields N orthogonal principal modes $\mathbf{m}_1, \mathbf{m}_2, \cdots, \mathbf{m}_N$. This gives the shape model where the possible shape variance $\mathbf{s}$ of pelvic bone is described by a linear combination of the leading k principal modes plus the mean B-spline coefficients $\bar{\mathbf{v}}$

$$\mathbf{s} = \bar{\mathbf{v}} + \alpha_1 \mathbf{m}_1 + \alpha_2 \mathbf{m}_2 + \cdots + \alpha_k \mathbf{m}_k. \tag{4.2}$$

As the deformation fields are obtained by aligning real pelvic bone structures, no further constraints are needed during PCA to ensure the principals modes represent anatomically reasonable shape variation. This is validated by deforming the reference image with deformation fields generated by multiplying principal modes with different mode coefficients. The first principal mode is found to be mainly changing the angle of the pelvis.

4.2.3 Application of model to identify pelvic bone in MR images

4.2.3.1 Deformable alignment guided by shape model

To identify pelvic bone on MR images, we deformed the corresponding MR data of a reference image $\mathbf{A}$ to match a target image $\mathbf{M}$. The bone regions on the reference MR image were identified and labeled according to the corresponding CT image. The deformation process was guided by the shape model developed in section 4.2.2 by constraining the search space to be within the space spanned by principal modes. The problem can be formulated as

$$\hat{\boldsymbol{\alpha}} = \arg\min_{\boldsymbol{\alpha}} D\big(\mathbf{T}(\mathbf{A}), \mathbf{M}\big)$$

$$s.t.\,\mathbf{T} = \mathbf{B}(\mathbf{v}), \mathbf{v} = \bar{\mathbf{v}} + \alpha_1 \mathbf{m}_1 + \alpha_2 \mathbf{m}_2 + \cdots + \alpha_k \mathbf{m}_k, \tag{4.3}$$

where D is the similarity metric between the target and reference images and $\mathbf{T}$ is the deformation field that is calculated from B-spline coefficients $\mathbf{v}$ at control points using the cubic B-spline interpolator $\mathbf{B}$. By using the subspace model Eq(4.2), we only need to optimize over the k model coefficients $(\alpha_1, \cdots, \alpha_k)$. As compared to deformable registration, our model has far fewer parameters involved in the optimization process.

In our approach, we chose the mean square intensity error between the reference and the target water images as the similarity metric D. The images were first normalized based on the mean intensity of the T1-weighted images from the same

acquisition that yielded the water images. Since we were interested in pelvic bone only, the value of the similarity metric was evaluated only in the regions where the deformed reference image indicated the existence of bone (i.e., fell within the labeled bone space on the reference image). The optimization problem was thus

$$\hat{\boldsymbol{\alpha}} = \underset{\boldsymbol{\alpha}_{\text{lower}} \leq \boldsymbol{\alpha} \leq \boldsymbol{\alpha}_{\text{upper}}}{\arg\min} \sum_{\mathbf{T}(\mathbf{A}(x)) \in \text{bone}} \Big(\mathbf{T}(\mathbf{A}(x)) - \mathbf{M}(x)\Big)^2$$

$$s.t. \mathbf{T} = \mathbf{B}(\mathbf{v}), \mathbf{v} = \bar{\mathbf{v}} + \alpha_1 \mathbf{m}_1 + \alpha_2 \mathbf{m}_2 + \cdots + \alpha_k \mathbf{m}_k, \tag{4.4}$$

with the binary operator introduced by the masked bone space, the optimization problem in Eq(4.4) is not differentiable. Given the small number of parameters, we may simply traverse a grid of parameter values to find the optimal solution. In this work, we determine the search range $\boldsymbol{\alpha}_{\text{lower}}$ and $\boldsymbol{\alpha}_{\text{upper}}$ based on the square-root of the sum of eigenvalues from PCA of B-spline coefficients, denoted as λ. The search grid is defined as -100λ to 100λ, with a grid spacing of 100. Furthermore, if we assume the uniqueness of minimum of the search space, a binary search scheme may be used to find the optimal coefficients efficiently. Algorithm 2 shows an illustrative example of using the binary search scheme to find the optimal coefficient, where the shape model consists of the first leading principal mode only.

4.2.4 Evaluation and validation

We evaluated the shape model by calculating the percentage of correctly identified bone voxels as well as the percentage of air being misclassified as bone, where air voxels were defined on corresponding CT images at a threshold of -400 HU. For each of the 17 patients that had both MR and CT scans, we constructed the shape model from the remaining 29 CT scans and applied the model to identify pelvic bone on the MR image of the patient whose CT was not used in the model development.

Algorithm 2 Binary Search

Initialize $\boldsymbol{\alpha} = [\boldsymbol{\alpha}(1), \boldsymbol{\alpha}(2)]$, convergence threshold ϵ, maximum number of iteration N.
Set $j = 1$. Set mean square error D_0 and difference between iterations d to be sufficiently large.

while $d > \epsilon$ and $j < N$
$\quad \mathbf{v}_1 = \bar{\mathbf{v}} + \boldsymbol{\alpha}(1)\mathbf{m}_1, \; \mathbf{v}_2 = \bar{\mathbf{v}} + \boldsymbol{\alpha}(2)\mathbf{m}_1, \; \mathbf{T}_1 = \mathbf{B}(\mathbf{v}_1), \; \mathbf{T}_2 = \mathbf{B}(\mathbf{v}_2)$

$\quad$ **if** $D(\mathbf{T}_1(\mathbf{A}), \mathbf{M}) > D(\mathbf{T}_2(\mathbf{A}), \mathbf{M})$
$\quad\quad \boldsymbol{\alpha} = [(\boldsymbol{\alpha}(1) + \boldsymbol{\alpha}(2))/2, \boldsymbol{\alpha}(2)]$
$\quad$ **else**
$\quad\quad \boldsymbol{\alpha} = [\boldsymbol{\alpha}(1), (\boldsymbol{\alpha}(1) + \boldsymbol{\alpha}(2))/2]$
$\quad$ **end**

$\quad d = D_0 - \min\left(D(\mathbf{T}_1(\mathbf{A}), \mathbf{M}), D(\mathbf{T}_2(\mathbf{A}), \mathbf{M})\right)$
$\quad D_0 = \min\left(D(\mathbf{T}_1(\mathbf{A}), \mathbf{M}), D(\mathbf{T}_2(\mathbf{A}), \mathbf{M})\right)$
$\quad j = j + 1$
end

if $D(\mathbf{T}_1(\mathbf{A}), \mathbf{M}) > D(\mathbf{T}_2(\mathbf{A}), \mathbf{M})$
$\quad$ output $\boldsymbol{\alpha}(2)$
else
$\quad$ output $\boldsymbol{\alpha}(1)$
end

To evaluate the complexity of the search space, we calculated the objective function values by brute force sampling of coefficients of the first three principal modes. To validate the choice of similarity metric, the model was fitted to CT data by directly maximizing the percentage of identified bone voxels, as the bone label map was already defined on CT images, and compared to the result when applying the model to MR images. Finally, the effect of dilation on the identified space was studied by evaluating both the percentage of correctly identified bone voxels and the percentage of air voxels mislabeled as bone.

4.3 Results

The B-spline coefficients within pelvic bone area ranged from -41.27 to 49.71, with a mean of -0.24, standard deviation of 8.33 and a mean absolute of 6.35. The percentage of total variance of deformations (B-spline coefficients) across patients explained by each mode is shown in Figure 4.8. The first leading mode accounts for
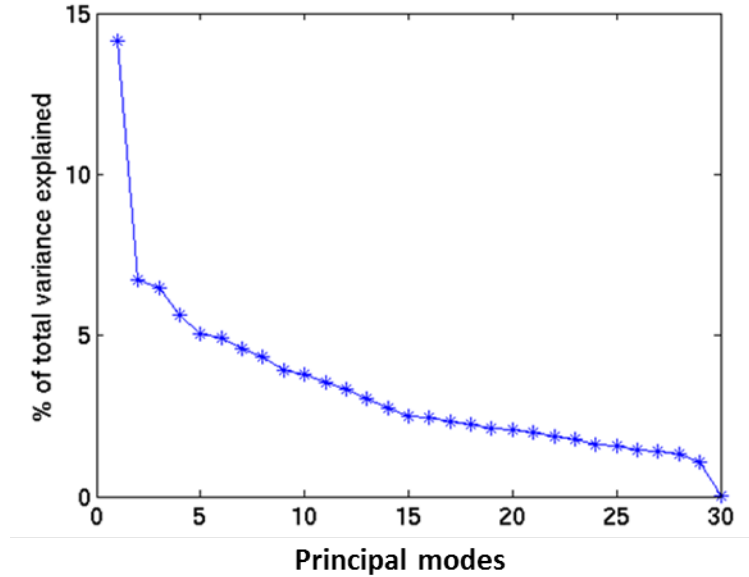
Figure 4.8: Percentage of variance explained by each mode.

14% of the total variance. Figure 4.9 shows the deformation field corresponding to the first principal mode and Figure 4.10 shows the reference image deformed by the first principal mode by varying the coefficients between 5 (red contour) and -15 (green contour) of the standard deviation of the population. We explored the properties of the search space by plotting the objective function values under various model coefficients, as shown in Figure 4.11. The objective function was observed to be free of local minima, thus justifying adoption of a binary search scheme (Algorithm 2) to find the optimal model coefficients more efficiently than brute force or more complex searching schemes.

Using the first leading principal mode only, the leave-one-out cross validation process achieved an averaged bone identification rate of 87% across 17 patients. The lowest coverage was 79% and the highest coverage was 94%. The standard deviation across the 17 patients was 6%. The averaged total volumes of unidentified bone were 118.33 cm^3 the standard deviation of unidentified bone volumes was 70.77 cm^3. The
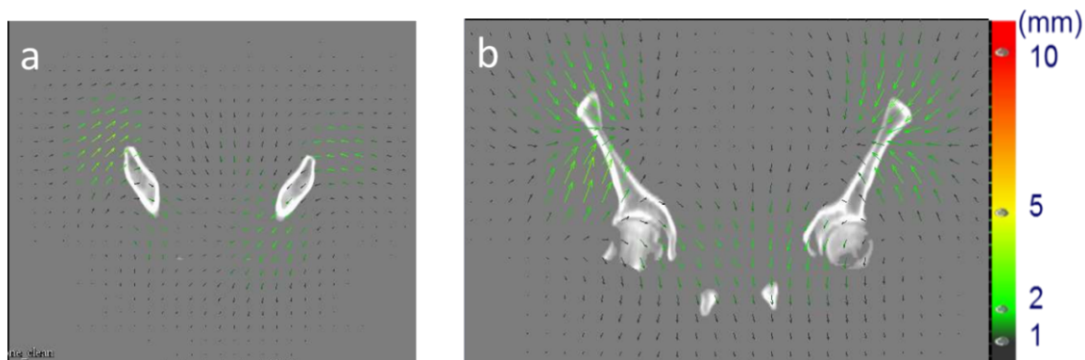
Figure 4.9: Deformation field corresponding to the first principal mode with the reference image in the (a) axial plane and (b) coronal plane.
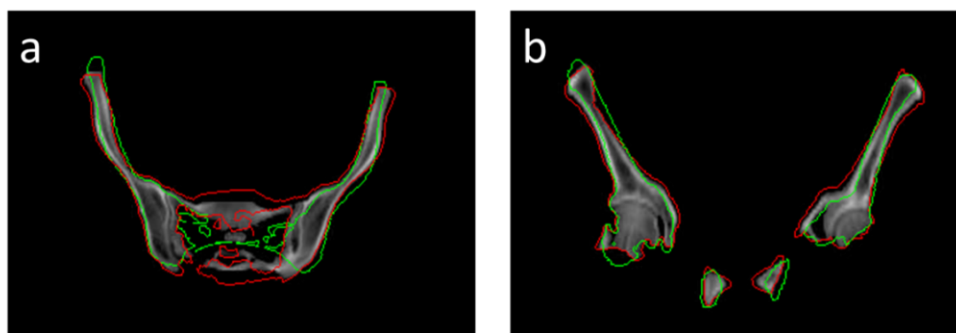
Figure 4.10: Contours of two deformation samples (red and green) using the first principal mode with the respective coefficients of 5 and -15 of the standard deviation of the population, with the reference CT images (grey) in the axial (left) and coronal plane (right).
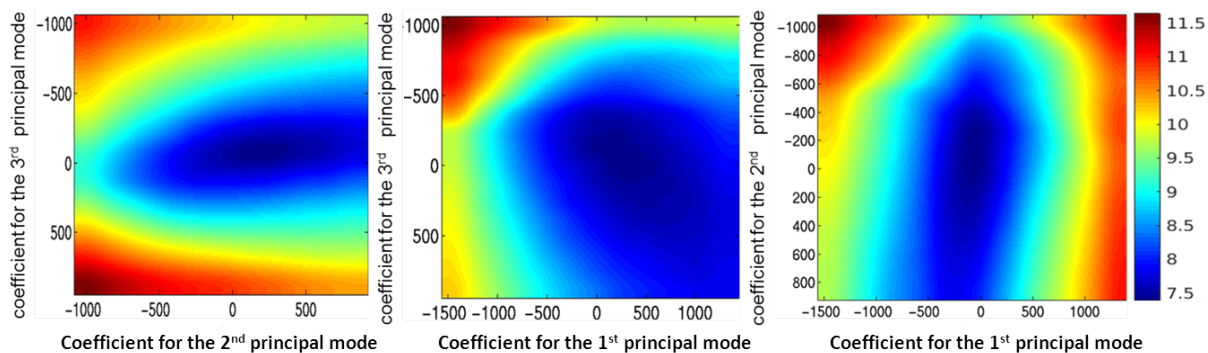
Figure 4.11: Plot of the objective function values with a) the coefficients of the 2^{nd} versus 3^{rd} principal mode at the optimal value for the coefficient of the 1^{st} principal mode, b) the coefficients of the 1^{st} versus 3^{rd} principal mode at the optimal value for the coefficient of the 2^{nd} principal mode, and c) the coefficients of the 1^{st} versus 2^{nd} principal mode at the optimal value for the coefficient of the 3^{rd} principal mode.

averaged percentage of air voxels misclassified as bone was 0.02% and the standard deviation across patients was 0.08%. The averaged total volumes of misclassified air were 0.23 cm^3 and the standard deviation of misclassified air volumes was 0.70 cm^3. For most patients, the largest portions of missing bone voxels were located at the bottom (femoral head) and top of the segmented pelvis. Figure 4.12 shows the averaged portion of missing bone voxels of the total amount of missing bone voxels in each slice (axial plane). The peaks appear around slice 81 and slice 175. Figure 4.12 also shows the two axial slices of the reference image as well as a coronal view with yellow lines indicating the position corresponding to the specific axial slice. This trend could be due to the bias introduced by the selected reference image and (in the case of the femoral heads) somewhat arbitrary means of cutting off the more distal bone from the segmentation. Future work might improve the model accuracy by iterative boosting of the reference image [62] or using multi-atlas fusion [60]. The spinal elements attached to the pelvis further contributed to misclassified bone voxels (20% of the total misclassified bone voxels) as the model was initially developed to describe pelvic bone shape only. Figure 4.13 shows the identified bone voxels overlapped with the ground truth CT images for 4 patients where the model achieves the best performance, average performance and worst performance.

The optimal model coefficient found for each of the 17 patients ranged from 87 to -350, with a mean of -28 and a standard deviation $\sigma = 211$. The averaged maximal decreases (when compared to individual optimizations) in bone coverage when varying the model parameter by $\pm 1.5\sigma$ and $\pm 1\sigma$ around the mean were 9.7% and 6.3% across these 17 patients, respectively. These added errors are large compared to the coverage with the first mode, and suggest that the variation across individuals warrants the use of some deformation to achieve a reasonable tradeoff of sensitivity and
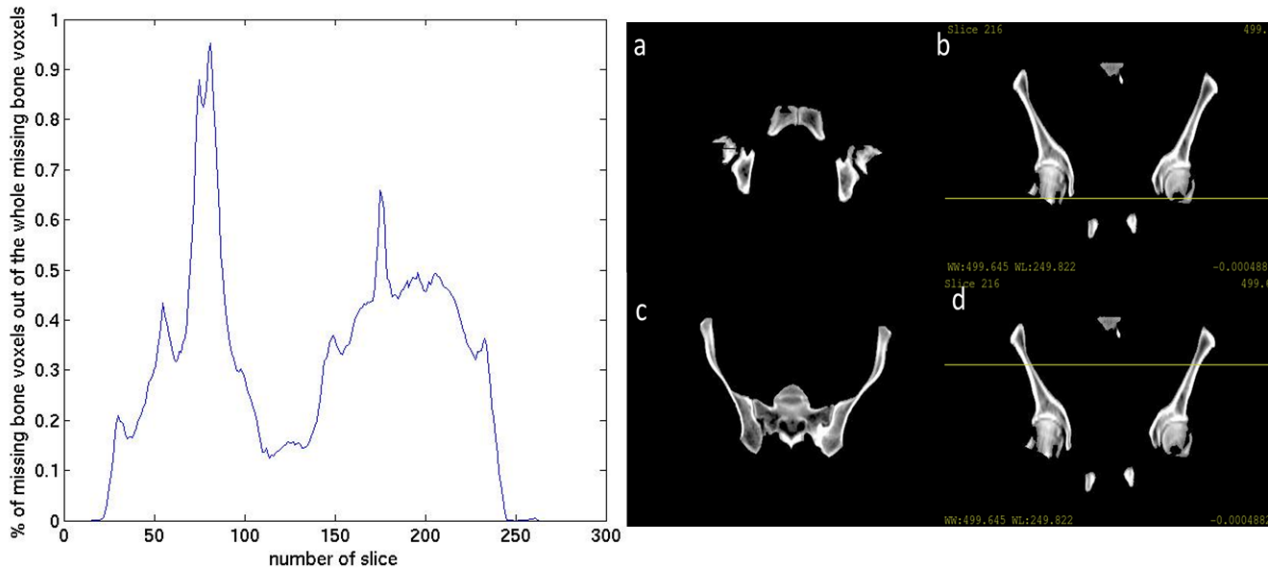
Figure 4.12: Analysis of the location of missing bone voxels. Left: Percentage of missing bone voxels in each axial slice, two peaks appear at slice 81 and slice 175. Right: (a) Axial slice 81 of the reference image. (b) Coronal plane with yellow line corresponding to axial slice 81. (c) Axial slice 175 of the reference image. (d) Coronal plane with yellow line corresponding to axial slice 175.
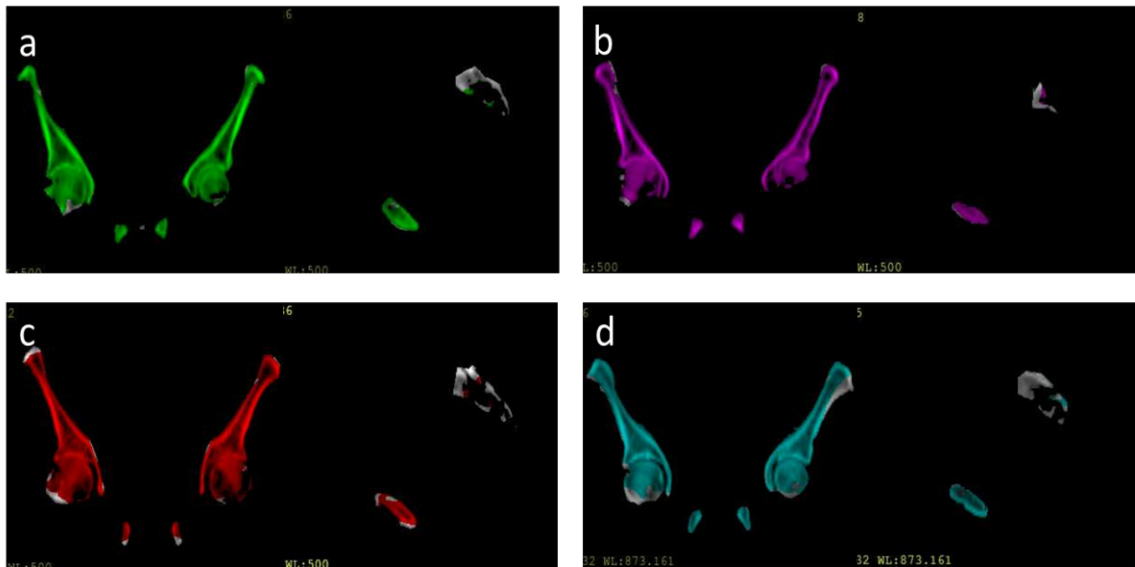
Figure 4.13: Example pelvic CT images with color wash indicating bone voxels identified by the deformed reference image for 4 patients that have different model performances: (a) patient that has bone coverage of 93%. (b) Patient that has bone coverage of 94%. (c) Patient that has bone coverage of 85%. (d) Patient that has bone coverage of 79%.

specificity in bone/air separation. On the other hand, varying the model parameter by $\pm 0.5\sigma$ around the optimal model parameter found for each patient decreased the model performance by 1.6% only, indicating the model is robust to small deviations from the true optimum. It is possible in the future to use this added knowledge to permit an even faster search for the optimal combination of deformation model and marginal expansion.

Fitting the model to CT data using the first leading mode, the leave-one-out cross validation achieved an average bone identification rate of 88%, which is similar to the result obtained from MRI data, thus validating our choice of the mean square error as the optimization criterion. Figure 4.14 shows the averaged percentage of correctly identified bone voxels by fitting the model to MRI data as well as the percentage of mislabeled air voxels versus radius of dilation of the identified bone space, with error bars specifying the standard deviation across the population. Dilating the space by 5mm improved the overlap with true bone to 96%, while the percentage of air voxels mislabeled as bone remained below 0.4%. Figure 4.15 shows an example of the bone label of the reference image overlapped with one target image, before and after the model-guided deformation. The coverage of bone voxels in the target image was improved after the deformation and dilation further improved the bone coverage.

