

Performance Evaluation and Algorithm Development for Real-Time Tumour Tracking Using Positron Emission Markers

by

Marc Chamberland, B.Sc.

A Thesis submitted to
the Faculty of Graduate Studies and Research
in partial fulfilment of
the requirements for the degree of
Master of Science

Ottawa-Carleton Institute for Physics

Department of Physics
Carleton University
Ottawa, Ontario, Canada

September 2009

Copyright ©

2009 - Marc Chamberland, B.Sc.

Abstract

I present a technique to track tumour location in real-time during radiotherapy treatments. Four scintillation detector modules are used to collect annihilation coincidence lines from fiducial positron emitting markers implanted in or around the tumour. The position of the markers can be tracked in real-time using an expectation-maximization (EM) clustering algorithm. I evaluated the accuracy and precision of the system by developing a Monte Carlo simulation of the detector design. The average precision of the system was found to be 0.49 mm and the average accuracy 0.77 mm. Testing of the tracking algorithm with a clinical PET system also yielded submillimetre precision and accuracy. I conclude this performance is sufficient for real-time tumour tracking. I also developed a method to initialize the EM algorithm and fail-safe methods to correct for failed localizations. The average error of two prediction methods was 1.1 ± 0.7 mm and 3.6 ± 2.3 mm respectively.

To Jezamine

Mahal na mahal kita hanggang sa wakas

Acknowledgments

I would like to thank my supervisor, Dr. Tong Xu, for his guidance and his help throughout my research. I look forward to continue my work under his supervision. I also extend my thanks to Dr. Richard Wassenaar from the Ottawa Hospital whose help with setting up and running the experiments was invaluable. Thanks to Nathan Churchill who provided me with the basic idea for the centre of mass initialization and to Benjamin Spencer for helping with the DAQ module. I would also like to thank Mojgan Soleimani-Marghmalecki for giving me data acquired with the prototype detector.

Thanks to my family for their support and understanding throughout my studies.

Finally, my very heartfelt thanks to my wonderful girlfriend, Jezamine, for her love, patience, encouragement, and understanding.

Contributions

This thesis presents work I have done in developing and assessing the performance of PeTrack. I developed a model of the PeTrack detector system in a Monte Carlo simulation. I analyzed the data from the simulation and calculated the accuracy and precision of PeTrack. My supervisor and I conducted two experiments at the Ottawa Hospital Cancer Centre to evaluate the performance of the tracking algorithm with data acquired with a clinical PET system. I was involved in the analysis of the data from these experiments. I converted the data from List Mode format to Cartesian coordinates and I assessed the performance of the tracking algorithm. I have also evaluated the performance of PeTrack using data that was recorded by a prototype of the PeTrack system. I was not involved in the design or operation of the PeTrack detector. I also improved the localization algorithm by implementing an initialization method based on the centre of mass of midpoints. The idea for this method came from a colleague. Most importantly, I added fail-safe methods that detect and recover lost markers. This addition makes the tracking algorithm more robust and suitable for clinical use. Finally, I implemented and tested two motion prediction methods with previously recorded data.

Table of Contents

- Abstract ii

- Acknowledgments iv

- Contributions v

- Table of Contents vi

- List of Tables x

- List of Figures xi

- List of Acronyms xiv

- List of Symbols xvi

- 1 Introduction 1**
 - 1.1 Respiratory Motion 1
 - 1.2 Tumour Tracking 2
 - 1.3 PeTrack 3

- 2 PeTrack Overview 5**
 - 2.1 Basic Physics 5
 - 2.1.1 Positron Emission 5

2.1.2	Annihilation	6
2.1.3	Scintillation Detector	6
2.2	PeTrack	7
2.2.1	Principle of Operation	7
2.2.2	Detector Design	8
2.2.3	Multimarker Localization Algorithm	9
2.3	Summary	13
3	Simulation Study	14
3.1	GATE	14
3.2	Detector	16
3.3	Electronics	16
3.4	Physics	17
3.5	Sources	18
3.6	Localization	19
3.7	Accuracy and Precision	20
3.8	Results	21
3.9	Discussion	26
3.10	Conclusion	31
4	Experimental Evaluation of PeTrack	32
4.1	Preliminary Evaluation of PeTrack	32
4.1.1	PET System	32
4.1.2	Sources	33
4.1.3	Compact List Mode Data Format	34
4.1.4	Conversion from Compact List Mode Data Format to Cartesian Coordinates	35
4.1.5	Localization	36

4.1.6	Results	38
4.1.7	Discussion	38
4.1.8	Conclusion	41
4.2	Assessment with an Improved Experimental Protocol	41
4.2.1	PET System	41
4.2.2	Sources	41
4.2.3	X-Y Plotter	42
4.2.4	Motion Patterns	42
4.2.5	Crystal ID List Mode Data Format	44
4.2.6	Conversion from Crystal ID List Mode Data Format to Cartesian Coordinates	46
4.2.7	Localization	47
4.2.8	Results	48
4.2.9	Discussion	50
4.2.10	Conclusion	53
4.3	Assessment with a Prototype PeTrack System	53
4.3.1	Methods	53
4.3.2	Results	54
4.3.3	Discussion	55
4.4	Conclusion	55
5	Algorithm Improvements	57
5.1	Algorithm Initialization	57
5.1.1	Initialization Around the Centre of Mass	57
5.1.2	Identifying Failed Markers	60
5.1.3	Testing for Bias	62
5.1.4	Results	62

5.1.5	Discussion	63
5.2	Motion Prediction	65
5.2.1	Linear Extrapolation	65
5.2.2	Circular Approximation	66
5.2.3	Results	66
5.2.4	Discussion	68
5.3	Conclusion	69
6	General Discussion	70
6.1	Sources of Errors	70
6.2	Limitations	72
7	Conclusion	74
7.1	Conclusions	74
7.2	Future Work	75
	List of References	76

List of Tables

3.1	Spatial configuration of the markers for each simulation run.	19
3.2	Localization error averaged over all localizations for each marker and each run.	24
3.3	Fraction of lines rejected by the algorithm and scatter and random fractions reported by GATE.	30
4.1	Characteristics of each scan of the experiment. A_x and A_y refer to the amplitude of motion along x and y , respectively.	45

List of Figures

2.1	After correctly assigning each coincidence line to its marker, the position of each marker minimizes the RMS distance to its coincidence lines.	8
2.2	PeTrack detector modules (labelled A1, A2, B1, and B2) mounted on the gantry of a linear accelerator 50 cm from the isocenter.	9
2.3	Logical flow of the expectation-maximization algorithm.	12
3.1	PeTrack detector design in GATE.	16
3.2	Effective location of a marker between $t = 300$ ms and $t = 400$ ms. . .	21
3.3	Histogram of the distribution of the absolute localization error for all runs.	22
3.4	Histogram of the distribution of the relative localization error for all runs.	23
3.5	Localized track of the markers (green stars) and real track of the markers (red crosses) for a sample run. All markers move in the same plane.	25
3.6	Systematic shrinking of the distances between the markers due to the curvature of the path of the markers.	27
3.7	To estimate the scale of the systematic error, the distance vector is projected onto the normalized radial vector.	28
3.8	Histogram of the projected distance vector onto the normalized radial vector.	29

4.1	Setup of the experiment.	34
4.2	Projection in the $x-y$ plane of the geometrical correspondence between Cartesian coordinates and List Mode indices.	37
4.3	Histogram of the distribution of the relative localization error for all scans.	39
4.4	The three dashed coincidence lines are binned into the single solid line.	40
4.5	X-Y plotter used in the experiment. The marker configuration is shown with coins on the picture.	43
4.6	Illustration of hysteresis. On the left, two sine waves are in phase. On the right, they are out of phase by 30°	44
4.7	2D path of the markers for scans #5 to #8. The actual amplitude of motion is different than shown.	45
4.8	Geometry of a detector block from the Philips Allegro PET system (not to scale).	47
4.9	Histogram of the distribution of the relative localization error for all scans.	48
4.10	PeTrack data of scan #4 fitted to a sine wave. Only the x component is shown.	49
4.11	Residuals of the fit of scan #4 to a sine wave.	50
4.12	PeTrack data of scan #8 fitted to the real breathing curve. Only the x component is shown.	51
4.13	Setup of the experiment with the PeTrack prototype.	54
5.1	Logic flow of the initialization method.	59
5.2	Projection of the normalized vector pointing from the localized marker to its true location onto the normalized vector pointing from the initial location to the true location of the marker.	63

5.3 x coordinate of a marker as a function of time for scan#5. From $t = 0$ to $t = 1.2$ s, the marker location fluctuates randomly. This corresponds to the initialization phase of the algorithm. The tracking recovers at approximately $t = 1.3$ s. 64

5.4 Comparison of predicted marker location (red crosses) and tracked path of the markers (blue stars) for scan #1 (projected onto the $x - y$ plane). 67

5.5 Comparison of predicted marker location (red crosses) and tracked path of the markers (blue stars) for scan #8 (projected onto the $x - y$ plane). 67

List of Acronyms

Acronyms	Definition
AAPM	American Association of Physicists in Medicine
BGO	bismuth germanate
COIL	coincidence line
CSDA	continuous slowing down approximation
CT	computed tomography
EM	expectation-maximization
FOV	field of view
FWHM	full-width at half-maximum
GATE	GEANT4 Application for Tomographic Emission
GSO	gadolinium oxyorthosilicate
ICRU	International Commission on Radiation Units and Measurements
IMRT	intensity modulated radiotherapy
LEP	low energy electromagnetic processes

PDF	probability density function
PET	positron emission tomography
RMS	root mean square
RMSE	root mean square error
SEP	standard energy electromagnetic processes
SPECT	single photon emission tomography

List of Symbols

Symbols	Definition
a_n	linear prediction coefficients
α	angle subtended by an arc at the centre of the circle
α_k	relative activity of marker k
A	mass number
A_x, A_y	amplitude of motion along x and y respectively
β^+, e^+	positron
$\vec{d}(\mathbf{L}_n, \vec{m}_k)$	minimum distance vector from coincidence line \mathbf{L}_n to marker location m_k
δ_a	crystal pitch in the axial direction of the scanner
δ_{com}	minimum distance between an arc and its centre of mass
δ_r	pitch between each radial bin
$\delta_x, \delta_y, \delta_z$	localization error of coordinates x , y , and z
e^-	electron

$G \left(\vec{d}(\mathbf{L}_n, \vec{m}_k) , \sigma_k \right)$	probability density function of the Gaussian distribution of the minimum distance between a marker and its coincidence lines
γ	photon (gamma ray)
i	iteration index
k	marker index
K	total number of markers
\mathbf{L}_n	coincidence line n
m_k	location of marker k
n	index of coincidence line
N	total number of coincidence lines OR neutron number
η	index of crystal in a detector block
ν	neutrino
$p_{n,k}$	probability of coincidence line \mathbf{L}_n belonging to marker k
ϕ	transverse angle index of a coincidence line with the x -axis
r	distance of a coincidence line to the isocentre
ρ	radius of circle
R	radius of scanner
σ_d	intrinsic spatial resolution of PeTrack detector

σ_k	standard deviation of the distribution of the minimum distance between a marker k and its coincidence lines
σ_m	precision of the localization of a marker
$\sigma_{r.l.e.}$	standard deviation of the distribution of relative localization error
t	time
Δt	system latency
$T_{\frac{1}{2}}$	half-life
θ	angle subtended by one detector block at the isocentre
\vec{V}_k	vector shift of marker k
w	crystal width
ω	angular velocity
$(x_1, y_1, z_1), (x_2, y_2, z_2)$	endpoints coordinates of a coincidence line
${}^A_Z X_N$	parent nucleus
\tilde{x}_{t_n}	predicted location of marker at time t_n
$x_{t_{n-1}}, x_{t_{n-2}}$	previous two locations of marker
${}^A_Z Y_N$	daughter nucleus
z_a, z_b	crystal indices in the axial direction
Z	atomic number

Chapter 1

Introduction

1.1 Respiratory Motion

Developments in radiotherapy have helped increase the delivery accuracy of treatments. For instance, intensity modulated radiotherapy (IMRT) can modulate the beam intensity in order to improve the conformity of the dose distribution to the tumour target volume. [1] This is particularly important for abdominal and thoracic tumours because of the numerous organs at risk in the vicinity.

The International Commission on Radiation Units and Measurements (ICRU) recommends that the uncertainty on dose delivery in radiotherapy not exceed 5%. [2] However, breathing limits the delivery accuracy of radiotherapy in the pulmonary region. Jacobs *et al.* measured a 4% variation in a patient's thickness during breathing. [3] This leads to a dose error of 2% which needs to be accounted for in the treatment planning. This error makes the 5% accuracy requirement of the ICRU hard to achieve in daily practice since other sources of errors contribute to the delivery accuracy (patient setup, tumour delineation, etc.). Furthermore, many organs, such as the lungs, liver, and pancreas, among others, move with breathing. This breathing motion can have adverse effects on the dose delivery to the tumour. The target volume can receive a lower dose than planned, or the normal tissues can receive

a higher dose than expected. Both of these effects need to be avoided. The American Association of Physicists in Medicine (AAPM) Task Group 76 [4] recommended the use of respiratory motion management if either of the following conditions occur:

- The tumour is observed to move by more than 5 mm in any direction;
- Significant normal tissue sparing is gained through the use of respiratory motion management.

1.2 Tumour Tracking

Respiratory gating can be applied in order to minimize the effects of breathing motion. Respiratory gated radiotherapy administers radiation only during a certain portion of the patient's breathing cycle. The timing of the gate with the respiratory cycle is determined by tracking the tumour motion with an external respiratory signal or with internal fiducial markers. Tracking tumours with an external surrogate of respiration, such as abdominal displacement, is risky. Hoisak *et al.* report that the relationship between abdominal displacement and tumour motion varied during and between treatment sessions, which might cause the tumour motion to be incorrectly tracked. [5]

Implanted gold markers can be tracked in real-time using a method called fluoroscopy. [6] The markers are spheres with a diameter of 1.5 to 2.0 mm. Four sets of diagnostic x-ray television systems constantly acquire images during the treatment. The identification of the radio-opaque markers on the images allows the tumour to be tracked. The mean 3D distance between the actual and the calculated position of the markers is approximately 0.5 mm. However, the estimated skin dose delivered by the x-ray fluoroscopy is between 0.3 and 1.2 Gy per hour of treatment. [7] As a reference, a typical treatment duration is 10 minutes per weekday for four weeks at

the Ottawa Hospital, but treatment length can vary with tumour size and several other factors. Efforts need to be made to manage and limit such radiation doses. Fluoroscopic tracking of smaller cylindrical markers with a length of 3 mm and a diameter of 0.8 mm has been reported with a maximum tracking error varying between 0.2 and 2.3 mm at a 95% confidence level. [8]

Electromagnetic transponders are also used to track tumour motion. The Calypso system uses an external magnetic source and receiver coil array to track the location of implanted wireless coils in the patient. The precision and accuracy depend on the distance between the transponders and the receiver array. In a phantom study, both precision and accuracy were well below 1 mm when the distance between the transponders and the array was 27 cm. [9] In a patient study, the 3D difference between the location of the transponders determined by the Calypso system and the radiographic localization of the markers was 1.5 ± 0.9 mm. [10] The main drawback of this technique is the large size of the transponders. The transponders have a length of 8 mm and a diameter of 1.85 mm. The size of the markers makes them difficult to implant in the lung due to high patient risk. The Micropos 4DRT system also uses electromagnetic transponders for real-time tumour tracking. [11] Current patient studies suggest an accuracy of 2.7 ± 1.2 mm. However, the transponders are currently wired to the system, which complicates implantation and may cause patient discomfort.

