

# **Motion Binning in Cone-Beam CT Without a Prior Model**

**by**

*Tracy Petrie*

**Submitted in accordance with the requirements  
for the degree of Doctor of Philosophy.**



**UNIVERSITY OF LEEDS**

**The University of Leeds  
School of Computing**

**November 2010**

**The candidate confirms that the work submitted is his own and that the appropriate credit has been given where reference has been made to the work of others.**

**This copy has been supplied on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement.**

*To Twila Petrie  
my most beloved companion  
on this grand adventure*

# Acknowledgements

My thanks must first go to Dr. Derek Magee, a supervisor whose creativity and knowledge is only exceeded by his unending patience. Working through challenges and brainstorming new ideas in his office was always an intensely exciting time of collaboration. More than anything else, I will miss that engagement.

Jonathan Sykes was also tremendously helpful and encouraging and I owe him much. His assistance in helping me acquire the data I needed from the CBCT machine at St. James's University Hospital was invaluable as was his patient education of a non-physicist on matters related to medical physics.

Doctors Di Gilson and Kevin Franks graciously provided me time to discuss the clinical aspects of cancer treatment in the UK and radiotherapy in particular. Amy Wooley and Helen Summers kindly let me watch them administer radiotherapy and helped me understand more about patient motion issues and the tools that are used in every day treatment.

The power of friendship was made more clear to me than ever on this project. New friends that I've made here and old friends back home have continually wished me success and cheered me on across the finish line. Two people I not only call friends but family, my sister Kim and my mother-in-law Trudy, deserve special mention.

Finally, and in the position of honour, no words can adequately express my gratitude for my wonderful wife. Without her constant encouragement, support, and occasional coaching, I could not have made it.

# Abstract

Cone-beam CT (CBCT), used to reduce setup error in radiotherapy, takes a sequence of about 670 two-dimensional X-rays acquired in a circular arc around the patient over two minutes and reconstructs a three-dimensional volume from these projections. Consequently, when tissues move significantly during the acquisition the resulting volume is blurry or contains streaks. The projections themselves, though, are sharp. One of the main areas of research with CBCT has been to attempt to reconstruct the motion from these projections by collecting them into respiratory-phase or amplitude bins and using these to reconstruct bin-specific volumes. A variety of mechanisms are employed to identify and record respiratory motion so that it can be correlated with the projections.

Not all motion that occurs in the body can be correlated with respiration. The research question pursued in this thesis is whether motion can be identified in a binning process without prior knowledge or models of the motion. Nomenclature describing motion classes and a specific type of artefact are introduced. The distinctiveness of this artefact class is demonstrated and methods to mitigate it are proposed and evaluated. Several techniques are then used to reduce an intractable search space to a computationally feasible one. A unique application of PCA to the reconstruction process allows new kinds of search approaches to be considered including an exhaustive search which requires a protocol change and a multiple-restart hill-climbing search that can be used with existing protocols. Experiments with three classes of phantoms, including a novel animated physical phantom, show the effectiveness of the two search methods which are finally compared with each other.

# Declarations

Some parts of the work presented in this thesis have been published in the following articles:

**Petrie, T., Magee, D. and Sykes, J.**, “Modeless Motion Binning in Cone Beam CT”, *International Symposium on Biomedical Imaging (ISBI)*, Rotterdam – The Netherlands, 2010.

**Petrie, T., Magee, D. and Sykes, J.**, “Binning Without a Model for Cone-Beam CT”, *Medical Image Understanding and Analysis (MIUA)*, Warwick – UK, 2010.

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Cancer, imaging, and computer science . . . . .	1
1.2	Thesis overview . . . . .	3
<b>2</b>	<b>Background</b>	<b>5</b>
2.1	Introduction . . . . .	5
2.2	Ionization of tissue . . . . .	6
2.2.1	Relevant physics . . . . .	6
2.2.2	Key aspects of radiobiology . . . . .	9
2.2.3	In summary . . . . .	10
2.3	X-ray imaging . . . . .	10
2.3.1	Radiographs . . . . .	11
2.3.2	Portal images . . . . .	11
2.3.3	Computed Tomography (CT) . . . . .	11
2.4	Physiological motion . . . . .	18
2.4.1	Motion in the lungs . . . . .	19
2.4.2	Motion in the liver . . . . .	19
2.4.3	Motion in the prostate . . . . .	20
2.4.4	In summary . . . . .	20
2.5	Radiotherapy . . . . .	21
2.5.1	Diagnosis . . . . .	21
2.5.2	Planning . . . . .	21
2.5.3	Therapy . . . . .	23
2.5.4	In summary . . . . .	25
2.6	Motion compensation . . . . .	25
2.6.1	Continuously periodic motion compensation: IMRT . . . . .	27
2.6.2	Shift motion compensation . . . . .	32
2.6.3	Discrete motion compensation: IGRT . . . . .	32

2.6.4	In summary . . . . .	33
2.7	Chapter summary and thesis problem . . . . .	34
<b>3</b>	<b>Materials</b>	<b>36</b>
3.1	Introduction . . . . .	36
3.2	CBCT geometries using the Elekta Synergy . . . . .	37
3.3	Simple virtual phantoms . . . . .	38
3.3.1	Description . . . . .	39
3.3.2	Projection construction . . . . .	39
3.3.3	Simple virtual phantom configurations . . . . .	40
3.4	Animated physical phantoms . . . . .	42
3.4.1	Description . . . . .	43
3.4.2	Acquisition details . . . . .	43
3.4.3	Construction . . . . .	43
3.4.4	Animated physical phantom configurations . . . . .	44
3.5	NCAT anthropomorphic phantoms . . . . .	46
3.5.1	Description . . . . .	46
3.5.2	Construction . . . . .	47
3.5.3	NCAT phantom configurations . . . . .	47
3.6	Summary . . . . .	48
<b>4</b>	<b>Projection Gaps in Reconstructions</b>	<b>49</b>
4.1	Introduction . . . . .	49
4.2	Artefacts in reconstruction . . . . .	50
4.2.1	Physics-based artefacts . . . . .	50
4.2.2	Patient-based artefacts . . . . .	51
4.2.3	Scanner-based artefacts . . . . .	51
4.2.4	Artefacts discussion . . . . .	51
4.3	Gap artefacts . . . . .	53
4.3.1	Experiments . . . . .	54
4.3.2	Results and discussion . . . . .	61
4.3.3	In summary . . . . .	68
<b>5</b>	<b>Filling Projection Gaps in CBCT</b>	<b>69</b>
5.1	Introduction . . . . .	69
5.2	Filling gaps using one scan . . . . .	72
5.2.1	Method . . . . .	72

5.2.2	Results of one-scan gap filling . . . . .	73
5.3	Filling gaps using two scans . . . . .	78
5.3.1	Method . . . . .	80
5.3.2	Results of two-scan gap filling . . . . .	82
5.4	Conclusions . . . . .	88
<b>6</b>	<b>Two Scan Searching</b>	<b>89</b>
6.1	Introduction . . . . .	89
6.2	Defining an assignment search space . . . . .	90
6.2.1	Search method . . . . .	92
6.2.2	Reducing the search space and computation cost . . . . .	93
6.3	Evaluation and results . . . . .	111
6.3.1	Using a ground-truth derived difference signal . . . . .	111
6.3.2	Using an estimated difference signal . . . . .	124
6.4	Conclusions . . . . .	134
<b>7</b>	<b>One Scan Searching</b>	<b>136</b>
7.1	Overview . . . . .	136
7.2	Method . . . . .	137
7.2.1	The search space . . . . .	137
7.2.2	Forward projections . . . . .	138
7.2.3	Gap filling during search . . . . .	139
7.2.4	Experiments . . . . .	139
7.3	Evaluation of results . . . . .	140
7.3.1	Simple phantom results . . . . .	141
7.3.2	Animated phantom results . . . . .	148
7.3.3	NCAT phantom results . . . . .	153
7.4	Conclusions . . . . .	157
<b>8</b>	<b>Conclusions and Future Work</b>	<b>160</b>
8.1	Summary and novel contributions . . . . .	160
8.2	General conclusions . . . . .	161
8.3	Future work . . . . .	163
	<b>Bibliography</b>	<b>165</b>



# List of Figures

2.1	X-ray spectrum . . . . .	7
2.2	Fan-beam geometries . . . . .	12
2.3	Fourier slice theorem . . . . .	13
2.4	Cone-beam geometry . . . . .	15
2.5	ICRU Report 50 volumes . . . . .	22
2.6	Three classes of motion . . . . .	25
3.1	Synergy geometry . . . . .	37
3.2	Elekta Synergy system . . . . .	38
3.3	Simple virtual phantom motions . . . . .	39
3.4	Simple phantom slices . . . . .	40
3.5	Quasar phantom . . . . .	42
3.6	Animated phantom slices . . . . .	44
3.7	Animated phantom illustration . . . . .	45
3.8	NCAT phantom slices . . . . .	46
4.1	Projection gaps by time . . . . .	52
4.2	Projection gaps by gantry angle . . . . .	52
4.3	Transaxial slices from reconstructions of Static phantom . . . . .	54
4.4	Slices from Simple phantom at three bin positions . . . . .	57
4.5	Motion magnitude plot of Simple phantom with bins 1,5, and 8 highlighted . . . . .	57
4.6	K-means clustering of NCAT phantom tumour position . . . . .	58
4.7	Slices from NCAT phantom at three bin positions . . . . .	58
4.8	Volume comparisons illustration . . . . .	59
4.9	ROI region examples of Simple and NCAT phantoms . . . . .	60
4.10	Reconstruction slices of Simple phantom at three bin positions . . . . .	61
4.11	Reconstruction slices of NCAT phantom at three bin positions . . . . .	62
4.12	Even vs binned projection distribution for three bins of an NCAT phantom configuration . . . . .	63

4.13	SSD errors for even and random binning of Static phantom . . . . .	65
4.14	SSD errors for even and random binning of Simple phantom . . . . .	66
4.15	SSD errors for even and random binning of NCAT phantom . . . . .	66
5.1	Four NCAT projections . . . . .	71
5.2	Before and after comparison of NCAT reprojection . . . . .	72
5.3	1-scan gap filling results, Simple phantom, sagittal slices, bin 1 . . . . .	74
5.4	1-scan gap filling results, Simple phantom, sagittal slices, bin 5 . . . . .	75
5.5	1-scan gap filling results, Simple phantom, sagittal slices, bin 8 . . . . .	76
5.6	1-scan gap filling results, NCAT phantom, sagittal slices, bin 1 . . . . .	77
5.7	1-scan gap filling results, NCAT phantom, sagittal slices, bin 5 . . . . .	78
5.8	1-scan gap filling results, NCAT phantom, sagittal slices, bin 8 . . . . .	79
5.9	NCAT projection differences . . . . .	83
5.10	Motion amplitude signals from two scans . . . . .	84
5.11	2-scan gap filling results, Simple phantom, bin 1, volume slices . . . . .	85
5.12	2-scan gap filling results, Simple phantom, bin 5, volume slices . . . . .	86
5.13	2-scan gap filling results, Simple phantom, bin 8, volume slices . . . . .	86
5.14	2-scan gap filling results, NCAT phantom, bin 1, volume slices . . . . .	87
5.15	2-scan gap filling results, NCAT phantom, bin 5, volume slices . . . . .	87
5.16	2-scan gap filling results, NCAT phantom, bin 8, volume slices . . . . .	88
6.1	Three examples of objective function scoring sequences . . . . .	93
6.2	Illustration of preliminary partitioning . . . . .	95
6.3	Partitioning differences between one-scan and two-scan approach . . . . .	98
6.4	Illustration showing direct and indirect reconstruction . . . . .	101
6.5	Simple phantom bucketing using ground truth difference signal . . . . .	112
6.6	2-scan Simple phantom results using ground truth difference signal . . . . .	114
6.7	2-scan qualitative results, Simple configuration 1 . . . . .	115
6.8	2-scan qualitative results, Simple configuration 2 . . . . .	115
6.9	2-scan qualitative results, Simple configuration 3 . . . . .	116
6.10	2-scan qualitative results, Simple configuration 4 . . . . .	116
6.11	2-scan qualitative results, Simple configuration 5 . . . . .	116
6.12	2-scan Animated phantom bucketing using ground truth difference signal . . . . .	117
6.13	2-scan Animated assignment results using ground truth difference signal . . . . .	118
6.14	2-scan qualitative results, Animated configuration 1 . . . . .	119
6.15	2-scan qualitative results, Animated configuration 2 . . . . .	119
6.16	2-scan qualitative results, Simple configuration 3 . . . . .	120

6.17	2-scan qualitative results, Simple configuration 4 . . . . .	120
6.18	2-scan NCAT phantom bucketing using ground truth difference signal . . .	121
6.19	2-scan NCAT phantom assignment results using ground truth difference signal . . . . .	122
6.20	2-scan qualitative results, NCAT configuration 1 . . . . .	123
6.21	2-scan qualitative results, NCAT configuration 2 . . . . .	123
6.22	2-scan qualitative results, Animated configuration 3 . . . . .	124
6.23	2-scan qualitative results, Animated configuration 4 . . . . .	124
6.26	2-scan Qualitative results of search on Simple configuration 1, estimated difference signal . . . . .	126
6.24	2-scan Simple phantom bucketing using estimated difference signal . . .	126
6.27	2-scan Qualitative results of search on Simple configuration 2, estimated difference signal . . . . .	127
6.25	2-scan Simple phantom assignment results using estimated difference signal	127
6.28	2-scan Qualitative results of search on Simple configuration 3, estimated difference signal . . . . .	128
6.29	2-scan Qualitative results of search on Simple configuration 2, estimated difference signal . . . . .	128
6.30	2-scan Qualitative results of search on Simple configuration 5, estimated difference signal . . . . .	128
6.31	2-scan Animated phantom bucketing using estimated difference signal . . .	129
6.32	2-scan Animated phantom assignment results using estimated difference signal . . . . .	130
6.33	2-scan Qualitative results of search on Animated configuration 1, esti- mated difference signal . . . . .	130
6.34	2-scan Qualitative results of search on Animated configuration 2, esti- mated difference signal . . . . .	131
6.35	2-scan Qualitative results of search on Animated configuration 3, esti- mated difference signal . . . . .	131
6.36	2-scan Qualitative results of search on Animated configuration 4, esti- mated difference signal . . . . .	131
6.37	2-scan NCAT phantom bucketing using estimated difference signal . . . .	132
6.39	2-scan Qualitative results of search on NCAT configuration 1. Sagittal slices are presented from a complete reconstruction (a), from the ground truth for bin 1(b), from an estimated bin 1(c), from the ground truth for bin 1 (d) and from the estimate for bin 2 (e). . . . .	133

6.38	2-scan NCAT phantom assignment results using estimated difference signal	133
6.40	2-scan Qualitative results of search on NCAT configuration 2.	134
6.41	2-scan Qualitative results of search on NCAT configuration 3.	134
6.42	2-scan Qualitative results of search on NCAT configuration 4.	134
7.1	1-scan assignment results, Simple configuration 1	141
7.2	Plot of mean of projections with assignment, no-fill process	142
7.3	1-scan assignment results, Simple configuration 2	142
7.4	1-scan assignment results, Simple configuration 3	143
7.5	1-scan assignment results, Simple configuration 4	143
7.6	1-scan assignment results, Simple configuration 5	144
7.7	1-scan qualitative results, Simple configuration 2	144
7.8	1-scan qualitative results, Simple configuration 1	145
7.9	1-scan qualitative results, Simple configuration 3	146
7.10	1-scan qualitative results, Simple configuration 4	146
7.11	1-scan qualitative results, Simple configuration 5	147
7.12	1-scan assignment results, Animated configuration 1	148
7.13	1-scan assignment results, Animated configuration 2	149
7.14	1-scan assignment results, Animated configuration 3	149
7.15	1-scan assignment results, Animated configuration 4	149
7.16	1-scan qualitative results, Animated configuration 1	150
7.17	1-scan qualitative results, Animated configuration 2	150
7.18	1-scan qualitative results, Animated configuration 3	151
7.19	1-scan qualitative results, Animated configuration 4	152
7.20	1-scan assignment results, NCAT configuration 1	153
7.21	1-scan assignment results, NCAT configuration 2	154
7.22	1-scan assignment results, NCAT configuration 3	154
7.23	1-scan assignment results, NCAT configuration 4	154
7.24	1-scan qualitative results, NCAT configuration 1	155
7.25	1-scan qualitative results, NCAT configuration 2	156
7.26	1-scan qualitative results, NCAT configuration 3	156
7.27	1-scan qualitative results, NCAT configuration 4	157

# List of Tables

4.1	Average projection counts and gap sizes per period by bin count . . . . .	56
4.2	Total number of projections per bin, Simple phantom . . . . .	56
4.3	Total number of projections per bin, NCAT phantom . . . . .	56
4.4	SSD, complete volumes, Simple phantom . . . . .	64
4.5	SSD, complete volumes, NCAT phantom . . . . .	65
4.6	SSD, ROI only, Simple phantom . . . . .	67
4.7	Sum of squared differences, constrained to region of interest (ROI), for the NCAT phantom. These are listed by distributions by bin as per the experiment description. . . . .	68
5.1	1-scan gap filling results, Simple phantom, complete volume comparison .	73
5.2	1-scan gap filling results, NCAT phantom, complete volume comparison .	75
5.3	1-scan gap filling results, Simple phantom, ROI comparisons . . . . .	77
5.4	1-scan gap filling results, NCAT phantom, ROI comparisons . . . . .	78
5.5	Projection differences between paired projections . . . . .	83
5.6	2-scan gap filling results, Simple phantom, volume comparisons . . . . .	84
5.7	2-scan gap filling results, Simple phantom, ROI comparisons . . . . .	84
5.8	2-scan gap filling results, NCAT phantom, volume comparisons . . . . .	84
5.9	2-scan gap filling results, NCAT phantom, ROI comparisons . . . . .	85
6.1	Summary of errors for phantoms using a ground-truth difference signal for partitioning . . . . .	125
6.2	Summary of percentage of errors for phantoms using an estimated difference signal for partitioning. . . . .	134
7.1	Percentage of errors for Simple phantoms by gap fill method . . . . .	148
7.2	Percentage of errors for Simple phantoms by gap fill method . . . . .	153
7.3	Percentage of errors for NCAT phantoms by gap filling method . . . . .	157

# Acronyms

CBCT:	Cone-Beam Computed Tomography
CT:	Computed Tomography
CTV:	Clinical Tumour Volume
FOV:	Field of View
GTV:	Gross Tumour Volume
IEC:	International Electrotechnical Commission
IGRT:	Image-Guided Radiotherapy
IMRT:	Intensity Modulated Radiotherapy
NURBS	Non-Uniform Rational B-Spline
NCAT:	NURBS-based CARDiac-Torso
ROI:	Region of Interest
PTV:	Planning Target Volume
OAR:	Organ at Risk
SSD:	Sum of Squared Differences
WHO:	World Health Organization

# Chapter 1

## Introduction

---

### 1.1 Cancer, imaging, and computer science

Cancer is one of the most frightening words in the English language. It describes a large group of diseases all of which are characterized by uncontrolled growth of a cell population resulting ultimately, if not eradicated, in the death of the host. In the US, according to the American Cancer Society, “men have slightly less than a 1 in 2 lifetime risk of developing cancer; for women, the risk is a little more than 1 in 3” [4]. In absolute numbers, more than 1,500 Americans die each day due to cancer which makes it the second most common cause of death. The GLOBOCAN database published by the World Health Organization shows that the US is not an outlier among the more developed nations [99]. The fact that the mantle “cancer” covers so many different kinds of diseases originating in different types of tissues makes it an exceedingly difficult problem to deal with and yet the prevalence and mortality rates reinforce the necessity of pursuing solutions.

From the clinical perspective, there are both therapeutic and research aspects of cancer. The therapeutic aspects cover both the ability to correctly diagnose and to effectively treat illness. The research aspects are what allow us to gain insight and understanding into the causes of cancer as well as the effectiveness of therapeutic approaches. Research also creates and validates new therapies. The principal modes of therapy in widespread use currently are surgery, radiation, and chemotherapy. Often, the treatment design will include some combination of these three.

Preceding any kind of treatment is some kind of diagnostic and planning process. In addition to any biological or chemical markers found with tests, images of the region of interest can be acquired to provide the diagnostician with more information. Different technologies exist having different strengths and weaknesses. X-ray imaging provides excellent bone contrast but poor soft tissue contrast and has risk associated with its use. Nuclear imaging techniques such as single photon emission computed tomography (SPECT) or positron emission tomography (PET) measure tissue function and physiology more than structure. These can provide better insight into the health of soft tissues but likewise have some risk associated with their use. They can be, however, difficult to “position” within a known reference frame. To compensate for this, they can be “fused” with other modalities such as X-ray to assist clinicians in seeing where a particular response to the nuclear material is occurring. Magnetic resonance imaging avoids the ionizing radiation risk of X-rays while providing excellent soft tissue contrast and good spatial resolution. It is among the most expensive imaging modalities and is sufficiently slow that motion poses a significant problem. On the other end of the cost scale, ultrasonic imaging is both portable and relatively inexpensive. It can also measure blood flow and is therefore useful in assessing heart function and vessel health. Poor soft-tissue contrast and the inability to “see” past bone and gas pockets are some of its limitations. Relatively new and exciting modalities such as magnetic resonance elastography (MRE) being pioneered at the Mayo Clinic, or molecular imaging using fluorescing markers in the near-infrared frequencies, offer continued hope for improvements in the tools available to clinicians in the future. Each of these imaging technologies, or “modalities” have trade-offs of noise, contrast, and spatial resolution both within a modality and between modalities. Different modalities will have different risk, cost, and clinical time trade-offs as well. Consequently, different modalities will have different usage models. Clinicians need to use the best method based on the information they most need.

In this work, the particular imaging technology or “modality” under study is cone-beam computed tomography (CBCT). It is a diagnostic X-ray modality integrated with a therapeutic linear accelerator that acquires 2D images used to reconstruct a 3D volumetric image. Typically, it is an adjunct technique used in conjunction with radiation therapy and will be described later in Chapter 2. In that context, it is used to evaluate the patient’s position with respect to the planning frame of reference and, if necessary, make adjustments before treatment begins. The acquisition protocol commonly used captures approximately 670 X-rays, or “projections,” over the course of 120 seconds in a single orbit covering 360 degrees. The two minutes (intentionally) used to capture the images implies two things: that the resulting reconstructed volume will be blurry where motion



has occurred and that, in some sense, that motion has been captured in the projections. After the reconstructed volume has been used to align the patient, it and the projections have no further use. This implies that the alignment assessment is the only benefit of CBCT at the dosage cost of the 670 X-rays. That is the motivation of the research in this thesis; the goal is to find a way to extract the motion information from the projections without requiring any additional model, planning CT, or captured signal.

## 1.2 Thesis overview

The main points of the thesis are presented here beginning with Chapter 2, the traditional background and relevant literature review. In that chapter, the basic biology and physics that relate to radiotherapy, the essential clinical practices, and the basic ideas behind CBCT and filtered backprojection reconstruction are presented. The work to date on identifying motion in CBCT is also highlighted. The chapter concludes with a new classification for motion types which is sufficiently general while clearly segmenting the motion types into distinct subgroups; this classification lays the foundation for more clearly elucidating the thesis problem.

In Chapter 3, the materials used in several of the subsequent chapters are presented. Three different kinds of phantom classes and specific configurations of them are described. Using phantoms lowers the cost of experimentation and the time required to run experiments. It minimizes risk to patients by first exploring experimental ideas on virtual or physical objects. A new type of phantom, the Animated phantom is presented in this chapter and is one of the novel contributions of the thesis.

In Chapter 4, a specific kind of artefact that is a particular consequence of the binning process as applied to CT and CBCT is defined and named gap artefacts. Several experiments are run to illustrate its distinctiveness and impact on reconstruction. In the literature, this class of artefact is alluded to but generally grouped with the larger class of artefacts it belongs to, namely *undersampling* artefacts. The contribution of this chapter is to explicitly call out the class as an important type that should be accounted for in filtered backprojection reconstructions that include a binning step. The use of k-means clustering on a known object's positions as an "oracle" for correct binning is described in this chapter and will be used in subsequent chapters.

In Chapter 5, having characterized the gap artefacts, two methods are proposed and evaluated for mitigating the impact of them on the reconstruction effort. Each of these methods in turn uses two different approaches, balanced fill and complete fill, to fill the gaps. These methods tie into and pave the way for the search methods presented in Chap-

ters 6 and 7.

Chapter 6 introduces the core contribution of the thesis: a search method for assigning a bin label to each projection without the use of any prior model or correlating signal/data. The search space is identified, an objective function to be maximized is proposed, and then a series of techniques are described to sufficiently reduce the size of the search space such that an exhaustive search can be performed. One of the novel proposals is the use of Principal Component Analysis (PCA) as a preparatory step to increase the computational performance of the filtered backprojection reconstruction step. The overhead required for this step is more than compensated for by the improvements in performance given the vast number of reconstructions that are necessary. A second novel contribution is the proposal to modify the scanning protocol to include two scans rather than the standard single scan. The advantages and disadvantages of this change are explained in the chapter.

Chapter 7 uses several of the ideas in Chapter 6 but applies them to the standard single-scan protocol. The advantages of the two-scan protocol are no longer available so a new search method is proposed, a random-restart hill climbing method. The justification for this approach is made, several experiments are performed, and the results presented and analysed. The two methods presented in Chapter 6 and Chapter 7 are subsequently compared.

The thesis concludes with a summary of the work including the strengths and weaknesses of the approaches presented. The elements of the thesis which are considered novel contributions are recapitulated, and several directions for future research are discussed.

# Chapter 2

## Background

---

### 2.1 Introduction

In Chapter 1, the importance of radiotherapy as one of the tools in the fight against cancer was introduced in the context of cancer types and prevalence. In this chapter, the precise nature of radiotherapy is introduced. With this understanding, current issues that form major thrusts of active research will be discussed. Finally, the specific problem that this research addresses will be identified.

The main imaging modality of interest, cone-beam computed tomography (CBCT) is described in Section 2.3.3.2 along with its diagnostic and planning cousin, Computed Tomography (CT). A natural progression when talking about computed tomography, and the one used here, is the dimensional advance from two-dimensional CT to three-dimensional CT and finally four-dimensional CT (3D + time).

Motion, the cause of the main problems of interest, is described in Section 2.4. A new classification of motion is introduced to more clearly distinguish the different problems that arise from motion. Methods to compensate for motion and the problems it causes will then be described in Section 2.6.

Some background discussion is deferred until future chapters where it makes more sense to introduce the material immediately in context with how it will be used. As an example, the background on artefacts is introduced in Chapter 4 where it forms a crucial part of the material in that chapter specifically. Likewise, a review of principal component

analysis (PCA) as used to facilitate searching in this research is introduced in the first chapter on search, Chapter 6.

One way to understand radiotherapy is to examine the timeline and the roles played by different clinicians in the treatment process. The problems caused by motion are also a function of these clinical practices and, as such, some comprehension of the process is necessary to understand the motivation and constraints behind much of the current research.

To reduce potential confusion, some nomenclature must be introduced which will be used when reporting the findings of different authors. With respect to anatomical position or motion direction, different terminology can be used to describe the same direction or anatomical coordinate system. For consistency, if needed, the results of a paper may be translated as follows. When referencing the direction from the head to the feet, the Superior-Inferior (SI) term will be used. This will be substituted whenever the term Cranial-Caudal (CC) is used. The term Left-Right (LR) will generally be used instead of Medial-Lateral (ML). Anterior-Posterior (AP) referring to the front and back seems to be universal and needs no translation.

## 2.2 Ionization of tissue

### 2.2.1 Relevant physics

X-rays are used to both image patients and to kill diseased tissue. The critical difference between the two is the *energy levels* involved. The energy levels of the X-rays, measured in kilo-electron volts (keV), typically range from 50 keV to 120 keV [21, p. 113] for imaging applications. Therapeutic X-rays range between 4 and 25 mega-electron volts (MeV) [89] though the distinction is not entire clear since MeV X-rays are sometimes used for portal imaging (see Section 2.3.2). The *intensity* of the X-ray—or the number of photons—is a function of the electrical current used to generate the X-rays.

In the literature reviewed here, the units  $kV$  or  $kV_{eff}$  or  $kV_p$  are used. These refer to the *potential voltage* applied between the cathode and anode in the X-ray tube, the *effective or modal energy* of a polychromatic spectrum, and the *peak photon energy* respectively though the last one,  $kV_p$ , seems to be interchangeably used as the peak voltage applied to the tube and the peak photon energy. What is important to point out is that the spectrum, when shown as  $keV$  values, cannot exceed the  $kV_p$  value used to generate the X-ray. Figure 2.1 shows these values and illustrates the polychromatic nature of X-rays. Included for completeness are the characteristic radiation lines of Tungsten which are important to

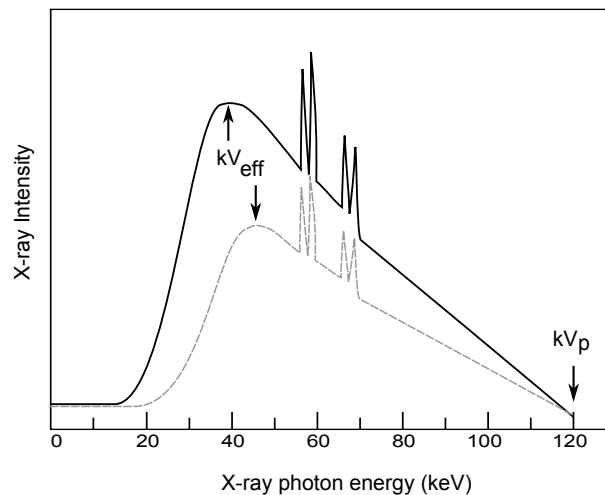


Figure 2.1: Example X-ray spectrum at 120 kV showing beam hardening (grey dashed line), effective kV and peak kV.

analytic chemistry but are irrelevant to the work in this thesis. The two spectra will be explained in more detail shortly.

Choosing voltage and current levels for imaging is complex because of the many trade-offs. Ultimately, the aim is to measure photons in some fashion by a detecting device. Contrast is achieved by the differences in absorption and scattering of the different tissue types. Noise in the detection of photons is modelled as a Poisson process [2] and therefore is a function of the expected number of photons. Given that the variance of a Poisson distribution is equal to the expectation, one sees that increasing the number of photons (by increasing the current) decreases the *relative* noise. Hence, the signal-to-noise ratio (SNR) is improved by increasing the current. Unfortunately, increasing current also increases dose which is undesirable. Ravenel et al. studied the effect of current change on image quality [65]. They found that for CT images generated at 120kV, the image quality increased as they increased the current from 40 mAs to 160 mAs but then plateaued and no significant increase in quality was observed up to their final value of 280 mAs.

Increasing electrical potential difference of the X-ray generation device increases the energy of the photon and hence its ability to penetrate tissue before being absorbed. However, it is this absorption (and scatter) which creates subject contrast. The higher the energy the lower the contrast-to-noise (CNR) so decreasing the  $keV_p$  increases the resulting contrast in images. The attenuating differences between fat, soft tissue, and bone is greatest between 10 keV and 30 keV [21, p. 134]. Unfortunately, low-energy X-rays are quickly absorbed or immediately attenuated by the glass enclosure on X-ray tubes. This is why, when imaging the soft tissues in mammography, breast compression is used to

reduce the distance the X-rays must travel and special tube enclosures are used so that lower-energy X-rays can be generated. Conversely, X-rays travelling through the torso, especially in larger patients, must have the effective energy levels of the photons increased through *beam hardening* so that low-energy X-rays, which contribute to dose but not to the image, are removed.

Both current and energy contribute to the ionization of molecules by X-rays. Ionization occurs in different ways though the probability of a given kind of interaction is a function of the energy. In the course of ionizing an atom or molecule, the X-ray photon loses energy. As X-rays continue through matter they continue to lose energy and continue to ionize molecules until they are ultimately absorbed. The speed with which this occurs is a function of the density of the material. The attenuation occurs as an exponential decay and is frequently described in terms of intensity changes, known as Lambert's law, using the equation [21, eq. 1.11]:

$$I = I_0 e^{-(\mu x)} \quad (2.2.1)$$

Where  $I_0$  is the initial intensity of the X-ray,  $\mu$  is the linear attenuation coefficient for material at the given X-ray energy level,  $x$  is the distance travelled, and  $I$  is the emerging intensity of the attenuated X-ray. This model indirectly describes a monochromatic X-ray because the attenuation coefficient is tuned to a specific energy level or X-ray frequency. As was mentioned above, X-rays used in imaging and those generated for therapeutic use by a linear accelerator ("linac") are polychromatic in nature. While the simple monochromatic equation is sufficient for most applications, understanding the polychromatic nature is necessary to understand when and why beam hardening is used. In Figure 2.1, two spectra are shown. The lower intensity plot is shown with a dashed grey line and represents the initial X-ray spectrum (shown with the solid black line) that has been filtered using a thin sheet of aluminium or copper for instance. These filters work by preferentially absorbing the lower energy photons, the easiest ones to absorb. Note also that beam hardening occurs automatically for very low energy levels via the X-ray tube enclosure. The purpose of beam hardening is to reduce the number of low energy photons which will only be absorbed by tissue and thus contribute to dose but not to the image in any way.

On an engineering note, less than 1% of the current used to generate X-rays are actually converted into X-ray photons [21, p. 39]; the remainder are converted into heat. Managing heat generation properly is a significant concern in the design of X-ray devices and is an indirect constraint on imaging ideas. Any suggested technique that requires leaving the X-ray tube on for long periods of time should therefore have a feasibility assessment.

### 2.2.2 Key aspects of radiobiology

The principal concern of radiobiology is the effect of ionizing radiation on biological processes. Ionization changes the chemical properties of the ionized atoms and molecules in ways that disrupt the normal cellular mechanisms, in particular DNA replication. For instance, the ionization of a water molecule begins a chain reaction of ionization that may end up in the production of hydrogen free radicals and hydrogen peroxide both of which are damaging to cellular function. In fact, given that tissues are 70–90% water, “It is the free radicals formed from water that are responsible for about 70% of the biological effects of radiation. . .” [97, p. 56].

Ionization of DNA, RNA, amino acids and proteins has a more direct effect on cellular structures. Unrepaired damage can lead to the death of the cell and/or the inability to correctly replicate. An example of this is found in the p53 tumour suppressing gene. When DNA is damaged, p53 either delays the normal replication cycle to give the cell time to repair the damage or, if the damage is too severe, to trigger cell suicide by apoptosis. However, if the p53 gene itself is damaged, these mechanisms may be thwarted leading to cellular reproduction with damaged DNA (as well as the damaged p53 gene) [97, p. 21].

Fortunately, cells have built-in repair systems though the repair mechanism of cells is not completely understood. What is known is that healthy cells are able to repair themselves more effectively than damaged cells (e.g. malignant tumour cells). This fact is ultimately the basis of the fractional dose delivery protocol which will be described later in the section on radiotherapy (Section 2.5). The relevant aspect is that radiation doses meant to kill tumours are split into “fractions” and delivered over several treatments rather than all at once [60, chap.16]. The timing between treatments is meant to maximize the repair time for healthy tissue while minimising the repair time for cancerous tissue. The introduction of fractionated radiotherapy plays a significant role in the research related to the work presented here.

Another characteristic of radiobiology that affects the therapy planning is the variation in intrinsic radiosensitivity of different tissue types [97, chap. 8]. Simply understood, certain tissues are more likely to be affected by a therapeutic dose than others. Hence, when developing a radiotherapy plan, the sensitivity of healthy tissues that will be irradiated needs to be taken into account. Similarly, in large tumours, hypoxic effects often cause the middle portion of tumours to be less sensitive to radiation and this too needs to be both planned for and, ideally, monitored.

When large tumours are successfully treated, *necrosis* (the death of the tumour cells) occurs and the cell debris is eliminated through *phagocytosis*. If the region where this occurs is large enough, the tumour can collapse or significantly reshape in response to

the structural change. Such changes may effectively invalidate the original treatment plan and unless detected and corrected, healthy tissue will be targeted and malignant tissues will be missed. This becomes another motivation for re-imaging patients during a course of treatments in order to verify that the original plan is still valid.

### 2.2.3 In summary

To recapitulate the essential elements of this section:

- Diagnostic X-rays create images as the attenuated photons are measured by detectors; contrast is generated by the different densities of the different tissue types.
- The energy level of the X-rays is a function of the goal of the imaging. Low energy levels are useful in mammograms, high-energy X-rays are needed for torso imaging.
- All X-rays in medical settings ionize matter. This is intentional in therapeutic X-rays and undesired in diagnostic X-rays.
- Both voltage (kV or MV) and current (mAs) contribute to dose but current is dominant.
- Sufficient levels of ionization lead to cell, and eventually tissue, death. Insufficient ionization allows cells to repair themselves.
- $I = I_0 e^{-(\mu x)}$

The application of X-rays in the medical context comes down to this: minimize the dose when imaging, minimize the dose to healthy tissue during treatment, and maximize the dose to target tissues during treatment.

## 2.3 X-ray imaging

The relevant research and the work presented in this thesis have their roots in the images created with X-rays. A basic background in the various modalities is therefore required and is now presented.



### 2.3.1 Radiographs

Radiographs are the “official name” for what commonly are called X-Rays. These are the ubiquitous two-dimensional images taken at the dentist or in the hospital to reveal broken bones or serious diseases, etc. They can be large images, in the case of a chest X-ray, or small images used to ascertain the health of a tooth. These images originally were created using a chemical film process which effectively established the use of light values to represent strong attenuation and dark values to represent weak attenuation.

### 2.3.2 Portal images

In addition to diagnostic and planning imaging using keV energy levels, radiographs using MeV energy levels are sometimes used during radiotherapy sessions. The main motivation for these energy levels is the pragmatism of having a ready MV source: the linac used to deliver the therapeutic radiation. Capturing two orthogonal MV images allows clinicians to verify a patient’s position on a treatment couch before delivering the radiation dose. The additional dose delivered by these “portal” images has historically been justified by the increased accuracy of therapeutic dose delivery [15].

### 2.3.3 Computed Tomography (CT)

Normal radiographs have a high information density. They contain information about an entire volume of tissue densities compressed into a two-dimensional projection or summation. Computed tomography, on the other hand, seeks to quantize that information by reconstructing the information inherently summed along the X-ray by taking many samples at different geometric positions (usually tracing a non-coplanar overlapping spiral or a coplanar arc).

#### 2.3.3.1 2D and 3D CT

In 1979 Godfrey Hounsfield and Allan Cormack were awarded the Nobel Prize in Medicine for their independent invention of X-ray computed tomography [97, p.34]. While Cormack’s work preceded Hounsfield’s by a few years, it was Hounsfield who patented his work and first successfully introduced it into clinical practice in 1972. In honour of this, the numerical units generated by CT systems are called Hounsfield units. Hounsfield’s initial work used an iterative algebraic technique to reconstruct the images. Later, Ramachandran and Lakshminarayanan introduced an analytical formulation using *filtered*




Figure 2.2: Parallel (left) and equispacial fan-beam (right) geometries.

*backprojection* to reconstruct 2D slices that were then stacked to form 3D volumes. Iterative algebraic techniques are no longer popular for standard reconstruction situations because they are slow and prone to numerical error [21, p.392]. Filtered backprojection, in contrast, is very fast and sufficiently accurate given enough projections. The independence of the projections means that the reconstruction can begin even before all the projections have been acquired which further improves its performance [57, 69], [32, p.275].

The work in this thesis is ultimately based on the filtered backprojection method. For details on iterative algebraic techniques, the reader is referred to Chapter 7 in Kak and Slaney’s book [32], the work by Mueller [57], or the description in Rit’s comparison of analytic and algebraic techniques with respect to cone-beam CT [69].

Filtered backprojection is best explained using the simple parallel-beam geometry. Older generations of CT machines used the parallel geometry but newer machines use a fan-beam variant. The most common fan-beam geometry is based on equiangular spacing on a curved detector and is used in most CT devices today. The equispacial geometry, with a flat detector, becomes the basis for the cone-beam CT geometry. It can be shown that either fan-beam geometry reduces to the parallel case through rebinning (and weighting in the equispacial case).

Filtered backprojection for the 2D parallel case begins with the understanding that the 1D projection through the 2D slice of the object (which can be alternatively thought of as a density function), is a line integral taken along the line perpendicular to the projection direction (See Figure 2.3). These 1D projections are taken uniformly at many angles around the object. The purpose of this is seen from the “Fourier slice theorem”, also known as the “central slice theorem” or the “projection-slice theorem” which essentially




Figure 2.3: Fourier slice theorem. The 1D Fourier transform of the projection becomes a slice in the 2D frequency domain.

shows that a *slice* through the 2D frequency domain is equal to the Fourier transform of the line projection through the spatial domain at that same angle. Figure 2.3 illustrates this.

Given enough 1D projections, the 2D space can be recovered by interpolating in 2D frequency space and then transferring back via a 2D inverse Fourier transform to the 2D spatial domain. In practice, interpolation errors are minimized if the interpolation is done in the spatial domain after the filtering takes place in the 1D frequency domain [32].

One very important assumption to note here is that because the projection is presumed to be a line integral, the values are presumed to be linearly summed along the projection line. In Equation 2.2.1, this was shown to not be true for X-rays. However, if the logarithm of the measured intensity is taken, then the assumption holds.

$$\begin{aligned} I &= I_0 e^{-\mu_1 x_1} e^{-\mu_2 x_2} \dots e^{-\mu_n x_n} \\ &= I_0 e^{-(\mu_1 x_1 + \mu_2 x_2 + \dots + \mu_n x_n)} \end{aligned} \quad (2.3.1)$$

becomes

$$-\ln\left(\frac{I}{I_0}\right) = \sum_{k=1}^n \mu_k x_k \quad (2.3.2)$$

The most common implementations of 3D CT construct 3D volumes either from stacks of reconstructed 2D slices, or by reconstructing and interpolating measurements

acquired sequentially in a spiral path along the rotation axis. Details of modern CT technologies can be found in [97, chap. 1], [21, chap. 14], and [91, chap. 1].

### 2.3.3.2 Cone-beam CT

Simplistically, one can say that CT takes one-dimensional projections and reconstructs a two-dimensional slice. Cone-beam CT accomplishes the same thing but increases the dimensionality: it takes a set of two-dimensional equispacial fan-beam projections and reconstructs a three-dimensional volume. Cone-beam CT is becoming increasingly popular as an adjunct to image-guided radiotherapy (IGRT) and is used to verify patient positioning and tumour/organ interfraction motion.