With the bone spaces identified by the shape model, the remaining dark spaces in MR images can be labeled as air and excluded from the following tissue classification process. The air spaces were identified by thresholding the MR images in non-bone spaces at intensity cutoffs of 300, 150 and 300 for the T1-weighted, water and fat images respectively. All images were normalized by the mean intensity of the T1-weighted images before thresholding. Examples of the thresholded MR images with and without the bone mask are shown in Figure 4.16. Without the bone identification
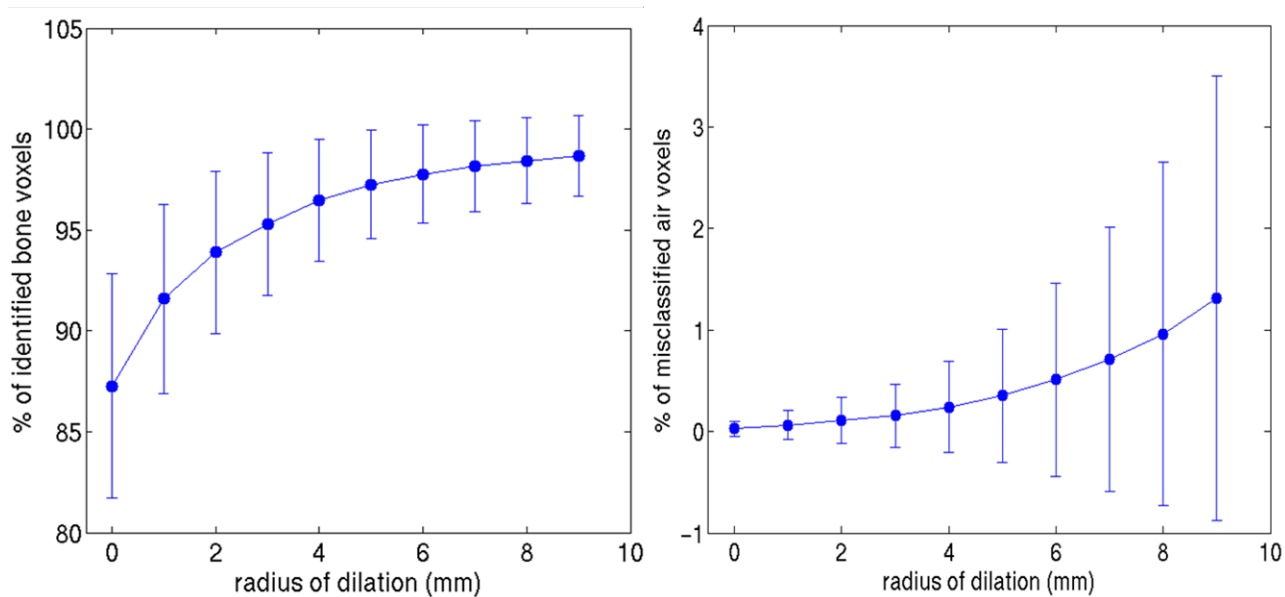
Figure 4.14: Left: fraction of correctly identified bone voxels under various radii of dilation. Right: fraction of air voxels mislabeled as bone under various radius of dilation. The length of the error bar corresponds to the standard deviation across the population.
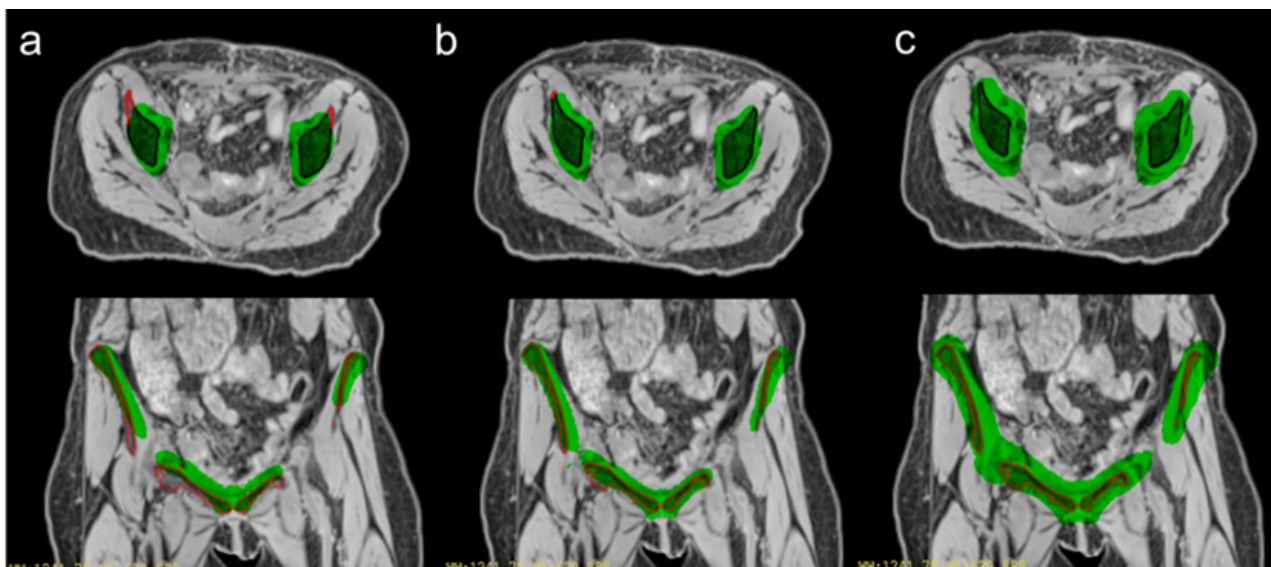
Figure 4.15: Labeled regions (green) of the reference image overlapped with the target image (with red contours for bone region) (a) before deformation (b) after deformation and (c) applying a 5 mm dilation of the labeled region after deformation.
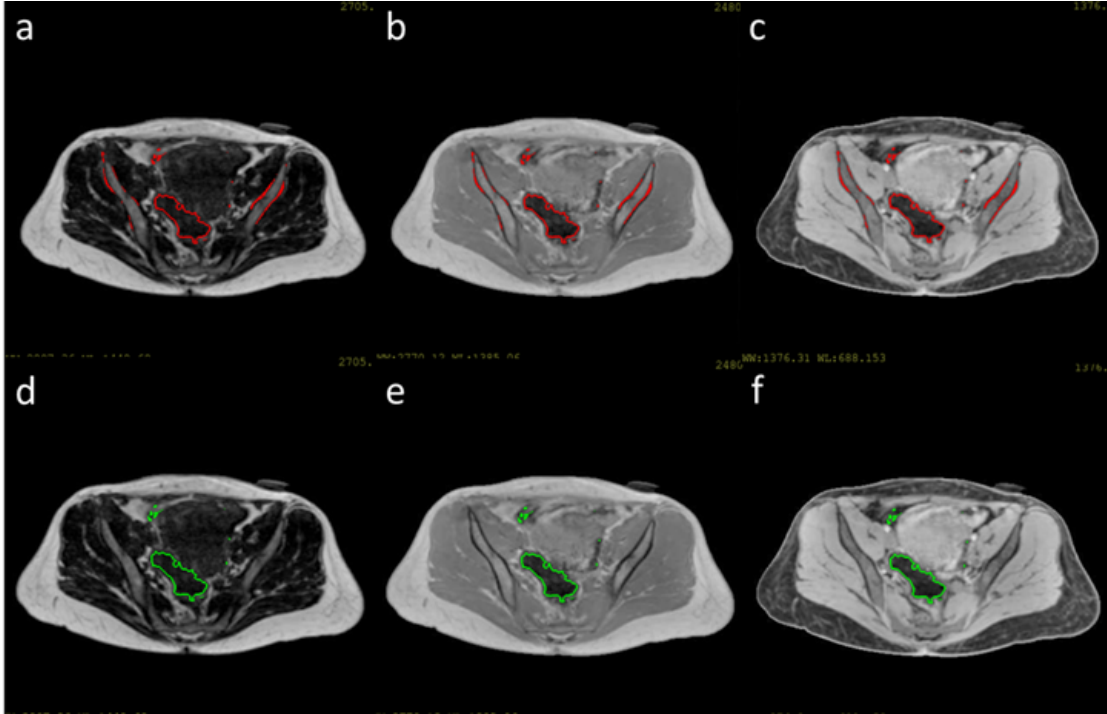
Figure 4.16: (a)-(b) Thresholded fat, T1-weighted and water image without bone mask. Threholded regions were contoured by red lines. (d)-(f) Thresholded fat, T1-weighted and water image after bone identification. Thresholded regions were contoured by green lines.

step, intensities of bone voxels and air voxels both fall below the threshold, resulting in mixture of air and bone. After the bone identification step, bone area was first masked out before thresholding and air regions were identified without mislabeling bone as air.

Including the 2nd and 3rd principal modes into the model resulted in an averaged improvement in coverage of bone of less than 2% prior to dilation. After the 5mm dilation was applied, the difference between the model with 3 modes and the model with 1 mode was only 0.06%. On the other hand, the model with 3 modes resulted in an air misclassification rate slightly higher than the model with 1 mode. The difference was 0.2% without dilation and 0.3% after the 5mm dilation. This result suggests that using the 1st leading mode may achieve a sufficiently accurate model for separation of bone and air while maintaining simplicity and computational efficiency.

4.4 Discussion

In this study, we investigated a female bone shape model to facilitate bone identification as the first step in the workflow of synthetic CT generation from pelvic MR data. By identifying a space that contains the pelvic bones plus a small expansion, a simple intensity threshold can be applied to the remainder of voxels within the pelvis to exclude air voxels from the tissue classification process, without mislabeling bone as air. This permits classification of bone without the use of ultrashort TE imaging in the pelvis.

While previous investigations of pelvic shape model [60] [61] [62] aimed at contouring bone structure accurately and evaluated the resulting models by surface distance, we focused on a simpler problem of identifying a region that contains bone and excludes air. Therefore we evaluated the model in terms of the coverage of bone and the overlap with air. Under this evaluation criterion, a single principal mode was observed to achieve sufficient accuracy, which reduced the model complexity as compared to most existing methods where more than 10 modes were typically involved. In comparison to results reported in existing work, our method gains computational efficiency where one segmentation takes 69 seconds on average on a shared 2.8 GHz CPU with 40 cores (i.e. we do not specifically allocate all cores as other jobs run with equal priority), as compared to 2.5 hours on an 8-core machine with a 3.2 GHz CPU as reported in [60]. Spatial dilation was further introduced to improve the coverage of bone, while the overlap with air remained reasonably small even with a 5mm dilation that covered 96% of pelvic bone voxels. By performing PCA on B-spline coefficients instead of deformation vectors, we were able to produce a more compact shape model that reduced computational burden for alignment with MR images.

With the constraint enforced by the shape model, local minima in the search space were removed effectively thus allowing an efficient binary search scheme to be used to find the optimal solution. The proposed method achieved reasonable accuracy on a heterogeneous dataset with variations in image quality, suggesting the model may be robust to potential imaging artifacts.

Our investigation used a randomly selected sample from the dataset as the reference image, which may introduce bias and systematical error in bone identification, such as the missing bone voxels at the top of pelvis, as shown in Figure 4.12. In the future, more sophisticated model construction schemes, such as multi-atlas fusion [60] or iterative boosting [62] may be explored to remove the potential bias and improve model performance and a larger training dataset may be investigated to construct a more representative shape model. The next chapter will incorporate this model into the entire workflow of synthetic CT generation and further validate its sufficiency in supporting treatment planning by comparing the final bone classification result on MR images to results on CT images as well as the dose calculated from synthetic CT images to the ones calculated from CT images.

4.5 Conclusion

This chapter described a female pelvic bone shape model for air/bone separation in support of pelvic synthetic CT generation for radiotherapy. Using the first principal component of shape variance learned from a training population, plus a 5mm dilation results in simple model that covers bone and excludes air with sufficient accuracy. This model can be incorporated into synthetic CT generation process to support MRI-only radiotherapy.

CHAPTER 5

Female Pelvic Synthetic CT Generation Based on Joint Intensity and Shape Analysis

5.1 Introduction

[1] Chapter 4, we developed a pelvic bone shape model to assist bone identification in MRI, as the first step towards pelvic synthetic CT (MRCT) generation. In this study, we first extend the pelvic bone shape model to cover spinal and femoral bones attached to the pelvis. Then we present a complete MRCT generation algorithm that incorporates this shape model for female pelvic radiotherapy patients, and evaluate its efficacy in supporting external beam radiation therapy treatment planning. By jointly analyzing intensity and shape features, our method is able to generate MRCT images using a single MR imaging sequence, which reduces the scanning time and avoids problems induced by tissue mobility in the pelvis confounding correspondence across successive scans.

5.2 Methods and Materials

5.2.1 Image acquisition

The dataset the same as the one in Chapter 4 was used for pelvic MRCT generation. Including the 30 CT image volumes for pelvic bone shape model construction and the 17 corresponding multi-contrast MRI dataset for MRCT generation. 7 of

---

[1]This chapter is based on publication [10] and [11].

the MR dataset were excluded from this investigation due to incomplete coverage of patient volumes in the axial plane.

We preprocessed MR and CT image volumes using previously described methods in Chapter 4. Here we briefly review the preprocessing steps. First we applied intensity inhomogeneity correction to MR images using the N4ITK algorithm [44]. Then we rigidly aligned MR images to CT images with the aim of aligning bone structures accurately without considering soft tissues, as soft tissues can deform across scanning sessions. All image volumes were reformatted to axial cuts with voxel size interpolated to $1 \times 1 \times 1$ mm^3. Finally we normalized the intensity of each MRI dataset with a scale factor that sets the mean intensity of the corresponding T1-weighted image volume to 1000.

5.2.2 Shape model construction for bone identification

To separate bone from air without using ultra-short TE imaging, we used the fact that bones have a somewhat predictable shape and built a bone shape model to facilitate bone identification. We extended the shape model presented in Chapter 4 to cover the spinal processes superior to the pelvis as well as to include a femur model based on connected component analysis. Figure 5.1 labels various bony structures that are of interest in this study. Figure 5.2 shows the flow chart for applying the bone shape model for bone identification and Figure 5.3 shows the corresponding example image of each step. First a rough mask (referred as the threshold mask) was generated by intensity thresholding, as described in Section 5.2.2.1. The pelvic bone shape model was next applied to separate pelvic bone and lumbar spine from other voxels that fall below the threshold, as described in Section 5.2.2.2. Finally, in Section 5.2.2.3 we incorporated a femur identification model to separate different tissues in the thighs. Each of these steps is detailed next.

Figure 5.1: Bony structures of interest for female pelvic radiotherapy. The previous pelvic bone (cyan) model was extended to cover femoral bones (yellow) and vertebrae (red) in this work
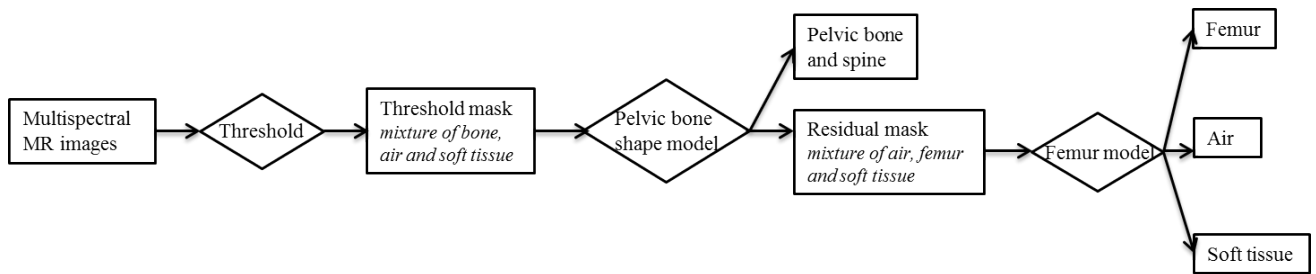
Figure 5.2: Overall scheme of separating bone, air and soft tissue
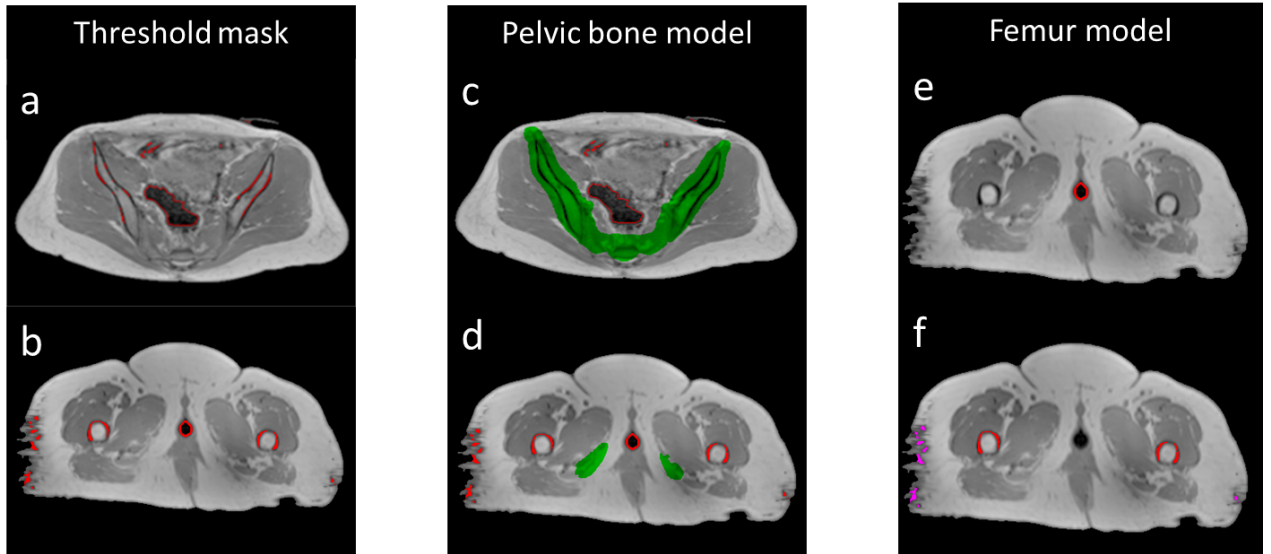
Figure 5.3: Example of applying the bone separation scheme on MRI data. Left: mask generated by thresholding (red contours) consists of a mixture of air and bone in the pelvis (a) as well as a mixture of bone, vaginal marker and soft tissue in the thighs (b). Middle: pelvic bone mask generated by the pelvic bone shape model (green) separates bone from air in the pelvis (c) but fails to cover femoral bones in the thighs (d). Right: applying the femur model separates vaginal markers (red contours in (e)), femoral bones (red contours in (f)) and soft tissue (magenta contours in (f)).

5.2.2.1 Generate the initial mask by thresholding

First, we applied intensity thresholds to the multi-contrast MRI data. Experimentally determined thresholds of 300, 300, and 150 were applied to normalized T1-weighted, fat and water images respectively. The intersection of masks from the three image volumes yielded a "threshold mask" (Figure 5.3(a) and Figure 5.3(b)). Both bone and air voxels were selected by this thresholding process as they both appear dark on MRI images, as shown in Figure 5.3(a). In addition, some soft tissue voxels in the thighs were also found to fall below the threshold due to imaging artifacts (e.g. image noise, peripheral signal loss), as shown in Figure 5.3(b). A vaginal marker used on some patients for clinical care was also selected by this thresholding process due to low image intensity.
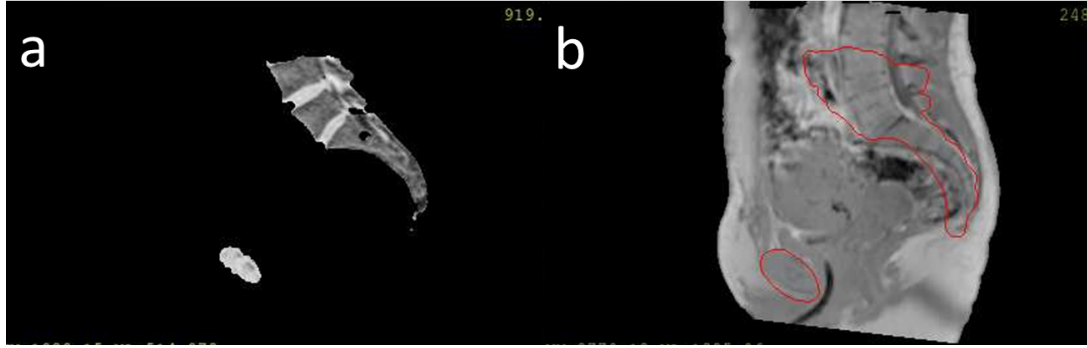
Figure 5.4: Mask (contoured in red) generated by deforming the atlas image with lower lumbar vertebrae maintained (a) overlapped with a target image (b). The coverage of the lumbar spine by the atlas image is limited cranially to roughly the upper third of the L4 vertebral level (a).

5.2.2.2 Apply pelvic bone model to the threshold mask

Next, a pelvic bone shape model was constructed from the CT scans of the 30 subjects, using the algorithm described in Chapter 4. The algorithm was primarily developed for localizing pelvic bone without considering attached lumbar vertebrae and femurs. To generate MRCT data suitable for treatment planning, we need also to classify such non-pelvic bones correctly. It is possible to build a separate shape model for those anatomical structures. However, experiments demonstrated that deforming the reference image with lumbar vertebrae maintained (Figure 5.4(a)) with the same parameters as in Chapter 4, plus the suggested 5mm dilation, was sufficient to cover the lower lumbar vertebrae in MR images with reasonable specificity (Figure 5.4(b)). Figure 5.3(c) shows an example of applying the bone mask to separate bone from air in the pelvis.