1.3 PeTrack

This thesis will present a technique to track tumours in real-time using multiple positron emission markers implanted inside the tumour volume. The technique is called PeTrack and provides some advantages over the previous methods. The markers used in PeTrack are small, which facilitates implantation in the tumour. Since the

radioactive markers are implanted in the tumour, the dose delivered to normal tissues will be lower than in x-ray fluoroscopy.

In Chapter 2, I will provide an overview of the proposed technique. Chapter 3 and 4 will report the performance of the proposed technique evaluated from simulation and experiments. Chapter 5 covers development of the PeTrack algorithm and presents relevant results. In Chapter 6, I will discuss the main sources of errors that affect PeTrack and the limitations of my work.

Chapter 2

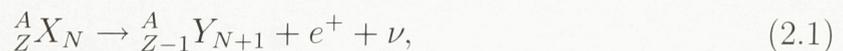
PeTrack Overview

This chapter gives a general overview of PeTrack. I start by reviewing some basic physics relevant to the technique. I describe the physical design of the PeTrack detectors. I conclude with a description of the tracking algorithm.

2.1 Basic Physics

2.1.1 Positron Emission

Positron emission, also called β^+ decay, occurs when a proton decays into a neutron inside a nucleus, ejecting a positron and a neutrino in the process. This decay can be expressed as follows:



where ${}^A_Z X_N$ is the parent nucleus (with mass number A , atomic number Z , and neutron number N) and ${}^A_{Z-1} Y_{N+1}$ is the daughter nucleus. The neutrino interacts very weakly with matter and is negligible in medical applications. However, its net effect is to carry away some energy from the decay. Thus, the positrons emitted from the nucleus have a continuous spectrum of energy, with the maximum kinetic energy corresponding to the energy of the decay. As a rule of thumb, the average energy of

the positrons is $\frac{1}{3}$ of the maximum kinetic energy. The β^+ decay occurs only if the parent's atomic mass exceeds the daughter's atomic mass by more than two electron rest masses (i.e. 1.022 MeV) due to energy conservation laws.

2.1.2 Annihilation

Since positrons are charged particles, they interact with matter in the same way as other light charged particles by ionizing and exciting nearby atoms and molecules. Eventually, an emitted positron will annihilate with an electron:



The distance from the source at which this occurs is called positron range. The annihilation process creates two gamma rays. Since the positron and electron will usually be close to rest at the moment of interaction, conservation laws dictate that the photons will be emitted back-to-back and each will have an energy of 511 keV. The left over energy and momentum of the positron and electron cause the gamma rays to be emitted at an angle smaller than 180° . This effect is called photon noncollinearity. The distribution of the deviation angle usually follows a Gaussian distribution centred at zero and its FWHM is $0.0044R$, where R is the radius of the scanner. [12]

2.1.3 Scintillation Detector

Positron emission tomography uses scintillation detectors to detect the gamma rays by-products of the β^+ decay. The annihilation photons deposit their energy in the scintillation material (typically an inorganic crystal). Some of this energy is re-emitted as visible light when the molecules and atoms de-excite. A photo-multiplier tube coupled to the crystal collects this visible light and generates an electrical signal. The finite size of the scintillation crystal is a limiting feature of detector resolution.

2.2 PeTrack

The following section is based on work by Xu *et al.* [13]

2.2.1 Principle of Operation

The detection of coincidence events from positron annihilation events forms the basis of PeTrack. A coincidence is defined as the simultaneous detection of two annihilation photons by opposite detectors. In practice, any two events detected within a certain interval of time in opposite detectors will be considered to be a coincidence event. This interval of time is called the coincidence timing window. It is usually set between 6 and 12 ns to account for variations in the crystal's response time to the interactions with light and for variations in the processing time of the signals by the electronics. [14] The 3D line that joins two detectors in which a coincidence was recorded is called a coincidence line.

It is possible to determine the location of a positron source using several coincidence lines. The position of the source is the point in 3D space that minimizes the root mean square (RMS) distance to all the coincidence lines. This principle can be extended to several sources as long as each coincidence line is correctly assigned to the correct source. The location of each source minimizes the RMS distance to their respective coincidence lines (see Figure 2.1). PeTrack uses this principle to achieve real-time tumour tracking. By implanting positron-emitting fiducial markers inside the tumour, four detector modules will track the motion of the tumour by collecting the coincidence events and determining the location of the markers in real-time.

The markers will consist of an active positron-emitting spherical core inside a gold, titanium, or tungsten cylindrical capsule. The core will have a diameter of 0.4 to 0.5 mm. The positron emission isotope can be ^{124}I ($T_{\frac{1}{2}} = 4.2$ days), ^{74}As ($T_{\frac{1}{2}} = 17.8$ days), or ^{84}Rb ($T_{\frac{1}{2}} = 32.7$ days) depending on the duration of the

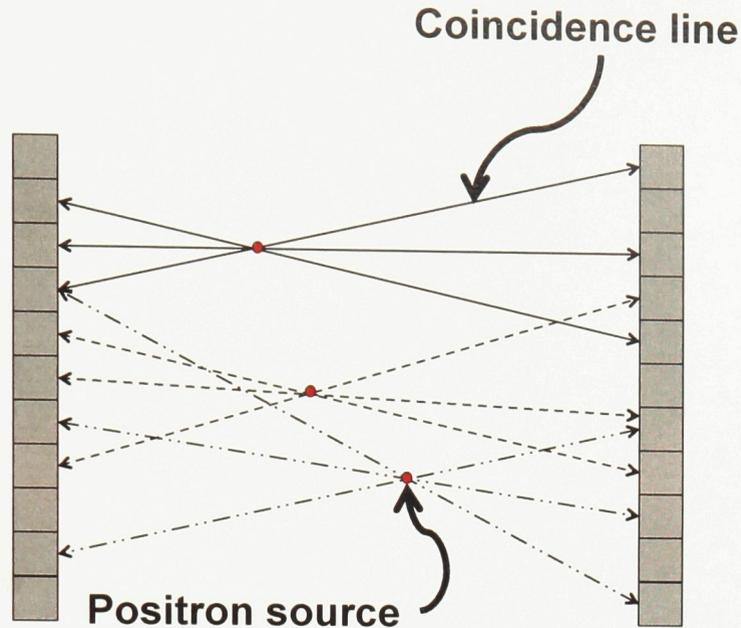


Figure 2.1: After correctly assigning each coincidence line to its marker, the position of each marker minimizes the RMS distance to its coincidence lines.

treatment of the patient. The capsule will have a length of 2 to 4 mm and a diameter of 0.5 to 0.8 mm. The small size of the markers will facilitate implantation and reduce patient risk. The relatively low activity of the markers will deliver a low dose to normal tissues since the markers can be implanted directly inside a tumour. The dose delivered to the tumour can be taken into account during treatment planning. The lifetime dose at 10 mm from a marker with an activity of 3.7 MBq is expected to be between 0.7 and 5.0 Gy, depending on the isotope used.

2.2.2 Detector Design

PeTrack will use position sensitive gamma ray detectors to collect the coincidence events from the markers. Two pairs of modules will be mounted at 90° on the gantry of a Linac (see Figure 2.2 [13]). The weight of the four PeTrack detectors will be less than 40 kg, which is comparable to current on-line imaging systems mounted on Linacs. Each module consists of a 3×3 array of detector blocks. Each detector block

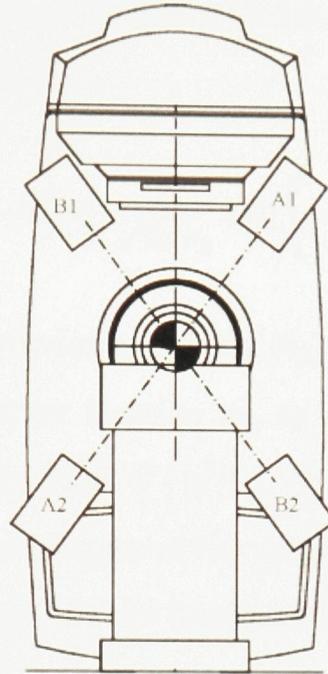


Figure 2.2: PeTrack detector modules (labelled A1, A2, B1, and B2) mounted on the gantry of a linear accelerator 50 cm from the isocenter. Reproduced by permission of the American Association of Physicists in Medicine.

in turn is made up of 13×13 bismuth germanate (BGO) crystals. The dimensions of each crystal are $4 \times 4 \times 20 \text{ mm}^3$. Since the scatter radiation from the Linac beam is considerably higher than the signal from the positron emission markers, the PeTrack detectors need to be gated off during the Linac pulse. Typically, the Linac pulse lasts 1 to $10 \mu\text{s}$. Coincidence events can be collected in the 2.5 to 10 ms interval between Linac pulses. This is sufficient for real-time tumour tracking. A data acquisition duty cycle of at least 80% is expected under these conditions.

2.2.3 Multimarker Localization Algorithm

The PeTrack system uses an expectation-maximization (EM) clustering algorithm to track the location of each of K different markers. Each marker is assigned a marker index k , which runs from $k = 1, \dots, K$. The distance from the centre of a marker to its coincidence lines is assumed to follow a Gaussian distribution with a mean of

zero and a standard deviation of σ_k . The probability density function (PDF) for this Gaussian distribution is given by:

$$G\left(|\vec{d}(\mathbf{L}_n, \vec{m}_k)|, \sigma_k\right) = \frac{1}{\sqrt{2\pi}\sigma_k} \exp\left(-\frac{|\vec{d}(\mathbf{L}_n, \vec{m}_k)|^2}{2\sigma_k^2}\right). \quad (2.3)$$

In this equation, k is the marker index and $\vec{d}(\mathbf{L}_n, \vec{m}_k)$ is the minimum distance vector that connects an estimated marker location m_k to one of the coincidence lines \mathbf{L}_n associated with this marker. The PDF $G\left(|\vec{d}(\mathbf{L}_n, \vec{m}_k)|, \sigma_k\right)$ is not a two-dimensional function. It is a one-dimensional function of the variable $\vec{d}(\mathbf{L}_n, \vec{m}_k)$, but the standard deviation, σ_k , is updated throughout the iterations of the algorithm. The notation $G\left(|\vec{d}(\mathbf{L}_n, \vec{m}_k)|, \sigma_k\right)$ brings attention to this fact.

The algorithm starts at iteration index $i = 0$ by making an initial estimate of the location of the markers. Chapter 5 describes a method I developed to initialize the starting points of the algorithm. The initial standard deviation is set to $\sigma_k = 5$ mm and the relative activity of each marker is set to $\alpha_k = \frac{1}{K}$. The following two steps are performed until the position of each marker changes by less than 0.05 mm:

1. Expectation step: the probabilities $p_{n,k}^{(i)}$ of each coincidence line \mathbf{L}_n , $n = 1, \dots, N$ (where N is the total number of coincidence lines used by the algorithm) belonging to each cluster $k = 1, \dots, K$ are calculated as follows:

$$p_{n,k}^{(i)} = \frac{\alpha_k^{(i)} G\left(|\vec{d}(\mathbf{L}_n, \vec{m}_k^{(i)})|, \sigma_k^{(i)}\right)}{\sum_{j=1}^K \alpha_j^{(i)} G\left(|\vec{d}(\mathbf{L}_n, \vec{m}_j^{(i)})|, \sigma_j^{(i)}\right)} \quad (2.4)$$

Each coincidence line \mathbf{L}_n has a probability $p_{n,k}^{(i)}$ of belonging to every marker k , i.e. each coincidence line is partly assigned to multiple markers. This is called a *soft clustering* scheme, as opposed to a *hard clustering* scheme where each coincidence line could only belong to one particular marker.

2. Maximization step: the relative activity of each marker, $\alpha_k^{(i)}$, and the standard deviation of the minimum distance between the coincidence lines and their corresponding markers, $\sigma_k^{(i)}$, are updated according to the probabilities obtained from the expectation step. The new parameters are given by:

$$\alpha_k^{(i+1)} = \frac{\sum_{n=1}^N p_{n,k}^{(i)}}{N}; \quad (2.5)$$

$$\sigma_k^{(i+1)} = \sqrt{\frac{\sum_{n=1}^N p_{n,k}^{(i)} \cdot |\vec{d}(\mathbf{L}_n, \vec{m}_k^{(i)})|^2}{\sum_{n=1}^N p_{n,k}^{(i)}}}. \quad (2.6)$$

The necessary vector shift to minimize the RMS distance of a marker to its coincidence lines is calculated for each previously estimated marker position. The vector shift is the result of the weighted sum of all $\vec{d}(\mathbf{L}_n, \vec{m}_k^{(i)})$, with each weight given by Equation (2.4):

$$\vec{V}_k^{(i)} = \frac{\sum_{n=1}^N p_{n,k}^{(i)} \cdot \vec{d}(\mathbf{L}_n, \vec{m}_k^{(i)})}{\sum_{n=1}^N p_{n,k}^{(i)}}. \quad (2.7)$$

The updated estimated position of the marker is thus:

$$\vec{m}_k^{(i+1)} = \vec{m}_k^{(i)} + \vec{V}_k^{(i)}. \quad (2.8)$$

Some coincidence lines may not be true coincidences: they might be scatter or random coincidences. Considering these events in the localization would affect adversely the accuracy of the algorithm. To prevent this situation, the probability of a coincidence line \mathbf{L}_n belonging to cluster k is set to zero for cluster k if the distance between the line \mathbf{L}_n and the estimated location of the marker $\vec{m}_k^{(i)}$ is larger than twice the standard deviation of the cluster: $|\vec{d}(\mathbf{L}_n, \vec{m}_k^{(i)})| > 2\sigma_k^{(i)}$.

Figure 2.3 illustrates the logical flow of the algorithm.

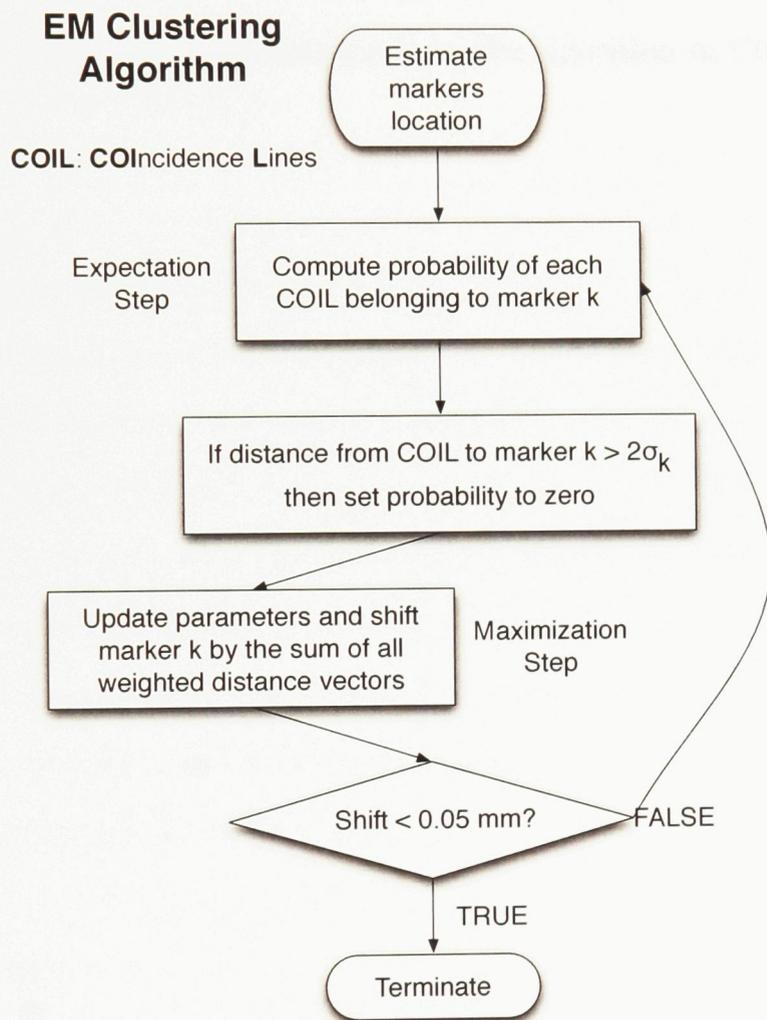


Figure 2.3: Logical flow of the expectation-maximization algorithm.

