

BOSTON UNIVERSITY
GRADUATE SCHOOL OF ARTS AND SCIENCES

Thesis

**REAL-TIME 4D TUMOR TRACKING AND MODELING FROM INTERNAL
AND EXTERNAL FIDUCIALS IN FLUOROSCOPY**

by

JOHANNA M. BREWER

B.A., Boston University, 2002

Submitted in partial fulfillment of the
requirements for the degree of
Master of Arts

2004

Approved by

First Reader

Margrit Betke, Ph.D.
Associate Professor of Computer Science
Boston University

Second Reader

Stan Sclaroff, Ph.D.
Associate Professor of Computer Science
Boston University

Third Reader

David Gierga, Ph.D.
Radiation Physicist, Massachusetts General Hospital
Instructor, Radiation Oncology, Harvard Medical School

Acknowledgments

This thesis is based on joint work with Margrit Betke, Jonathan Lombardi, David P. Gierga, Gregory C. Sharp, Steve B. Jiang and George T. Y. Chen. Thanks to the Image and Video Computing Group at Boston University, especially William Mullally and John Isodoro. For the support of Kate and Jeffry Mullen I am ever grateful. Let them be thanked. Funding was provided by The Whitaker Foundation, the National Science Foundation (DGE-0221680, IIS-0093367, IIS-0308213 and EIA-0202067), and the Office of Naval Research (N000140110444).

**REAL-TIME 4D TUMOR TRACKING AND MODELING FROM INTERNAL
AND EXTERNAL FIDUCIALS IN FLUOROSCOPY**

JOHANNA M. BREWER

Abstract

Fluoroscopy is currently used in treatment planning for patients undergoing radiation therapy for lung and abdominal tumors. Radiation oncologists seek to maximize the radiation the tumor receives, and minimize the amount delivered to the surrounding tissues, a task made difficult by movement caused by the breathing of the patient. A model of the tumor motion could greatly improve dose calculation and delivery. There exists a system which directly estimates the location of the tumor, but this system is very costly and not widely deployed. In this thesis a model which can be derived using less specialized equipment is presented. First, a method of tracking the two-dimensional (2D) motion of internal markers (surgical clips) placed around the tumor is presented and examined. A ground truth is established by visual inspection of 10 data sets of 5 patients to evaluate the tracker. The root mean squared error in estimating 2D marker position was 0.47 mm on average. Using two orthogonal sequentially obtained fluoroscopic image sequences, a method for calculating a model of the average or maximum three-dimensional (3D) motion of the clips is presented, examined, and compared to the direct estimation system. On average, the error was 3.0 mm for four pairs of trajectories. If imaging is possible during treatment, this modeled motion can be used for beam-guided radiation, otherwise, the modeled motion

can be correlated to a set of external markers for use in respiratory gating.

Contents

1	Introduction	1
2	Related Work	6
2.1	Radiation Therapy	6
2.1.1	Breath-hold Techniques	7
2.1.2	Gated Therapy	8
2.1.3	Correlation	9
2.1.4	Image-guided Therapy	11
2.2	Tracking	13
2.2.1	Radiation Therapy	14
2.2.2	Computer Vision and Medical Imaging	15
2.3	3D Modeling, Transformation, and Deformation	17
2.3.1	Medical Imaging	18
2.4	Summary	19
3	Materials and Methods	21
3.1	2D Clip Tracking	24
3.1.1	Initialization	25

3.1.2	Tracking Algorithm	26
3.2	Estimated Model of 3D Clip Motion	28
3.2.1	Finding the Maximum and Minimum of Each Breathing Cycle	28
3.2.2	Calculating the Parameters of a 1D Model	32
3.2.3	Calculating the Parameters of the Final 3D Model	34
3.3	Fourier Analysis	34
3.4	Correlation with External Markers	35
4	Results	37
4.1	2D Clip Tracking	37
4.2	Estimated Model of 3D Clip Motion	39
4.3	Correlation With External Markers	42
5	Discussion and Conclusion	48
5.1	2D Clip Tracking	48
5.2	Estimated Model of 3D Clip Motion	50
5.3	Correlation with External Markers	53
5.4	Future Work	56
	Bibliography	61

List of Tables

4.1	Clip Localization Error	39
4.2	1D Estimated Model Parameters: Patient Jack	40
4.3	3D Estimated Model Parameters: Patient Jack	40
4.4	1D Estimated Model Parameters: Patient Nancy	40
4.5	3D Estimated Model Parameters: Patient Nancy	40
4.6	3D Estimated Model Error	41
4.7	1D Frequency Parameter Comparison: Patient Jack	42
4.8	1D Frequency Parameter Comparison: Patient Nancy	42
4.9	External Marker Correlation	47

List of Figures

1.1	Example of a linear accelerator	2
1.2	Diagram of a linear accelerator	3
3.1	Coordinate system of patient setup	21
3.2	High-level Overview of Goals of Method	22
3.3	Detailed Overview of Method	24
3.4	Spine crossing problem	25
3.5	Inhalation-induced darkening problem	26
3.6	Smoothing to find local maximum	30
3.7	Modified sliding window technique	31
3.8	Average motion shown in fluoroscopic image	35
3.9	External markers	36
4.1	2D trajectories produced by tracker	38
4.2	Original trajectories compared to models of average 3D motion	41
4.3	External markers and models of average 3D motion	43
4.4	External marker correlation plots 1	45
4.5	External marker correlation plots 2	46

5.1 Clips in CT slices 58

List of Abbreviations

1D	one-dimensional
2D	two-dimensional
3D	three-dimensional
4DCT	four-dimensional computed tomography
ABC	Active Breathing Control
ATT	average tumor trajectory
CT	computed tomography
CTV	clinical target volume
DIBH	Deep Inspiration Breath-Hold
DMLC	dynamic multileaf collimator
DVH	dose-volume histogram
IMRT	intensity-modulated radiotherapy
IRIS	Integrated Radiotherapy Imaging System
linac	linear accelerator
MAX-T	motion adaptive X-ray therapy
MR	magnetic resonance
PTV	planning target volume
RMS	root mean squared
RTRT	Real-time Tumor Tracking Radiotherapy
SMART	Synchronized Moving Aperture Radiation Therapy

Chapter 1

Introduction

Despite all of the advances of modern science the treatment of cancer is still an open problem. According to the Cancer Statistics 2004 report from the American Cancer Society [6], this year about 563,700 Americans are expected to die of cancer. This report also states that the 5-year survival rates for lung/bronchial and pancreatic cancer during 1974-1976 were 12% and 3% respectively. Between 1992-1999 the 5-year survival rates were only 15% and 4% respectively. For other forms of cancer the mean increase in the 5-year rate of survival was 13 percent points. Lung/bronchial and pancreatic cancer have the lowest 5-year survival rates and these rates have been the slowest to improve. One reason is that these types of cancer are very difficult to treat. The tumors are surrounded by critical structures, and move considerably when the patient breathes. Thus, knowledge of tumor location as it moves during breathing is integral to radiation therapy.

Fluoroscopy and computed tomography (CT) scans are typically used in the pre-treatment planning phase to discern the location of the tumor. Fluoroscopy is an imaging technique in which X-rays continually strike a fluorescent panel that is coupled to a video monitor. In a fluoroscopic image tumors lack sufficient contrast with the surrounding tissue, so in preparation for treatment metal clips are often implanted around the tumor. Since these clips are radio-opaque they are visible in both fluoroscopic images and CT scans. In fluoroscopy, they provide a way to observe the tumor as it changes position due to various

rigid and non-rigid body movements.

During the treatment phase patients receive a dose of high-energy X-ray from a linear accelerator (linac), an example of which (the Integrated Radiotherapy Imaging System (IRIS) of Massachusetts General Hospital) is shown in Fig. 1.1. The radiation beam is emitted by a part of the linac called the gantry which has the ability to rotate around the patient, and the patient lies on a movable couch underneath the gantry. For lung and abdominal tumors, radiotherapy is particularly complicated by tumor motion due to patient breathing. In order to compensate for this, the clinical target volume (CTV), the tumor itself, is often expanded by a margin to form the planning target volume (PTV), the volume of the patient which will be irradiated. This is done so that the tumor will receive sufficient dose but it leads to undesirable radiation of healthy tissue surrounding the tumor.



Figure 1.1: Example of a linear accelerator. Photo credit: Massachusetts General Hospital.

Usually the linac is accompanied by a fluoroscopic imaging system for pre-treatment

simulation. When this is the case it is usually a single-panel imager, meaning an image sequence can only be taken from one viewpoint at a time. Some linac machines have multiple-panel imaging systems (up to four panels). With these machines it is usually the case that the imaging system can be used not only for pre-treatment simulation, but for online imaging during treatment as well. The linac shown in Fig. 1.1 is one such machine, and Fig. 1.2 shows a diagram of the setup of a machine with a two-panel imaging system. Additionally, linac machines are also outfitted with a laser alignment system to reposition the patients in the same way during all sessions. Although setup error is a possibility, we will assume for the purposes of this thesis that the patient is aligned in the linac machine in the same way during every session.




Figure 1.2: Diagram of the front (on left) and side (on right) views of the setup of a linac machine with a two-panel imaging system. Diagram credit: Massachusetts General Hospital.

The fluoroscopic imaging system mounted on the linac machine can aid in more precise irradiation of the patient's tumor in several ways. It can be used for pre-treatment

simulation by radiation oncologists to visually observe the location of the patient's tumor. This information can be obtained while the patient is holding his or her breath to determine parameters for breath-hold techniques [14, 39, 34, 27]. These techniques assume that the patient's tumor is stationary when the patient is not breathing, and that when the patient is at a certain level of inspiration the tumor location will be consistently reproducible. During treatment the patient is brought to the same level of inspiration and the tumor is irradiated as if it were static. However, patients often cannot hold their breath due to the severity of their illness.

The fluoroscopic imaging system can also be used in pre-treatment simulation to observe how the motion of the patient's tumor moves with respect to some external indicator of the breathing signal, such as lung air flow or a physical marker placed on the patient's abdomen [23, 20, 36, 26, 8, 37]. A correlation between the motion of the tumor and the external indicator can be established, and then used during treatment to infer tumor position from the position of the external indicator. In gated radiation therapy, when the tumor is determined by observation of the external indicator to be within a specified range of positions, the radiation beam is turned on. When the tumor moves out of this range the beam is turned off. In gated therapy treatment time becomes an issue because the radiation beam is not always on, which could be considered an inefficient use of the machine, and radiation oncologists would like to speed up treatment in order to attend to as many patients as possible. Additionally, patient breathing varies over time and this can cause the correlation between the motion of the external marker motion and the motion of

the tumor, which gated therapy relies on, to break down.