5.2.2.3 Identification of femoral bones

While applying the bone shape model covered the majority of pelvic bone and lower lumbar spine voxels, a significant volume of femoral voxels in the scanned volumes remained uncovered (referred as the residual mask) and would be mislabeled
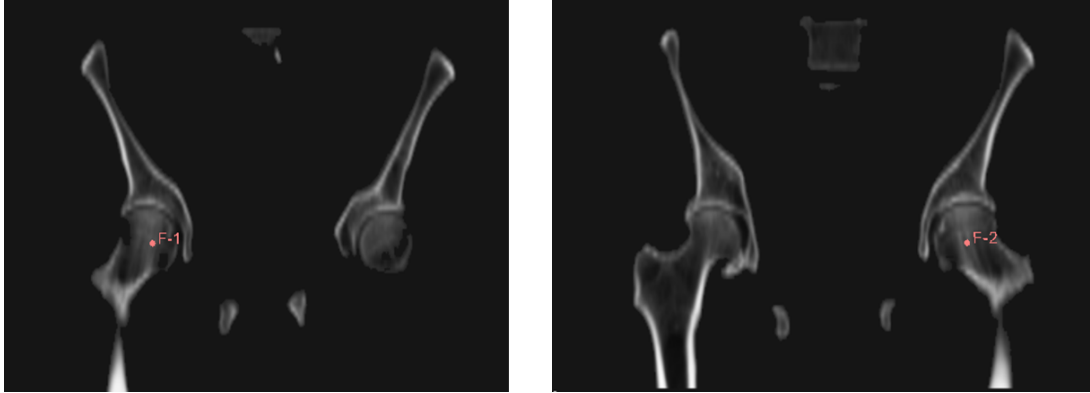
Figure 5.5: Landmark points (F-1 and F-2) placed on the femoral heads in the atlas CT image.

as air voxels, as shown in Figure 5.3(d). Extension of the shape model to cover femurs was not considered practicable as the position variations of femurs are large across patients. To address this, we extended the model based on the fact that air does not exist in the thighs. Landmark points were placed on each of the femoral heads in the atlas CT image, as shown in Figure 5.5 with positions tracked during the deformation process. Image volumes inferior to the deformed landmarks were considered as candidate space that contains potential femoral bones.

As air does not exist in the thighs, a search was done across the candidate space of the residual mask, where voxels outside the pelvis were excluded from being identified as air. A voxel was defined as inside the pelvis if there were both voxels on its left and right that belonged to the pelvic bone mask and were in the same axial slice. Figure 5.3(e) shows an example of such voxels that are inside the pelvis. The remaining voxels in the residual mask are a mixture of femur bone voxels and soft tissue voxels with low intensities. As soft tissue voxels in the residual mask appear as scattered noise while femur bone voxels have a regular pattern, we separated femur bones from soft tissues by performing 3D connected component analysis on voxels in the residual mask and identified the left/right femurs as the 2 largest

connected components, as shown by red contours in Figure 5.3(f). The remaining small connected components were treated as soft tissue voxels, as shown by magenta contours in Figure 5.3(f).

5.2.3 Fuzzy c-means classification with a shape prior

The intensity-based fuzzy c-means classification used for whole brain MRCT generation, as described in Chapter 3, is formulated as

$$\hat{\mathbf{u}}, \hat{\mathbf{v}} = \arg\min \sum_{s=1}^{d} \sum_{i=1}^{c} \sum_{k=1}^{N} u_{ik}^{m} \|x_{ks} - v_{is}\|^2 + \alpha \sum_{s=1}^{d} \sum_{i=1}^{c} \sum_{k=1}^{N} u_{ik}^{m} \|\bar{x}_{ks} - v_{is}\|^2$$

$$s.t. \sum_{i=1}^{c} u_{ik} = 1, \quad \forall k = 1, \cdots, N, \tag{5.1}$$

where $\{x_{ks}\}_{k=1,s=1}^{N,d}$ is the acquired d MRI volumes from the same object, N is the number of voxels in each image and c is the pre-determined total number of tissue classes presented. By minimizing the objective function in Eq(5.1), we get the probabilistic membership u_{ik} of the k^{th} voxel belonging to the i^{th} tissue class as well as the intensity centroid of the i^{th} class on the s^{th} image, v_{is}.

The intensity distributions of pelvic MRI data however, present challenges for this classification scheme. Firstly, volumes of different tissue types are highly unbalanced in the pelvis. For example, the volume of fat far exceeds that of pelvic bone. Standard fuzzy c-means will favor large clusters over small clusters [63]. The estimation of the bone class will be negatively affected by the soft tissue class, resulting in inaccurate bone classification. Secondly, the limited spatial resolution and large field of view lead to significant partial volume effects that cause the intensity distributions at the interface between fat and other soft tissues to be different from soft tissues, and overlapped significantly with bone marrow. Figure 5.6 shows example ROIs of fat interface and bone marrow as well as their intensity histograms (normalized to have the same peak). The intensities of the fat interface and bone marrow appear to be
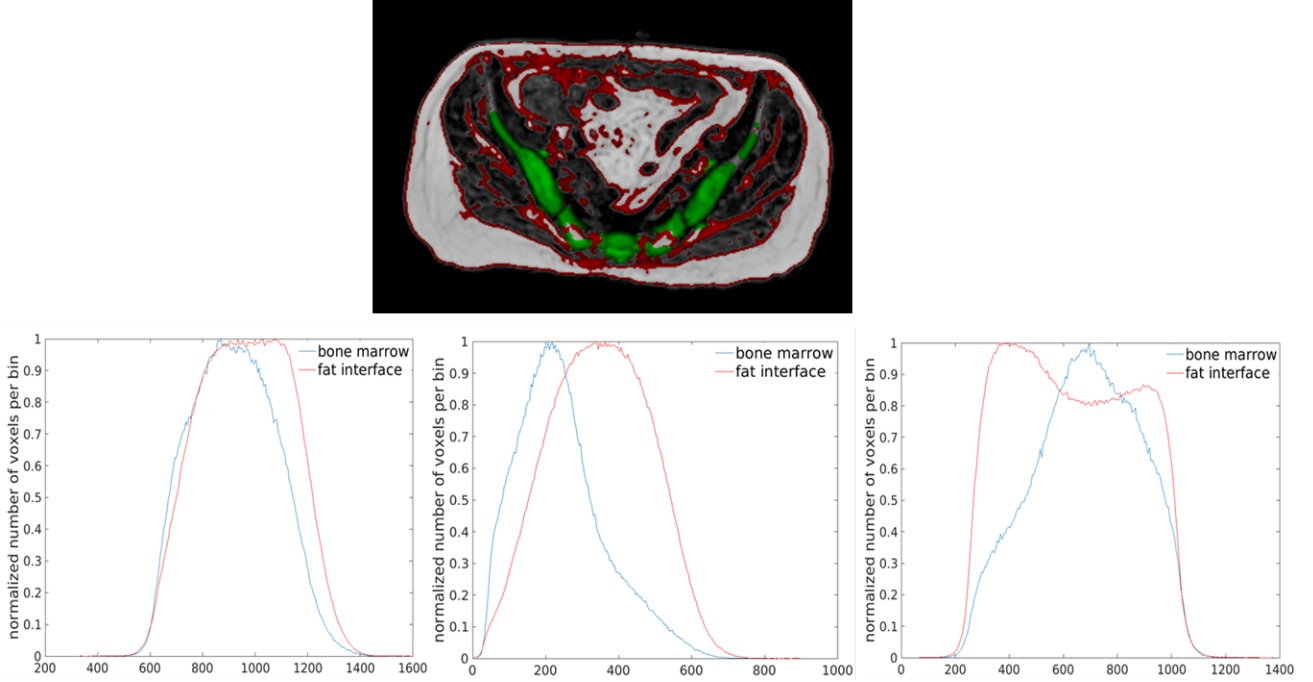
Figure 5.6: Intensity overlap between fat and soft tissues interface and bone marrow. Top: Example ROIs of interfaces between fat and other soft tissues (red) and bone marrow (green) on a fat image. Bottom: Normalized intensity histograms of fat interfaces (red) and bone marrow (blue) on T1-weighted (right), water (middle) and fat (left) images from an example patient.

non-separable. Therefore, intensity information alone will not suffice to accurately classify different tissue types from pelvic MRI data, and standard fuzzy c-means classification needs to be modified to incorporate shape information.

We added shape information to the standard fuzzy c-means classification formula by introducing a regularization term based on the pelvic bone shape model we constructed. As our bone shape model covers the majority of bone voxels [8], we discouraged voxels outside the bone mask generated by the model from being classified as bone. Mathematically, assuming the bone class is the l^{th} class, we regularized the classification on bony tissue $u_{lk}, k = 1, \cdots, N$ with the binary bone mask $\mathbf{b}$ as

$$L(\mathbf{u}, \mathbf{b}) = \sum_{k=1}^{N} \mathbf{1}(u_{kl} > 0 \quad \& \quad b_k = 0), \tag{5.2}$$

where $b_k = 0$ indicates voxel k is not covered by the bone mask and we modified the

fuzzy c-means classification with the regularization term as

$$\hat{\mathbf{u}}, \hat{\mathbf{v}} = \arg\min \sum_{s=1}^{d}\sum_{i=1}^{c}\sum_{k=1}^{N} u_{ik}^m \|x_{ks} - v_{is}\|^2 + \alpha \sum_{s=1}^{d}\sum_{i=1}^{c}\sum_{k=1}^{N} u_{ik}^m \|\bar{x}_{ks} - v_{is}\|^2 + \lambda L(\mathbf{u}, \mathbf{b})$$

$$s.t. \sum_{i=1}^{c} u_{ik} = 1, \quad \forall k = 1, \cdots, N, \tag{5.3}$$

where λ controls the impact of the regularization term on the entire classification process. Basically, the regularization term penalizes voxels to have membership in the bone class ($u_{lk} > 0$) outside the bone mask ($b_k = 0$). Since our bone mask covers a major portion of bony tissues, we set $\lambda = +\infty$. In this way, we enforced a hard constraint on the classification scheme where any voxel outside the bone mask will have zero probability of belonging to the bone class, which will also prevent the estimation of the intensity centroid of bone class from being affected by non-bone voxels.

To solve for Eq(5.3), we first initialized a rough estimation of intensity centroids of each class v_{is}. The corresponding optimal estimation of u_{ik}, without the regularization term, is given by

$$u_{ik} = \frac{\sum_{s=1}^{d}\left(\|x_{ks} - v_{is}\|^2 + \alpha\|\bar{x}_{ks} - v_{is}\|^2\right)^{-1/(m-1)}}{\sum_{s=1}^{d}\sum_{j=1}^{c}\left(\|x_{ks} - v_{js}\|^2 + \alpha\|\bar{x}_{ks} - v_{js}\|^2\right)^{-1/(m-1)}}, \tag{5.4}$$

with the hard constraint on the bone class, the bony membership u_{lk} is adjusted by

$$\tilde{u}_{lk} = \begin{cases} u_{lk} & \text{if } b_k > 0 \\ 0 & \text{otherwise} \end{cases}. \tag{5.5}$$

To make sure $\mathbf{u}$ is a valid probability distribution, we need $\sum_{i=1}^{c} u_{ik} = 1, \forall k = 1, \cdots, N$. Here we do not assume any prior knowledge on other tissue classes and distribute the residual probability ($u_{lk} - \tilde{u}_{lk}$) evenly to all remaining tissue classes

$$\tilde{u}_{ik} = u_{ik} + \frac{u_{lk} - \tilde{u}_{lk}}{c - 1}, \quad \forall i = 1, \cdots, c, i \neq l. \tag{5.6}$$
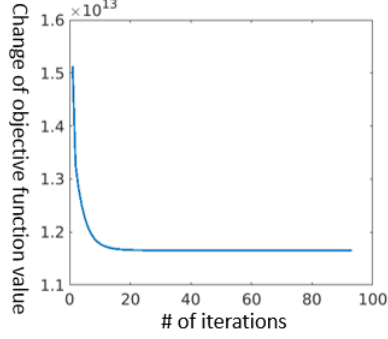
Figure 5.7: Updates of objective function value.

After calculating $\tilde{u}_{ik}$, we updated the centroid of each tissue class by

$$v_{is} = \frac{\sum_{k=1}^{N} \tilde{u}_{ik}^m (x_{ks} + \alpha \bar{x}_{ks})}{(1 + \alpha) \sum_{k=1}^{N} \tilde{u}_{ik}^m}.$$

(5.7)

We performed the above calculations Eq(5.4) through Eq(5.7) iteratively until a convergence criteria was met. In our implementation, we terminated iterations when the decrease of the objective function Eq(5.3) value was below a threshold. Roughly 100 iterations were needed for convergence across patients, which took about 18 minutes on a shared 2.5 GHz CPU with 12 cores. Figure 5.7 plots the change of objective function values during iterations.

5.2.4 MRCT generation

We generated MRCT image volumes for 9 of the 10 patients who had MRI scans (except for the one whose CT-extracted data was used as the reference image for bone model construction). First, the skin surface was extracted by thresholding the normalized T1-weighted images at 300. This surface was then cleaned up using morphologic operations (3mm dilation, filling holes and 3mm erosion). Next, air masks and bone masks were generated using the thresholding scheme together with the shape model and the identified air voxels were excluded from tissue classification. Next, the fuzzy c-means classification with a shape prior was performed on the multi-

contrast MRI data. Optimization was performed over 5 classes including compact bone, fat, muscle and the combination of fat interfaces and bone marrow (which were assigned two classes during optimization). Bone marrow and fat interfaces were separated retrospectively after the fuzzy c-means classification using the shape model, where voxels presenting bone marrow/fat interfaces were treated as bone marrow if inside the bone mask and fat interfaces otherwise. The fuzzy degree m was 1.5 and the weight for the spatial constraint term α was 3.8, the same as used previously in the head [4]. After the fuzzy c-means classification, MRCT volumes were generated by assigning each tissue class a CT number, multiplying the probability of each voxel belonging to each tissue class (u_{ik}) with the assigned CT number of that class and summing over all classes. To decide CT numbers of each tissue class, we drew ROIs that contained primarily the corresponding tissue on the reference atlas CT image. Rounding the mean intensity of each ROI to the nearest ten yielded the CT number of the class. The CT numbers assigned to fat, muscle, bone marrow and bone were -100 HU, 30 HU, 150 HU and 800 HU respectively. Fat interfaces were assigned the same CT number as fat. The CT number assigned to the identified air regions was -1000 HU.

5.2.5 MRCT evaluation

To evaluate the usefulness of the MRCT volumes for treatment planning, both their intensity correlations with corresponding CT image volumes as well as accuracy for supporting treatment planning dose calculations were evaluated. The mean absolute error (MAE) in intensity between MRCT and CT images was calculated on various ROIs drawn at different locations of the pelvis. ROIs encompassing solid bone and marrow were generated from CT images by first thresholding the images at 150 HU, followed by morphologic operations (3mm dilation, filling holes and 3 mm
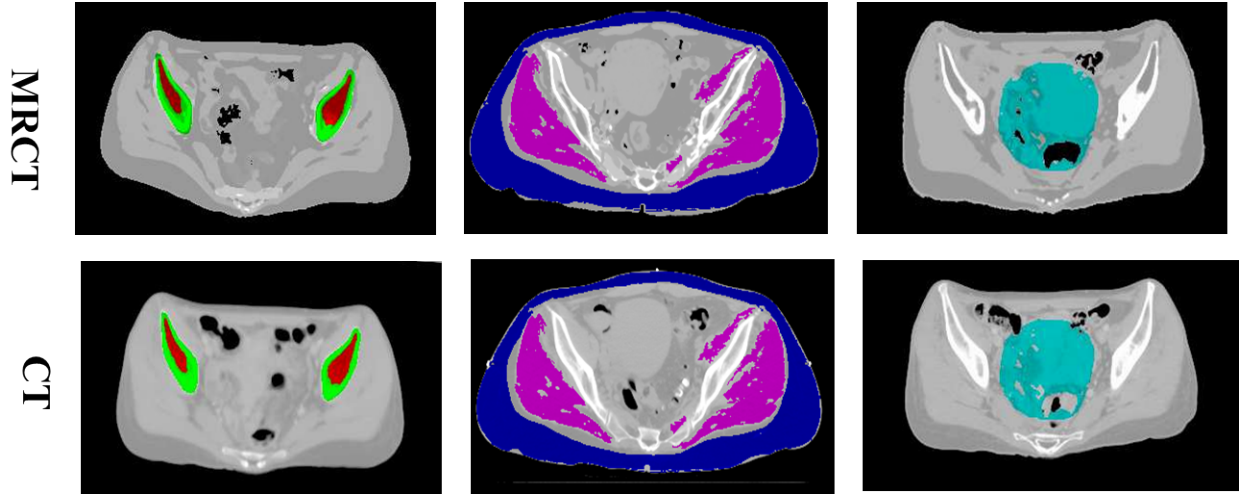
Figure 5.8: Example ROIs on MRCT (top) and CT (bottom) images. Left: ROIs for solid bone (green) and bone marrow (red). Middle: ROIs for muscle (magenta) and external fat (blue). Right: ROIs for intra-pelvic soft tissues (cyan).

erosion). Bone voxels were defined as voxels within the ROIs whose intensities were above 250 HU and marrow voxels were defined as those below 250 HU. Before evaluating MAEs for soft tissues, deformable registration was performed using commercial software (Velocity) to align the MRCT with CT images, to attempt to resolve the natural variations in patient position and internal anatomic configuration between CT and MR imaging sessions. ROIs for soft tissues (muscle, external fat and intra-pelvic soft tissues) were then manually drawn on regions where the overlap between MRCT and CT image volumes was reasonably sufficient through visual evaluation. Figure 5.8 shows example ROIs.

Both the deformably aligned MRCT and their corresponding CT image volumes were imported into a commercial treatment planning system (Eclipse 11.0, Varian, Palo Alto CA). Each patient's clinically defined structures from their actual treatment plans were used for treatment planning. Volumetric modulated arc therapy (VMAT) plans were then optimized using density grids derived from each MRCT image set. Similar to previous investigations [64], beam fluences from each of the
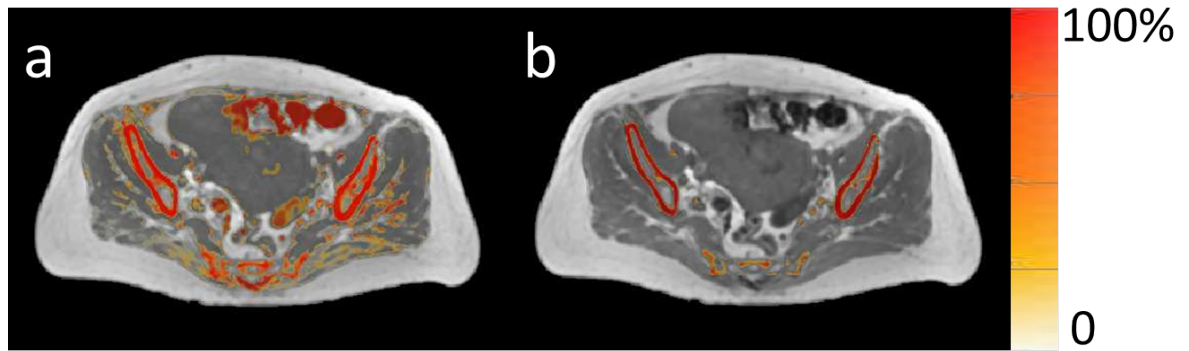
Figure 5.9: Comparison of classification results of bones (a) without and (b) with a shape prior. Probability maps of voxels belonging to the bone class (colorwash) are overlaid in their corresponding T1-weighted image.

MRCT-optimized plans were transferred to the associated CT-derived density grids, and the dose subsequently recalculated. These transposed MRCT (tMRCT) dose distributions were used to more directly evaluate the impact of density grid selection on dose calculation. For each patient, calculated dose and volume metrics were compared between MRCT optimizations and tMRCT calculations for planning target volumes (PTVs) and comparable structures including bowel, pelvis, rectum, sacrum and femur.