2.3 Summary

This chapter presented an overview of PeTrack, a real-time tumour tracking technique using positron emission markers. The design of the position sensitive gamma ray detectors and the EM clustering algorithm were both described.

The following chapters will present performance evaluation of the tracking algorithm. I will also detail improvements made to the algorithm in Chapter 5.

Chapter 3

Simulation Study

The preliminary study of PeTrack by Xu *et al.* evaluated the localization accuracy of the tracking algorithm. [13] They developed a Monte Carlo simulation model of the detector design and radioactive markers. They localized four stationary ^{124}I sources using the localization algorithm described in Section 2.2.3. The mean localization error was 0.55 ± 0.27 mm using an average of 62 coincidence lines per marker. The accuracy of the tracking algorithm now needs to be assessed with multiple moving markers since the previous study used static markers.

This chapter presents a simulation study of PeTrack using three moving fiducial markers in various spatial configurations.

3.1 GATE

The simulation was designed in GATE, the GEANT4 Application for Tomographic Emission, developed by the OpenGATE collaboration. [15] GATE is a Monte Carlo simulation toolkit for nuclear medicine applications, specifically Positron Emission Tomography (PET) and Single Photon Emission Tomography (SPECT). It is based on the GEANT4 libraries to model the interaction of particles with matter. These are validated and widely used in a range of areas such as high-energy physics and space

science. It benefits from a worldwide collaboration and is constantly being improved and refined. [16]

GATE models scanner geometry, detector electronics, and relevant particle physics. Scanners can be constructed by assembling several basic geometrical shapes together to achieve the desired geometry. Each volume can also be assigned a material from a list of custom-defined materials. GATE can model the signal processing chain of the detector, from the detection of individual events in the scintillating crystals to the production of coincidences by the electronics. Each step in this digitization chain can be individually controlled. GATE allows the user to define which types of particle are to be tracked and which interaction processes are to be simulated (Compton scattering, photoelectric effect, etc.). Thresholds for the production of secondary particles such as electrons, x-rays and δ -rays, can also be set. These allow simulations to run faster while preserving the accuracy of photon detection in PET. [15] GATE can also describe time-dependent phenomena such as moving detector elements and source decay. It uses a global virtual clock to model continuous time flow. In particular, the noncollinearity of annihilation photons is modelled with a 0.58° FWHM Gaussian blur. Finally, GATE implements a scripting language that allows the end user to control each feature by using simple keywords at the command line. These script commands can also be saved into one or several macro files for ease of use.

The OpenGATE collaboration website lists several studies that have validated simulated models of commercial PET and SPECT scanners with real data obtained from those scanners. These studies show general agreement between the results generated by GATE and the real-life scanner (see for example Lamare *et al.* and Jan *et al.*). [17,18] This demonstrates that GATE can accurately simulate various detector designs.

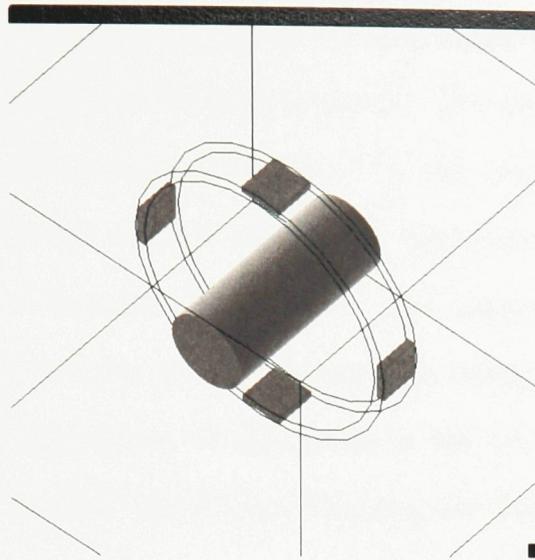


Figure 3.1: PeTrack detector design in GATE.

3.2 Detector

The PeTrack detector design implemented in GATE consisted of two pairs of modules that were 90° apart from each other and 0.5 m away from the isocentre. Each module was made up of a 3×3 array of detector blocks and each detector block comprised 13×13 BGO crystals. The crystals had a surface area of $3.9 \times 3.9 \text{ mm}^2$ and a depth of 30 mm. There was a 0.1 mm gap between each crystal, so each detector block had an effective surface area of $52 \times 52 \text{ mm}^2$. A cylindrical water phantom with a diameter of 30 cm and a length of 70 cm was placed in the centre of the detector and was aligned with the detector axis. Figure 3.1 shows the design of the system in the GATE display window.

3.3 Electronics

A photon can interact multiple times with a single detector crystal. For instance, several Compton interactions could occur before a photoelectric absorption. Each interaction is called a hit in GATE.

Thus, the first step of the digitization process is the adder module. The adder

produces a single pulse from all the hits made by a single photon in a given crystal since the electronics measure an integrated signal. The position of the pulse is the centre of mass of all the hit positions, weighted by the energy deposited in each hit. The energy of the pulse is the sum of the energy deposited in each hit. Readout of the pulses is made at the detector block level. The pulses from all crystals in the block are added together to form a single pulse with energy equal to the sum of the energy of all pulses. The position of this pulse is set to the position of the pulse with the largest energy. Spatial and energy blurring are then applied to simulate the detector's response.

The energy resolution of the system is set to 25% at 511 keV to simulate blurring of the energy spectrum of the pulses. A thresholder module rejects all photons with energy lower than 420 keV. Similarly, an upholder module rejects all photons with energy higher than 600 keV. Finally, two events are considered in coincidence if they are detected within 10 ns of each other.

3.4 Physics

Annihilation photons deposit their energy into a scintillator crystal through two types of interactions: Compton scattering and the photoelectric effect. I use the Standard Energy Electromagnetic Processes (SEP) GEANT4 libraries for these two effects, since the SEP simulates interactions at energies higher than 10 keV and the annihilation gamma rays have an energy of 511 keV. Rayleigh scattering, which does not deposit energy in the crystal, uses the Low Energy Electromagnetic Processes (LEP) libraries. This package extends the validity of the electromagnetic interactions of the SEP down to an energy of 250 eV.

3.5 Sources

Three moving markers were modelled in the simulation. The isotope ^{124}I was used because its half-life of $T_{\frac{1}{2}} = 4.18$ days would allow it to be used in a clinical setting. I used a custom-defined type for the isotope: fast ^{124}I . This is a specialized type of source in GATE. It uses a simplified decay scheme of the original positron-emitting ^{124}I . Positrons are emitted, but not neutrinos. Nuclear recoil and atomic de-excitation, such as x-rays and Auger electrons, are not simulated. These simplifications still retain the essential features of ^{124}I while speeding up the simulation. The branching ratio of 74.4% electron capture and 25.6% positron emission is respected. The initial activity per source was set to $135 \mu\text{Ci}$.

The markers were simply modelled as three 0.5-mm diameter spherical active cores. No capsules surrounding the active cores were modelled. This decision was made since the management of moving volumes and sources in GATE is cumbersome. The absence of capsules confining the active cores should increase the positron range and decrease the performance of the system. This issue is addressed in Section 3.9.

The markers were positioned in a triangular pattern and moved in a circular motion around an arbitrary axis. All markers rotated in the same plane. Each marker had a different radius of rotation, ranging from 5 to 31 mm. The speed of the markers was set to $120^\circ/\text{s}$. This corresponds to a breathing cycle of three seconds. By comparison, Nehmeh *et al.* measured the breathing cycle period to be around five seconds, whereas Chen *et al.* measured an average breathing cycle period of 2.78 seconds. [19, 20] The linear velocities of the spatial configurations were in the range of known lung tumour speeds. [21] I ran the simulation four times with different marker configurations, i.e. different initial coordinates, and axis of rotation and radius of rotation. The axis of rotation was selected randomly at the start of each simulation and was kept constant for the duration of the simulation. Table 3.1 shows the spatial

Table 3.1: Spatial configuration of the markers for each simulation run.

Run #	Marker #	Initial coordinates (mm)	Linear velocity (mm/s)	Axis of rotation
1	1	(-3.2, 0.0, 3.8)	10.5	(0.663, -0.5, 0.557)
1	2	(6.4, 0.0, -7.7)	20.9	(0.663, -0.5, 0.557)
1	3	(6.9, 15.6, 5.8)	37.7	(0.663, -0.5, 0.557)
2	1	(12.0, -12.0, 0.0)	35.6	(0, 0, 1)
2	2	(27.0, 0.0, 0.0)	56.5	(0, 0, 1)
2	3	(0.0, 31.0, 0.0)	64.9	(0, 0, 1)
3	1	(5.8, -7.7, -14.0)	35.6	(0.321, -0.766, 0.557)
3	2	(23.4, 0.0, -13.5)	56.5	(0.321, -0.766, 0.557)
3	3	(11.9, 19.9, 20.6)	64.9	(0.321, -0.766, 0.557)
4	1	(-8.5, 3.2, 9.3)	27.2	(-0.321, -0.766, 0.557)
4	2	(23.4, 0.0, -13.5)	56.5	(-0.321, -0.766, 0.557)
4	3	(11.9, 19.9, 20.6)	64.9	(-0.321, -0.766, 0.557)

configuration of the markers for each simulation run. The simulated scan ran for 18 seconds. This represents six full cycles given the angular velocity of the markers.

3.6 Localization

I localized the markers every 100 ms and repeated the localization 175 times. The localization time interval of $t = 100$ ms was selected so that, given the initial activity of the markers, approximately 100 coincidence lines per marker would be recorded during each localization time interval. This target number of coincidence lines was selected based on the work of Xu *et al.* They showed that an accuracy of approximately 0.5 mm could be achieved using 100 coincidence lines per marker per localization. [13] The tracking algorithm was initialized with the real initial position of the markers, since

I wanted to assess the accuracy of the tracking component of the algorithm only. I evaluate the performance of an initialization method in Chapter 5.

3.7 Accuracy and Precision

Accuracy and precision are two closely related concepts that need to be distinguished. Accuracy is a measure of the degree to which a measurement or a result agrees with the correct or accepted value. Accuracy is a measure of both random and systematic errors. Precision on the other hand is the degree to which repeated measurements agree with each other when performed under similar conditions. Precision does not include systematic errors.

I can describe the performance of PeTrack in terms of absolute and relative accuracy. The accuracy is equivalent to the localization error. The absolute accuracy corresponds to the 3D distance between the position of a marker localized by the algorithm and the position of the marker in the middle of the localization time interval. For instance, if the markers are localized every 100 ms, I compare the localized markers with the position of the markers in the middle of the 100 ms. I call this distance the absolute localization error since it is always positive.

The relative accuracy corresponds to the difference between the distance between a pair of localized markers and the true distance between the markers. I call this the relative accuracy (or relative localization error) because it can be negative (if the distance is shorter) or positive (if the distance is larger).

Ideally, the localization error of each individual coordinate (x , y , and z) could be evaluated for every marker. I will show that for this particular simulation, the value obtained in this manner does not correctly reflect the accuracy of the system.

During treatment, the PeTrack system will constantly be acquiring coincidence lines, except during the very short Linac pulses (approximately 10 μ s each, spaced

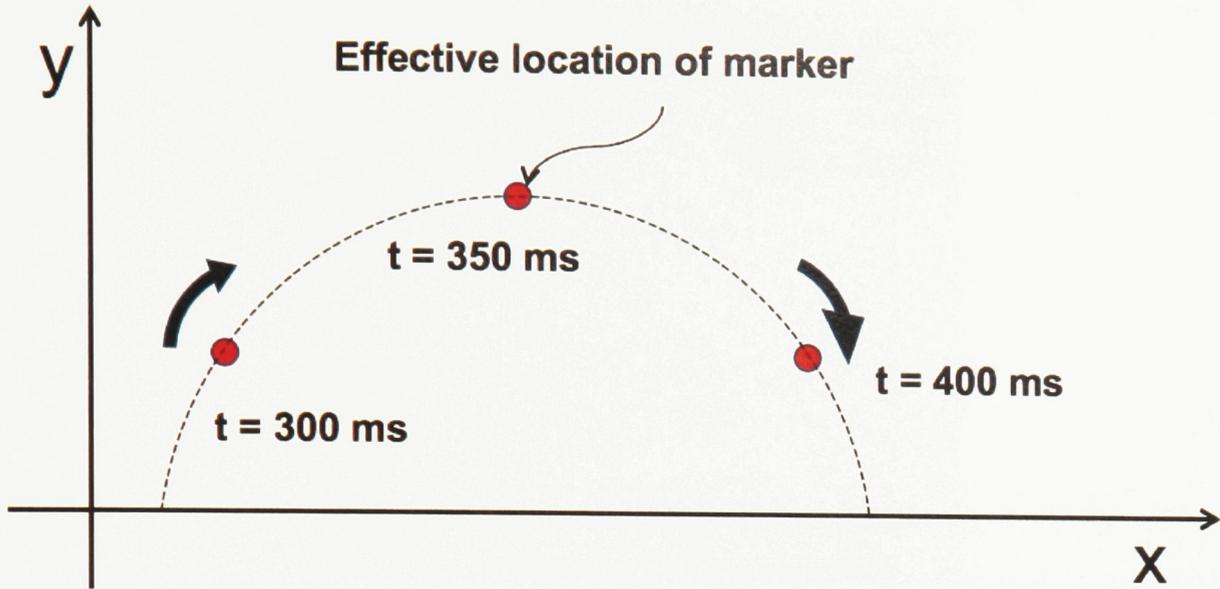


Figure 3.2: Effective location of a marker between $t = 300$ ms and $t = 400$ ms.

apart by 2.5 to 10 ms in time). The algorithm will localize the marker using all coincidence lines recorded since the last localization. The interval of time between each localization is called the localization time interval. However, during that interval of time, a marker moves along the arc of a circle. Assuming the marker maintains a constant velocity during the localization interval of time, a point in the middle of the track should approximately minimize the RMS distance to all coincidence lines recorded in the interval of time. Thus, the algorithm should localize the marker near the middle of the track. I define the effective location of a marker as the location of the marker halfway into the localization time interval. For example, the effective location of a marker between $t = 300$ ms and $t = 400$ ms is defined as the location of the marker at $t = 350$ ms (see Figure 3.2).

3.8 Results

On average, 81 coincidence lines per marker were used for the localization after rejecting scatter and random coincidences. The absolute localization error was calculated

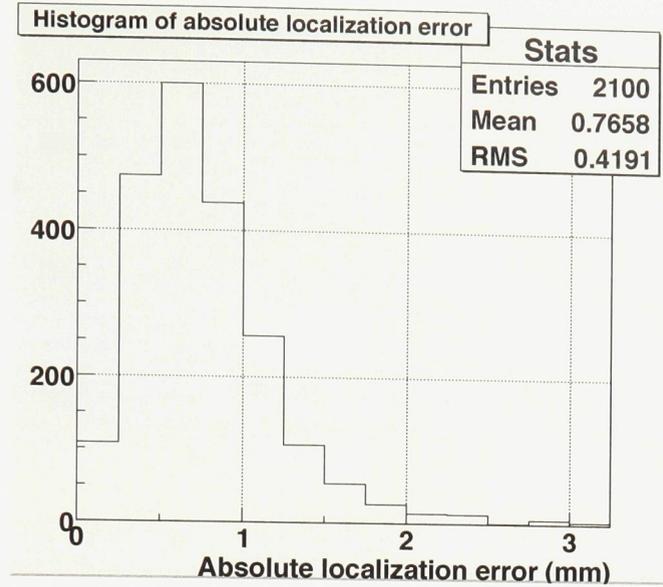


Figure 3.3: Histogram of the distribution of the absolute localization error for all runs.

for each marker for each of the 175 localization runs. Figure 3.3 reports the distribution of the absolute localization error of all three markers for each simulation run. The average absolute accuracy among all markers and runs is 0.77 ± 0.42 mm.