The linac machines which have fluoroscopic imaging systems that can be used during treatment allow the tumor to be observed more directly while the patient is being irradiated [16, 25, 31, 32, 33, 30, 26]. However, these machines are costly and very few exist.

The problem is then to devise a way to recover the position of the tumor volume while the patient is breathing during treatment as accurately as possible so that the PTV can be reduced and more healthy tissue can be spared. This needs to be accomplished for scenarios with varying technological capabilities. The contribution of this thesis is to address the problem of tumor motion during radiation therapy in a rigorous manner by (1) formulating and analyzing a method for the tracking of fiducial markers in fluoroscopy in 2D, (2) formulating and analyzing a method to model the motion of the fiducials in 3D when only 2D data is available, (3) analyzing the correlation of the motion of internal and external fiducial markers, and (4) proposing a method to model motion of the tumor volume, the CTV, itself.

Chapter 2

Related Work

The methods employed in the proposed system draw from a diverse body of related work in the fields of medical imaging, radiation therapy, computer vision, computer graphics and computer science and contribute to these areas. This chapter highlights the most relevant work in each of those areas.

2.1 Radiation Therapy

There are many different techniques in radiation therapy which seek to address the problems that tumor motion can cause during treatment. Radiation oncologists want to minimize the amount of radiation delivered to healthy tissue around the tumor, while maximizing the radiation the tumor itself receives. The latter of these two takes precedence, and in the case of tumors which move because of patient breathing the simplest solution is to expand the PTV to attempt to compensate for the motion. Essentially, a large enough volume of the patient is irradiated so that the tumor receives enough dose even if it is moving. But the greater the range of tumor motion, the more this naive method, which we will call the standard approach, affects healthy tissue. Clearly, it would be better to have a way to deliver radiation only to the tumor. Intensity-modulated radiotherapy (IMRT) is a technique used to deliver radiation conformally to complex 3D target volumes while

sparing surrounding critical structures. However, Gierga et al. [11] have shown that tumor motion in the abdomen could significantly degrade the planned dose-volume histogram (DVH) of an IMRT plan. Thus, both objectives will be compromised. The following methods have been developed with the goal of fulfilling both objectives in mind.

2.1.1 Breath-hold Techniques

Breath-hold techniques seek to eliminate tumor motion entirely. Deep Inspiration Breath-Hold (DIBH) methods [14, 27] require the patient to achieve a level of deep inspiration which is consistently reproducible. This is done through verbal coaching and the use of a spirometer to monitor the patient's level of respiration. Active Breathing Control (ABC) techniques [39, 34], on the other hand, immobilize patient breathing at a certain point in the respiratory cycle via an occlusion valve.

By stopping the respiration at a reproducible level, the movement of the tumor is also halted. The patients can be imaged in a pre-treatment phase under the breath-hold conditions. Then, during treatment the position of the tumor will be known with greater accuracy and the PTV can be reduced. Breath-hold techniques have shown to help decrease the amount of dose delivered to healthy tissue. However, the patients undergoing this treatment are often very ill and their ability to hold their breath is greatly diminished. Thus, it has been shown that these techniques are not always feasible for the patient [18]. Additionally, these techniques require a longer treatment time per patient than the standard approach because the radiation beam is not always on.

2.1.2 Gated Therapy

Gating methods allow patients to breathe while receiving radiation, and seek to compensate for the motion by activating the radiation beam only when the tumor is at a predetermined position. This position can be detected in several ways. One method involves indirectly detecting the position by monitoring external markers. The external markers have taken the form of an infrared light-emitting diode [23] and now more commonly an infrared reflective marker [20, 36, 8, 37] which the commercially available Varian Real Time Position Management (RPM) system makes use of. All of these studies analyzed the relationship between the motion of the external marker and the motion of the diaphragm in fluoroscopy and found a strong correlation. Vedam et al. [37] found that based on the respiration signal produced by the RPM the motion of the diaphragm could be predicted within 1 mm.

In addition to infrared light-emitting diodes Ozhasoglu and Murphy [26] also investigated other external indicators of the breathing signal, namely lead fiducials and lung air flow, to indirectly determine tumor position. Their study analyzed the correlation between these signals and the motion of internal fiducials implanted in the tumors.

It has been noted that a phase shift between external and internal motion may occur under free breathing conditions [7, 20, 36, 37, 26]. Kini et al. [17] showed that improving the reproducibility of the amplitude and frequency of breathing via coaching is possible and that this might increase the accuracy of gated therapy.

Position can also be directly detected by means of online fluoroscopic imaging of

internal fiducial markers [32, 33, 30, 26]. Online imaging removes the problems that relying on a correlation between internal and external signals can cause. However, the technology required to image during treatment is only found in a few hospitals and the imaging results in the undesired effect of an increase in the amount of radiation the patient receives.

Gated therapy in general is more time consuming than the standard approach because the radiation beam is not always on. Due to large patient volume at treatment clinics, there is pressure to decrease the time of a treatment session, but this comes at the cost of increasing the size of the gating window (the time when the beam is on) and thus more healthy tissue is irradiated. This trade-off is difficult to balance and becomes even more problematic in the typically longer sessions of IMRT.

2.1.3 Correlation

In most of the indirect gating methods correlation between external and internal motion was assumed. Although studies have attested to the correlation of an external marker with 2D internal motion of the diaphragm [23, 20, 8, 37] and with 2D internal motion of fiducial markers placed around or in the tumor [26, 12, 10], Ozhasoglu and Murphy [26] pointed out many flaws in the assumptions and analysis techniques of these studies. By imaging patients for 1-10 minutes they found complexities of the breathing pattern which would not always be noticed when the imaging was done over the shorter, standard time frame of 30-60 seconds. Vedam et al. [36] described a method to analyze and compensate for

one of these problems, phase shift. However, the study done by Ozhasoglu and Murphy showed that the phase shifts are usually transient or time-changing. This is because the elastically coupled components (e.g., organs, tissues, internal and external markers) of a mechanical system which are subjected to a periodic driving force (e.g., breathing) and are experiencing simple harmonic motion, cannot sustain a fixed phase difference. Thus, a method which does not adapt to changes in the breathing pattern, e.g., [36], could actually lead to serious errors in estimating marker position.

Ozhasoglu and Murphy pointed out several other often overlooked issues. The placement of the external marker is very important as correlation with either the chest or the abdomen yielded different results, and one cannot choose *a priori* which position would yield the highest correlation or if any correlation exists at all. Thus, longer periods of observation are necessary. However, more observational imaging means more radiation is being delivered to the patient and because of this there is an extreme reluctance to perform such extended studies.

They also showed that tumor trajectories cannot be assumed to be simple. This necessitates measurement of the motion of the tumor in three dimensions. In practice, this sort of information is typically sparse. Since tumors cannot be easily tracked in fluoroscopy the internal fiducials are often tracked and taken to represent the motion of the tumor.

In summary, Ozhasoglu and Murphy found three significant problems with the current state of gated therapy which utilizes external markers:

1. The observations of breathing motion that are used to set the gating con-

ditions tend to last 1 min or less, which is not long enough to characterize respiration or establish that equilibrium conditions exist;

2. The tumor is never observed directly;

3. The gate-triggering algorithm does not accommodate non-stationary aspects of breathing motion, as for example time dependence of phase or amplitude differences.

2.1.4 Image-guided Therapy

Image-guided radiation therapy seeks to address the problems of gating by moving the radiation beam in synchronization with the tumor. It allows for shorter treatment session times and, when coupled with online imaging, avoids the correlation problem. The method was first introduced in robotic radiosurgery [1, 28], and later adopted for motion-adaptive radiotherapy [16, 25, 31].

Motion-adaptive radiotherapy involves synchronously adapting the radiation beam using a dynamic multileaf collimator (DMLC) to follow the motion of the tumor. Keall et al. [16] showed that the dose delivered by motion adaptive X-ray therapy (MAX-T) to a moving target was equivalent to the dose a static target received from a static beam, and so the concept is feasible. They propose that the leaf sequence be based upon the motion of an external breathing signal. However, as Ozhasoglu and Murphy pointed out, this is more complicated than it seems. Thus, the benefits of using image-guided radiation might be lost due to the correlation inaccuracies discussed previously. Keall et al. also point out

the fact that studies usually assume that only the tumor is moving. In reality, however, the surrounding organs and tissues move as well. This is very important because in planning treatment it is critical to avoid delivering radiation to certain structures, and if the position of those structures changes they could be unintentionally irradiated.

Neicu et al. [25] proposed a system called synchronized moving aperture radiation therapy (SMART). During treatment planning the tumor motion in three dimensions is measured and used to derive an average tumor trajectory (ATT). In the treatment phase the tumor motion is monitored (either directly or indirectly) and the beam moves according to the ATT unless the tumor's trajectory deviates from the ATT, in which case the beam is turned off. If the tumor's trajectory resynchronizes with the ATT, the beam is turned on again. No attempt is made to change the leaf sequence online, and so this method is in a way a combination of image-guided therapy and a more sophisticated version of gating. They acknowledge that this system is then reliant on the regularity of the patient's breathing pattern in order to be usefully efficient. Again, according to Ozhasoglu and Murphy this is not a reasonable assumption. Furthermore, as with the system of Keall et al. [16], using an external marker might negate the benefits of image-guided therapy.

Sharp et al. [31] introduce a method which does not rely so strongly on such an assumption. They propose to image the patient in three dimensions during treatment and to use this information along with prediction to guide the beam. Because there is a mechanical latency between imaging time and the movement of the beam, Neicu et al. [25] relied on a predetermined form of predictability in the tumor's trajectory. Sharp et al. [31]

suggest doing the prediction during each session, and so the patient's breathing need not be perfectly regular from day to day. This technique could be thought of as tailoring the ATT to a breathing period of a few minutes. Sharp et al. [31] acknowledge that within a given treatment session they do not update the prediction parameters, and that the longer the prediction is relied upon the more error will be introduced due to the non-stationary nature of breathing. Overall their system worked well when the imaging rate was below 10 Hz and the latency of the beam was greater than 33 ms. The low imaging rate helps to reduce unwanted radiation delivered to the patient. However, with all techniques that require online imaging of the patient in three dimensions it is important to keep in mind that this technology is limited to only a few hospitals around the world.

2.2 Tracking

Direct gating and image-guided methods which use online tracking have additional weaknesses. The robustness of tracking software, which is of paramount importance, is often insufficiently addressed. Furthermore, these systems often rely on only one internal marker to represent 3D tumor motion, but this is not adequate enough to accurately capture tumor motion and cannot account for tumor deformation. Tracking is a relatively new concept in radiation therapy, but it has been extensively studied in other communities. In this section we will discuss how radiation therapy has utilized tracking thus far and highlight a few relevant examples of tracking in computer vision and medical imaging.