5.3   Results

5.3.1   Tissue classification with a shape prior

Figure 5.9 compares probability maps of voxels belonging to the bone class from fuzzy c-means without and with a shape prior, overlapped with the corresponding T1-weighted image. Without the bone mask, not only were air voxels classified as bone, certain soft tissue voxels were also assigned bone memberships, which would have caused large errors in MRCT images generated.
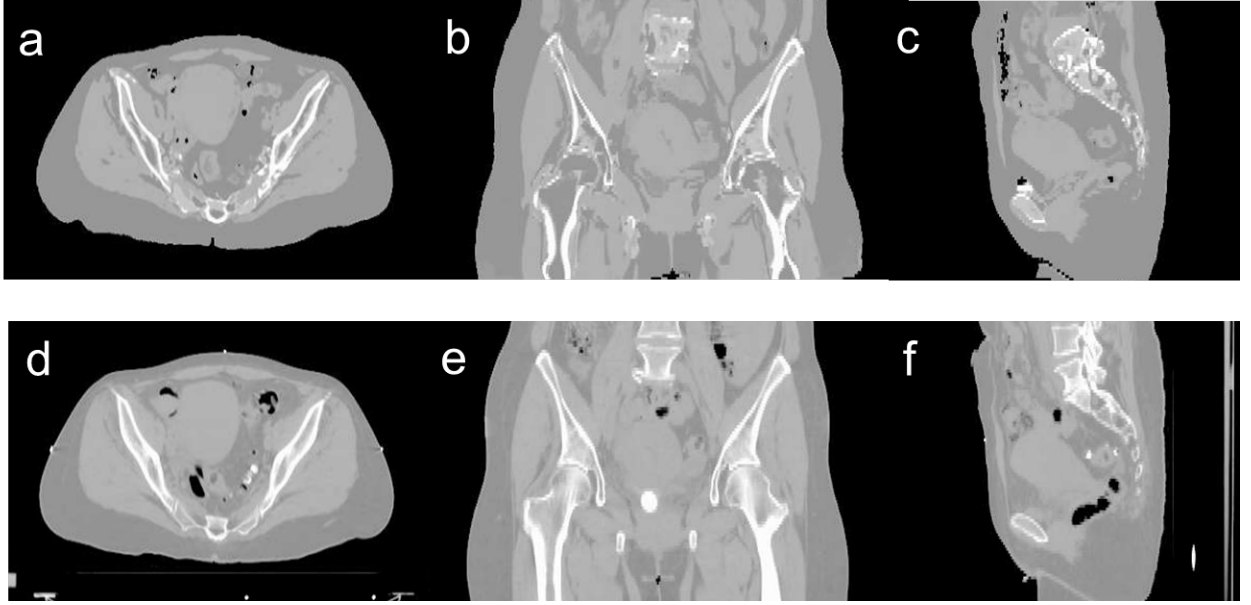
Figure 5.10: Axial (a), coronal (b) and sagittal (c) images of a MRCT image volume and corresponding cuts through the same patient's CT image volume (d), (e) and (f).

5.3.2 MRCT generation and accuracy evaluation

Figure 5.10 shows example MRCT images and their corresponding CT images. MRCT image volumes present contrast similar with CT image volumes with the exception of the superior region of the lumbar spine, which was classified as soft tissue. This is due to the limited coverage of the atlas image, as has been discussed in 5.2.2.2. Table 5.1 summarizes the statistics of MAE of the 9 patients. The average/standard deviation of MAE across 9 patients was 13.7/1.8 HU for muscle, 15.9/2.8 HU for fat, 49.1/17.8 HU for intra-pelvic soft tissues, 129.1/29.2 HU for marrow and 274.4/26.9 HU for bones.

Table 5.2 presents the mean and standard deviation of differences between treatment planning objectives evaluated using doses calculated on MRCT and tMRCT plans across all patients. Figure 5.11 shows statistics of dose differences of PTVs and organs at risk (OARs) across patients. Both absolute and relative dose differ-

65

Table 5.1: Mean absolute HU difference between MRCT and CT images across patients

	Mean	Standard deviation	Range
Muscle	13.7	1.8	9.8 to 17.4
External fat	15.9	2.8	12.0 to 19.8
Intra-pelvic soft tissue	49.1	17.8	25.5 to 75.3
Bone marrow	129.1	29.2	92.8 to 170.0
Solid bone	274.4	26.9	226.4 to 314.3

ences between MRCT and tMRCT calculations are small compared to the prescribed doses (45-58.25 Gy), with a maximal mean difference smaller than 0.3Gy/0.5%. Figure 5.12 shows statistics of volume differences of OARs for clinically relevant dose values. The volume differences are small for all OARs with a noted exception of the rectum (V45Gy), where the mean volume difference is 2.5%. From the box plot in Figure 5.12, there are outliers in the comparison of rectum V45Gy. Two patients were found to have much larger volume differences than others (9.3% and 3.9 % respectively). After excluding these two patients, the mean differences/variances drop from to 2.5%/3.3% to 0.4%/0.2%. The large difference is mostly due to the mobility of air in the rectum between MR and CT scans of these two patients which could not be fully resolved using deformable alignment, as shown in Figure 5.13. The CT image volumes show larger air pockets in the rectum while little air presents in the MR image volumes. The dose volume histograms however are very similar between MRCT and tMRCT for both outlier patients, as shown in Figure 5.14. Volume differences of rectum at other dose levels (V20Gy, V35Gy and V50Gy) are marginal (less than 0.6%) for these two outlier patients.

5.4 Discussion

This study investigated an algorithm for pelvic MRCT generation using joint shape and intensity features to differentiate and classify tissues in MR images. The

Table 5.2: Mean and standard deviation of differences between extracted dose metrics

Dose to location	Mean differences	Standard deviation of differences
PTV		
D0.5cc (Gy)	0.21	0.42
D0.5cc (% of prescribed dose)	0.4	0.8
D99% (Gy)	0.29	0.33
D99% (% of prescribed dose)	0.5	0.6
Femur		
V30 Gy(%)	0.4	0.4
Pelvis		
Mean dose (Gy)	0.10	0.10
Mean dose (% of prescribed dose)	0.2	0.2
V10 Gy(%)	0.0	0.1
V20 Gy(%)	0.1	0.1
Rectum		
Mean dose (Gy)	-0.03	0.15
Mean dose (% of prescribed dose)	-0.1	0.3
V45 Gy(%)	2.5	3.3
	0.4 with outliers removed	0.2 with outliers removed
Sacrum		
V10 Gy(%)	-0.2	0.3
V20 Gy(%)	0.0	0.1
Bowel		
D1cc (Gy)	0.18	0.40
D1cc (% of prescribed dose)	0.3	0.7
D5cc (Gy)	0.18	0.35
D5cc (% of prescribed dose)	0.3	0.6
V55Gy(cc)	0.22	0.64
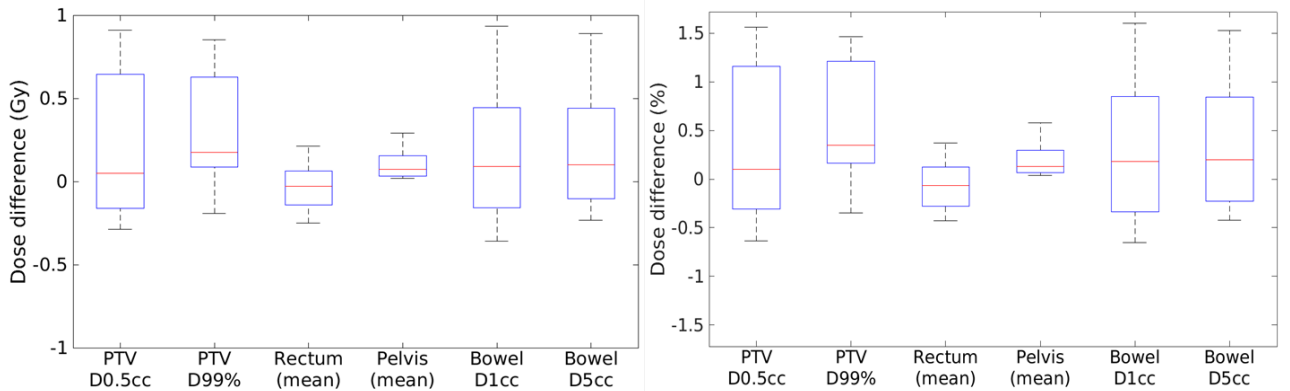
Figure 5.11: Box plot of absolute (left) and relative (right) dose differences of PTV and OARs. Red line indicates the median. Bars indicate the maximum and minimum.
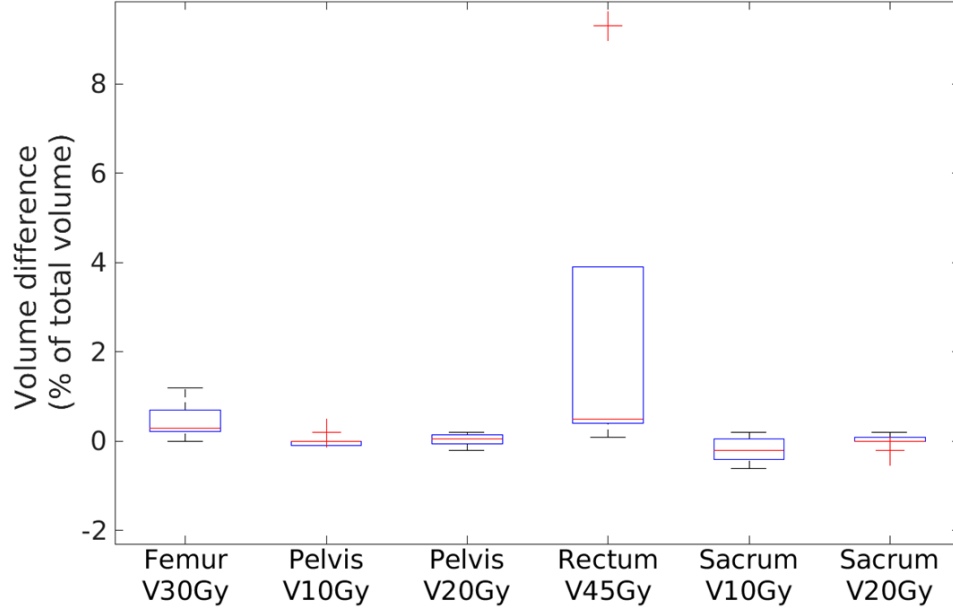
Figure 5.12: Box plot of volume difference of OARs. Red line indicates the median. Bars indicate the maximum and minimum and red crosses indicate outliers. (Data points are defined as outliers if they are greater than $q_3 + 1.5(q_3 - q_1)$ or less than $q_3 - 1.5(q_3 - q_1)$, q_3 and q_1 are the 75^{th} and 25^{th} percentiles of the sample data, respectively.)
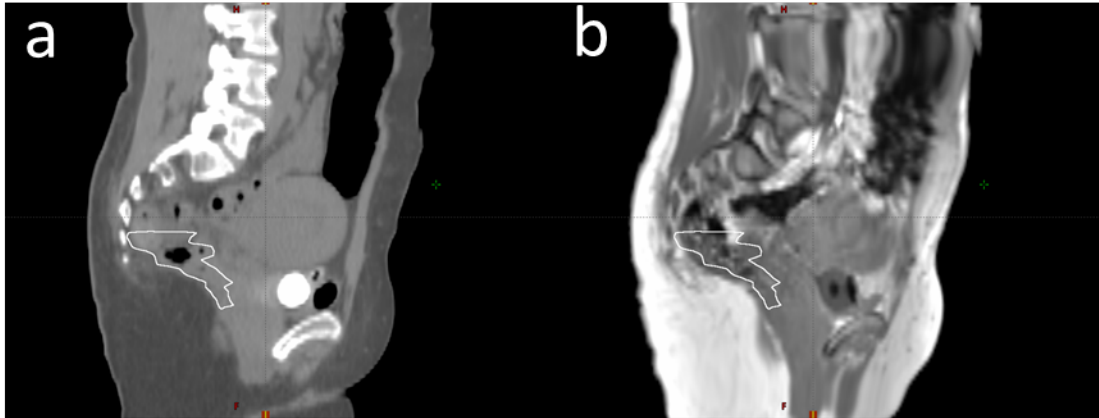
Figure 5.13: Mobility of air in the rectum (white contours) between CT scans and MR scans. (a) The CT image shows a larger air pocket in the rectum as compared to (b) the MR image (T1-weighted image after applying the deformable transformation that aligned the MRCT image to the CT image).
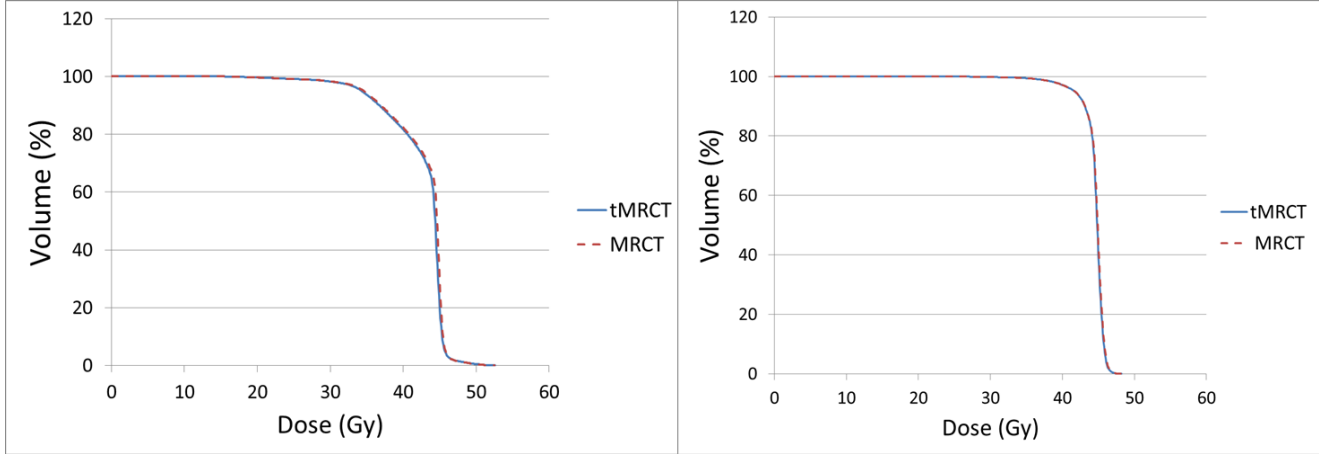
Figure 5.14: Dose volume histograms of rectum show similar results between MRCT and tMRCT plans for the two patients with larger rectum V45Gy variations.

shape patterns of pelvic bones were combined with a femoral bone extraction method to assist the separation of bone from other tissues. A regularization term was formulated using the bone shape model and added to an intensity-based fuzzy classification scheme. Regularized classification was found to reduce the misclassification of non-bone tissues effectively. The resulting MRCT images presented contrast sufficiently close to CT images to support treatment planning in radiotherapy, as validated by both direct intensity comparison and by performing treatment planning using MRCT image volumes and comparing the dose distributions on MRCT image volumes and CT image volumes. The idea of generating MRCT using both shape and intensity information is not only applicable to pelvis, but also other body parts, as has been evaluated and validated in abdominal MRCT generation [65].

The presented method differs from previous studies [33] [34] [35] [36] in that it classifies each patients data individually and assigns each voxel attenuation values based on the probability membership of the voxel belonging to a specific tissue class, rather than relying on a fixed intensity atlas. By introducing a bone shape model to the intensity-based classification scheme, our method is able to classify tissues

accurately without manual contouring of bones [52] [53] [54] [55]. The MRCT images were generated from MRI data using a single imaging sequence without ultra-short TE imaging [4] [37] [38], thus the total scan time is much shorter and the problem of patient motion, which is more significant in the pelvis than in the head, is largely avoided. Compared to Bayesian approaches based on deformable alignment between MR images [66] and exact bone segmentation through deformable registration [67], our method incorporates the shape information by simply defining a binary space that covers bone and excludes air, which can be found by a binary search algorithm [8], and thus is simpler and more efficient for clinical implementation.

Both image intensity, as well as calculated dose comparisons between MRCT and CT image volumes, show acceptably small variations. Although the MAE values for bone and marrow appear somewhat large, the dose calculation studies demonstrated that these differences did not impact dose calculation significantly. This is easily understood due to the relatively short pathlengths over which such intensity differences are sampled for beams traversing the pelvis for treatment planning. Future investigations will include possible recalibration of baseline intensity assignments to the bone and marrow classes to further improve agreement. Uncertainty in aligning MR images to CT images, which can be challenging in the pelvis due to the large deformation of soft tissues, variations in air distribution, and different filling status of organs such as bladder and rectum, complicates the direct evaluation of differences in intensity and dose. Also, the MRI scans used for this study had a shorter longitudinal field of view (FoV) than CT scans and the coverage of lumbar spine by the reference MR images was limited, and thus the bone model was not extended far superior to the pelvis. Further research will extend the FoV of MR scans as well as the coverage of lumbar spines by the bone shape model, as well as evaluate the usefulness

of MRCT image volumes as references for image-guided patient positioning.

5.5 Conclusion

A method that generates pelvic MRCT using joint shape and intensity features from MR images has been presented and evaluated. Adding shape information to the intensity based fuzzy c-means classification scheme was shown to improve the classification accuracy effectively and eliminate the need for multiple imaging sequences, including the ultra-short TE sequence previously used for air identification in the head. The presented method has the potential to provide an accurate estimation of CT information and support MRI-only radiotherapy.

Accelerated High b-value DWI for Higher-order Diffusion Analysis via a Phase-constrained Low-rank Tensor Model

6.1 Introduction

[1] Diffusion-weighted MRI (DWI) and the parameters derived from it are useful in various aspects of cancer treatment, including tumor evaluation and grading [68], target volume delineation for treatment planning [12] and treatment response evaluation [14] [28]. This chapter describes our work on improving DWI reconstruction techniques to extract information useful for radiotherapy, while maintaining a clinically feasible scan time.

Conventionally, the "Apparent Diffusion Coefficient" (ADC), as described in 2.3.2 is the most commonly used parameter derived from DWI [14] [28] [68]. ADC is usually extracted from DWI images acquired under a b-value of 1000 s/mm^2. Recent interest however, is emerging in high b-value (> 1000 s/mm^2) DWI. Various studies show the advantage of using b-values higher than conventional b-values, including better differentiation of tumor from benign tissues in prostate [69] and more accurate grading and delineation of cerebral gliomas [15] [70]. However, the signal-to-noise ratio (SNR) of high b-value DWI is very poor. Repeated acquisition and signal averaging are performed in practice to improve SNR at the price of a longer scan

---

[1]This chapter is based on publication [17] and [18].

time. Besides, the calculation of ADC assumes mono-exponential decay behavior while tissues actually exhibit multi-exponential decay behavior [71]. When extending the range of b-values (from $b = 0$ s/mm^2 to $b > 3000$ s/mm^2) in high b-value DWI, multi-exponential decay behavior has been reliably demonstrated [16]. To model such multi-exponential decay behavior, higher-oder diffusion models, such as bi-exponential model [72] and stretched model [73] have been proposed. Such higher-order diffusion models require more data points for accurate fitting, which will further prolongate the image acquisition time.

Most work in accelerated DWI acquisition focuses on undersmapling the k-q space [74, 75, 76, 77, 78], where images are acquired using multiple diffusion directions and two b-values. While in cancer imaging, such as in studies of brain glioma [79, 80, 15], usually only 3 orthogonal directions are sampled using a range of b-values. Besides, most existing work uses a moderate b-value at around 1000 s/mm^2, while in high b-value DWI, the extension of b-values to 2000 s/mm^2 and higher, plus the limited sampling of diffusion directions, significantly degrades the SNR and poses challenges for image reconstruction. Although [75] presents reconstruction results using $b = 2000$ s/mm^2 and 64 diffusion directions, that method was reported to fail at an acceleration factor of 8 when undersampling k-space only.

Low-rank tensor models have been applied to accelerated MRI and the results are promising [81, 82, 83, 84]. However, most methods do not consider the coil dimension, yet multichannel acquisition plays a key role in clinically used acceleration schemes such as parallel imaging. Although [84] builds a tensor of the form space×coil×time and exploits the dependencies between coils, acceleration using multi-channel acquisition can use not only the dependencies between coils, but also the dependencies between neighboring k-space samples. Therefore, exploiting coil dependencies only

may not fully utilize the benefits of multichannel acquisition. In [85], the authors studied tensor models for parallel imaging, but with single image acquisition instead of image series acquisition.

This chapter proposes a new low-rank tensor model that exploits both the global low-rank structure of DWI that results from the strong dependencies between diffusion signals of voxels and the local low-rank structure that results from the dependencies between neighboring k-space samples. The method also includes a phase constraint to account for the large phase variations between b-values and handles partial Fourier acquisition naturally. We evaluate our method using both patient and simulated data and compare our method to the clinically used parallel imaging only method, as well as another low-rank tensor model-based method [84]. The proposed method achieves an acceleration factor of 8 and shows reduced noise and aliasing, as proved on both simulated and patient data.

6.2   Theory

6.2.1   Low-rank tensor model construction

In high b-value DWI, a series of 2D DWI images are acquired using N_b different b-values. Each b-value image is acquired using N_c coils. Denoting the size of the imaging matrix as $N_x \times N_y$, we record a collection of k-space samples $\mathcal{D} \in \mathbb{C}^{N_x \times N_y \times N_c \times N_b}$. We propose to build a 3D tensor from $\mathcal{D}$ that will allow us to exploit both local and global dependencies of this high dimensional dataset.

6.2.1.1   Block-Hankel matrix construction for local low-rankness

First, we model the dependencies between neighboring k-space samples across coils. For each b-value, we have the multi-coil k-space signals received by coils as

$\mathcal{A}_b \in \mathbb{C}^{N_x \times N_y \times N_c}$, $b = 1, \cdots, N_b$. $\mathcal{A}_b$ is a 3D tensor, with

$$
\mathcal{A}_b^{(c)} = \begin{pmatrix} a_{1,1} & a_{1,2} & \cdots & a_{1,N_x} \\ a_{2,1} & a_{2,2} & \cdots & a_{2,N_x} \\ \vdots & \vdots & \ddots & \vdots \\ a_{N_y,1} & a_{N_y,2} & \cdots & a_{N_y,N_x} \end{pmatrix}, c = 1, \cdots, N_c \tag{6.1}
$$

Parallel imaging techniques such as GRAPPA [86] suggest that a k-space sample $a_{x,y,c}$ can be estimated as a linear combination of its neighbors across coils:

$$
a_{x,y,c} = \mathbf{g}^T \mathbf{a}_{x,y}, \tag{6.2}
$$

where $\mathbf{a}_{x,y} = [a_{x-w,y-w,1}, \cdots, a_{x+w,y+w,1}, a_{x-w,y-w,2}, \cdots, a_{x+w,y+w,N_c}] \in \mathbb{C}^{(2w+1)^2 N_c}$, $\mathbf{g}$ is the GRAPPA kernel and w is the kernel size.