The relative localization error was calculated for each pair of markers for every localization run. Figure 3.4 shows the distribution of the relative localization error for each simulation run. The average relative accuracy among all markers and runs is -0.28 ± 0.69 mm.

The distribution of relative localization error in Figure 3.4 gives an estimate of the precision of the tracking. A narrow distribution indicates high precision. Therefore, the average standard deviation of the relative localization error, $\sigma_{r.l.e.}$, is the precision of the tracking. Let the precision associated with the localization of a marker be given by σ_m . Relative accuracy is associated to the distance between a pair of markers. Since each marker contributes equally to the average standard deviation of the relative

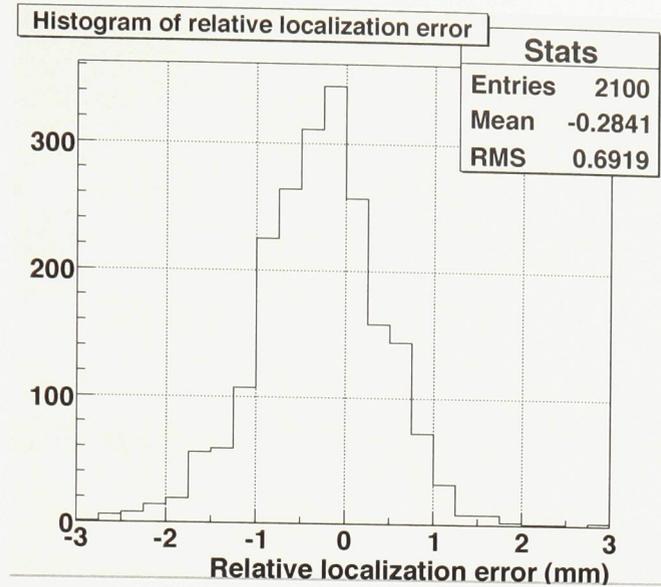


Figure 3.4: Histogram of the distribution of the relative localization error for all runs.

localization error, we have:

$$\begin{aligned}\sigma_{r.l.e.} &= \sqrt{\sigma_m^2 + \sigma_m^2} = \sqrt{2}\sigma_m \\ \frac{\sigma_{r.l.e.}}{\sqrt{2}} &= \sigma_m\end{aligned}\tag{3.1}$$

Thus, the standard deviation of the distribution of relative localization error needs to be divided by $\sqrt{2}$ in order to obtain the precision of the localization of an individual marker. The average precision among all markers and runs is 0.49 mm.

I calculated the localization error of coordinates x , y , and z separately for all markers after each localization. I averaged the error of all localization runs. Table 3.2 reports the mean localization error of the Cartesian coordinates for all markers and runs. The average error is 0.005 ± 0.048 mm on x , 0.02 ± 0.05 mm on y , and -0.007 ± 0.042 mm on z .

Figure 3.5 shows both real and localized tracks of the three markers for a sample simulation run.

Table 3.2: Localization error averaged over all localizations for each marker and each run.

Run #	Marker #	δ_x (mm)	δ_y (mm)	δ_z (mm)
1	1	0.017	0.014	-0.042
1	2	0.031	0.032	-0.013
1	3	0.021	0.055	0.005
2	1	0.132	-0.054	-0.011
2	2	-0.006	0.128	0.005
2	3	-0.010	-0.117	0.000
3	1	-0.036	-0.013	0.032
3	2	-0.015	0.083	0.005
3	3	-0.007	-0.002	0.007
4	1	-0.013	0.011	0.018
4	2	-0.047	0.041	-0.032
4	3	-0.050	0.006	0.004
Mean		0.001 ± 0.048	0.02 ± 0.06	-0.002 ± 0.020

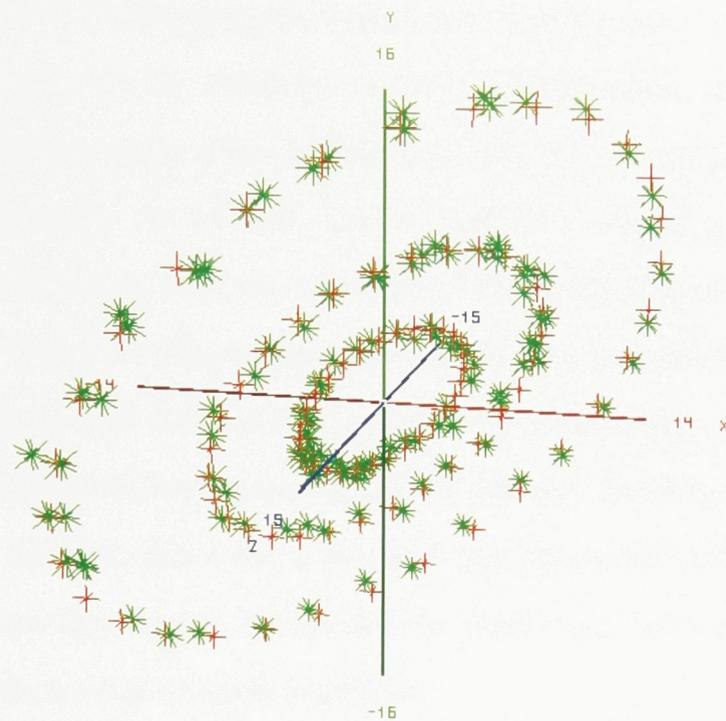


Figure 3.5: Localized track of the markers (green stars) and real track of the markers (red crosses) for a sample run. All markers move in the same plane.

3.9 Discussion

Though the localization error of the three coordinates is in the range of -0.007 to 0.02 mm when averaged among all runs (see Table 3.2), this result should be taken with a caveat. It does not reflect the actual accuracy of the system. Figure 3.4 shows that the mean of the relative localization error is negative for all simulation runs. This suggests that a systematic error in the localization of the markers is bringing the localized markers closer together than they actually are. I believe this is due to the circular motion of the markers during the localization. During each 100 ms interval, the markers move by a few millimetres along an arc and generate coincidence lines all along this track. As a result, the localization is biased towards the centre of mass of the arc, which is located on the inside of the track. Because all markers share the same centre of rotation (the origin of the coordinate system), the net effect of this localization bias is to pull the markers towards the same point. This systematically leads to a smaller distance between each pair of markers (see Figure 3.6). For a given localization time interval, this error scales with the curvature of the arc and with the marker speed. Since the motion is circular, the systematic error gets averaged out of the mean localization error of each coordinate.

Nevertheless, the mean localization error of the coordinates displays an interesting feature: the error on x and y for the second simulation run is noticeably higher than for other runs. I believe this represents a special case. The four detector modules are aligned with the x and y axis and in this instance, all three markers are rotating in the $x-y$ plane. In this configuration, two markers are more likely to end up approximately on the same line that joins two opposite detector modules. This scenario would make assigning the coincidence lines to the correct marker more difficult. For example, if two markers are approximately lined up on the x axis, their y coordinates might be harder to resolve since coincidence lines are being wrongly assigned between the two

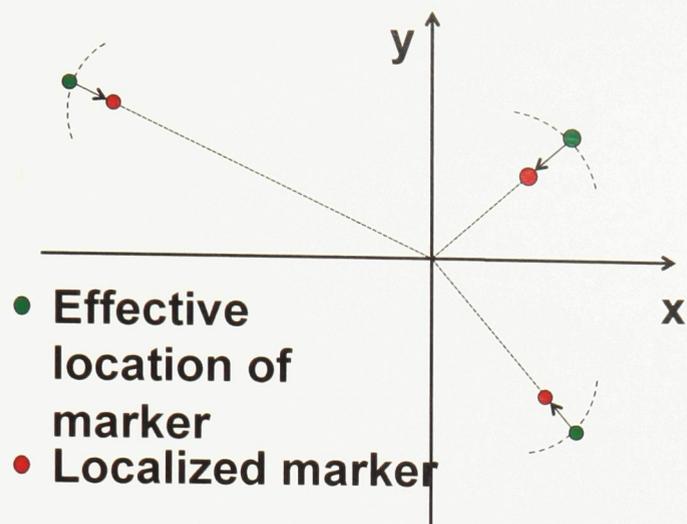


Figure 3.6: Systematic shrinking of the distances between the markers due to the curvature of the path of the markers.

markers.

The scale of the systematic error introduced by the circular motion of the markers can be estimated by constructing the distance vector, i.e. the vector that points from the position of a localized marker (i.e. the location of the marker as calculated by the localization algorithm) to its effective location along the middle of the track and projecting this vector onto the radial vector, i.e. the normalized vector that points from the effective location of the marker to the centre of rotation. Figure 3.7 shows the vectors described above and Figure 3.8 shows the distribution of the projection of the distance vector onto the radial vector for one simulation run. The systematic error is in the range of -0.10 mm to -0.37 mm for all markers and simulation runs. I expect a similar range of systematic error if the localization is applied to real tumour tracking because the range and speed of the simulated motion is similar to the range and speed of real tumour motion.

Results from all four simulation runs show that the PeTrack algorithm achieves submillimetre accuracy when using 81 coincidence lines per marker. Xu *et al.* had achieved 0.55 ± 0.27 mm absolute accuracy with only 62 coincidence lines per

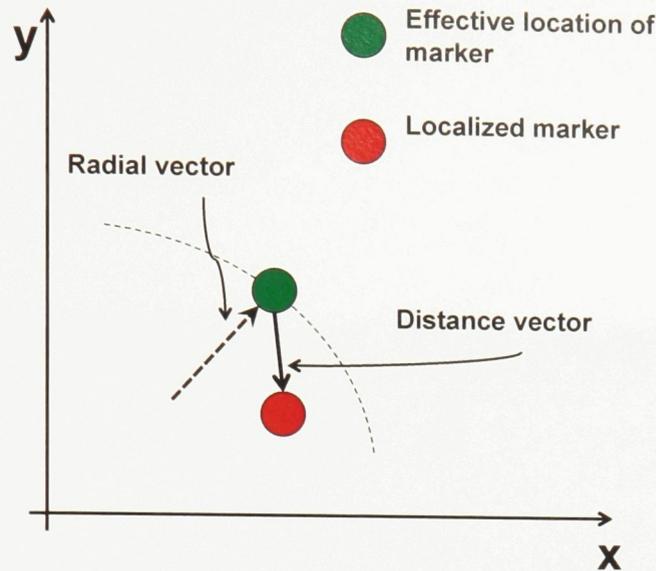


Figure 3.7: To estimate the scale of the systematic error, the distance vector is projected onto the normalized radial vector.

marker. [13] Using more lines should improve the accuracy, but in this case the accuracy is lower. This may be explained by the fact that the simulation did not model a capsule around the markers. In the simulation, the markers were placed directly in water. In Xu *et al.*, the markers were encapsulated in a 0.2-mm thick gold shell. The average maximum energy (weighted by the branching ratio) of positrons emitted by ^{124}I is 1.811 MeV. The average energy of the positrons is approximately 604 keV, or $\frac{1}{3}$ of the average maximum energy. The continuous slowing down approximation (CSDA) range of positrons with the average energy is 0.22 mm in gold and 2.27 mm in water. I conclude that positron range in water is not negligible for ^{124}I and degrades the accuracy of the localization.

Scatter and random coincidences can introduce significant errors in the localization of the markers. I can estimate how well the algorithm rejects scatter and random coincidences by looking at the fraction of lines rejected in each localization over the total number of lines collected in each localization time interval. This fraction should be approximately equal to the sum of scatter and random coincidences (these number

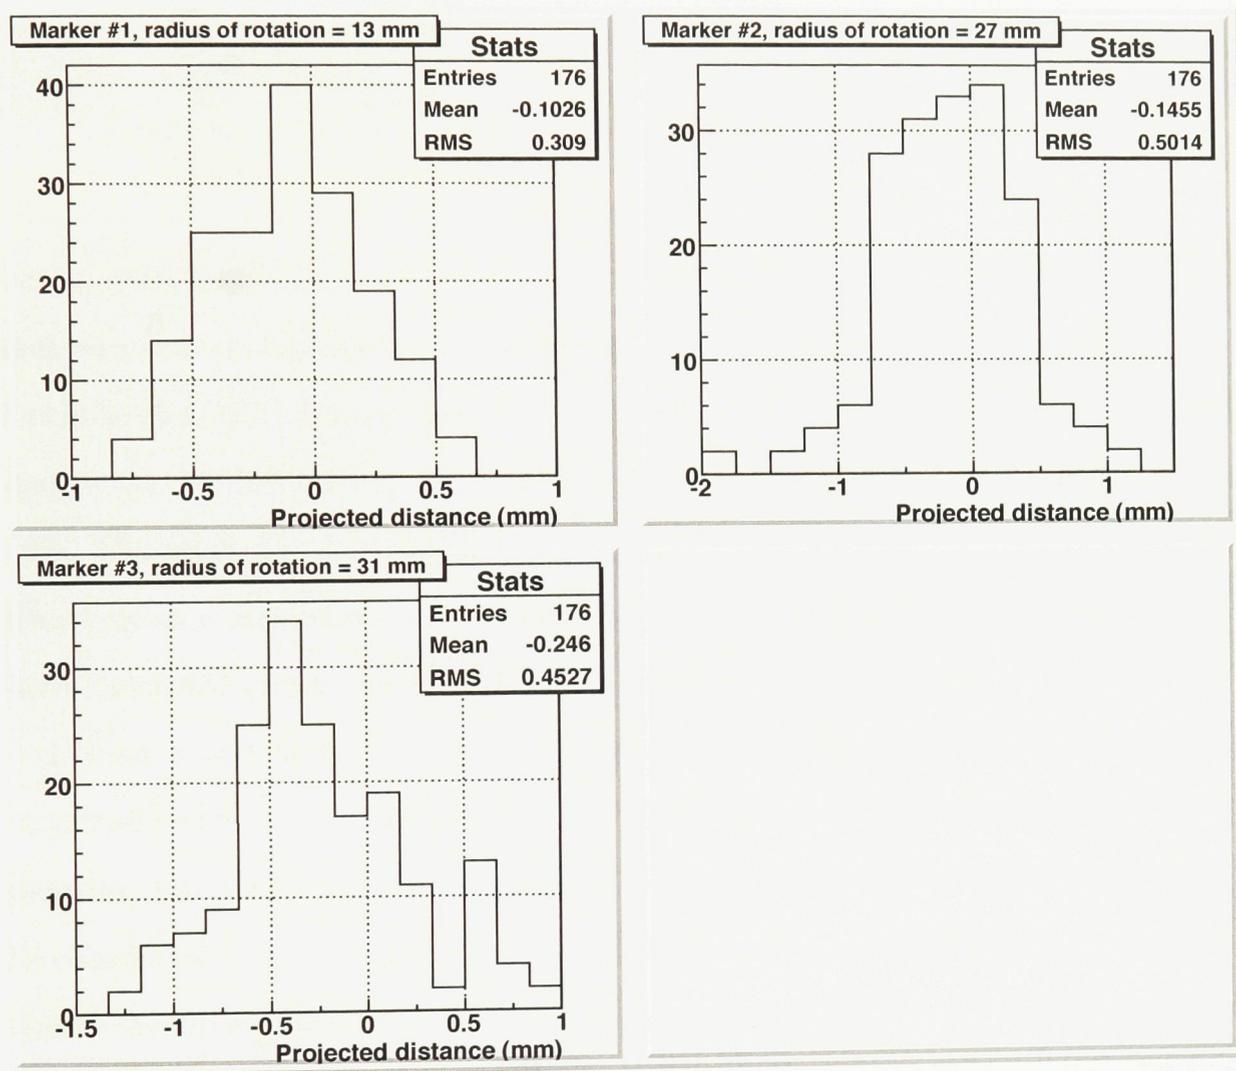


Figure 3.8: Histogram of the projected distance vector onto the normalized radial vector.