2.2.1 Radiation Therapy

Several groups developed software to track the motion of the diaphragm in fluoroscopy [20, 8, 37]. Ford et al. [8] only examined the motion of the diaphragm in the z -direction. Although it is true that the motion in the z -direction is typically the greatest, the motion in the other two directions is often substantial and therefore should not be overlooked. Tracking the motion of the diaphragm is clearly better than having no information about how the internal structures of the patient are moving. However, the diaphragm is not a perfect representative of tumor motion, and the tumor must be studied directly according to Ozhasoglu and Murphy [26]. Additionally, the accuracy of the diaphragm trackers used in these studies was not reported. Furthermore, Berbeco et al. [2] concluded that 3D imaging is necessary because of the irregularity and three dimensionality of tumor trajectories, and because the diaphragm is not tracked in 3D it cannot be an adequate surrogate for tumor motion.

Shirato et al. [32, 33] developed a system called the real-time tumor tracking radiotherapy (RTRT) system to track the 3D motion of a 2 mm gold marker implanted in or near a patient's tumor at a rate of 30 Hz. They report an accuracy of 1.5 mm. It should be noted that this is the same system used by Seppenwolde et al. [30], Neicu et al. [25], and Sharp et al. [31]. The ability to track the marker in 3D is beneficial, however, only one marker is used to represent tumor motion, and a single marker cannot adequately represent tumor rotation or deformation. To uniquely determine the pose of a 3D object, at least three points are necessary. Furthermore, determining the relative pose between subsequent

frames is problematic because the tumor is not a rigid object. The fiducials often do not move as a rigid body, and according to Murphy [24], using a fourth fiducial so that the pose problem is overdetermined greatly improves precision. This is because if only one fiducial moves non-rigidly with respect to the group it can be ignored.

In general, the main problem with these methods is that the tracked internal markers are still only somewhat representative of tumor motion, and do not reflect the actual motion of the tumor.

2.2.2 Computer Vision and Medical Imaging

There has been significant work in the computer vision community on tracking [9], and we highlight only one method here, the Kalman filter [38]. As many tracking methods, the Kalman filter uses a type of feedback control to estimate the tracking process. The filter continually estimates the process state at a given point in time and then obtains measurements to use as feedback for the estimation. The discrete Kalman filter is used when the relationship between the process and the measurements, or the process itself, is linear. However, if the breathing motion is modeled as a sinusoid where amplitude, frequency, and the DC component must be estimated simultaneously the process is non-linear. The extended Kalman filter (EKF) can address such non-linear processes. Choosing the parameters which govern the filter is problematic for this application for two reasons. First, although breathing is roughly sinusoidal it is non-stationary and can possibly change drastically if the patient inhales sharply. In this case a filter that is based on a fixed

amplitude, frequency, and DC component may not be beneficial. Secondly, the time a Kalman filter needs to choose the parameters accurately would often be unacceptably long, due to the excess radiation that would be delivered to the patient.

There exist other methods from the field of signal processing to estimate the parameters of a sine wave. Fourier analysis is one technique which can be used to find the dominant frequency of a signal. The results of this method are dependent on the frequency at which the data is sampled as well as the number of samples. Because the data we are working with is sampled at a relatively low frequency, 30 Hz, for a short amount of time, 5-10 s, the set of frequencies which the method can identify is typically more limited than desired. Furthermore, radiation oncologists place more weight on the peaks and troughs of the breathing cycle and consider these points to be more stable. Therefore it would be desirable to have a method that takes this practice into account. This is why we developed a method which relies more heavily on the maxima and minima of the breathing cycle than Fourier analysis to estimate the frequency.

There has also been much work done in the medical imaging community on tracking, in particular in the areas of fMRI and cardiac motion (e.g., [13]) and blood flow analysis (e.g., [29]). This work is usually done with higher resolution imaging techniques and also often involves a large amount of prediction. Additionally, these methods often incorporate intense computation when estimating the motion and so with current computing technology they cannot feasibly run in real time.

2.3 3D Modeling, Transformation, and Deformation

Because tracking tumors in 3D is a relatively new area of radiation therapy it is important to study how 3D information is used in other areas for which 3D modeling methods have been developed. There has been ample work in the computer graphics and the computer vision communities with respect to 3D representations of objects [9, 15]. Bookstein [5] proposed to use pairs of thin-plate splines to model biological shape change as deformation. A thin-plate spline is the 2D analog of a cubic spline. The pair of splines acts as an interpolation map which relates two sets of landmark points. According to Bookstein, “The spline maps decompose, in the same way as the spline surfaces, into a linear part (an affine transformation) together with the superposition of principal warps, which are geometrically independent, affine-free deformations of progressively smaller geometrical scales.” He gives examples how this can be used for biological data, and explains how it can apply to 3D data.

Another method for dynamic deformable models was proposed by Metaxas [22]. This method involves using solid primitives which are allowed to deform kinematically as well as undergo global deformations to produce realistic animations. These animations display physically correct behavior.

2.3.1 Medical Imaging

Much work has been done in the medical imaging community with respect to 3D models. For the most part this work concerns non-rigid registration between a hand segmented volume and other scans of the same organ which has undergone a change in position or shape. In this section we will discuss one recent example. McLeish et al. [21] studied the effects, namely motion and deformation, that respiration had on the heart via magnetic resonance (MR) images. They took several scans of each patient at different levels of inspiration. In one of these scans they segmented the heart volume by hand. Then an automated procedure was used to calculate a free-form deformation by iteratively altering a uniform array of B-spline control points which were each 10 mm apart. These deformation fields were also used to create a motion model. The chief use of this work to is to correct for the artifacts caused by heart motion during free breathing scans. Similar models have been developed for the liver (e.g., [4]), but the chief aim of these models is to help treatment planning for radiotherapy.

There has also been much work in registering 2D images (X-ray) to 3D images (MR/CT) [35, 19]. This is usually done for the purpose of image-guided therapy which allows surgeons to use the pre-operative scans (MR/CT) during surgery. Because the patient cannot be repositioned in the exact same way as when the pre-operative scan was taken, an X-ray is often taken during treatment and used to compute a transformation between the pre-operative and intra-operative coordinate systems. This allows the surgeon access to the high quality MR/CT data which would not otherwise be available during surgery.

2.4 Summary

There has been significant work in the radiation therapy community [14, 39, 1, 28, 34, 27, 32, 33, 23, 16, 20, 36, 8, 30, 26, 37, 25, 12, 10, 31] which attempts to address the problems that the motion of lung and abdominal tumors during treatment can cause. Breath-hold techniques [14, 39, 34, 27] seem to make the tumor reproducibly static, but they are often too difficult for the patient to endure. Studies of indirect gating methods [23, 20, 36, 8, 26, 37, 12, 10] have shown correlation between internal markers and external indicators of the breathing signal over short (30 s) periods of observation. However, these studies only examine internal marker motion in two dimensions. The motion of the internal markers must be studied in 3D and longer periods of observation are required. Gating and image-guided methods [1, 28, 32, 33, 16, 30, 26, 25, 31] which use online tracking give good estimates of internal marker position in 3D, but they are often too costly to implement and can expose patients to excess radiation. Furthermore, these methods often rely on only one internal marker to represent tumor motion [16, 25, 31].

Tracking has been more extensively studied in the computer vision community. Although the extended Kalman filter has the capability to estimate non-linear processes like breathing, it might not be beneficial in the case where breathing is modeled as a sinusoid and the parameters for amplitude, frequency and the DC component are fixed.

There has also been much work in the area of 3D modeling [5, 22, 4, 21, 35, 19], but it has yet to be used in relation to internal marker tracking in radiation therapy. Additionally,

it is not feasible to apply some of the methods to the problems discussed in this thesis because the methods can require more information than is available from the fluoroscopy.

Chapter 3

Materials and Methods

Fluoroscopic imaging provides a two-dimensional (2D) projection of the density values of the imaged body. The surgically implanted internal markers (clips) can be detected and tracked in the fluoroscopic image sequences since metal is radio-opaque and has a higher density than the surrounding tissue.

During treatment planning, radiation oncologists typically request fluoroscopy from two views, the Anterior-Posterior (taken along the y -axis) and the Lateral (taken along the x -axis) views. Because these views are orthogonal, we can combine the tracking data to recover how the positions of the clips change in three dimensions (3D) over time. The setup of the imaging system is such that the Anterior-Posterior and Lateral views share the z -axis of the respective images, which is the Cranio-Caudal axis of the patient, and the isocenter of the patient is at the center of the image (Fig. 3.1).

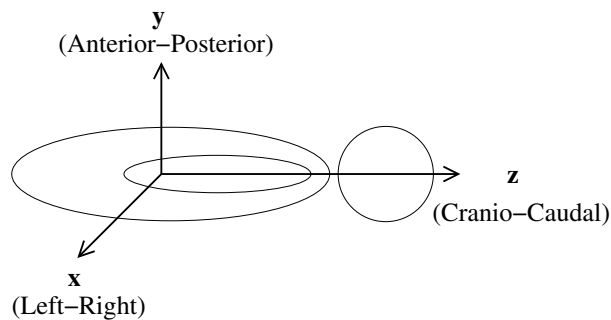


Figure 3.1: Illustration of the coordinate system of fluoroscopic views. The patient is lying on his or her back with the head towards the positive z -coordinate axis.

During treatment planning, radiation oncologists also request a CT scan. The tumor is then contoured manually on the CT by the radiation oncologists. This expert knowledge can be used to our advantage by first determining a relationship between the clips, which are visible in the scan, and the contour. Each clip moves rigidly but the collection of clips generally moves non-rigidly with respect to itself. If we have a sufficient number of clips, we can infer from the rigid motion of the clips in 3D the non-rigid motion of the tumor volume in 3D by using the contour and clips to generate a model of the motion.

The motion of the clips observed in the fluoroscopy along with the CT scan tumor contour are the inputs from which a model of the 3D motion of the patient's tumor can be built for use during treatment. From the radiation oncologists' standpoint the overview of the system is shown in Fig. 3.2.




Figure 3.2: High-level overview of goals of method from a radiation oncologist's point of view.

One of the goals of this thesis is to develop a system that will work in hospitals with different technological capabilities. The main difference in technology is the type of fluoroscopic imaging system which is mounted on the linac machine. If the imaging system has multiple panels which can be used simultaneously, then the location of the clips in 3D can be directly estimated from the fluoroscopy. This information can then be immediately

coupled with the tumor contour to produce a model of 3D tumor volume motion. This model can then be either correlated with external markers for use in gated therapy, or during image-guided radiation when online imaging is available. Most linac machines, however, have only single-panel imaging systems which cannot be used during treatment. This means that the 3D position of the internal markers cannot be directly estimated. Instead, we propose to estimate a model of the 3D clip motion from two sequentially obtained orthogonal fluoroscopic image sequences. Although an internal marker's motion in the z -direction is not exactly the same in both image sequences, it is typically similar. This is because each of the orthogonal fluoroscopic image sequences (Anterior-Posterior and Lateral) are taken only minutes apart, during which the patient's breathing and anatomy do not change drastically.