Historically, while exact methods for reconstructing 2D images from 1D slices were well understood, the obvious idea of reconstruction 3D images from 2D projections was hindered by the tension between the theoretical and the pragmatic. The work of Tuy [93] and Smith [85] rigorously describe conditions that are necessary and sufficient for exact reconstructions from cone-beam or 2D projections. Tuy's condition is frequently described in visual terms as requiring every possible plane through the imaged volume to intersect the path traced by the cone-beam source at some point. These conditions show that exact reconstructions are not possible without more than one independent acquisition path. More specifically, using the existing linear accelerator gantry mechanisms which describe a source orbit in a single plane to acquire projections was proven to be incomplete. Implementing exact solutions would require a major change to the design of existing medical equipment.

Nevertheless, this inexact impediment did not stop a team of researchers from proposing an inexact but adequate method or, as the title put it, a 'practical cone-beam algorithm' in 1984. The method, by Feldkamp, Davis, and Kress [22], is often referred to as the FDK algorithm and is the basis for most modern cone-beam reconstruction implementations despite other algorithm or geometry proposals that fulfilled the Tuy condition [16, 85, 100]. This dense paper is elucidated in the oft referred to tome of Kak and Slaney [32] as well as the thesis by Turbell [91, chap. 2].

More recent work, by Katsevich for example [33], has introduced an exact solution by combining the orbit with an axial traversal in the form of a spiral source traversal. Though computationally intensive it can be dramatically sped up with parallel implementations [17]. However, the work in this thesis does not require an exact solution and uses a simple property of the FDK algorithm to accomplish its goals so Katsevich-type algorithms are not considered further.

In one of the first reported implementations [12], a modification of the FDK algorithm




Figure 2.4: Cone-beam geometry. The shaded “fan” is shown tilted away from the plane containing the X-ray source.

is suggested to deal with large field-of-views (FOVs). The problem arises from most flat-panel detectors (FPDs) being smaller than the human torso. In order to compensate, they shift the detector laterally and then weight the projection data appropriately before reconstructing. Another suggested approach to this problem is to virtually extend the detector and compensate for missing data by mirroring the edge data appropriately [62].

The most published consequence of the FDK algorithm is the increased error that occurs along the  $z$ -axis (rotation axis) away from the isocentre. This is a direct consequence of violating the Tuy conditions and an effort to improve this was proposed by Mori et al. [56] which they call combination-weighted FDK (CW-FDK).

The pioneering work of Jaffray et al. [28] on CBCT has increased its popularity as a replacement for portal imaging. Because of this widespread popularity and its use in this thesis, the basic geometry and mathematics behind the most popular reconstruction algorithm, the FDK algorithm as described by [32, chap. 3] is briefly recapitulated here. The seemingly trivial property that forms a significant feature of this research is then highlighted.

The FDK algorithm takes the parallel filtered backprojection algorithm, recasts it into the equispacial fan-beam geometry, and then extends it three-dimensionally by tilting the fan-beam geometry up the  $z$ -axis (Figure 2.4). Kak and Slaney summarize the cone-beam reconstruction algorithm as a sequence of three generalized steps:

1. Weight the projection to compensate for the fan and tilt geometry changes
2. Filter the projection
3. Backproject the weighted, filtered projection through the volume

The derivations of these steps are well documented and need not be repeated here. What is important with respect to the application of PCA to the search problem in Chapter 6 is step three. Beginning with Kak and Slaney's backprojection equation [32, p.107]:

$$g(t, s, z) = \int_0^{2\pi} \frac{D_{so}^2}{(D_{so} - s)^2} Q_{\beta} \left( \frac{D_{so} t}{D_{so} - s}, \frac{D_{so} z}{D_{so} - s} \right) d\beta \quad (2.3.3)$$

where  $t, s, z, x, y$ , and  $\beta$  can be seen in Figure 2.4 and

$$t = x \cos\beta + y \sin\beta \quad (2.3.4)$$

$$s = -x \sin\beta + y \cos\beta \quad (2.3.5)$$

To simplify further, the geometric constant  $D_{so}$  is removed, the weighting fraction is made a function  $w(s)$ , and  $Q$  is made a function of the angle  $\beta$  and the rotated coordinates.

$$g(t, s, z) = \int_0^{2\pi} w(s) Q(s, t, z, \beta) d\beta \quad (2.3.6)$$

Then, because this is implemented using discrete voxels and a limited number of projections, and because the initial equation only gives the value at one point in space, the  $\beta$  value is replaced with a function,  $p(i)$  extracting  $\beta$  and  $D(i)$  which extracts  $\Delta\beta$  from an index of projections. The summation now ranges over the index.

$$g(t, s, z) = \sum_{i=1}^N w(s) Q(s, t, z, p(i)) D(i) \quad (2.3.7)$$

This equation is a weighted average. The weight  $D(i)$  is the relative amount of the total arc covered by a given projection. Pragmatically, this can be treated uniformly and moved outside the summation.

$$g(t, s, z) = \frac{1}{N} \sum_{i=1}^N w(s) Q(s, t, z, p(i)) \quad (2.3.8)$$

Finally, the three-dimensional volume can just as easily be considered a very high dimension vector with  $s, t$ , and  $z$  being functions of the vector index  $j$  yielding:

$$v(j) = \frac{1}{N} \sum_{i=1}^N w(s(j)) Q(s(j), t(j), z(j), p(i)) \quad (2.3.9)$$

or even more simply

$$\mathbf{v} = \frac{1}{N} \sum_{i=1}^N w(j) Q(j, p(i)) \quad (2.3.10)$$

$$\mathbf{v} = \frac{1}{N} \sum_{i=1}^N \mathbf{w}_j \mathbf{q}_j \quad (2.3.11)$$

$$\mathbf{v} = \sum_{i=1}^N \left( \frac{\mathbf{w}_j}{N} \right) \mathbf{q}_j \quad (2.3.12)$$

This last equation allows something slightly subtle to be more easily described which the mathematical notation uncharacteristically fails to capture. From an algorithmic perspective, the most efficient way to implement 2.3.12 is to allocate space for vectors  $\mathbf{v}$  and  $\mathbf{q}$ , iteratively construct each  $\mathbf{q}_j$ , weight it, and add it to  $\mathbf{v}$ . Equally valid, but seemingly serving no purpose, is the idea of constructing *and allocating memory for* each  $\mathbf{q}$  before performing the averaging step. The motive for this will become clear in Chapter 6 when the retained  $\mathbf{q}$  vectors are dimensionally reduced using principal component analysis. It is shown in that chapter that reconstruction can be performed in feature eigenspace by averaging the dimensionally reduced vectors and that this provides a computational boost to the search algorithms proposed in this thesis.

### 2.3.3.3 MV CBCT vs kV CBCT

Cone-beam CT usually involves reconstructing keV energy projections. In light of the fact that all linacs are equipped by definition with an MV source, and that it is much easier to add a detector than a detector and another X-ray source, some work has been done on generating mega-voltage cone-beam CT (MV-CBCT). Investigations by [10] show the viability of it but they acknowledge the fact that soft tissue contrast is not very good and that MV-CBCT, in contrast to kV-CBCT results in nearly four times the dose. Because of these limitations and because more vendors are introducing kV-CBCT equipped linacs, it is unlikely that MV-CBCT will gain widespread acceptance.

### 2.3.3.4 In summary

To recapitulate the essential elements of this section:

- X-ray imaging typically uses keV energy levels, though MeV energy levels are sometimes used to create portal images.
- Traditional radiographs are imaged orthogonally to the direction of the X-ray. Computed tomography images are reconstructed from multiple samples resulting in images roughly parallel to the X-rays.
- 3D CTs are constructed by stacking 2D CT slices; the slices are either acquired axially with a “step and repeat” process or interpolated from a continuous spiral scanning path.
- Cone beam CT reconstructs 3D volumes directly from 2D projections. The FDK algorithm is the most widely adopted reconstruction implementation. Some studies have reported extending portal imaging to construct MV CBCTs but due to extreme doses and the increasing availability of kV CBCT, there is little adoption of this idea.

## 2.4 Physiological motion

The most obvious biological property that impacts all aspects of radiotherapy is motion. Motion comes in many forms. The most dominant motions are in the abdominal region [36] and are caused by respiratory and cardiac motion. These motions also have the advantage of being somewhat predictable although there can be great variation even within the same patient [98].

In the pelvic region and lower abdomen, bladder filling and gastrointestinal effects can create motion but are not periodic in nature and less (if at all) predictable. Even though they can occur during treatment, they are not as much an intrafractional motion concern as they are an interfractional concern. If a plan is constructed based on a diagnostic CT with the bladder full and the patient arrives for treatment with an empty bladder, the plan may be invalidated. Research has attempted to quantify the range of motions due to these effects [50], but compensation is essentially limited to dietary regulation and endorectal balloons [94].

In addition to internal organ/tissue changes and motion, there is the issue of patient placement. All treatment plans assume the patient can be placed in precisely the same position in a reference frame. This of course is a naïve assumption and is explicitly compensated for in the treatment plan as specified by ICRU Reports 50 and 62 [26, 96].



In a review of intensity modulated radiation therapy [98], Webb summarizes the motion characteristics of tumours in various organs as reported by many authors. In preparation for the discussion on motion compensation techniques later, certain tumour motions will be included here.

### 2.4.1 Motion in the lungs

In a study on 24 lung patients, Alasti et al. [3] found that motion is principally in the SI direction. Tumours in the lower lobes incur the greatest motion whilst mid and upper lobes experience less motion. They also found no significant difference between left and right lungs nor between male and female patients.

Seppenwoolde et al. [81] studied 20 patients using implanted gold markers in or near tumours. A real-time tracking system using fluoroscopy was used to determine tumour trajectories. They quantified the SI motion as being greatest in the lower lobes, when unattached, with an average amplitude of  $12 \pm 2\text{mm}$ . AP and LR direction motion was small ( $2 \pm 1\text{mm}$ ). They also documented hysteresis effects and reaffirmed that more time is spent in the exhalation state than the inhalation state.

Lujan uses a widely followed formula based on even powers of a cosine to model respiration [49]. While useful, it must be noted that Rietzel [68] mentions respiration irregularities which manifest themselves as frequency, amplitude, and general waveform differences. Sharp reports similar variation in amplitude among 14 lung tumour patients with standard deviations between 0.8 mm (mean 9.1 mm) and 5.3 mm (mean 27.5 mm) [82].

A comprehensive discussion of lung motion with respect to radiotherapy can be found in [35] which “describes the magnitude of respiratory motion, discusses radiotherapy specific problems caused by respiratory motion, explains techniques that explicitly manage respiratory motion during radiotherapy and gives recommendations in the application of these techniques for patient care, including quality assurance (Q) guidelines for these devices and their use with conformal and intensity modulated radiotherapy.”

### 2.4.2 Motion in the liver

In [13], results of a study on 79 patients showed a median tumour volume of  $294\text{cm}^3$  with average (maximum) intrafraction motion ranges of  $17\text{mm}(29\text{mm})$ ,  $9\text{mm}(18\text{mm})$ , and  $8\text{mm}(13\text{mm})$  in the SI, AP, and RL directions respectively. They note that the liver deforms during respiratory motion and suggest that deformable registration may have more value than standard rigid registration techniques currently in use with IGRT. Giergta’s

team [24] found smaller motion in 7 patients with pancreas or liver tumours reporting an average magnitude of 7.4 mm in the SI direction and 3.8 mm in the AP direction.

### 2.4.3 Motion in the prostate

Using gold markers implanted in the prostate, Kitamura et al. principally reported on interfraction motion differences as a function of prone versus supine patient positioning (supine is better). For supine positioning, the average amplitude of prostate movement, in millimetres, was  $0.1 \pm 0.1$ ,  $0.3 \pm 0.2$ , and  $0.3 \pm 0.4$  in the LR, SI, and AP directions respectively. Their conclusion is that prostate motion is affected by the respiratory cycle and bowel movement and the principal motion is in the SI direction. In contrast, Millender et al. [55] reported on a study of three morbidly obese men, positioning errors occurred predominantly in the LR direction with a mean magnitude of 11.4 mm compared to 7.2 mm for the SI direction and 2.6 mm in the AP direction. Other authors report slightly different results but the numbers remain relatively small, the SI direction is always dominant, and the conclusion is that intrafraction is insignificant but interfraction motion is not [50]. For this reason, motion compensation research with respect to prostate tumours is focused on interfraction changes, not on intrafraction motion.

### 2.4.4 In summary

Perhaps an ideal summary of anatomical motion comes from Diez [20]:

“The physiologic organ motion can be classified, according to their temporal behaviour, in: (i) non-periodic motion, produced by the filling status of structures such as the bladder or the rectum, (ii) periodic motion, due to breathing and cardiac motion (these movements are repeated many times during a single treatment session), and (iii) quasi-periodic motion, like the peristaltic movement of the stomach.”

To recapitulate the essential elements of this section:

- Respiratory motion is often modelled as a periodic function though in reality it can have significant changes even between cycles within the same patient.
- Respiratory motion is the most significant cause of motion within the patient. It not only impacts lung tissue but the liver (and other organs) as well.
- Respiratory motion is principally in the SI direction.

- Prostate motion, such as it is, is not as influenced by the periodic motion of the lungs (though there is some debate about this) but rather by the non-periodic or quasi-periodic motions caused by filling events in the rectum or bladder and peristalsis.

## 2.5 Radiotherapy

In Chapter 1, radiotherapy was introduced as one of the three main therapies currently used to treat cancer. Modern radiotherapy is sometimes called ‘conformal radiotherapy’ and conveys the idea that the delivered radiation dose conforms to the tumour or target tissue shape. The goal is to spare healthy tissue and completely cover the target. To enable this goal, increasingly complex protocols are being introduced into both the planning and treatment stages and the line between these two stages is starting to become blurred as re-planning during treatment gains momentum.

In this section, the process and clinical roles involved in radiotherapy are detailed. These further explain the problems of motion from the clinical perspective as well as ways in which motion is currently being included in the planning and execution of radiotherapy. The differences between curative, adjuvant, and palliative therapy are ultimately not relevant to the imaging issues discussed in the rest of the thesis and will not be considered here.

### 2.5.1 Diagnosis

At some point in a patient’s care, a physician may suspect the development of some form of cancer and order tests to confirm (or not) the suspicion. Should the diagnostic tests confirm that the patient has cancer, a clinical team will then examine the results. The clinical team often consists of specialists or consultants in oncology, surgery, pathology, radiology, and in some cases may include organ-specific specialists. This team discusses the prognosis and decides on the best course of treatment which may include radiation possibly in conjunction with surgery and/or chemotherapy.

### 2.5.2 Planning

If radiotherapy is prescribed as part of the treatment, a planning 3D CT is usually created. In some cases, a 4D CT scan (3D plus time) is preferred. Additionally, X-ray CT volumes can be fused with other modalities such as MRI or PET when those modalities provide better delineation of diseased tissue. In order to position a patient in the same reference frame for radiotherapy as they were for the planning CT scan, markers of some kind are

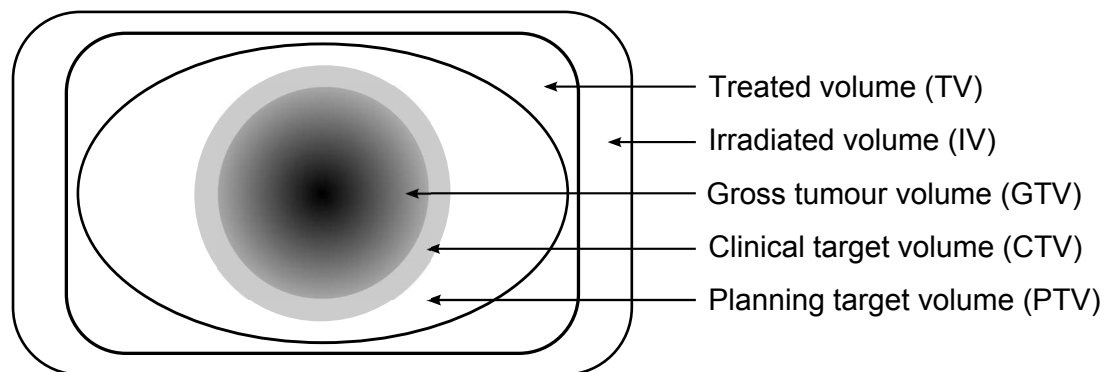


Figure 2.5: ICRU Report 50 abstract representation of volumes

placed on the patient. Often these are small alignment markings tattooed on the sides and stomach of the patient. These are later used to position the patient with respect to treatment room lasers that identify the reference frame origin.

As well as providing diagnostic capabilities, an X-ray CT has the advantage of revealing bony structures better than other modalities. Bony structures such as the vertebrae and pelvis are less subject to motion during image acquisition and therefore provide better landmarks for image registration algorithms. This, in conjunction with the fact that most imaging at dose delivery time is also done with X-rays, makes it useful for *image-guided radiotherapy* described in Section 2.6.3.

Once the planning CT has been acquired, the oncologist delineates the Gross Tumour Volume (GTV) and adds a Clinical Tumour Volume (CTV) using, usually, a pre-defined margin. The GTV is the part of the image where the cancerous tissue is distinguishable to the oncologist. In some cases, the CTV is defined directly because the GTV is not clearly visible. The CTV is where the oncologist expects the cancerous tissue to actually be. Benign tumours usually have distinct edges but malignant tumours often have cellular tendrill-like structures which invade the surrounding health tissue. These structures are often invisible on X-ray images yet it is critical that they be killed and so a margin is added to the GTV based on the knowledge of the oncologist about the tumour type, its location, etc. The oncologist may also delineate important organs that may be especially vulnerable to radiation; these are called Organs at Risk (OARs) and the therapy plan must be devised to minimise their exposure to radiation.

Once the oncologist has marked the CTV and any OARs, a *dosimetrist* develops a conformal treatment plan. The first step is to add a margin to the CTV based on expected intrafractional organ motion and interfractional setup error. This new expanded volume is called the Planning Target Volume (PTV). Often this is performed automatically by the software used in the planning process. This is because in many protocols, the PTV mar-

gins are simply prescribed in adopted standards and it is easier for software to uniformly add 10mm to a CTV than it is for a dosimetrist. To conform the deadly radiation dose to the target tissue, and spare healthy tissue, a plan consisting of a number of different *fields* is created. These fields describe the individual dosages delivered by the megavoltage linac X-ray beams at different gantry angles and possibly at different couch positions as well. On most machines, these fields are coplanar though in some newer devices such as the Cyberknife®<sup>®</sup>, they can be non-coplanar. A field more specifically includes any collimation of the beam, the voltage, the current, the angle or position of delivery, any couch position changes, and the duration of the beam. The terms ‘field’ and ‘beam’ are sometimes used interchangeably in the literature but in this thesis, they will be used as described above. These plans are becoming increasingly sophisticated and can now include enhancements such as gating and dynamic leaf collimation (DMLC). These are detailed further in the section on motion compensation (Section 2.6).

It is worth pointing out some ways in which pragmatic realities can affect the research process. Radically new processes imply new machines which in turn imply training. Both the machines and the training are expensive both in terms of money and in terms of time. Also, when dealing with dangerous or insufficiently understood methods, a sufficient number of studies with the attendant ethical evaluations must be performed before it will even be considered for deployment. Hence a new methodology may take years to reach widespread deployment. A further pragmatic constraint is the limited availability of skilled clinicians. Marking GTVs and CTVs is tedious and time consuming. To ask an oncologist with a heavy patient load to mark up GTVs and CTVs on a 4DCT creates a very serious resource trade-off decision. Some of the research proposing 4DCT as part of their process recognise this and discuss their attempts at automatically creating 4DCT versions from a 3DCT PTV [48].

### 2.5.3 Therapy

Once the plan has been created and verified by the attending oncologist, the patient can begin a course of treatment. Delivery of the treatment is a collaboration between *radiologists* and *medical physicists*. Radiologists actually oversee the dose delivery using machines that are installed, calibrated, and maintained by the medical physicists. Often radiologists will consult with the physicists as well as with oncologists in unusual circumstances.

As mentioned in the section on the biology of radiation, tissue can usually recover from radiation if the dose is sufficiently small and healthy tissue can recover generally

faster than cancerous tissue. This fact is exploited in two ways. In the planning phase, the use of multiple fields allows dose to accumulate where the fields overlap while reducing the dosage in areas of single field radiation. Hypothetically, if four fields were used with a uniform dosage per field, the cumulative effect where they overlap would be 100% while the non-overlapping regions would only incur a 25% dose. Of course real plans are more complicated than this, but the principle is the same. The second way this biological fact is exploited is in temporally dividing the treatment. If the intended dose for a tumour is 2Gy, this can be delivered in two 1Gy treatments. If those treatments are separated in time such that the healthy cells have repaired themselves but the malignant cells have not fully recovered, then the goal of sparing healthy cells while targeting the cancer is furthered. This is the main idea behind *fractional* delivery wherein a patient is scheduled for multiple radiation treatments often called “fractions” or “fractionated radiotherapy” [60, chap. 16].

Multiple treatments pose a significant problem however: patients need to be positioned exactly as they were when the planning CT was acquired each time they have a treatment. One way of attempting this is to line up reference marks made at the planning time with laser beams that represent the same position in space for both the planning CT machine and the radiotherapy linac. Unfortunately over the course of the fractional treatments, patients may lose weight and skin may sag causing the markers to shift around slightly. Likewise, if the treatment period is long enough, the tumour may shrink and shift position. Finally, if the tumour is in the pelvic region, the bladder and bowel changes may affect the position significantly. One way to compensate for these changes is to make the margins for the PTV sufficiently generous. This of course causes additional healthy tissues to be unnecessarily irradiated.

Another way to compensate for tumour position changes and setup errors is to re-image the patient immediately prior to the dose delivery and to reposition the patient if the tumour has moved beyond some prescribed threshold. There is some evidence, at least in prostate patients, that this shifting is ineffective [63]. Conversely, Nelson et al. conclude that setup error plays a more important role than respiratory motion [59]. Such use of imaging during the course of treatment is called image-guided radiotherapy (IGRT). An excellent overview of the state of IGRT can be found in [14].

In addition to setup errors, it is desirable to compensate for intrafractional motion. Many compensation methods have been investigated and these will be detailed in the section specifically on motion compensation immediately below. To summarize in advance, the therapeutic forms of compensation can be categorized as attempting to suppress the motion, waiting to irradiate until the tumour has moved to a specific position, or tracking and irradiating the tumour dynamically.




Figure 2.6: Three classes of motion (a) continuously periodic, (b) discrete and, (c) shift.

Purdie's team [64] experienced changes in tumour motion between the planning 4DCT and the respiratory correlated CBCT reporting a discrepancy of 10 mm in one case. This agrees with the results of two other authors they reference. Sonke, however, reports contradicting results claiming that tumour trajectory shape is stable [86] but reaffirms the problem of setup error.

#### 2.5.4 In summary

Radiotherapy involves a clinical team which devises and implements a fractionated delivery plan of radiation treatments. Delivering radiotherapy in fractions requires aligning the patient to the original planning frame of reference which results in setup errors. Likewise, anatomical changes can occur during the waiting period between fractions resulting in additional sources of misalignment.

## 2.6 Motion compensation

Motion is positional change in time. In radiotherapy, target motion can be compartmentalized into the motion that occurs during imaging or treatment and the motion that occurs between treatments. As a reminder, these are typically called *intrafraction* and *interfraction* motion respectively.

The physiological motions described earlier (Section 2.4) and summarized by Diez have motion types that fall into both categories. However, virtually all of the compensation techniques for physiologically-based motions are only concerned with the periodic forms. This is because motion models can be built and easily correlated with surrogate measurements.

To clarify this situation, an extended motion classification system is proposed and used here. In Figure 2.6, three motion classes are illustrated. The solid line represents some 1D description of motion, e.g. the SI motion. The shaded boxes represent the acquisition times for four representative cone-beam scans. The motion of (a) is continuously periodic.

The motion of (b) does not occur during the scan but rather appears when comparing two scans at different times, e.g. interfraction motion. The motion of (c) represents motion observed during the acquisition but which is not continuously periodic. Note that the motion in (b) could just as easily be represented conceptually as two instantaneous measurements. Motion (a) represents the movement caused by respiration and cardiac cycles. Motion (b) represents changes such as bladder and bowel filling or emptying, tumour shrinkage, and setup error. Motion (c) is what Diez calls “quasi-periodic” motion though this can be extended beyond peristalsis and catastalsis to include random effects such as gas shifting or an uncontrolled patient shift. These three classes of motion will henceforth be termed “continuously periodic”, “shift”, and “discrete” for types (a), (c), and (b) respectively.

Motion, if not compensated for, can create a variety of errors [23]. All of these errors result in the same consequence: target tissues may be insufficiently dosed and re-grow while healthy tissue may be overdosed leading to loss of critical function and/or secondary cancers. In the planning stages, tumour motion during imaging creates blurry images making the delineation of the GTV difficult. The simplest possible compensation method, for therapy, is to enlarge the GTV with a margin thus creating the PTV and to irradiate the entire PTV. If the GTV is incorrectly identified, the recurring theme of administering an incorrect dose is raised. Therefore, if the GTV can be more accurately delineated, the PTV will consequently be more accurately defined. Otherwise, uncertainty in the GTV may necessitate an increase in margins required to construct the PTV.

Compensation, then, attempts to do any or all of the following: improve the images used in radiotherapy, improve the accuracy and precision of dose calculations, and improve the accuracy and precision of delivered therapy.

The ideal solution is to eliminate or minimize motion and this is attempted with various restraining methods. The breath-hold technique is commonly used both for capturing the planning CT and during treatment. This technique is self-described remarkably well: a patient literally holds their breath while the relevant activity ensues. Ironically, many of the patients who would benefit from this technique have some form of lung disease which makes breath control a challenge and therefore this technique may not be an option. In some regions of the body - head, neck, and liver in particular - it is sometimes useful to physically restrain the patient with various devices such as vacuum pillows [61] or custom fabricated moulds [44]. This branch of motion compensation is not relevant to the remainder of the thesis and will not be discussed further.

Another interesting branch of research that plays an important role in both diagnostic and planning stages is the fusion of different imaging modalities to better assess diseased



tissue location. However, as with motion restraint ideas, this is not particularly relevant to the work yet to be described and so only the research that somehow relates to motion identification or compensation in CBCT will be reviewed below.

## 2.6.1 Continuously periodic motion compensation: IMRT

Most of the work in this area ultimately falls under the catchphrase “intensity-modulated radiation therapy” or IMRT. Technically, IMRT covers everything from using more than one field to irradiate the patient to advanced dynamic multi-leaf collimation linked with surrogate signals to track tumour motion in real-time. This section will not cast its net so widely but will restrict itself to those ideas that influence or impact the research in cone-beam CT. For the broader picture, a thorough review of IMRT can be found in Webb [98]. In particular, Figure 1 from that paper illustrates the various tumour position identification techniques being used. As in the section on discrete motion below (Section 2.6.3), the simplest approach to compensation for motion is to increase the margin used to construct the PTV. This is both the historical and current practice. Newer and better ideas are being evaluated of which several will be described in this section.

### 2.6.1.1 Gating

Gating the motion: for the most significant cause of motion — respiration — monitoring the breathing pattern and waiting for a specific phase is frequently used in both planning and treatment. For planning purposes, image slices or projections can be acquired only during the specific phase as identified by a variety of external devices (respirometers, reflective markers, etc.) For treatment purposes, the same triggers can lead to turning the X-ray beam on and off again essentially irradiating the tissue only when it reaches a certain position. The principle downside to this strategy is that significant additional time is required both for treatment and the initial planning CT acquisition. Gating during acquisition or therapy is often distinguished from post-acquisition gating of projections by the term “prospective gating” versus “retrospective gating” for the latter process.

Li introduces something they call dose shaping (DS) based on convolving a 4D motion probability model, obtained with 4DCT (discussed next), with the PTV [41]. They compare the results with both gating and conventional margin extension and conclude that DS performs better.

The main problem with gating, especially in therapy, is that it increases the time required to deliver the fraction. Especially in respiratory gating, much of the actual time is spent waiting for the breathing to reach the desired state. This increases the probability

that the patient will shift or move thus necessitating a re-evaluation of the setup and a possible repositioning before continuing. Dynamically gating the motion using dynamic multi-leaf collimators (dMLCs) is one attempt to remove the time penalty of gated radiotherapy. If a sufficiently accurate 4D model can be developed that correlates well with the patient's breathing, then a surrogate can be used to continuously track the tumour by proxy while dynamically collimating the radiation beam to match the hypothesized tumour motion [11, 34]. This method requires a consistent and reproducible respiratory pattern. It also assumes the surrogate is a good indicator of actual tumour motion [98].

The fundamental idea behind gating, tracking motion in time based on its periodically repeating property, is also the main trick used to add the fourth dimension to CT.

### 2.6.1.2 4D CT

Four-dimensional CT, sometimes referred to as “ $3D + t$  CT” where  $t$  is time, is not as straight-forward as one might imagine. A video sequence is analogously a 3D object in that it consists of  $2D + t$  data. The subtle difference is that each sample in time – each frame – of a video sequence is complete whereas each volume in a  $3D + t$  sequence must be first reconstructed from lower-dimensional samples *acquired*. The acquisition of these 1D samples and the subsequent reconstructions into 2D slices used to construct 3D volumes requires a relatively long period of time. The result is that CT volumes cannot be constructed in the same “instantaneous” manner analogous to the capture of the 2D images in a video sequence. Consequently, the goal of 4DCT is not a true capture of 3D volumes over the recording timeframe; rather the goal is to construct a model of motion over an appropriate period of time. For respiratory motion, this appropriate period would be a full respiratory cycle. For cardiac motion, the period would be a complete cardiac cycle.

The ideal way to accomplish this is dependent on the method of image acquisition but all forms effectively rely on some common principles. The period of interest is subdivided into phases (or amplitudes) and then samples are acquired at each *position* for each *phase/amplitude*. For instance, with CT a complete set of slices for a breathing cycle are acquired at each slice position. These slices are then sorted by phase and the phase volumes can be constructed by stacking slices from the same phase together [67, 95]. Respiratory cycles are typically subdivided into between eight and ten phases resulting in eight to ten times the number of acquisitions needed to build a motion model. Many authors continue to use phase-binning though [1, 47, 102] assert that amplitude binning is more accurate than phase binning.

Assigning slices to a subdivided phase, henceforth referred to as “binning”, is not

always easy. For cardiac reconstructions, the electrical signal of the heart measured in electrocardiograms is sufficient for binning cardiac slices. Breathing has no universally accepted direct measure of phase and so surrogates are used. Two very popular external surrogates are spirometers which measure tidal volume [44, 45, 47] and optically tracked markers positioned on the chest [3, 53]. The most general internal surrogate is the diaphragm which requires proper automatic identification in order to correlate with phase [64, 87]. This method, however, is challenged by [59] who claim that the diaphragm moves air, not tumours, and that tissue moves and deforms differently from respiration.

One interesting refinement of this general idea is proposed by McClelland et al. [53]. In addition to the initial 4DCT dataset, they acquire a relatively static volume using a breath-hold technique. The 4D volumes are then non-rigidly registered with the static volume and a continuous motion model is then constructed by fitting a temporal B-spline to the deformation parameters. Rietzel [66] does something similar but does not require the static volume step.

Zeng et al. [102] attempt to eliminate the use of external surrogates and propose an iterative technique that estimates two initial extreme reference volumes from slice centroids. Their method was comparable to surrogate-guided methods in three out of five tested cases.

Sarrut et al. [75] and Schreibmann [77] proposed a dose-saving idea of collecting only maximum and minimum inspiration volumes and constructing a 4D approximating motion model from them. Sarrut added a prior lung density model to their densely deformable registration algorithm and claim improved accuracy. Both report encouraging results. Keall et al. [34] demonstrate an application of such a model with respect to planning for dynamic multileaf collimation (dMLC).

While the use of 4D CT has the potential to increase the accuracy of dose planning, Rietzel et al. explain that it is limited by the increased workload required to delineate volumes [67]. They propose solving this by nonrigidly registering the 4D CT volumes to construct a motion model and then apply this to 3D planning volumes to automatically generate 4D plans and dose assessments.

One particularly interesting line of research is found in [101, 102]. They estimate deformation parameters on a B-spline model by projecting deformed reference CT and minimizing the least square error between measured and estimated. This is done in an iterative fashion and has relevance to one of the techniques proposed in Chapter 6.

Regardless of the details, all of these ideas essentially take the oversampled CT slices, bin them to create sequential volumes at different stages of motion, then build (for some) a motion model by registering these volumes, usually using some form of B-spline registration algorithm.

### 2.6.1.3 4D reconstruction of cone-beam

It is natural to consider applying the idea behind 4D-CT to cone-beam CT to create 4D-CBCT. In 4D-CT, reconstructions are generated at each slice for each phase of the periodic motion. Given the typical 8-10 phase subdivision, to accomplish the same thing in CBCT would require acquiring 8-10 volumes. There are three reasons why this is not feasible. Firstly, the clinical time required is too great. Typical CBCT volumes require approximately two minutes to acquire; a twenty minute protocol would be too consuming of a limited resource. Secondly, CBCT X-ray tubes cannot run continuously and ten volumes could only be acquired with sufficient pauses to allow the X-ray tube to cool thus exacerbating the time-resource problem as well as contributing significantly to mechanical wear. Thirdly, the dose increase makes such an implementation too risky. The dose for CBCT can be estimated from Islam's reporting [27] as approximately 46 mGy which is roughly in agreement with Lu's reported dose of 57 mGy [46]. Lu then goes on to compare CBCT with 4D-CT and concludes that one CBCT scan is approximately equal to .75 4D CT scans. If this is true, then constructing 4D-CBCT from ten volumes is equivalent to more than seven 4D-CT scans delivered back-to-back. Dose calculations are very difficult but given the on-going debate over the frequency with which 4DCTs should be administered, this naïve 4D-CBCT method would clearly be too much dose. Li attempts to circumvent the dose problem by acquiring CBCT images at lower currents to lower the dose [40], however they have not addressed the expected decrease in CNR in their work and they acknowledge the clinical throughput issues. In consequence, most proposed 4D-CBCT simply sort projections from a single scan using the same kinds of surrogates used for 4D-CT [19,40,46,87, and many others]. This approach creates a variety of problems due to the need to reconstruct phase volumes with significant numbers of missing projections. The consequent artefacts are discussed in detail in Chapter 4.

Cone-beam CT has from its inception been seen as a useful tool for compensating for setup errors (see Section 2.6.3 below). Recently, to compensate for the lack of projections, several authors have proposed applying prior motion models derived from 4D-CT to deform the out-of-phase CBCT volumes. Rit shows results compensating for respiratory motion for both analytical and algebraic reconstruction techniques [69]. Li, Koong, and Xing propose using the planning CT as a baseline volume and then registering the initial phase-binned CBCT volumes to it to construct a motion model [38]. Hypothetically, this model is then used to deform the 2D CBCT projections before reconstructing to the baseline phase. They demonstrate this in a 2D simulation.

Just as with 4D-CT, the question of binning by phase or binning by amplitude is important. Rit [71] describes the amplitude/phase binning trade-off by concluding that, for

cone-beam, amplitude binning introduces more artefacts while phase binning introduces more blur.

One concern of using the 4D-CT motion model to compensate for deficiencies in CBCT is that, over the course of the fractionated delivery, the motion model may become invalid. Purdie asserts their results show the relationship between external surrogate and tumour motion does not hold over multiple fractions [64]. Sonke seems to disagree by reporting rather the opposite [87]. Clearly, further investigation into this possibility is required before it can be relied upon in a clinical setting.

In 4D-CT, the diaphragm is used as a surrogate [64, 87] and Rit uses the same idea to correlate cone-beam projections [69, 70]. Use of the diaphragm has the advantage of not requiring any additional equipment though it involves an additional pre-processing step. Thus it is applicable to retrospective studies where surrogate signals have not been captured and where real-time requirements do not exist. Perhaps the only criticism of this technique is the requirement that CBCT projections contain enough of the diaphragm to successfully be tracked. As hospitals and regulatory agencies become increasingly concerned with radiation dose, the desire is growing to collimate CBCT projections to just the portion that windows the tumour.

Grangeat et al. [25] propose speeding up the rotation of the gantry and collecting multiple sequences effectively treating CBCT like 4D CT. They partially compensate for the increased dose by acquiring only multiple half-turns rather than multiple full rotations. Their motion compensation method relies on a dynamic particle model. Their work is of particular interest because they propose the multiple rotation idea, albeit for completely different reasons, which is proposed in Chapter 6. They also make use of a region of interest (ROI) as is done in this work, though they do not use the PTV as the justification and basis.

#### **2.6.1.4 C-arm**

Recently, a portable form of CBCT called “C-arm” scanners has been introduced principally in cardiac applications. Motion compensation is therefore a critical concern in that context. Most of the methods used for compensation make use of the fact that injected contrast agents are commonly used in cardiac environments. The use of such agents facilitates the identification of vascular structures which are then used as proxies for cardiac motion in general. The relative sparseness and distinctiveness of these vascular trees allow gradient based optimizers to be used in finding solutions to parameter sets that drive the motion compensation.

For instance, Rohkohl et al. [72] use a 4D B-spline model to warp individual backpro-

jection volumes in much the same way that Rit and others have done with CBCT (See [48] and [69]). Rather than linking the volumes to reference 4D CT volumes, as is commonly done with CBCT, their objective function compares forward projections of the best estimate with the original images. This work extends the earlier work by Blondel [7, 8].

Zhang et al. [103] present ideas similar to Rohkohl but add a region of interest (ROI) to improve performance. They also reorder projections from multiple cardiac cycles into a single normalized cycle. This makes the assumptions that the periodic motion of the heart is consistent and that additional respiratory can be removed with a breath-hold technique.

Metz et al. [54] mimic the work done in CBCT by registering the C-arm projections with a pre-operative CTA data. They do so by using a novel 2D+t to 3D+t registration approach. The segmented vascular tree aids in the deformable registration.

## 2.6.2 Shift motion compensation

The extent of compensation for this type of motion is essentially to attempt to measure it and then include knowledge of the motion in the planning volumes. As was discussed in the section on prostate motion (Section 2.4.3), it is generally acknowledged, though not universally agreed, that respiratory motion has no significant impact on the prostate. Therefore, continuously periodic models and some of the compensation ideas described in the prior sub-section are of no consequence.

Shift motion is, by its very definition, unpredictable. Therefore, any approach that attempts to correlate motion from some model at time  $t_1$  with motion at time  $t_2$  later cannot succeed. The only reported ‘surrogates’ for motion are implanted fiducial markers which is a highly invasive approach [51] and which is susceptible to displacement [59].

The result is this: there is no reported method for binning either CT slices or cone-beam projections into some kind of motion bin for unpredictable, non-periodic motions.

## 2.6.3 Discrete motion compensation: IGRT

Much of the work in this area falls under the catchphrase “image-guided radiotherapy” or IGRT. Fundamentally, IGRT seeks to mitigate the effects of external positioning errors caused by sagging skin, weight loss, etc. and the consequences of internal tumour shape and position changes [14].

Positioning errors have historically been “compensated for” by simply adding a setup error margin to the PTV at planning time. A more active approach, and one that some research has shown to be promising, is to actually reduce setup error margins through the use of IGRT [15].

The two most common imaging methods used to guide radiotherapy are the older portal imaging systems using orthogonal 2D MV X-ray images which are registered with simulated projections taken from the planning CT, and the newer kV cone-beam CT volume which is rigidly registered with planning CT volume directly. Borst et al. compared portal imaging with CBCT imaging and concluded that CBCT was more accurate at identifying and correcting setup errors [9].

A problem with CBCT in IGRT is that motion occurs during acquisition of the pre-treatment image volume resulting in streaking artefacts [37, others]. These artefacts cause difficulties for registration algorithms. Consequently systems such as the Synergy machine deployed at St James Hospital in Leeds, England, are sometimes equipped with different registration algorithms including a purely manual option. One of the outcomes of Rit's work referred to earlier ([69]) is a better reference CBCT for registering with the planning CT. This improved reference CBCT is created by applying the 4D-CT generated motion model to the individual CBCT phase volumes via a deformation field which results in all CBCT phase volumes behaving as if they came from the same reference phase. These can then be combined which reduces the streak artefacts caused both by the motion and missing projections.

#### 2.6.4 In summary

- An extended motion classification system is introduced using three classes: continuously periodic, shift, and discrete.
- Motion compensation involving restraints, though commonly used, are not discussed.
- Fusing alternate modalities with X-ray imaging is also explicitly removed from discussion.
- IMRT is introduced.
- Gating, both prospective and retrospective, is introduced as both a planning and therapeutic adjunct. Gating paves the way for 4D-CT.
- 4D CT is the construction, from samples acquired over many periods, of a single period's motion.
- Respiration motion is generally subdivided into eight or ten phases or amplitudes. Such subdivision is called "binning" in this thesis. Research tends to use phase binning though several papers report that amplitude binning is better.

- External surrogates such as spirometers and markers and internal surrogates such as the diaphragm are used to correlate the 4D model of the motion in a representative period with the expected tumour position.
- 4D CBCT uses the binning ideas from 4D CT but not the multiple sampling ideas. Projections are binned by phase or amplitude using a surrogate of some kind. Binned projections are reconstructed into a phase-specific volume.
- shift motion is poorly studied and typically requires implanted gold markers to study the motion.
- Cone-beam is rapidly becoming the modality of choice for IGRT.

## 2.7 Chapter summary and thesis problem

In this chapter, the nature of X-rays from the standpoint of imaging and clinical therapy has been discussed. The biology of radiation has likewise been reviewed to the level sufficient for understanding the goals and problems of radiotherapy. Radiotherapy as a form of treatment for cancer has been reviewed with special emphasis on the GTV, CTV and PTV volumes and the nature of fractionated delivery. The problems created by motion on imaging, planning, and treatment in radiotherapy have been reviewed. Because of the many different kinds of motion, a new way of categorizing motion was proposed to better clarify one aspect of the problem space this thesis addresses. As a precursor to discussing motion compensation, the principal forms of X-ray imaging have been quickly described with an emphasis on cone-beam CT, the modality used by the research in this thesis. With the groundwork laid, the main ideas currently proposed to compensate for motion were then reviewed.

As has been shown, the principal technique when dealing with motion in CBCT is to bin the projections by phase or amplitude. This technique implicitly presumes periodicity for the motion and relies on some external or internal signal surrogate to assign, either actively or retrospectively, the projections to some bin. Reconstruction is then performed using only the projections assigned to a given bin though some methods seek to fill missing projections with projections deformed by a motion model derived from a 4DCT sequence. In the following chapters, a novel method for binning projections without any model will be presented. This method has the following advantages. It allows CBCT projections to be collimated to the PTV which may exclude the diaphragm and thus reduces dose relative to methods that use the diaphragm as an internal surrogate. It



is performed without the use of externally created motion models thus eliminating setup errors or errors from the degradation of the motion model over the course of the fractions. It creates the potential for capturing a third class of physiological motions not addressed in any other method.