Table 3.3: Fraction of lines rejected by the algorithm and scatter and random fractions reported by GATE.

Run #	Rejected lines	Scatter	Random	Scatter+Random
1	0.179	0.149	0.014	0.162
2	0.014	0.143	0.013	0.156
3	0.198	0.155	0.015	0.170
4	0.236	0.158	0.014	0.171
Mean	0.157	0.151	0.014	0.165

are reported by GATE at the end of each simulation). Table 3.3 shows the fraction of lines rejected by the algorithm and the actual number of scatter and random coincidences from GATE. On average, 15.7% of lines were unused by the algorithm. GATE reports the scatter fraction to be on average 15.1% and the random fraction to be 1.4%. The total of scatter and random coincidences is thus 16.5%. The PeTrack algorithm rejects a very reasonable amount of scatter and random coincidences; therefore these corrupted events should not impact adversely the accuracy.

I bring attention to the fact that run #2 rejected a very small fraction of lines compared to the other runs. I believe this is due to the motion of the markers being constrained to the $x-y$ plane. As explained above for the relative localization error of the coordinates, this configuration can make it hard to correctly assign the coincidence lines to the right marker. In this case, I believe that scatter and random coincidences generated by a marker are getting incorrectly assigned to another marker. This is a special case and seems to happen only when motion is constrained to the $x-y$ plane.

This simulation includes some limitations. Some elements of the detector design are not accurate, notably the electronics. A single value for blurring is modelled while in reality each crystal could have a different energy resolution. Scintillating light from a crystal could also spill over to neighbouring crystals, causing crosstalk.

This effect is absent from the simulation. Other effects such as dead time and pile up are also absent. Random noise was also omitted. The absence of these factors could have caused the intrinsic resolution of the detector to be better than it would be in a clinical setting.

3.10 Conclusion

This chapter presented a simulation study of three moving positron sources. The events were recorded by a model of the PeTrack detector design. On average, the accuracy and the precision of the tracking is below 1 mm and will not be degraded by either scatter or random coincidences. A systematic error is introduced when the markers move along a curved path.

Chapter 4

Experimental Evaluation of PeTrack

The previous chapter suggested that the PeTrack algorithm can achieve submillimetre accuracy and precision. However, limitations in the modelling of the detector electronics motivate a study with a real-life PET scanner. In this chapter, I present three experiments that my supervisor and I conducted to assess the accuracy and precision of the PeTrack system using real-life detectors and more diverse movement patterns.

4.1 Preliminary Evaluation of PeTrack

This section presents an experimental study conducted at the Ottawa Hospital Cancer Centre. The goal is to do a preliminary experimental evaluation of the accuracy and precision of the tracking algorithm.

4.1.1 PET System

The Philips GEMINI PET system on the premises of the Ottawa Hospital Cancer Centre was used to conduct the experiment. This whole-body scanner is comprised of both a computed tomography (CT) and a PET scanner. The latter is a Philips Allegro PET system. This scanner is a full-ring detector made up of 28 detector blocks, each

with 22×29 gadolinium oxyorthosilicate (GSO) crystals. Each crystal has a surface area of $4 \times 6 \text{ mm}^2$ and a thickness of 20 mm. The ring diameter is 86.41 cm and the field of view (FOV) is 18 cm in the direction of the detector axis and 57.72 cm in the transverse direction. The timing window for coincidences is 7.5 ns. The spatial resolution near the centre is approximately 5.53 mm in the transverse direction and 5.64 mm in the axial direction. [22]

4.1.2 Sources

Ideally I would have used ^{124}I in this study. Unfortunately, its relatively short half-life makes it ill-suited to conduct multiple experiments over the span of several months. My research group's supply of ^{124}I would need to be replenished constantly. Instead, I used three ^{22}Na sources as markers. The isotope ^{22}Na has a half-life of 2.6 years, which makes it more suitable for experimental studies conducted every few months. Each source consists of an active core shaped as a 1-mm diameter sphere. The core is sealed in the centre of a cylindrical acrylic disk with a radius of 25.4 mm and a height of 0.635 mm. The activities of the sources on the day of the experiment were 18.1, 18.8 and 20.9 μCi .

The three markers were placed in a triangular pattern on a plastic platform. A small motor made the platform rotate around its central axis. The design is shown in Figure 4.1. The radius of rotation of the markers and the speed of the motion were chosen so that the linear velocities of the markers would match the known range of speed of tumour motion. I performed three scans, each with a different speed of motion. The linear velocities of the markers varied between 14.6 mm/s and 53.6 mm/s. Each scan lasted 50 seconds. During the first 10 seconds of the data acquisition, the markers were kept motionless in order to facilitate the initialization of the tracking algorithm (see Section 4.1.5).

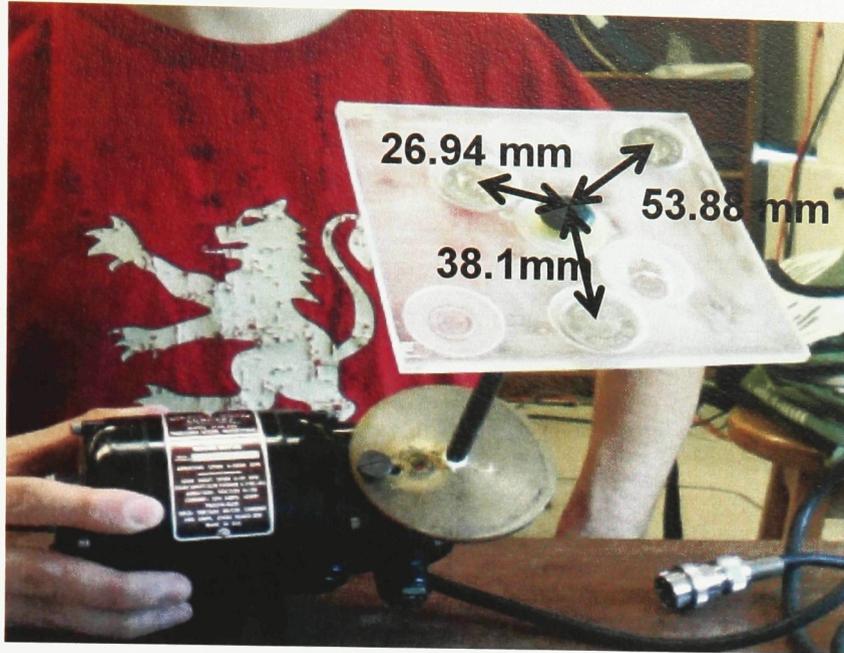


Figure 4.1: Setup of the experiment.

4.1.3 Compact List Mode Data Format

The PeTrack algorithm uses the endpoints coordinates of the recorded coincidence lines in order to localize the markers, but PET scanners do not usually store the Cartesian coordinates of the coincidence lines. PET system manufacturers opt for different ways to store the data, such as 2D or 3D histograms (known as sinograms), or List Mode format. In this experiment, I used the List Mode format to record the data. In List Mode, a number of integer indices uniquely identify each coincidence line. For this particular experiment, I selected the Philips Compact Acquisition List Mode Format, which is a proprietary format available on the Philips system used. Four indices stored in this format are relevant for each coincidence line.

Indices z_a and z_b are the two crystal numbers in the axial direction. There are 29 crystal rows in the axial direction, so z_a and z_b run from 0 to 28. Index ϕ is the transverse angle of a coincidence line with respect to the x axis and index r is the perpendicular distance of a coincidence line to the isocentre. In addition, an elapsed time stamp is recorded in the List Mode file. This identifies the time from the start

of the acquisition.

Once the data were acquired, I used MATLAB routines to read the proprietary Philips List Mode data and extract the relevant information: the four indices described above, the time stamp for each coincidence line, and header files that contain critical information about the scanner characteristics.

4.1.4 Conversion from Compact List Mode Data Format to Cartesian Coordinates

The implementation of the localization algorithm requires the Cartesian coordinates, (x_1, y_1, z_1) and (x_2, y_2, z_2) , of the two endpoints of each coincidence line. The z_a and z_b indices were converted to z -coordinates by using the crystal pitch in the axial direction. This value was obtained from the main header file of the List Mode format. The equations for the z -coordinates (in mm) are:

$$z_1 = \delta_a(z_a - 14); \quad (4.1)$$

$$z_2 = \delta_a(z_b - 14); \quad (4.2)$$

where $\delta_a = 6.3$ mm is the crystal pitch in the axial direction. This ensures that a value of 14 for the z_a and z_b indices corresponds to $z = 0$ (the centre of the scanner) in the Cartesian coordinate system. The index ϕ takes 161 different integer values, from $\phi = 0$ to $\phi = 160$, and covers 180° since a coincidence line angled at 181° is equivalent to a coincidence line angled at 1° . Thus, the conversion from the index ϕ to radians is given by:

$$\phi \text{ (in radians)} = \pi\phi/161. \quad (4.3)$$

The Philips List Mode format uses 295 radial bins, r , that run from 0 to 294. The radial bins cover the system's transverse FOV of 57.6 cm. The conversion from the

index r to millimetres is given by:

$$r \text{ (in mm)} = \delta(r - 147), \quad (4.4)$$

where $\delta = 2.15$ mm corresponds to the pitch between each radial bin. Equation 4.4 ensures that a value of 147 for the radial bin corresponds to a coincidence line that passes through the isocentre. The indices ϕ (in radians) and r (in mm) were converted to Cartesian coordinates, (x_1, y_1, z_1) and (x_2, y_2, z_2) , by considering the geometry shown in Figure 4.2, where R is the radius of the scanner (in mm). The correspondence between indices and Cartesian coordinates x and y (in mm) is:

$$\begin{pmatrix} x_1 \\ y_1 \end{pmatrix} = R \begin{pmatrix} \cos(\phi + \psi) \\ \sin(\phi + \psi) \end{pmatrix}, \quad (4.5)$$

$$\begin{pmatrix} x_2 \\ y_2 \end{pmatrix} = R \begin{pmatrix} \cos(\phi - \psi) \\ \sin(\phi - \psi) \end{pmatrix}, \quad (4.6)$$

where $\psi = \arccos\left(\frac{r}{R}\right)$.

4.1.5 Localization

The markers were localized every 15 ms for 2000 runs. The time interval was set so that approximately 80 coincidence lines would be used per marker per localization after rejecting scatter and random coincidences. This corresponds to approximately the same number of coincidence lines used in the simulation study. The tracking algorithm was initialized with the estimated initial location of the markers. The initial position of the markers was determined by looking at the images recorded by

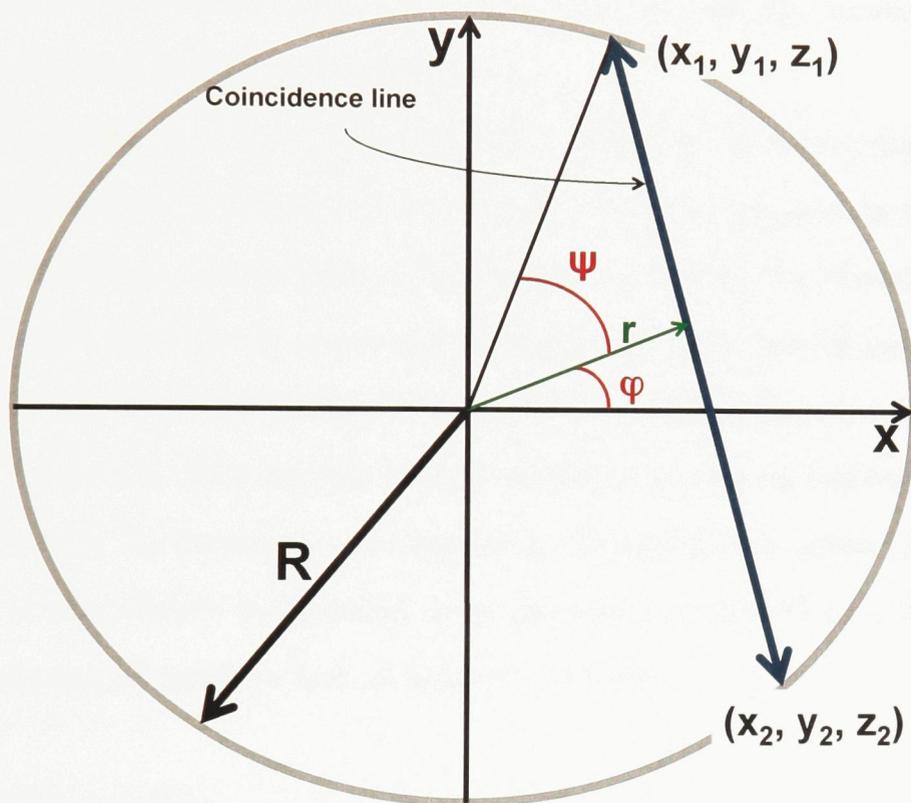


Figure 4.2: Projection in the $x - y$ plane of the geometrical correspondence between Cartesian coordinates and List Mode indices.

the PET camera during the scans. These images were reconstructed by the Philips Allegro system's software. During the first 10 seconds of the acquisition, the markers were motionless and appear as hot spots on the reconstructed images. I used the pixel pitch of 4 mm to estimate the initial location of the markers.

4.1.6 Results

On average, 86 coincidence lines per marker were used for the localization after rejecting scatter and random coincidences.

In this experiment, there is no easy way to know the true absolute location of a marker in the frame of reference of the scanner. Only the distance between each pair of markers is known. Accordingly, I cannot give a value for the absolute accuracy.

The relative localization error was calculated for each pair of markers for every localization run. The distribution of the relative localization error for each scan is shown in Figure 4.3. The average relative accuracy among all markers and all scans is 0.6 ± 1.8 mm. As mentioned in Chapter 3, the precision is given by the standard deviation of the relative localization error distribution divided by $\sqrt{2}$. The average precision among all markers and all scans is 1.25 mm.

4.1.7 Discussion

Compared to the simulation study presented in Chapter 3, the results of this experiment show degradation in both relative accuracy and precision of the tracking, even though the same number of coincidence lines per marker were used in the localization. Simulation results suggested a relative accuracy of -0.28 ± 0.69 mm whereas the relative accuracy evaluated in this experiment is 0.6 ± 1.8 mm. The value of the precision increased to 1.25 mm compared to the 0.49 mm precision obtained in the simulation.

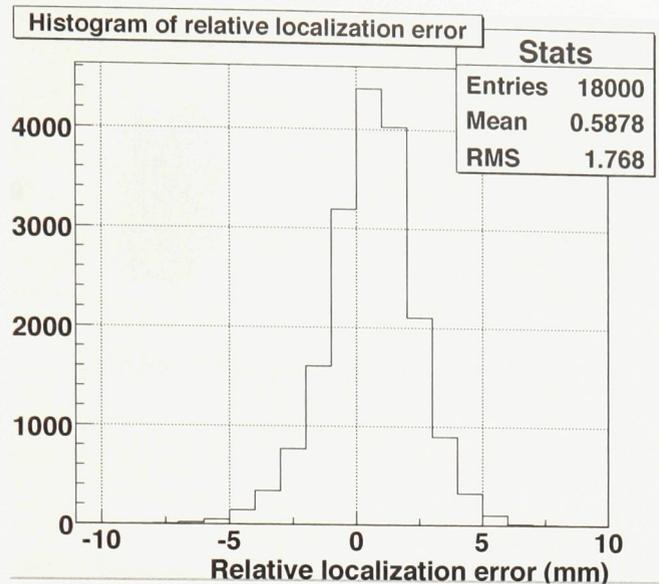


Figure 4.3: Histogram of the distribution of the relative localization error for all scans.