When only sequentially obtained fluoroscopic image sequences are available, certain assumptions must be made in order to build a model. Generally, breathing is well described by a sinusoid for short sequences of time when the patient is breathing regularly. Because the modeling algorithm will be used on sequences of approximately 30 s we believe this is a reasonable initial assumption, and so we estimate the parameters of a sine wave model based on the two sequentially obtained fluoroscopic image sequences. The objective is to model the average and maximum range of 3D clip motion. Once this is done, the tumor contour can be used to build a final model of the 3D tumor volume motion. This can then be correlated with external markers for use in gated therapy.

The model building section of the system can therefore take two paths. Either the 3D

position of the internal markers can be directly estimated if a multiple-panel fluoroscopic imaging system is available, or the 3D position of the internal markers must first itself be modeled if only a single-panel imager is used. The overall view of the steps which the system takes is shown in Fig. 3.3.

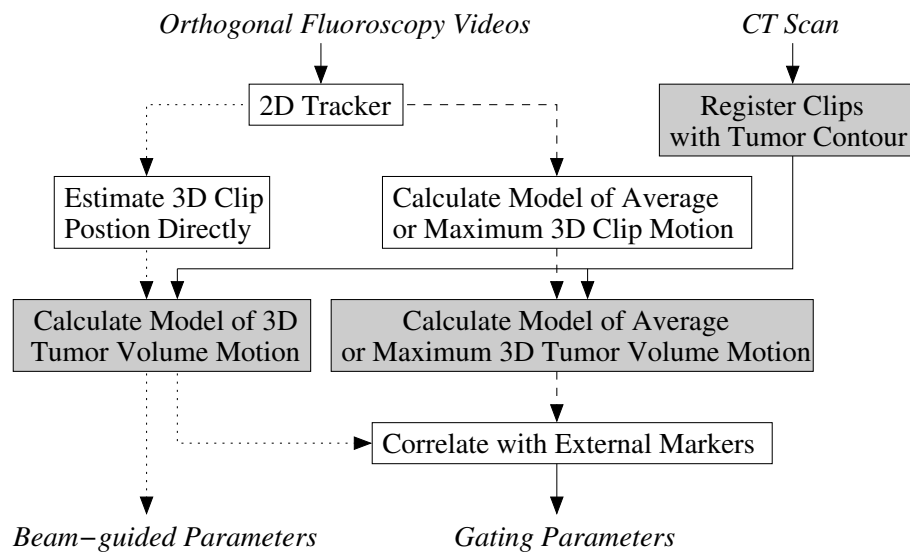


Figure 3.3: Detailed overview of method. Dotted lines indicate steps taken when a multiple-panel fluoroscopic imaging system is available. Dashed lines indicate steps taken when only a single-panel imaging system can be used. Solid lines are steps taken in both cases. Gray boxes have not yet been implemented.

3.1 2D Clip Tracking

Tracking the motion of the clips implanted in the patient can be difficult. Although the metal which the clips are made of has a higher density than the structures and tissues of the body, the clips are not always easily visible. This is because the surrounding tissue may also appear dark at times or contain edges due to high-density bone structures such as the

spine (Fig. 3.4) and the images themselves can become dark during inhalation (Fig. 3.5).




Figure 3.4: Subsection of a fluoroscopic image containing a section of the spine with the vertebrae edges indicated by white arrows and a surgical clip marked by a white rectangle. Clip and spinal section have similar intensity values.

Below the real-time algorithm for tracking the 2D motion of a clip in fluoroscopy which is based on previous work done by our group [3] is described. The following procedure is performed in parallel for all of the clips in a given fluoroscopic image sequence.

3.1.1 Initialization

The tracking method is initiated by manual selection of a rectangular region r containing a clip in an initial fluoroscopic image I . In order to find a minimal rectangle containing each clip, the largest “dark” connected component in each region is found by first binarizing the image according to an automatically computed ptile threshold for the region. A ptile threshold is one in which a threshold is chosen such that $p\%$ of the image area has grayscale values less than the threshold, and the rest of the image area has grayscale values greater than the threshold. The percentage p is fixed, and the threshold is computed online.




Figure 3.5: Two subsections of fluoroscopic images taken during inhalation and exhalation, respectively, clips marked by white rectangles. The image on the right is much darker, and the clips have changed position relative to each other.

When manually selecting r users are instructed to make the rectangular region as small as possible. Given this instruction and the average size of a clip (5 mm), it was found that the algorithm performed well when p was set at 20%. The rectangular region is used as a $(w \times h)$ grayscale template T of the clip. The location of the template in I provides the starting coordinates for tracking the clip in subsequent fluoroscopic frames.

3.1.2 Tracking Algorithm

The normalized correlation coefficient is used to find the position of the clip in subsequent image frames. The value of the normalized correlation coefficient at position (x', y') is

$$R_{I,T}(x', y') = \frac{\sum_{x,y} (I(x' + x, y' + y) - \bar{I})(T(x, y) - \bar{T})}{\sqrt{\sum_{x,y} (I(x' + x, y' + y) - \bar{I})^2 \sum_{x,y} (T(x, y) - \bar{T})^2}}, \quad (3.1)$$

where \bar{T} and \bar{I} are the respective mean intensities within the template and image window. The tracking algorithm searches for the best match of the clip template T with a region of the image I . This is done by shifting the template T through the image to various points (x', y') and correlating it with each $(w \times h)$ sub-image of I . The location (x', y') which maximizes $R_{I,T}$ is taken to be the new clip location.

Searching over all positions (x', y') in I is computationally expensive. Because we know that the clips do not move much from frame to frame, it is possible to restrict the size of the sub-region of I which will be searched. The apparent velocity of the clip's movement in the image is calculated to predict the clip location in the next frame. Velocity $(u, v) = (\frac{dx}{dt}, \frac{dy}{dt})$ is approximated in terms of the rate of change in x and y from the past frame to the current frame: $(u_{t+1}, v_{t+1}) = (x_t - x_{t-1}, y_t - y_{t-1})$. We assume that the velocity is constant for two consecutive frame pairs, and use (u_{t+1}, v_{t+1}) as an offset from the current position (x_t, y_t) to determine the center $(x_t + u_{t+1}, y_t + v_{t+1})$ of a 5×5 region of interest (ROI) to search over in the subsequent frame, $t + 1$. When the clip reaches a minimum or maximum of the breathing cycle in its trajectory and reverses its direction, this assumption does not negatively affect the tracking. Because the clip typically does not move by more than one pixel per frame, the clip will remain within the 5×5 search window. By not weighting the ROI with probabilities as a Kalman filter would, all locations within the ROI are equally valid and so the correct position will not be negatively weighted because it is in seemingly violation of the model. Within a frame the velocity will be reversed to reflect the new trajectory.

3.2 Estimated Model of 3D Clip Motion

The 2D tracking software produces a set of time-indexed 2D coordinates for each fluoroscopic view. The tracker output of two sequentially obtained fluoroscopic image sequences will then give us four trajectories: $x(t)$, $y(t)$ and two trajectories for $z(t)$. In general the algorithm works by first processing the two $z(t)$ trajectories. To calculate a model for one of these 1D trajectories, we compute the DC component, amplitude and frequency of a sine wave which models this motion. In order to do that, we need to find the maximum and minimum of each breathing cycle. After the model is computed for both $z(t)$ trajectories, the models of other two trajectories $x(t)$ and $y(t)$ are calculated. Finally, the results of the calculations of the 1D models for all four trajectories are combined to form a final 3D model of the average or maximum range of motion of the clip.

We present the description of the algorithm in three parts (1) finding the maximum and minimum of each breathing cycle for a given trajectory (2) computing the DC component, amplitude and frequency of a sine wave which is a 1D model a given trajectory (3) combining the four 1D models into a 3D model.

3.2.1 Finding the Maximum and Minimum of Each Breathing Cycle

It is difficult to determine the maximum and minimum of a breathing cycle for several reasons. First, there is often no change of position of the clip in subsequent frames of the time-indexed trajectory because of the rests a person takes at full inhale and exhale. Sec-

only, the clips move slowly and the motion appears as either no change or a small change in the image. These conditions result in a step-like discretizing effect in the trajectory, and makes the moments in time when the maximum and minimum of each breathing cycle occur difficult to identify.

After a smoothing operation the maximum and minimum of each breathing cycle can be identified uniquely even if their positions in time were not originally unique (Fig. 3.6). The trajectory is smoothed using a 1D Gaussian kernel with a support of $j=10$ frames (or 1/3 second) and a standard deviation $\sigma=1$ frame. This kernel G is generated by constructing an array of $2j + 1$ entries, whose i th value is:

$$G_i = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(i-j-1)^2}{2\sigma^2}\right). \quad (3.2)$$

Because a typical human breathing cycle is about 150 frames (or 5 seconds), we chose a support of 10 frames to adequately smooth over the plateaus which result from the stepping effect in the trajectory.

After smoothing, a modified sliding-window technique is applied to find the maximum and minimum of each breathing cycle of the smoothed trajectory $z'(t)$ as follows. The global maximum g_{max} and minimum g_{min} of the entire data set are determined. A threshold τ is set to be equal to 110% of the difference between g_{max} and g_{min} :

$$\tau = 110\% \cdot (g_{max} - g_{min}). \quad (3.3)$$

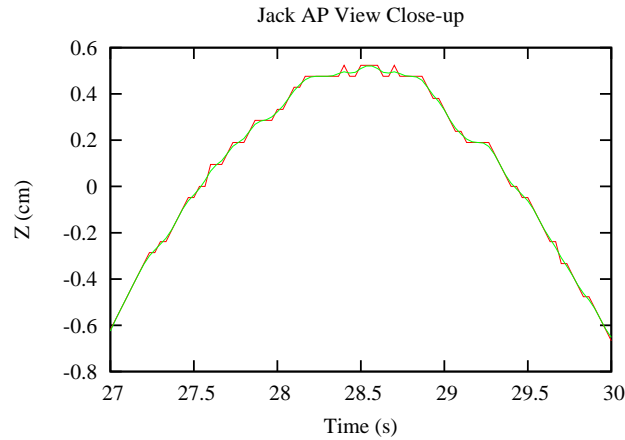


Figure 3.6: A close-up view of one of the peaks of the breathing cycle. The original trajectory produced by the 2D tracker (red) has several instances of the maximum, but after smoothing (green) only one instance of the maximum exists.

This threshold represents an aggregate distance counter (Fig. 3.7). By making the threshold slightly greater than the largest amount of distance between g_{max} and g_{min} we ensure that we are thoroughly searching the trajectory because the distance searched is greater than the distance traversed in one half cycle. Also, we are providing room for some jitter in the data.