# Chapter 3

## Materials

---

### 3.1 Introduction

In this thesis, phantoms are used to test and evaluate several different methods. In medical imaging studies, a phantom can be either a physical device or a digitally synthesized object. Both of these represent some aspect of the human body and facilitate experimentation without the necessary approval processes needed when using patients and without any risk to patients when new protocols are explored. Phantoms have the advantages of providing ground truth measurements, of being able to precisely control experimental parameters, and — in the case of digitally synthesized phantoms — being inexpensive. The disadvantage of phantoms is that they are usually much simpler than real human data and algorithms that work on phantom-derived data must eventually be tested on human data. However, by testing ideas on phantoms, bad ideas can be inexpensively and quickly removed from consideration while good ideas can be refined before any human trials. In this work, three different phantoms of increasing complexity are used to explore the continuously periodic and shift motions described in Chapter 2. Experiments are performed using a virtual phantom (Section 3.3) constructed of simple mathematical solids. A physical phantom (Section 3.4) combined with a novel animation method is used to study simple motions using real projections. Finally, an anthropomorphic digital phantom, the NCAT 2.0 (Section 3.5), is used to study the more complex structures of the human anatomy.

Each of these phantom types can be “configured” according to the needs of the experiment and the unique capabilities of the phantom. A core component of each configuration is the motion that is both simulated and recorded for later use. The purpose of digital phantoms is ultimately to generate a set of projections that are sufficiently close approximations of the cone-beam projections that would be generated if the phantom were actually scanned.

These phantoms and the configurations constructed are described in this chapter.




Figure 3.1: Geometry of small-field Elekta Synergy system. All measurements in mm.

## 3.2 CBCT geometries using the Elekta Synergy

In Section 2.3.3.2, the general geometry of cone-beam CT is documented for the purpose of understanding FDK reconstruction. For the purposes of the experiments in this thesis, the specific cone-beam CT geometry used with the Elekta Synergy image guided radiation therapy systems is described. This configuration forms the basis for the geometries used with the virtual phantoms. The Synergy consists of a kV X-ray source that can be statically collimated and pre-filtered and is located 1000 mm away from the isocentre. A square flat panel detector consists of  $1024^2$  detectors measuring 409.6 mm on an edge and is located 535 mm away from the isocentre. This is illustrated in Figure 3.1. For most clinical scans, the projections are downsampled to  $512 \times 512$  pixels. For small field scans, the type used in the experiments here, the detector’s centre is in line with the isocentre and the source. Small position variations due to mechanical motion are recorded as part of the maintenance calibration of the machine and incorporated into the reconstruction process as angle-specific scanner offsets.



Figure 3.2: Elekta Synergy system at the St. James's University Hospital, Leeds, England.

### 3.3 Simple virtual phantoms

The idea and some of the code used to create the synthetic phantoms comes from the work of Jens Müller-Merbach [58] who constructed a static cone-beam geometry X-ray simulator. This work was extended by Henrik Turbell [90] who added more general source-trajectory paths and included additional geometric objects. Only a subset of the features provided by the TAKE simulator, as it is named, were needed while the ability to incorporate motion and noise was missing. The C source code was significantly modified to fit in with the custom written C++ framework implementing the FDK algorithm.

Using the simple virtual phantom framework, models are constructed from geometric shapes and projections are generated analytically. Motions of objects are sinusoid-based and composed of independent motions in the  $x/y/z$  directions. Simulated projections are constructed analytically. Of the three phantom types, this phantom gives the greatest flexibility in terms of simulated tumour shape, contrast, and motion. It also eliminates the structure “noise” introduced with the other phantom types used in this research. For example, the physical phantom captures the couch support grid in many projections while the anthropomorphic virtual phantom includes virtual ribs, etc. Its principal weaknesses

are its perfectness: object surfaces, shapes, and densities are completely smooth; motion is too regular; the aforementioned structures that do exist in real-world projections are not present; and the analytically generated projections simulate a monochromatic energy source with no beam hardening or characteristic radiation.

### 3.3.1 Description

Three geometric object types can be instantiated when constructing virtual phantom models. These are: ellipsoids, cylinders, and boxes. Each object has position parameters, scaling factors, rotation parameters, and a density factor. For each of these parameters, an independent motion can be attached.

Motions are designed to represent the motions under study: the continuously periodic respiratory motion and the shift motion captured during a CBCT scan. The simple motions used for this phantom type are:

**Lujan** : Lujan [49] has proposed a  $z(t) = z(0) - A \cos^{2n}(\frac{\pi t}{\tau} - \phi)$  model that has been widely used [19, 41, 86, etc.]. A simplified version of this motion is driven by a  $A \cos^{2n}(\frac{\pi t}{\tau})$  function of time  $t$  controlled by period  $\tau$ , amplitude  $A$ , and exponent multiplier  $n$ . Combining two such motions of the same frequency, one in the SI direction and one in the AP direction, provides a simple simulation of actual lung tumour motion.

**Shift** : this motion creates a simple motion shift of a given amplitude at a specific point (percentage) of the scan.




Figure 3.3: Simple virtual phantom motions: (a) Lujan simplified respiratory model shown with  $\cos^2$  (solid line),  $\cos^4$  (dashed line) and  $\cos^6$  (dotted line) forms; (b) quantum-shift motion occurring at different points during a scan.

### 3.3.2 Projection construction

Synthetic phantoms are used to construct projections that simulate the Elekta Synergy® CBCT machine described above in Section 3.2. Normal clinical protocol involves acqui-

sition along a 360-degree arc over the course of two minutes collecting approximately 670 projections during one scan. Usually, in a patient setting, a medium field of view is used which means only part of the body is imaged with each projection. This is necessary to reconstruct a volume which is larger than what can be reconstructed using the natural geometry given the detector size. For the purposes of this thesis, the simpler, natural reconstruction is used which is referred to as a “small field” acquisition. In addition, a bow-tie filter matched to the small or medium field is usually used when acquiring data from the Synergy machine. Experiments confirm that this is automatically compensated for by the machine and so is not included in our simulation.

As noted in Section 2.2.1, the X-ray attenuation for monochromatic sources can be modelled as an exponential decay via Lambert’s law which is then linearized:

$$-\ln\left(\frac{I}{I_0}\right) = \sum_{k=1}^n \mu_k x_k \quad (3.3.1)$$

An unoccluded image acquired from the Synergy machine is used to create a reference point for  $I_0$  which is then used in the synthetic simulations. A virtual detector pixel is calculated by summing the individual attenuation contributions of each object in the phantom that intersects the ray from the source to the pixel and taking the exponential of its negative. Each object’s contribution is simply the length of the intersection of the ray with that object times its attenuation parameter.

### 3.3.3 Simple virtual phantom configurations




Figure 3.4: Transverse, coronal, and sagittal slices from an example reconstructed Simple phantom configuration.

Phantom configurations have been constructed to reflect both the respiratory motion and the sudden-shift motion that is unique to this research. Ideally, the controlling parameters should mimic known physiological behaviour as closely as possible. Given that

uniqueness is claimed for studying sudden-shift motion, very rough estimates are constructed from documented interfraction motions. Descriptions of lung tumours and their motion in the literature is inconsistent. For instance, Seppenwoolde [81] records the results from 20 patients where the average tumour size is  $53\text{cm}^3$ , the standard deviation is  $72\text{cm}^3$ , the median size is  $20\text{cm}^3$ , the minimum size is  $2\text{cm}^3$  and the maximum size is  $257\text{cm}^3$ . Breaking this down into lower, middle, and upper lobe regions, the average tumour sizes are (respectively)  $69\text{cm}^3$ ,  $22\text{cm}^3$ , and  $48\text{cm}^3$  and after removing outliers, these become  $25\text{cm}^3$ ,  $23\text{cm}^3$  and  $25\text{cm}^3$ . Li [41] reports the average size for 6 patients, ignoring one outlier tumour size of  $20\text{cm}^3$ , is  $206\text{cm}^3$ . This represents nearly an order of magnitude of difference between two reported studies. Motion magnitude is likewise varied. In Seppenwoolde, the average magnitude is  $5.3\text{mm}$  which is quite close to Langen's [36] reported  $5.2\text{mm}$  value. In contrast, Li reports a  $12\text{mm}$  average magnitude while a different Li [39] used displacements with a magnitude of  $57\text{mm}$  for their lung phantom motion. Reported respiratory periods are somewhat more consistent ranging from 3.7 seconds to 6.5 seconds. Chang et al. [10] best summarizes this by observing for that study on 8 patients, "the breathing periods for these patients were 4 to 6 s." Some of the variability of reported breathing patterns can be attributed to different clinical protocols. Shallow breathing is encouraged in some cases while normal breathing is supported in others. The simple phantom configurations used in this thesis are constructed to represent a cross section of these various parameters. Tumour size is either *average* (18 mm sphere) or *large* (34 mm sphere). Motion amplitude is either *average* (15 mm and 8 mm in the SI, AP directions) or *short* (5.5 mm and 2.2 mm in the SI, AP directions). The respiration period is usually 4-seconds but one 6-second configuration is included.

For the shift motion in configuration 5, an average prostate size and the displacements are chosen from Langen's survey of organ motion [36]. The "prostate" is modelled using an ellipsoid measured 30, 25, and 20 mm along the three ellipsoid axes. This is meant to represent a wide variety of tumours in the pelvic region where shift-motions are important.

1. *Average* lung tumour size, average motion amplitude, Lujan motion in SI and AP directions, using a 4-second period.
2. *Average* lung tumour size, short motion amplitude, Lujan motion in SI and AP directions, using a 4-second period.
3. *Large* lung tumour size, short motion amplitude, Lujan motion in SI and AP directions, using a 4-second period.
4. *Large* lung tumour size, average motion amplitude, Lujan motion in SI and AP

directions, using a 6-second period.

5. 10mm Shift motion, 50%, high contrast.

### 3.4 Animated physical phantoms

Animated physical phantoms are constructed from actual projections of a Quasar® physical phantom acquired in the Synergy CBCT machine. These projections are acquired in individual steps as the position of the phantom is physically changed. This allows motion to be simulated in a fashion similar to how stop motion animation is used in motion photography. Animated phantoms remove the weaknesses associated with simple virtual phantoms but introduce the constraint that motion can only occur along the one path traced by the physical placement of the phantom. The motion along this path, however, can be completely arbitrary and is not restricted to the set of sinusoid-based functions used with the simple virtual phantoms.




Figure 3.5: Quasar® phantom with cedar insert.



### 3.4.1 Description

The Quasar respiratory motion phantom was used to generate the projections needed to construct the animated phantoms. This phantom has an acrylic body with two cylindrical insert openings and an optional respiratory motion platform (see Figure 3.5). Because motion is generated through animation, the motion platform was not used. In the outer position, a neutral acrylic insert was positioned. A cedar insert, which emulates lung densities and contains a spherical object to simulate a lung tumour, was placed in the middle insert position.

### 3.4.2 Acquisition details

The cedar insert was manually moved inferiorly along the axis of rotation in 1 mm increments over the course of 16 mm and a full set of projections was acquired for each static position. The increment size was chosen because of the limitations of moving the insert literally by hand and because of the availability of a reliable mm ruler. The 16 mm span was chosen because it includes all but the most extreme ranges of motion observed in prostate studies recorded by [36]. This allows a projection to be selected from any of the 16 positions and for any chosen angle within a small variance. By selecting the correct set of projections, any kind of motion along the axis of rotation can be simulated constrained by the granularity of individual steps and the overall distance.

Each scan was acquired at 120 kV and 20 ma over the course of an average of 124 seconds resulting in an average of 678 projections per scan. Due to a “warming up” effect of the X-ray source, the first few projections were often not usable since the resulting attenuation was significantly greater than nearby comparable projections obtained after sufficient warm-up. Papers that report projection counts tend to use either 670 projections [9,87] or 660 projections [51] for 360-degree scans. For this thesis, the 670 figure is used.

### 3.4.3 Construction

A configuration file for this kind of phantom consists of a set of projection angles and the position number to draw the projection from. This configuration file can be generated manually, programmatically using some parametric process, or via a simple drawing utility that translates hand drawn motion into the configuration file. For this work, the configuration files are generated programmatically. In contrast with both the simple virtual and the NCAT phantoms, projections for the animated physical phantom are not synthesized from the phantom configuration but copied from the sets of pre-acquired scans. For

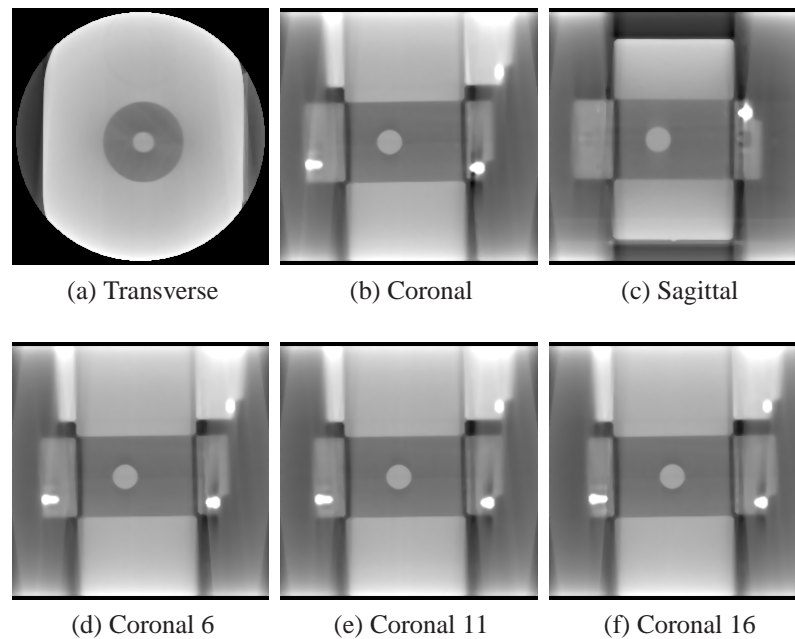


Figure 3.6: Top row: transverse, coronal, and sagittal slices from an example animated physical phantom configuration. Bottom row: coronal slices from three static positions of the phantom.

a given projection at position  $N$  (where  $N$  ranges from 1 to 16 representing the 16 1-mm steps), the projection in that set whose acquisition angle is the closest to the specified angle is copied to the new animated projection set. This is illustrated in 3.7.

### 3.4.4 Animated physical phantom configurations

The goal of this phantom class is to evaluate the effectiveness and limitations of the motion binning methods proposed in Chapters 6 and 7 for the shift motion type. As described in Section 2.6, this type of motion is not well understood so estimates of motion range are based on reported interfractional motion [10, 36, 86] and private communications. Given the method that will be introduced in Chapters 6 and 7, one of the important parameters to study is the point where the shift occurs. Two points are chosen: the 50% position, the 70% position where percentage is the portion of time during the scan(s) preceding the prescribed motion shift. To that end, the following configurations are used in the experiments:

1. Single 16 mm motion shift occurring 50% through the acquisition: this configuration represents a best-case scenario wherein the motion is large and there is an equal




Figure 3.7: Reconstruction from animated sets using a 1D motion model.

distribution of the sampling of the two motion states.

2. Single 12 mm motion shift occurring 50% through the acquisition: this configuration retains the best-case distribution but begins to test the range of detectable motion.
3. Single 6 mm motion shift occurring 50% through the acquisition: this configuration approaches what may be a realistic average motion amplitude while keeping the distribution of projections between the motion states constant for this sequence of tests.
4. Single 16 mm motion shift occurring at the 70% acquisition point: this configuration maximizes the motion while examining the impact of fewer projections being available to construct the motion of the second position.
5. Single 16 mm motion shift occurring at the 90% acquisition point: this configuration maximizes the motion while examining the impact of significantly fewer projections being available to construct the motion of the second position.

### 3.5 NCAT anthropomorphic phantoms




Figure 3.8: Transverse, coronal, and sagittal slices from an example NCAT phantom configuration.

The 4D NCAT phantom is an extension of the earlier work [78–80] by Paul Segars on the 4D XCAT phantom. The 2.0 version has been available for academic research which was used with the kind permission of the Johns Hopkins Medical Institutes. The NCAT phantom is an anthropomorphic 4D phantom of the body beginning mid-femur and encompassing the head. Surfaces are modelled using Non-Uniform Rational B-Splines (NURBS) and include major organ and bone structures that move in a realistic fashion. The two motion drivers, cardiac and respiration, are parametrically controlled. For this research, volumes are generated at specific time points and projections are numerically generated at specific angles to simulate a CBCT acquisition. Like the simple virtual phantom, a weakness of the use here of the NCAT phantom is the simplistic model of monochromatic X-ray energies. The principal advantage is the human-realistic structures which allow experimentation with the impact of high density bone structures and realistic tissue deformations.

#### 3.5.1 Description

Phantom configuration with the 4D NCAT is extremely flexible. At a higher level, parameters exist which dictate the duration of the respiration cycle, the period of the heart beat, which motions to include, what part of the phantom to generate, how extended the diaphragm motion should be, how much detail should be generated for the lung brachia, etc. Well over a hundred such global parameters exist. Individually, there are parameters to control the activity response (MRI perspective) or density (CT perspective) in indi-

vidual anatomical structures.. Individual volumes representing the state of the phantom at specific moments in time are subsequently generated based on these parameters. Tumour simulation is achieved by specifying a tumour location and diameter. All tumours generated by the NCAT system are spherical. Using the same parameters that dictate the phantom, individual tumour volumes are generated with the motion of the tumour occurring as if it were embedded in the anatomy. Construction of a phantom containing the tumour(s) is a simple matter of adding voxels in volumes paired by time. The density of the tumour is controlled by the process that merges the tumour volumes with the anatomical volumes. Tumours placed anywhere within the left lung, the right lung, or the liver are “moved” with the same motion controlling those organs. Sizes for *average* and *very large* have been derived from the reporting of Seppenwoolde [81] and Li [41].

### 3.5.2 Construction

Given a two-minute simulation, constructing two minutes worth of volumes (i.e. approximately 670) is both too time intensive and too disk intensive. Instead, a single representative breathing cycle is generated and logically concatenated. For convenience, the location of the phantom being generated and its size are modified to fit into our small-field geometry. Projections are then synthesized using a simple forward-projection method. This method constructs a “ray” from the source to a detector and samples values in the space along the ray. The sample size is one quarter the width of a reconstruction voxel so it oversamples the voxels at least four times. Sampling is done using bicubic interpolation.

### 3.5.3 NCAT phantom configurations

The goal of this phantom class is to evaluate the effectiveness and limitations of the motion binning methods proposed in Chapters 6 and 7 for the continuously periodic motion type. Specifically, the aspect of this motion type that is of interest is the motion that, if collimated in the cone-beam projections, would wholly or partially exclude the diaphragm.

1. Upper lobe tumour of *average* size: upper lobe tumours do not move as much as middle and lower lobe tumours but, when collimated, are likely to exclude visibility of the diaphragm.
2. Middle lobe tumour of *average* size: this region includes greater motion and will likely exclude the diaphragm at least partially.

3. Lower lobe tumour of *average* size: the greatest motion for lungs occurs in the lower lobe. This is also most likely to include diaphragm visibility in collimated projections.
4. Lower lobe tumour, *very large*: the motion of a large tumour in the lower lobe will be large but relative to the tumour size will be less than an average sized tumour. This is likely to negatively affect the algorithms in Chapters 6 and 7 so this configuration is used to better understand that impact.

### 3.6 Summary

To facilitate the experiments in the remaining chapters, three classes of phantoms are used to first test the ideas (with the Simple phantoms), then to test impact of using actual CBCT projections (with the Animated phantoms) and finally to see if the various proposals work on more human-realistic data using the NCAT phantoms.

A novel idea presented in this chapter is the use of a hybrid process to virtually create a very large set of new phantom out of an existing small sequence of acquired scans. In the case presented, 16 scans were captured with an average of 670 scans which can then be used to generate any of  $16^{670}$  different possible results.

# Chapter 4

## Projection Gap Artefacts in Filtered Backprojection Reconstructions

---

### 4.1 Introduction

The use of binning as a way to separate object motion into states in the reconstruction process was reviewed in Chapter 2. In the case of standard CT, this is accomplished with slice binning. In CBCT, it is accomplished with projection binning. In both cases the slice or projection are assigned to a specific bin where the bin represents some distinct state of motion. Each collection of slices or projections is then used to construct a volume for the associated bin. In many cases (see Section 2.6), the subsequent volumetric images form the basis of some kind of motion compensation.

The reconstruction process is, unsurprisingly, imperfect. Various processes contribute to corruption of the resulting image and these errors are called *artefacts*. In this chapter, the most common artefacts and their causes are reviewed prior to describing a new kind of CBCT-specific artefact induced by the binning process. This artefact, which will be subsequently called a *gap artefact*, has not been sufficiently described nor evaluated in the existing literature.

In this chapter, the effects of gap artefacts on filtered backprojection reconstruction will be demonstrated as an independent type of artefact and then two proposed methods for mitigating its effects will be described and evaluated.

## 4.2 Artefacts in reconstruction

In their survey of artefacts in CT [6], Barrett and Keat begin by mentioning one of the most significant problems with artefacts: “Artefacts can seriously degrade the quality of computed tomographic (CT) images, sometimes to the point of making them diagnostically unusable.” This consequence justifies much of the research dedicated to preventing or mitigating artefacts. They subsequently describe the following classes of artefacts:

### 4.2.1 Physics-based artefacts

These artefacts are caused by the properties of X-rays discussed in Section 2.2.1, namely *beam-hardening* and *scatter*. Beam hardening manifests itself in cupping artefacts due to stronger attenuation in the middle of the patient. Cupping artefacts in CT images appear as slightly darker regions in the middle of the image because a higher percentage of low-energy photons have been absorbed relative to the thinner edges of the patient [84]. Beam hardening also appears as streaks and dark bands when strong density differences occur (e.g. between bone and soft tissue). Scatter manifests itself as voxel noise and is frequently called “mottle” [65]. Another source of streaking is *photon starvation* which is effectively beam-hardening to the limit when all photons have been absorbed resulting in shadows and streaking. Also included in the class of physics-based artefacts, but not described by Barrett and Keat, is the noise caused by X-ray scatter; this is reported specifically for CBCT by Siewerdsen and Jaffray [83]. *Partial-volume* effects can occur when an object is geometrically not present in all ‘views’ of the region being imaged. In the case of CT, this can be caused when slices are sufficiently wide and objects move only partially into the slices as the imaging occurs. In the case of CBCT, this can occur if the image is outside the fan angle on one side of the object and inside the fan angle as the gantry rotates to the other side. In both cases, large voxel sizes can result in incorrect reconstruction values. In all cases, the result is inaccurate values for the voxels. *Under-sampling* can either refer to an insufficient density of detectors or, more commonly, to an insufficient number of samples used to reconstruct the image. The assumption that the angular distribution of projections (or samples) is even is usually implicit when this class of artefact is discussed. This assumption is particularly incorrect with respect to CBCT binning and is discussed at length beginning with Section 4.3.



### 4.2.2 Patient-based artefacts

*Metallic materials* cause severe streaking artefacts in CT reconstruction due to the substantial density difference with the surrounding tissues. *Incomplete projections* are generated when portions of the patient lie outside the source-detector geometry; these too result in shadows and streaks [18, 62]. Of principal significance to the research in this thesis is *patient motion* induced artefacts resulting in streaks and blurred regions. In Section 2.6.3, the relevant research on image guided radiotherapy which is concerned with identifying and compensating for both gross patient setup positioning errors and the internal tissue motion that occurs during CT or CBCT imaging was reviewed.

### 4.2.3 Scanner-based artefacts

Imaging systems are mechanical systems and as such require calibration. During calibration, a scanner's offsets are measured and compensated for during reconstruction [29]. When a scanner becomes out of alignment, ring artefacts can be generated. As with any sensing device, regions or elements of the scanner may have noisy or faulty sensors. Both of these artefacts are usually corrected or compensated for by the machines prior to any reconstruction.

### 4.2.4 Artefacts discussion

The list of artefact types presented in this section serves to demonstrate various artefact causes and effects in X-ray computed tomography imaging, particularly those using some form of filtered backprojection reconstruction. With respect to CBCT, Li et al. [37] note that when "CBCT is used in imaging the thorax or abdomen of a patient, respiration induced artefacts such as blurring, doubling, streaking and distortion are observed, which heavily degrade the image quality... These artefacts are much more severe than those found in conventional CT examinations." In Li's work with Xing and Munro [40], they partially explain the cause for the severity of the artefacts by mentioning that each bin ("phase group") contains fewer projections than a full CBCT scan which is true by definition. Sonke et al. [87] also explain the cause of the artefacts as being "due to the limited number of projections acquired per breathing phase." While true, this "limited number" explanation is not complete as is shown later in Section 4.3.

One of the problems also alluded to by [6] are the consequent problems with image registration as a result of these artefacts. Specifically, the intensity-based similarity measures used in deformable based registration algorithms are sensitive to artefacts [74]. The

streaks and bands introduced by many of the artefacts create false structures that registration algorithms attempt to align. These false structures are particularly noticeable with binned CBCT reconstructions (See Figure 4.3). One of the causes of these streaks is usually classed with motion-induced artefacts because they occur simultaneously with motion artefacts and are indeed tightly coupled with them. However, in Section 4.3 they are shown to exist purely as a function of the binning process even in the absence of motion. These specific artefacts are henceforth called *gap artefacts*.




Figure 4.1: Partial plot (only 4 respiratory cycles) of displacement over acquisition time or gantry angle. Projections positioned above the displacement marked by the dashed line are binned together. Only those projections identified by the shaded areas are included in this bin, the remainder form the gaps in the projection set.

Figure 4.1 illustrates how binning induces gaps in set of projections. In the diagram, the displacement of the object of interest is plotted as a function of time or alternately the gantry angle of the projections. The dashed line demarcates the bin containing, for instance, the projections captured at the maximum inhalation state. The shaded area connects this region with the projections and where there is an absence of shading, there is a gap in the projection set used to reconstruct that specific bin. Figure 4.2 also shows the gaps in a slightly different way. The wheel represents a complete 360-degree acquisition path or gantry rotation. During that acquisition, projections for a given bin will be captured and then a gap will occur while projections for different bins are captured. The captured projections for some such bin are shown as the white wedges whereas the remaining grey area shows the resulting gaps.

It should be pointed out that most of the manifestations of artefacts as streaks and bands have a secondary cause in the choice of reconstruction. For filtered backprojection reconstruction, it is the backprojection stage which generates the actual artefact. If an




Figure 4.2: Example gaps as a function of gantry angle. In this diagram, white wedges represent the angles of a given displacement bin's projections. The grey remaining regions show the gaps in the projection set visualized by gantry angle.

algebraic technique is used, the streaks and bands are less pronounced and blurring becomes the more dominant expression of the artefact cause [69]. The penalty when using an algebraic technique is time. These methods are much more computationally expensive which explains why filtered backprojection methods are the method of choice in clinical X-ray computed tomography products. In the search methods introduced in Chapters 6 and 7, principal component analysis (PCA) will be used to introduce a requisite significant performance improvement in the actual reconstruction process. As is explained in Section 6.2.2.3, the successful application of PCA in this way necessitates the use of the filtered backprojection method.

In summary: artefacts in computed tomography reconstruction have a wide variety of well documented causes. These artefacts create difficulties for diagnosticians as well as registration algorithms. A particular type of artefact, the gap artefacts, has not been sufficiently explored to date yet it contributes significantly to the difficulties of binned CBCT filtered backprojection reconstructions. This is shown in the following section.

### 4.3 Gap artefacts

In this section, gap artefacts are examined as an independent source of error in reconstructions. To do this, two sets of experiments are performed to illustrate the impact of binning as a source of artefacts. In the first experiment, a completely static phantom is used to generate projections and these are grouped firstly as if they were evenly binned and then in an equally spaced fashion. An even binning makes a simplifying assumption that the motion is uniform and therefore each volume for each bin will be reconstructed from the same number of projections. These projections will be grouped together in each period as opposed to the equally spaced reconstructions where the individual projections will, as the name implies, be spaced equally apart. The results show the impact of the gap-inducing projection clustering of the binning method without introducing artefacts from either motion or scatter noise.

In the second experiment, two of the digital phantom configurations described in Chapter 3 are used. As in the first experiment, scatter noise is not introduced into the experiment. These phantom configurations are designed to simulate the kind of motion typically addressed in 4D CBCT research: the respiratory motion called continuously periodic in Section 2.6. To eliminate the impact of this motion on the experiment in order to demonstrate the impact of gap artefacts as opposed to the more general *projection undersampling artefacts*, a method (described in the subsection below) is introduced to generate comparative static phantom configurations.

### 4.3.1 Experiments

#### 4.3.1.1 Experiment 1: Evenly binned comparison

In the first experiment, a new simple virtual phantom configuration is created for this experiment alone. It is similar to the simple virtual phantoms in Section 3.3 except that a moving spheroid is replaced with static cylinders. This phantom is called the Static phantom for reference. CBCT projections are synthesized as per Section 3.3.2. From the projections synthesized by this configuration, three different sets of reconstructions are generated to illustrate the impact of projection set choice on reconstruction.

The first set contains randomly chosen projections. These projections are chosen in increments of 20 beginning with 40 projections and finishing with 340 projections. This range is sufficient to cover the number of projections needed to evenly construct 2 bins




Figure 4.3: Transaxial slices from reconstructions of Static phantom. Reconstructions from evenly distributed projections and from evenly binning of projections are shown. Columns show reconstructions either from 2, 4, 6, 8, and 10 bins or from using the same number of projections evenly distributed. Top row: evenly distributed projections. Second row: evenly binned reconstructions. Third row: highlighted region from evenly distributed reconstructions. Fourth row: highlighted region from binned reconstructions.

up through 10 bins. Because of the random nature, each sub-test is repeated ten times. Each reconstructed volume is compared to a volume constructed from the complete set of projections. Because there is no modality change here, a standard sum of squared difference (SSD) metric is used to measure the error. Use of SSD is further supported by the fact that the mechanism generating the projections and reconstructing the volumes is identical for the compared volumes and noise is not added so changes in voxel values are driven by the only changing key variables: the introduction of gaps in the set of projections or the addition of motion. Furthermore, the number of voxels compared in any given experiment is constant for that particular phantom/experiment even though region of interest sizes may differ. For different phantoms, different voxel value ranges may occur relative to other phantom classes. This leads to the observation that values shown in the tables reporting SSD results are not as important as the relative changes shown by the processes and/or bins being evaluated. For this reason, all tables are adjusted so that the maximum value in the table is represented as a normalized significant, all other values are adjusted to have the same exponent, and the exponent is then dropped. This removes the distraction of units or absolute values and emphasizes the relationship between results.

The second set of reconstructions simulate the binning that would occur in a perfectly linear fashion if a typical breathing period (five seconds) is applied. In other words, the simulation treats each bin as having the same number of projections. Given a 120-second acquisition time used throughout this thesis for single-scan acquisition protocols, this yields 24 periods per acquisition. The number of sequential projections, and therefore the number of sequential missing projections in a gap, for a single bin for any given period is, on average:

$$\text{Average Sequential Projections} = \frac{\text{Projection Count}}{(\text{Periods})(\text{Bins})} \quad (4.3.1)$$

$$\text{Average Gap Size} = \frac{\text{Projection Count}}{\text{Periods}} \left( 1 - \frac{1}{\text{Bins}} \right) \quad (4.3.2)$$

Using a minimum of two bins and a maximum of ten bins common in the literature [1,49], a projection count of 660 (which is close to the “normal” figure of 670 but which has the greatest number of common divisors for the bin counts used), the values shown in Table 4.1 are generated from Equations 4.3.1 and 4.3.2.

These nine reconstructed volumes are, like the random-projection volumes, compared to the complete reconstruction volume using SSD.

The third set of reconstructions simulate the results of an evenly distributed set of projections given the same total number as exists in the corresponding bin set. For example, in the case of a two-bin test, 330 projections are evenly distributed resulting in alternate projections being combined to reconstruct the resulting volume. As with the prior two tests, the resulting volumes are compared to the complete reconstruction using SSD.

#### 4.3.1.2 Experiment 2: Realistically binned comparison

The Simple virtual phantom configuration 1 (Chapter 3, Section 3.3) and NCAT configuration 3 (Section 3.5) phantom configurations are used for the experiment in this section. These will be referred to subsequently as the Simple and NCAT phantoms in this chapter. To eliminate an unnecessary variable, the generation of noise is eliminated from the synthesis process. As part of the initial generative process, the centroids of an object of interest, the tumour, are recorded for each projection. The centroids of these objects are subsequently clustered using the k-means algorithm to generate a binning based on amplitude, not phase [1]. The cluster assignments form the “oracle” for deciding which projections belong to which bin. For the purpose of this experiment, eight bins are used which is consistent with several other authors’ choice [34, 87, 95]. Figure 4.6 shows such a clustering for the NCAT phantom.

Bins	2	3	4	5	6	7	8	9	10
Projections	13.75	9.17	6.88	5.50	4.58	3.93	3.44	3.06	2.75
Gap Size	13.75	18.33	20.63	22.00	22.92	23.57	24.06	24.44	24.75

Table 4.1: Comparison of average number of projections per period and consequent number of missing projections creating a gap by number of bins. These values presume 660 projections with 24 evenly divided periods.

Bin	[1]	2	3	4	[5]	6	7	[8]	
Count	225	76	59	53	46	51	60	100	total: 670

Table 4.2: Total number of projections used to reconstruct each bin volume for the Simple phantom. These values are from 670 projections with 4 second simulated breathing resulting in 30 periods captured over 120 seconds.

Bin	[1]	2	3	4	[5]	6	7	[8]	
Count	115	69	92	69	69	92	69	95	total: 670

Table 4.3: Total number of projections used to reconstruct each bin volume for the NCAT phantom. These values are from 670 projections with 5.2 second simulated breathing cycles resulting in approximately 23 periods captured over 120 seconds.




Figure 4.4: Sagittal slice from reconstructions of the Simple virtual moving phantom that has been “frozen” at three different bin positions. (a) is the complete reconstruction showing the motion, (b) is the extreme position representing maximum exhalation, (c) is a mid-point position (d) is the extreme position representing maximum inhalation.




Figure 4.5: Motion magnitude of Simple phantom “tumour” by time. Projections corresponding to bins 1, 5, and 8 (the two ends and a middle) are identified with red dots, green X’s, and blue dots respectively.

Because this experiment is attempting to show that gap artefacts are significant and exist independently of motion, the motion normally captured in the simulated CBCT projections from these phantom configurations must be eliminated. To achieve this, the secondary output of the k-means clustering process, the set of centroids of the clusters, is used to generate phantom configurations “frozen” at the point where the “tumour” is closest to the k-means cluster centroid. Figure 4.6 shows the cluster centroid positions as ‘Xs’ amidst the clustered data positions for the NCAT phantom. These “frozen” phantoms serve both to provide ground truth for evaluating the results and to simulate projections without motion. Given that realistic motion binning results in an *unevenly distributed*

number of projections for each bin (See Figure 4.12), three of the eight bins are chosen to capture the two endpoints of motion and a midpoint. Figure 4.5 identifies the three bins for the Simple phantom as marks on a plot of the motion magnitude for each projection. These three reconstructions also serve to reveal the results of having only a few projections and of having relatively many projections in a bin. Figure 4.4 shows sagittal slices from the normal Simple phantom reconstruction and from the “three frozen” states. Figure 4.7 shows sagittal slices from the normal NCAT phantom reconstruction and from the “three frozen” states.

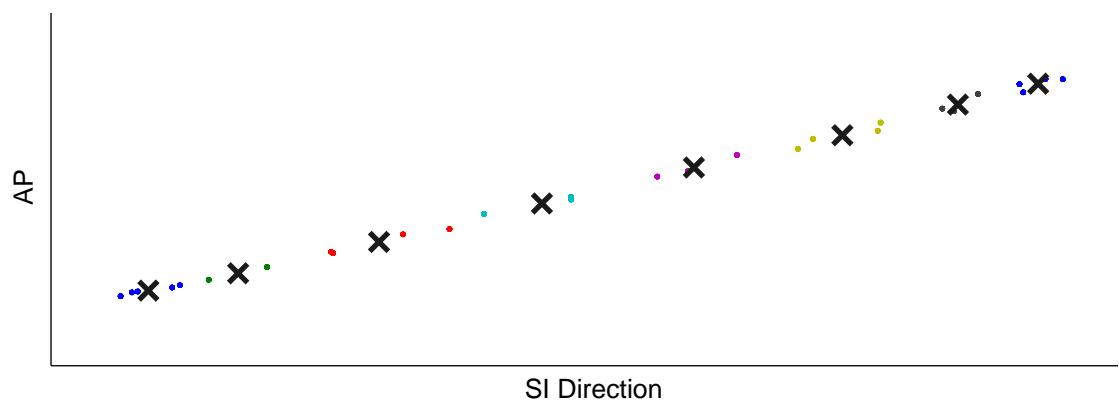


Figure 4.6: K-means clustering of NCAT phantom tumour motion in the sagittal plane. Centroids are shown as X's, cluster association is shown by colour.




Figure 4.7: Sagittal slice from reconstructions of the NCAT virtual moving phantom that has been “frozen” at three different bin positions. (a) is the complete reconstruction showing the motion, (b) is the extreme position representing maximum exhalation, (c) is a mid-point position, and (d) is the extreme position representing maximum inhalation.

The motion is principally used here to create realistic bin assignments. Once this is accomplished and the bin centroids are calculated, the three frozen phantom configurations are defined. From these configurations, CBCT projections are synthesized and reference volumes are reconstructed. Using the previous oracle-generated bin assignments from



the moving phantom and the projections from the frozen phantoms, bin volume reconstructions can be generated with resulting significant artefacts arising from the gap effect and not from motion. Likewise, the projections from the frozen phantoms can be used to generate evenly distributed reconstructions without fear of using a projection from the “wrong” bin. These two volumes, created from the binned distribution and from the even distribution, can subsequently be compared to the frozen reference volumes reconstructed from the complete set of projections.

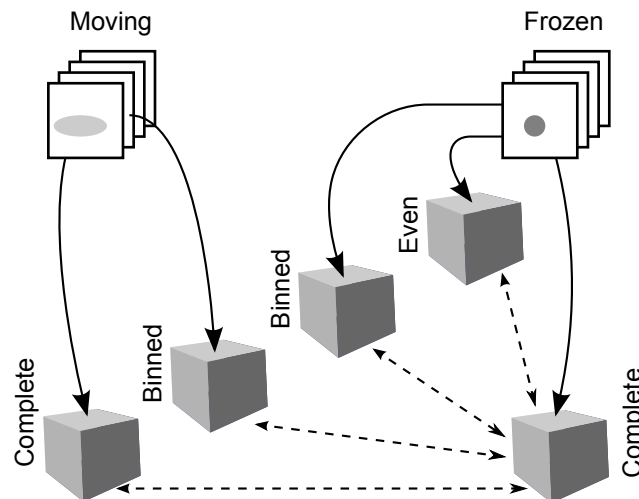


Figure 4.8: Illustration of the volume comparisons used in Section 4.3. The top left “Motion” projections are generated normally. The top right “Frozen” projections are generated as described in the section. From these projection sets, volumes are constructed from the oracle-defined bin assignments. From the Frozen projection set, a volume is reconstructed from evenly distributed projections and the reference volume is generated from the complete set. The two binned volumes and the even volume are then compared to the reference volume.

So far, the experiment only demonstrates the presence of gap artefacts independent from motion. In reality, gap artefacts never exist isolated from motion-generated artefacts. The purpose of binning is to isolate the motion; without motion, there is no need for binning. The relationship between artefacts caused by motion and gap artefacts must be understood. To test this, volumes are constructed using the original projections and not projections from the frozen versions of the phantoms. However, to fully understand the relationship a new constraint based on clinical practices is also introduced into the experiment. In Section 2.5.2, the practice of adding margins to clinical target volume (CTV) to generate the planning target volume (PTV) was documented. The PTV is meant to compensate for motion as well as setup error. Worded slightly differently, the PTV *completely contains* the moving tumour. By simulating a PTV, the measurement of error

can be constrained to just the part of the volume where motion of interest is occurring. This is important because the number of voxels in the volume whose intensity values are changed purely because of motion is much smaller than the number of voxels changed because of gap artefacts. Thus the error induced by gap artefacts across the entire volume is greater than the error induced by the motion. What is more relevant is the comparison of motion in the region of interest (ROI) with the gap artefacts in that same ROI. The PTV is the ideal ROI selector. Figure 4.9 shows sagittal slices from both the Simple and NCAT phantoms where the simulated PTV is shown in red. Note that the voxels (pixels in the slices) are not the same size physically but in both cases the PTV has been calculated to add a 10 mm margin which is a commonly used value. For completeness, four volumes are constructed and compared to the ground-truth volume. These volumes are then clipped to the ROI and re-compared. The four volumes, illustrated in Figure 4.8, are the volume reconstructed from the complete set of moving projections, the volume reconstructed from the binned set of moving projections, the volume reconstructed from the evenly distributed set of frozen projections, and the volume reconstructed from the binned set of frozen projections. These are all compared, using SSD as the metric, with the ground-truth volume reconstructed from the complete set of frozen projections.




Figure 4.9: ROI region of Simple and NCAT phantoms highlighted in red on a representative sagittal slice. ROI regions represent the standard 10 mm expansion of a CTV; pixels shown are not the same dimension.

### 4.3.2 Results and discussion




Figure 4.10: Transaxial slices from reconstructions of a simple virtual moving phantom that has been frozen at three different bin positions. Reconstructions from evenly distributed projections and from even binning of projections are shown. Columns show reconstructions from the two extreme positions (a) and (c), and a middle position (b). Evenly distributed reconstructions use the same number of projections as the corresponding binned reconstructions. Top row: evenly distributed projections. Second row: evenly binned reconstructions. Third row: highlighted region from evenly distributed reconstructions. Fourth row: highlighted region from binned reconstructions.




Figure 4.11: Transaxial slices from reconstructions of the NCAT phantom that has been frozen at three different bin positions. Reconstructions from evenly distributed projections and from even binning of projections are shown. Columns show reconstructions from the two extreme positions (a) and (c), and a middle position (b). Evenly distributed reconstructions use the same number of projections as the corresponding binned reconstructions. Top row: evenly distributed projections. Second row: evenly binned reconstructions. Third row: highlighted region from evenly distributed reconstructions. Fourth row: highlighted region from binned reconstructions.