There are a number of factors influencing the accuracy and precision of the tracking in this experiment. Though the simulation study did not accurately model the electronics of a real-world scanner, I doubt the effects of dead time and pile up could have such a significant effect on the localization of the markers. Data rebinning and the conversion from List Mode format to coincidence lines may have played a larger role.

According to the main header file extracted from the List Mode data, the acquisition format places a virtual crystal between each detector block. The total number of crystals is $28 \times (22 + 1) = 644$. There are more than 300 crystal combinations, but only 161 angular bins are used to cover 180° . This indicates that the data are being rebinned, i.e. four coincidence lines shared between two adjacent crystals are being counted only as one coincidence line. This is illustrated in Figure 4.4. The rebinning leads to a poorer resolution since information about the original coincidence lines is lost.

In this experiment, the markers are rotating around the central axis of the platform on which they rest. This motion is identical to the one used in the simulation. The

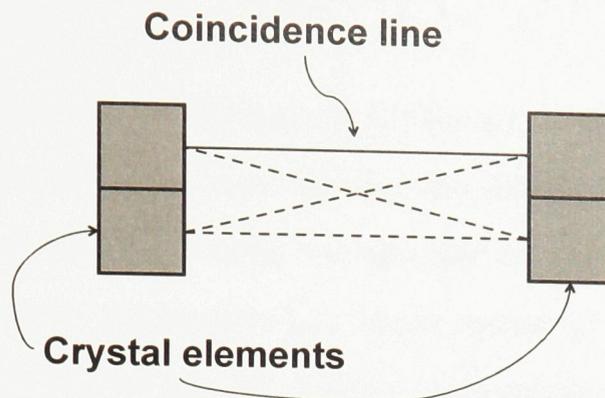


Figure 4.4: The three dashed coincidence lines are binned into the single solid line.

relative accuracy is expected to be negative, indicating the systematic error caused by the curvature of the path of the markers. Instead, the positive relative accuracy obtained indicates the localized markers are further apart than the true markers are. I believe this is due to the conversion from indices ϕ and r of the List Mode format to Cartesian coordinates x , y , and z . It is difficult to find the true relationship between the indices and the Cartesian coordinates since this is proprietary information owned by the manufacturer of the scanner. I assumed that the radial bins were evenly spaced, though current results suggest this might not be the case.

Another drawback of this experiment is the large amplitude of motion compared to real tumour motion. For example, our maximum amplitude of motion was 107.76 mm whereas in a study by Shirato *et al.*, the maximum amplitude observed was 24.6 mm. [21] Also, since the markers should be implanted at the tumour site, the distance between the three markers should be of the scale of the size of lung tumours. In this experiment, the distance between the markers varied between 60.2 and 85.2 mm. However, in a study by Harada *et al.*, the average tumour size among their patients varied between 9 and 38 mm, which is much less than half the size that I used. [23]

4.1.8 Conclusion

This section presented an experimental study conducted at the Ottawa Hospital Cancer Centre. The results obtained were significantly different than results from the simulation study. The difference could be explained by an error in the conversion from the Philips Compact Acquisition List Mode format to Cartesian coordinates. Furthermore, the amplitude of motion and the configuration of markers were unrealistic. In order to obtain more reliable results, an improved experimental protocol needs to be established.

4.2 Assessment with an Improved Experimental Protocol

This section presents a second experimental study of the tracking algorithm. The goal is to improve on the first experimental protocol, notably with regard to the motion of the markers and their configuration in space.

4.2.1 PET System

I used the same PET system as in the previous experiment. The system specifications and design is described in Section 4.1.1.

4.2.2 Sources

I used the same three ^{22}Na sources from the previous experiment (see Section 4.1.2). The activity of the sources on the day of the experiment was 14.3, 14.9, and 16.5 μCi respectively.

4.2.3 X-Y Plotter

One of the shortcomings of the previous simulation and experimental studies is the idealistic circular motion of the markers. To address this point, I used an X-Y plotter (Model HR-93, Houston Instrument, USA) to generate more complex movement patterns. An X-Y plotter is an instrument made up of two orthogonal axes (the x and y axis) that are capable of independent translation on a flat plotting surface. A computer controlled digital/analog input/output module (USB-1208FS, Measurement Computing Corporation, USA) connected to a computer was used to generate a custom analog signal that was fed to the input of the X-Y plotter. This allowed us to drive the axes in very specific patterns. The amplitude and frequency of motion were controlled from the computer.

The markers were on a Lucite platform resting on the intersection of the two axes of the plotter. The markers were positioned in a tightly packed triangular formation so that the distance between each pair of markers was effectively the same as the diameter (25.4 mm) of each capsule containing the active core. Figure 4.5 shows the experimental setup.

The X-Y plotter contains several metallic components inside its casing. The motors are the largest elements and are located near the edges of the surface. The plotter will increase the scatter and attenuate the coincidence lines. This is desirable as organs and bony structures have similar effects in a real clinical setting.

4.2.4 Motion Patterns

Several scans were recorded with the platform and the markers moving together in different patterns. Four scans involved patterns composed of a sine wave driving each axis.

If both axes are moving in phase with each other, the markers will follow a straight

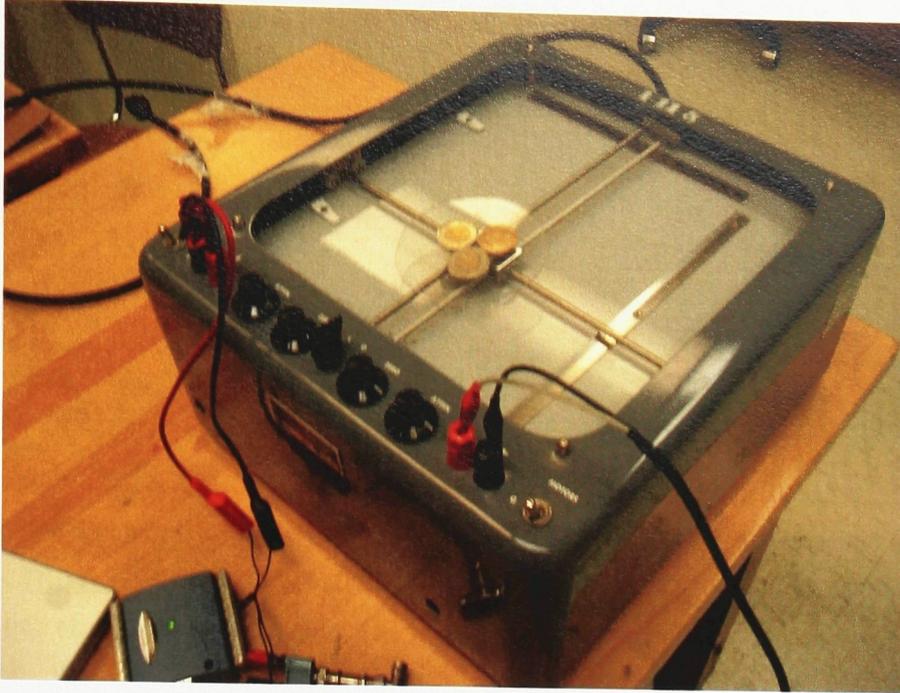


Figure 4.5: X-Y plotter used in the experiment. The marker configuration is shown with coins on the picture.

path. However, if a phase difference is introduced between the axes, one axis will lag behind the other and the markers will move on an elliptical path (see Figure 4.6). This corresponds to what is called hysteresis in a patient. Hysteresis is defined in Keall *et al.* as the lagging of the tumour motion behind the muscular contraction. [4] This causes the tumour to move along a different path during exhalation and inhalation. Seppenwoolde *et al.* found hysteresis in tumour motion in 10 out of 20 patients. [24] A phase difference was introduced between the two axes of the X-Y plotter in order to simulate hysteresis in the tumour motion.

Four more scans were done with different motion patterns. Two of these scans (#5 and #6) were done with the Lucite platform moving along a path generated by the intersection of two arctangent curves. The last two scans (#7 and #8) were done with the Lucite platform moving according to a real breathing pattern obtained from an animal study (Department of Radiology, University of California, Irvine). The X axis was driven to match the displacement of the lungs. Since I wanted the x and y

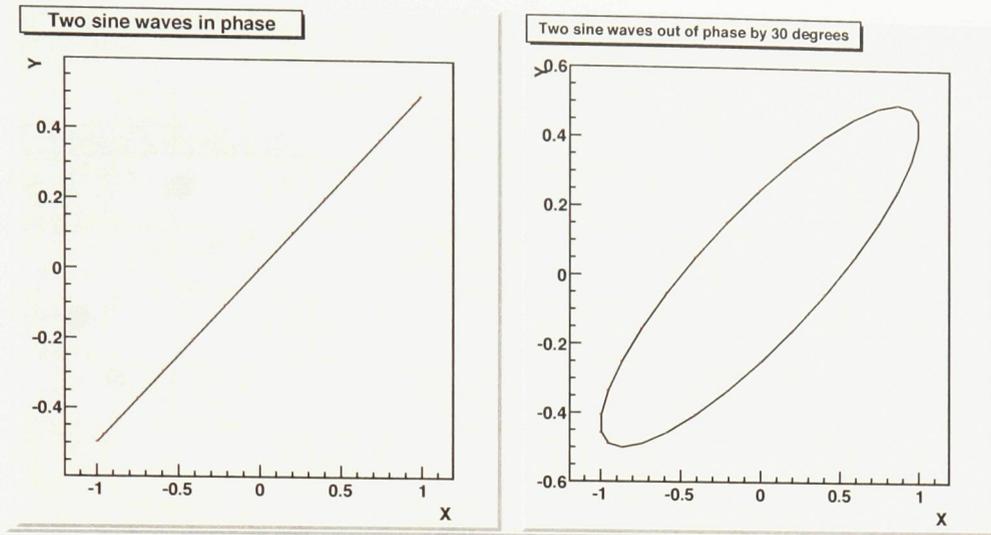


Figure 4.6: Illustration of hysteresis. On the left, two sine waves are in phase. On the right, they are out of phase by 30° .

components of the motion to be different, I fed the Y axis with a motion matching the exhaled breath temperature profile of the animal. The temperature profile was measured by a thermocouple inserted in the tracheal tube. To simulate hysteresis, I introduced a phase difference between both axes. The resulting 2D paths for scans #5 to #8 are shown in Figure 4.7.

The frequency of the motion was set to 0.3 Hz. The peak-to-peak amplitude of motion ranged from approximately 20.2 to 61.8 mm, which is more clinically relevant to tumour motion than the amplitude of the previous experiment. Every scan was 30 seconds long.

Table 4.1 summarizes the characteristics of each scan.

4.2.5 Crystal ID List Mode Data Format

In the previous experiment, I recorded data using the Philips Compact Acquisition List Mode Format (see Section 4.1.3). This format provides poor angular sampling and the exact conversion from List Mode to Cartesian coordinates is unknown.

For this experiment, I opted to acquire data using the Philips Research Crystal

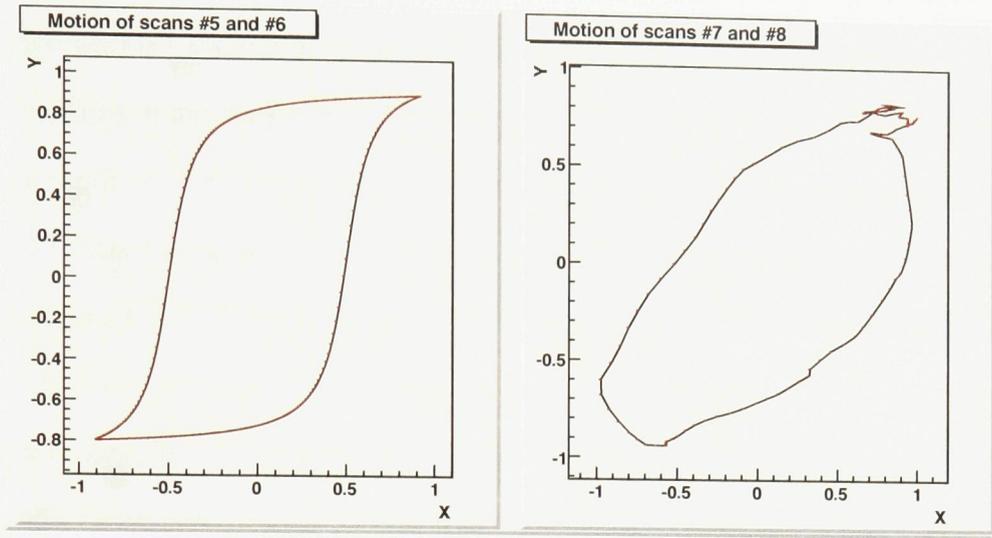


Figure 4.7: 2D path of the markers for scans #5 to #8. The actual amplitude of motion is different than shown.

Table 4.1: Characteristics of each scan of the experiment. A_x and A_y refer to the amplitude of motion along x and y , respectively.

Scan #	Motion	A_x (mm)	A_y (mm)	Phase difference
1	Sinusoidal	61.8	60.6	60°
2	Sinusoidal	61.8	60.6	30°
3	Sinusoidal	31.0	30.2	30°
4	Sinusoidal	31.0	30.2	60°
5	Arctangent	31.0	30.2	NA
6	Arctangent	20.6	20.2	NA
7	Real breathing	20.6	20.2	NA
8	Real breathing	31.0	30.2	NA

Acquisition List Mode Format. In this mode, four indices are associated to each coincidence line. The indices can be used to uniquely identify each crystal element. This is more precise than the Philips Compact Acquisition List Mode Format where coincidence lines from adjacent crystals are binned together. Two indices indicate the crystal rows in the axial direction in which the coincidence event was detected. As in the Philips Compact Acquisition List Mode Format, these indices are denoted by z_a and z_b and run from 0 to 28. Two other indices indicate the crystal number in a given row of the two crystals that detected the event. These run from 0 to 615 because there are $28 \times 22 = 616$ crystals per row. I used MATLAB routines to extract the four indices above as well as the time stamp for each coincidence line in the List Mode file.

4.2.6 Conversion from Crystal ID List Mode Data Format to Cartesian Coordinates

Converting crystal ID indices to the Cartesian coordinates, (x_1, y_1, z_1) and (x_2, y_2, z_2) , of the two endpoints of each coincidence line is relatively easy since the geometry of the Allegro PET scanner is known accurately.

The conversion between indices z_a and z_b to the Cartesian z coordinate is the same as Equations (4.1) and (4.2). Crystal ID numbers can be converted to x and y coordinates by analysing the geometry of the detector blocks (see Figure 4.8). Coordinates x and y correspond to the centre of each crystal element. The crystal pitch is 4.3 mm, the crystal depth is 20 mm, the radius R of the detector is 432.05 mm, and the crystal index η starts at 0. The angle θ is the angle subtended by one detector block at the centre of the scanner. In the case of detector block #1, x and y (in mm) are given by:

$$x = 4.3(\eta - 10.5); \quad (4.7)$$

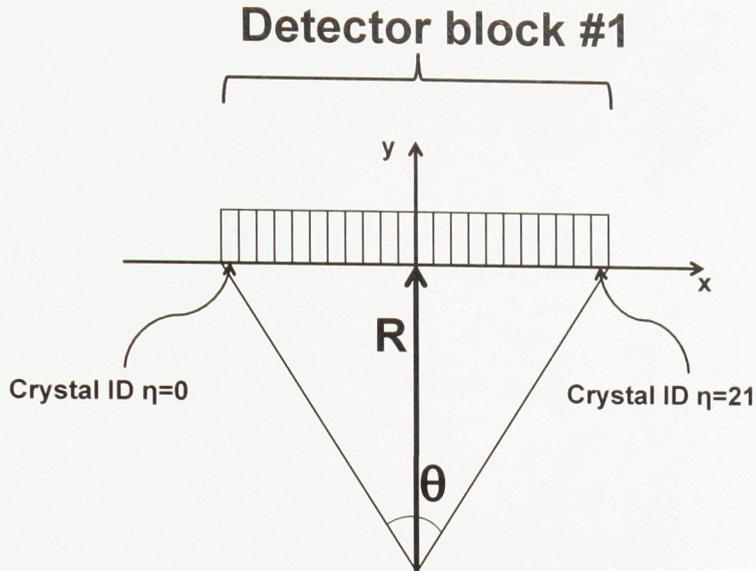


Figure 4.8: Geometry of a detector block from the Philips Allegro PET system (not to scale).

$$y = R + 10 = 442.05. \quad (4.8)$$

Coordinates for crystals belonging to other detector blocks are easily obtained since an arbitrary detector block M has simply been rotated by $(M - 1)\theta$ degrees around the origin, from the position of detector block #1.