The initial point of the data set could be anywhere in the breathing cycle, so the algorithm first determines if the patient is in an inspiration or expiration phase by comparing the first point of the trajectory $z'(0)$ with a point 20 frames later in the trajectory $z'(20)$. We choose 20 frames here because it is far enough to overcome any jitter in the data. If $z'(20) > z'(0)$ then the patient is in an inspiration phase and the algorithm will search for a maximum. Otherwise the patient is in an expiration phase and the algorithm will search for a minimum. The algorithm steps along the trajectory and keeps track of the distance d (in the z dimension) travelled, by aggregating the difference between subsequent points $z'(t)$

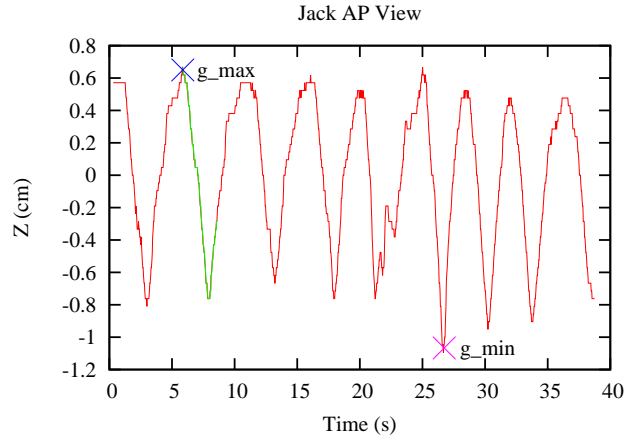


Figure 3.7: An illustration of how the modified sliding-window technique works with the one dimensional trajectory is shown in red. In an initial phase g_{max} and g_{min} are found, and the threshold τ is computed. In the second phase the maxima and minima of the breathing cycles are found. Here the aggregate distance count (green) begins at the point where the most recently processed local maximum was found and ends when the threshold τ is reached. It represents the amount of distance searched over for the next local minimum.

and $z'(t + 1)$:

$$d_k = d_{k-1} + |z'(t + 1) - z'(t)|. \quad (3.4)$$

When searching for the maximum of a breathing cycle, the algorithm records the time index t_c , the value of the point $z'(t_c)$, and the distance travelled d_{t_c} when its value exceeds the last recorded value $z'(t'_c)$ of a candidate for the maximum. When $d > T$ the information about the last recorded candidate for maximum is reported. The distance is updated as follows: $d_k = d_k - d_{t_c}$. This has the effect of resetting the aggregated distance counter to the time the maximum occurred. Then, the algorithm begins looking for a minimum. The procedure for this is identical, except that the algorithm records the time index and value of a point when its value is less than the last recorded value.

This procedure yields the time indices of the maxima and minima l_1, \dots, l_m , where m is the total number maxima and minima. These time indices can then be used to index into the original unsmoothed trajectory $z(t)$.

3.2.2 Calculating the Parameters of a 1D Model

After the maxima and minima of the breathing cycles have been identified, the algorithm calculates the parameters of a sine wave which models the data. We can calculate a model of both the average motion and the worst case (maximum) motion. The only difference in these models will be their amplitude; the DC component and frequency will remain the same. Assuming that a maximum is found first for the sake of notation, the average amplitude a_{avg} is given by half the difference between the average maximum and the average minimum of the breathing cycles:

$$a_{avg} = \frac{1}{m} \left(\sum_{i=1}^m z(l_{2i-1}) - \sum_{i=1}^m z(l_{2i}) \right). \quad (3.5)$$

The maximum amplitude is given by:

$$a_{max} = \frac{g_{max} - g_{min}}{2}. \quad (3.6)$$

The period ρ is given by twice the average time of a half cycle:

$$\rho = \frac{2}{m-1} \sum_{i=2}^m (l_i - l_{i-1}), \quad (3.7)$$

The frequency f is the inverse of the period: $f = 1/\rho$. We can determine the DC component o of the z motion by finding the average value

$$o = \frac{1}{l_{m-1} - l_1} \sum_{i=l_1}^{l_{m-1}} z(i), \quad (3.8)$$

over all points in the trajectory, where S is the number of frames in the trajectory.

Finally, we can then define the average or maximum pattern of z motion as a sine wave given by:

$$o + a \sin(2\pi ft), \quad (3.9)$$

where a is either a_{avg} or a_{max} .

We assume that each clip follows a roughly linear trajectory, and so the maximum and minimum of each breathing cycle for the trajectories of the motion in the x -direction and the y -direction must occur at the same times as those of the motion in the z -direction. Thus, when calculating the models for the motion in the x -direction and the y -direction, we do not need to recalculate the maximum and minimum of each breathing cycle. It is sufficient to reuse the time indices calculated for the model of the motion in the z -direction, evaluate $x(t)$ and $y(t)$ at the those times, and perform the calculations outlined above. This also ensures that the model will remain in phase with itself.

3.2.3 Calculating the Parameters of the Final 3D Model

After the 1D models are computed for the four trajectories we have four sets of parameters: $\{o_x, a_x, f_x\}$, $\{o_y, a_y, f_y\}$, $\{o_{z1}, a_{z1}, f_{z1}\}$, $\{o_{z2}, a_{z2}, f_{z2}\}$. These parameters are processed together to calculate the final 3D model as follows. The corresponding parameters of the two models of the motion in the z -direction are averaged. The frequency parameter of the models of the motion in the x -direction and the y -direction are updated to use the average frequency calculated for the model of the motion in the z -direction. This results in a parametric equation which describe a model of the motion of the clip in 3D:

$$\mathbf{m}(t) = \mathbf{o} + \mathbf{a} \sin(2\pi \mathbf{f}t), \quad (3.10)$$

where $\mathbf{o} = (o_x, o_y, \frac{o_{z1}+o_{z2}}{2})$, $\mathbf{a} = (a_x, a_y, \frac{a_{z1}+a_{z2}}{2})$, and $\mathbf{f} = (\frac{f_{z1}+f_{z2}}{2}, \frac{f_{z1}+f_{z2}}{2}, \frac{f_{z1}+f_{z2}}{2})$.

Fig. 3.8 shows the range of motion in an actual fluoroscopic image that the estimated model of average 3D motion yields.

3.3 Fourier Analysis

To evaluate the modeling algorithm Fourier analysis was used to estimate the dominant frequency for each of the 1D trajectories. A discrete Fast Fourier Transform was used to do this:

$$F(k\Delta f) = \sum_{n=0}^{N_s-1} f(n\Delta t) e^{-i(2\pi k\Delta f)(n\Delta t)}, \quad (3.11)$$




Figure 3.8: Two subsections of fluoroscopic views of patient Jack (Anterior-Posterior on the left, Lateral on the right) taken during the peak of a breathing cycle (patient is in full inspiration). Clips are marked by rectangles. Clip 3 in the AP view and clip 1 in the Lateral view are the same clip. The average range of motion from the estimated 3D model is shown in red.

for $k = 0, 1, 2, \dots, N_s - 1$, where N_s is the number of discrete samples, T_s is the total sampling time, $\Delta t = T_s/N_s$, and $f_s = N/T$ is the sampling rate of the original signal. The dominant frequency was obtained by finding the frequency with the highest magnitude in the first $N_s/2$ samples of the spectrum (only the first $N_s/2$ frequencies are useful due to the Nyquist criterion). It should be noted that this implementation of the Fast Fourier Transform restricts N_s to powers of 2.

3.4 Correlation with External Markers

In previous work [12, 10], we showed that the 2D motion of internal markers correlates with the motion in the y -direction of a set of external markers. This study, as well as the previous one, differs from other work dealing with external markers [23, 20, 36, 8, 26, 37] in that instead of having only one external marker, we use a chain of beads, each of which is a single marker, placed on the patient's abdomen. This allows us to examine the external

motion at several different points on the abdomen of the patient simultaneously. Fig. 3.9 shows a fluoroscopic image containing both the internal and external markers.




Figure 3.9: Lateral fluoroscopic image of patient Nancy with three external markers, which are beads marked 3, 4, and 5, resting on the patient's abdomen (right), and two internal markers (clips) marked 1 and 2.

We have extended our previous work [12, 10] to examine how the 3D estimated average motion of a clip is correlated with the motion in the y -direction of an external marker. The correlation is calculated between the motion of a clip in each of the three dimensions and the motion of the external marker in the y -direction. The chain of beads results in many external markers, and for this study we chose three beads on the chain with which to correlate the internal marker. The correlation is found by comparing the positions of the external and internal marker at each point in time, and finding the regression line which fits that data.

Chapter 4

Results

Fluoroscopic videos were collected with a Varian Ximatron radiotherapy simulator with a resolution of 640×480 pixels with 8 bits per pixel at 30 frames per second. The CT scans were taken with a General Electric Lightspeed scanner at 140 kVp. The computer used to process the data was an Intel Xeon 1.7 GHz with 1 GB RAM and an ATI All-In-Wonder 9800 graphics card.

4.1 2D Clip Tracking

Examples of trajectories produced for four patients for both the Anterior-Posterior and Lateral views are shown in Fig. 4.1.

A ground truth study involving 5 patients was done. Ground truth was established for the motion of five clips, one per patient, using the Anterior-Posterior views. We manually recorded the centroid of each clip for every tenth frame of the image sequence, or every $1/3$ second. The ground truth positions were then compared with the trajectory produced by the 2D tracker. Table 4.1 shows the averages and standard deviations of the root mean squared (RMS) and the maximum error in centroid position of each clip's motion in the x -direction and the z -direction. On average, the mean and standard deviation of the RMS error of detecting each of the 10 motion trajectories were under 0.5 mm (≈ 1 pixel).




Figure 4.1: 2D motion of clips of four patients in Anterior-Posterior and Lateral Fluoroscopy.

Table 4.1: Clip Localization Error

Patient	Direction of Motion	Root Mean Squared Error (RMS)		Max Error (mm)
		Mean (mm)	Std. Dev. (mm)	
Alice AP View	<i>x</i> -direction	0.61	0.58	2.31
Alice AP View	<i>z</i> -direction	1.21	0.58	2.00
Doug AP View	<i>x</i> -direction	0.40	0.15	0.74
Doug AP View	<i>z</i> -direction	0.39	0.38	1.50
Eve AP View	<i>x</i> -direction	0.25	0.25	0.74
Eve AP View	<i>z</i> -direction	0.23	0.26	1.00
Frank AP View	<i>x</i> -direction	0.25	0.23	0.74
Frank AP View	<i>z</i> -direction	0.41	0.33	1.00
Gary AP View	<i>x</i> -direction	0.48	0.32	1.48
Gary AP View	<i>z</i> -direction	0.47	0.24	1.00
Average		0.47	0.33	1.25

4.2 Estimated Model of 3D Clip Motion

This study involved 2 patients. In order to evaluate the accuracy of the estimated model of average 3D motion we compared the trajectories of average motion with the corresponding four original trajectories obtained from the 2D tracker (see Fig. 4.2). The 1D estimated model parameters for the two patients are shown in Figs. 4.2 and 4.4 and the final 3D estimated model parameters are shown in Figs. 4.3 and 4.5.