The model in Eq(6.2) implies a linear dependency between neighboring k-space samples. To model such linear dependency into a rank-deficient matrix, we construct a block-Hankel matrix $\mathbf{A}_b$ from $\mathcal{A}_b$, following the SAKE [87] method:

$$
\mathbf{A}_b = [\mathbf{a}_{w,w}, \mathbf{a}_{w+1,w}, \cdots, \mathbf{a}_{N_x-w,N_y-w}] \in \mathbb{C}^{(2w+1)^2 N_c \times (N_x-2w+1)(N_y-2w+1)}. \tag{6.3}
$$

An illustration of $\mathbf{A}_b$ is shown in Figure 6.1. $\mathbf{A}_b$ is constructed by sliding a window of size $(2w+1) \times (2w+1)$ across the k-space samples, and vectorizing the k-space samples selected by the window into one column of $\mathbf{A}_b$. Consider one column of $\mathbf{A}_b$, denoted as $\mathbf{a}$, from Eq(6.2), each element $a_{x,y,c}$ in $\mathbf{a}$ can be expressed as $a_{x,y,c} = \mathbf{g}^T \mathbf{a}_{x,y}$, $\mathbf{a}_{x,y}$ is another column in $\mathbf{A}_b$. Therefore, $\mathbf{a}$ can be expressed as a linear combination of other columns in $\mathbf{A}_b$ so $\mathbf{A}_b$ is rank-deficient. Like the SAKE [87] method, our method exploits this low-rank property of $\mathbf{A}_b$ to improve image reconstruction results.

6.2.1.2 3D tensor construction for global low-rankness

The block-Hankel matrix above helps exploit the low-rank property locally within a k-space neighborhood. There are further redundancies we can exploit. When we
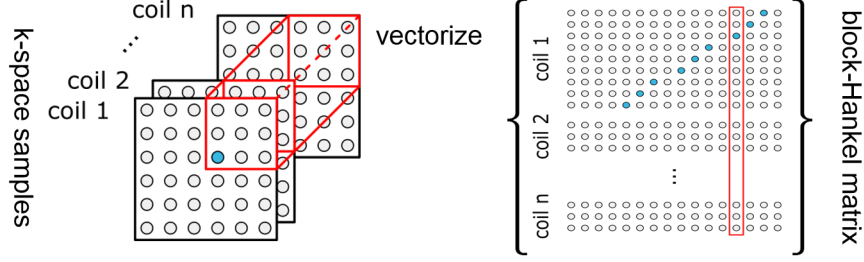
Figure 6.1: Illustration of the block-Hankel matrix construction

acquire multiple diffusion-weighted images of the same object, each spatial location (x, y) will have a signal decay curve $\mathbf{d} \in \mathbb{C}^{N_b \times 1}$ induced by diffusion. Putting all signal decay curves together forms a high-dimensional matrix

$$\mathbf{D} = [\mathbf{d}_1, \cdots, \mathbf{d}_{N_x N_y}] \in \mathbb{C}^{N_b \times N_x N_y} \tag{6.4}$$

A limited number of tissue types are present in the object. The number of tissue types is much smaller than the number of voxels. Assuming each tissue has a signature signal decay curve due to its own diffusion property, the signal decay at each voxel can be viewed as a linear combination of these signature decay curves, depending on the tissue composition of that voxel. Therefore, the high dimensional matrix $\mathbf{D}$ has a low-rank structure.

Multiplying the signal decays with coil sensitivity maps will only change the absolute magnitude of the signal, but not the decay rate. Therefore multi-coil diffusion-weighted images should still have the low-rank structure. Further more, applying Fourier transform to $\mathbf{D}$ to get the corresponding k-space data matrix will not change the low-rank property due to the orthogonality of Fourier transform. Therefore, if we put the block-Hankel matrix in section 6.2.1.1 at each b-value together and stack them into a 3D tensor $\mathcal{X} \in \mathbb{C}^{(2w+1)^2 N_c \times (N_x - 2w+1)(N_y - 2w+1) \times N_b}$, with $\mathcal{X}^{(b)} = \mathbf{A}_b$, $\mathcal{X}$ will be a low-rank tensor.

6.2.1.3 Phase constraint and partial Fourier reconstruction

The above discussion assumes image phases remain the same between acquisitions, which holds in some cases, such as T2 mapping that acquires images with varying TEs [88]. However, DWI sequence is very sensitive to motion, and phase variations are large between different b-values [24]. If uncorrected, such phase variations would invalidate the assumption of global low-rankness. As phase maps are slow-varying spatially, we can estimate phase information using the center of the k-space only, and use the estimated phase maps to remove or model phase variations.

Consider one coil image and denote the phase maps estimated under different b-values of that coil image as $\mathbf{P}_1, \cdots, \mathbf{P}_{N_b}$. To correct for phase variations of the k-space data $\mathbf{M}_1, \cdots, \mathbf{M}_{N_b} \in \mathbb{C}^{N_x \times N_y}$, we may calculate the phase difference between different b-values as $\mathbf{P}_b^{\text{diff}} = \mathbf{P}_1 \mathbf{P}_b^*$, $b = 1, \cdots, N_b$, and correct the phase variation as

$$\mathbf{M}_b^{\text{correct}} = \mathcal{F} \mathbf{P}_b^{\text{diff}} \mathcal{F}^* \mathbf{M}_b, \tag{6.5}$$

where $\mathcal{F}$ and $\mathcal{F}^*$ are FFT and IFFT operators respectively. In this way, images at all b-values will have a consistent phase map, namely $\mathbf{P}_1$.

Equivalently, we can also define $\mathbf{x}_b = \text{mag}(\mathcal{F}^* \mathbf{M}_b) \in \mathbb{R}^{N_x \times N_y}$, and calculate $\mathbf{M}_b^{\text{correct}}$ as

$$\mathbf{M}_b^{\text{correct}} = \mathcal{F} \mathbf{P}_1 \mathbf{X}_b, \tag{6.6}$$

The advantage of using Eq(6.6) for phase correction is that it can also fill up missing k-space data when partial Fourier acquisition is applied. By constraining $\mathbf{X}$ to be real, we can simply multiply the phase maps back with $\mathbf{X}$ and perform FFT to fill up the unacquired k-space due to partial Fourier sampling, which is exactly the POCS algorithm [89] for partial Fourier reconstruction.

6.2.1.4 Problem formulation

Combining the low-rank tensor model with the phase constraint scheme, we formulate a model-based image reconstruction algorithm for DWI as

$$\hat{\mathbf{y}}, \hat{\mathbf{x}}, \hat{\mathcal{X}} = \underset{\mathbf{y}, \mathbf{x}, \mathcal{X}}{\arg \min} \|\mathbf{d} - \Omega \mathbf{y}\|_2^2 + \lambda \mathbf{R}(\mathcal{X})$$

$$\text{s.t.} \quad \mathbf{y} = \mathcal{F} \mathbf{P} \mathbf{x}, \quad \mathcal{X} = \mathcal{H} \mathcal{F} \mathbf{P}_1 \mathbf{x}, \quad \mathbf{x} \in \mathbb{R}^{N_x N_y N_c N_b}, \tag{6.7}$$

where $\mathbf{d}$ is the vectorization of sampled k-space data, Ω is the k-space sampling operator, $\mathcal{F}$ is the (full) Fourier transform operator and $\mathcal{H}$ is the operator that constructs block-Hankel tensor as described in section 6.2.1.1 and section 6.2.1.2. The tensor $\mathcal{X}$ is constructed from phase corrected k-space data $\mathcal{F} \mathbf{P}_1 \mathbf{x}$, and is a 3D tensor of the size $(2w+1)^2 N_c \times (N_x - 2w + 1)(N_y - 2w + 1) \times N_b$. $\mathbf{P} \in \mathbb{C}^{N_x N_y N_c N_b \times N_x N_y N_c N_b}$ is a diagonal matrix that contains the phase information estimated for each coil/b-value image. $\mathbf{P}_1$ is of the same size of $\mathbf{P}$, and is the coil phase information, estimated from the phase maps at $b = 0$ s/mm^2, but replicated N_b times.

The regularizor $\mathbf{R}$ encourages the low-rank structure of tensor $\mathcal{X}$. Low-rank tensor decomposition is an N-P hard problem. In this work we chose the regularizor $\mathbf{R}$ to be a hard constraint on the n-rank [90] of the tensor $\mathcal{X}$ such that

$$(\text{rank}(\mathbf{X}_{(1)}), \text{rank}(\mathbf{X}_{(2)}), \text{rank}(\mathbf{X}_{(3)})) \leq (r_1, r_2, r_3), \tag{6.8}$$

where $\mathbf{X}_{(i)}$ denotes the ith order matrix unfolding of tensor $\mathcal{X}$. We enforce this hard constraint efficiently by performing truncated multilinear singular value decomposition (SVD)[91][92] that reduces the problem size significantly, which is beneficial given the large size of the tensor $\mathcal{X}$.

6.2.2 Algorithm

We propose an ADMM algorithm to solve problem(6.7) efficiently. All subproblems have closed-form solutions. The augmented Lagrangian function can be written as

$$L(\mathbf{y}, \mathbf{x}, \mathcal{X}, \mathbf{u}_1, \mathbf{u}_2) = \|\mathbf{d} - \Omega\mathbf{y}\|_2^2 + \lambda\mathbf{R}(\mathcal{X}) + \mu_1(\|\mathbf{y} - \mathcal{F}\mathbf{P}\mathbf{x} + \mathbf{u}_1\|_2^2 - \|\mathbf{u}_1\|_2^2)$$
$$+ \mu_2(\|\text{vec}(\mathcal{X} - \mathcal{H}\mathcal{F}\mathbf{P}_1\mathbf{x} + \mathbf{u}_2)\|_2^2 - \|\text{vec}(\mathbf{u}_2)\|_2^2),$$

$$(6.9)$$

where vec denotes vectorization. We minimize (6.9) by updating splitting variables iteratively.

The update of $\mathbf{y}$ is

$$\mathbf{y}^{k+1} = (\Omega^*\Omega + \mu_1\mathbf{I})^{-1}\big(\Omega^*\mathbf{d} + \mu_1(\mathcal{F}\mathbf{P}\mathbf{x}^k - \mathbf{u}_1^k)\big), \qquad (6.10)$$

by matrix inversion lemma

$$(\Omega^*\Omega + \mu_1\mathbf{I})^{-1} = \frac{\mathbf{I}}{\mu_1} - \frac{1}{\mu_1}\Omega^*(\mathbf{I} + \frac{\Omega\Omega^*}{\mu_1})^{-1}\Omega\frac{1}{\mu_1}$$
$$= \frac{1}{\mu_1}\big(\mathbf{I} - \frac{\Omega^*\Omega}{1 + \mu_1}\big), \qquad (6.11)$$

as $\Omega\Omega^* = \mathbf{I}$. Putting (6.11) back into (6.10), after some simplifications, we can show

$$\mathbf{y}^{k+1} = \frac{\Omega^*\mathbf{d} + \mu_1^k\Omega^*\Omega(\mathcal{F}\mathbf{P}\mathbf{x}^k - \mathbf{u}_1^k)}{1 + \mu_1} + (\mathbf{I} - \Omega^*\Omega)(\mathcal{F}\mathbf{P}\mathbf{x}^k - \mathbf{u}_1^k), \qquad (6.12)$$

which can be calculated easily by updating sampled locations of $\mathbf{y}$ using the weighted-average of k-space samples and $(\mathcal{F}\mathbf{P}\mathbf{x}^k - \mathbf{u}_1^k)$, and unsampled locations using $(\mathcal{F}\mathbf{P}\mathbf{x}^k - \mathbf{u}_1^k)$. The update of $\mathbf{x}$ is

$$\mathbf{x}^{k+1} = \text{Real}\big(\mu_1\mathbf{P}^*\mathcal{F}^*(\mathbf{y}^{k+1} + \mathbf{u}_1^k) + \mu_2\mathbf{P}_1^*\mathcal{F}^*\mathcal{H}^*(\mathcal{X}^k + \mathbf{u}_2^k)\big)/(\mu_1 + \mu_2), \qquad (6.13)$$

where $\mathcal{H}^*$ is an operator that averages antidiagonal entries in the block-Hankel matrix (which are replicates of the same k-space samples) and puts the average back into appropriate locations [87]. By construction, $\mathcal{H}^*\mathcal{H}$ is an identity operator. The update of the low-rank tensor is calculated as

$$\mathcal{X}^{k+1} = \text{mlsvd}(\mathcal{H}\mathcal{F}\mathbf{P_1}\mathbf{x}^{k+1} - \mathbf{u}_2^k), \tag{6.14}$$

where mlsvd is the truncated multilinear singular value decomposition, as described in section 6.2.1.4, with a core tensor size of (r_1, r_2, r_3). The code is available online [93]. The updates of auxiliary variables are

$$\mathbf{u}_1^{k+1} = \mathbf{u}_1^k + (\mathbf{y}^{k+1} - \mathcal{F}\mathbf{P}\mathbf{x}^{k+1}) \tag{6.15}$$

$$\mathbf{u}_2^{k+1} = \mathbf{u}_2^k + (\mathcal{X}^{k+1} - \mathcal{H}\mathcal{F}\mathbf{P_1}\mathbf{x}^{k+1}) \tag{6.16}$$

The above process was iterated until the gap between the two splitting variables $\mathbf{y}$ and $\mathcal{F}\mathbf{P}\mathbf{x}$ is below a threshold. Experimentally we observed that after 20 iterations the gap between the two variables stabilizes. We ran the algorithm 50 iterations to guarantee convergence.

6.3 Methods and Materials

6.3.1 Image acquisition

Under institution review board approval, one brain glioma patient was scanned with a Diffusion-weighted Echo Planar Imaging (EPI) sequence on a Siemens 3T scanner with 20-channel coil arrays and 3 orthogonal diffusion directions. Eleven b-values ranging from 0 to 2500 s/mm^2 were used with an increment of 250 s/mm^2. The voxel size was $1.25 \times 1.25 \times 5.2$ mm^3, with flip angle 90°, TE = 93 ms and TR = 9300 ms. The total number of slices was 34 and the imaging matrix size was 192×192.
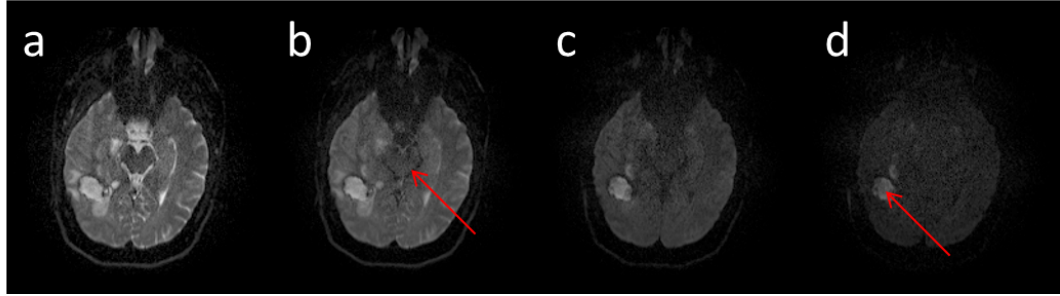
Figure 6.2: DWI images for different b-values: (a) b = 0 s/mm^2, (b) b = 250 s/mm^2, (c) b = 1000 s/mm^2, (d) b = 2500 s/mm^2. Fluid (red arrows in (b)) signals decay fast and are undetectable for a b-value of 250 s/mm^2 while tumor (red arrows in (d)) signals decay slowly and are still visible for a b-value of 2500 s/mm^2.

As the echo planar imaging sequence is prone to geometric distortion, especially for a long acquisition time, parallel imaging (GRAPPA) was used during the acquisition with a subsampling rate of 4 to ensure reasonable geometric accuracy as well as total scan time. The center of the k-space (81 lines) at b-value equals zero was fully sampled to calculate the GRAPPA interpolation kernel. Partial Fourier acquisition was also applied with 25% of k-space not sampled. Figure 6.2 shows example images for different b-values. It can be seen that different tissues have very different diffusivities. Signals of fluid decay very fast and are undetectable for a b-value of 250 s/mm^2 while signals from tumor decay very slowly and are still visible for a b-value of 2500 s/mm^2. Also, without repeated acquisition and signal averaging, high b-value images show very poor SNR. Figure 6.3 compares the DWI image for a b-value of 2000 s/mm^2, without and with signal averaging.

6.3.2 Simulation setup

To maintain a reasonable total scan time for patients, no repeated acquisition was done. As a result, the dataset is very noisy at high b-values, as shown in Figure 6.3. To provide noise-free ground truth for quantitative reconstruction error analysis, we simulated a DWI dataset using imaging parameters the same as the clinical
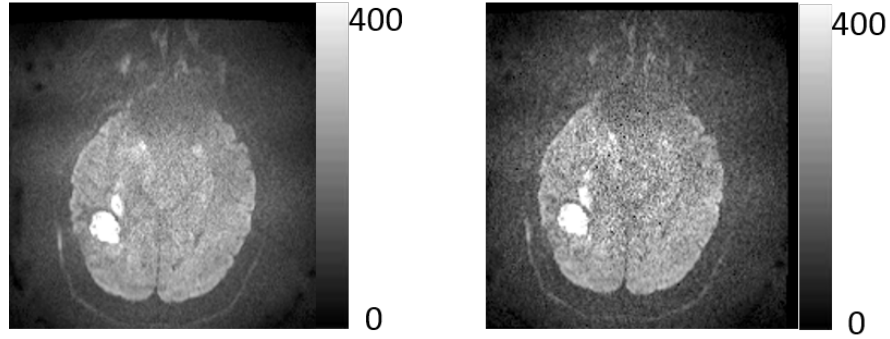
Figure 6.3: DWI image for a b-value of 2000 s/mm^2, with (left) and without (right) signal averaging. By repeating the acquisition 4 times and averaging the acquired image, the image on the left shows improved SNR than the one on the left.

scan. Using the BrainWeb [94] phantom, we considered three tissue types: white matter (WM), gray matter (GM) and CSF. Table 6.1 shows relaxation and diffusion properties of each tissue type. The relaxation properties were from BrainWeb. To simulate diffusion properties, we referred to [95] and used a bi-exponential decay model. The imaging parameters were flip angle 90°, TE = 93 ms, TR = 9300 ms with eleven evenly spaced b-values from 0 to 2500 s/mm^2. The phantom was reformatted to a voxel size of $1 \times 1 \times 5$ mm^3, for consistency with the clinical scan. Figure 6.4 shows example simulated images under different b-values.

	Proton density	T1(ms)	T2(ms)	D1(10^{-3}mm^2/s)	D2(10^{-3}mm^2/s)	f
WM	770	500	70	0.17(0.01)	1.3(0.08)	0.43(0.02)
GM	860	833	83	0.33(0.04)	1.74(0.18)	0.32(0.03)
CSF	1000	2569	329	NA	2(0.01)	NA

Table 6.1: Relaxation and diffusion parameters used for simulation. D1 and D2 are the diffusion coefficients of the two compartments in the bi-exponential model. f denotes the fraction of the slow compartment. The diffusion parameters for each voxel were generated from a Gaussian distribution, with mean and variance (in the parentheses) specified in the table.

To generate multi-coil k-space data, we scanned a uniform water phantom. Coil sensitivity maps were estimated from the center of the k-space from the uniform water phantom scan. The brain phantom was scaled to be smaller than the water phantom so that we can have a reliable coil sensitivity estimation at each location within
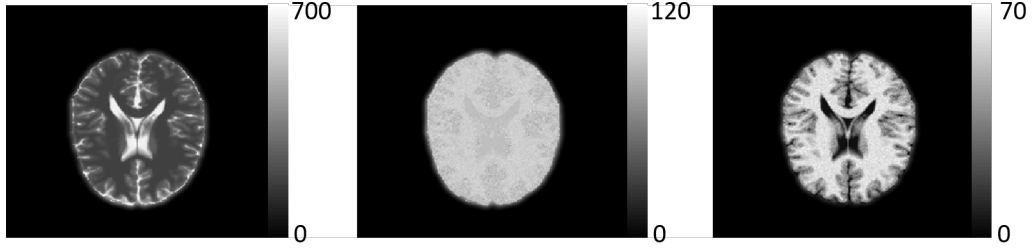
Figure 6.4: Example simulated images at left: b = 0 s/mm^2; middle: b = 1000 s/mm^2; right: b = 2500 s/mm^2.

the brain region. We further simulated the motion-induced phase variations across different b-values. Following [96], we assumed a rigid body motion and linear phase variations. We then added complex Gaussian white noise to the multi-coil k-space data. We estimated covariance matrix of noise distribution from noise calibration data from the clinical scan. Figure 6.5 shows the signal-to-noise ratio of images reconstructed using sum of square reconstruction from the simulated multi-coil k-space data, estimated as the mean signal within the brain region, divided by the standard deviation of the background. Figure 6.5 also plots the signal-to-noise ratio of one patient image, reconstructed using 4-fold GRAPPA [86]. To compensate for the noise amplification due to parallel imaging, the signal-to-noise ratio of the patient data was scaled by $\sqrt{4} = 2$. The simulated data and the scaled clinical data exhibit similar signal-to-noise ratios.