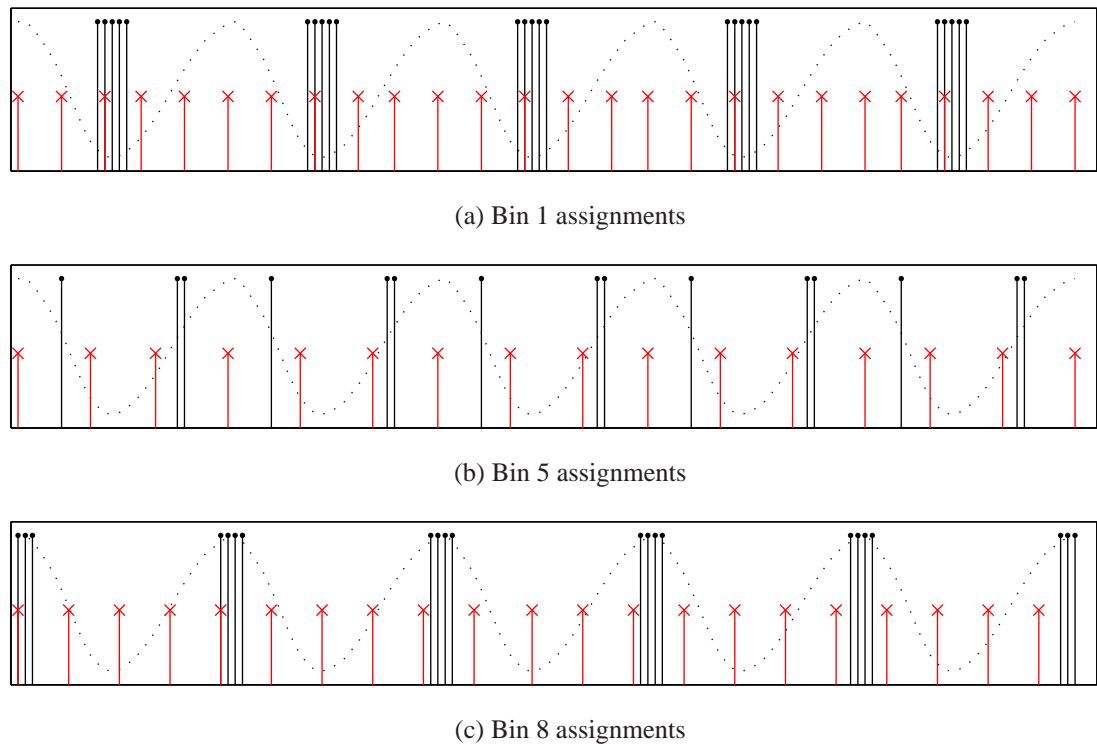


Figure 4.12: Illustration of assignments caused by evenly distributing projections vs the distribution created by binning. Five periods from an 8-binning process of the NCAT phantom, bins 1, 5, and 8, shown. Approximately 30 periods over 120 seconds with 670 acquired projections are the parameters used to construct the assignments. Spikes imply use of that projection in a reconstruction. The signal is shown as a dashed black line, the binned projection assignments are the tall black spikes, and the evenly distributed projection assignments are the shorter red spikes.

Figure 4.3 shows qualitatively what is also demonstrated quantitatively with the graph in Figure 4.13; for the Static phantom, the artefacts generated by reconstructing from binned projections appear even in the 2-bins case and very gradually get worse as the number of bins increases. The artefacts generated by reconstructing from evenly distributed projections are almost non-existent until around the 8-bins case and even by the 10-bins case are not as great as the binning-induced artefacts at the best-case 2-bins configuration. Looking at Figure 4.13, one also sees that the binning-induced gap artefacts are also worse than randomly assigned projections until the 4-bins case where some random assignments have a worse error or the 5-bins case where the average random error is worse. This makes sense given the claim that evenly distributed projections generate lower errors. As the number of bins increases, the number of projections in each bin decreases until, at the extreme case, the number of projections used equals the number of periods captured. When this occurs, each period (in this uniform scenario) will have

one projection and these will be equally spaced, assuming equal-sized periods. This is partly illustrated with the 5-bin assignment illustrated in Figure 4.12b. In that case, 69 projections are distributed over approximately 30 periods (See Table 4.3). Note in the figure that the projections are additionally split between the inhale and exhale phases of the period for mid-phase bins. This is a consequence of amplitude binning vs phase binning and further serves to evenly distribute the projections as the number of projections per bin (and thus per period) decreases. Nevertheless, the key message from Figure 4.13 is that the errors generated by gap artefacts are worse than the errors generated by general “undersampling” represented by the even distribution errors and that the difference in errors only increases as the number of bins used increases.

Having demonstrated that gap artefacts are more severe than generic undersampling, especially when large numbers of bins are used, using a Static phantom, the question must be asked: does this result hold true when more realistic data is used? The answer comes from experiment 2 using the Simple and NCAT phantoms. In these examples, the number of projections per bin varies. Rather than evaluate the error as a function of the number of bins, as was done with the Static phantom, a relatively conservative choice is made for the number of bins (eight) and the error is evaluated for representative bins as explained in Section 4.3.1 above. Figures 4.14 and 4.15 give the quantitative results. Consistent with the Static phantom results, the evenly distributed errors are lower than the gap artefact errors. Also consistent with the Static phantom results, the average error for random assignments is lower when there is a large number of projections and higher when there are fewer projections. In Figure 4.14, the actual projection counts per bin are shown in Table 4.2; for Figure 4.15, the projection counts per bin are shown in Table 4.3.

Figure 4.10 shows the qualitative results of 8-bins binning for bins 1, 5, and 8 of the Simple phantom. It is again evident that the artefacts induced by binning gaps are worse than evenly distributed projections and slightly less obvious that gap artefacts caused by small number of projections are worse than those caused by having larger numbers of projections. Figure 4.11 shows similar results for the NCAT phantom.

Comparison	Bin 1	Bin 5	Bin 8
Evenly distributed	0.004	0.333	0.066
Binned (frozen)	0.732	0.903	1.384
Binned (moving)	0.732	0.903	1.382
Complete (moving)	0.004	0.005	0.010

Table 4.4: Sum of squared differences, complete volume comparison, Simple phantom for the listed distributions by bin as per the experiment description.




Figure 4.13: Sum of squared difference error between a “best” reconstruction from a complete set of projections and the test reconstructions using the Static phantom. Blue dots show the randomly generated volume experiments by projection count. The magenta circles identify the average random error by projection count. The black diamonds show the error created by binning. Red crosses show the error created by even distribution of the same number of projections used in binning.

Tables 4.4, 4.5, 4.6 and 4.7 provide the basis for discussing the difference between gap artefacts errors and motion errors. In Tables 4.4 and 4.5, the error when comparing

Comparison	Bin 1	Bin 5	Bin 8
Evenly distributed	0.039	0.160	0.068
Binned (frozen)	1.349	0.573	1.551
Binned (moving)	1.347	0.590	1.551
Complete (moving)	1.197	0.429	1.015

Table 4.5: Sum of squared differences, complete volume comparison, NCAT phantom for the listed distributions by bin as per the experiment description.




Figure 4.14: SSD error of the reconstructions from the three selected bins for the Simple phantom. The green squares are from the evenly distributed projection reconstructions, the red diamonds are from the binned projection reconstructions, the blue dots are from the randomly generated assignments containing the same number of projections as the binned and evenly distributed reconstructions, and the magenta circled dots are the average errors of the random reconstructions.

entire volumes for both the Simple and NCAT phantoms is presented. Looking at the two ‘Binned’ results in the middle rows of the table, it is observed that the values are remarkably close. This is due to the fact that both volumes are generated in the same way so the gap artefact errors will be similar. The difference comes from the fact that the second volume is generated from the motion projections. However, the motion within a given bin is substantially less than the motion for a complete reconstruction (this is the fundamental motivation behind binning). Therefore, the dominant source of errors




Figure 4.15: SSD error of the reconstructions from the three selected bins for the NCAT phantom. The green squares are from the evenly distributed projection reconstructions, the red diamonds are from the binned projection reconstructions, the blue dots are from the randomly generated assignments containing the same number of projections as the binned and evenly distributed reconstructions, and the magenta circled dots are the average errors of the random reconstructions.



Comparison	Bin 1	Bin 5	Bin 8
Evenly distributed	0.001	0.021	0.003
Binned (frozen)	0.086	0.195	0.141
Binned (moving)	0.086	0.326	0.152
Complete (moving)	3.717	4.715	8.933

Table 4.6: Sum of squared differences, constrained to region of interest (ROI), for the Simple phantom. These are listed by distributions by bin as per the experiment description.

when binning – at least when looking at the entire volume – are gap artefacts. This assertion is likewise supported by comparing the errors caused by binning with the errors associated with reconstructed volumes created from evenly distributing the same number of projections as shown in the first line of the tables. Evenly distributing the projections causes the error to decrease significantly.

The columns in the table show the results on a binning basis. Referencing Tables 4.2 and 4.3, one sees the projections counts are inversely correlated with the error (correlation coefficients of  $-0.84$  and  $-0.97$  respectively though these are only for sample sizes of 3). As has been remarked earlier, this is intuitive since a complete set of projections is the best volume that can be constructed while the worst reconstruction will be some single backprojection.

Examining the last line of the tables, the error from the volume reconstructed from the complete set of moving projections for the Simple phantom shown in Table 4.4 is quite small compared to all but the Bin 1 evenly distributed reconstruction. This too is unsurprising because the entire volume is being compared and the error from the moving part of the volume is relatively small. For the Simple phantom, only the virtual tumour is moving. For the NCAT phantom, the situation is different. The complexity of the phantom causes much more variation when motion is induced on the phantom. Consequently, the complete volume comparison with the frozen reference volume will result in a substantially higher error. Nevertheless, the reconstructions from binning are still worse than the blurry moving reconstruction. In this case, the ROI comparisons are more helpful. Looking at tables 4.6 and 4.7, the first observation is that evenly distributed reconstructions again have the lowest error rate. The important point from these two tables comes from the last line. In contrast to the complete volume comparisons, the ROI-clipped volume comparisons show that the error rate from the moving volume is the highest error. This will become an important feature in Chapters 6 and 7, the chapters on search.

Comparison	Bin 1	Bin 5	Bin 8
Evenly distributed	0.018	0.107	0.023
Binned (frozen)	0.465	0.215	0.510
Binned (moving)	0.458	0.267	0.433
Complete (moving)	1.764	0.774	1.465

Table 4.7: Sum of squared differences, constrained to region of interest (ROI), for the NCAT phantom. These are listed by distributions by bin as per the experiment description.

### 4.3.3 In summary

The term “gap artefacts” has been introduced in the context of other classes of artefacts. It has been experimentally demonstrated that gaps in the projection set caused by a binning process create larger errors than evenly distributing the same number of missing projections. The evenly missing projections pattern represents the general meaning of “undersampling” and therefore careful and explicit characterization of this source of error should be included in discussions of 4D CBCT. It has also been shown that gap artefacts cause more error than voxel intensity changes caused by motion alone but that by clipping the volumes to the ROI which contains the motion, the PTV in clinical terms, the relative motion error with respect to gap artefacts can be increased.

# Chapter 5

## Filling Projection Gaps in CBCT Reconstructions

---

### 5.1 Introduction

Previously, in Chapter 4, the presence of gap artefacts as a specific subclass of artefact prevalent in binned CBCT filtered backprojection reconstructions was demonstrated. It was shown that volumes reconstructed using binned projections created these gap artefacts as a consequence of grouping projections together while leaving large empty gaps in the projection set. A corollary observation was that uniformly distributed projections always generate better reconstructions than random or binned reconstructions.

In this chapter and in the prior chapter, the assignment of bins to projections is performed with the aid of a k-means oracle. In Chapters 6 and 7, two search methods will be presented which attempt to find assignments without the oracle knowledge. These search methods generate ‘trial’ reconstructions which potentially induce gap artefacts, one from a form of binning and the other from a form of random assignment. The other principal form of error generated by incorrectly assigned projections is blur. Ideally, one would like to reduce the gap artefacts so that the search evaluation could be performed on the blur generated by motion in some kind of minimizing fashion. To that end, in this chapter two potential methods for mitigating the gap artefacts are proposed and evaluated.

Fundamentally, the main idea behind the gap filling methods presented here is that

if we have a collection of objects of some kind that can somehow be averaged together, then adding additional ‘average’ objects to the set and re-averaging should not change the resulting average. Filtered backprojection is an averaging process on individually back-projected filtered projections, as explained in Chapter 2, Section 2.3.3.2. Consequently, if one can find an appropriate “average projection” for a given missing projection, then hypothetically this can be substituted in the reconstruction process without making the reconstruction any worse than the original reconstruction generated from the complete set of projections.

A simple one-dimensional toy analogy helps make this clear. Suppose some signal is measured for the one-dimensional position of an object:

$$5 \ 6 \ 5 \ 1 \ 0 \ 1 \ 0 \ 0 \ 1 \ 6 \ 5 \ 6 \quad \text{Average} = 3.0$$

The object appears to have two roughly stationary positions at 5.5 and 0.5 while the mean position is 3. If these measurements were binned into two bins, those two bins would contain the values:

$$\begin{array}{r} 5 \ 6 \ 5 \ \_ \ \_ \ \_ \ \_ \ \_ \ 6 \ 5 \ 6 \quad \text{Average} = 5.5 \\ \_ \ \_ \ \_ \ 1 \ 0 \ 1 \ 0 \ 0 \ 1 \ \_ \ \_ \ \_ \quad \text{Average} = 0.5 \end{array}$$

From Chapter 4, a fundamental difference between this toy example and CBCT is that one cannot take the mean of a subset without introducing gap artefact errors. If the gaps in the toy analogy are “filled” with the average, then the resulting binned signals contain the values:

$$\begin{array}{r} 5 \ 6 \ 5 \ 3 \ 3 \ 3 \ 3 \ 3 \ 3 \ 6 \ 5 \ 6 \quad \text{Average} = 4.25 \\ 3 \ 3 \ 3 \ 1 \ 0 \ 1 \ 0 \ 0 \ 1 \ 3 \ 3 \ 3 \quad \text{Average} = 1.75 \end{array}$$

Two important observations must be made at this point. The first observation is that the resultant average values for the binned signals are not accurate. That accuracy is influenced by the number of average values used to replace missing measurements (projections, in the CBCT case). Obviously, the more average values are used, the closer to the global average the binned average will become. In the CBCT case, the more the “average” projections are used, the blurrier and more like the global reconstruction the resultant will be. The second observation is that the two binned averages are still apart from each other - they are still distinct. This property becomes very important in the search chapters.

What constitutes an “average” projection? A naïve approach might be to take the average value of all projections and substitute the global value in place of a missing projection

pixel. However, this would merely shift the intensity of all the values in the reconstruction by some constant value. Another variation might be to average all projections together and use this as a single global average replacement image. Unfortunately, as Figure 5.1 shows projections taken at different angles are very different from each other so averaging these together and then backprojecting them where gaps exist would just create a general haze in the reconstructed image. Yet another idea is to create an atlas of projections and construct a representative average from these for each projection angle. Unlike the brain, where this idea has merit, the diversity among patients in the abdominal region is so great that this approach is questionable at best. Something better is preferred. In this chapter, two possible average projections that are both angle-specific and patient-specific are considered. In Section 5.2, a method for filling gaps in the projection space using the existing medical protocol for CBCTs is presented. As described earlier in Chapter 2, Section 2.3.3.2, the existing clinical protocol uses a single rotation of the CBCT gantry over 360 degrees in approximately two minutes collecting approximately 670 projections. In Section 5.3, a novel clinical protocol using two scans instead of one is proposed. The purpose of this proposed protocol change is made clear in Chapter 6. In both sections, the phantom configurations and ground truth described in Chapter 4 are used to evaluate the methods.




Figure 5.1: Four example projections from the NCAT phantom acquired at the represented angles (degrees).




Figure 5.2: Sample NCAT projection before and after forward reprojection process.

## 5.2 Filling gaps using one scan

Reconstructions from projections containing moving objects/tissues are visibly blurry. This is an expected outcome of the averaging nature of reconstruction. This also presents an opportunity. If one synthesizes a projection by forward projection through a reconstructed (averaged) volume, then one can create an *angle-specific average* projection. That is the basis of this proposed method.

### 5.2.1 Method

For this method, projections from the Simple and NCAT phantoms are synthesized as has been previously described in Chapter 4, Section 4.3.1.2. These projections are then used to reconstruct a volume. The simple forward-projection method used to generate the original NCAT projections, and described in Section 3.5.2, is used to create a new set of “average” projections from this reconstructed volume at the same set of angles as the original projection set. Figure 5.2 shows an example NCAT phantom projection before and after this process. The blurring around the diaphragm, liver, and chest wall is particularly evident.

The k-means oracle binning used in Chapter 4 is also used here. Once again, the results from bins 1, 5, and 8 are used as representative bins from an 8-bin process for both the Simple and NCAT phantoms. Likewise, as per Chapter 4, both complete volume and region of interest (ROI) sub-volumes are evaluated using the sum of squared difference (SSD) metric.

The method explored here actually has two variations: a *complete fill* variant and a *balanced fill* variant. With the complete fill variation, the original projections assigned to a given bin are combined with the reprojected average projections such that a complete projection count is obtained. In other words, all original projection angles have some contributing projection. For our phantom projection count of 670 projections, bin 5 has 69 projections assigned to it (see Table 4.3) so 69 original projections will be combined with 601 averaged projections.

With the balanced fill variation, one of the principal observations from Chapter 4 is used: that uniformly distributed projections generate fewer errors than clustered projections. This idea is used to mitigate the main disadvantage of the complete fill technique which is that the original ‘signal’ of the binned original projections gets lost or swamped by the average projections. In the bin 5 case just mentioned, only 10% of the reconstructed signal comes from the phantom at the state identified by the binning process. In our toy analogy used earlier, this would result in estimated positions of 3.25 and 2.75 instead of 5.5 and 0.5. Therefore, the balanced fill method takes the same number of average projections as is assigned to a bin and uniformly distributes those projections throughout the gaps. In the 5 bin case, 69 average projections, evenly selected from the 601 missing projections, are combined with the 69 binned projections to reconstruct the volume.

## 5.2.2 Results of one-scan gap filling

When looking at the results of filling the projection gaps in the one-scan protocol, regardless of the phantom used one observes that gap filling indeed mitigates the error caused by gap artefacts when the entire volume is considered. Equally interesting is the result that balanced fill gap filling produces better results than complete fill gap filling.

Comparison	Bin 1	Bin 5	Bin 8
No filling	0.732	0.903	1.383
Balanced fill	0.230	0.394	0.365
Complete fill	0.369	0.710	0.599
No binning	0.818	0.818	0.823

Table 5.1: Volume comparisons of results of 1-scan protocol gap filling for Simple phantom.

In Table 5.1, the top row shows the SSD errors when comparing oracle-binned projections against the frozen reference volumes for each bin. The second row shows the errors from the balanced fill method which generates the lowest errors in the table. The third row shows the results from the complete fill method which generates lower errors

than just binning, but which are still higher than the balanced fill method. The fourth row shows the error from a reconstruction using all (reprojected) projections. Interestingly, it is higher than the error created by binning without gap filling for bin 1, and lower for bins 5 and 8. This is of interest because it reveals the complex interaction between isolating the motion (reducing blur) and reducing gap artefacts. Referring back to Table 4.2, bin 1 has substantially more projections (225) than bins 5 (46) and 8 (100). Figure 4.5 shows that bin 1, as shown by the red dots at the bottom, captures less motion than bins 5 and 8. These two reasons combine when reconstructing from the binned projections for bin 1 to create a volume with fewer gap artefacts and minimal blur. Figure 5.3 shows this visibly. The binned reconstruction in the left column shows some gap artefacts errors but also a clearly isolated tumour object. The columns showing the fill reconstructions show minimal gap artefacts errors but visible tumour blur.




Figure 5.3: One-scan fill results, Simple phantom, bin 1. Top row shows sagittal slices, bottom row shows transaxial slices. Columns: (a) standard binning, (b) fill using balanced approach, (c) fill using complete set.

Figure 5.4 shows a different outcome. Using fewer projections creates very visible gap artefacts and despite the lack of tumour blur that is still evident in the fill columns images, the binned volume error for bin 5 now exceeds the error from the complete reprojected reconstruction. The difference between the balanced fill and complete fill methods are also apparent in this figure. The tumour object is blurry in the complete fill case (right) but is relatively distinct in the balanced fill case (middle) without introducing the larger



gap artefacts apparent in the binned volume on the left.




Figure 5.4: One-scan fill results, Simple phantom, bin 5. Top row shows sagittal slices, bottom row shows transaxial slices. Columns: (a) standard binning, (b) fill using balanced approach, (c) fill using complete set.

Figure 5.5 shows a similar outcome to Figure 5.4. Surprisingly, though bin 8 has almost double the number of projections of bin 5, the gap artefacts are more severe as is also shown by the third column in Table 5.1. This is possibly explained by the simplistic nature of the phantom which contains just the single moving sphere object. In the position captured by bin 8, the object has moved to its furthest point from the axis of rotation and the gap artefacts are slightly less symmetric.

Comparison	Bin 1	Bin 5	Bin 8
No filling	1.347	0.590	1.551
Balanced fill	0.668	0.297	0.596
Complete fill	1.082	0.480	0.830
No binning	1.198	0.430	1.015

Table 5.2: Volume comparisons of results of 1-scan protocol gap filling for NCAT phantom.

The one-scan fill results for the NCAT phantom are shown in Table 5.2 with very similar results to the Simple phantom. The top row again shows the SSD errors when comparing oracle-binned projections against the frozen reference volumes for each bin.




Figure 5.5: One-scan fill results, Simple phantom, bin 8. Top row shows sagittal slices, bottom row shows transaxial slices. Columns: (a) standard binning, (b) fill using balanced approach, (c) fill using complete set.

The second row shows the errors from the balanced fill method which, as for the Simple phantom, generates the lowest errors in the table. The third row shows the results from the complete fill method which also generates lower errors than just binning, but are still higher than the balanced fill method.

Looking at the NCAT results shown in Figures 5.6, 5.7, and 5.8, the blur of the tumour object in the column (c) showing the complete fill method reconstructions is obvious when compared to columns (a) binned and (b) balanced fill. Also evident in the middle balanced fill columns are the balance between a sharp tumour and gap artefacts. While some gap artefacts are present, they are not as prevalent as the binned reconstruction on the left. Similarly, though there is some blur in the reconstruction, it is not as severe as the blur in the complete fill reconstruction on the right.

A further observation that can be made from the image slices of the reconstructed volumes is that gap artefacts are more substantial at the outer regions of the volume than in the central regions. This is a convenient feature since the object of interest, the tumour, is usually placed at the isocentre (also the volume center) during CBCT scans. Therefore, our region of interest (ROI) which has been previously defined to be the planning target volume (PTV), will likely incur lower gap artefacts induced errors. Likewise, by restricting the number of voxels compared to the ROI, the ratio of moving or blurred voxels to




Figure 5.6: One-scan fill results, NCAT phantom, bin 1. Top row shows sagittal slices, bottom row shows transaxial slices. Columns: (a) standard binning, (b) fill using balanced approach, (c) fill using complete set.

Comparison	Bin 1	Bin 5	Bin 8
No filling	0.09	0.33	0.15
Balanced fill	1.35	1.36	2.55
Complete fill	2.39	4.42	6.93
No binning	4.77	5.03	9.21

Table 5.3: ROI-clipped comparisons of results of 1-scan protocol gap filling for Simple phantom.

the total comparison region is significantly increased from 0.004% to 11.7%. For these reasons, it is not surprising that the SSD errors are substantially lower overall and that for the extremely simple case (Table 5.3), the binned reconstruction is substantially better because of the increased weighting on moving voxels. It must be noted, though, that the two filled methods still record lower errors than the complete projection volume errors in the fourth line of the table.

In the more complex and more realistic NCAT phantom results, it is observed that the binned reconstruction errors are now on par with the balanced fill technique. Both fill techniques continue to outperform the global reconstruction approach. These results support both the ROI clipping methodology and the balanced fill approach to mitigating gap artefacts.




Figure 5.7: One-scan fill results, NCAT phantom, bin 5. Top row shows sagittal slices, bottom row shows transaxial slices. Columns: (a) standard binning, (b) fill using balanced approach, (c) fill using complete set.

### 5.3 Filling gaps using two scans

If, instead of capturing projections during the course of one gantry rotation, one captures them with two rotations, then an interesting property of the projections can be exploited. This property is that two projections, taken at nearly the same angle, but from two different scans, will produce essentially the same image except for the motion that may have occurred (See Figure 5.9). Since the purpose of the averaged projections is effectively to contain the averaged motion in each projection, this can be accomplished by taking the projection closest to some established angle from each of the two scans and averaging them together. That is the basis behind this proposed method.

Comparison	Bin 1	Bin 5	Bin 8
No filling	0.458	0.267	0.433
Balanced fill	0.699	0.259	0.446
Complete fill	1.519	0.531	1.125
No binning	1.764	0.774	1.465

Table 5.4: ROI-clipped comparisons of results of 1-scan protocol gap filling for NCAT phantom.




Figure 5.8: One-scan fill results, NCAT phantom, bin 8. Top row shows sagittal slices, bottom row shows transaxial slices. Columns: (a) standard binning, (b) fill using balanced approach, (c) fill using complete set.

An additional potential benefit of this method is the likely increase in the uniformity of the projection distribution. For continuously periodic scenarios, the correct bin assignment on any given rotation at some angle  $\theta$  will be random with a probability distribution derived from the breathing motion. More time is spent in minimum inhalation than in maximum inhalation and, on some breath cycles, the maximum inhalation state is not even achieved. Therefore the probability of one state is not the same as the probability of another state. Nevertheless, there is a likelihood that, given eight to ten bins, the bin assignments for two projections at the same location acquired on independent scans will be different. Thus when both sets of projections from the two independent scans are binned, the collective set of projections for any given bin will have a more uniform distribution than that observed from a single scan. From the previous chapter, it was shown that increasing the uniformity of the projection set decreases the gap artefacts. This is effectively what Li et al. [40] attempt using both a “slow gantry rotation” protocol and a “multiple gantry rotation” protocol. To compensate for the increased number of projections these methods require, they lower the current which degrades the projection quality but, because it mitigates the gap artefacts, improves the reconstruction. They refer to gap artefacts as “view-aliasing artefacts” and describe the cause as “insufficient angular sampling.” However they acknowledge the clinical constraint of requiring too much

time which limits the viability of the protocols and later propose a B-spline deformable registration solution using a planning CT as the baseline target for registering the binned volumes [38].

Chang et al. [11] also propose both a slow (five minute) continuous acquisition protocol and a slow gated protocol for MV CBCT likewise acknowledging “streak artefacts.” Both their work and the earlier work of Li’s team essentially solve the problem of “under-sampling” by increasing the number of samples.

The two scan protocol is similar in nature to the Opposite Ray Algorithm (ORA) proposed by Linney [42]. While that method only works with fan beam CT projections, both that idea and this proposed protocol are conceptually similar. Both seek to acquire the same view at two different points in time. The ORA does this by recognizing the fact that as the fan beam projections can be rebinned as parallel projections and that, when viewing a stationary object, the projections of a parallel projection offset by 180 degrees will be identical. They will, however, be at different temporal points and therefore can capture motion. Because of the off-plane geometry of cone beam CT projections, the same method cannot be used. Instead, as has been described, the same purpose is achieved by acquiring two scans at nearly the same spatial position but different temporal positions.

### 5.3.1 Method

For this method, projections from the Simple and NCAT phantoms are synthesized with slight changes to the normal protocol. Rather than generate one set of 670 projections over 120 seconds, two sets of projections are constructed. Each set is constructed simulating a hypothetical 60 second scan which is the maximum scan rate currently allowed for CBCT machines by the IEC [39]. To keep dosage the same, only half the projections (335) are acquired on each simulated scan so the total number of projections remains the same as a conventional CT at the same kV and mA levels thus keeping the hypothetical total dosage the same. Likewise, the incremental time required by the system to stop the first scan and then start the second scan is small enough to effectively not change the duration of the total scan from the patient’s or hospital’s perspective.

The k-means oracle binning used in Chapter 4 is also used here. Once again, the results from bins 1, 5, and 8 are used as representative bins from an 8-bin process for both the Simple and NCAT phantoms. Likewise, as per Chapter 4, both complete volume and region of interest (ROI) sub-volumes are evaluated using the sum of squared difference (SSD) metric. An important difference with this proposed protocol is that each scan will have its own binning assignment.

The earlier assertion that the differences between similarly located (in terms of angle) projection pairs from the two scans are principally caused by motion is based on the idea that projection differences caused geometric changes, i.e. changes in angle (approximately half a degree), are small while projection differences due to motion can be large. If a projection in one scan at some angle  $\theta$  is captured at the maximum inhalation state while the projection in the second scan at approximately  $\theta + \Delta\theta$  is captured at the minimum inhalation state, then the intuition is that the motion-induced differences will be greater than the geometry-induced differences. However, this is a probabilistic assumption. The two projections may, in the case of periodic motion, have come from the same phase at a given angle. They may both be from, say, the maximum inhalation state. Periodic motion is principally respiratory-based and for real-world patients this means some level of irregularity will always exist. The hypothetical worst case is that patients somehow mimic exactly the same breathing pattern for each scan resulting in effectively duplicate projections for each angle, but this is extremely unlikely. Nevertheless, even this worst case scenario would be equivalent to a single-scan protocol collecting only half the usual projections. For the purpose of these experiments, a best case scenario (two scans 180-degrees out of phase) is used. More realistic scenarios are explored in the next chapter. Nevertheless, the assertion that motion differences are greater than geometric differences needs to be validated. To do so, the frozen phantom configurations are used again to isolate the motion changes from the geometry changes. Three such phantoms have been generated for bins 1, 5, and 8 so the validation method consists of finding pairs of projections from the two scans (where motion has occurred), taking their absolute difference, and comparing this to the absolute difference obtained from two projections from the frozen phantom corresponding to the same angles as the two scan projections. In pseudocode, this can be described as follows:

The fill method is different from the one-scan fill methods because there are more choices. In the one-scan case, the projection for a given angle can either come from the bin-assigned set or from the average set (or from neither in the balanced fill case). In the two-scan case, the projection for a given angle can come from the bin-assigned set of the first scan, the bin-assigned set of the second scan, or the averaged set. However, if the projections from the two scans for a given angle are both assigned to the bin, then they should both be used and the best way to do this is to average them together. But this is exactly what is done to create the average set. Therefore, the filling method consists of taking the projections *exclusively* assigned to the bin from the first scan, the projections *exclusively* assigned to the bin from the second scan, and then filling the remaining slots with projections from the average set.

```

Input: projections from two scans, projections from frozen phantom
Output: basic statistics of differences from motion vs geometry
foreach Bin b in SampleBins do
  foreach Projection p1 in scan1Projections do
    if BinAssignment(p1) == b then
      /* find projection from 2nd scan closest angle-wise to p1 */
      p2 = Nearest(scan2Projections,p1.angle)
      frozen1 = Nearest(frozenProjections,p1.angle)
      frozen2 = Nearest(frozenProjections,p2.angle)
      /* calculate SADs between moving and frozen projection pairs */
      motionDiff = SumOfAbsoluteDifference(p1,p2)
      geometryDiff = SumOfAbsoluteDifference(frozen1,frozen2)
      motionDiffs.bin(b).append(motionDiff)
      geometryDiffs.bin(b).append(geometryDiff)
    end
  end
  /* calculate statistics from pair differences */
  averageGeometryDiffs.bin(b) = Average(geometryDiffs.bin(b))
  sdGeometryDiffs.bin(b) = StdDev(geometryDiffs.bin(b))
  avgMotionDiffs.bin(b) = Average(motionDiffs.bin(b))
  sdMotionDiffs.bin(b) = StdDev(motionDiffs.bin(b))
end

```

**Algorithm 1:** Measuring differences caused by gantry rotation vs physiological motion.

### 5.3.2 Results of two-scan gap filling

The first set of experiments to be examined are those pertaining to the question of physiological motion-induced projection differences versus those induced by changes in the gantry position. Figure 5.9 visibly demonstrates the effects of geometry changes for the small angles involved and motion changes caused by the temporal shift. On the left are two representative difference images generated by changing the angle of acquisition on the frozen phantom by .54 degrees (360 degrees / 670 projections). On the right are two difference images at roughly the same angular separation but now from the two different scans thus representing both the geometry and motion changes. The effect of the motion on the projections is most pronounced along the diaphragm wall and, in the top right image, the anterior chest wall.

The quantitative results of the motion differences versus geometry differences validation are seen in Table 5.5. This table records the average projection differences and standard deviation for the three sample bins. For each of the bins, it is apparent that motion contributes significant additional differences. If one assumes, however, that the






Figure 5.9: Sample NCAT projection differences. On the left (a),(c), the absolute differences between two successive projections from the same fast scan from a 2-scan protocol is shown. On the right (b),(d), the absolute differences between two projections at nearly the same angle but separated in time, using the 2-scan protocol, are shown.

Comparison	Bin 1	Bin 5	Bin 8
Mean difference (motion)	6.23	8.96	5.41
Standard dev. (motion)	1.50	2.71	1.74
Mean difference (geometry)	3.48	3.50	3.54
Standard dev. (geometry)	0.42	0.46	0.52

Table 5.5: Differences, by sample bins, between projections. The top two rows are the mean and standard deviation induced by motion differences between projections at nearly the same angle but different times, the bottom two rows are the mean and standard deviation caused by geometry changes alone.

comparisons made between projections containing motion also contain the differences caused by geometric changes, then *on average* the changes induced by motion are a little less than the changes caused by geometry. What is less obvious, but can be inferred from the standard deviations lines, is that the motion-induced differences are much more variable than the variation in geometry-change differences. This is caused by the fact that these phantoms have approximately sinusoidal motion and that the motion between the two scans is out of phase. Therefore, as the two motions “cross” the motion difference

will be negligible. Where the two motions peak, the motion difference will be maximal. Figure 5.10 illustrates this. This characteristic will play an important role in the next chapter.




Figure 5.10: Motion amplitude signals from two scans. The motion, perfectly out of phase, is the same where the signals cross and is maximally different where one signal peaks and the other is at rest.

Comparison	Bin 1	Bin 5	Bin 8
Binned (moving)	1.52	8.71	9.44
Filled	0.06	0.05	0.08
Averaged (all)	0.05	0.06	0.10

Table 5.6: Volume SSD error comparisons with reference volumes of results of 2-scan protocol gap filling for Simple phantom.

Comparison	Bin 1	Bin 5	Bin 8
Binned (moving)	0.03	0.42	0.23
Filled	0.52	3.74	5.56
Averaged (all)	4.77	5.03	9.21

Table 5.7: ROI SSD error comparisons with reference volumes of results of 2-scan protocol gap filling for Simple phantom.

Comparison	Bin 1	Bin 5	Bin 8
Binned (moving)	1.916	1.322	3.000
Filled	1.213	0.477	1.072
Averaged (all)	1.213	0.478	1.071

Table 5.8: Volume comparisons of results of 2-scan protocol gap filling for NCAT phantom.

Comparison	Bin 1	Bin 5	Bin 8
Binned (moving)	1.216	0.595	1.339
Filled	1.583	0.881	1.688
Averaged (all)	1.764	0.774	1.465

Table 5.9: ROI-clipped comparisons of results of 2-scan protocol gap filling for NCAT phantom.

The results for the Simple phantom shown in Tables 5.6 and 5.7 are consistent with earlier results from the one scan filling approach. Volume errors are much less for gap filled reconstructions than for binned reconstructions but are greater when clipped to the ROI. Likewise, the ROI clipped error is still less than the standard complete reconstruction error. The results for the NCAT phantom (Tables 5.8 and 5.9) are also consistent with the results observed for the Simple phantom.

Visually, the difference between filling and binning for the Simple phantom for the three bins used in the experiment can be seen in Figures 5.11, 5.12, and 5.13. Note the increased gap artefacts visible in the transaxial slices and the increased blurring in the filled sagittal slices.




Figure 5.11: Two-scan fill results, Simple phantom, bin 1. Top row shows sagittal slices, bottom row shows transaxial slices. Columns: (a) standard binning, (b) complete fill.




Figure 5.12: Two-scan fill results, Simple phantom, bin 5. Top row shows sagittal slices, bottom row shows transaxial slices. Columns: (a) standard binning, (b) complete fill.




Figure 5.13: Two-scan fill results, Simple phantom, bin 8. Top row shows sagittal slices, bottom row shows transaxial slices. Columns: (a) standard binning, (b) complete fill.




Figure 5.14: Two-scan fill results, NCAT phantom, bin 1. Top row shows sagittal slices, bottom row shows transaxial slices. Columns: (a) standard binning, (b) complete fill.




Figure 5.15: Two-scan fill results, NCAT phantom, bin 5. Top row shows sagittal slices, bottom row shows transaxial slices. Columns: (a) standard binning, (b) complete fill.




Figure 5.16: Two-scan fill results, NCAT phantom, bin 8. Top row shows sagittal slices, bottom row shows transaxial slices. Columns: (a) standard binning, (b) complete fill.

As with the Simple phantom, the difference between filling and binning for the NCAT phantom can be seen for the three bins used in the experiment in Figures 5.14, 5.15, and 5.16. Note the increased gap artefacts visible in the transaxial slices and the increased blur in the filled sagittal slices.

## 5.4 Conclusions

In this chapter, two methods for reducing the errors introduced by gap artefacts were proposed and evaluated experimentally. Both methods attempt to fill the gaps introduced by the binning process with a patient-specific, angle-specific average projection. It was seen that both methods do reduce the errors when comparing entire volumes. However, it was also shown that when volumes are clipped to a region of interest, the total errors caused by the change in position of the object of interest in the ROI exceed the errors caused by gap artefacts (within the ROI). In the one-scan gap filling method, a “balanced” variation was introduced which showed error rates similar to the normal binned ROI error rates.

# Chapter 6

## Searching for CBCT Projection Bin Assignments Using a Two Scan Protocol

---

### 6.1 Introduction

The goal of the previous chapters has been to pave the way for this chapter and Chapter 7 which seek to satisfy the ultimate goal: generating reconstructed volumes that contain different states of motion. In order to get to this goal, the projections acquired in the CBCT scan need to be assigned to bins which represent these states of motion. In Chapter 5, this was accomplished using an oracle. This oracle used k-means to classify the position information that is generated with the projections. Given that the paths of the motions of interest are not complex, this method of constructing an oracle is effective and yields a set of  $k$  motion centroids and a classification of the projections into motion bins. Such an oracle, unfortunately, is not available in reality and so some method of finding assignments must be constructed.

In Chapter 2, the background chapter, various methods for linking different kinds of respiratory markers acquired simultaneously with the CBCT projections in order to bin the projections were presented. These methods all show good results and are especially strong when they can be coupled with a 4D CT scan, usually using some form of B-spline deformation model. These solutions handle what has been referred to in this thesis as continuously periodic motion. They do not solve the other kinds of motion described in

Chapter 2, namely shift and discrete motions.

In this chapter a novel approach to identifying motion, but one which requires a protocol change in how CBCTs are acquired, is presented. This is done by first defining a search space (Section 6.2) which is computationally intractable and then using a series of ideas to reduce the computational task to the point of feasibility. Results on the phantom configurations defined in Chapter 3 are presented and evaluated.

## 6.2 Defining an assignment search space

What is needed is a precise definition of a space to search and a way to search it. In optimization parlance, this is a function  $f : \mathbb{R}^n \rightarrow \mathbb{R}$  to be maximized (or minimized) over a given set  $\mathcal{D} \subset \mathbb{R}$ . The function  $f$  is the *objective function*, and the set  $\mathcal{D}$  is the *constraint set* [88, p.74]. Ultimately, the goal is a set of reconstructed volumes, one for each bin. These volumes are in turn reconstructed from the set of projections. One way to represent these volumes, then, is to define a mapping from projections to bins. Such a mapping is often called an *assignment*. The actual volumes themselves follow from this mapping by reconstructing each bin volume using those projections assigned to that bin.

An objective function for the space defined above can then proceed in two distinctly different directions. It can assess the quality of the individual projection assignments and perform some aggregating operation (mean, sum, etc.) on those values, or it can assess the quality of the reconstructions generated by the assignment. For binned reconstructions based on respiratory motion, and which allow the observation and identification of the diaphragm in the projections, the first approach is reasonable: each projection can be evaluated and assigned to a bin. In Chapter 2, it was established that this is not always possible or, in terms of X-ray exposure, desired. Likewise, it still leaves the question of non-respiratory motion identification open. If one assumes that, having no prior model of motion, there is nothing to measure a given projection against then finding an objective function in this way becomes difficult.

One approach to the difficulty of having no reference model is to treat this as a *monocular structure-from-motion* problem. The problem is made slightly easier by accurately knowing the motion of the “camera” but is made significantly more difficult because the structures are themselves moving. This makes it a *multibody structure and motion* (MSaM) problem. A further difficulty with this proposition is that the moving structures are (to X-rays) transparent. This makes any depth-queueing from occlusion impossible and compounds the tracking problem when features cross over each other. Recent work by Schindler and Suter [76] reinforces the difficulty of this kind of problem; their



work is constrained to “a small number of motions.” Likewise, the work by Avidan and Shashua [5] points out that, when attempting to identify the position in a 3D space from a monocular image, “if the point is moving generally then the task of triangulation is not feasible. . . Knowledge of the camera ego-motion does not change the feasibility of the problem.” For these reasons, both search mechanisms suggested in this chapter and in Chapter 7 will pursue the second option: evaluating reconstructed volumes.

If projections are reconstructed into volumes according to some candidate bin assignment, an interesting possibility arises which yields a new approach. If motion has occurred (implicitly to a degree one cares about), and is observable, then the reconstructed volumes will be different when the bin assignment is correct. This difference creates an initial basis for an objective function: looking for volumes that are the most different from each other. As was described in 4.3.1.1, a sum of squared differences metric is proposed and implemented. This choice of a metric is supported by several factors. Firstly, the comparison that occurs is relative and not absolute: what matters is that the two volumes are different, not that they are different by some specified threshold. Secondly, the number of compared voxels for any particular test in the same so comparisons are not biased by one volume having more voxels than another. Thirdly, the reconstruction process is identical for the compared volumes as is the source data generation process so the values in one volume will not be biased or shifted as might be the case if samples were acquired in different ways or at different times or were processed in different ways. Finally, a fundamental property of the SSD metric with respect to the methods described here is its ability to be applied in eigenspace. A problem arises with this idea, though. Assuming the simple case (and the case that will be used for the remainder of the thesis) of using only two bins, one can construct two volumes wherein the “brightest” projections are assigned to one bin and the “darkest” projections are assigned to a second bin. In this case the two reconstructed volumes will be significantly different from each other without the cause being from motion.

A further evaluation metric that performs a form of regularization on the difference metric is to assess whether the two volumes maintain the same essential statistical property of the original complete reconstruction. A basic *mean* statistic is proposed whereby the mean grey-levels of the two reconstructed candidate volumes are compared with the mean grey-level of the original reconstructed volume. If for instance, for a given region of interest (ROI), the tumour contained by the ROI moves (but stays within the ROI), the overall mean of the two reconstructed volumes remain similar to each other and to the original complete reconstruction while the differences between the two volumes increases. As was the case with the SSD metric, use of this metric alone poses a problem.

Hypothetically, the reconstruction from a random draw of projections as a subset of the overall population, will have a mean close to the mean of the volume reconstructed from the complete population of projections. This creates a tension: while the SSD metric wants to “pull the volumes apart” the mean metric wants to “keep them together.”