Table 4.6 shows the averages and standard deviations of the root mean squared (RMS) and the maximum error in centroid position of each clip's motion in the *x*-direction, the *y*-direction, and the *z*-direction. On average, the mean and standard deviation of the RMS error of the model were under $0.3 \text{ cm} = 3.0 \text{ mm}$ (≈ 7 pixels). The period ρ is included for the sake of comparison.

Table 4.2: 1D Estimated Model Parameters: Patient Jack

	x	$z1$	y	$z2$
o (mm)	4.16	0.02	4.17	0.35
a (mm)	-0.18	0.69	-0.17	0.53
ρ (s)	4.46	4.46	4.29	4.29
f (Hz)	0.22	0.22	0.23	0.23

Table 4.3: 3D Estimated Model Parameters: Patient Jack

	x	y	z
o (mm)	4.16	4.17	0.18
a (mm)	-0.18	-0.17	0.61
ρ (s)	4.38	4.38	4.38
f (Hz)	0.23	0.23	0.23

Table 4.4: 1D Estimated Model Parameters: Patient Nancy

	x	$z1$	y	$z2$
o (mm)	-7.36	-1.36	-1.72	-1.55
a (mm)	0.10	1.08	-0.46	0.97
ρ (s)	5.27	5.27	4.35	4.35
f (Hz)	0.19	0.19	0.23	0.23

Table 4.5: 3D Estimated Model Parameters: Patient Nancy

	x	y	z
o (mm)	-7.36	-1.72	-1.45
a (mm)	0.10	-0.46	1.02
ρ (s)	4.81	4.81	4.81
f (Hz)	0.21	0.21	0.21




Figure 4.2: 2D motion of clips of 2 patients (red) and their respective estimated models of average 3D motion in Anterior-Posterior and Lateral Fluoroscopy (green).

Table 4.6: 3D Estimated Model Error

Patient	Direction of Motion	Root Mean Squared Error (RMS)		Max Error (cm)
		Mean (cm)	Std. Dev. (cm)	
Jack Lateral View	y -direction	0.01	0.02	0.24
Jack Lateral View	z -direction	0.07	0.11	0.58
Jack AP View	x -direction	0.02	0.03	0.14
Jack AP View	z -direction	0.25	0.31	1.87
Nancy Lateral View	y -direction	0.33	0.27	1.23
Nancy Lateral View	z -direction	1.49	1.22	5.43
Nancy AP View	x -direction	0.01	0.01	0.06
Nancy AP View	z -direction	0.27	0.38	2.37
Average		0.30	0.29	1.49

Tables 4.8 & 4.7 show comparisons between the frequency f estimated by the modeling algorithm and the dominant frequency estimated by Fourier analysis for each of the 1D trajectories.

Table 4.7: 1D Frequency Parameter Comparison: Patient Jack

Estimation Method	f (Hz)			
	x	$z1$	y	$z2$
Modeling Algorithm	0.22	0.22	0.23	0.23
Fourier Analysis	0.23	0.23	0.21	0.21

Table 4.8: 1D Frequency Parameter Comparison: Patient Nancy

Estimation Method	f (Hz)			
	x	$z1$	y	$z2$
Modeling Algorithm	0.19	0.19	0.23	0.23
Fourier Analysis	0.23	0.23	0.21	0.21

4.3 Correlation With External Markers

This study involved 2 patients. The only 3D motion data currently available to us that includes external markers is the 3D estimated model of average motion. Fig. 4.3 shows the trajectories of the external markers in the y -direction along side the 3D estimated models of average motion in each dimension.

We studied the correlation between the 3D estimated model of the average motion of the clip in each dimension and the motion in the y -dimension of each of three markers for two patients. By way of comparison, we included the correlation between the tracked




Figure 4.3: Motion in the y -direction of three external markers (red, green, and blue) shown with the average motion in the x , y and z -directions of the estimated models of 2 patients (pink).

motion of the clip in two dimensions and the external markers as well. The correlation was measured by plotting the motion of the internal marker against the motion of the external marker and then computing the linear regression line which fits the data. The sum of the squares of the vertical deviations from the line (R^2) was used as the measure of correlation. Table 4.9 summarizes the results and Figs. 4.4 and 4.5 show the correlation plots for both the tracked motion and the 3D estimated model of the average motion for the two patients.

The correlation with the tracked motion was 0.88 on average while the correlation with the 3D estimated model of the average motion was 0.53. The slopes of the regression lines are included for purposes of comparison. For the patient Jack, the slopes of the regression lines of the correlation which uses the 3D estimated model of average motion are very similar to the slopes of the regression lines of the correlation which uses the tracked motion. However, for the patient Nancy, this is not the case.




Figure 4.4: Motion in the y -direction of external marker 2 (blue diamonds), 3 (pink squares), and 4 (yellow triangles) of patient Jack plotted against the 3D average motion in the x , y and z -directions of the estimated models (A, C, E) and against the tracked motion in the y and z -directions (B, D).




Figure 4.5: Motion in the y -direction of external marker 3 (blue diamonds), 4 (pink squares), and 5 (yellow triangles) of patient Nancy plotted against the 3D average motion in the x , y and z -directions of the estimated models (A, C, E) and against the tracked motion in the y and z -directions (B, D).

Table 4.9: External Marker Correlation

Patient	Direction of Clip Motion	External Marker Number	R^2		Slope	
			Tracked Motion	Modeled Motion	Tracked Motion	Modeled Motion
Jack	<i>x</i> -direction	2	-	0.64	-	1.34
Jack	<i>x</i> -direction	3	-	0.62	-	1.75
Jack	<i>x</i> -direction	4	-	0.58	-	1.70
Jack	<i>y</i> -direction	2	0.82	0.64	1.32	1.27
Jack	<i>y</i> -direction	3	0.80	0.62	1.71	1.65
Jack	<i>y</i> -direction	4	0.71	0.58	1.62	1.61
Jack	<i>z</i> -direction	2	0.94	0.64	-4.14	-4.58
Jack	<i>z</i> -direction	3	0.90	0.62	-5.34	-5.96
Jack	<i>z</i> -direction	4	0.83	0.58	-5.14	-5.82
Nancy	<i>x</i> -direction	3	-	0.41	-	0.46
Nancy	<i>x</i> -direction	4	-	0.35	-	0.27
Nancy	<i>x</i> -direction	5	-	0.32	-	0.15
Nancy	<i>y</i> -direction	3	0.86	0.41	3.11	-2.16
Nancy	<i>y</i> -direction	4	0.93	0.61	2.10	-1.36
Nancy	<i>y</i> -direction	5	0.95	0.88	1.24	-0.60
Nancy	<i>z</i> -direction	3	0.91	0.41	-6.34	4.79
Nancy	<i>z</i> -direction	4	0.97	0.35	-4.24	2.85
Nancy	<i>z</i> -direction	5	0.99	0.32	-2.51	1.60
Average			0.88	0.53		

Chapter 5

Discussion and Conclusion

Our method can be used in settings where the latest technology, i.e., simultaneous orthogonal fluoroscopy, is not available. During the course of this study it became apparent that equipment shortcomings were not the only problem we faced during the data acquisition phase. Because work in this field is relatively new, radiation oncologists do not always collect data which is as usable to us as possible (e.g., a given clip is not always visible in both fluoroscopic image sequences). It is often not clear if the data is usable until after it has been analyzed. The possibilities of the kinds of information which can be recovered are still emerging. One goal of our work is to develop and demonstrate the ways in which information can be recovered and establish a dialogue with the radiation oncologists to explain ways to improve data collection methods.

5.1 2D Clip Tracking

Estimating the location of a tumor accurately is crucial in providing effective and safe radiation therapy. The ground truth study of our clip tracking system shows that the average RMS error and standard deviation is 0.47 mm and 0.33 mm, respectively. This is a considerable improvement of the 1.5 mm error of the RTRT system [32, 33]. This is underscored by the fact that the internal markers our system was tested on were typically

5 mm long while the RTRT system uses 2 mm markers. An error of 1.5 mm when tracking a 2 mm object misses the object by 75%, whereas an error of 0.47 mm when tracking a 5 mm object only misses the object by 9.4%. The RTRT system uses 1024×1024 pixel images with 8 bits per pixel. Whereas our system was only tested on a resolution of 640×480 pixels with 8 bits per pixel. Our system achieves more accurate performance than the RTRT system on lower resolution data, and would most likely do even better given higher quality images. Additionally, our method has the capability to track more than one marker which is necessary for designing a method to recover 3D tumor position and deformation accurately. The RTRT system, which in the reported studies did not use tracking information for multiple internal markers for the purpose of estimating tumor position, is currently used by most of the groups doing internal marker tracking [32, 33, 30, 25, 31].

Although error statistics are given for the RTRT system, there is no detailed explanation of how the error was calculated [32, 33, 30, 25, 31]. By manually marking clip position in the images and comparing that with tracked positions we have explained how we calculated the error of our system. If the information produced by the tracker is the first step of a system which models tumor motion its error should be analyzed. By calculating the error of the measurements which the rest of our system relies on we can give a more accurate estimate of how reliable the system will be overall.

5.2 Estimated Model of 3D Clip Motion

Due to regulations of Massachusetts General Hospital we were unable to gather additional fluoroscopic data from the same patients with which to compare the model, but we were able to assess if the assumptions made by the modeling algorithm were reasonable. This was done by comparing the frequency parameter obtained by the modeling algorithm with the dominant frequency obtained from Fourier analysis and by comparing the model with the trajectories used to produce it.

The Fourier analysis yielded similar dominant frequencies of the 1D trajectories to the modeling algorithm (see Tables 4.8 & 4.7). The fluoroscopy is taken at 30 Hz for a short amount of time, 5-10 s. This means that there are usually between 2^9 and 2^{10} samples. When using Fourier analysis the set of frequencies is then discretized with an interval between 0.06 Hz and 0.03 Hz. The dominant frequency resulting from the Fourier analysis differed from the frequency resulting from the modeling algorithm by at most 0.04 Hz and by as little as 0.01 Hz. If the dominant frequencies resulting from the Fourier analysis were used to compute the frequency for the 3D model the result would only be 0.01 Hz different from the result produced by using the frequencies obtained by the 1D modeling algorithm.