6.3.3 Ghosting correction, k-space sampling and phase estimation

EPI sequence is very sensitive to eddy currents, due to its zig-zag trajectory of k-space sampling. Eddy currents will result in a shift in the center of the k-space. In EPI acquisition, as even and odd acquisition lines are traversed in the opposite direction, the shift is not uniform across k-space samples and will result in the "N/2" ghosting artifacts in the reconstructed images. Before image reconstructions, we first
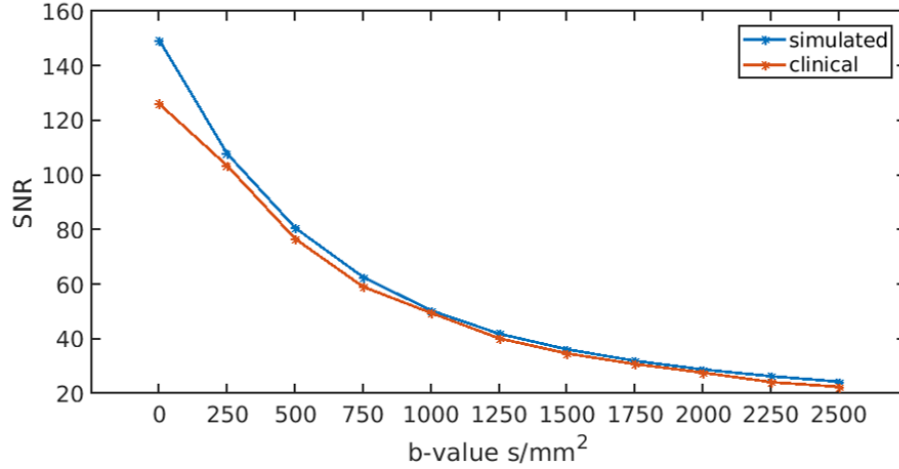
Figure 6.5: Signal-to-Noise ratio (SNR) of the simulated and clinical data, under different b-values. The simulated data shows SNR similar with the clinical data.

corrected for eddy currents. Three calibration lines $\mathbf{l}_1, \mathbf{l}_2, \mathbf{l}_3 \in \mathbb{C}^{N_x}$ were acquired. Among them $\mathbf{l}_1, \mathbf{l}_3$ traversed k-space in the same direction, and $\mathbf{l}_2$ in the opposite direction.

We first calculated $\mathbf{l}_{\mathrm{odd}} = (\mathbf{l}_1 + \mathbf{l}_3)/2$ and $\mathbf{l}_{\mathrm{even}} = \mathbf{l}_2$. We reversed $\mathbf{l}_{\mathrm{even}}$ and calculated the phase difference as $\mathbf{g} = \mathrm{angle}(\mathcal{F}^* l_{\mathrm{odd}} \odot \mathrm{conj}(\mathcal{F}^* l_{\mathrm{even}}))$, where $\odot$ denotes point-wise production. Figure 6.6 plots example phase shift $\mathbf{g}$. We fitted a linear model to $\mathbf{g}$, using the central part (61 data points) of $\mathbf{g}$ with good signal-to-noise ratio only. The estimated phase shift due to eddy currents is $\hat{\mathbf{g}}(x) = a + bx$, for $-N_x/2 < x \leq N_x/2$. Figure 6.6 also plots the estimated phase shift. We corrected an even acquisition line $\mathbf{l}$ by first reverting it as $\mathbf{l}_{\mathrm{rev}}$. The corrected k-space sample is $\mathbf{l}_{\mathrm{correct}} = \mathcal{F}(\exp(i\hat{\mathbf{g}}) \odot \mathcal{F}^* \mathbf{l}_{\mathrm{rev}})$

We retrospectively undersampled the k-space of both datasets, where we fully sampled the center of the k-space, and randomly undersampled the peripheral part of k-space, along the phase-encoding direction only. Note that the patient dataset was undersampled by 4 with a quarter plane of k-space not acquired. Our sam-
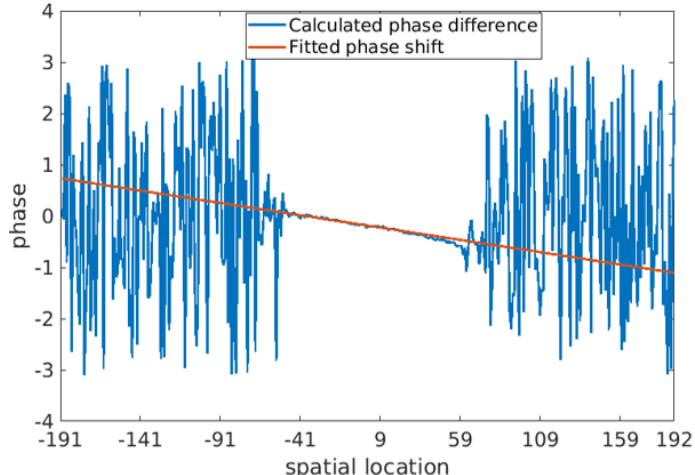
Figure 6.6: Calculated phase difference and estimated phase shift for ghost correction.

pling scheme further undersamples this dataset. For consistency in comparison, we undersampled the simulated dataset the same way as the patient dataset. Define the acceleration factor as M/P, where M is the total number of data points of the k-space and P is the number of sampled k-space data points, both undersamplings of patient and simulated dataset achieve an acceleration factor of 8.0. Figure 6.7 shows example sampling scheme of the patient data.

The sampling scheme above does not fully sample the k-space center at b-values greater than 0 mm^2/s. To estimate phase information at each b-value, which is necessary for correcting motion-induced phase variations, we first calculated a GRAPPA [86] kernel from the auto-calibration region at $b = 0$ s/mm^2 and used it to fill up the regularly undersampled k-space center for other b-values. The phase map was then calculated from the GRAPPA-filled center (21 lines) of k-space. To avoid Gibbs ringing, we applied a Hanning window along both frequency and phase encoding directions, before estimating the phase map. Figure 6.8 shows example phase estimations from the center of the k-space.
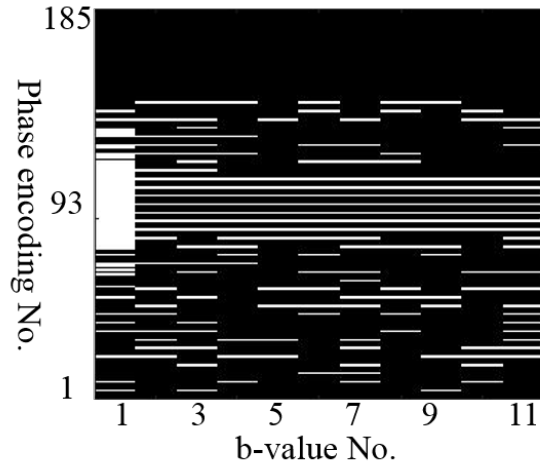
Figure 6.7: Sampling scheme of patient data. White lines indicate sampled readouts.
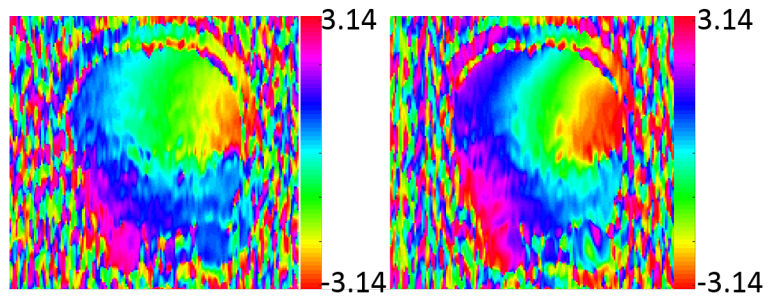
Figure 6.8: Phase map estimated from the center of the k-space of the same coil at different b-values. Phase maps vary between b-values as the DWI sequence is very sensitive to motion.

6.3.4 Evaluation

We compared our proposed method with another low-rank tensor model-based algorithm, denoted as the LRT algorithm, originally developed for accelerated dynamic MRI [84]. This algorithm dose not consider the local low-rankness within k-space neighborhoods, and constructs the tensor by reshaping the 4D k-space data $\mathcal{D} \in \mathbb{C}^{N_x \times N_y \times N_c \times N_b}$ into a 3D tensor of size $N_x N_y \times N_c \times N_b$. To account for phase variations in diffusion-weighted MRI, we added phase constraints, as described in section 6.3.3 into this algorithm,

For the simulated dataset, we calculated reconstruction errors for different algorithms, defined as $\|\hat{\mathbf{y}} - \mathbf{y}\| / \|\mathbf{y}\|$, where $\hat{\mathbf{y}}$ and $\mathbf{y}$ are reconstruction and noise-free groundtruth respectively. We also evaluated different algorithms in support of higher-order diffusion analysis. We fitted a bi-exponential decay model

$$S(b)/S(0) = f \times (\exp(-bD_1) + (1 - f) \times (\exp(-bD_2), \qquad (6.17)$$

where $S(b)$ and $S(0)$ are the signal magnitude at b-values equal b and 0 mm^2/s respectively, to each voxel using both reconstructed and ground truth images.

Eq(6.17) is a nonlinear equation and fitting can be challenging with the presence of noise. To improve fitting accuracy, we assume that some side information is available to us. First, we assume that we have Regions of Interest (ROI) drawn for each tissue types. In the clinic, physicians will usually contour different tissues before treatment planning. The contours can be drawn by combining diffusion-weighted images with other images (for example, T1-weighted images or T2-FLAIR) to improve contouring accuracy. Second, we assume the average diffusion properties of different tissues are known to us. With this prior knowledge, we initialize the non-linear least-square fitting with mean D_1, D_2 and f, as provided in table 6.1, according to the major

tissue classes presented in each voxel. We also set the lower and upper bounds for the three parameters as $[0\text{mm}^2/\text{s}, 0\text{mm}^2/\text{s}, 0]$ and $[3\text{mm}^2/\text{s}, 3\text{mm}^2/\text{s}, 1]$. The constrained nonlinear least-square fitting problem is solved using an interior-point method [97].

For the clinical dataset, due to the noise amplification of parallel imaging, it is not clear whether the difference between model-based reconstructions and GRAPPA reconstructions is due to reconstruction errors or noise reductions. The low signal-to-noise ratio and the more complicated diffusion properties with the presence of brain tumor also make voxel-wise comparison of diffusion parameters infeasible. Instead, we used physician drawn contours of normal tissues (white matter and gray matter), and fitted the bi-exponential model to signals averaged within the drawn ROIs. By averaging signals over a large number of voxels, we hope to reduce noise and evaluate if our algorithm introduces systematic differences in reconstruction when compared to GRAPPA reconstruction.

6.4   Results

6.4.1   Validation of the low-rank structure and the choice of rank

Figure 6.9 and figure 6.10 plot the singular value distributions of the unfolded tensors constructed from 8-fold undersampled k-space data, using both the proposed method and the LRT algorithm [84]. Comparing to the LRT algorithm, the proposed block-Hankel tensor exhibits sharper drop in singular values when unfolded along the coil dimension. The hard constraint on the rank number, (r_1, r_2, r_3) is determined by visually inspecting the singular value distributions. The singular value truncation points are marked on Figure 6.9 and figure 6.10. The hard constraints for the proposed algorithm and the LRT algorithm are $(35, 75, 4)$ and $(13, 30, 4)$, corresponding to coil, spatial and b-value dimensions respectively.

To evaluate if the low-rank structure changes during iterations, we also plotted
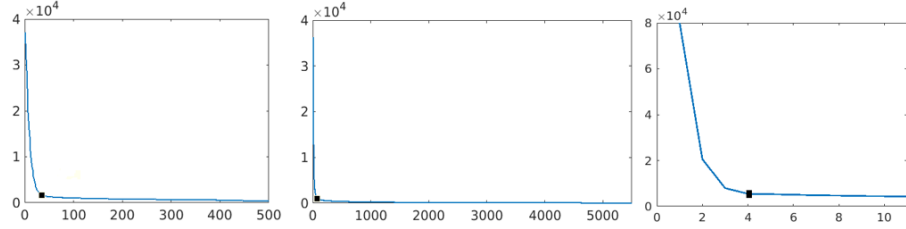
Figure 6.9: Singular value distributions of the proposed tensor, unfolded along coil, spatial and b-value dimensions respectively. Black dots mark the singular value truncation point.
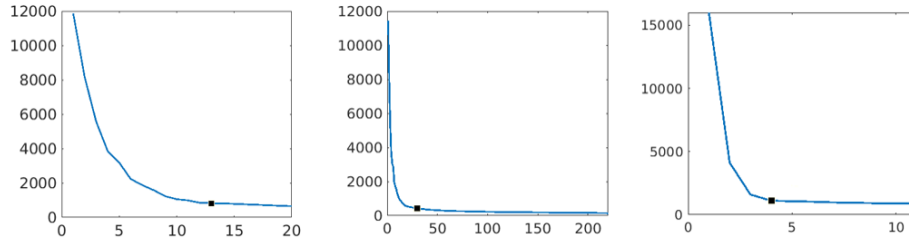
Figure 6.10: Singular value distributions of the tensor constructed using the LRT algorithm, unfolded along coil, spatial and b-value dimensions respectively. Black dots mark the singular value truncation point.

the singular value distributions of the unfolded tensor constructed from the final estimated k-space data, as shown in Figure 6.11. The distributions of singular values do not change significantly, while the magnitude of singular values after the truncation points appear smaller, as compared to Figure 6.9.

We also looked at the structure of the core tensor $\mathcal{S} \in \mathbb{C}^{r_1 \times r_2 \times r_3}$. Figure 6.12 shows (magnitude) frontal slices of $\mathcal{S}$ constructed from 8-fold undersampled k-space data with $r_1 = 35, r_2 = 75, r_3 = 4$. We observed a nearly diagonal structure in the
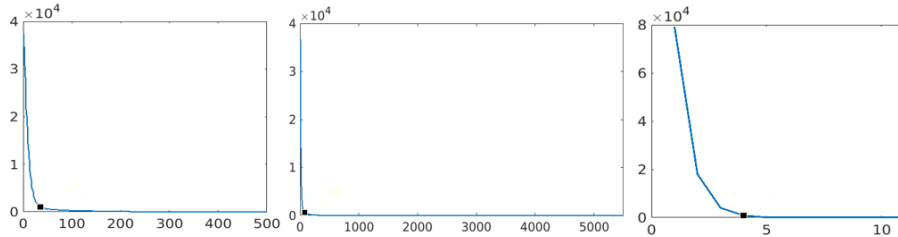
Figure 6.11: Singular value distributions of the proposed tensor constructed from the final estimation of k-space data, unfolded along coil, spatial and b-value dimensions respectively. Black dots mark the singular value truncation point.
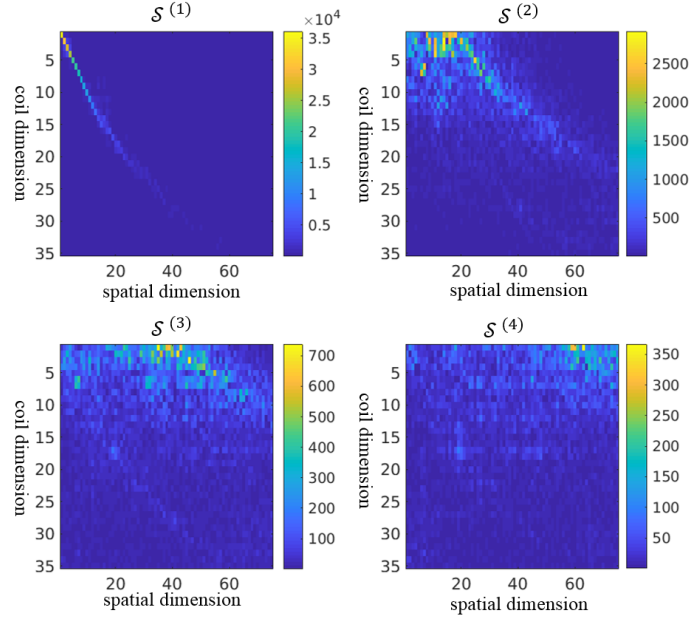
Figure 6.12: Frontal slices of the core tensor $\mathcal{S}$.

first slice of the core tensor, which combines the strongest components calculated from tensor decomposition. We also tested a sparsity constraint on the core tensor, by adding a penalty term $\|\text{vec}(\mathcal{S})\|_1$ to the objective function in Eq(6.9), but no significant difference is observed in the image reconstruction results with and without this sparsity constraint.

Figure 6.13 visualizes the spatial basis vectors estimated from the final filled-up k-space samples, by reshaping the vectors into 2D matrices and performing inverse Fourier transform on the matrices. It is clear that the first few components represent most of the image information. On the other hand, as we change the spatial rank constraint from 75, which is the "cut-off" point in truncated SVD to 80, and visualize the 76th to 80th basis vectors, we do not observe a clear change from anatomical structures to random noise, as shown in figure 6.13b and figure 6.13c. This suggests that a more comprehensive study on the choice of low-rank constraint may be needed instead of visual inspection of singular value distributions.
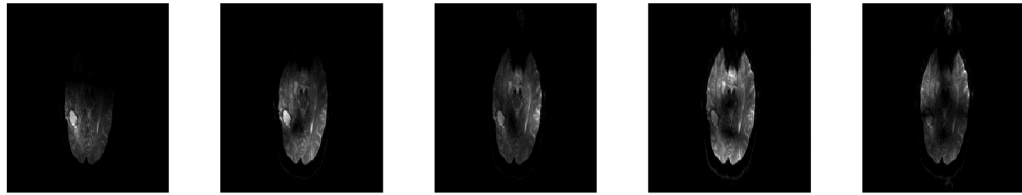
(a) The 1st to the 5th spatial basis
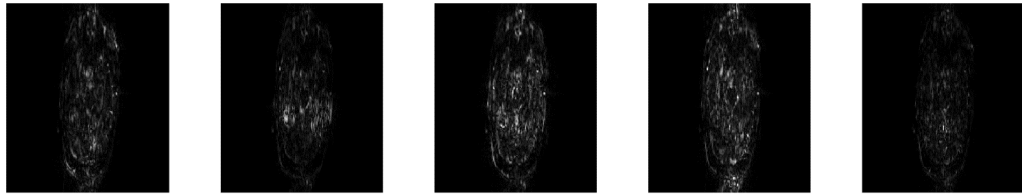
(b) The 71st to the 75th spatial basis
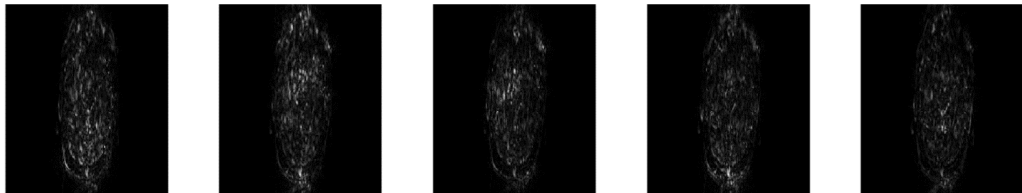
(c) The 76th to the 80th spatial basis

Figure 6.13: Visualization of selected spatial basis

6.4.2 Results on the simulated dataset

Figure 6.14 shows example reconstructed images (b-value equals 2500 mm^2/s), using 4-fold GRAPPA (denoted as the PI method), the LRT algorithm with 8-fold acceleration, and our proposed algorithm with 8-fold acceleration. Reconstruction error maps, as compared to the groundtruth are also included. Although both the LRT algorithm and our method show improved signal-to-noise ratio as compared to the GRAPPA reconstruction, the reconstruction error maps reveal that the LRT reconstruction has more aliasing than our method. Such aliasing implies the need to exploit additional model constraints, such as local low-rankness. Figure 6.15 plots the reconstruction errors of different methods. The proposed method achieves the lowest reconstruction errors for all b-values.

Figure 6.16 shows the diffusion parameter mapping results, by fitting the bi-exponential model described by Eq(6.17) to reconstructed images. Although the aliasing artifacts in LRT reconstructions are not obvious visually, they have an impact on higher order diffusion analysis, which is very sensitive to reconstruction errors due to the non-linear fitting process, as indicated by the red arrow in Figure 6.16.

6.4.3 Results on the patient dataset

Figure 6.17 shows example reconstruction images using the clinical data. All three methods produce good reconstruction images at $b = 0$ s/mm^2, possibly due to the high SNR and the fully sampled auto-calibration region. However, at a high b-value ($b = 2500$ s/mm^2), GRAPPA reconstruction results in poor SNR, and the LRT method shows aliasing (indicated by the red arrow), which is consistent with the results using the simulated data, while our method produces clean reconstruction.