Reality unfortunately imposes certain complications which result in possible limitations to this technique. Tissue intensity differences between the tumour and the surrounding tissue must be sufficiently large to create a useful measurement. This difference must also be larger than the differences caused by noise and other artefacts. This is a possibility in the case of liver or prostate tumours when no contrast is used or when no markers (natural or implanted) are present. Also, if motion changes are too small, i.e. if the tumour moves very little, the total difference in the reconstructed volumes caused by noise and other artefacts may exceed the differences caused by the motion shift. These boundary conditions are represented in the phantom configurations used to test this method to verify both cases where the technique works reasonably well and where it under performs.

### 6.2.1 Search method

The search space is defined as the set of possible bin assignments for the projection set and, more specifically for this chapter and Chapter 7, the set of two-bin assignments. In set notation,  $\mathcal{D} = A^n \mid A = \{0, 1\}$  where  $n$  is the number of projections. Elements in the search space (of which the solution is a member) will subsequently be referred to as projection assignments. Furthermore, the choice has been made to evaluate the projection assignment using the two reconstructed bins it generates. Given the search space, an objective function  $f$  must be chosen. Three significant barriers impede progress at this point. The first is the size of the search space. Given a binary assignment and recognizing that it is irrelevant which volume is defined as “bin 1” and which is defined as “bin 2”, there exist  $2^{n-1}$  possible assignments or, for the average projection count used throughout this thesis of 670 projections,  $2^{669}$  assignments. This is computationally intractable.

A second significant barrier compounds this limitation: filtered backprojection reconstructions currently take between seconds and minutes depending on optimizations and hardware. The use of cloud computing or local many-core servers can reduce this significantly because of the ease with which the problem is parallelized. Regardless, reconstruction of a typical  $256^3$  sized volume from 670 projections each containing  $512^2$  voxels which are filtered and then backprojected (averaged across the volume) requires a significant amount of computation, memory, and number of disk accesses.

A third significant barrier which affects the search possibilities is the fact that the




Figure 6.1: Three example objective function scoring sequences. Each sequence is obtained by first randomly generating a complete assignment and then changing each of the individual projection bin assignments and rescored with the objective function. The plots show the change in score as a function of the particular projection assignment change.

search space is not smooth. Ideally, one would like some kind of monotonically changing function which allows a gradient descent/ascent search to be performed. This is not the case for binned volume comparisons. Assignment changes result in very “spiky” cost changes using the proposed objective function as can be seen in Figure 6.1. In the plots, the most fine-grained change possible, a single projection assignment switch, leads to substantial changes in the score outcome. This makes a gradient ascent approach unrealistic. Consequently, since “hill climbing” is a poor option, the idea of an exhaustive search must be considered. Yet clearly, this is *not* an option given these barriers. Something must be done to reduce the size of the search space in some reasonable fashion and reduce the computation cost of reconstruction. The proposed method here performs an exhaustive search on a significantly reduced search space using an extremely efficient, after an initial overhead, reconstruction method.

## 6.2.2 Reducing the search space and computation cost

Three methods are now proposed as mechanisms for reducing the search space and the computation cost. The first method groups projections together to reduce the number of

objects that must be assigned a bin to reduce the search space. The second method applies principal component analysis (PCA) to reduce the dimensionality of vectors lowering the computation cost of reconstruction. The third method - when necessary - subdivides the search space into two search spaces which reduces the search space. These are each explained in the subsections below.

### 6.2.2.1 Bucketing

To reduce the search space in a reasonable fashion, the observation is first made that temporally adjacent projections, i.e. projections  $n$  and  $n + 1$  in a given scan, are more likely to belong to the same bin than to different bins. This leads to the idea of grouping projections together and assigning them in unison to one of the two bins. This idea of partitioning the projection set into sequential groups of projections requires some way of “throwing the switch” and changing the bin assignment. Otherwise, the grouping assumption will just traverse the projection set and assume inductively that each  $n + 1$  projection belongs to the same bin as the  $n^{th}$  projection’s bin assignment.

Presuming such a partitioning mechanism exists, the worst-case reduction in the search space can be obtained from the fastest respiratory period typically used in the literature: 4 seconds (See Chapter 2, Section 2.4). Using two bins means that a 4-second respiratory cycle will result in two groups of projections over the course of the 4 seconds, one group or bucket (the term which will subsequently be used) for each. Partitioning projections in this way over the course of a typical 120-second scan will result in:

$$\left( \frac{D \frac{\text{seconds}}{\text{scan}}}{T \frac{\text{seconds}}{\text{period}}} \right) \frac{2 \text{ buckets}}{\text{period}} = B$$

Where  $D$  is the duration of the scan,  $T$  is the period of a single breathing cycle, and  $B$  is the total number of required buckets. When  $D = 120$ ,  $T = 4$ , then  $B = 60$ . Thus this simple partition step reduces the search space from approximately  $2^{670}$  to approximately  $2^{60}$  which is significantly smaller but still intractable.

The next reduction technique is to introduce the two-scan protocol change. This change requires scanning the patient twice but with a scan protocol that uses half the time and half the number of projections. Each scan collects approximately 335 projections and is acquired in about 60 seconds (this is within the allowed parameters of the International Electric Commission recommendation; see [39]). Thus the total scanning time is nearly the same, the number of projections is the same at the same voltage and current levels, and therefore the patient dose is the same. This results in  $D = 60$ ,  $T = 4$  still, and  $B$  is now 30. However, this is done for two scans so our total number of objects

to assign to a given bin is still 60.

The most significant consequence of the two-scan protocol change already discussed in Section 5.3 but briefly repeated here is that two sets of projections are collected with each set containing projections at approximately the same angle (within  $670 \text{ projections} / 360 \text{ degrees} = .54 \text{ degrees}$ ) but *at different points in time*. The two projections from each scan that are closest to each other in terms of angle separation will henceforth be called paired projections. If paired projections are compared, the principal difference should be from motion and noise and not from geometry changes. This was confirmed in Chapter 5, Section 5.3. To compensate for potential noise differences in the method used, the two projections are first median filtered with a  $3 \times 3$  neighbourhood to minimize the noise while retaining edge features. In a worst case scenario, the two scans could be acquired with patient breathing perfectly in synch but this is highly unlikely given the variation within patient breathing.

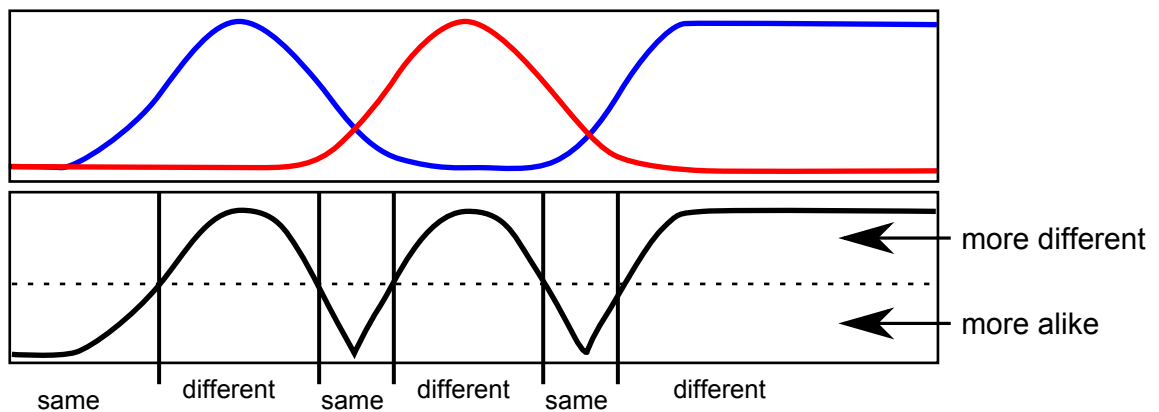


Figure 6.2: Illustration of the effect of partitioning based on paired projections differences and constructing buckets with same-different assignments. The top figure shows two one dimensional motion signals, one in blue and one in red. The bottom figure shows the difference between these motions and the partitioning that aggregates paired projections that are more alike and more different.

The outcome of this is that the differences between paired projections can be treated as a difference signal over time. This signal can then be processed to generate both a partition and very simple preliminary comparative classification of the groups in the partition. This classification states that the paired projections in a given bucket are either *more* or *less* alike. If a bucket contains paired projections that, on the aggregate, are more alike, then the bucket is preliminarily labelled as containing “similar” projections. If a bucket contains paired projections that, on the aggregate, are less alike, then the bucket is preliminarily labelled as containing “different” projections. This process of partitioning and creating a preliminary bucket assignment is henceforth called bucketing. Figure 6.2

illustrates this idea. Samples of a one dimensional motion are shown for two different hypothetical acquisitions, the difference signal is then shown, the partitioning based on the difference signal is shown and finally a similar/different classification is shown.

It should be pointed out that it is entirely possible to capture paired projections that are quite similar but do not, in fact, represent similar motion states. If the motion occurs along the path of the X-rays, it will be undetectable. The impact of such a mislabelling is benign; no information is “lost” because the same situation would eliminate the possibility of capturing the motion in a bin to begin with. In other words, to the reconstruction the same object looks the same in a given projection (minus small magnification effects) regardless of where along the path from source to detector it is positioned.

In the experiments, two difference signals will be used. The first is one derived from the ground truth. In all experiments, the position of the moving object is recorded and thus the actual differences between the object’s position in the two scans can be measured and transformed into a difference signal based on Euclidean distance. This allows the experiment to focus on the efficacy of the search method assuming the separate and independent difference signal generation process is perfected. A second difference signal shows an initial effort at identifying such a signal from the projections. The algorithm for identifying the signal is shown here:

```

Input: Projections from two independent scans
Output: differenceSignal
foreach Projection p1 in scan1Projections do
    p2 = Nearest(scan2Projections,p1.angle)
    /* See Section 6.2.2.4 for ROI masking description */
    m1 = ROIMaskTheProjection(p1)
    m2 = ROIMaskTheProjection(p2)
    abs = SumOfAbsoluteDifference(m1,m2)
    differenceSignal.append(abs)
end

```

**Algorithm 2:** Difference signal construction. Each projection is masked using an angle-specific projection of the ROI and the sum of absolute differences for all pixels in the projection masked region is calculated.

Having identified both the projection partitioning and a similar/different classification, the search space is modified slightly because the buckets in the independent scans are no longer themselves independent. From the pre-processing step, it is known whether the projections from scan 1 contained in bucket N belong to the same bin as the projections from scan 2, or whether they belong to a different bin. The original four possibilities for

each bucket are:

1. projections from scan 1 in bin 1, projections from scan 2 in bin 1
2. projections from scan 1 in bin 2, projections from scan 2 in bin 1
3. projections from scan 1 in bin 1, projections from scan 2 in bin 2
4. projections from scan 1 in bin 2, projections from scan 2 in bin 2

With the addition of an assessment of whether the projections in a bucket are from the same bin or from different bins, the decision changes. For the buckets identified as containing similar paired projections, the average of these paired projections are used in the reconstruction and no decision is required. For the buckets identified as containing different paired projections, there are now only two possibilities for each bucket labelled as containing different projections:

1. projections from scan 1 in bin 1, projections from scan 2 in bin 2
2. projections from scan 1 in bin 2, projections from scan 2 in bin 1

Thus, for buckets containing similar paired projections, no decision needs to be made. For buckets containing different paired projections, only two possible assignments can be made. Once the choices have been made for all the buckets, the resulting assignment can be used to construct a projection assignment and will itself be called a bucket assignment in contrast to the earlier stage preliminary bucket assignment. In some circumstances this may reduce the search space because of phase overlap in the two scans but in the ideal case where there is maximum phase difference, the number of buckets actually doubles. Figure 6.3 illustrates this. In the top figure, three buckets must be assigned ideally to bin 1 but this must be determined via a different kind of search (see Chapter 7) because the motion signal does not exist in a one-scan protocol. In the second figure, the preliminary steps construct a partition with six buckets (identified with the green plot), each of which will be assigned one of two states.

To summarize this reduction step, the size of the search space is significantly reduced by bucketing the projections and splitting the scan into two scans. This provides a basis for automatically generating the partition for the buckets and using the information contained in the difference signal to reduce the bin decision choices from 4 to 2. Instead of an assignment which tags each projection with a bin, this step generates a preliminary partition of the paired projections with a same/different assignment for each partition.




Figure 6.3: Partitioning differences between one-scan and two-scan approach. In the top figure, the motion of interest is characterized as a one-dimensional plot (black line) with the green plot identifying the projections belonging to one bin and the red plot identifying the projections assigned to the second bin. In the bottom figure, the motion of two scans are identified by the solid and dashed black lines. The green plot shows the projections that are labelled as belonging to different bins while the red plot shows projections labelled as belonging to both bins.

Subsequently, this will be used to generate a true *bucket assignment* where the projections grouped in a bucket are assigned to either the first bin, the second bin, or both (i.e. the reconstruction step uses a projection constructed from the average of the paired projections at that position).

### 6.2.2.2 Subdividing the search space

Even using the reductions from Section 6.2.2.1, the search space size is still (potentially) too large. Taking the worst case scenario, the continuously periodic respiration with a short breathing cycle of four seconds, and assuming the desired case of non-overlap of breathing cycles between scans, the resulting number of buckets to assign a value to in the search is still approximately 30. This results in, since the assignment is a binary assignment into different bins,  $2^{30}$  assignments which, even given the performance improvements still to be outlined, is too large.

One of the main results from Chapter 5 is used here, namely the fact that uniformly distributed projections create fewer artefacts than random or grouped projections. By taking an assignment space and uniformly removing every other bucket to be assigned (and replacing with projections from an average set as was outlined in Section 5.1), the search is first performed on this subspace and then, once an optimum solution is found,



the search is renewed by searching for the previously ignored assignments but adding them to the optimum solution obtained in the first pass. For instance, suppose a bucketing partition has the *preliminary* assignment:

s d s d s d s d s d s d s d s d s d s d s d s d s d s d

where ‘s’ identify buckets that contain projections that should be assigned to the same bin, and ‘d’ represent buckets that contain projections which should be assigned to opposite bins (either 1,2 or 2,1). A first pass search can be performed making assignments to only half of the candidate buckets:

s d s - s d s - s d s - s d s - s d s - s d s - s d s -

where dashes that are actually “different” buckets are temporarily treated as “same” buckets, i.e. they are filled with average projections during the reconstruction process. Now let ‘1’ denote a “different” *preliminary* assignment resulting in a *final* bucket assignment wherein the projections from pass 1 are assigned to bin 1 and the projections from pass 2 are assigned to bin 2. The marker ‘2’ then defines the opposite wherein the projections from pass 1 are assigned to bin 2 and the projections from pass 2 are assigned to bin 1 for that given bucket. Assume the optimum solution returned is:

s 1 s - s 1 s - s 2 s - s 1 s - s 2 s - s 2 s - s 1 s -

The second pass now seeks to fill in the missing “different” preliminary assignments (represented by the dashes) by trying alternative final bucket assignments while using and protecting the final assignments from the first pass. The following sequences show the first set of assignments that would generated and evaluated:

s 1 s 1 s 1 s 1 s 2 s 1 s 1 s 1 s 2 s 1 s 2 s 1 s 1 s 1  
s 1 s 2 s 1 s 1 s 2 s 1 s 1 s 1 s 2 s 1 s 2 s 1 s 1 s 1  
s 1 s 1 s 1 s 2 s 2 s 1 s 1 s 1 s 2 s 1 s 2 s 1 s 1 s 1  
s 1 s 2 s 1 s 2 s 2 s 1 s 1 s 1 s 2 s 1 s 2 s 1 s 1 s 1  
...

Initially, all the remaining “different” preliminary bucket assignments are replaced with a final assignment ‘1’ as described above. This is evaluated against a change in the first unprotected bucket, in position 4, which is modified to the assignment ‘2’. The next unprotected bucket is in position 8 which is changed while returning position 4 back to ‘1’ and so on. Given this approach, the  $2^{30}$  assignment space is reduced to a  $2^{16}$  assignment space ( $2 \times 2^{15}$ ) which is, when combined with the remaining enhancements, computationally tractable.

### 6.2.2.3 PCA of individual backprojections

Despite the bucketing reduction, if reconstruction could even be reduced to one second, the time required to exhaustively search the space defined above by generating two test reconstructions for each candidate assignment (one for bin 1, one for bin 2) is over 18 hours. Certainly in some scenarios, such as off-line batch processing for a retrospective study, this may be acceptable but it clearly puts severe limitations on the method. It is at this point that a classic trade-off in computer science is made. To reduce the computational time required, the memory requirements are increased.

In Chapter 2, it was established that filtered backprojection is effectively an averaging process. After each projection is filtered, it is “smeared” through the reconstruction volume by adding the pixel values to the voxel values in the volume as the geometry dictates. This step is repeated for each projection usually with a pre-weighting of the projections such that the resultant volume is averaged once the final projection’s values have been added to it.

The novel idea here is to perform that backprojection for each projection, *but on individual backprojection volumes*. In other words, rather than ending up with a single backprojection volume containing the averaged values from all of the projections, the same number of volumes as projections are generated with each volume containing just that single projection’s values “smeared” through it. If these volumes are then averaged together, the result is identical to averaging them “in place” to generate the single backprojection volume. Figure 6.4 illustrates this showing the equivalence of the two different methods.

The goal of introducing this memory overhead is that the projections, after backprojecting, can now be treated as very high dimensional vectors that can be averaged to create a reconstruction volume. By doing so, principal component analysis (PCA) can be applied to the set to reduce the dimensionality. It is worth commenting that this is one of two ways PCA is used to reduce the dimensionality of a data set. When the number of samples greatly exceeds the number of features in each sample, the dimensionality can be reduced by using PCA to find the most important features and ignore the others. This is a “lossy” form of PCA. The other form used here, when the number of features is greater than the number of samples, is lossless. PCA provides a mechanism for finding an optimal basis for  $N$  vectors, regardless of their size, of dimension  $N - 1$ . In the case of this problem, the backprojected volumes are the samples and each voxel represents a feature. Given an average reconstruction volume with  $256^3$  voxels, the 670 backprojected volumes can be transformed into 670 vectors each 669 elements long. This is a vast improvement over the 16,777,216 elements that would otherwise be required for each vector.

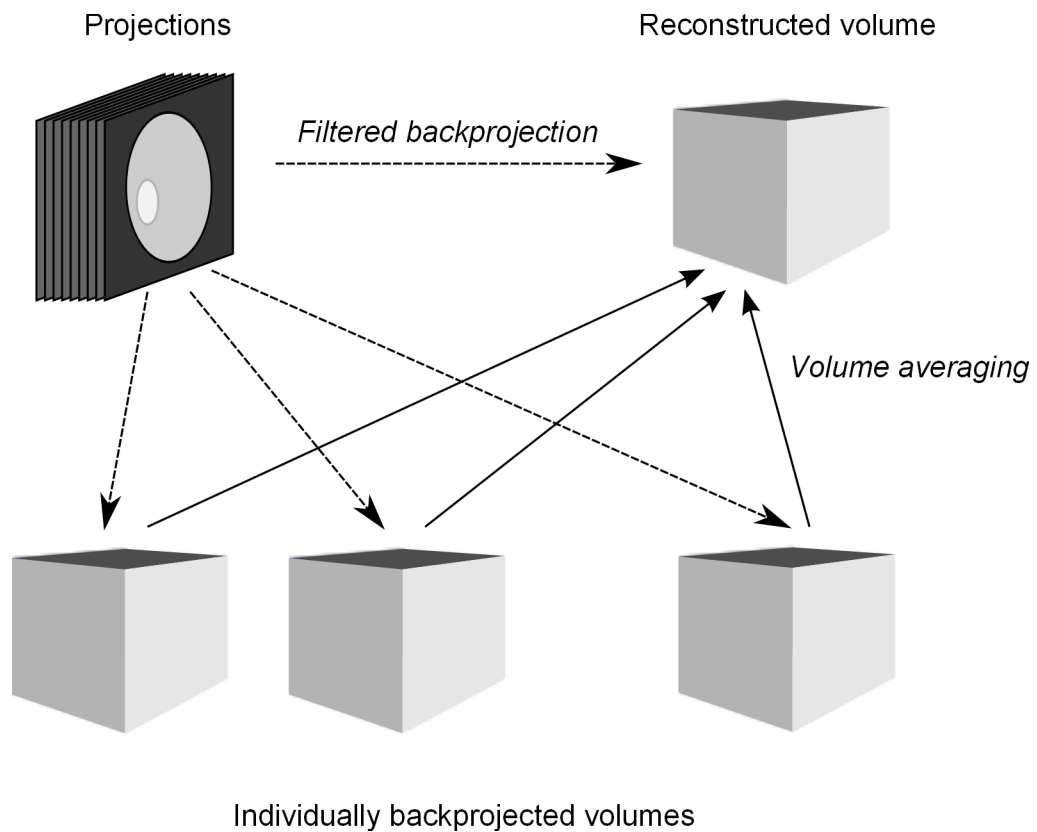


Figure 6.4: A reconstructed volume constructed two ways. Along the top, the projections are reconstructed directly into a volume using a filtered backprojection algorithm. On the bottom, the projections are individually backprojected into separate volumes which are then averaged together.

More importantly, the averaging process has just changed from adding 16,777,216-length vectors together to adding 670-length vectors together (and then dividing by 670). This, ignoring memory caching issues, requires less than 0.004% of the computation effort.

PCA is not a trivial process, however. Constructing a covariance matrix of the *features* of the set of vectors in order to find the eigenvectors can require a large amount of memory. It must be done in stages using large amounts of disk space unless the technique popularized in the computer vision community by Turk and Pentland [92] is used. This clever trick is used to find the *feature* eigenvectors by first finding the *sample* eigenvectors and deriving them from there. Very simply, this is shown by:

$$Av = \lambda v \quad (6.2.1)$$

$$B^T Bv = \lambda v \quad (6.2.2)$$

$$B^T BV = V\Lambda \quad (6.2.3)$$

$$(BB^T)BV = BV\Lambda \quad (6.2.4)$$

$$W = BV \quad (6.2.5)$$

$$X = W^T B \quad (6.2.6)$$

Equation 6.2.1 is the definition of an eigenvector  $v$  with eigenvalue  $\lambda$ . Let  $B$  be the array composed of the clipped backprojection volumes as column vectors with the empirical mean of the dataset subtracted from each vector. Equation 6.2.2 shows an eigenvector for the sample covariance matrix and Equation 6.2.3 defines the complete set of eigenvectors  $V$  and their eigenvalues in a diagonal  $\Lambda$  matrix. Multiplying both sides by  $B$  again (Equation 6.2.4) shows that the eigenvectors of the sample covariance matrix, when multiplied by the original matrix, become the eigenvectors of the feature covariance matrix (Equation 6.2.5) which is the objective. This method achieves the objective by calculating the eigenvectors of a matrix many, many orders of magnitude smaller (around  $10^{-9}$  if complete volumes were used) than the straightforward samples covariance matrix. Once these eigenvectors are constructed, the original data vectors can be dimensionally reduced by taking them into the row space of  $W$  as is shown in Equation 6.2.6.

Each vector in  $X$  is a dimensionally reduced version of the original backprojection volume. Reconstruction consists now of averaging together the vectors, *or a particular subset of the vectors*, projecting them back to the original space, and adding the mean back. For the purposes of the proposed search, these last two steps are not needed because the evaluation of a given proposed assignment is constructed as a function of the dimensionally reduced vectors.

Consider the first metric used in the objective function: the sum of squared differences between two reconstructed volumes. This metric is exactly the same as the Euclidean distance (squared) between the averaged transformed vectors. As a computational aside, this also has the nice property of being most efficiently accomplished as a dot product of the difference vector with itself on modern CPUs and GPUs which include very fast instructions for this kind of operation. The transformation matrix  $W$  generated by PCA is composed of orthonormal vectors whose rank is equal to the rank of  $X$ . Thus, the distance between two vectors is preserved after being transformed so the sum of squared difference metric can be calculated in eigenspace. This allows the optimization function to

calculate the score for a candidate assignment without requiring an inverse transformation back into the original vector space. For the expected number of assignments generated in an exhaustive search, avoiding this inverse transformation is an important part of the computational cost savings of this method.

The application of the mean metric is somewhat less straightforward. A short proof is presented here which shows that the mean of some vector in the original space is equal to the dot product of the means of the columns of the  $W$  matrix and the projected original vector. First, to disambiguate between the dual uses of the averaging idea, the term “average of vectors” refers to the vector addition of some set of vectors multiplied by the scalar reciprocal of the number of vectors. This is different than the “vector mean” which will refer to the sum of the components of a given vector divided by the total number of components. Furthermore, in the following discussion if a vector (in bold) is shown with two subscripts, the first subscript refers to an element of the vector while the second subscript refers to a particular vector from some set of vectors. With that in mind, some preliminary equations are described:

$$\mathbf{b} \in \text{columns}(B) \quad (6.2.7)$$

$$W^T \mathbf{b} = \mathbf{x} \quad (6.2.8)$$

$$W^T B = X \quad (6.2.9)$$

$$WW^T B = WX \quad (6.2.10)$$

$$B = WX \quad (6.2.11)$$

$$WW^T \mathbf{b} = W\mathbf{x} \quad (6.2.12)$$

$$\mathbf{b} = W\mathbf{x} \quad (6.2.13)$$

These equations describe the relationship between the original data space (with average of vectors subtracted) and the transformed space. Vector  $\mathbf{b}$  is in the column space of  $B$  either as a column vector or as a linear combination of column vectors belonging to  $B$ . Vector  $\mathbf{x}$  belongs to the column space of the dimensionally reduced transform space  $X$  which is itself created by multiplying  $B$  by the principal components in  $W^T$  which is an orthogonal matrix so the transpose is the inverse. Using this property, the  $B$  matrix or a  $\mathbf{b}$  vector is easily recovered by multiplying the  $X$  matrix or  $\mathbf{x}$  vector by  $W$ . Now let  $\bar{\mathbf{b}}$  be some average of vectors in  $B$ . The transformed version can be derived either by multiplying  $\bar{\mathbf{b}}$  by the principal component vectors or by averaging the same set of vectors in the transformed space. If this is done, the  $\bar{\mathbf{b}}$  vector can be obtained from the  $\bar{\mathbf{x}}$  vector by

multiplying it by  $W$ . This step can also be described as multiplying each of the column vectors in  $W$  by the correct element in  $\bar{\mathbf{x}}$  and adding them together. Averaging the values in the vector, however it is obtained, yields the scalar vector mean which is the desired value. More precisely, this is expressed as:

for the  $k^{th}$  column vector  $\mathbf{x}^k$ , the  $j^{th}$  element of the  $\mathbf{b}^k$  vector is

$$\mathbf{b}_j^k = \sum_{i=1}^N \mathbf{x}_i^k \mathbf{W}_{j,i} \quad (6.2.14)$$

the average of elements of the vector are

$$\bar{b} = \frac{1}{M} \sum_{j=1}^M \mathbf{b}_j^k \quad (6.2.15)$$

or, by substitution

$$\bar{b} = \frac{1}{M} \sum_{j=1}^M \sum_{i=1}^N \mathbf{x}_i^k \mathbf{W}_{j,i} \quad (6.2.16)$$

and rearranging

$$\bar{b} = \sum_{i=1}^N \mathbf{x}_i^k \left( \frac{1}{M} \sum_{j=1}^M \mathbf{W}_{j,i} \right) \quad (6.2.17)$$

but the averages of the column vectors of  $W$  can be expressed as

$$\bar{\mathbf{w}} = \frac{1}{M} \sum_{j=1}^M \mathbf{W}_{j,i} \quad (6.2.18)$$

therefore

$$\bar{b} = \bar{\mathbf{w}} \cdot \mathbf{x}^k \quad (6.2.19)$$

$$(6.2.20)$$

The last step, Equation 6.2.20 is derived from the fact that the summation on the

individual elements of the  $\mathbf{w}$  vectors is independent of the summation with respect to  $j$  and so the summation can be moved inside of the summation involving the elements of the  $\mathbf{x}$  vector. This, if written out in long form as an extended set of additions of multiples, would be equivalent to factoring out the  $\mathbf{x}$  components. A closer look at this inner summation reveals that it is just the vector mean of the individual column vector  $\mathbf{w}$ . Therefore, during the PCA process, if the vector means of the columns of  $W$  are calculated once and stored in row vector, then the calculation of the mean of the averaged reconstruction vector  $\bar{\mathbf{b}}$  is simply the dot product of the  $W$  mean row vector and the averaged transformed vector  $\bar{\mathbf{x}}$ .

A further reduction even before applying the PCA comes from the realization described earlier in the discussion of objective functions that the motion of interest is found in the region of interest. If the reconstructed volumes are first clipped to this ROI, and the PCA is performed on that, then typical workstations (as of the time of this thesis) are sufficiently powerful to perform PCA on the datasets. In all the experiments performed, disk access proved to be the performance limiting factor with memory requirements and processor load being insignificant in comparison.

In the previous paragraphs, the discussion of PCA suggested it be applied to the original set of CBCT projections which have been backprojected into individual volumes. In fact, the number of projections, and therefore volumes, is 50% greater. The method relies on filling the gaps in assignments with “average” projections, namely the projections generated by averaging paired projections together. Thus, instead of 670 backprojections, 1005 backprojections are used to calculate the feature eigenspace.

#### 6.2.2.4 Review of the complete bin assignment method

Each of the prior subsections presents a part of the overall process. They are collected here and presented as algorithms so that the interactions and connections can be better comprehended. One step that has not been described elsewhere is the generation of 2D masks for each projection. These are constructed from the ROI volume by forward projecting a shadow of the ROI onto the projection plane at the same angle as each of the projections in the scan set.

The preparatory process for the two-scan search involves first constructing a set of “filler” projections by averaging together paired projections. These projections are used when a bucket from the bucket assignment is labelled as containing similar projections. Each projection is then individually filtered and backprojected. These backprojected volumes are then clipped to the ROI and these clipped voxels are concatenated to form a data array of column vectors containing the samples with rows containing the features

```

Input: Projections from two independent scans
Output: PCA feature eigenspace and mean vector
foreach Projection p1 in scan1Projections do
    |  $p2 = \text{Nearest}(\text{scan2Projections}, p1.\text{angle})$ 
    |  $p3 = (p1 + p2)/2$ 
    |  $p3.\text{angle} = (p1.\text{angle} + p2.\text{angle})/2$ 
    |  $\text{avgdProjections.add}(p3)$ 
end
foreach Projection p1 in scan1Projections do
    |  $bp = \text{Backproject}(p1)$ 
    |  $B.\text{append}(\text{ROIClipTheVolume}(bp))$ 
end
foreach Projection p2 in scan2Projections do
    |  $bp = \text{Backproject}(p2)$ 
    |  $B.\text{append}(\text{ROIClipTheVolume}(bp))$ 
end
foreach Projection p3 in avgdProjections do
    |  $bp = \text{Backproject}(p3)$ 
    |  $B.\text{append}(\text{ROIClipTheVolume}(bp))$ 
end
 $B_{\text{avg}} = \text{VectorAverage}(B)$ 
foreach Vector v in B do
    |  $v = v - B_{\text{avg}}$ 
end
 $B = \text{PCA}(B)$ 

```

**Algorithm 3:** PCA on backprojected, clipped volumes.

– the voxels. The average column vector of the data array is then calculated and subtracted from the original vectors to create data vectors whose feature means are each zero. Finally, principal component analysis using the covariance method described earlier in Section 6.2.2.3 is performed.

Reviewing the steps discussed so far, the paired projections are first bucketed (Section 6.2.2.1) creating a preliminary bucket assignment. Then the original sets from the two scans as well as the averaged set of projections are clipped to the ROI. These clipped regions are treated as high-dimensional vectors and PCA is applied in a lossless fashion to reduce their dimensionality. Once these steps have been performed, an exhaustive search is used to find a bucket assignment that maximizes an objective function which can then be used to construct a projection assignment.



```

Input: Eigenspace vectors from PCA process (Algorithm 3)
Output: Best scoring binning assignment
differences = ScanDifferences(scan1Projections,scan2Projections)
bkts = Partition(differences)
repeat
    a = GenerateRandomAssignment()
     $\mathbf{v}_1, \mathbf{v}_1$  = ReconstructAssignment(a)
    avg1 = MeanOfVector( $\mathbf{v}_1$ )
    avg2 = MeanOfVector( $\mathbf{v}_1$ )
    avgmax = Max(avg1,avg2,avgmax)
until 5000 times
if Count(buckets == different) > 15 then
    | firstSet,secondSet = SplitBuckets(bkts)
    | passes = 2
end
else
    | firstSet = buckets
    | passes = 1
end
/* This is the start of the actual exhaustive search */
bestScore = 0
bestAssignment = []
permutationCount =  $2^{\text{Count}(\text{firstSet} == \text{different}) - 1}$ 
for permutation=0 to permutationCount-1 do
    | a = GenerateAssignmentFromPermutation(permutation,firstSet)
    |  $\mathbf{v}_1, \mathbf{v}_1$  = ReconstructAssignment(a)
    | avg1 = MeanOfVector( $\mathbf{v}_1$ )
    | avg2 = MeanOfVector( $\mathbf{v}_2$ )
    |  $SSD = (\mathbf{v}_1 - \mathbf{v}_2)^T \cdot (\mathbf{v}_1 - \mathbf{v}_2)$ 
    | score = ObjectiveFunction(avg1,avg2,avgmax,SSD)
    | if score > bestScore then
    | | bestScore = score
    | | bestAssignment = a
    | end
end
if passes == 2 then
    | /*Same search as prior loop only the assignment is generated by combining the
    | | found bestAssignment with the secondSet pre-assignment.*/
end

```

**Algorithm 4:** Two-scan search.

```

Input: PCA feature eigenspace and assignment
Output: Two eigenspace vectors containing the two reconstructions
 $\mathbf{v}_1 = 0$ 
 $\mathbf{v}_2 = 0$ 
 $v1\_count = 0$ 
 $v2\_count = 0$ 
/* Add vectors assigned to bins 1, 2 */
foreach EigenspaceBackProjection ebp associated with scan1 do
    if ebp is assigned to bin 1 then
         $\mathbf{v}_1 = \mathbf{v}_1 + \mathbf{ebp}$ 
        Increment( $v1\_count$ )
    end
    if ebp is assigned to bin 2 then
         $\mathbf{v}_2 = \mathbf{v}_2 + \mathbf{ebp}$ 
        Increment( $v2\_count$ )
    end
end
foreach EigenspaceBackProjection ebp associated with scan2 do
    if ebp is assigned to bin 1 then
         $\mathbf{v}_1 = \mathbf{v}_1 + \mathbf{ebp}$ 
        Increment( $v1\_count$ )
    end
    if ebp is assigned to bin 2 then
         $\mathbf{v}_2 = \mathbf{v}_2 + \mathbf{ebp}$ 
        Increment( $v2\_count$ )
    end
end
/* Fill in gaps with averaged vectors */
foreach EigenspaceBackProjection ebp associated with averaged set do
    if ebp is assigned to neither bin 1 nor bin 2 then
         $\mathbf{v}_1 = \mathbf{v}_1 + \mathbf{ebp}$ 
        Increment( $v1\_count$ )
         $\mathbf{v}_2 = \mathbf{v}_2 + \mathbf{ebp}$ 
        Increment( $v2\_count$ )
    end
end
/* Divide by number of added vectors to create averages */
 $\mathbf{v}_1 = \mathbf{v}_1 / v1\_count$ 
 $\mathbf{v}_2 = \mathbf{v}_2 / v2\_count$ 

```

Algorithm 5: ReconstructAssignment

```

Input: difference signal  $S$ , local max threshold  $th$ 
Output: partition  $P$ , preliminary bucket assignment  $A$ 
 $S = \text{Normalize}(S)$ 
 $S = \text{Smooth}(S, \text{3-element box kernel})$ 
 $L = \text{LocalMinimums}(S)$ 
 $lmcount = L.\text{count}()$ 
for  $lm = 0$  to  $lmcount - 1$  do
     $xstart = L(lm)$ 
     $xend = L(lm + 1) - 1$ 
    if some value  $x$  between  $xstart$  and  $xend$  is greater than  $th$  then
         $p1 = \text{FindFirst}(x > .5 \text{ between } xstart \text{ and } xend)$ 
         $p2 = \text{FindLast}(x > .5 \text{ between } xstart \text{ and } xend)$ 
         $index = P.\text{addPartition}(p1, p2)$ 
         $A.\text{appendDifferentBucket}(index)$ 
    end
end

```

**Algorithm 6:** Partition algorithm.

One way to represent the bucket assignments is with a vector containing a binary value for each of the “different” buckets indicating whether the assignment is from scan 1 to bin 1 and scan 2 to bin 2 or vice versa. Another way to represent it is as a single number from the permutation set of  $2^n$ . A complete exploration of this set is redundant because there is no meaning to “bin 1” versus “bin 2”; there are just two bins. Thus an assignment of projections 1..m to bin 1 and projections m+1..n to bin 2 is equivalent to an assignment of projections 1..m to bin 2 and projections m+1..n to bin 1. Both result in the same two volume estimates but with swapped names. A complete set of permutations would test for both cases so the second “swapped name” set of assignments is dispensed with. This reduces the size of the permutation space by half to  $2^{n-1}$ .

Assignments are created by converting a number  $0 \leq n < 2^{n-1}$  into a sequence of numbers corresponding to the bucketing indicating which set the projections/volumes identified by a bucket should be drawn from. These bucket assignments are then converted into projection assignments using the partition information constructed during the bucketing process. This projection assignment is then used to construct two candidate vectors, one for each bin. The candidate vectors consist of the vector average of the (transformed) projections belonging to the bin being estimated and the (transformed) filler projections from the averaged set. The mean of each of these is then calculated as is the SSD between them. These are passed to the objective function to be scored along with the estimate of a maximum vector mean obtained from a random trial involving an empirically determined

```

Input: Preliminary assignment A
Output: A1, A2, Two partial preliminary assignments
A1 = A1 = A
toggle = 0
for n = 0 to A.count() do
  if A(n) is labelled "different" then
    if toggle is even then
      label A1(n) "same"
    else
      label A2(n) "same"
    end
    Increment(toggle)
  end
end

```

**Algorithm 7:** SplitBuckets

number of samples (5000). The assignment with the highest score is retained and doubled. In the case where the number of buckets exceeds the empirically determined threshold of 15, the iterative approach described earlier is applied.

```

Output: score
avg_biggest = max(avg1, avg2)
avg_score = max(0, (1 - (avg_biggest / avg_max)))
score = (vector1 - vector2)T · (vector1 - vector2) * avg_score

```

**Algorithm 8:** Objective scoring function

The objective scoring function is very simple:  $score = SSD * Penalty(\mu)$ . It seeks volumes whose SSD is the largest while retaining the mean characteristic of the original reconstructed volume. Since the mean has been removed from the data as part of the PCA process, the new mean target is zero. The implication is that the vector mean of some candidate volume is closest to the original complete projection reconstruction mean when it is closest to zero. This property is used to construct a simple penalty function as a function of  $\mu$  the largest of the two vector means. By constructing a linear ramp with the highest value at  $\mu = 0$  and terminating it at some estimated globally maximal  $\mu_{max}$ , and setting any found  $\mu$  values that might exceed this estimate to 0, a crude but very efficient and, in the experiments, effective way of constructing a penalty function is created.

## 6.3 Evaluation and results

Two novel ideas have been described in this chapter, both of which are necessary for automatic binning from data. The search space with its objective function and the extraction of a difference signal from the paired projections both impact the results of the search and poor results on either component can mask good results on the other. Consequently, the experiments below have been run with a ground-truth version of the difference signal, to isolate the objective function aspect of the search method, and they have also been run with a simple difference signal extraction method.

The experiments consist of performing the search on two scans simulated using the phantom configurations described in detail in Chapter 3. Each phantom class contains specific configurations designed to test particular aspects and operational boundaries of the method and results are presented in the context of each individual result and then summarized at the end.

### 6.3.1 Using a ground-truth derived difference signal

In these experiments, the actual position of the moving object of interest is used to generate the difference signal. For each scan, the Euclidean distance between the position of the objects at the two times associated with each projection in the paired projections is calculated and this is considered the difference.

In Figure 6.5, the bucketing performed using Algorithm 6 is shown for the Simple phantom configurations used in these tests. In each sub-figure, the difference signal derived from the ground-truth measurements of the position of the object is shown as a function of the sequence of paired projections with a black line. The *preliminary bucket assignment* is expanded and shown as blue and red bars covering the paired projections sequence. The upper blue bars identify paired projections that should be considered as belonging to different bins. The lower red bars identify paired projections that should be considered as belonging to both bins. Note in the upper three configurations, the signal has been normalized so the amplitude appears the same even though absolute difference ranges will be different due to different motions. The period difference in configuration 4 is evident from the signal difference relative to the top three configurations. Configuration 5 shows a shift motion covering 50% of the total scan time. For convenience, this has been accomplished by essentially switching the position of the object between the two separate scans. The transition is smooth and thus the tails appear at either end of the paired projections sequence. The preliminary same-different bucketing using the ground-truth position information appears to work well in the Simple case. This focuses

the results of the search on the search mechanism itself, and not on poorly partitioned paired projections.




Figure 6.5: Ground-truth derived difference signals and resulting bucketing for Simple phantom configurations. The top blue bar shows projections that are grouped into “different” buckets; the bottom red bar shows projections grouped into “same” buckets. The data is normalized in each case before displaying; higher values represent greater differences. Bucketing bars are placed at the 33% and 67% positions for clarity.

### 6.3.1.1 Simple virtual configuration 1, ground truth difference signal

These first sets of results are obtained using the Simple class of phantom described in Chapter 3, Section 3.3. This particular configuration mimics an average sized lung tumour with average motion amplitude and a *continuously periodic motion* (Lujan) in both the SI and AP direction. In Figure 6.6, the first row contains the results of the search. The search actually returns back a bucket assignment and this has been expanded using the partitioning information in the bucketing to show the projection assignment. For this case, the search returned a perfect (relative to the ground truth) assignment.

Figure 6.7 shows images from the reconstructed bin volumes. Along the top row are sagittal slices that best reveal the motion of the object. These images are representative

of all of the Simple phantoms and are not repeated in other Simple configuration figure results. The second row contains magnified regions where the motion has occurred. The first image shows the results that would have been obtained from a volume reconstructed from the complete 670 projections. Blur is evident but, because of the Lujan motion, the object was more commonly located in the upper left of the image (the “rest” phase of the simulated respiratory cycle) and this is apparent. The next two images show the ground truth and the estimated volumes, respectively, for one bin (arbitrarily called “bin 1”). The final two images show the ground truth and estimated volumes, respectively, for the other bin. No differences are evident and this is to be expected since the search returned a nearly equivalent binning assignment (one projection assignment difference) as the ground-truth derived assignment. Note that the slices from the bin on the right reveal where the tumour is positioned for the majority of the time. This information is potentially useful for evaluating treatment plans.