4.2.7 Localization

The markers were localized every 100 ms for 250 runs. Approximately 2000 coincidence events were collected per localization interval. This is higher than in the simulation study (~ 300 coincidences per localization), even though the activity of the markers was 15% of the activity in the simulation. This higher count rate is accounted for by the higher geometrical efficiency of the PET scanner and the absence of a water phantom in the experiment. A higher count rate should result in an improved accuracy and precision. In order to make the following results comparable to the simulation study, I randomly selected only a fraction of the total coincidence lines per localization interval. I aimed to use the same average number of lines per

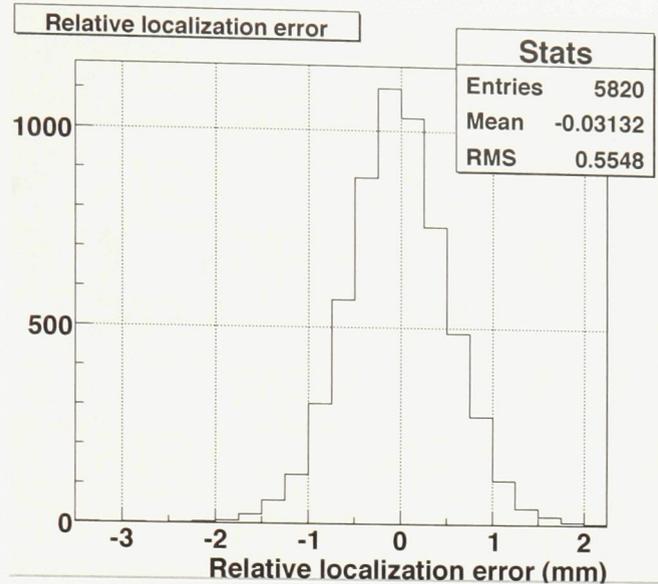


Figure 4.9: Histogram of the distribution of the relative localization error for all scans.

marker per localization after rejecting scatter and random coincidences, i.e. approximately 80. The tracking algorithm was initialized with the method to be introduced in Chapter 5.

4.2.8 Results

On average, 88 coincidence lines per marker were used for the localization after rejecting scatter and random coincidences. Once again, the true position of the markers in the frame of reference of the scanner is unknown for this experiment. Therefore, the absolute accuracy cannot be estimated.

The relative localization error was calculated for each pair of markers for every localization run. The distribution of the relative localization error for each scan is shown in Figure 4.9. The average relative accuracy among all markers and all scans is -0.03 ± 0.55 mm. The average precision among all markers and all scans is 0.39 mm.

The motion in the first four scans was generated by sinusoidal waveforms driving each axis. I can fit the coordinates of the localized markers to a sine wave. I used the data analysis software Origin (OriginLab Corporation, USA) and fitted separately

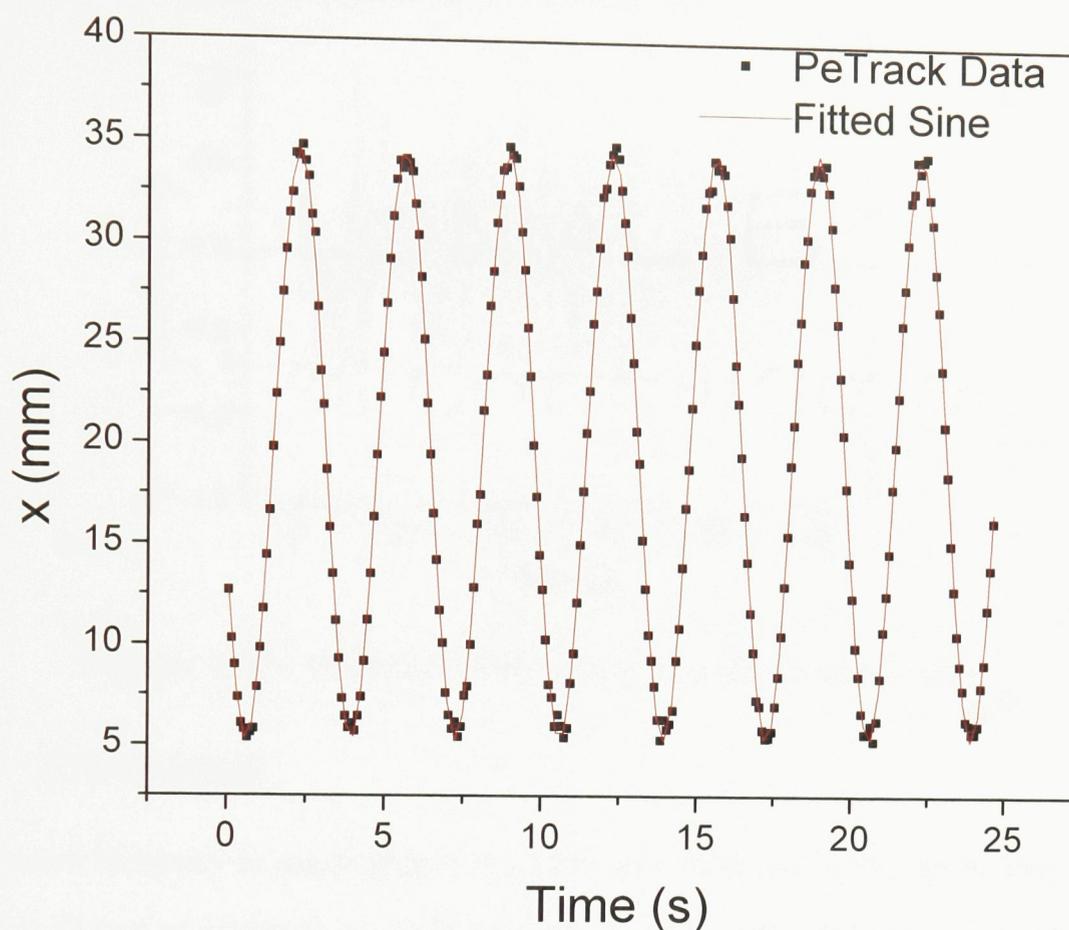


Figure 4.10: PeTrack data of scan #4 fitted to a sine wave. Only the x component is shown.

the x and y coordinates versus time of a marker from scan #4. Figure 4.10 shows the experimental data and the fitted curve for coordinate x . Figure 4.11 displays the residuals of the fit. The mean of the residuals was approximately 0. The root mean square error (RMSE) was 0.43 mm in x and 0.46 mm in y . The amplitude of the waveform was found to be 14.67 ± 0.04 mm in x and 14.77 ± 0.05 mm in y .

I also fitted the x coordinate of a marker from scan #8 against the animal breathing data used to generate the motion. Figure 4.12 shows the experimental data and the fitted curve. The RMSE was 0.64 mm and the mean of the residuals was approximately 0.

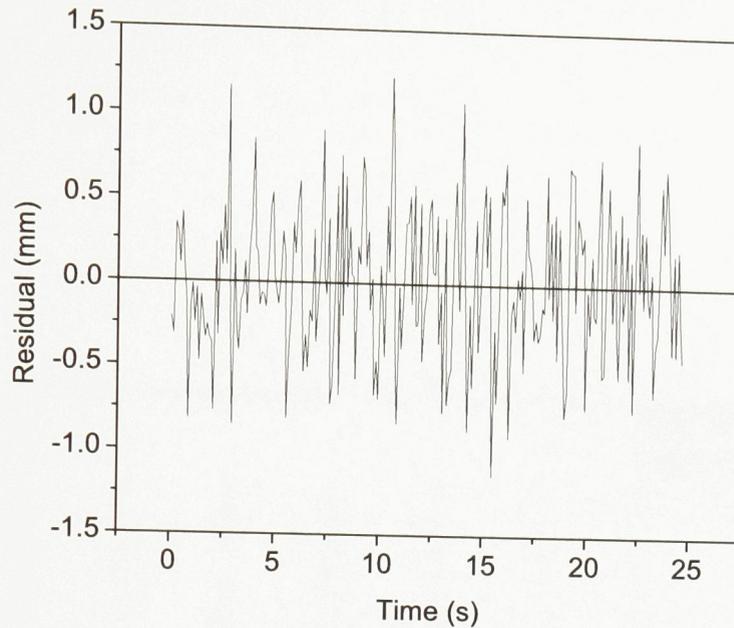


Figure 4.11: Residuals of the fit of scan #4 to a sine wave.

4.2.9 Discussion

The relative accuracy is much higher than the one obtained in the simulation study: -0.03 ± 0.55 mm as opposed to -0.28 ± 0.69 mm. In Chapter 3, I explained that there is a systematic error introduced by the curvature of the path of the markers. This error is also present in this experiment, however its effect is unnoticeable on the relative localization error. In the simulation study, the platform onto which the markers rested rotated around its central axis. Consequently, each marker was pulled towards the centre of rotation. The net effect was to shrink the distance between the markers. In this experiment, the platform itself moves in an orbit. The markers resting on it move in sync with each other. At any given time, all three marker localizations are biased in the same direction by the same distance. Thus, the distance between each pair of localized markers is unchanged. The relative localization error seems to indicate the absence of any other systematic error.

The precision obtained is better than the one obtained in the simulation study:

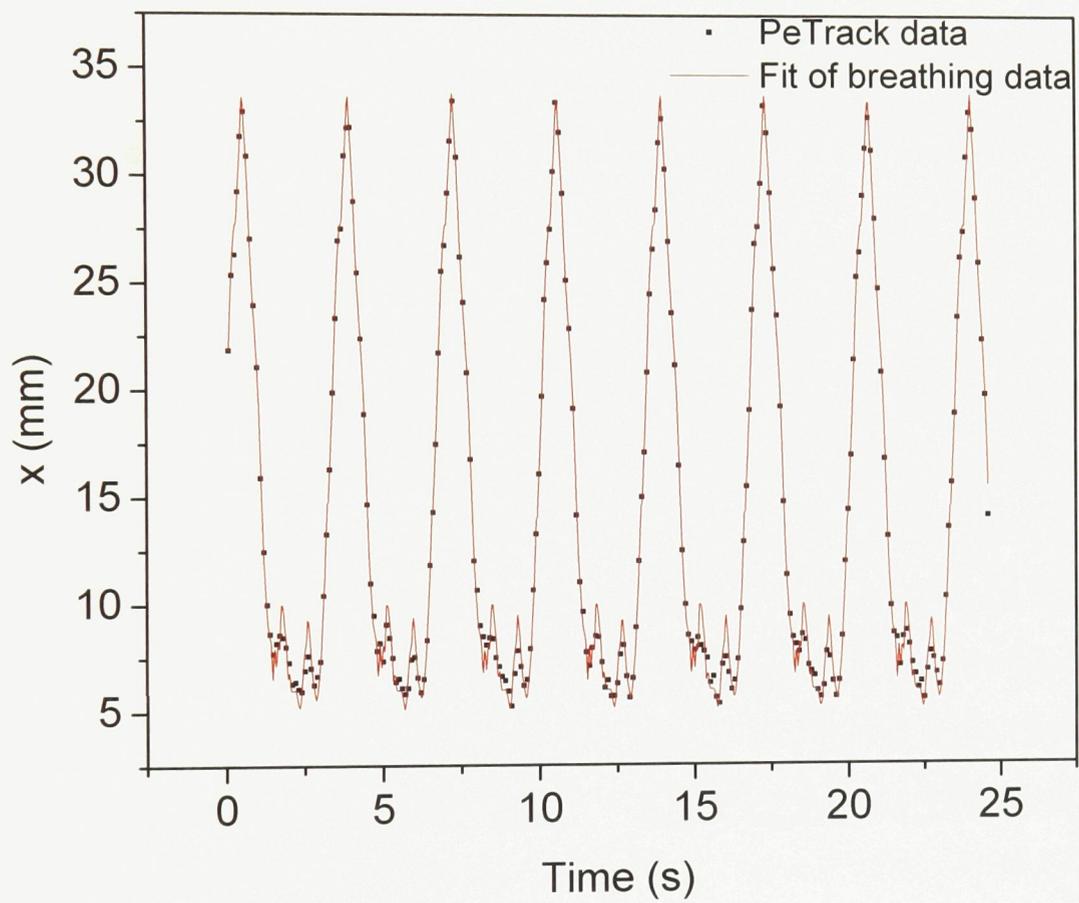


Figure 4.12: PeTrack data of scan #8 fitted to the real breathing curve. Only the x component is shown.

0.39 mm as opposed to 0.49 mm. I believe that this is due to the shorter positron range. In this experiment, I used ^{22}Na sources. The average energy of positrons emitted by ^{22}Na is 182 keV, as opposed to 604 keV for positrons emitted by ^{124}I (used in the simulation study). Thus, the CSDA range of ^{22}Na positrons for a given material is smaller than the CSDA range of ^{124}I positrons. I conclude that the positron range is smaller in this experiment and this effect has improved the precision compared to the simulation study.

The RMSE of the fit of scan #4 (0.46 mm in x , 0.43 mm in y) is in the expected range for the precision of the tracking, i.e. 0.39 mm as determined above and 0.49 mm in the simulation study.

The amplitude of the motion was determined from the fit to be 14.77 ± 0.05 mm in x and 14.67 ± 0.04 mm in y . According to our calibration of the X-Y plotter, the actual amplitude of the motion in x and y was 15.46 mm and 15.13 mm respectively. The discrepancy between the amplitude of the sinusoidal fit and the real amplitude is -0.69 mm and -0.46 mm and in x and y respectively. This indicates a systematic error. Several factors could be contributing to this error. The systematic error due to the curvature of the path of the markers would effectively decrease the amplitude of the motion. Also, the calibration of the X-Y plotter has uncertainties itself. The precision of the calibration is of the order of 0.3 mm. Finally, a slight mechanical anomaly in the motion of the X-Y plotter was noticed. The net effect of this anomaly was to faintly flatten the tip of the elliptical path. This would cause the amplitude of the motion to be smaller than its theoretical value.

Figure 4.12 shows that the localization picked up all the peaks in the breathing motion of scan #8, though they were not as sharp as in the theoretical curve. This may be explained by the sampling rate of the localization. Using a localization interval of 100 ms, it is difficult to track such fast changes in the motion of the markers. Nevertheless, the path of the localized markers strongly suggest that PeTrack can

track tumour motion in real-time.

4.2.10 Conclusion

I presented an improved experimental protocol to test the accuracy and precision of the tracking algorithm with a clinical scanner. The results suggest that the algorithm can achieve submillimetre precision of localization in real-time tumour tracking.

4.3 Assessment with a Prototype PeTrack System

The previous two experiments were done with a full-ring PET system. As mentioned in Chapter 2, the PeTrack system will consist of only two pairs of detector modules 90° apart from each other. This section presents preliminary results obtained from a prototype of the PeTrack system.