Our algorithm places more weight on the maxima and minima of the breathing cycles because radiation oncologists rely on these points more heavily. However, the results from the Fourier analysis are very similar, so it is difficult to assert which method is more successful. The modeling algorithm might perform better if there were even fewer samples

because the Fourier method would have even larger intervals between possible frequencies. Longer imaging times might be able to confirm which method is more accurate. Additionally, the two dominant frequencies of the 1D trajectories in a given fluoroscopic sequence computed with Fourier analysis are the same. This confirms that our assumption that they are the same is reasonable.

For one patient, Jack, the motion of a clip in two sequentially obtained fluoroscopic image sequences was similar enough to allow us to produce a useful model of the 3D motion of the clip. However, for the patient Nancy, the modeling algorithm did not perform as well. In order to recover 3D motion of a clip directly or to calculate a model to describe the motion it is necessary that the clip be present in both of the orthogonal image sequences. However, for all nine of the patients in this study, this condition only held for two patients. Even for those two patients, only one clip of each patient was visible in both image sequences. Thus, we had very limited data with which to test our 3D clip motion modeling algorithm. One goal of our work is to motivate the radiation oncologists to provide more usable data. It is desirable for the clip to be visible in both image sequences, and so we hope that this work will open a dialogue between the fields of computer vision and radiation therapy to increase the availability of usable data for testing. With more data we could hopefully determine if the model will work during subsequent treatment sessions and to improve it if necessary. Allowing the patient to undergo a little extra radiation during observation will hopefully result in a large reduction of radiation delivered to healthy tissue during treatment because the tumor can be more accurately located.

With respect to the data that the system was tested on, the RMS error of 3.0 mm although reasonable is not ideal. If we look only at the RMS error for the estimated model of 3D clip motion for the patient Jack, we see a dramatic decrease to only 0.8 mm. If the original trajectories from which the models were produced are examined it is clear that the sequentially obtained image sequences of the patient Jack are better suited for the modeling algorithm than that of the patient Nancy. This is because the patient Jack breathed at approximately the same frequency during both imaging times, 0.22 Hz and 0.23 Hz. However, the patient Nancy breathed noticeably faster in the Lateral image sequence, 0.23 Hz, as compared to the Anterior-Posterior sequence, 0.19 Hz. If a person breathes at a different rate this can alter the pattern of motion. This result indicates that breath coaching might be useful if a 3D model of motion is desired for treatment planning.

If breath coaching is not a possibility, another way to obtain a more accurate model could be to omit outlier breathing cycles when computing the parameters of the 3D model. The breathing pattern of the Anterior-Posterior image sequence of the patient Nancy had two breathing cycles with longer periods than the other breathing cycles of the image sequence. If these were omitted in calculating the 3D model, it might increase the accuracy of the model. However, this might not be a safe assumption. When there are so few breathing cycles to work with (in this study the lowest was 4) there is sometimes not enough information to accurately classify a breathing cycle as an outlier. Longer observation periods would be one way to solve this problem. For the patient Nancy it seemed that the main cause of the problems with the model were not caused by outlier breathing cycles but

rather the fact that the two image sequences had significantly different average frequencies as mentioned before.

We hope to see if the model will work well for additional patients when the data becomes available. If so, the next step would be to test the model on subsequent fluoroscopy sessions to evaluate the day to day accuracy of the model. Hospitals do not like to do more imaging than they deem necessary because it involves more patient risk due to the increase of radiation received. However, if the model is successful for additional patients, the potential to more accurately discern the tumor position during treatment might be justification enough to warrant additional pre-treatment imaging to test the validity of the model from day to day so that it can be cleared for use during treatment.

5.3 Correlation with External Markers

The study of the correlation between the motion of the external markers and the estimated average 3D motion of the internal markers was likewise limited to two patients. The results were nonetheless informative. The study yielded information about correlation between the motion of internal and external markers and about the modeling algorithm. The correlation between the motion of the external markers and the 3D estimated model of the average motion for patient Jack, which was 0.61, was less than the correlation with the tracked motion, which was 0.83. This is not surprising since the model is a more generalized version of the tracked motion, and when it was compared to the tracked motion of the external markers it was less strongly correlated. This is due to the fact that the model

smoothes over some of the small differences between breathing cycles.

Although the correlation of the motion of the external markers with the modeled motion was not as strong as the correlation of the motion of the external markers with the tracked motion, the slopes of the regression lines of both correlations were very similar in the case of patient Jack. To date, studies have only looked at the correlation between the motion of an external marker in one dimension, typically in the y -direction, and the motion of an internal marker in two dimensions, y and z . This is because of the setup of most fluoroscopic imaging systems. In order to compute the correlation between the motion of the external marker and the internal markers their motions should be recorded simultaneously. The easiest way to do this is to record fluoroscopy in such a way that all of the markers are visible within the frame. Since most fluoroscopic imaging systems can only be used from one viewpoint at a time, the motion of the internal marker can then only be observed in two dimensions. It is desirable to have information about the correlation of the external marker with the motion of the internal marker in the unobserved x -dimension. Thus, when we look at the correlation with the model of motion in the x -direction for patient Jack we could conclude that the slope of the regression line of the data would be similar to the correlation with the tracked motion in the x -direction if it were available because the slopes of the regression lines for data in the y and z -directions are similar. We hope to confirm this with tests on additional patients. If it is the case, the modeling algorithm could provide correlation information about motion in the x -direction that was previously unavailable.

The correlation between the motion of the external markers and the 3D estimated model of the average motion for the patient Nancy, which was 0.45, was much poorer than for the patient Jack, which was 0.61. This is most likely due to the problems with the model outlined in the previous section. The poor correlation further brings out the fact that improvements in the similarity of the frequency of the breathing pattern between two sequentially obtained image sequences are necessary. In Fig. 4.3 the frequency of the breathing pattern of the model of motion for patient Nancy does not stay synchronized with the external marker, whereas for patient Jack the model and the external marker stay synchronized. This problem illustrates well that if two patterns of motion slowly become desynchronized due to a transient phase shift, the correlation will break down.

Because this study was only performed on two patients the results are inconclusive. However, the results obtained for patient Jack are promising and we are hopeful that tests on more patient data will show that when the correlation between the motion of the model and the external marker is reasonably strong, and if the slope of the regression line for the data is similar to that of the tracked motion, we can recover information about the motion of the internal marker in the x -direction. Currently, the regression lines are used by radiation oncologists to predict the motion of the internal markers. This is done by using the observed position of the external marker during treatment as an index into the regression line to predict the position of the internal marker. If a regression line for the correlation between the motion of the external marker and the motion of the internal marker in the x -direction is computed then this would allow us to use the external markers to predict the

motion of the internal markers in all three dimensions.

The results from patient Nancy confirm that breathing is non-stationary and attention must be drawn to this. One option to account for the fact that breathing is non-stationary is to use breath coaching [17]. A more thorough course of action could be to use longer pre-treatment observation times to generate a more reliable model. This could be done by using the modeling algorithm to generate several different models to model different sections of the longer breathing pattern. It seems that breathing is regular over short intervals, so if the modeling algorithm were applied to small windows of time, this would be reasonable. Then during treatment the external markers could be monitored to choose which of the models should be used. This would allow for more flexible and robust modeling during treatment. Ideally the pre-treatment imaging would expose patients to substantially lower amounts of radiation than the radiation that the patient's healthy tissue would be exposed to during treatment, and our goal is to motivate more studies to this effect. If more accurate 3D motion models could be generated with longer pre-treatment imaging sessions as proposed than the extra radiation might be justified in order to benefit patients during treatment.

5.4 Future Work

There is much future work to be done. We hope that this work will motivate an even closer partnership between the fields of computer vision and radiation therapy. We would like to further test and subsequently improve the modeling methods. One method of improvement would be to no longer rely on the assumption that clips have linear trajectories. This could

be done by applying the section of the modeling algorithm which finds the maxima and minima of the breathing cycles to all of the trajectories. Creating a more dynamic model once longer pre-treatment image sequences are obtained is another key goal. This work serves to show that if Ozhasoglu and Murphy's [26] call for longer observation periods is fulfilled then gating techniques which account for the non-stationary aspects of breathing will be more easily developed. We hope that we will be able to realize this goal once more data is obtained.

It is known that tumors deform and that the deformation has the possibility of greatly altering the radiation both the tumor and the surrounding tissue receive. 4DCT is a new imaging technology which may be used to address the problems caused by tumor deformation. It allows CT scans to be taken at many steps during the respiratory cycle. However, it exposes the patient to more radiation, is generally more costly, and is not widely available. It would be beneficial to have a method which allows for a similar type of information recovery but works with technology which is not as advanced and therefore more prevalent. Ozhasoglu and Murphy [26] stressed the importance of monitoring the tumor directly, but this is difficult given current technological limitations.

Radiation oncologists often hand-contour the outline of the tumor in each slice of the CT scan obtained for treatment planning. The contours provide 3D information about the shape of the tumor that could be used in conjunction with fluoroscopy to extrapolate 4D information. Two example slices are shown in Fig. 5.1. The following describes a theory which we will implement in order to attempt to model the motion of the tumor volume.




Figure 5.1: Two slices of a CT scan. The tumor contour is marked yellow. One clip in each image is visible as a thick white mark outside of the contour as indicated by the red arrow. The pink cross marks the center of the image.

We propose the following method to model the motion and deformation of the tumor volume using the internal fiducial markers. First, transform the contours into a thin-plate spline representation as described by Bookstein [5]. In doing this, make sure that the clips act as control points on the splines. If the clip does not lie directly on the tumor, then assume that the clip moves rigidly with respect the closest point on the contour, and use this point as a control point. There will be additional control points as well, and they should be chosen intermittently along each spline between the clip control points. The more of these points that there are, the less influence each of the clip control points will have on the deformation of the tumor. The number of non-clip control points will need to be determined with the help of radiation oncologists. Using the 3D motion of the clips a model of the 3D tumor motion can be generated. The new clip positions represent the new positions of the control points along the splines and a transformation and deformation can be calculated.

If the only available data is the 3D estimated model of average clip motion an average model of tumor motion can still be generated. If tracked motion in 3D is available the model can be estimated directly. Although calculating the transformation and deformation of a 3D object is computationally expensive, the model could still be used with online tracking as follows. During the pre-treatment phase data can be collected and the model can be generated offline. A lookup table of possible positions can be created. Then for any given set of tracked clip positions during treatment instead of generating the model, it (or the most similar model) can simply be looked up. This method could provide a way to model the motion of the tumor directly. We hope to implement and test this method when suitable data becomes available.

An additional potential benefit of this method comes from examining how the clips move relatively with one another. By combining the models of the clips the overall error of the system could be reduced. This is because the additional information of the relative motion of the clips will be used. In computing the 3D model of tumor motion, if a clip seems to be deviating too much from its expected trajectory it could be classified as an outlier. This could account for mistakes made by the 2D clip tracker and the 3D internal marker modeling algorithm. We hope to explore this in future work and to see how the overall error of the system will be affected.