#### **6.3.1.2 Simple virtual configuration 2, ground truth difference signal**

This configuration mimics an average sized lung tumour with short motion amplitude and Lujan motion in both the SI and AP direction. This particular configuration tests the effect of very little motion which may be the case in lung tumours located in the upper lung or in regions of the abdomen only indirectly affected by the diaphragm motion. Referring to the second line of Figure 6.6, the performance is not as good as with the first configuration. This is expected because the energy contained in the differences caused by the motion are much less than in the first configuration and therefore the signal to noise ratio is lower (the noise is the same in both data sets). The error, as a function of the number of incorrectly assigned projections is about 7.8%. Looking at the qualitative results in Figure 6.8, it is apparent when comparing the estimates with the ground truth that even with the given error rate, the results still “look” good. This is especially true with the second bin figures which shows the most significant location of the tumour.

#### **6.3.1.3 Simple virtual configuration 3, ground truth difference signal**

This configuration mimics a large lung tumour with short motion amplitude and Lujan motion in both the SI and AP direction with a minimum (4 seconds) breathing period. As with the second configuration above, the error rate is greater because of the lack of motion information caused by the smaller motion. In this case, the error rate is approximately 11.3%. Given that the error rate is relatively high despite changing the size of the tumour, the idea that smaller motion ranges are potentially problematic is suggested. Once again,

looking at the qualitative results in Figure 6.9, it can be seen that the results are still quite acceptable.

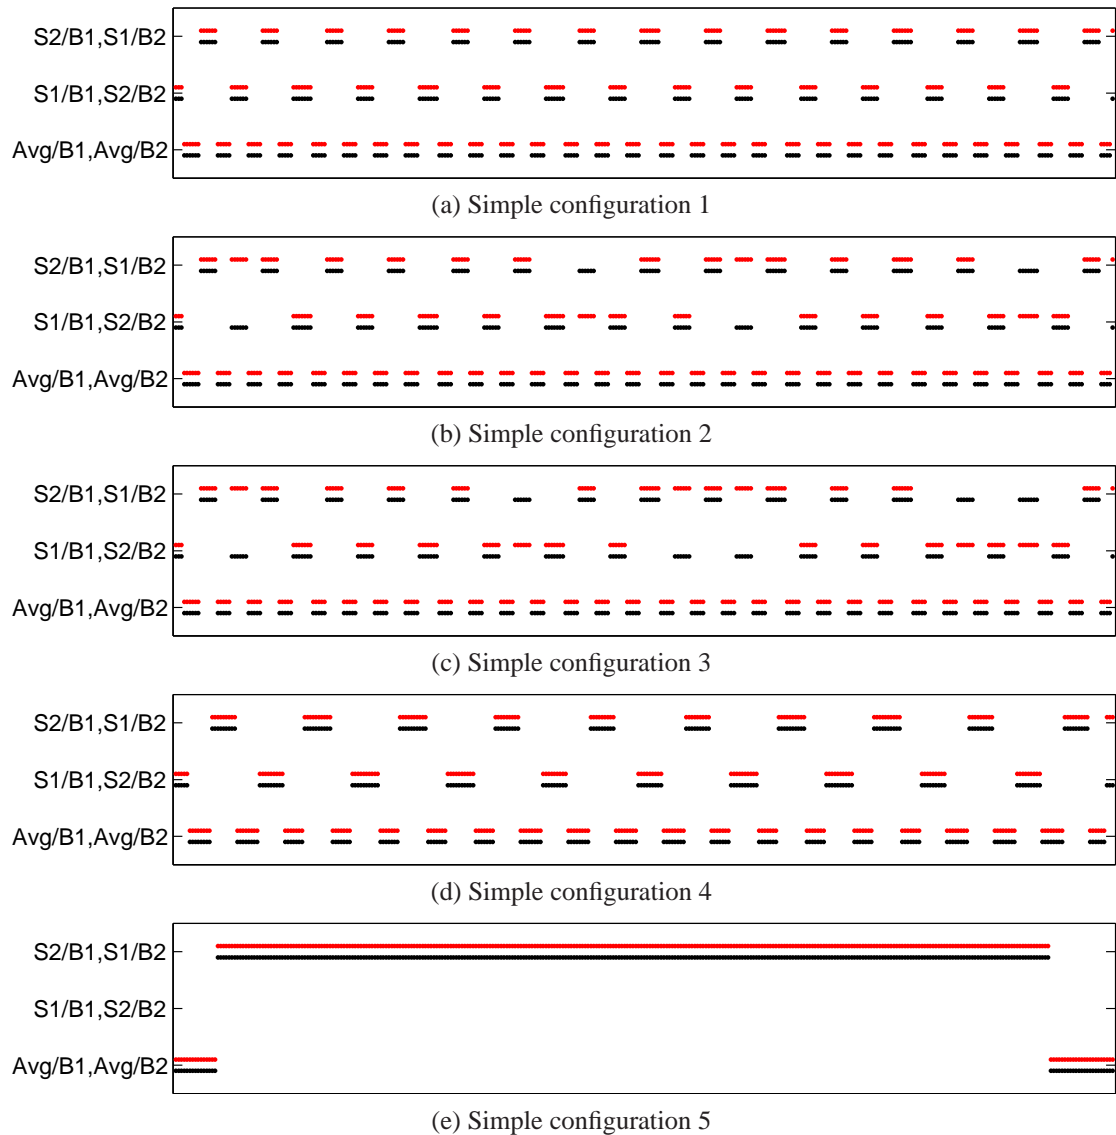


Figure 6.6: Ground-truth derived difference signal based search results for Simple phantom configurations. The black dots each identify a projection pair and its correct assignment. Red dots identify the resulting assignment found by the search. The lower set of dots refer to projections from the averaged set which are added to both of the bin reconstructions. The middle row shows the paired projections for which the projection from the first scan is assigned to bin 1 and the projection from the second scan is assigned to bin 2. The top row shows the paired projections for which the projection from the first scan is assigned to bin 2 and the projection from the second scan is assigned to bin 1. See Table 6.1 for error percentages.






Figure 6.7: Qualitative results of search on Simple configuration 1, the averaged-sized tumour travelling an average distance, as define in Section 3.3.3, in a Lujan motion over a 4 second period. Top row shows sagittal slices from a complete reconstruction (a), from a ground truth bin 1(b), from the estimated bin 1(c), from the ground truth bin 2 (d), and from the estimated bin 2 (e). The second row contains magnifications of the regions from the images above.




Figure 6.8: Qualitative results of search on Simple configuration 2, an averaged sized tumour travelling a short distance in a Lujan motion over a 4 second period. Shown are magnified regions of a sagittal slice from a complete reconstruction (a), from a ground truth bin 1(b), from a the estimated bin 1(c), from the ground truth bin 2 (d), and from the estimated bin 2 (e).

#### 6.3.1.4 Simple virtual configuration 4, ground truth difference signal

This configuration mimics a large lung tumour with average motion amplitude and Lujan motion in both the SI and AP direction with a longer (6 seconds) breathing period. A greater range of motion has dropped the error rate back to around 0.9% supporting the idea that, when compared with the results from configuration 3 above, size is not the issue, the range of motion is. With the smaller error rate, Figure 6.9 shows the qualitative results from representative sagittal slices are quite good.




Figure 6.9: Qualitative results of search on Simple configuration 3, a large tumour travelling a short distance in a Lujan motion over a 4 second period. See Figure 6.8 for explanation.




Figure 6.10: Qualitative results of search on Simple configuration 4, a large tumour travelling a large distance over a 6 second period. See Fig. 6.8 for the explanation.

### 6.3.1.5 Simple virtual configuration 5, ground truth difference signal

This configuration mimics a prostate-sized, ellipsoid-shaped tumour with a 10 mm shift at the 50% point which, for the two scan protocol, means most of the motion in the first position is obtained in the first scan and then the shift occurs for the second scan. As in the earlier cases (which differed as well in their use of spherical tumour shapes), the motion is sufficiently large to reduce the error rate. In this case there are zero errors so the ground truth and estimated reconstructions for the two bins are identical.




Figure 6.11: Qualitative results of search on Simple configuration 5, 10 mm shift motion 50% of the total scan time. See Fig. 6.8 for the explanation.

### 6.3.1.6 Animated configuration 1, ground truth difference signal

These next results are obtained using the Animated phantom class derived from real CBCT data and discussed in Chapter 3, Section 3.4. In Figure 6.12, the bucketing performed using Algorithm 6 is shown for the Animated phantom configurations used in these tests. See Figure 6.5 for an explanation of the plots.

This first Animated configuration contains a single shift of 16 mm occurring half-way through the total scan cycle. Using the ground truth difference signal, the resulting search method returns an assignment identical to the ground truth assignment. This is seen in the




Figure 6.12: Ground-truth derived difference signal based search results for Animated phantom configurations. See Fig. 6.5 for the explanation.

comparison of the first row of Figure 6.13. The qualitative results shown in Figure 6.14 are therefore identical for both the ground truth and estimated bin volumes.

### 6.3.1.7 Animated configuration 2, ground truth difference signal

This configuration contains a single shift of 12 mm occurring half-way through the total scan cycle. The search results return two wrong paired projections assignments out of the possible 335 for an error rate of approximately 0.5%. The qualitative results in Figure 6.15 are consequently very good.

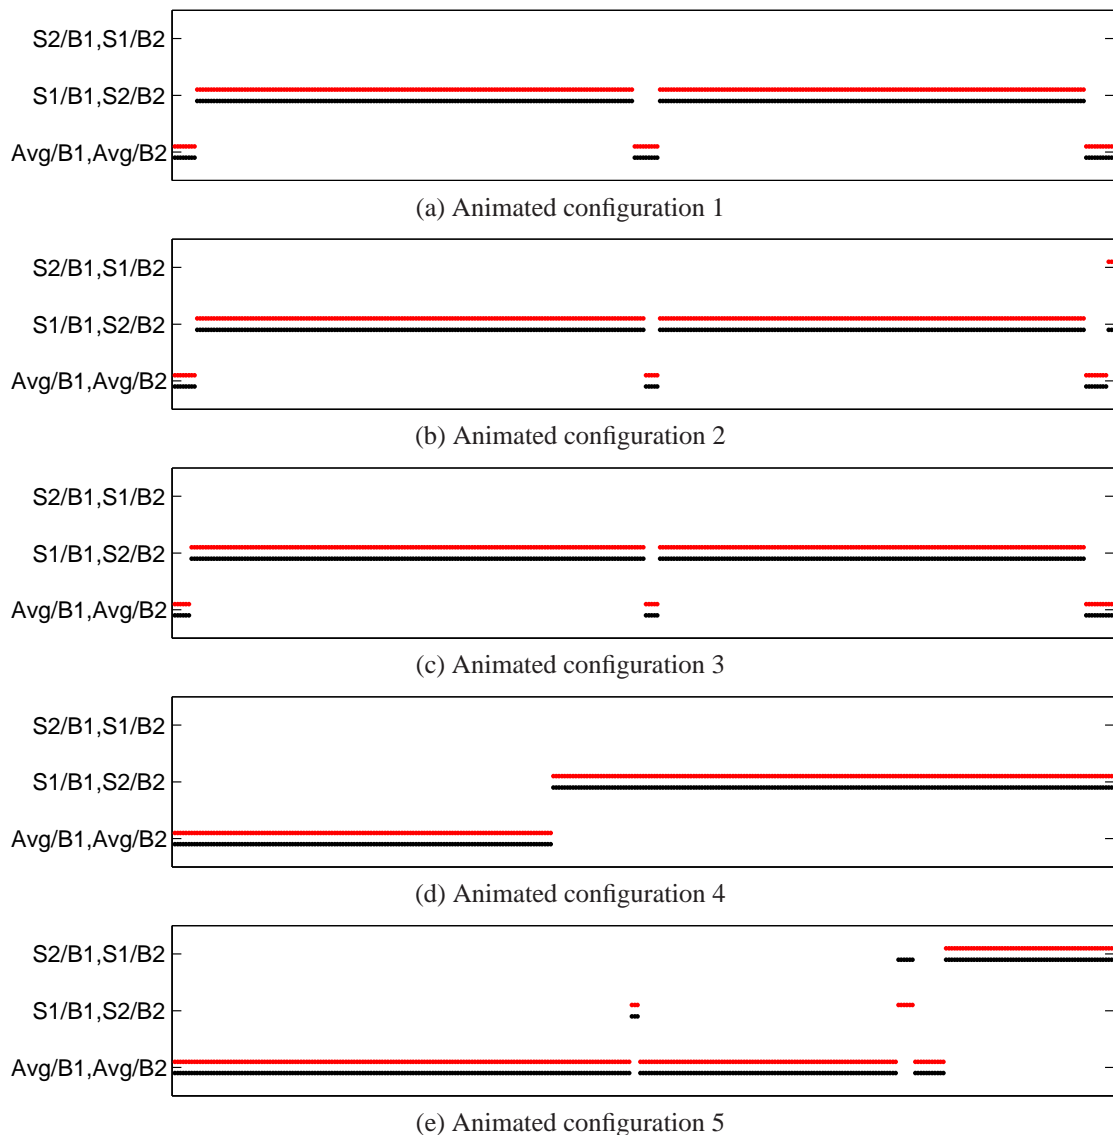


Figure 6.13: Ground-truth derived difference signal based search results for Animated phantom configurations. See Fig. 6.6 for the explanation and Tab. 6.1 for error percentages.




Figure 6.14: Qualitative results of search on Animated configuration 1. Top row shows sagittal slices from a complete reconstruction (a), from a ground truth bin 1(b), from the estimated bin 1(c), from the ground truth bin 2 (d), and from the estimated bin 2 (e). The second row contains close-ups of the regions from the images above.




Figure 6.15: Qualitative results of search on Animated configuration 2. Shown are magnified regions of a sagittal slice from a complete reconstruction (a), from a ground truth bin 1(b), from a the estimated bin 1(c), from the ground truth bin 2 (d), and from the estimated bin 2 (e).

### 6.3.1.8 Animated configuration 3, ground truth difference signal

This configuration contains a single shift of 6 mm occurring half-way through the total scan cycle. The earlier results with the Simple configurations containing small-motions suggests that the error rate should be higher for this configuration than its contemporary configurations. This, in fact, is not what is observed. Like all the cases other than the second case, using the ground truth difference signal, the search results are as perfect as can be achieved with the bucketing. This introduces a second significant factor in determining the effectiveness of the search: motion type. Comparing this motion range with the comparable motion range in the Simple cases suggests that the Lujan motion

creates more difficulty than a shift motion. This is actually very reasonable. The Lujan motion “spreads” the motion information through a range of positions whereas the shift motion concentrates it in only two positions.




Figure 6.16: Qualitative results of search on Animated configuration 3. Shown figures are as per the previous diagrams.

#### 6.3.1.9 Animated configuration 4, ground truth difference signal

This configuration contains a single shift of 16 mm occurring at the 70% point in the total scan cycle (composed of the two sequential scans). It is similar to virtual configuration 1 but rather than having an equal distribution of projections, one bin has over twice as many projections that (should be) assigned to it. This case represents observing a motion only part-way into the second scan. The fourth row of Figure 6.13 shows the success of the search even in this sub-optimal case, given the ground truth difference signal. Qualitatively, the success is seen in Figure 6.17




Figure 6.17: Qualitative results of search on Animated configuration 4.

#### 6.3.1.10 NCAT configuration 1, ground truth difference signal

These next results are obtained using the NCAT phantom class discussed in Chapter 3, Section 3.5. These more realistic phantoms introduce much greater structural complexity consistent with what would be expected in a real clinical setting. Figure 6.18, as with the other phantom types, shows the bucketing performed using Algorithm 6 for the NCAT

phantom configurations used in these tests. See Figure 6.5 for an explanation of the plots. These seem quite reasonable for the ground truth difference signals shown.

The first NCAT configuration contains an averaged size lung tumour in the upper lobe of the right lung. The first row of Figure 6.19 shows the resulting assignment with the dot plot used with each of the phantom configurations. An error is observed with one of the bucket assignments in the middle and with the beginning and ending buckets. This results in 16 paired projections being mislabelled creating a 4.8% error rate. Even at that error rate, the qualitative results appear quite good. Note especially the diaphragm in the first bin for both the ground truth (b) and estimated (c) volumes.

### 6.3.1.11 NCAT configuration 2, ground truth difference signal

The second NCAT configuration contains an averaged size lung tumour in the middle lobe of the right lung. The second row of Figure 6.19 shows the resulting assignment with the dot plot. An error is observed with one of the bucket assignments in the middle and with two of the ending buckets. This results in 26 paired projections being mislabelled creating a 7.8% error rate. Even at that error rate, the qualitative results in Figure 6.21 still appear good. Note especially the diaphragm in the first bin for both the ground truth (b) and

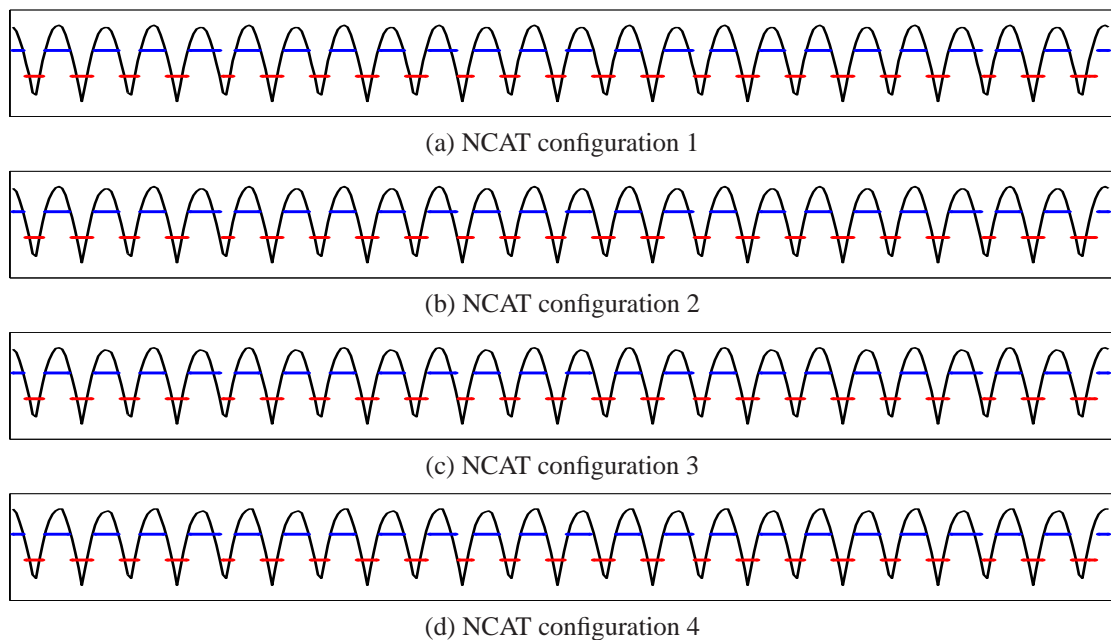


Figure 6.18: Ground-truth derived difference signals and resulting preliminary bucketing for NCAT phantom configurations. The top blue bar shows projections that are grouped into “different” buckets; the bottom red bar shows projections grouped into “same” buckets.

estimated (c) volumes.

### 6.3.1.12 NCAT configuration 3, ground truth difference signal

The third NCAT configuration contains an averaged size lung tumour in the lower lobe of the right lung. The third row of Figure 6.19 shows the resulting assignment with similar bucketing errors observed earlier. This results again in 16 paired projections being mislabelled creating a 4.8% error rate. The qualitative results in Figure 6.22 again appear quite good.

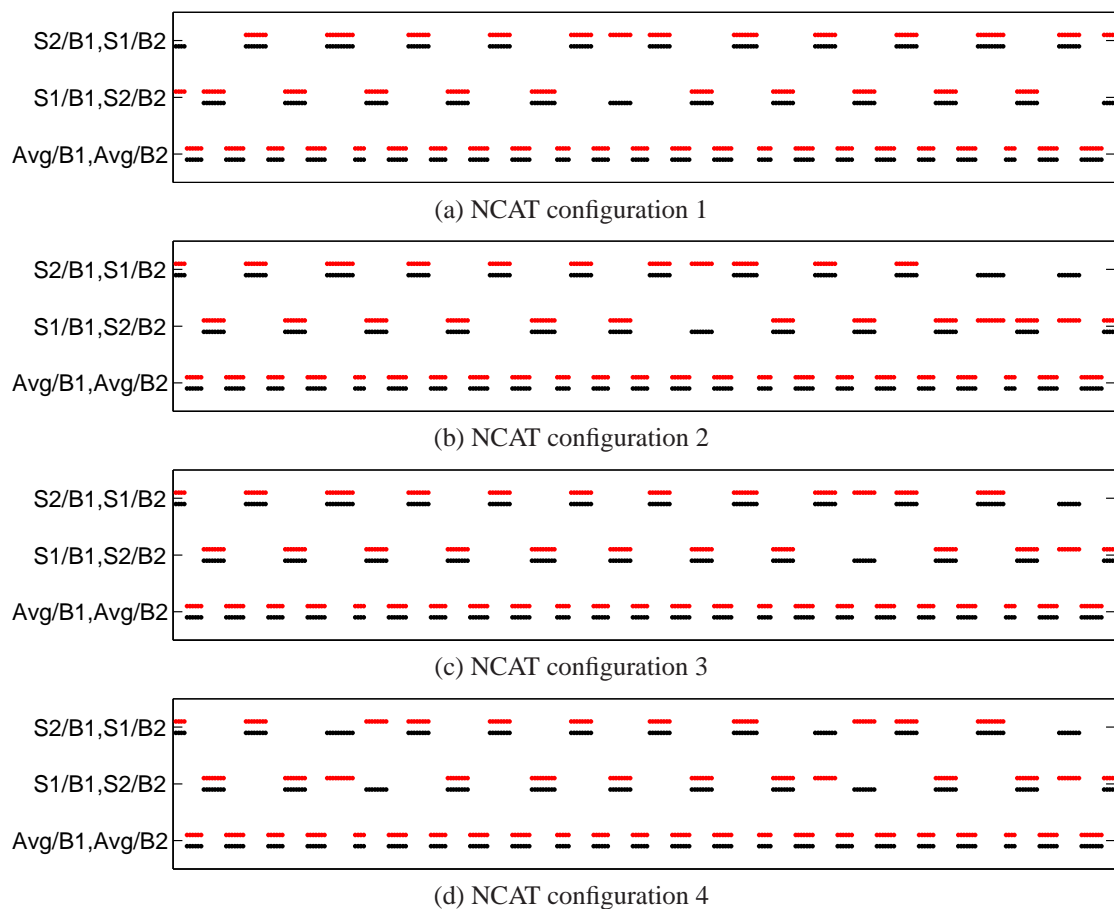


Figure 6.19: Ground-truth derived difference signal based search results for NCAT phantom configurations. The black dots each identify a projection pair and its correct assignment. Red dots identify the resulting assignment found by the search. See Table 6.1 for error percentages.






Figure 6.20: Qualitative results of search on NCAT configuration 1. Top row shows sagittal slices from a complete reconstruction (a), from a ground truth bin 1(b), from the estimate for bin 1(c), from the ground truth bin 2 (d), and from the estimated bin 2 (e). The second row contains close-ups of the regions from the images above.




Figure 6.21: Qualitative results of search on NCAT configuration 2. See Fig. 6.20 for the explanation of the images.

### 6.3.1.13 NCAT configuration 4, ground truth difference signal

The fourth NCAT configuration contains a large size lung tumour in the lower lobe of the right lung. The fourth row of Figure 6.19 shows five bucketing errors resulting in 42 paired projections being mislabelled creating a 12.5% error rate. As with the third case of the Simple phantom set, this relatively large error rate still generates very reasonable




Figure 6.22: Qualitative results of search on NCAT configuration 3. See Fig. 6.20 for the explanation of the images.

qualitative results seen in Figure 6.23.




Figure 6.23: Qualitative results of search on NCAT configuration 4. See Fig. 6.20 for the explanation of the images.

### 6.3.2 Using an estimated difference signal

In these experiments, the difference signal is estimated from the differences between the 2D ROI clipped paired projections. This separates the effect of under performing due to the difference signal processing from the results of the search algorithm proposed.

Phantom class	1	2	3	4	5
Simple:	0.3	7.8	11.3	0.9	0.0
Animated:	0.0	0.6	0.0	0.0	0.0
NCAT:	4.8	7.8	4.8	12.5	–

Table 6.1: Summary of percentage of errors for phantoms using a ground-truth derived difference signal for partitioning.

In Figure 6.24, the bucketing performed using Algorithm 6 is shown for the Simple phantom configurations used in these tests. The signals and partitioning should be compared with those obtained from a ground truth difference signal shown in Figure 6.5. The estimated difference is shown as a function of the sequence of paired projections with a black line. The blue and red bars are as described in Figure 6.5. The preliminary same-different bucketing using the ground-truth position information appears to work acceptably well in the Simple case for the periodic Lujan type of motion. The shift motion shows an increased error.

Figure 6.25 shows the results of the search assignments in dot plot form. The red dots are, as in previous plots, the “found” bin assignment and the black dots are the ground truth derived assignment. Comparing the first four rows, the efficacy of the partitioning algorithm seems good; the bucketing is nearly identical. In the shift motion case, row 5, significant errors begin to appear suggesting the particular algorithm is not a good fit with shift style motion.

### 6.3.2.1 Simple virtual configurations, estimated difference signal

The first row of Figure 6.25 shows 7 paired projections being mislabelled creating a 2.1% error rate for Simple configuration 1. No noticeable differences from the ground truth can be observed in the sagittal slices shown in Figure 6.26.

The second row of Figure 6.25 shows 30 paired projections being mislabelled creating a 9.0% error rate for Simple configuration 2. No significant differences from the ground truth can be observed in Figure 6.27.

The third row of Figure 6.25 shows 39 paired projections being mislabelled creating a 11.6% error rate for Simple configuration 3. Differences from the ground truth in Figure 6.28 are very slight.

The fourth row of Figure 6.25 shows 4 paired projections being mislabelled creating a 1.2% error rate. Again, no noticeable differences from the ground truth can be observed in Figure 6.29.

The Fifth row of Figure 6.25 shows 73 paired projections being mislabelled creating

a 21.8% error rate. Differences from the ground truth are quite noticeable and can be observed in Figure 6.30. Even with this error, the tumour shift and shape can still be seen.




Figure 6.26: Qualitative results of search on Simple configuration 1 using an estimated difference signal.

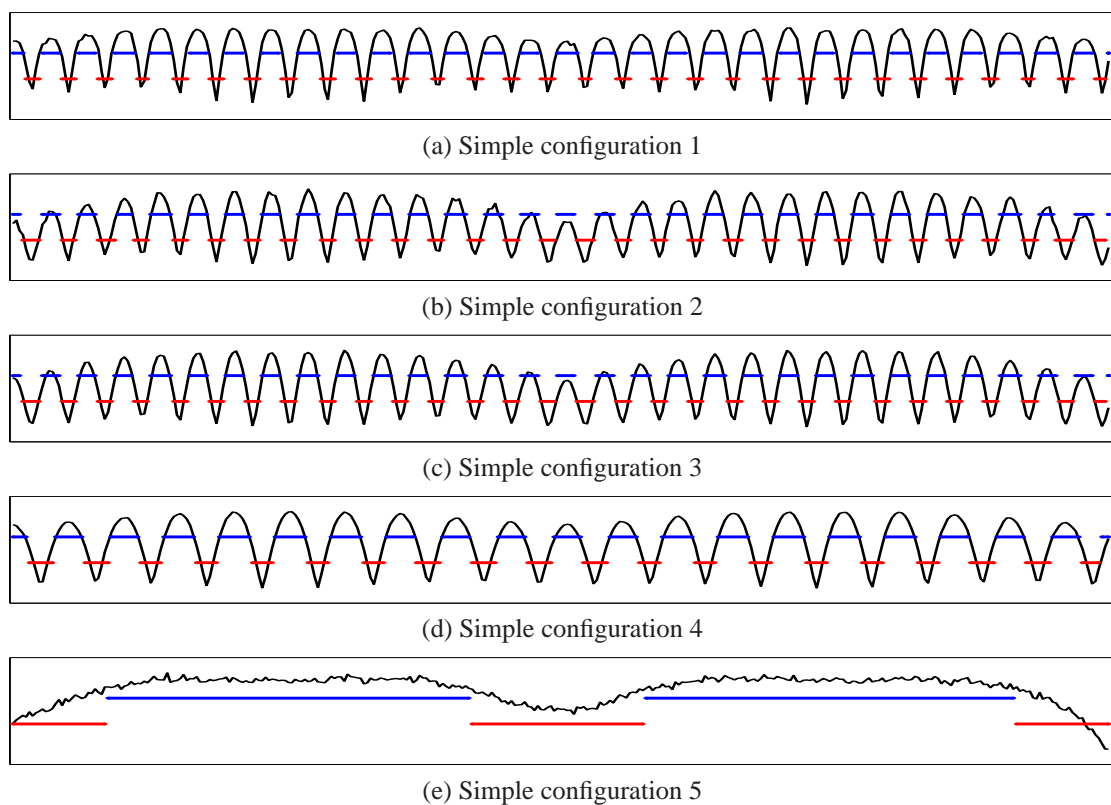


Figure 6.24: Estimated difference signals and resulting bucketing for Simple phantom configurations. The top blue bar shows projections that are grouped into “different” buckets; the bottom red bar shows projections grouped into “same” buckets.




Figure 6.27: Qualitative results of search on Simple configuration 2. See Figure 6.7 for explanation.

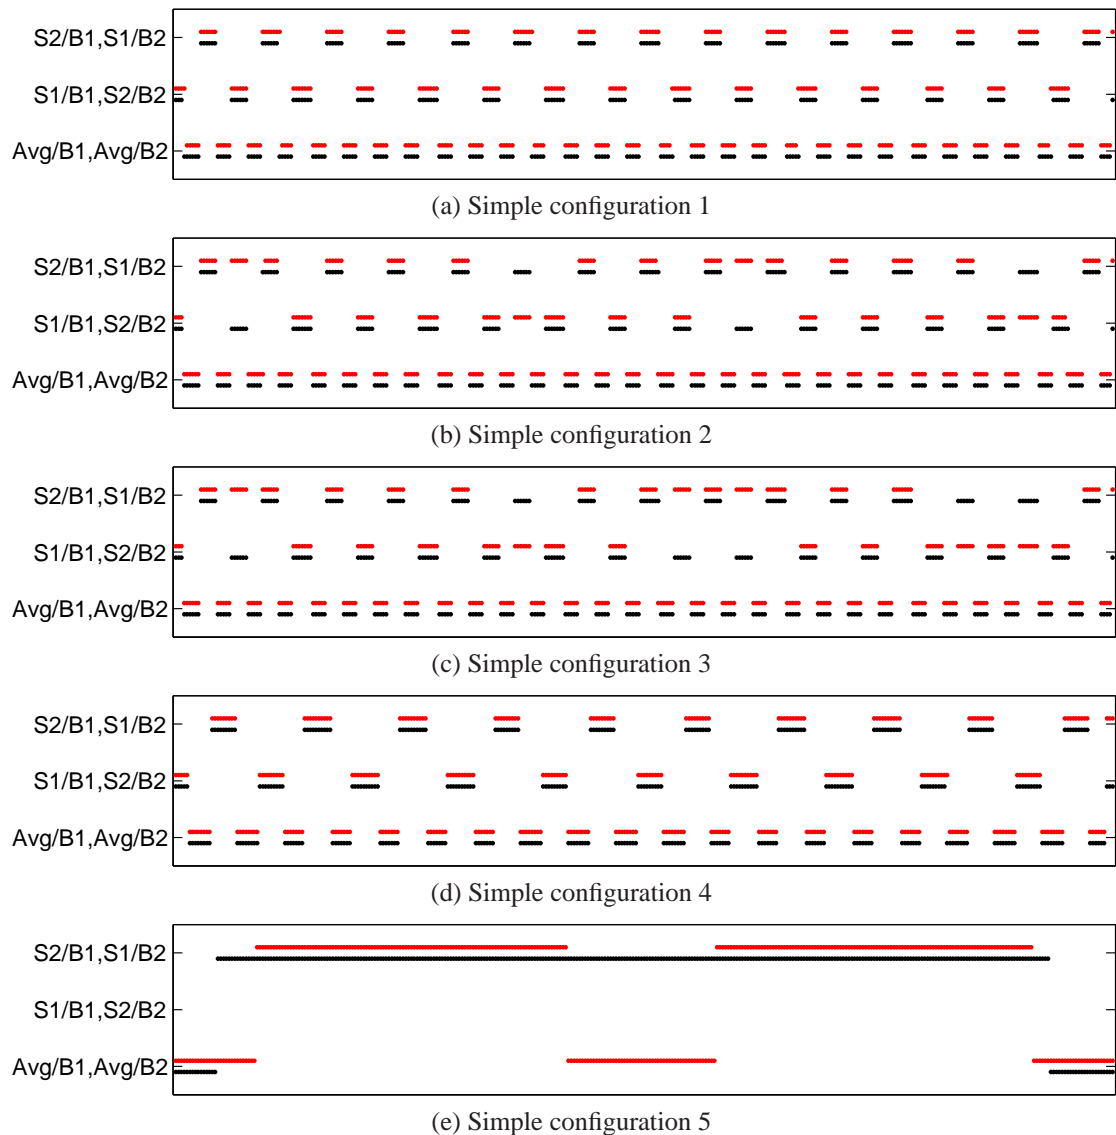


Figure 6.25: Estimated difference signal based search results for Simple phantom configurations. The black dots each identify a projection pair and its correct assignment. Red dots identify the resulting assignment found by the search. See Table 6.2 for error percentages.




Figure 6.28: Qualitative results of search on Simple configuration 3. See Figure 6.7 for explanation.




Figure 6.29: Qualitative results of search on Simple configuration 4. See Figure 6.26 for explanation.




Figure 6.30: Qualitative results of search on Simple configuration 5. See Figure 6.26 for explanation.

### 6.3.2.2 Animated configurations, estimated difference signal

In Figure 6.31, the bucketing performed using Algorithm 6 is shown for the Animated phantom configurations used in these tests. The signals and partitioning should be compared with those obtained from a ground truth difference signal shown in Figure 6.12. It is evident from the resulting bucketing that the partitioning algorithm is insufficient for correctly partitioning this kind of motion.

Figure 6.32 shows the results of the search assignments in dot plot form. Along with




Figure 6.31: Estimated difference signals and resulting bucketing for Animated phantom configurations.

the previous figure, these results reinforce the idea that the partitioning algorithm performs poorly on this kind of motion.

The first result of Figure 6.32 shows 195 paired projections being mislabelled creating a 58.2% error rate which is easily visible in Figure 6.33. At this error rate, the results are becoming unusable. Here, as in most of the remaining NCAT cases, the qualitative results will continue to support the assertion that a different algorithm is required for partitioning shift motion difference signals.

The second row of Figure 6.32 shows 148 paired projections being mislabelled creating a 44.2% error rate which is easily visible in Figure 6.34. At this error rate, the results are becoming unusable.

The third row of Figure 6.32 shows 137 paired projections being mislabelled creating a 40.9% error rate. Artefacts are clearly visible in Figure 6.41 but the overall shape and position of the tumour is still apparent.

The fourth row of Figure 6.32 shows 77 paired projections being mislabelled creating a 23.0% error rate. Performance relative to the ground truth in Figure 6.41 is still acceptable.

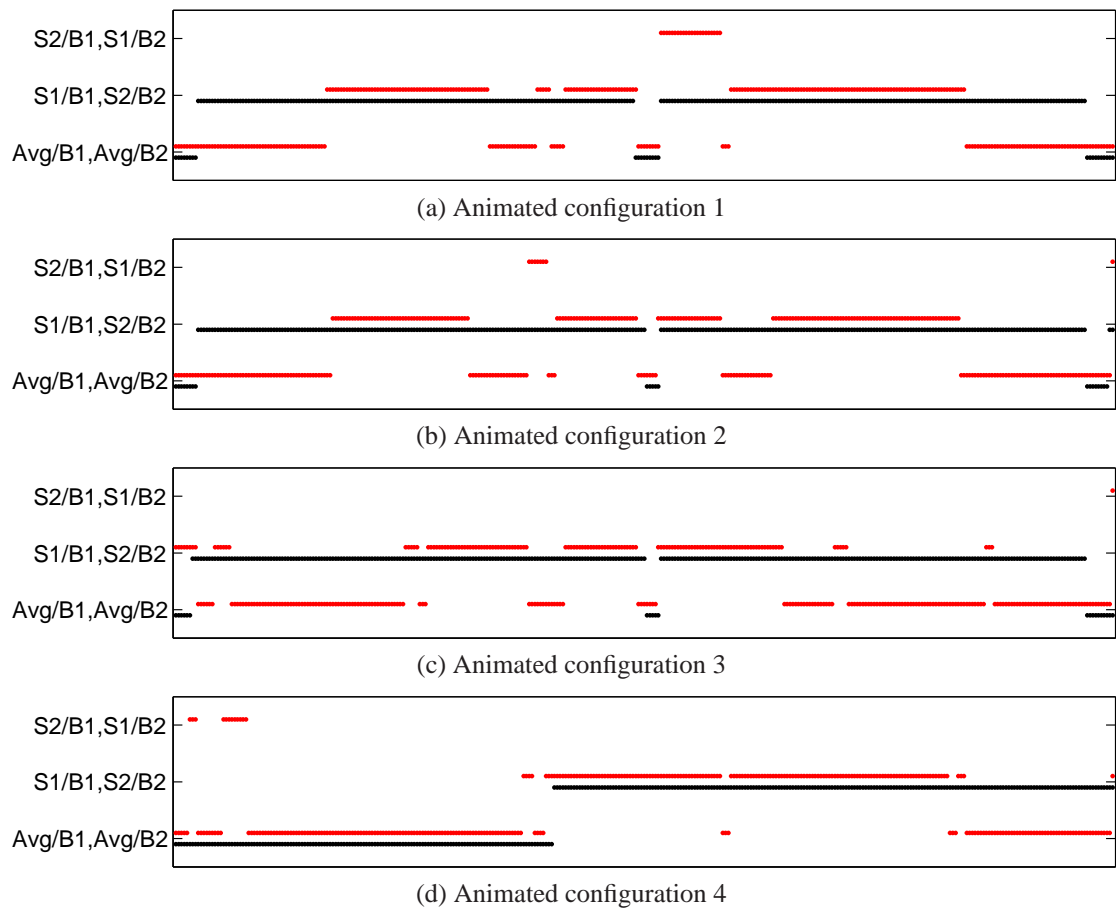


Figure 6.32: Estimated difference signal based search results for Animated phantom configurations.




Figure 6.33: Qualitative results of search on Animated configuration 1.






Figure 6.34: Qualitative results of search on Animated configuration 2.




Figure 6.35: Qualitative results of search on Animated configuration 3.




Figure 6.36: Qualitative results of search on Animated configuration 4.

### 6.3.2.3 NCAT configurations, estimated difference signal

In Figure 6.37, the bucketing performed using Algorithm 6 is shown for the NCAT phantom configurations used in these tests. The signals and partitioning should be compared with those obtained from a ground truth difference signal shown in Figure 6.18. The estimated difference is shown as a function of the sequence of paired projections with a black line. The blue and red bars are as described in Figure 6.5. The preliminary same-different bucketing using the ground-truth position information looks reasonable and the experiments will confirm that it is.

Figure 6.38 shows the results of the search assignments in dot plot form. Comparing the rows, the efficacy of the partitioning algorithm seems good; the bucketing is nearly

identical though the sizes of the partitions show differences. In the first set a missing partition can be observed in the middle of the set in Figure 6.37 and the effect of this can be seen in the middle of the assignments on the first row in Figure 6.38.

The missing partition and other search errors result in 44 mislabelled paired projections for a cumulative error of 13.1%. Nevertheless, it is difficult to observe the error in the qualitative result shown in Figure 6.39.

The second row of Figure 6.38 shows 53 paired projections being mislabelled creating a 15.8% error rate. Slightly more blurring relative to the ground truth can be observed for the second bin estimated image (e) of Figure 6.40.

The third row of Figure 6.38 shows 38 paired projections being mislabelled creating a 11.3% error rate. No significant differences from the ground truth can be observed in Figure 6.41.

The fourth row of Figure 6.38 shows 68 paired projections being mislabelled creating a 20.3% error rate. Demarcation of the two main positions of tumour is still very clear though some minor additional warping and blurring can now be observed in Figure 6.42.

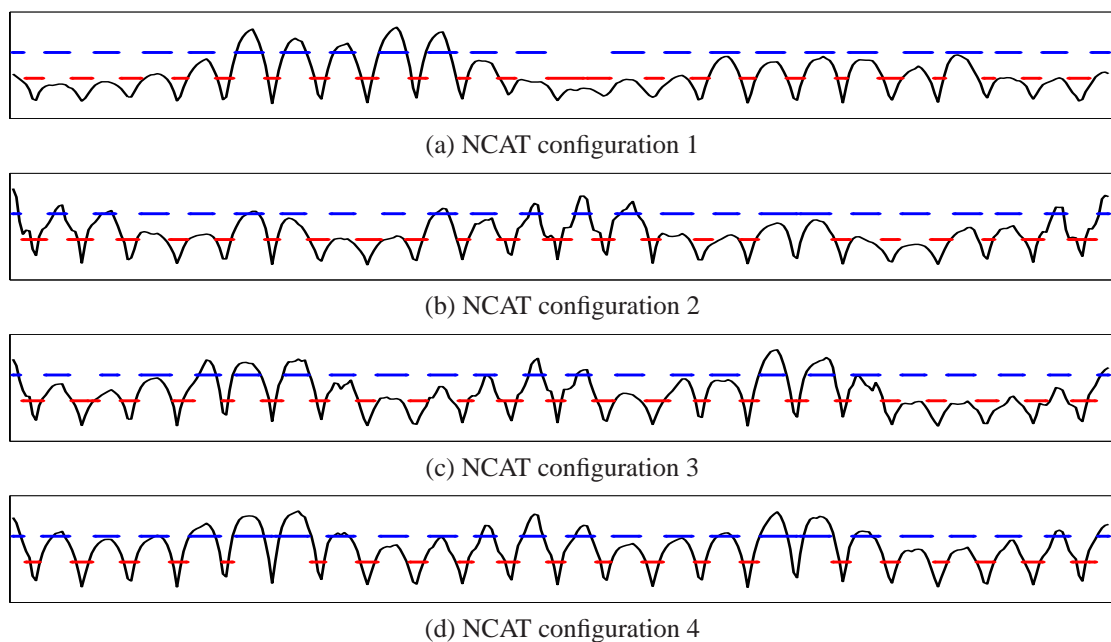


Figure 6.37: Estimated difference signals and resulting bucketing for NCAT phantom configurations.




Figure 6.39: 2-scan Qualitative results of search on NCAT configuration 1. Sagittal slices are presented from a complete reconstruction (a), from the ground truth for bin 1(b), from an estimated bin 1(c), from the ground truth for bin 1 (d) and from the estimate for bin 2 (e).