Figure 6.18 shows example ROIs of white matter (magenta) and gray matter
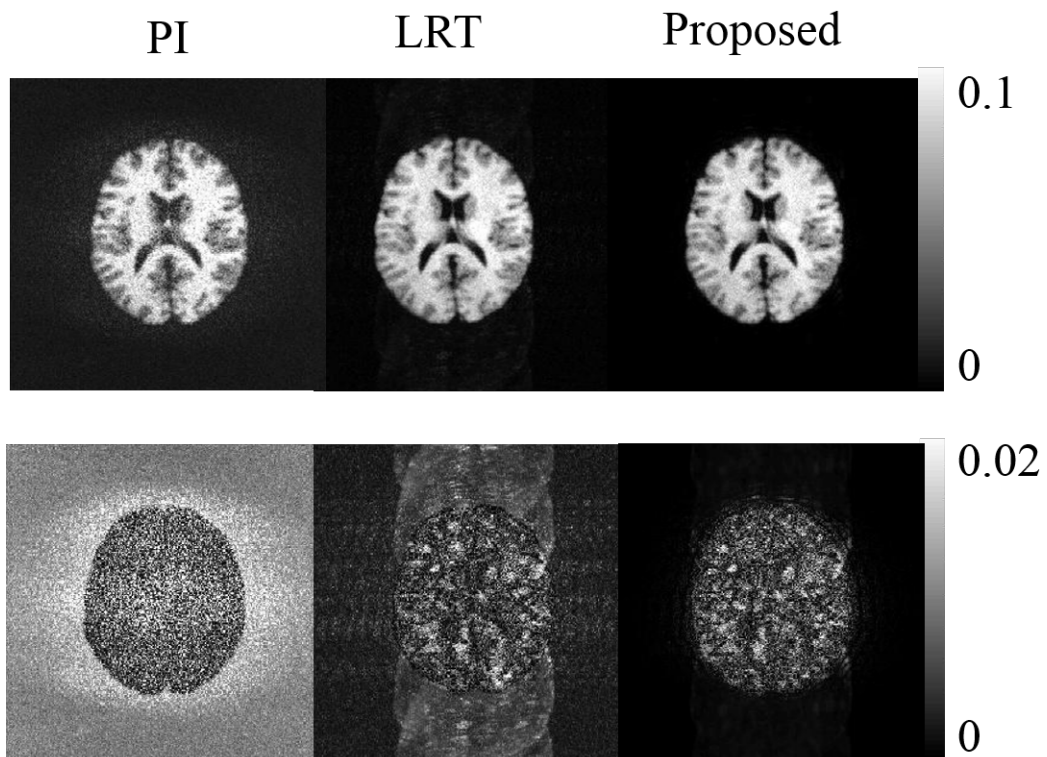
Figure 6.14: Upper: Reconstructed images (b = 2500 mm^2/s). Lower: reconstruction errors as compared to the groundtruth, using GRAPPA reconstruction (PI method); low-rank tensor model without local low-rankness (LRT method) and the proposed method.
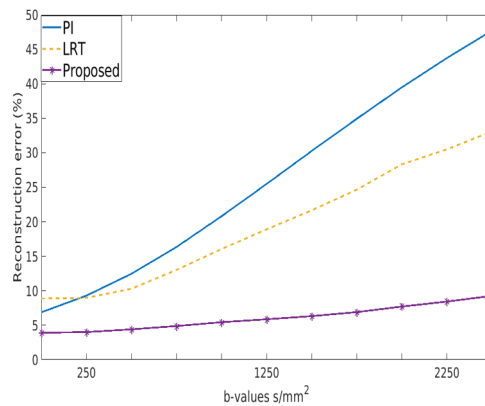
Figure 6.15: Plot of reconstruction errors of different methods for each b-value.
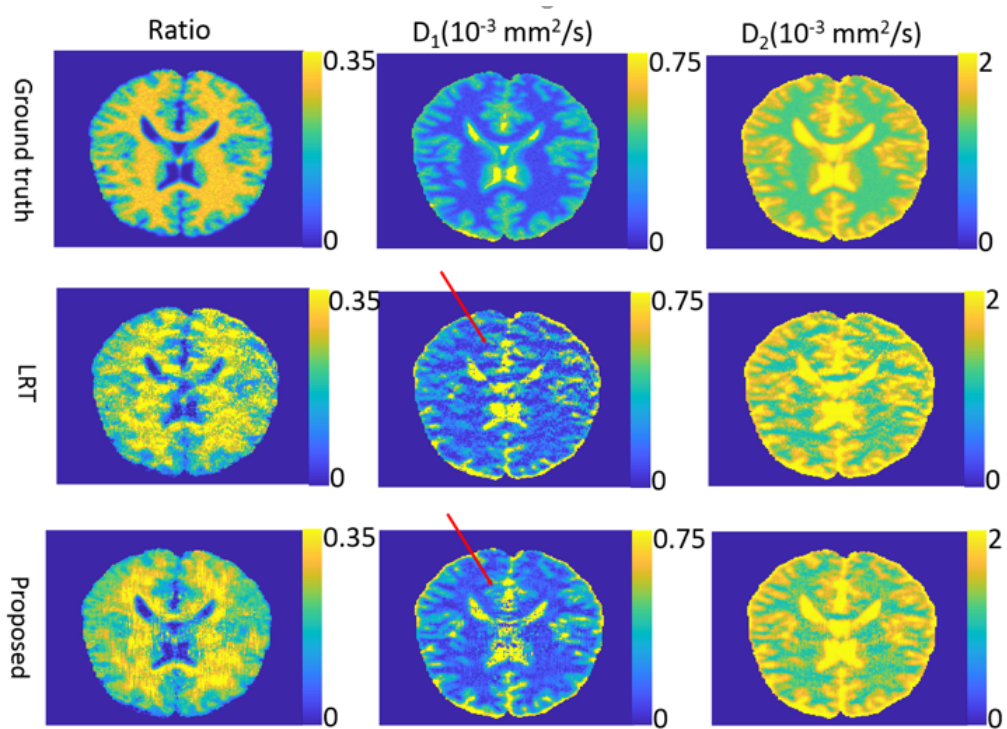
Figure 6.16: Parameter mapping results, using groundtruth, LRT reconstruction and proposed reconstruction. Aliasing artifacts (red arrow) are presented in parameter maps estimated using LRT reconstruction, and are reduced using the proposed reconstruction.
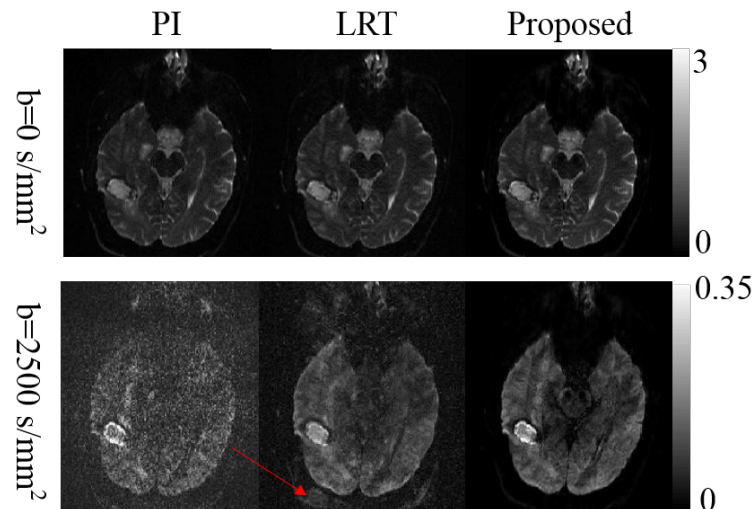
Figure 6.17: Reconstruction results from different methods using the clinical data.
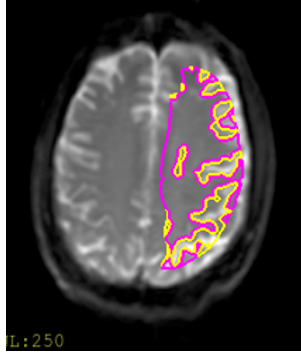
Figure 6.18: Example ROIs of white matter (magenta) and gray matter (yellow) contoured by physicians.

(yellow) contoured by physicians, overlapped with our reconstruction result. We fitted the bi-exponential model to averaged signals within the white matter and gray matter ROIs respectively. Table 6.2 summarizes the parameter mapping results using the GRAPPA reconstruction and the proposed reconstruction. The difference between the two sets of parameters is very small, which suggests that our method does not introduce systematic errors into reconstructed images.

	D1(10^{-3}mm^2/s)	D2(10^{-3}mm^2/s)	f
PI	0.39	2.09	0.48
Proposed	0.41	2.04	0.49

Table 6.2: Parameter mapping results for white matter and gray matter, using the GRAPPA reconstruction and the proposed reconstruction.

6.5   Discussion

In this chapter, we described a low-rank tensor model for accelerated high b-value diffusion-weighted MRI. The proposed model exploits both local low-rank structure and global low-rank structure simultaneously through a block-Hankel tensor. Phase constraint is further applied to correct phase variations across acquisitions and to accommodate partial Fourier acquisition.

The proposed method achieves an acceleration factor of 8, as compared to full

k-space sampling. Both reconstruction results using simulated data and patient data show improved SNR as compared to parallel imaging method and reduced aliasing as compared to other low-rank tensor based algorithms. Higher-order diffusion analysis using our reconstruction results validates the efficacy of our method in reducing aliasing.

Our method differs from previous studies using low-rank tensor model [81, 82, 83, 84] in that it proposes a new data structure that allows us to utilize dependencies across both k-space samples in a small neighborhood and all voxels across the imaging object. Therefore our method is a more natural combination of multi-coil and multi-parameter acquisition. We chose the SAKE [87] method to model local low-rankness, as it models dependencies between neighboring k-space samples from multiple coils, as compared to other local low-rank models, such as LORAKS [98]. By constructing a block-Hankel tensor, we can enforce both local and global low-rankness simultaneously, without the need to calculate a calibration kernel explicitly, which is required in other reconstruction algorithms such as SPIRiT [99]. Comparing to existing work in accelerated diffusion-weighted MRI [74, 75, 76, 77, 78], we look at a special problem of high b-value DWI for cancer imaging. The number of diffusion directions is much smaller (3 orthogonal directions only), as compared to most work in diffusion tensor imaging where more that 10 diffusion directions may be sampled. The limited diffusion directions sampled reduces the data redundancy, thus making accelerated imaging more challenging. The extended b-value range and the degraded SNR at high b-values further present challenges for accurate image reconstruction.

Future work will perform a more systematic evaluation of rank constraint selection. Different rank constraints will be tested on the simulated data to decide the optimal rank constraint for the problem. Other forms of tensor decomposition, such

as CANDECOMP decomposition [100], may also be evaluated. Decomposition of large sized tensor can be time consuming, so parallel computing [101] may be used to make the algorithm more efficient for clinical use. More complete evaluation of the method in support of higher-order diffusion analysis will also be performed. We will evaluate the proposed method on a multiple patient datasets to better determine if there is a systematic difference between the proposed reconstruction and parallel imaging-based reconstruction. The low SNR and long acquisition time make it difficult to get clinical data that is clean enough to serve as a ground truth. Phantom study may be an alternative choice for evaluation.

6.6 Conclusion

A low-rank tensor model-based image reconstruction algorithm is proposed for accelerated high b-valued DWI, that reconstructs images from sparsely sampled k-space data. By enforcing local and global low-rankness simultaneously, reconstruction results show improved signal-to-noise ratio and reduced aliasing. Higher-order diffusion analysis using the reconstructed images show accuracy comparable to the one used clinically.

CHAPTER 7

Future Work

MRI is playing an increasingly important role in radiotherapy, and MRI-guided linear accelerator (MRI-Linac) is being introduced at various medical centers. To make MRI more suitable for radiotherapy, various aspects remain to be improved, which present interesting signal processing problems to be solved.

7.1 Extended synthetic CT model for pelvic patients

MRI volumes used in our synthetic CT study were acquired for tumor and organ contouring initially and did not cover the entire treatment region. For most female pelvic patients, their treatment fields go up to thoracic spine (T9). Acquiring MRI with such a large field of view is challenging. We are currently experimenting a method called "stitching" that acquires the upper part and the lower part of the scanning volume separately with an overlap in the middle of the target. The overlap region serves as an reference for combining the upper and lower scans together. Figure 7.1 shows an example image using the stitching method. We are currently evaluating our synthetic CT generation algorithm on the new MRI dataset. With an extended scan range, the intensity distributions of tissues may be more complicated and require further investigation. Also, as we move upwards, respiratory motion may cause imaging artifacts (red arrow) and present challenges to accurate tissue
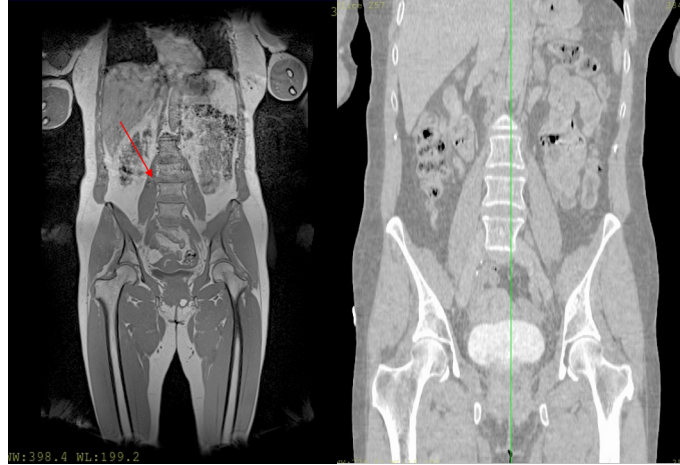
Figure 7.1: MRI volumes acquired using stitching (right) have a field of view comparable to the CT scans (left). Red arrow indicates imaging artifacts that are possibly due to respiratory motion.

classification.

We are also applying the idea of joint intensity and shape analysis for synthetic CT generation to male prostate patient. Initial results, as shown in Figure 7.2 are promising. Dosimetric analysis using the synthetic CT will further evaluate its usefulness in support of MRI-based treatment planning. Other aspects of MRI-guided treatment planning, such as prostate targeting also require future study. Prostate position is usually determined on CT images by inserting gold seeds into the prostate. Gold seeds will cause unique artifacts on CT images which can be used to locate the prostate. However, gold seeds on MRI are mostly dark and difficult to be distinguished from air pockets. Shape information of the gold seeds may be used to improve its identification on MRI [102]. Other fiducial markers, other than gold seeds that are more easily identified on MRI may be an alternative solution.

Current synthetic CT generation algorithm requires identifying several landmarks manually on the pelvic bone, to initialize the alignment of the atlas image to a target image. To make the algorithm more suitable for clinical use, we want to minimize the amount of manual input. Automatic landmark detection, such as femur heads
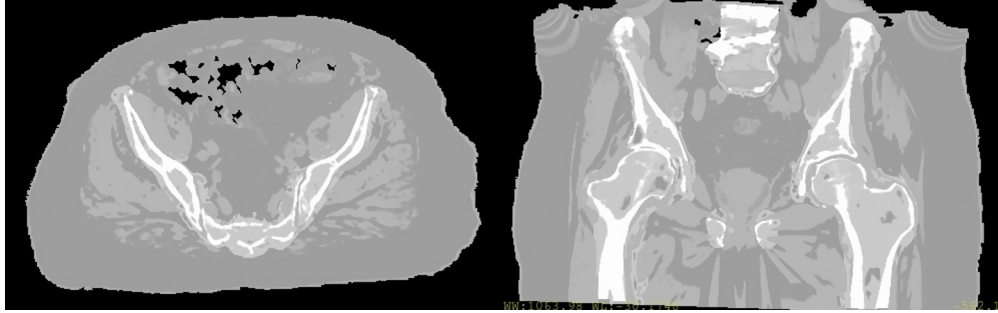
Figure 7.2: Synthetic CT generated for male prostate patients

identification, may be an interesting problem to look at.

7.2 Improve low-rank tensor model for high b-value DWI reconstruction

Tensor decomposition is an area of active research, and various algorithms have been proposed. For example, [103] proposed a robust tensor decomposition algorithm that may be of interest to our problem, given the low SNR at high b-values. The algorithm is designed based on the assumptions of linear dependencies between k-space samples, future work will explore if non-linear embedding will help to improve results. Parallelizing the algorithm [101] will also be of interest given the large size of the tensor we built. The proposed algorithm corrects phase variations between acquisitions, which is a key component in multi-shot EPI [104]. Multi-shot EPI divides the k-space trajectory into several segments (shots) along the readout direction. In this way, TE can be reduced, which reduces geometric distortion and allows acquisitions of high-resolution DWI volumes. Future work will test this algorithm on multi-shot EPI data. A 2D navigator signal is acquired in multi-shot EPI for phase correction [105]. Phase maps estimated from the navigator signal can be easily incorporated into our problem formulation for image reconstruction. With single-shot EPI, random sampling is only feasible along the phase-encoding direction. With multi-shot EPI, random sampling can be performed along two dimensions, and sam-

pling patterns between different shots may also vary, thus allowing a more incoherent sampling, which may provide additional space for acceleration.

7.3 Direct diffusion parameter mapping from sparsely sampled k-space data

Skipping the image reconstruction step and directly estimating parameter maps from k-space samples is a potentially attractive alternative formulation. It will also allow the use of signal models as a prior knowledge to regularize the iterative process of estimating missing k-space samples. The major challenge for this approach is the complexity of the signal model. Firstly, the phase variations between acquisitions are due to random motions and are difficult to model. Secondly, the diffusion properties of tissues under high b-values remains an open problem. It is commonly acknowledged that a multi-exponential decay is demonstrated, yet different exponential decay curves are not linearly independent, making accurate fitting to a multi-exponential model challenging. In [106], the author proposed to regularize the fitting procedure with a neighborhood constraint. However, the neighborhood constraint only enforces the smoothness over fractions of different exponential decay components, and overlooks the smoothness of maps of diffusion coefficients, while neighboring voxels are likely to have similar diffusion coefficients. Our initial studies suggest that better regularization is needed for accurate multi-exponential fitting, given the low SNR and limited resolution of high b-value DWI.

BIBLIOGRAPHY

[1] CancerResearchUK. What radiotherapy is. `http://www.cancerresearchuk.org/about-cancer/cancers-in-general/treatment/radiotherapy/about/what-radiotherapy-is`.

[2] CancerResearchUK. What image guided radiotherapy is. `http://www.cancerresearchuk.org/about-cancer/cancers-in-general/treatment/radiotherapy/external/image-guided-radiotherapy-igrt`.

[3] L.A. Dawson and D.A. Jaffray. Advances in image-guided radiation therapy. *Journal of clinical oncology*, 25(8):938–946, 2007.

[4] S. Hsu, Y. Cao, K. Huang, M. Feng, and J.M. Balter. Investigation of a method for generating synthetic CT models from MRI scans of the head and neck for radiation therapy. *Physics in medicine and biology*, 58(23):8419, 2013.

[5] D. Wang, W. Strugnell, G. Cowin, D.M. Doddrell, and R. Slaughter. Geometric distortion in clinical MRI systems: Part I: evaluation using a 3D phantom. *Magnetic resonance imaging*, 22(9):1211–1221, 2004.

[6] J. Wang, M. Qiu, Q.X. Yang, M.B. Smith, and T.R. Constable. Measurement and correction of transmitter and receiver induced nonuniformities in vivo. *Magnetic resonance in medicine*, 53(2):408–417, 2005.

[7] L. Liu, S. Hsu, J.M. Balter, and Y. Cao. A joint classification and bias field estimation algorithm for synthetic CT generation. In *2014 MR in RT symposium*, 2014.

[8] L. Liu, Y. Cao, J.A. Fessler, S. Jolly, and J.M. Balter. A female pelvic bone shape model for air/bone separation in support of synthetic CT generation for radiation therapy. *Physics in medicine and biology*, 61(1):169, 2015.

[9] L Liu, Y Cao, J Fessler, and J Balter. SU-F-303-14: Investigation of a pelvic bone shape model in support of bone classification for synthetic CT generation. *Medical physics*, 42(6):3541–3541, 2015.

[10] L. Liu, S. Jolly, Y. Cao, K. Vineberg, J.A. Fessler, and J.M. Balter. Female pelvic synthetic CT generation based on joint intensity and shape analysis. *Physics in Medicine & Biology*, 62(8):2935, 2017.

[11] L. Liu, S. Jolly, Y. Cao, K. Vineberg, J. Fessler, and J. Balter. SU-D-207A-01: Female pelvic synthetic CT generation based on joint shape and intensity analysis. *Medical Physics*, 43(6):3343–3343, 2016.

[12] C. Tsien, Y. Cao, and T. Chenevert. Clinical applications for diffusion magnetic resonance imaging in radiotherapy. In *Seminars in radiation oncology*, volume 24, pages 218–226. Elsevier, 2014.

[13] P.A. Hein, C.J. Eskey, J.F. Dunn, and E.B. Hug. Diffusion-weighted imaging in the follow-up of treated high-grade gliomas: tumor recurrence versus radiation injury. *American Journal of Neuroradiology*, 25(2):201–209, 2004.

[14] Y. Mardor, R. Pfeffer, R. Spiegelmann, Y. Roth, S.E. Maierand, O. Nissim, R. Berger, A. Glicksman, J.Baram, A. Orenstein, et al. Early detection of response to radiation therapy in patients with brain malignancies using conventional and high b-value diffusion-weighted magnetic resonance imaging. *Journal of clinical oncology*, 21(6):1094–1100, 2003.

[15] P.P. Pramanik, H.A. Parmar, A.G. Mammoser, L.R. Junck, M.M. Kim, C.I. Tsien, T.S. Lawrence, and Y. Cao. Hypercellularity components of glioblastoma identified by high b-value diffusion-weighted imaging. *International Journal of Radiation Oncology* Biology* Physics*, 92(4):811–819, 2015.

[16] J.H. Burdette, D.D. Durden, A.D. Elster, and Y. Yen. High b-value diffusion-weighted MRI of normal brain. *Journal of computer assisted tomography*, 25(4):515–519, 2001.

[17] L. Liu, A. Johansson, J.M. Balter, J.A. Fessler, and Y. Cao. Accelerated high b-value diffusion-weighted MRI for higher-order diffusion analysis using a phase-constrained low-rank tensor model. In *Proceedings of the 26th Annual Meeting of ISMRM, Paris, France*, 2018.

[18] L. Liu, A. Johansson, J.M. Balter, Y. Cao, and J.A. Fessler. Accelerated high b-value diffusion-weighted MR imaging via phase-constrained low-rank tensor model. In *Proceedings of 2018 IEEE International Symposium on Biomedical Imaging*, 2018.

[19] Y. Long. *Statistical Image Reconstruction and Motion Estimation for Image-Guided Radiotherapy*. PhD thesis, The University of Michigan, 2011.

[20] H. Sun. *Topics in steady-state MRI sequences and RF pulse optimization*. PhD thesis, The University of Michigan, 2015.