If further tests show that the modeling of the motion of internal markers in 3D is successful, using this information in conjunction with the proposed method to model the motion of the tumor itself we could come even closer to fulfilling Ozhasoglu and Mur-

phy's [26] request for direct tumor observation. We look forward to working with radiation oncologists to realize these goals and help improve the treatment of cancer further.

Bibliography

- [1] J. R. Jr. Adler, M. J. Murphy, S. D. Chang, and S. L. Hancock. Image-guided robotic radiosurgery. *Med Phys*, 44:1299–1306, 1999.
- [2] R. I. Berbeco, S. B. Jiang, G. C. Sharp, G. T. Y. Chen, H. Mostafavi, and H. Shirato. Integrated radiotherapy imaging system (IRIS): design considerations of tumor tracking with linac gantry-mounted kV x-ray systems. *Phys Med Biol*, 49:243–255, 2004.
- [3] M. Betke, J. Ruel, G. C. Sharp, S. B. Jiang, D. P. Gierga, and G. T. Y. Chen. Tracking and prediction of tumor movement in the abdomen. 2003, submitted.
- [4] J. M. Blackall, A. P. King, G. P. Penney, A. Adam, and D. J. Hawkes. A statistical model of respiratory motion and deformation of the liver. In W. J. Niessen and M. A. Viergever, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2001*, pages 1338–1340, Utrecht, The Netherlands, October 2001. Springer-Verlag, Berlin.
- [5] F. L. Bookstein. Principal warps: thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(6):567–585, June 1989.
- [6] Cancer Statistics 2004, The American Cancer Society. <http://www.cancer.org>.
- [7] Q. Chen, M. S. Weinhaus, F. C. Deibel, J. P. Ciezki, and R. M. Macklis. Fluoroscopic study of tumor motion due to breathing: facilitating precise radiation therapy for lung cancer patients. *Med Phys*, 28(9):1850–1856, September 2001.
- [8] E. C. Ford, G. S. Mageras, E. Yorke, K. E. Rosenzweig, R. Wagman, and C. C. Ling. Evaluation of respiratory movement during gated radiotherapy using film and electronic portal imaging. *Int J Radiat Oncol Biol Phys*, 52(2):522–531, 2002.
- [9] D. A. Forsyth and J. Ponce. *Computer vision, a modern approach*. Prentice Hall, NJ, 2003.
- [10] D. P. Gierga, J. Brewer, G. C. Sharp, M. Betke, C. G. Willett, and G. T. Y. Chen. The correlation between internal and external markers for abdominal tumors: implications for respiratory gating. *Submitted*, May 2004.
- [11] D. P. Gierga, G. T. Y. Chen, J. H. Kung, M. Betke, J. Lombardi, and C. G. Willett. Quantification of respiration-induced abdominal tumor motion and its impact on IMRT dose distributions. *Int J Radiat Oncol Biol Phys*, 58(5):1584–1595, 2004.

- [12] D. P. Gierga, G. Sharp, J. Brewer, M. Betke, C. G. Willett, and G. T. Y. Chen. Correlation between external and internal markers for abdominal tumors: implications for respiratory gating. In *Proceedings of the 45th Annual Meeting of the American Society for Therapeutic Radiology and Oncology*, Salt Lake City, Utah, October 2003.
- [13] I. Haber, D. N. Metaxas, and L. Axel. Using tagged MRI to reconstruct a 3D heartbeat. *Computing in Science and Engineering*, 2(5):18–30, 2000.
- [14] J. Hanley, M. M. Debois, D. Mah, G. S. Mageras, A. Raben, K. Rosenzweig, B. Mychalczak, L. H. Schwartz, P. J. Gloeggler, W. Lutz, C. C. Ling, S. A. Leibel, Z. Fuks, and G. J. Kutcher. Deep inspiration breath-hold technique for lung tumors: the potential value of target immobilization and reduced lung density in dose escalation. *Int J Radiat Oncol Biol Phys*, 45(3):603–611, 1999.
- [15] D. Hearn and M. P. Baker. *Computer graphics with OpenGL*. Prentice Hall, NJ, 2004.
- [16] P. J. Keall, V. R. Kini, S. S. Vedam, and R. Mohan. Motion adaptive x-ray therapy: a feasibility study. *Phys Med Biol*, 46(1):1–10, 2001.
- [17] V. R. Kini, S. S. Vedam, P. J. Keall, S. Patil, C. Chen, and R. Mohan. Patient training in respiratory-gated radiotherapy. *Med Dosim*, 28(1):7–11, 2003.
- [18] H. D. Kubo, P. M. Len, S. Minohara, and H. Mostafavi. Breathing-synchronized radiotherapy program at the university of california davis cancer center. *Med Phys*, 27:346–353, 2000.
- [19] H. Livyatan, Z. Yaniv, and L. Joskowicz. Gradient-based 2-D/3-D rigid registration of fluoroscopic X-Ray to CT. *IEEE Trans Med Imag*, 22(11):1395–1406, November 2003.
- [20] G. S. Mageras, E. Yorke, K. Rosenzweig, L. Braban, E. Keatley, E. Ford, S. A. Leibel, and C. C. Ling. Fluoroscopic evaluation of diaphragmatic motion reduction with a respiratory gated radiotherapy system. *J Appl Clin Med Phys*, 2(4):191–200, 2001.
- [21] K. McLeish, D. L. G. Hill, D. Atkinson, J. M. Blackall, and R. Razavi. A study of the motion and deformation of the heart due to respiration. *IEEE Trans Med Imag*, 21(9):1142–1150, September 2002.
- [22] D. Metaxas and D. Terzopoulos. Dynamic deformation of solid primitives with constraints. In *ACM Special Interest Group on Computer Graphics and Interactive Techniques – SIGGRAPH 1992*, pages 309–312, Chicago, IL, July 1992. ACM Press, New York.
- [23] S. Minohara, T. Kanai, M. Endo, K. Noda, and M. Kanazawa. Respiratory gated irradiation system for heavy-ion radiotherapy. *Int J Radiat Oncol Biol Phys*, 47(4):1097–1103, 2000.

- [24] M. J. Murphy. Fiducial-based targeting accuracy for external-beam radiotherapy. *Med Phys*, 29:334–344, 2002.
- [25] T. Neicu, H. Shirato, Y. Seppenwoolde, and S. B. Jiang. Synchronized Moving Aperture Radiation Therapy (SMART): Average tumor trajectory for lung patients. *Phys Med Biol*, 48:587–598, February 2003.
- [26] C. Ozhasoglu and M. J. Murphy. Issues in respiratory motion compensation during external-beam radiotherapy. *Int J Radiat Oncol Biol Phys*, 52(5):1389–1399, 2002.
- [27] K. E. Rosenzweig, J. Hanley, D. Mah, G. Mageras, M. Hunt, S. Toner, C. Burman, C. C. Ling, B. Mychalczak, Z. Fuks, and S. A. Leibel. The deep inspiration breath-hold technique in the treatment of inoperable non-small-cell lung cancer. *Int J Radiat Oncol Biol Phys*, 48(1):81–87, 2000.
- [28] A. Schweikard, G. Glosser, M. Bodduluri, M. J. Murphy, and J. R. Adler. Robotic motion compensation for respiratory movement during radiosurgery. *Comput Aided Surg*, 5(4):263–277, 2000.
- [29] A.M. Seifalian, D. J. Hawkes, C. R. Hardingham, A.C. Colchester, and J.F. Reidy. Validation of a quantitative radiographic technique to estimate pulsatile blood flow waveforms using digital subtraction angiographic data. *J Biomed Eng*, 3(3):225–233, May 1991.
- [30] Y. Seppenwoolde, H. Shirato, K. Kitamura, S. Shimizu, M. Van Herk, J. V. Lebesque, and K. Miyasaka. Precise and real-time measurement of 3D tumor motion in lung due to breathing and heartbeat, measured during radiotherapy. *Int J Radiat Oncol Biol Phys*, 53(4):822–834, 2002.
- [31] G. C. Sharp, S. B. Jiang, S. Shiimizu, and H. Shirato. Prediction of respiratory tumor motion for real-time image-guided radiotherapy. *Phys Med Biol*, 49:425–440, January 2004.
- [32] H. Shirato, S. Shimizu, K. Kitamura, T. Nishioka, K. Kagei, S. Hashimoto, H. Aoyama, T. Kunieda, N. Shinohara, H. Dosaka-Akita, and K. Miyasaka. Four-dimensional treatment planning and fluoroscopic real-time tumor tracking radiotherapy for moving tumor. *Int J Radiat Oncol Biol Phys*, 48(2):435–442, September 2000.
- [33] H. Shirato, S. Shimizu, T. Kunieda, K. Kitamura, M. van Herk, K. Kagei, T. Nishioka, S. Hashimoto, K. Fujita, H. Aoyama, K. Tsuchiya, K. Kudo, and K. Miyasaka. Physical aspects of a real-time tumor-tracking system for gated radiotherapy. *Int J Radiat Oncol Biol Phys*, 48(4):1187–1195, November 2000.

- [34] J. S. Stromberg, M. B. Sharpe, L. H. Kim, V. R. Kini, D. A. Jaffray, A. A. Martinez, and J. W. Wong. Active breathing control (ABC) for Hodgkin's disease: reduction in normal tissue irradiation with deep inspiration and implications for treatment. *Int J Radiat Oncol Biol Phys*, 48(3):797–806, 2000.
- [35] D. Tomaževič, B. Likar, T. Slivnik, and F. Pernuš. 3-D/2-D registration of CT and MR to X-Ray images. *IEEE Trans Med Imag*, 22(11):1407–1416, November 2003.
- [36] S. S. Vedam, P. J. Keall, V. R. Kini, and R. Mohan. Determining parameters for respiration-gated radiotherapy. *Med Phys*, 28(10):2139–2146, October 2001.
- [37] S. S. Vedam, V. R. Kini, P. J. Keall, V. Ramakrishnan, H. Mostafavi, and R. Mohan. Quantifying the predictability of diaphragm motion during respiration with a noninvasive external marker. *Med Phys*, 30(4):505–513, April 2003.
- [38] G. Welch and G. Bishop. An introduction to the Kalman filter. <http://www.cs.unc.edu/welch/kalman/kalmanIntro.html>.
- [39] J. W. Wong, M. B. Sharpe, D. A. Jaffray, V. R. Kini, J. M. Robertson, J. S. Stromberg, and A. A. Martinez. The use of active breathing control (ABC) to reduce margin for breathing motion. *Int J Radiat Oncol Biol Phys*, 44(4):911–919, 1999.