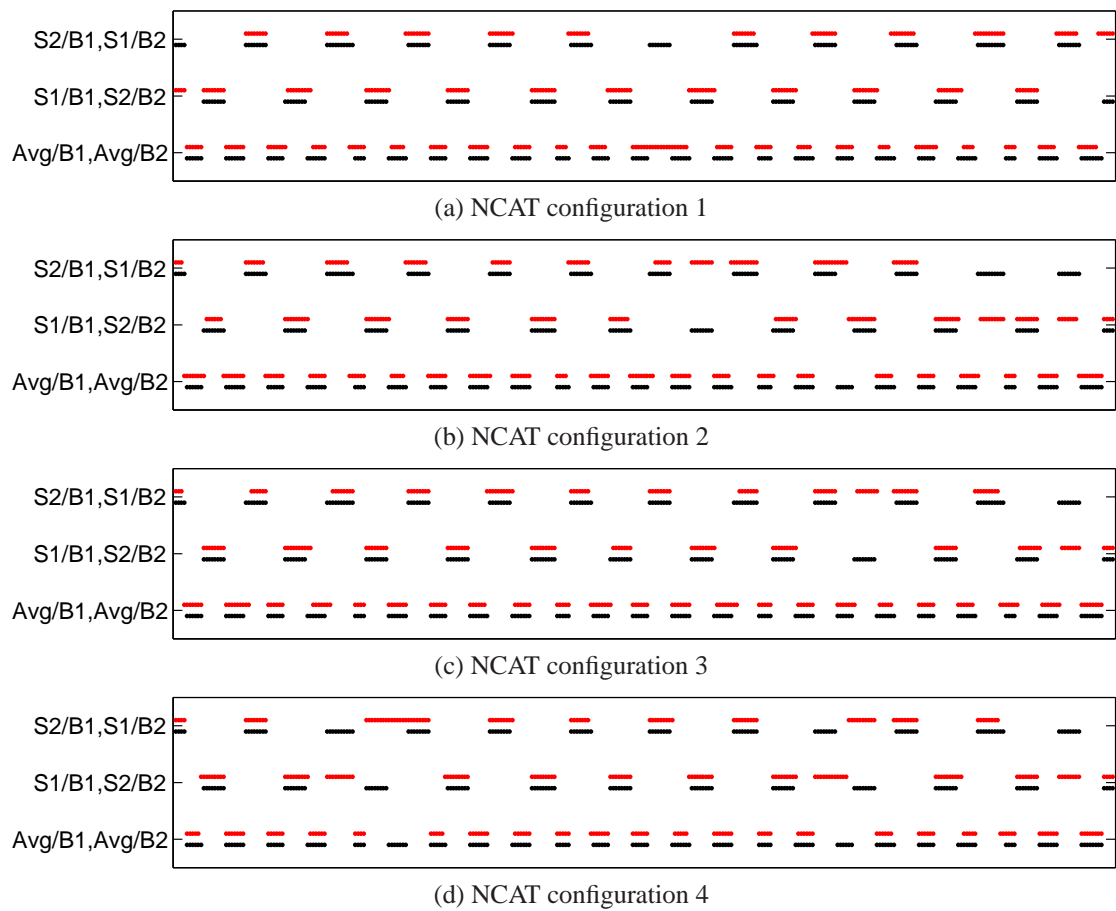


Figure 6.38: Estimated difference signal based search results for NCAT phantom configurations.




Figure 6.40: 2-scan Qualitative results of search on NCAT configuration 2.




Figure 6.41: 2-scan Qualitative results of search on NCAT configuration 3.




Figure 6.42: 2-scan Qualitative results of search on NCAT configuration 4.

Phantom class	1	2	3	4	5
Simple:	2.1	9.0	11.6	1.2	21.8
Animated:	58.2	44.2	40.9	23.0	
NCAT:	13.1	15.8	11.3	20.3	–

Table 6.2: Summary of percentage of errors for phantoms using an estimated difference signal for partitioning.

## 6.4 Conclusions

The goal of this chapter was to ascertain if the proposed two-scan approach to searching is a viable method for finding bin assignments. Two-scan binning requires two related steps:

obtaining and partitioning what has been referred to as a “difference signal”, and using this to exhaustively search the space of possibilities. These were both described in some detail in the text and various properties necessary for the PCA dimensional reduction were shown. The search space and an objective function for evaluating solutions were defined.

To better isolate the impact of the search ideas, a form of ground truth was used to generate the difference signal and a bucketing. The search was then tested on a number of different phantom configurations from three different classes described in Chapter 3. The results consistently supported the search approach and showed it to be generally successful and robust when the paired projections assignment accuracy was between 80 and 100%. Many more tests would be required to establish a curve to better estimate an error rate threshold.

The experiments were then performed again but the ground truth bucketing process was replaced with a simple automated bucketing approach. This approach proved to be reasonably good with *continuously periodic* type motion incurring an average error rate of 6.9% with the Simple phantom class experiments and an average error rate of 15.1% for the more realistic NCAT phantoms. The approach failed for *shift* type motion using the initial attempt at a partitioning algorithm, but the clear success of the search algorithm (nearly perfect search in all cases) with accurate partitioning shows the problem to be with the partitioning algorithm, and not with the search algorithm.

In both sets of experiments, with ground truth derived difference signals and calculated difference signals, another factor that increased errors was found to be small motions. Without sufficient motion information captured by scans, the search algorithm had difficulty finding accurate assignments. Nevertheless, the assignments it found were still acceptable. Fortunately, as the motion decreases to the point where the search algorithm fails, it also ceases to be necessary for the algorithm to succeed.

A remaining problem with this approach is the necessity of instituting a change in well established clinical protocols. This also precludes the use of this technique on existing data sets. A method which can take the core ideas from the search method and apply it to the single-scan protocol is desired and forms the basis for the next chapter.

# Chapter 7

## Searching for CBCT Projection Bin Assignments Using a One Scan Protocol

---

### 7.1 Overview

In Chapter 6, a method for finding an assignment of projections to two bins was presented and evaluated in a set of experiments on three classes of phantoms. The principal advantage of this method was to extract information whereby a (nearly) exhaustive search could be performed on the search space. The principal disadvantage of the proposed method was that it required implementing a new two-scan imaging protocol. While the protocol change is not unreasonable, *any* protocol change in a clinical setting requires strong evidence to support the benefit of such a change and it requires time consuming and expensive studies to validate the impact of the change. A method using the existing protocol circumvents much of this effort. It also has the advantage of being able to be retrospectively applied to any retained CBCT data sets because the method is not coupled with either an existing 4D planning CT or any respiratory measurement data.

In Chapter 6, paired projections were used to construct a difference signal which formed the basis for automatically finding a partition of the set into buckets and for finding a set of preliminary labels for those buckets containing “different” or “similar” paired projections. Without the two-scan protocol, both of these capabilities are lost and the search space resumes an intractable size for exhaustive searching. Consequently, while

this chapter retains the objective function described in Chapter 6, it proposes a different method for searching the bucket assignment space using only the traditionally acquired CBCT projections.

## 7.2 Method

The one-scan search method can be summarized as grouping projections into uniformly sized buckets and filling the search space with average projections derived from the initial reconstructed volume using a forward projection method. The details of these steps are contained in this section.

### 7.2.1 The search space

As was discussed in Section 6.2, the correct search space for a two-bin problem with 670 projections contains  $2^{669}$  possibilities. In Chapter 6, the idea of bucketing was proposed as a reasonable approximation of a correct projection assignment by assigning contiguous projections to the same bin. The boundaries of these buckets was determined by examining the signal constructed from the absolute differences of the 2D regions of interest (obtained by projecting a “shadow” of the 3D ROI which, in clinical settings, would be the planning target volume) of the paired projections. This signal is not available without having the second scan so an alternate form of bucketing is proposed in this chapter.

Taking the respiratory motion as the “fastest” motion to be encountered - i.e. excluding the cardiac motion from consideration - the standard range for a breathing period in the literature is between four and six seconds. If a bucket size is chosen arbitrarily to cover the number of projections acquired in half of a second, then a given bucket will contain approximately  $(670 \text{ projections} / 120 \text{ seconds})(1 \text{ second} / 2 \text{ buckets})$  or 2.8 projections per bucket. If one makes a simplistic assumption that projections will be evenly split between bins (i.e. that for  $N$  projections,  $N/2$  will be assigned to each of the two bins), and if the partition lines up ideally in the four-second case, four partition’s worth of projections will belong to one bin and four will belong to the other. The worst case scenario is one where the partitions are offset in such a way that half of the projections in a bucket belong to one bin and the other half, within the same bucket, belong to the opposite bin. However, even in this case, six of the buckets can still be correctly assigned and the other two will effectively become averaged buckets and their effect will be uniformly distributed which has been established in Chapter 4 as a useful artefact reduction method.

Given this new bucketing approach, the size of the search space for a 120 second scan

(the standard scan duration in clinical practice at the time of this thesis), is approximately  $2^{240}$ . This it is still intractable for an exhaustive search. As was described in Section 6.2.1 and shown in Figure 6.1, a single bucket assignment change can create a significant score change causing normal gradient ascent methods to get trapped in local maxima. A common response to this problem is to “hill climb” to the local maxima then restart the search randomly in a different place. This is known as “random restart hill climbing.” Another approach, referred to as “local beam search” by Russell and Norvig [73], starts with several randomly generated states then generates all successors from these states and retains the best  $k$  successors from all the parent states as the new parent states. In the specific search space under discussion, a given state is the current bucket assignment and all successor states are those resulting from switching the assignment of an individual bucket. Both the random restart hill climbing and the local beam search algorithms were implemented with similar overall results. The simpler random restart hill climbing is reported in the results section.

## 7.2.2 Forward projections

A fundamental theme of this thesis is the importance of filling projection gaps with some form of angle and patient specific averaged projection. In the two-scan case, this was easily accomplished by averaging the paired projections. In the one-scan case, a new approach is taken. Given the property of filtered backprojection reconstructions being an average of the projections as described in Chapter 2, the ideal way to generate an average projection is to reproject the needed projections from this volume. This method is called “forward projection” and is a necessary step in iterative algebraic reconstruction techniques [30, 32, 43].

Turbell, in his thesis [91], compares four different forward projection techniques: Sidon’s method, Joseph’s method, Kohler’s method, and a Simple method. His conclusions suggest the use of Joseph’s method because it performs approximately the same as the Kohler method but is less complex. Matej et al. propose a Fourier based forward projection method [52] while Long et al. [43] suggest using something they call a separable footprint projector. These are each interesting but following Einstein’s maxim that, “everything should be made as simple as possible, but no simpler” the simple method was chosen to generate projections from the reconstructed volume to create the average projections needed for the search procedure in this chapter.

Simple forward projection constructs a ray from the source to the virtual detector pixel and then samples the values in the discrete volume along that ray using tri-linear



interpolation. The step width of the sampling is the principal parameter to be adjusted. In these experiments, the sampling width was set to one quarter of the width of a voxel (voxels are uniform and cubic in the data sets).

### 7.2.3 Gap filling during search

The three different gap filling methods described in Chapter 5 are used in this chapter as well. In that chapter, they were used *after* a projection assignment had been determined through external means as would be the case with a spirometer or some other gating mechanism. In this chapter, it is used to drive a search algorithm. The intuition is that it will smooth the scoring results by reducing the SSD errors induced by gap artefacts. This hypothesis is tested in this chapter by using the “average” projections, as described in the previous section, to fill the gaps. As a quick review of the material in Chapter 5, the three filling mechanisms used are:

**Unfilled** : reconstruct using only the original projections specified in the candidate projection assignment.

**Balanced fill** : reconstruct using the original projections specified in the candidate projection assignment and fill in the gaps uniformly using up to but not more than the number of projections specified. For instance, if a binning assignment specifies 200 buckets (out of 240) should be in bin 1, then the maximum number of fill projections will be 40, the remaining set. If, however, the binning assignment specifies only 50 buckets for a bin, then 50 buckets worth of fill projections will be added uniformly in the gap spaces in the projection set.

**Complete fill** : reconstruct using the original projections specified in the candidate projection assignment and fill in any missing projections from the reprojected “average” set.

From the results in Chapter 5, the expectation is that searching without filling will perform more poorly than searching with filling and that the balanced filling method will perform the best.

### 7.2.4 Experiments

The phantom configurations for the three classes of phantoms, as defined in Chapter 3, are used in the experiments in this chapter. For each configuration, the ideal two-bin assignment is searched for using the unfilled, balanced fill and complete fill methods

described immediately above. The error is calculated as a function of correct buckets assigned in the found bucket assignment as compared with a ground truth derived bucket assignment. The ground truth, as elsewhere in this thesis, is constructed using a k-means clustering on the tumour positions recorded during the synthesis of each phantom. Errors are reported as a percentage of misassigned buckets.

### 7.3 Evaluation of results

As in Chapter 6, the phantoms and specific configurations described in Chapter 3 are used to test the proposed search method. Each configuration is searched using the three different filling techniques. The results for each class of phantom are reported before a final summary and conclusion is presented.

As in the previous chapter, a dot plot comparing the ground truth bucket assignment with the found bucket assignment is shown. For each phantom, three such plots can be shown, one for each type of fill. Because of the consistency of the results, only one or two representative configurations show all three results graphically. The complete filling approach, which ends up being the search variant which generates the most accurate projection assignments, is shown for all configurations. The error rates for all results are presented in tables for each phantom class.

Sagittal slices, as in Chapter 6, are shown for each result to visually present in a qualitative form the achievements of the search variations. Again because of the consistency between the results, only a representative (or especially interesting) configuration from each class has slices from all three search methods.

What is immediately seen is that the complete fill method generates the best solutions. Unfortunately, a reconstructed volume containing fill projections where gaps exist blurs the volumetric image in the same way that an unbinned volume is blurred. Observing the utility of the ground truth volumes which are constructed from an assignment created by the same oracle used throughout the thesis, the realization is that one need not reconstruct the final *visualized* binned volumes using the fill projections, one only needs them for the *search* process. For this reason, the bin volumes that would be reconstructed for clinical visualization are constructed using the projection assignment found with the complete fill variation, but without performing any filling while reconstructing the volumes. These are presented next to the ground truth volumes for each of the phantoms.

### 7.3.1 Simple phantom results

The Simple phantom configurations all test mostly continuously periodic kinds of motion but include a shift motion for case 5. See Chapter 3, Section 3.3 for detailed descriptions of the configurations.

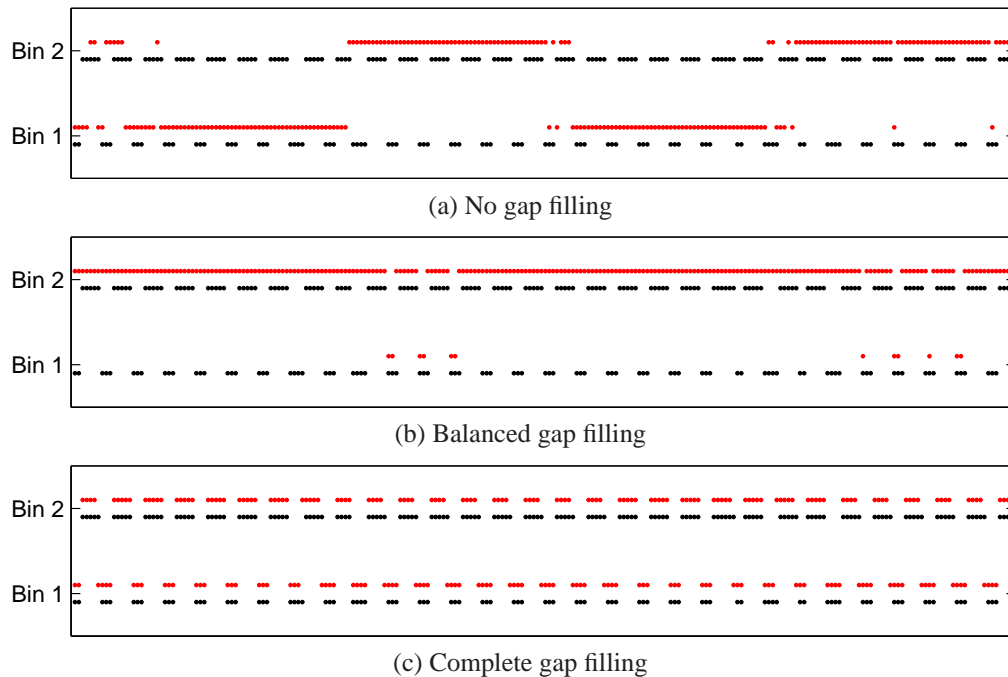


Figure 7.1: Dot plot of bucket assignments for Simple configuration 1. Each of the three fill types are shown and for each result, the black dots show the ground truth bucket assignment while the red dots show the bucket assignment resulting from the search.

In Figure 7.1, the results show that using the complete gap filling method generates very good results. In fact, only 19 buckets have been mislabelled which is about an 8% error rate. On the other hand both the balanced gap filling method and not filling gaps at all had significant errors of 33% and 49% respectively.

The assignment shown in (a) seems to be driven by some global sinusoidal function. If one considers the principal shape of the object being imaged (a “squashed” cylinder with unequal axes) and the circular path of the imaging geometry, then the attenuation will be stronger when the (virtual) X-rays cross the wider part of the cylinder and weaker when they traverse the smaller axis direction. This effect happens twice as the gantry rotates about the object. One possible hypothesis for the result is that the search is trying to bin “bright” projections in one bin and “dark” projections in the other bin to maximize the SSD term in the objective function. This was anticipated in the previous chapter, Section 6.2 which motivated the inclusion of the mean probability term. Examining a plot

of the mean of (a central section of) the projections in this configuration along with the assignment, Figure 7.2, shows something interesting. The projections have been grouped roughly into four macro-buckets and each macro-bucket contains a balance of “dark” and “light” projections which generates an assignment whose reconstructed bin volumes have a mean close to the global mean value. This is the “rewarded” outcome in the objective function. These macro-buckets also create large gap artefacts thus maximizing the SSD component as well which is the other rewarded outcome in the objective function. Results for the other configurations show the same effect for this search variation.

The assignment shown in (b) for the balanced fill approach also reveals a problem with the balancing function as it has been defined above (7.2.3). This method is binning the vast majority of the projections into one bin leaving only a very few projections, along with a very few fill projections, to form the second bin. The volume reconstructed from the first bin will obviously have a mean value similar to the global mean as only a small number of “samples” have been removed from it. The second volume is reconstructed from the ideal small set of projections such that the mean of the volume is again close to the global mean, by virtue of the careful choice of projections, but the gap artefacts will be very severe again leading to a large SSD value.





Figure 7.2: Plot of mean of projections along with the assignment found by the no-fill search process.



(a) Complete gap filling

Figure 7.3: Dot plot of bucket assignments for Simple configuration 2. Only the best result from the complete fill type is shown. Black dots show the ground truth bucket assignment while the red dots show the bucket assignment resulting from the search.

Figures 7.3, 7.4, and 7.5 show the same good results for configurations 2-4 as was seen in Figure 7.1 (c) for configuration 1. They also manifest the search pathologies for the




Figure 7.4: Dot plot of bucket assignments for Simple configuration 3. Only the best result from the complete fill type is shown.




Figure 7.5: Dot plot of bucket assignments for Simple configuration 4.

unfilled and balanced fill methods discussed with Simple configuration 1. The plots for those cases are not shown here but the error rates are recorded in Table 7.1. In all the complete fill cases, the error rates are under 8%. Each of these cases share an important characteristic: the motion modelled is continuously periodic. In Chapter 6, the two-scan search protocol performed reasonably well with this type of motion but under performed with shift type motion when the ground truth difference signal was not available.

Figure 7.6 shows that, for the Simple phantom class at least, shift motion is handled by the complete fill method with an error rate less than 1%, the best error rate of the class. Interestingly, given the completely different type of motion, the effects described before for Simple configuration 1 shown in 7.1 (b) and (c), are also present with this configuration.

Figure 7.8 shows a complete set of sagittal slices for the first Simple configuration results. These slices visually confirm the results shown in the dot plots of the assignments above. In particular, comparing the ground truth images on the top row with the reconstructions in the bottom row, one sees very good agreement though the bin represented in the left images has more blur. These volumes are generated without filling in the gaps but using the best assignment found using the complete fill method. In Chapter 5, it was discovered that such reconstructions, while containing the gap artefacts that can be problematic to search and registration algorithms, are generally better from a clinical perspective because the objects of main interest will have less blur. This can be seen by comparing the images in the bottom two rows. In both cases, the same assignment has

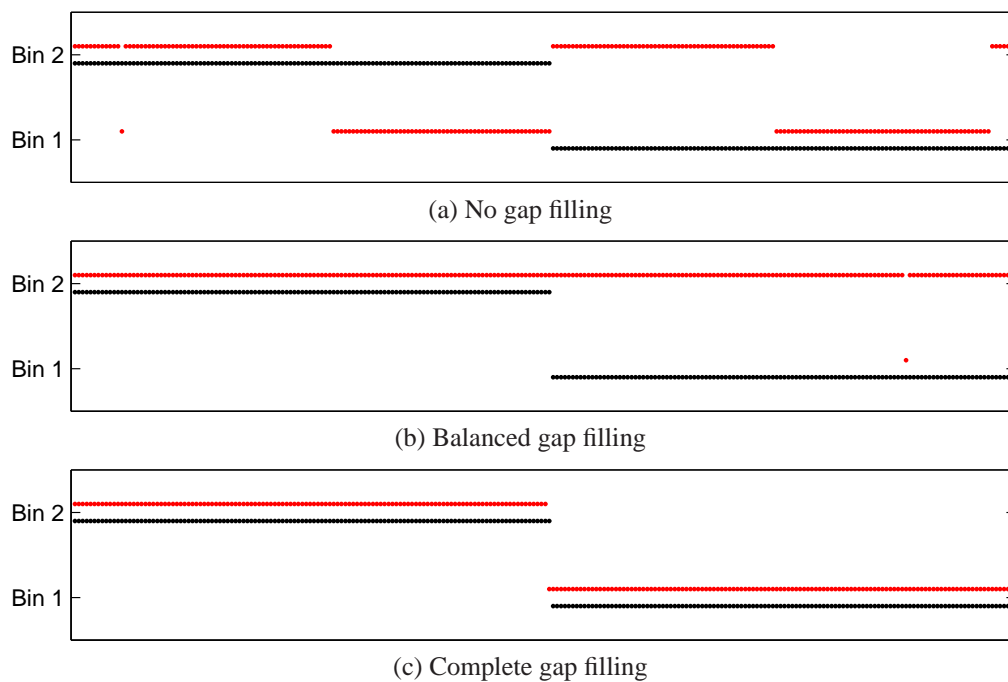


Figure 7.6: Dot plot of bucket assignments for Simple configuration 5. All three fill types are shown because this represents a new type of motion, the shift motion.

been used to drive the binned volume reconstructions, but in row 4, the gaps have been filled while in row 5, they have not been filled.

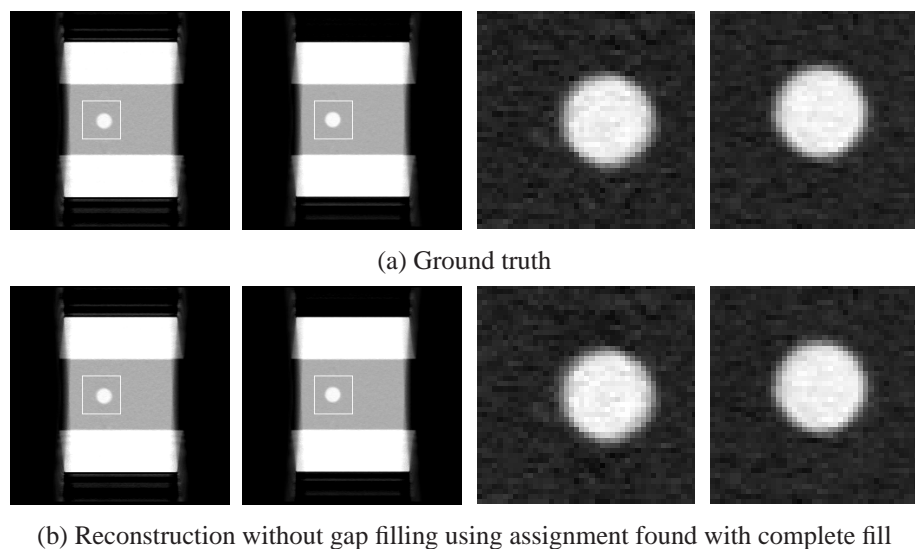


Figure 7.7: Representative sagittal slices from the results for Simple configuration 2. See Figure 7.8 for a description of the columns. The rows here correspond to the top and bottom rows in that figure.

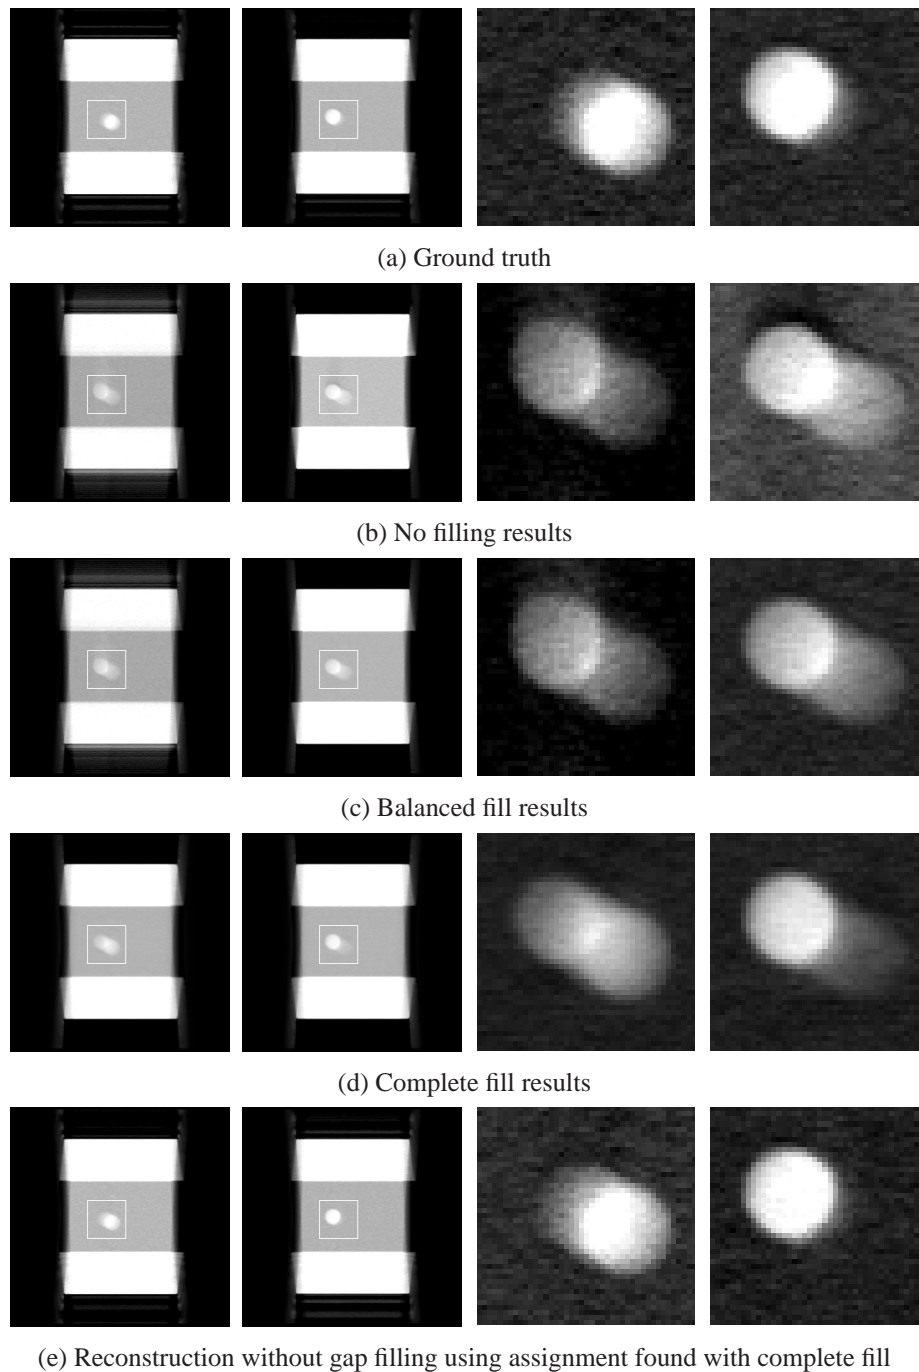




Figure 7.8: Representative sagittal slices from the results for Simple configuration 1. Each row contains, from left to right, the full sagittal slice for bins 1 and 2 then the magnified sub-slices (identified with the light square). The top row contains the ground truth obtained from the recorded positions of the moving object of interest. The second row contains the results of searching without any filling process. The third row contains the results of attempting to use a balanced filling method. The fourth row contains the results from using the complete fill method. The final row shows the results of using the assignment found using the complete fill method, but reconstructing the binned volumes without the fill projections.




(a) Ground truth




(b) Reconstruction without gap filling using assignment found with complete fill

Figure 7.9: Representative sagittal slices from the results for Simple configuration 3. See Figure 7.8 for a description of the columns.



(a) Ground truth



(b) Reconstruction without gap filling using assignment found with complete fill

Figure 7.10: Sagittal slices from the volumes generated by the results from the Simple configuration 4 experiment.

Figures 7.8, 7.7, 7.9 and 7.10 qualitatively show the results for the one-scan search on Simple configurations 2, 3, and 4. These show the search can be successful on continuously periodic motion for both large and small tumour sizes and for large and small motion ranges.

Figure 7.11 shows the interesting set of sagittal slices for Simple configuration 5 which models the shift motion of the object 50% of the way through the scan. This class of motion caused the two-scan search method in Chapter 6 to perform poorly, but here it



performs well with an error rate under 1%. In this figure, the top row showing the ground truth, the bottom row showing the best reconstruction representation, and a middle row showing one of the under performing methods (balanced fill) are shown. Visually, the reconstruction using the results from the complete fill are clearly superior to the results from the balanced fill search method.




Figure 7.11: Sagittal slices from the volumes generated by the results from the Simple configuration 5 experiment. See Figure 7.8 for a description of the sub-figures.

Table 7.1 gives the tabulated results as percentage errors of the three different search variations for the Simple phantom set. The advantage of the complete fill approach shown in the fourth column is clear from this table. The highest error for that gap filling technique is 7.9% whereas the lowest error percentage for the other two methods is 32.9%.

Phantom configuration	No fill	Balanced	Complete
Simple 1	48.8	32.9	7.9
Simple 2	46.7	37.5	7.9
Simple 3	47.9	38.3	6.7
Simple 4	45.4	43.3	5.4
Simple 5	50.0	48.8	0.4

Table 7.1: Percentage of errors for Simple phantoms by gap fill method.

### 7.3.2 Animated phantom results

The Animated phantom configurations all show different kinds of shift motion occurring at different positions and amplitudes. See Chapter 3, Section 3.4 for detailed descriptions of the configurations.





Figure 7.12: Dot plot of bucket assignments for Animated configuration 1.

The bucket assignments shown in Figure 7.12 are consistent with what is seen in Figure 7.1. In particular, the results generated by the unfilled and balanced gap filling methods show the same kind of errors. These errors are, as in the case with the Simple phantoms, consistent for the Animated phantoms as is shown in Table 7.2.



(a) Complete gap filling

Figure 7.13: Dot plot of bucket assignments for Animated configuration 2.



(a) Complete gap filling

Figure 7.14: Dot plot of bucket assignments for Animated configuration 3.



(a) Complete gap filling

Figure 7.15: Dot plot of bucket assignments for Animated configuration 4.

Figures 7.13, 7.14, and 7.15 all show good agreement with the ground truth assignment though the Animated configuration 4 starts to incur some errors. This configuration was explicitly included as a boundary test case to test the robustness of decreasing bin projection percentages. As fewer and fewer projections in the complete scan are attributable to a second motion state, the reconstructed volume for that bin increasingly looks more like the global reconstruction than one capturing a distinct state. With the complete fill method, implicitly the projections that do not come from the original data set will come from the averaged (forward projected) set. In this particular case, 70% of the projections for one of the bins are coming from the average set and only 30% are coming from projections that truly capture the motion in that particular state. The impact of that is seen qualitatively in Figure 7.19 and in the sudden jump in the error rate in Table 7.2.

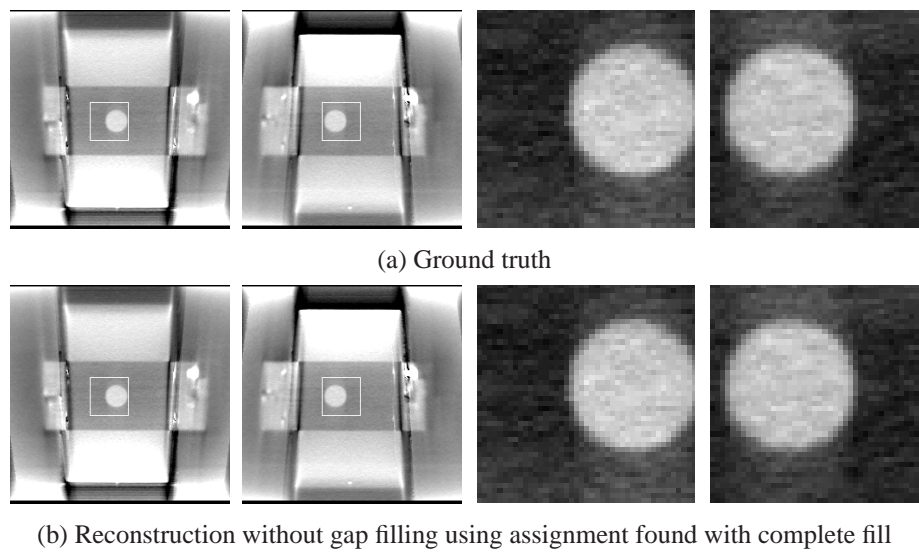


Figure 7.16: Sagittal slices from the volumes generated by the results from the Animated configuration 1 experiment.

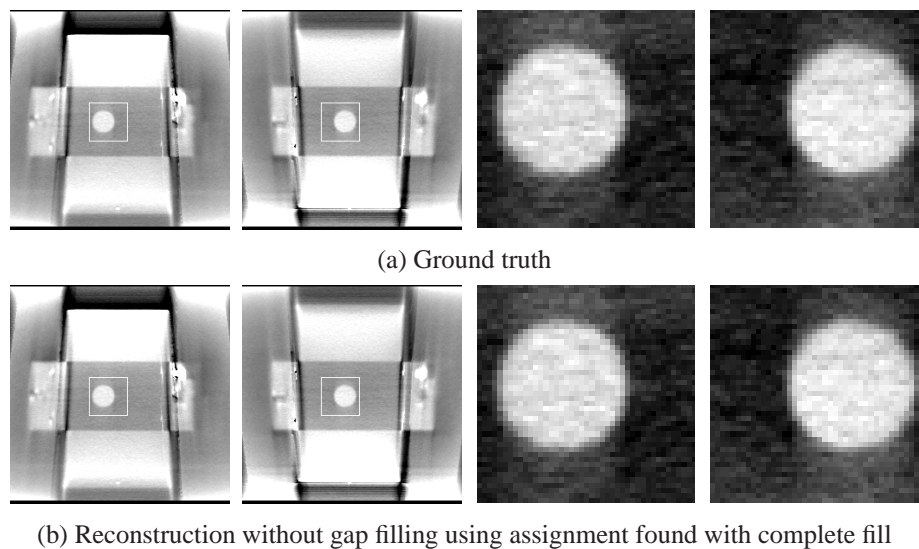


Figure 7.17: Sagittal slices from the volumes generated by the results from the Animated configuration 2 experiment.

Figures 7.16, 7.17 and 7.18 all qualitatively show good results from the search. Of particular note is configuration 3 which shows a small movement in the object. One of the early concerns with the approach that forms the basis of the search in both this chapter and Chapter 6 was that small movement would not generate enough difference information for the SSD metric to be of use. These results, in conjunction with the results for the Simple phantom cases 2 and 3, show that this is not the case at least for these phantom classes.

Figure 7.19 qualitatively shows what happens when there is an insufficient number of projections belonging to a bin. The bin slice on the left, represented with 70% of the projections, is well defined while the slice from the other bin on the right contains a mix of projections and with only 30% of the projections from the correct bin contributing to its reconstruction, any erroneous projection assignments have a relatively greater impact on the reconstruction.

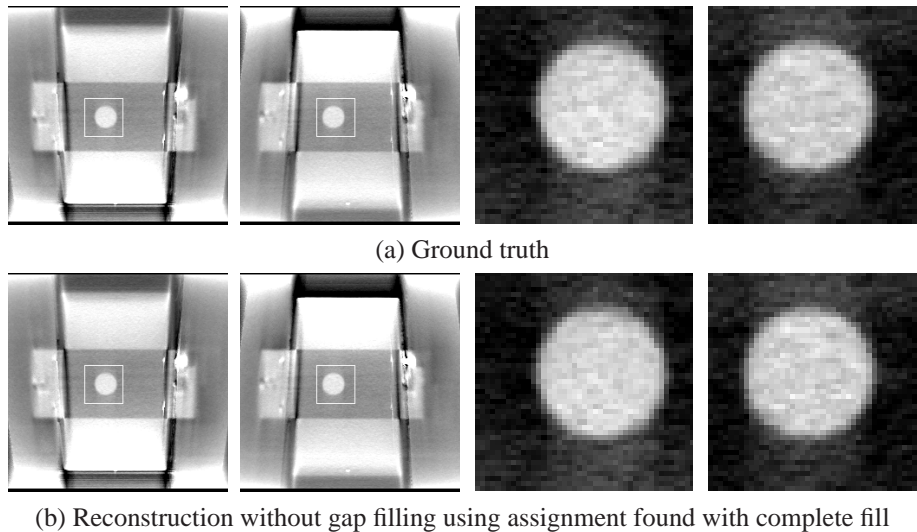





Figure 7.18: Sagittal slices from the volumes generated by the results from the Animated configuration 3 experiment.




(a) Ground truth




(b) No filling results



(c) Balanced fill results



(d) Complete fill results



(e) Reconstruction without gap filling using assignment found with complete fill

Figure 7.19: Sagittal slices from the volumes generated by the results from the Animated configuration 4 experiment.

Phantom configuration	No fill	Balanced	Complete
Animated 1	42.1	50.0	0.0
Animated 2	39.6	49.6	0.4
Animated 3	49.2	49.2	0.4
Animated 4	33.8	30.4	14.6

Table 7.2: Percentage of errors for Simple phantoms by gap fill method.

The results seen in Table 7.2 above confirm the utility of the search method using real projections from a CBCT system. They also demonstrate that as the number of projections from one motion state decreases, the fidelity of the results also decreases. In the the configuration 4 test case, this resulted in a clear reconstruction of one bin but a poor reconstruction of the second bin.

### 7.3.3 NCAT phantom results

The NCAT phantom configurations all show different kinds of continuously periodic motion occurring at different anatomical positions. See Chapter 3, Section 3.5 for detailed descriptions of the configurations.

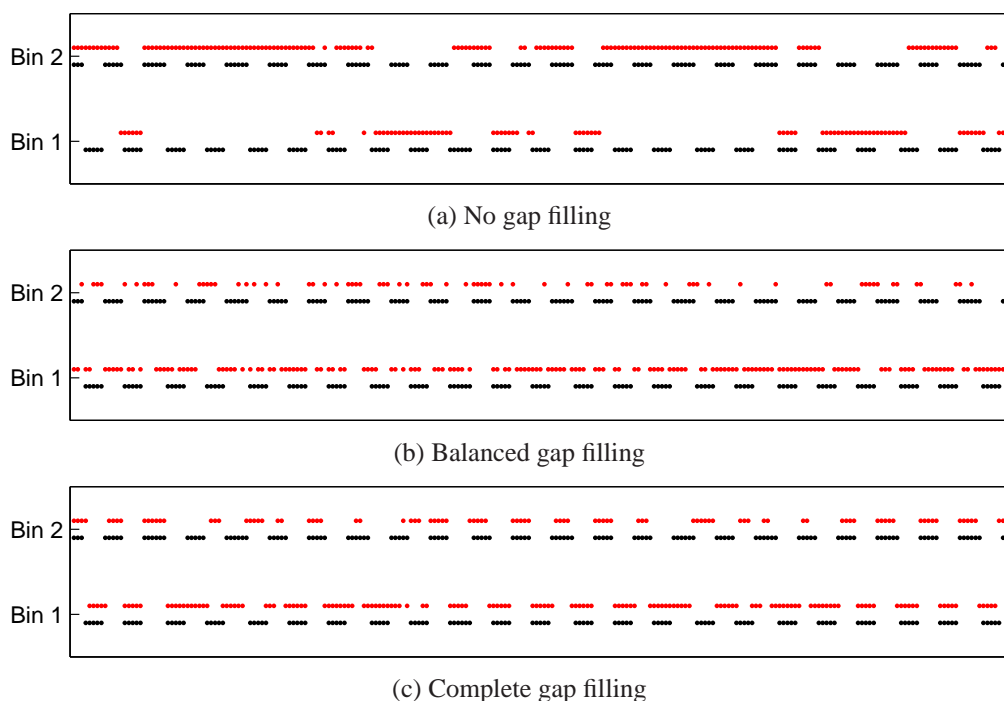




Figure 7.20: Dot plot of bucket assignments for NCAT configuration 1. Each of the three fill types are shown and for each result, the black dots show the ground truth bucket assignment while the red dots show the bucket assignment resulting from the search.

The bucket assignments shown in Figure 7.20 are no longer consistent with what is seen in earlier dot plots for the other phantom types. In particular, the results generated by the balanced fill methods show a more uniform scattering of bucket assignments instead of the kind of assignments seen with the Simple and Animated phantoms wherein most buckets are assigned to one bin and a small number are assigned to the other bin. The same kind of assignment exists for the balanced fill and unfilled searches in the other NCAT phantoms and will be reported numerically in Table 7.3 but not shown in the dot plot figures. The previous best case, from the complete fill variant, shows some regions of correctness but an overall error of almost 29% shows that the search algorithm is not accurate for this particular configuration. This is visually demonstrated in Figure 7.24 as well. One observation to be made is that the total range of motion is shown to be small in the upper portion of the lung and this is in agreement with the literature. The earlier assessment of success with small motion ranges for the Simple and Animated test cases does not hold in this case.




(a) Complete gap filling

Figure 7.21: Dot plot of bucket assignments for NCAT configuration 2. Only the best result from the complete fill type is shown.



(a) Complete gap filling

Figure 7.22: Dot plot of bucket assignments for NCAT configuration 3.



(a) Complete gap filling

Figure 7.23: Dot plot of bucket assignments for NCAT configuration 4.