[21] M.D. Robson, P.D. Gatehouse, M. Bydder, and G.M. Bydder. Magnetic resonance: an introduction to ultrashort TE (UTE) imaging. *Journal of computer assisted tomography*, 27(6):825–846, 2003.

[22] D.M. Grodzki, P.M. Jakob, and Heismann B. Ultrashort echo time imaging using pointwise encoding time reduction with radial acquisition (PETRA). *Magnetic resonance in medicine*, 67(2):510–518, 2012.

[23] J. Crank. *The mathematics of diffusion*. Oxford university press, 1979.

[24] T.L. Chenevert. Principles of diffusion-weighted imaging (DW-MRI) as applied to body imaging. In *Diffusion-Weighted MR Imaging*, pages 3–17. Springer, 2010.

[25] A. Szafer, J. Zhong, and J.C. Gore. Theoretical model for water diffusion in tissues. *Magnetic resonance in medicine*, 33(5):697–712, 1995.

[26] D. Yang, Y. Korogi, T. Sugahara, Mika. Kitajima, Y. Shigematsu, L. Liang, Y. Ushio, and M. Takahashi. Cerebral gliomas: prospective comparison of multivoxel 2D chemical-shift imaging proton MR spectroscopy, echoplanar perfusion and diffusion-weighted MRI. *Neuroradiology*, 44(8):656–666, 2002.

[27] D. Le Bihan, E. Breton, D. Lallemand, P. Grenier, E. Cabanis, and M. Laval-Jeantet. MR imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders. *Radiology*, 161(2):401–407, 1986.

[28] T.L. Chenevert, J.A. Brunberg, and J.G. Pipe. Anisotropic diffusion in human white matter: demonstration with MR techniques in vivo. *Radiology*, 177(2):401–405, 1990.

[29] L.B. Denis, J-F. Mangin, C. Poupon, C.A. Clark, S. Pappata, N. Molko, and H. Chabriat. Diffusion tensor imaging: concepts and applications. *Journal of magnetic resonance imaging*, 13(4):534–546, 2001.

[30] A.L. Alexander, K. Hasan, G. Kindlmann, D.L. Parker, and J.S. Tsuruda. A geometric analysis of diffusion tensor measurements of the human brain. *Magnetic Resonance in Medicine*, 44(2):283–291, 2000.

[31] D. Le Bihan, E. Breton, D. Lallemand, M.L. Aubin, J. Vignaud, and M. Laval-Jeantet. Separation of diffusion and perfusion in intravoxel incoherent motion MR imaging. *Radiology*, 168(2):497–505, 1988.

[32] R. Turner, D. Le Bihan, and A. Scott Chesnicks. Echo-planar imaging of diffusion and perfusion. *Magnetic resonance in medicine*, 19(2):247–253, 1991.

[33] C. Siversson, F. Nordström, T. Nilsson, T. Nyholm, J. Jonsson, A. Gunnlaugsson, and L.E. Olsson. Technical note: MRI only prostate radiotherapy planning using the statistical decomposition algorithm. *Medical physics*, 42(10):6090–6097, 2015.

[34] J. Uh, T.E. Merchant, Y. Li, X. Li, and C. Hua. MRI-based treatment planning with pseudo CT generated through atlas registration. *Medical physics*, 41(5):051711, 2014.

[35] J.A. Dowling, J. Lambert, J. Parker, O. Salvado, J. Fripp, A. Capp, C. Wratten, J.W. Denham, and P.B. Greer. An atlas-based electron density mapping method for magnetic resonance imaging (MRI)-alone treatment planning and adaptive MRI-based prostate radiation therapy. *International Journal of Radiation Oncology* Biology* Physics*, 83(1):e5–e11, 2012.

[36] P.B. Greer, J.A. Dowling, J.A. Lambert, J. Fripp, J. Parker, J.W. Denham, C. Wratten, A. Capp, and O. Salvado. A magnetic resonance imaging-based workflow for planning radiation therapy for prostate cancer. *Medical Journal of Australia*, 194(4):S24, 2011.

[37] J.H. Kim, J.S. Lee, I. Song, and D.S. Lee. Comparison of segmentation-based attenuation correction methods for PET/MRI: evaluation of bone and liver standardized uptake value with oncologic pet/ct data. *Journal of Nuclear Medicine*, 53(12):1878–1882, 2012.

[38] A. Johansson, M. Karlsson, and T. Nyholm. Ct substitute derived from MRI sequences with ultrashort echo time. *Medical physics*, 38(5):2708–2714, 2011.

[39] S. Hsu, Y. Cao, T.S. Lawrence, C. Tsien, M. Feng, D.M. Grodzki, and J.M. Balter. Quantitative characterizations of ultrashort echo (UTE) images for supporting air–bone separation in the head. *Physics in medicine and biology*, 60(7):2869, 2015.

[40] J. Chiou, C.B. Ahn, L.T. Muftuler, and O. Nalcioglu. A simple simultaneous geometric and intensity correction method for echo-planar imaging by EPI-based phase modulation. *IEEE transactions on medical imaging*, 22(2):200–205, 2003.

[41] H. Watanabe, N. Takaya, and F. Mitsumori. Non-uniformity correction of human brain imaging at high field by RF field mapping of B1+ and B1. *Journal of Magnetic Resonance*, 212(2):426–430, 2011.

[42] U. Vovk, F. Pernus, and B. Likar. A review of methods for correction of intensity inhomogeneity in MRI. *IEEE transactions on medical imaging*, 26(3):405–421, 2007.

[43] J.G. Sled, A.P. Zijdenbos, and A.C. Evans. A nonparametric method for automatic correction of intensity nonuniformity in MRI data. *IEEE transactions on medical imaging*, 17(1):87–97, 1998.

[44] N.J. Tustison, B.B. Avants, P.A. Cook, Y. Zheng, A. Egan, P.A. Yushkevich, and J.C. Gee. N4ITK: improved N3 bias correction. *IEEE transactions on medical imaging*, 29(6):1310–1320, 2010.

[45] W. Chen and M.L. Giger. A fuzzy c-means (fcm) based algorithm for intensity inhomogeneity correction and segmentation of MR images. In *Biomedical Imaging: Nano to Macro, 2004. IEEE International Symposium on*, pages 1307–1310. IEEE, 2004.

[46] C. Li, C. Xu, A.W. Anderson, and J.C. Gore. MRI tissue classification and bias field estimation based on coherent local intensity clustering: A unified energy minimization framework. In *International Conference on Information Processing in Medical Imaging*, pages 288–299. Springer, 2009.

[47] A. Liew and H. Yan. An adaptive spatial fuzzy clustering algorithm for 3-D MR image segmentation. *IEEE transactions on medical imaging*, 22(9):1063–1075, 2003.

[48] M.N. Ahmed, S.M. Yamany, N. Mohamed, A.A. Farag, and T. Moriarty. A modified fuzzy c-means algorithm for bias field estimation and segmentation of MRI data. *IEEE transactions on medical imaging*, 21(3):193–199, 2002.

[49] M.J. Allison, S. Ramani, and J.A. Fessler. Accelerated regularized estimation of MR coil sensitivities using augmented lagrangian methods. *IEEE transactions on medical imaging*, 32(3):556–564, 2013.

[50] H. Nien and J.A. Fessler. Fast x-ray CT image reconstruction using a linearized augmented lagrangian method with ordered subsets. *IEEE transactions on medical imaging*, 34(2):388–399, 2015.

[51] S. Lee, G. Wolberg, and S. Shin. Scattered data interpolation with multilevel B-splines. *IEEE transactions on visualization and computer graphics*, 3(3):228–244, 1997.

[52] J. Kim, C. Glide-Hurst, A. Doemer, N. Wen, B. Movsas, and I.J. Chetty. Implementation of a novel algorithm for generating synthetic CT images from magnetic resonance imaging data sets for prostate cancer radiation therapy. *International Journal of Radiation Oncology* Biology* Physics*, 91(1):39–47, 2015.

[53] J. Korhonen, M. Kapanen, J. Keyriläinen, T. Seppälä, and M. Tenhunen. A dual model HU conversion from MRI intensity values within and outside of bone segment for MRI-based radiotherapy treatment planning of prostate cancer. *Medical physics*, 41(1):011704, 2014.

[54] M. Kapanen and M. Tenhunen. T1/T2*-weighted MRI provides clinically relevant pseudo-CT density data for the pelvic bones in MRI-only based radiotherapy treatment planning. *Acta Oncologica*, 52(3):612–618, 2013.

[55] L. Chen, T. Nguyen, É. Jones, Z. Chen, W. Luo, L. Wang, R.A. Price, A. Pollack, and C. Ma. Magnetic resonance–based treatment planning for prostate intensity-modulated radiotherapy: creation of digitally reconstructed radiographs. *International Journal of Radiation Oncology* Biology* Physics*, 68(3):903–911, 2007.

[56] T.F. Cootes, C.J. Taylor, D.H. Cooper, and J. Graham. Active shape models-their training and application. *Computer vision and image understanding*, 61(1):38–59, 1995.

[57] J. Schmid and N. Magnenat-Thalmann. MRI bone segmentation using deformable models and shape priors. In *International conference on medical image computing and computer-assisted intervention*, pages 119–126. Springer, 2008.

[58] A. Tsai, A. Yezzi, W. Wells, C. Tempany, D. Tucker, A. Fan, W.E. Grimson, and A. Willsky. A shape-based approach to the segmentation of medical imagery using level sets. *IEEE transactions on medical imaging*, 22(2):137–154, 2003.

[59] M.E. Leventon, W.E.L. Grimson, and O. Faugeras. Statistical shape influence in geodesic active contours. In *Computer Vision and Pattern Recognition, 2000. Proceedings. IEEE Conference on*, volume 1, pages 316–323. IEEE, 2000.

[60] Q. Gao, P. Chang, D. Rueckert, S. M. Ali, D. Cohen, P. Pratt, E. Mayer, G. Yang, A. Darzi, et al. Modeling of the bony pelvis from MRI using a multi-atlas AE-SDM for registration and tracking in image-guided robotic prostatectomy. *Computerized Medical Imaging and Graphics*, 37(2):183–194, 2013.

[61] S. Thompson, G. Penney, D. Buie, P. Dasgupta, and D. Hawkes. Use of a CT statistical deformation model for multi-modal pelvic bone segmentation. In *Medical Imaging*, page 69141O. International Society for Optics and Photonics, 2008.

[62] G. Chintalapani, L.M. Ellingsen, O. Sadowsky, J.L. Prince, and R.H. Taylor. Statistical atlases of bone anatomy: construction, iterative improvement and validation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 499–506. Springer, 2007.

[63] J.C. Noordam, W. Van Den Broek, and L.M.C. Buydens. Multivariate image segmentation with cluster size insensitive fuzzy c-means. *Chemometrics and intelligent laboratory systems*, 64(1):65–78, 2002.

[64] E. Paradis, Y. Cao, T.S. Lawrence, C. Tsien, M. Feng, K. Vineberg, and J.M. Balter. Assessing the dosimetric accuracy of magnetic resonance-generated synthetic CT images for focal brain VMAT radiation therapy. *International Journal of Radiation Oncology* Biology* Physics*, 93(5):1154–1161, 2015.

[65] J.S. Bredfeldt, L. Liu, M. Feng, Y. Cao, and J.M. Balter. TU-AB-BRA-01: Abdominal synthetic CT generation in support of liver SBRT dose calculation. *Medical Physics*, 43(6):3733–3733, 2016.

[66] M. Gudur, W. Hara, Q. Le, L. Wang, L. Xing, and R. Li. A unifying probabilistic bayesian approach to derive electron density from MRI for radiation therapy treatment planning. *Physics in medicine and biology*, 59(21):6595, 2014.

[67] D.H. Paulus, H.H. Quick, C. Geppert, M. Fenchel, Y. Zhan, G. Hermosillo, D. Faul, F. Boada, K.P. Friedman, and T. Koesters. Whole-body PET/MR imaging: quantitative evaluation of a novel model-based MR attenuation correction method including bone. *Journal of Nuclear Medicine*, 56(7):1061–1066, 2015.

[68] T. Sugahara, Y. Korogi, M. Kochi, I. Ikushima, Y. Shigematu, T. Hirai, T. Okuda, L. Liang, Y. Ge, Y. Komohara, et al. Usefulness of diffusion-weighted MRI with echo-planar technique in the evaluation of cellularity in gliomas. *Journal of Magnetic Resonance Imaging*, 9(1):53–60, 1999.

[69] C. Kim, B. Park, and B. Kim. High-b-value diffusion-weighted imaging at 3T to detect prostate cancer: comparisons between b values of 1,000 and 2,000 s/mm2. *American Journal of Roentgenology*, 194(1):W33–W37, 2010.

[70] H.S. Seo, K.H. Chang, D.G. Na, B.J. Kwon, and D.H. Lee. High b-value diffusion (b= 3000 s/mm2) MR imaging in cerebral gliomas at 3T: visual and quantitative comparisons with b= 1000 s/mm2. *American Journal of Neuroradiology*, 29(3):458–463, 2008.

[71] S.C. Grant, D.L. Buckley, S. Gibbs, A.G. Webb, and S.J. Blackband. MR microscopy of multicomponent diffusion in single neurons. *Magnetic resonance in medicine*, 46(6):1107–1112, 2001.

[72] T. Niendorf, R.M. Dijkhuizen, D.G. Norris, M. van Lookeren Campagne, and K. Nicolay. Biexponential diffusion attenuation in various states of brain tissue: Implications for diffusion-weighted imaging. *Magnetic Resonance in Medicine*, 36(6):847–857, 1996.

[73] K.M. Bennett, K.M. Schmainda, D.B. Rowe, H. Lu, J.S. Hyde, et al. Characterization of continuously distributed cortical water diffusion rates with a stretched-exponential model. *Magnetic resonance in medicine*, 50(4):727–734, 2003.

[74] C. Liao, Y. Chen, X. Cao, S. Chen, H. He, M. Mani, M. Jacob, V. Magnotta, and J. Zhong. Efficient parallel reconstruction for high resolution multishot spiral diffusion data with low rank constraint. *Magnetic resonance in medicine*, 77(3):1359–1366, 2017.

[75] M. Mani, M. Jacob, A. Guidon, V. Magnotta, and J. Zhong. Acceleration of high angular and spatial resolution diffusion imaging using compressed sensing with multichannel spiral data. *Magnetic resonance in medicine*, 73(1):126–138, 2015.

[76] J.P. Haldar and K. Setsompop. Fast high-resolution diffusion MRI using gSlider-SMS, interlaced subsampling, and SNR-enhancing joint reconstruction. In *Proceedings of the 25th Annual Meeting of ISMRM, Honolulu, HI, USA*, 2017.

[77] H. Gao, L. Li, K. Zhang, W. Zhou, and X. Hu. PCLR: Phase-constrained low-rank model for compressive diffusion-weighted MRI. *Magnetic resonance in medicine*, 72(5):1330–1341, 2014.

[78] X. Shi, X. Ma, W. Wu, F. Huang, C. Yuan, and H. Guo. Parallel imaging and compressed sensing combined framework for accelerating high-resolution diffusion tensor imaging using inter-image correlation. *Magnetic resonance in medicine*, 73(5):1775–1785, 2015.

[79] T.C. Kwee, C.J. Galbán, C. Tsien, L. Junck, P.C. Sundgren, M.K. Ivancevic, T.D. Johnson, C.R. Meyer, A. Rehemtulla, and B.D. Ross. Comparison of apparent diffusion coefficients and distributed diffusion coefficients in high-grade gliomas. *Journal of Magnetic Resonance Imaging*, 31(3):531–537, 2010.

[80] B.A. Hoff, T.L. Chenevert, M.S. Bhojani, T.C. Kwee, A. Rehemtulla, D. Le Bihan, B.D. Ross, and C.J. Galbán. Assessment of multiexponential diffusion features as MRI cancer therapy response metrics. *Magnetic resonance in medicine*, 64(5):1499–1509, 2010.

[81] S.F. Roohi, D. Zonoobi, A.A. Kassimand, and J.L. Jaremko. Dynamic MRI reconstruction using low rank plus sparse tensor decomposition. In *Proc. ICIP*, pages 1769–1773. IEEE, 2016.

[82] D. Banco, S. Aeron, and W.S. Hoge. Sampling and recovery of MRI data using low rank tensor models. In *Proc. EMBC*, pages 448–452. IEEE, 2016.

[83] J. He, Q. Liu, A. Christodoulou, C. Ma, F. Lam, and Z. Liang. Accelerated high-dimensional MR imaging with sparse sampling using low-rank tensors. *IEEE transactions on medical imaging*, 35(9):2119–2129, 2016.

[84] J.D. Trzasko and A. Manduca. A unified tensor regression framework for calibrationless dynamic, multi-channel MRI reconstruction. In *Proceedings of the 21st Annual Meeting of ISMRM, Salt Lake City, Utah, USA*, pages 603–603, 2013.

[85] Y. Liu, J. Cao, M. Lyu, and E.W. Wu. Calibrationless parallel imaging reconstruction using hankel tensor completion (HTC). In *Proceedings of the 25th Annual Meeting of ISMRM, Honolulu, HI, USA*, 2017.

[86] M.A. Griswold, P.M. Jakob, R.M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magnetic resonance in medicine*, 47(6):1202–1210, 2002.

[87] P.J. Shin, P.E. Larson, M.A. Ohliger, M. Elad, J.M. Pauly, D.B. Vigneron, and M. Lustig. Calibrationless parallel imaging reconstruction based on structured low-rank matrix completion. *Magnetic resonance in medicine*, 72(4):959–970, 2014.

[88] B. Zhao, W. Lu, T.K. Hitchens, F. Lam, C. Ho, and Z. Liang. Accelerated MR parameter mapping with low-rank and sparsity constraints. *Magnetic resonance in medicine*, 74(2):489–498, 2015.

[89] E.M. Haacke, E.D. Lindskogj, and W. Lin. A fast, iterative, partial-Fourier technique capable of local phase recovery. *Journal of Magnetic Resonance*, 92(1):126–145, 1991.

[90] L. De Lathauwer, B. De Moor, and J. Vandewalle. A multilinear singular value decomposition. *SIAM journal on Matrix Analysis and Applications*, 21(4):1253–1278, 2000.

[91] N. Vannieuwenhoven, R. Vandebril, and K. Meerbergen. A new truncation strategy for the higher-order singular value decomposition. *SIAM Journal on Scientific Computing*, 34(2):A1027–A1052, 2012.

[92] N. Halko, P. Martinsson, and J.A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.

[93] N. Vervliet, O. Debals, L. Sorber, M. Van Barel, and L. De Lathauwer. Tensorlab 3.0, Mar. 2016. Available online.

[94] D.L. Collins, A.P. Zijdenbos, V. Kollokian, J.G. Sled, N.J. Kabani, C.J. Holmes, and A.C. Evans. Design and construction of a realistic digital brain phantom. *IEEE transactions on medical imaging*, 17(3):463–468, 1998.

[95] R.V. Mulkern, H.P. Zengingonul, R.L. Robertson, P. Bogner, K.H. Zou, H. Gudbjartsson, C.R. Guttmann, D. Holtzman, W. Kyriakos, and F.A. Jolesz. Multi-component apparent diffusion coefficients in human brain: relationship to spin-lattice relaxation. *Magnetic resonance in medicine*, 44(2):292–300, 2000.

[96] A.T. Van, D.C. Karampinos, J.G. Georgiadis, and B.P. Sutton. k-space and image space combination for motion artifact correction in multicoil multishot diffusion weighted imaging. In *Proc. EMBC*, pages 1675–1678. IEEE, 2008.

[97] S. Mehrotra. On the implementation of a primal-dual interior point method. *SIAM Journal on optimization*, 2(4):575–601, 1992.

[98] J.P. Haldar. Low-rank modeling of local k-space neighborhoods (LORAKS) for constrained MRI. *IEEE transactions on medical imaging*, 33(3):668–681, 2014.

[99] M. Murphy, M. Alley, J. Demmel, K. Keutzer, S. Vasanawala, and M. Lustig. Fast l1-spirit compressed sensing parallel imaging MRI: Scalable parallel implementation and clinically feasible runtime. *IEEE transactions on medical imaging*, 31(6):1250–1262, 2012.

[100] J.D. Carroll and J. Chang. Analysis of individual differences in multidimensional scaling via an N-way generalization of Eckart-Young decomposition. *Psychometrika*, 35(3):283–319, 1970.

[101] W. Austin, G. Ballard, and T.G. Kolda. Parallel tensor compression for large-scale scientific data. In *Parallel and Distributed Processing Symposium, 2016 IEEE International*, pages 912–922. IEEE, 2016.

[102] S. Ghose, J. Mitra, D. Rivest-Hénault, A. Fazlollahi, P. Stanwell, P. Pichler, J. Sun, J. Fripp, P.B. Greer, and J.A. Dowling. MRI-alone radiation therapy planning for prostate cancer: Automatic fiducial marker detection. *Medical physics*, 43(5):2218–2228, 2016.

[103] D. Goldfarb and Z. Qin. Robust low-rank tensor recovery: Models and algorithms. *SIAM Journal on Matrix Analysis and Applications*, 35(1):225–253, 2014.

[104] C. Liu, M.E. Moseley, and R. Bammer. Simultaneous phase correction and SENSE reconstruction for navigated multi-shot DWI with non-cartesian k-space sampling. *Magnetic resonance in medicine*, 54(6):1412–1422, 2005.

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

[106] D. Kim, E.K. Doyle, J.L. Wisnowski, J.H. Kim, and J.P. Haldar. Diffusion-relaxation correlation spectroscopic imaging: A multidimensional approach for probing microstructure. *Magnetic resonance in medicine*, 78(6):2236–2249, 2017.