The dot plots of the bucket assignments in Figures 7.21 and 7.22 again showed poor performance with error rates of about 43% and 50% respectively. The sagittal slice samples in Figures 7.25 and 7.26 visually display the impact of the poor search results.

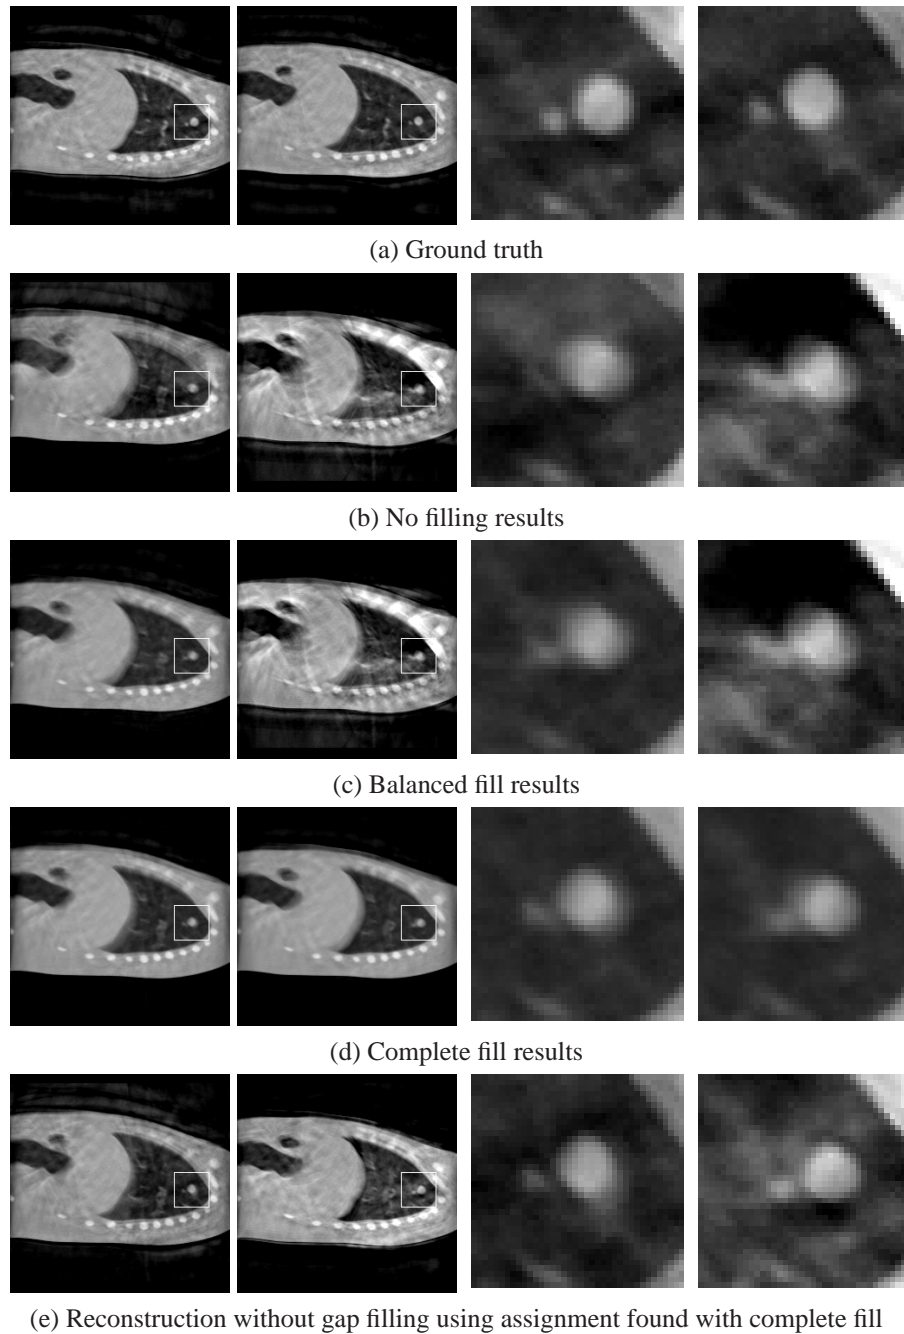




Figure 7.24: Representative sagittal slices from the results for NCAT configuration 1. See Figure 7.8 for a description of the images.




(a) Ground truth




(b) Reconstruction without gap filling using assignment found with complete fill

Figure 7.25: Representative sagittal slices from the results for NCAT configuration 2.



(a) Ground truth



(b) Reconstruction without gap filling using assignment found with complete fill

Figure 7.26: Representative sagittal slices from the results for NCAT configuration 3.

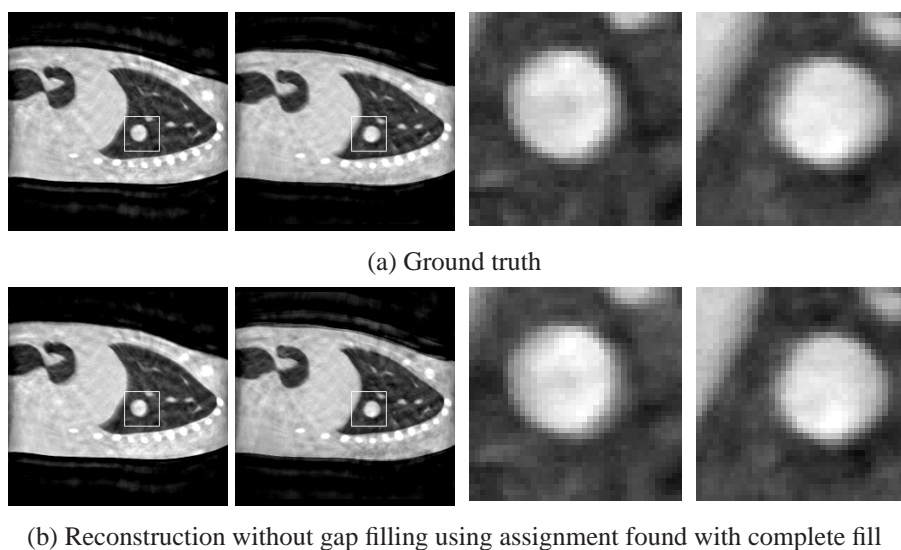


Figure 7.27: Representative sagittal slices from the results for NCAT configuration 4.

NCAT configuration	No fill	Balanced	Complete
NCAT 1	45.0	49.6	28.8
NCAT 2	32.5	45.0	42.9
NCAT 3	49.2	48.3	49.6
NCAT 4	47.9	20.0	8.3

Table 7.3: Percentage of errors for NCAT phantoms by gap filling method.

The results seen in Table 7.3 shows that one-scan search is not ready yet for clinical testing. The complete filling method achieves good results in the fourth case where a large tumour is simulated, and shows better accuracy than the other methods in the first case. Its performance on cases 2 and 3, however, is close to that achieved with a purely random assignment.

## 7.4 Conclusions

The initial hypothesis, that gap-filling would improve the search performance when compared with not filling the gaps, was shown to be true for two of the phantom classes. The surprise was that the superior method in Chapter 5, balanced fill, was not the superior method here. From the perspective of the search algorithm used here, the complete fill method was clearly superior. The third phantom class, the digital anthropomorphic NCAT phantom, yielded unacceptable results however.

If the NCAT results are compared with the results for the same configurations from Chapter 6, something interesting emerges. For the complete fill method of searching in the one-scan case, the effect is similar to the search method in the two-scan case: both construct candidate volumes from a complete set of projections derived from the original projections and a set containing angle-specific, patient-specific and even treatment-specific averages. The differences are in how the average projections are constructed and how the projection space is partitioned. In the two-scan case, the average projections are exactly that. They are pixel-wise averages of the paired projections. In the one-scan case, a more complex method is involved requiring a filtered backprojection reconstruction followed by the forward projection process. The filtered backprojection reconstruction of cone-beam projections is not exact [93] especially as one moves away from the central plane containing the orbit of the X-ray source [32]. Coupled with perhaps an overly simple forward projection method which generates aliasing artefacts and the decision to median filter the projections to reduce scatter noise, the overall system appears to be producing insufficient averaged projections in the NCAT case. Even so, the search method is very successful in the other two phantom cases which is promising. Furthermore, if one looks at the results of using the true averaged projections in the two-scan tests, it is shown that the search idea itself is viable. This lends credence to the hypothesis of insufficient average projections for the one-scan case, at least for complex configurations.

The other principal difference between the reasonably successful results for the NCAT phantom using the two-scan approach and the results obtained with the one-scan approach is in the partitioning. Given the observation that the uniform partitioning of the one-scan approach succeeded on the Animated phantom set while the two-scan partitioning algorithm requires additional work, one can conclude that the partitioning method is much less likely to be the cause of high error rates for the NCAT phantom set in this chapter. It also suggests that some variation of the uniform partitioning algorithm used for the one-scan approach may solve the problems encountered with partitioning on the data sets involving shift motions in the two-scan approach.

A further consideration is in the difference in complexity between the Simple and Animated phantom sets and the NCAT phantom set. In the Simple and Animated cases, a single high-contrast object is moving against a uniform background. In the NCAT case, other moving objects (heart) and high density objects (ribs, diaphragm) have the potential to disrupt the search algorithm. Techniques exist to partially compensate for these effects (see [2,31]) but the fact that the two-scan approach was reasonably successful implies that they may be unnecessary. Further work is needed to verify that the principal issue is the forward projection process and to experiment with other more robust techniques to see if

the error rate for the NCAT test suite can be reduced before clinical testing is warranted.

Comparing the two search methods, the one-scan method presented here and the two-scan method in Chapter 6, there is no clear “winner.” Given an ideal bucketing, the two-scan method outperforms the one-scan method in the majority of the test cases. With the current bucketing algorithm, it is on par with the one-scan method for the continuously periodic Simple test cases and outperforms it on the average tumour size NCAT phantoms. It is significantly faster than the one-scan method because the averaging step is just a pixel-wise average whereas the one-scan method requires a reconstruction and then each of the 670 projections must be constructed with a forward projection. Furthermore, the two-scan search method is exhaustive or semi-exhaustive (when it has to split the search space) while the one-scan relies on a random restart to find the optimal solution. The two-scan search is therefore theoretically more reliable than the one-scan as a search method. In contrast, the one-scan search method is successful on shift motions which was identified in Chapter 2 as the missing motion in current compensation methods and therefore the main opportunity for this research. The Animated data set always under performed in the two-scan search with its current partitioning method whereas the one-scan search was successful on this data. The principal advantage of the one-scan approach is its use of the existing CBCT protocol.

# Chapter 8

## Conclusions and Future Work

---

### 8.1 Summary and novel contributions

In Chapter 2, cone beam computed tomography was introduced in the context of radiotherapy. A review of the relevant biology and physics necessary to understand the modality and some of the problems was presented and then followed with an overview of the clinical concepts and terminology of radiotherapy in general and CBCT specifically. Motion compensation is a significant area of research and there is much overlap with work that has been done in 4D CT; important work done in this area was reviewed and an unaddressed gap in the research was identified. To precisely define this research gap, new terminology was introduced to distinctly classify different types of motion. These classes were coined continuously periodic motion, shift motion, and discrete motion.

In Chapter 3, three different phantom classes were discussed and specific test configurations were described. These phantom configurations were then used throughout the thesis. A digital phantom constructed analytically was called a Simple virtual phantom. A unique hybrid form of phantom was introduced in this chapter, the Animated physical phantom. This novel phantom consists of a physical phantom that is sampled at different positions and which can be virtually animated in an arbitrary fashion along the path the samples are acquired. The benefit of such a phantom, especially when used for CBCT studies, is that a small number of scans can be used to generate an extremely large number of virtual scans ( $n^p$  where  $n$  is the number of samples and  $p$  is the number of projections).

A final type of phantom, the digital anthropomorphic NCAT phantom developed at John Hopkins Medical Institutions was also described.

In Chapter 4, the existence and causes of various kinds of artefacts were reviewed in preparation for a discussion of a particular type of artefact. This artefact, coined gap artefacts, is a significant problem observed in filtered backprojection reconstructions when projections are binned. The artefact was defined as a specific subclass of under sampled artefacts and experiments were performed which quantify the induced errors. In support of these experiments, the use of k-means clustering on a known object's positions as an "oracle" for correct binning was introduced.

In Chapter 5, two methods are proposed and evaluated for mitigating the impact of them on the reconstruction effort. One method used the standard single-scan protocol while the second method introduced a novel two-scan protocol. Each of these methods in turn used two different approaches, balanced fill and complete fill, to fill the gaps. The projections used in both approaches consisted of patient, angle, and treatment-specific averages.

Chapter 6 introduced the core contribution of the thesis: a search method for assigning a bin label to each projection without the use of any prior model or correlating signal/data. The search space was identified, an objective function to be maximized was proposed, and then a series of techniques were described and used to sufficiently reduce the size of the search space. PCA was used in a novel fashion as a preparatory step to increase the overall computational performance of the filtered backprojection reconstruction step. A second novel contribution was the proposal to modify the scanning protocol to include two scans rather than the standard single scan. The advantages and disadvantages of this change were explained in the chapter.

Chapter 7 used several of the ideas in Chapter 6 but applied them to the standard single-scan protocol. This required the development of a new search method. A random-restart hill climbing method was chosen and justified, several experiments are performed, and the results were presented and analysed. The two methods presented in Chapter 6 and Chapter 7 were subsequently compared.

## 8.2 General conclusions

The principal goal of this work was to find a mechanism for identifying motion states in CBCT projections without the use of a prior model or correlated signals. This has been partially achieved in the two-binning case. The general search method used in Chapters 6 and 7 balances the overall difference between the volumes generated at each search step

against their fundamental statistical property, the similarity of their means with the mean of the globally reconstructed volume. Where it performs sub-optimally, comparisons with other cases indicated that the cause was not with the search but with some ancillary step such as the partition estimation in the two-scan method or the forward-projection in the one-scan method.

The partition estimation step in the two-scan search was shown to be effective for continuously periodic motion but ineffective for shift motions. The forward-projection step in the one-scan search did not cause any problems with the Simple and Animated phantom classes, but is hypothesized to be a cause of the underperforming results with the NCAT phantom.

For shift motion, the beginning of poor performance was seen in the 70/30% split case. Small motion changes of approximately 5 mm for both shift and continuously periodic motion types were detected by the one-scan search method.

In both the one-scan and two-scan approaches, searching for a binning assignment is made possible by the use of the computational gains when reconstructing candidate volumes. In standard backprojection reconstruction methods, which are much faster than algebraic methods and hence form the basis for most deployed equipment, a single volume reconstruction requires on the order of minutes to complete. Reconstructing volumes by averaging together sets of eigenspace vectors can be done in Matlab on a computer with a 2.80 GHz Intel Core i7-860 CPU and 8 GB of RAM in less than 0.004 seconds. This comes at an initial overhead cost to read in the previously backprojected volumes of 19.2 minutes from an internal hard drive (one-scan case) or 42.5 minutes from an external USB 2.0 hard drive (two-scan case). Once read in and clipped to the ROI region, the PCA process takes approximately 5 seconds. For the sake of experimentation, complete backprojection volumes were stored on the hard drives. The backprojection process takes approximately 26 minutes, the majority of which is the disk access time. In a real application, only the ROI-clipped sub-volume would be stored on the hard disk or possibly retained in RAM which would significantly speed up the initial overhead.

Using the NCAT phantom configuration 3, the performance of the one-scan and two-scan search methods can be compared. For the one-scan case, the search duration was 285.4 seconds. For the two-scan case, the search duration was 11.3 seconds. In each case, many thousands of reconstructions are needed (more for the one-scan random search method, less for the difference-signal informed two-scan method). Without the computational performance improvements generated by ROI clipping and PCA, the search methods would require many hundreds of hours. With these efficiencies, the worst-case search, including the construction of individual backprojections, can be performed in less



than two hours. Optimizing the methods by keeping clipped backprojections in RAM and parallelizing both the backprojection and search processes would dramatically improve performance.

### 8.3 Future work

The most promising route of continued research is with the one-scan search method. Implementing a more sophisticated forward projection algorithm and evaluating the need for the median filtering step are the principal activities required. Should these be insufficient for successfully finding search solutions to the NCAT phantom test cases, techniques to remove the high density objects should be considered.

If the one-scan search improvements are still considered insufficient then the uniform partitioning used with that method should be applied to the two-scan search method to determine its sufficiency with respect to shift motion.

If the improvements to either the one-scan or two-scan searches enable them to perform adequately on all data sets, then an iterative approach can be considered with respect to the number of bins. By first binning into two bins, each of these bins may possibly support further binning into two more bins for a total of four bins. Splitting the projection set into 4 bins implies an upper bound on the smallest bin count of 25% of the projections. The Animated test configuration 4, where shift motion occurred at the 70% point, showed some degradation in accuracy so lowering the projection count from 30% to 25% (best case for at least one of the bins) may prove problematic. In that case, the techniques in common use with other research described in Chapter 2, namely performing a B-spline registration and then using that registration to deform the projections from bin 1 to fit bin 2 may prove useful.

One class of tumours that are likely to cause both search methods to underperform are the low-contrast tumours visualized without any contrast medium. As part of the validation step for any performance improved search algorithm, these test cases should be added with varying degrees of contrast evaluated.

The algorithms implemented in this thesis made heavy use of hard disk resources as a way of saving state and facilitating the extraction of necessary plots and images to document performance and accuracy. These disk accesses are the bottleneck to the overall system performance and most of the largest disk accesses can be removed. Additionally, most of the pre-processing on the projections is independent of other projections and can be parallelized very easily. Likewise, the multi-restart nature of the one-scan search and the sequential nature of the two-scan search make them both candidates for parallelization.

Given the push for “cloud computing” applications, the ideas expressed here are natural candidates for such an implementation and results needed in the timeframe required for clinical protocols is achievable.

Finally, the application goal of all of this work is to enable collaborative studies with a team of clinicians so that discovered tumour motion can be compared with expected tumour motion. The knowledge gained by this clinical research may help to better quantify the margins used in planning the radiotherapy and better margins will lead to reduced dose to healthy tissues and increased dose to the target tissues. Reducing the dose to healthy tissues reduces the risk of DNA damage which in turn reduces the risk of radiation-induced secondary cancers. Increasing the dose to the target tissue increases the likelihood of successfully killing the tissue and avoiding reoccurring cancer. Richard Hamming is quoted as saying, “the purpose of computing is insight, not numbers.” I will conclude the reporting of this work by paraphrasing him: the purpose of this research is better health, not numbers.

# Bibliography

- [1] A F Abdelnour, S A Nehmeh, T Pan, J L Humm, P Vernon, H Schöder, K E Rosenzweig, G S Mageras, E Yorke, S M Larson, and Y E Erdi. Phase and amplitude binning for 4D-CT imaging. *Physics in medicine and biology*, 52(12):3515–29, June 2007.
- [2] Mónica Abella and J.A. Fessler. A new statistical image reconstruction algorithm for polyenergetic X-ray CT. In *Proceedings of the Sixth IEEE international conference on Symposium on Biomedical Imaging: From Nano to Macro*, pages 165–168. IEEE Press, 2009.
- [3] H Alasti, Y B Cho, A D Vandermeer, A Abbas, B Norrlinger, S Shubbar, and A Bezjak. A novel four-dimensional radiotherapy method for lung cancer: imaging, treatment planning and delivery. *Physics in medicine and biology*, 51(12):3251–67, June 2006.
- [4] American Cancer Society. *Cancer Facts & Figures 2007*. American Cancer Society, Atlanta, 2007.
- [5] Shai Avidan and Amnon Shashua. Trajectory triangulation: 3D reconstruction of moving points from a monocular image sequence. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 22(4):348–357, 2002.
- [6] Julia F Barrett and Nicholas Keat. Artifacts in CT: recognition and avoidance. *Radiographics : a review publication of the Radiological Society of North America, Inc*, 24(6):1679–91, 2004.
- [7] C Blondel. 4D deformation field of coronary arteries from monoplane rotational X-ray angiography. *International Congress Series*, 1256(2003):1073–1078, June 2003.

- [8] Christophe Blondel, Régis Vaillant, Grégoire Malandain, and Nicholas Ayache. 3D tomographic reconstruction of coronary arteries using a precomputed 4D motion field. *Physics in Medicine and Biology*, 49(11):2197–2208, June 2004.
- [9] Gerben R Borst, Jan-Jakob Sonke, Anja Betgen, Peter Remeijer, Marcel van Herk, and Joos V Lebesque. Kilo-voltage cone-beam computed tomography setup measurements for lung cancer patients; first clinical results and comparison with electronic portal-imaging device. *International journal of radiation oncology, biology, physics*, 68(2):555–61, June 2007.
- [10] Jenghwa Chang, Gig S Mageras, Ellen Yorke, Fernando De Arruda, Jussi Sillanpaa, Kenneth E Rosenzweig, Agung Hertanto, Hai Pham, Edward Seppi, Alex Pevsner, C Clifton Ling, and Howard Amols. Observation of interfractional variations in lung tumor position using respiratory gated and ungated megavoltage cone-beam computed tomography. *International journal of radiation oncology, biology, physics*, 67(5):1548–58, April 2007.
- [11] Jenghwa Chang, Jussi Sillanpaa, Clifton C. Ling, Edward Seppi, Ellen Yorke, Gikas Mageras, and Howard Amols. Integrating respiratory gating into a megavoltage cone-beam CT system. *Medical Physics*, 33(7):2354, 2006.
- [12] P S Cho, R H Johnson, and T W Griffin. Cone-beam CT for radiotherapy applications. *Physics in medicine and biology*, 40(11):1863–83, November 1995.
- [13] Laura A Dawson, Cynthia Eccles, and Tim Craig. Individualized image guided iso-NTCP based liver cancer SBRT. *Acta oncologica (Stockholm, Sweden)*, 45(7):856–64, January 2006.
- [14] Laura A Dawson and David A Jaffray. Advances in image-guided radiation therapy. *Journal of clinical oncology : official journal of the American Society of Clinical Oncology*, 25(8):938–46, March 2007.
- [15] Laura A Dawson and Michael B Sharpe. Image-guided radiotherapy: rationale, benefits, and limitations. *The lancet oncology*, 7(10):848–58, October 2006.
- [16] M Defrise and R Clack. A cone-beam reconstruction algorithm using shift-variant filtering and cone-beam backprojection. *IEEE transactions on medical imaging*, 13(1):186–95, January 1994.

- [17] Junjun Deng, Hengyong Yu, Jun Ni, Tao He, Shiyong Zhao, Lihe Wang, and Ge Wang. A Parallel Implementation of the Katsevich Algorithm for 3-D CT Image Reconstruction. *The Journal of Supercomputing*, 38(1):35–47, October 2006.
- [18] S Diederich, H Lenzen, R Windmann, Z Puskas, T M Yelbuz, S Henneken, T Klaiber, M Eameri, N Roos, and P E Peters. Pulmonary nodules: experimental and clinical studies at low-dose CT. *Radiology*, 213(1):289–98, October 1999.
- [19] Lars Dietrich, Siri Jetter, Thomas Tücking, Simeon Nill, and Uwe Oelfke. Linac-integrated 4D cone beam CT: first experimental results. *Physics in medicine and biology*, 51(11):2939–52, June 2006.
- [20] Sergio Díez, Javier García, and Francisco Sendra. Analysis and evaluation of periodic physiological organ motion in radiotherapy treatments. *Radiotherapy and oncology : journal of the European Society for Therapeutic Radiology and Oncology*, 73(3):325–9, December 2004.
- [21] David Dowsett, Patrick A. Kenny, and R. Eugene Johnston. *The Physics of Diagnostic Imaging*. Hodder Arnold, New York, second edition, 2006.
- [22] L. A. Feldkamp, L. C. Davis, and J. W. Kress. Practical cone-beam algorithm. *Journal of the Optical Society of America A*, 1(6):612, June 1984.
- [23] Isabelle M. Gagne and Don M. Robinson. The impact of tumor motion upon CT image integrity and target delineation. *Medical Physics*, 31(12):3378, 2004.
- [24] David P Gierga, George T Y Chen, Jong H Kung, Margrit Betke, Jonathan Lombardi, and Christopher G Willett. Quantification of respiration-induced abdominal tumor motion and its impact on IMRT dose distributions. *International journal of radiation oncology, biology, physics*, 58(5):1584–95, April 2004.
- [25] Pierre Grangeat, Anne Koenig, Thomas Rodet, and Stéphane Bonnet. Theoretical framework for a dynamic cone-beam reconstruction algorithm based on a dynamic particle model. *Physics in medicine and biology*, 47(15):2611–25, August 2002.
- [26] ICRU. *ICRU report 50: prescribing, recording, and reporting photon beam therapy*. International Commission on Radiation Units and Measurements, 1993.
- [27] Mohammad K. Islam, Thomas G. Purdie, Bernhard D. Norrlinger, Hamideh Alasti, Douglas J. Moseley, Michael B. Sharpe, Jeffrey H. Siewerdsen, and David A. Jaffray. Patient dose from kilovoltage cone beam computed tomography imaging in radiation therapy. *Medical Physics*, 33(6):1573, 2006.

- [28] D. A. Jaffray, D. G. Drake, M. Moreau, A. A. Martinez, and J. W. Wong. A radiographic and tomographic imaging system integrated into a medical linear accelerator for localization of bone and soft-tissue targets. *International journal of radiation oncology, biology, physics*, 45(3):773–89, October 1999.
- [29] David A Jaffray, Jeffrey H Siewerdsen, John W Wong, and Alvaro A Martinez. Flat-panel cone-beam computed tomography for image-guided radiation therapy. *International journal of radiation oncology, biology, physics*, 53(5):1337–49, August 2002.
- [30] P M Joseph. An Improved Algorithm for Reprojecting Rays through Pixel Images. *IEEE transactions on medical imaging*, 1(3):192–6, January 1982.
- [31] PM Joseph and RD Spital. A method for correcting bone induced artifacts in computed tomography scanners. *Journal of computer assisted tomography*, 1978.
- [32] Avinash C. Kak and Malcolm Slaney. *Principles of Computerized Tomographic Imaging*. SIAM, 2001.
- [33] A. Katsevich. Theoretically exact filtered backprojection-type inversion algorithm for spiral CT. *SIAM Journal on Applied Mathematics*, 62(6):2012–2026, 2002.
- [34] Paul J. Keall, Sarang Joshi, S. Sastry Vedam, Jeffrey V. Siebers, Vijaykumar R. Kini, and Radhe Mohan. Four-dimensional radiotherapy planning for DMLC-based respiratory motion tracking. *Medical Physics*, 32(4):942, 2005.
- [35] Paul J. Keall, Gig S. Mageras, James M. Balter, Richard S. Emery, Kenneth M. Forster, Steve B. Jiang, Jeffrey M. Kapatoes, Daniel A. Low, Martin J. Murphy, Brad R. Murray, Chester R. Ramsey, Marcel B. Van Herk, S. Sastry Vedam, John W. Wong, and Ellen Yorke. The management of respiratory motion in radiation oncology report of AAPM Task Group 76. *Medical Physics*, 33(10):3874, 2006.
- [36] K M Langen and D T Jones. Organ motion and its management. *International journal of radiation oncology, biology, physics*, 50(1):265–78, May 2001.
- [37] T Li, E Schreibmann, Y Yang, and L Xing. Motion correction for improved target localization with on-board cone-beam computed tomography. *Physics in medicine and biology*, 51(2):253–67, January 2006.

- [38] Tianfang Li, Albert Koong, and Lei Xing. Enhanced 4D cone-beam CT with inter-phase motion model. *Medical Physics*, 34(9):3688, 2007.
- [39] Tianfang Li and Lei Xing. Optimizing 4D cone-beam CT acquisition protocol for external beam radiotherapy. *International journal of radiation oncology, biology, physics*, 67(4):1211–9, March 2007.
- [40] Tianfang Li, Lei Xing, Peter Munro, Christopher McGuinness, Ming Chao, Yong Yang, Bill Loo, and Albert Koong. Four-dimensional cone-beam computed tomography using an on-board imager. *Medical Physics*, 33(10):3825, 2006.
- [41] Xiang Li, Pengpeng Zhang, Dennis Mah, Richard Gewanter, and Gerald Kutcher. Novel lung IMRT planning algorithms with nonuniform dose delivery strategy to account for respiratory motion. *Medical Physics*, 33(9):3390, 2006.
- [42] N C Linney and P H Gregson. Organ motion detection in CT images using opposite rays in fan-beam projection systems. *IEEE transactions on medical imaging*, 20(11):1109–22, November 2001.
- [43] Yong Long, JA Fessler, and JM Balter. A 3D Forward and Back-Projection Method for X-Ray CT Using Separable Footprint. In *10th International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine*, pages 146–149, 2009.
- [44] Daniel A. Low, Michelle Nystrom, Eugene Kalinin, Parag Parikh, James F. Dempsey, Jeffrey D. Bradley, Sasa Mutic, Sasha H. Wahab, Tareque Islam, Gary Christensen, David G. Polite, and Bruce R. Whiting. A method for the reconstruction of four-dimensional synchronized CT scans acquired during free breathing. *Medical Physics*, 30(6):1254, 2003.
- [45] Daniel A Low, Parag J Parikh, Wei Lu, James F Dempsey, Sasha H Wahab, James P Hubenschmidt, Michelle M Nystrom, Maureen Handoko, and Jeffrey D Bradley. Novel breathing motion model for radiotherapy. *International journal of radiation oncology, biology, physics*, 63(3):921–9, November 2005.
- [46] Jun Lu, Thomas M. Guerrero, Peter Munro, Andrew Jeung, Pai-Chun M. Chi, Peter Balter, X. Ronald Zhu, Radhe Mohan, and Tinsu Pan. Four-dimensional cone beam CT with adaptive gantry rotation and adaptive data sampling. *Medical Physics*, 34(9):3520, 2007.

- [47] Wei Lu, Parag J. Parikh, James P. Hubenschmidt, Jeffrey D. Bradley, and Daniel A. Low. A comparison between amplitude sorting and phase-angle sorting using external respiratory measurement for 4D CT. *Medical Physics*, 33(8):2964, 2006.
- [48] Weiguo Lu, Gustavo H Olivera, Quan Chen, Ming-Li Chen, and Kenneth J Ruchala. Automatic re-contouring in 4D radiotherapy. *Physics in medicine and biology*, 51(5):1077–99, March 2006.
- [49] Anthony E. Lujan, Edward W. Larsen, James M. Balter, and Randall K. Ten Haken. A method for incorporating organ motion due to breathing into 3D dose calculations. *Medical physics*, pages 715–720, 1999.
- [50] Dennis Mah, Gary Freedman, Bart Milestone, Alexandra Hanlon, Elizabeth Palacio, Theresa Richardson, Benjamin Movsas, Raj Mitra, Eric Horwitz, and Gerald E Hanks. Measurement of intrafractional prostate motion using magnetic resonance imaging. *International journal of radiation oncology, biology, physics*, 54(2):568–75, October 2002.
- [51] T E Marchant, A M Amer, and C J Moore. Measurement of inter and intra fraction organ motion in radiotherapy using cone beam CT projection images. *Physics in medicine and biology*, 53(4):1087–98, March 2008.
- [52] Samuel Matej, Jeffrey A Fessler, and Ivan G Kazantsev. Iterative tomographic image reconstruction using Fourier-based forward and back-projectors. *IEEE transactions on medical imaging*, 23(4):401–12, April 2004.
- [53] Jamie R. McClelland, Jane M. Blackall, Segolene Tarte, Adam C. Chandler, Simon Hughes, Shahreen Ahmad, David B. Landau, and David J. Hawkes. A continuous 4D motion model from multiple respiratory cycles for use in lung radiotherapy. *Medical Physics*, 33(9):3348, 2006.
- [54] Coert T Metz, Michiel Schaap, Stefan Klein, Lisan A Neefjes, Ermanno Capuano, Carl Schultz, Robert Jan van Geuns, Patrick W Serruys, Theo van Walsum, and Wiro J Niessen. Patient specific 4D coronary models from ECG-gated CTA data for intra-operative dynamic alignment of CTA with X-ray images. In *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, volume 12, pages 369–76, January 2009.



- [55] LE Millender, M Aubin, J Pouliot, K Shinohara, and M Roach. Daily electronic portal imaging for morbidly obese men undergoing radiotherapy for localized prostate cancer. *Journal of Radiation*, 59(1):6–10, 2004.
- [56] Shinichiro Mori, Masahiro Endo, Shuhei Komatsu, Susumu Kandatsu, Tomoyasu Yashiro, and Masayuki Baba. A combination-weighted Feldkamp-based reconstruction algorithm for cone-beam CT. *Physics in medicine and biology*, 51(16):3953–65, August 2006.
- [57] K Mueller, R Yagel, and J J Wheller. Anti-aliased three-dimensional cone-beam reconstruction of low-contrast objects with algebraic methods. *IEEE transactions on medical imaging*, 18(6):519–37, June 1999.
- [58] J Muller-Merbach. Simulation of x-ray projections for experimental 3d tomography. *Report LiTH-ISY-R-1866*, 1996.
- [59] Christopher Nelson, George Starkschall, Peter Balter, Rodolfo C Morice, Craig W Stevens, and Joe Y Chang. Assessment of lung tumor motion and setup uncertainties using implanted fiducials. *International journal of radiation oncology, biology, physics*, 67(3):915–23, March 2007.
- [60] A.H.W. Nias. *An introduction to radiobiology*. John Wiley & Sons, Ltd., Chichester, England, second edition, 1998.
- [61] Uwe Oelfke, Thomas Tücking, Simeon Nill, Annete Seeber, Bernd Hesse, Peter Huber, and Christoph Thilmann. Linac-integrated kV-cone beam CT: technical features and first applications. *Medical dosimetry : official journal of the American Association of Medical Dosimetrists*, 31(1):62–70, January 2006.
- [62] B Ohnesorge, T Flohr, K Schwarz, J P Heiken, and K T Bae. Efficient correction for CT image artifacts caused by objects extending outside the scan field of view. *Medical physics*, 27(1):39–46, January 2000.
- [63] Nigel P. Orton and Wolfgang A. Tome. The impact of daily shifts on prostate IMRT dose distributions. *Medical Physics*, 31(10):2845, 2004.
- [64] Thomas G Purdie, Douglas J Moseley, Jean-Pierre Bissonnette, Michael B Sharpe, Kevin Franks, Andrea Bezjak, and David A Jaffray. Respiration correlated cone-beam computed tomography and 4DCT for evaluating target motion in Stereotactic Lung Radiation Therapy. *Acta oncologica (Stockholm, Sweden)*, 45(7):915–22, January 2006.

- [65] J G Ravenel, E M Scalzetti, W Huda, and W Garrisi. Radiation exposure and image quality in chest CT examinations. *AJR. American journal of roentgenology*, 177(2):279–84, August 2001.
- [66] Eike Rietzel and George T. Y. Chen. Deformable registration of 4D computed tomography data. *Medical Physics*, 33(11):4423, 2006.
- [67] Eike Rietzel, George T Y Chen, Noah C Choi, and Christopher G Willet. Four-dimensional image-based treatment planning: Target volume segmentation and dose calculation in the presence of respiratory motion. *International journal of radiation oncology, biology, physics*, 61(5):1535–50, April 2005.
- [68] Eike Rietzel, Tinsu Pan, and George T. Y. Chen. Four-dimensional computed tomography: Image formation and clinical protocol. *Medical Physics*, 32(4):874, 2005.
- [69] Simon Rit, David Sarrut, and Laurent Desbat. Comparison of analytic and algebraic methods for motion-compensated cone-beam CT reconstruction of the thorax. *IEEE transactions on medical imaging*, 28(10):1513–25, October 2009.
- [70] Simon Rit, David Sarrut, and Chantal Ginestet. Respiratory signal extraction for 4D CT imaging of the thorax from cone-beam CT projections. *Medical image computing and computer-assisted intervention : MICCAI ... International Conference on Medical Image Computing and Computer-Assisted Intervention*, 8(Pt 1):556–63, January 2005.
- [71] Simon Rit, David Sarrut, and Serge Miguet. Gated cone-beam CT imaging of the thorax: a reconstruction study. *Proceedings of SPIE*, (0):651022–651022–10, 2007.
- [72] C Rohkohl, G Lauritsch, L Biller, M Prümmer, J Boese, and J Hornegger. Interventional 4D motion estimation and reconstruction of cardiac vasculature without motion periodicity assumption. In *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, volume 14, pages 687–94, October 2009.
- [73] S.J. Russell and P. Norvig. *Artificial intelligence: a modern approach*. Prentice hall, second edition, 2009.
- [74] David Sarrut. Deformable registration for image-guided radiation therapy. *Zeitschrift für medizinische Physik*, 16(4):285–97, January 2006.

- [75] David Sarrut, Vlad Boldea, Serge Miguet, and Chantal Ginestet. Simulation of four-dimensional CT images from deformable registration between inhale and exhale breath-hold CT scans. *Medical Physics*, 33(3):605, 2006.
- [76] Konrad Schindler and David Suter. Two-view multibody structure-and-motion with outliers through model selection. *IEEE transactions on pattern analysis and machine intelligence*, 28(6):983–995, 2006.
- [77] Eduard Schreibmann, George T Y Chen, and Lei Xing. Image interpolation in 4D CT using a BSpline deformable registration model. *International journal of radiation oncology, biology, physics*, 64(5):1537–50, April 2006.
- [78] W.P. Segars, D.S. Lalush, and B.M.W. Tsui. A realistic spline-based dynamic heart phantom. *IEEE Transactions on Nuclear Science*, 46(3):503–506, June 1999.
- [79] W.P. Segars, D.S. Lalush, and B.M.W. Tsui. Modeling respiratory mechanics in the MCAT and spline-based MCAT phantoms. *IEEE Transactions on Nuclear Science*, 48(1):89–97, 2001.
- [80] W.P. Segars, B.M.W. Tsui, E.C. Frey, and E.K. Fishman. Extension of the 4D NCAT phantom to dynamic X-ray CT simulation. *2003 IEEE Nuclear Science Symposium. Conference Record (IEEE Cat. No.03CH37515)*, pages 3195–3199, 2004.
- [81] Yvette Seppenwoolde, Hiroki Shirato, Kei Kitamura, Shinichi Shimizu, Marcel van Herk, Joos V Lebesque, and Kazuo Miyasaka. Precise and real-time measurement of 3D tumor motion in lung due to breathing and heartbeat, measured during radiotherapy. *International journal of radiation oncology, biology, physics*, 53(4):822–34, July 2002.
- [82] Gregory C Sharp, Steve B Jiang, Shinichi Shimizu, and Hiroki Shirato. Prediction of respiratory tumour motion for real-time image-guided radiotherapy. *Physics in Medicine and Biology*, 49(3):425–440, February 2004.
- [83] Jeffrey H. Siewerdsen and David A Jaffray. Cone-beam computed tomography with a flat-panel imager: Magnitude and effects of x-ray scatter. *Medical Physics*, 28(2):220, 2001.
- [84] Jeffrey H. Siewerdsen, D. J. Moseley, B. Bakhtiar, S. Richard, and David A Jaffray. The influence of antiscatter grids on soft-tissue detectability in cone-beam

- computed tomography with flat-panel detectors. *Medical Physics*, 31(12):3506, 2004.
- [85] BD Smith. Image reconstruction from cone-beam projections: necessary and sufficient conditions and reconstruction methods. *IEEE Trans. Med. Imaging*, 1985.
- [86] Jan-Jakob Sonke, Joos Lebesque, and Marcel van Herk. Variability of four-dimensional computed tomography patient models. *International journal of radiation oncology, biology, physics*, 70(2):590–8, February 2008.
- [87] Jan-Jakob Sonke, Lambert Zijp, Peter Remeijer, and Marcel van Herk. Respiratory correlated cone beam CT. *Medical Physics*, 32(4):1176, 2005.
- [88] R.K. Sundaram. *A first course in optimization theory*. Cambridge Univ Pr, 1996.
- [89] D I Thwaites, A R DuSautoy, T Jordan, M R McEwen, A Nisbet, A E Nahum, and W G Pitchford. The IPEM code of practice for electron dosimetry for radiotherapy beams of initial energy from 4 to 25 MeV based on an absorbed dose to water calibration. *Physics in medicine and biology*, 48(18):2929–70, September 2003.
- [90] Henrik Turbell. New Functionality in take version 2 . 1, 1999.
- [91] Henrik Turbell. *Cone-beam reconstruction using filtered backprojection*. PhD thesis, Linkopings universitet, 2001.
- [92] Matthew Turk and Alex Pentland. Eigenfaces for recognition. *Journal of cognitive neuroscience*, 3(1), 1991.
- [93] Heang K. Tuy. An Inversion Formula for Cone-Beam Reconstruction. *SIAM Journal on Applied Mathematics*, 43(3):546, 1983.
- [94] Emile N J Th van Lin, Lisette P van der Vight, J Alfred Witjes, Henkjan J Huisman, Jan Willem Leer, and Andries G Visser. The effect of an endorectal balloon and off-line correction on the interfraction systematic and random prostate position variations: a comparative study. *International journal of radiation oncology, biology, physics*, 61(1):278–88, January 2005.
- [95] S S Vedam, P J Keall, V R Kini, H Mostafavi, H P Shukla, and R Mohan. Acquiring a four-dimensional computed tomography dataset using an external respiratory signal. *Physics in medicine and biology*, 48(1):45–62, January 2003.

- [96] A Wambersie and T Landberg. ICRU Report 62: Prescribing, recording, and reporting photon beam therapy (supplement to ICRU report 50). *Bethesda, MD: International Commission on Radiation*, 1999.
- [97] Andrew Webb. *Introduction to biomedical imaging*. John Wiley & Sons, Inc., Hoboken, New Jersey, 2003.
- [98] S Webb. Motion effects in (intensity modulated) radiation therapy: a review. *Physics in medicine and biology*, 51(13):R403–25, July 2006.
- [99] WHO. GLOBOCAN. <http://globocan.iarc.fr/>, Accessed November 2010.
- [100] G L Zeng, R Clack, and G T Gullberg. Implementation of Tuy’s cone-beam inversion formula. *Physics in medicine and biology*, 39(3):493–507, March 1994.
- [101] Rongping Zeng, Jeffrey A. Fessler, and James M. Balter. Respiratory motion estimation from slowly rotating x-ray projections: Theory and simulation. *Medical Physics*, 32(4):984, 2005.
- [102] Rongping Zeng, Jeffrey A. Fessler, James M. Balter, and Peter A. Balter. Iterative sorting for four-dimensional CT images based on internal anatomy motion. *Medical Physics*, 35(3):917, 2008.
- [103] Chong Zhang, Mathieu De Craene, Maria-Cruz Villa-Uriol, Jose M Pozo, Bart H Bijnens, and Alejandro F Frangi. Estimating continuous 4D wall motion of cerebral aneurysms from 3D rotational angiography. In *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, volume 12, pages 140–7, January 2009.