4.3.1 Methods

The prototype of the PeTrack system consisted of one pair of modules, each comprising only one crystal block with 169 BGO crystals arranged in a 13×13 matrix. The crystal dimensions were $4 \times 4 \times 20$ mm³. One Hamamatsu H8500 photo-multiplier tube (Hamamatsu Photonics K.K., Japan) was coupled to each crystal block. The data from the detector was read out by a mesytec MPET-8 module (mesytec GmbH & Co. KG, Germany).

The prototype was designed with the two detector blocks 60 cm apart, directly facing each other. Two sealed ²²Na sources (described in Section 4.1.2) were placed between the two modules on a plastic platform, approximately 25.4 mm away from each other. Two pairs of modules 90° apart are required in order to correctly localize markers. Since only one pair of modules were built, the data acquisition was split

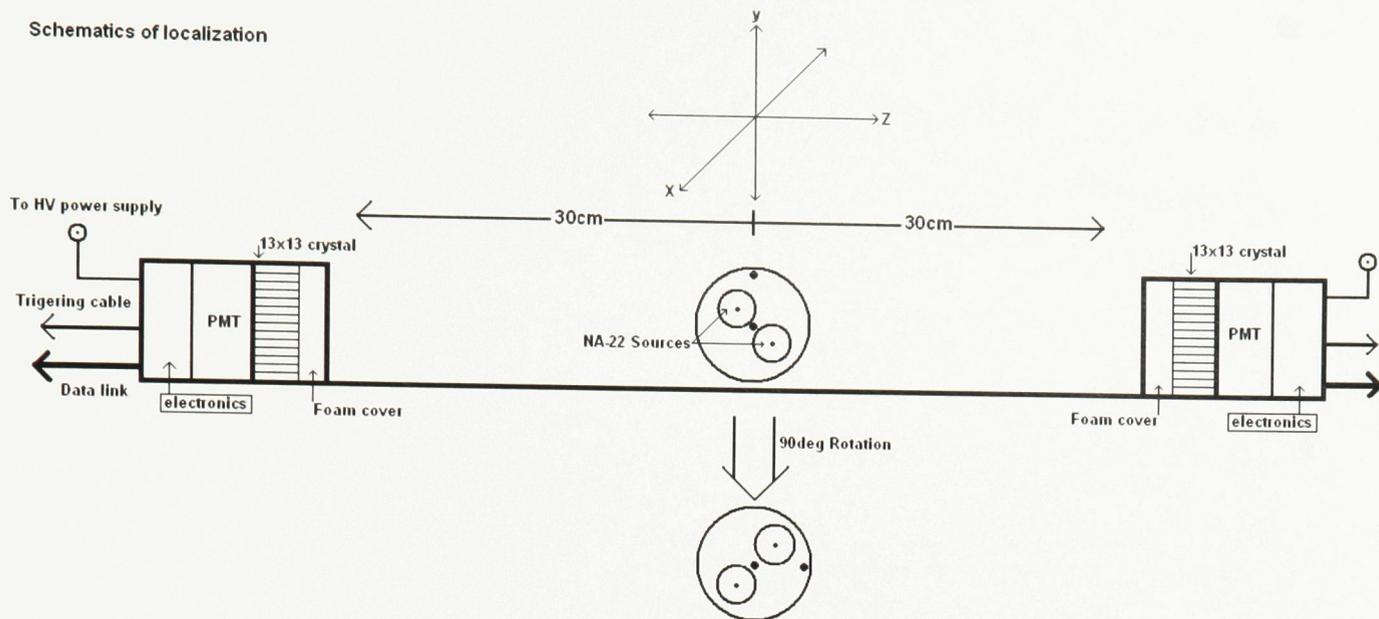


Figure 4.13: Setup of the experiment with the PeTrack prototype.

into two halves of equal duration in time. To simulate a second pair of modules, the markers were rotated by 90° counter-clockwise between the first and second half of the data acquisition. In total, approximately 5500 coincidence events were collected. The order of the coincidence lines was randomly shuffled. The experimental setup is shown in Figure 4.13. I localized the two markers by using 200 total coincidence lines per localization run.

4.3.2 Results

On average, 88 true coincidence lines per marker were used for the localization. The remaining 12% of coincidence lines were rejected by the algorithm as corrupted events (scatter or random). The average relative accuracy over all runs was -0.66 ± 0.26 mm. Since the markers are stationary in this experiment, I can evaluate the precision of the tracking on each direction. The average precision was 0.16 mm on x , 0.20 mm on y and 0.21 mm on z .

4.3.3 Discussion

The results from the prototype PeTrack are very encouraging. The average relative accuracy was negative, which indicates that the distance between the markers was systematically smaller than its expected true value of 25.4 mm. This discrepancy is due to uncertainties in the geometrical setup of the prototype since it has not yet been fully calibrated.

The average precision on each coordinate was between 0.16 and 0.21 mm. The precision on y and z was slightly worse than the precision on x . This was expected since y and z were both normal to the detector surface. However, this effect is small.

The precision was better than in previous experiments. At least two factors can explain this improvement. First, photon noncollinearity degrades precision, but scales with the distance between two opposite detector blocks. The detector blocks of the prototype were only 60 cm away from each other, compared to 86.41 cm in the case of the Philips Allegro system and 100 cm in the case of the simulation study. Secondly, the crystals were smaller in size in the prototype ($4 \times 4 \times 20 \text{ mm}^3$) than in either the clinical PET system ($4 \times 6 \times 20 \text{ mm}^3$) or in the simulation ($4 \times 4 \times 30 \text{ mm}^3$). Both of these factors contributed to the precision observed with the prototype system.

4.4 Conclusion

This chapter presented three experiments conducted to assess the performance of the tracking algorithm. The protocol of the first experiment had limitations. In addition, the conversion from List Mode format to coincidence lines is difficult to infer from simple geometric arguments. A new experimental protocol was developed with a more relevant configuration of the markers and more realistic motion patterns. The precision and accuracy were found to be below 1 mm in all cases. Finally, I evaluated the performance of the tracking using two stationary sources and a prototype of the

PeTrack system. The results showed once again that submillimetre accuracy and precision is achievable.

Chapter 5

Algorithm Improvements

Previous chapters presented results from simulation and experimental studies of the PeTrack tracking algorithm. However, two key components of the algorithm are missing in order to apply the technique to real-time tumour tracking. This chapter presents improvements made to the algorithm to address two specific issues.

5.1 Algorithm Initialization

The tracking algorithm presented in Chapter 2 is based on an expectation-maximization clustering algorithm. The solution to the EM clustering algorithm can depend strongly on the values of the initial parameters. [25] Several methods exist to select initial parameters. [26–28] The tracking algorithm needs an initial estimate of the position of the markers. I present a novel initialization method inspired by a clustering technique proposed by Gundogdu. [29]

5.1.1 Initialization Around the Centre of Mass

I select the initial starting points of the algorithm by randomly initializing around the centre of mass of the first set of collected coincidence lines. The centre of mass is found by using midpoints. A midpoint is defined as the point of intersection of two

coincidence lines or the halfway point on the minimum distance vector that joins two coincidence lines. Since the RMS distance between coincidence lines should be minimized at a marker location, the density of midpoints is expected to be considerably higher in the vicinity of a marker. Hence, the initialization proceeds as follows:

1. Randomly select three coincidence lines from the set of events recorded in the first localization interval. Typically, 300 coincidence lines are collected per localization for three markers.
2. Calculate the minimum distance between each pair of coincidence lines. If all three coincidence lines are not within 5 mm of each other, it is highly probable that they did not originate from the same marker; in this case, repeat step 1.
3. Calculate the three midpoints between each pair of coincidence lines.
4. Repeat steps 1 through 3 500 times.
5. Take the centre of mass of the 1500 midpoints obtained from steps 1 to 4.
6. Initialize the EM algorithm randomly around the centre of mass for each of K different markers. The random range can be set to the average distance between the markers.

A diagram of the logical flow of this method is shown in Figure 5.1.

I should note that I also explored a different method of initializing the algorithm. The spatial distribution of the midpoints shows a higher density around markers. If these regions of high densities can be identified, the centre of each of these regions should be close to the position of each of the K different markers. Identifying these regions by computational means is difficult. I have explored the use of a clustering algorithm that groups the midpoints into a pre-specified number of clusters (K clusters in this case, corresponding to K different markers). However, this method performed

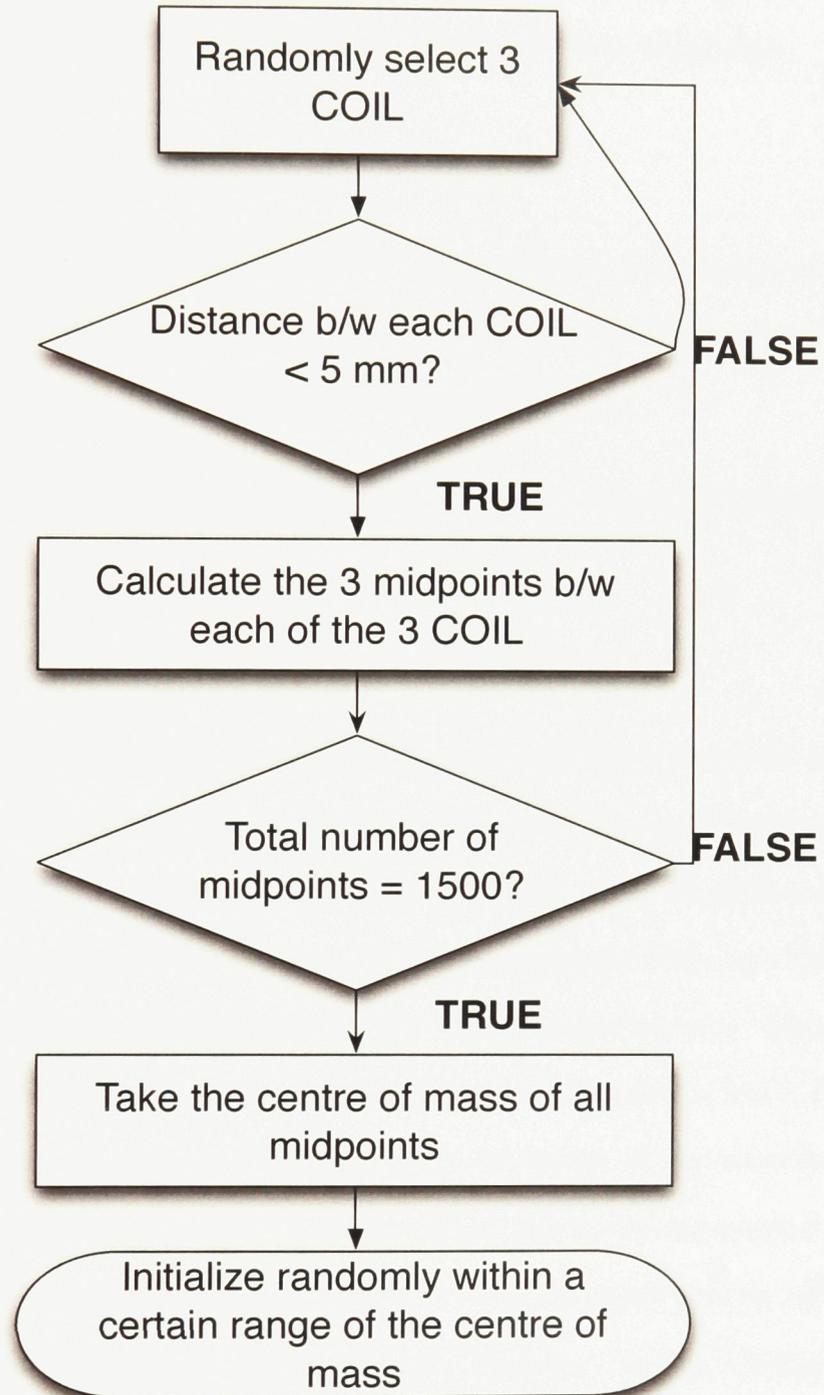
COIL: COIncidence Lines

Figure 5.1: Logic flow of the initialization method.

poorly in terms of computational time. Also, the spatial distribution of midpoints sometimes contained more regions of high density than there were markers, leading to a poor clustering of the midpoints. This in turn led to a poor estimate of the initial position of the markers and to frequent failures in the localization algorithm.

5.1.2 Identifying Failed Markers

It is likely that the random initialization around the centre of mass will not be accurate enough for the tracking to correctly localize all markers on the first run. Failed localizations can be identified by any of the following conditions:

- Condition 1: All three markers are assigned to the same cluster of coincidence lines.
- Condition 2: Two markers are assigned to the same cluster.
- Condition 3: An unusually high number of coincidence lines are used in the localization of a marker or the standard deviation of the distance between the coincidence lines and their respective marker is unusually high. For three markers of equal activity, each marker is expected to use approximately one third of the total coincidence lines detected during the localization interval. Hence, if a marker is using more than half the total number of coincidence lines, I consider the localization as a failure. The standard deviation of the distribution of the minimum distance of coincidence lines to the corresponding marker for a successfully localized marker was 3.2 ± 0.3 mm, when averaged among all data sets from the second experiment at the Ottawa Hospital Cancer Centre (see Chapter 4). The standard deviation was considered abnormally high when it was higher than 10 mm.

Whenever one or more conditions are met, the algorithm is re-initialized and the localization is repeated until all markers are successfully localized. To this end, I have implemented three fail-safe methods to address each of the above conditions:

- Method 1: If all three markers were assigned to the same cluster (condition 1), this signifies that one marker was successfully localized, while the other two markers were wrongly assigned to the same location as the successful marker. In this case, I randomly re-initialize the two failed markers at a set distance around the localized marker. The re-initialization distance should be of the order of the distance between the true locations of the markers.
- Method 2: If only two markers were assigned to the same cluster (condition 2), this means that two markers were successfully localized while the third marker was incorrectly assigned to the same location as one of the two successful markers. In this case, I randomly re-initialize the failed marker at a set distance around the half-way point of the distance vector that joins the two successful markers and on the perpendicular plane to that distance vector. The distance should be of the order of the distance between the true locations of the markers.
- Method 3: If condition 3 is met, I simply re-initialize this failed marker randomly around its current location and move on to the next set of collected coincidence lines.

These methods are based on the scenario where only the approximate distance between markers is known. It would be possible to include prior knowledge of the position of the markers relative to each other to make the initialization phase shorter.

5.1.3 Testing for Bias

The initialization method might produce a systematic effect on the localization. To verify this, I used one of the simulation data sets from Chapter 3 to test whether the localization is biased towards the location selected to initialize the tracking. For each localization, I constructed the normalized vector pointing from the localized marker to its true location. I projected this vector on the normalized vector that points from the initial location used in the localization to the true location of the marker. I repeated this procedure for all markers over all localizations.

5.1.4 Results

I tested the initialization method on the data sets acquired during the second experiment at the Ottawa Hospital Cancer Centre (see Chapter 4). I first tested the centre of mass initialization method, without checking and recovering for failed localizations. In this situation, all three markers were localized successfully in two out of the eight data sets. Two markers were successfully localized in five out of the eight data sets. Finally, all three markers were assigned to the same cluster in the remaining data set.

When adding re-initialization for the failed localizations, all three markers were eventually successfully localized in all data sets. Since the re-initialization is random, a few attempts are usually needed before the three markers are correctly localized. I call this the initialization phase of the algorithm (see Figure 5.3). Among all eight data sets, the initialization phase took on average 750 ms from the start of the data acquisition and never lasted more than 1.6 seconds before correctly finding all three markers. These times are data acquisition times and not computational times (which are on the order of a few milliseconds on a 2.2 GHz Intel®Core 2 Duo processor).

Figure 5.2 shows the results of the bias test described in the previous section. The average value is -0.002, which indicates that the localization is not biased towards the

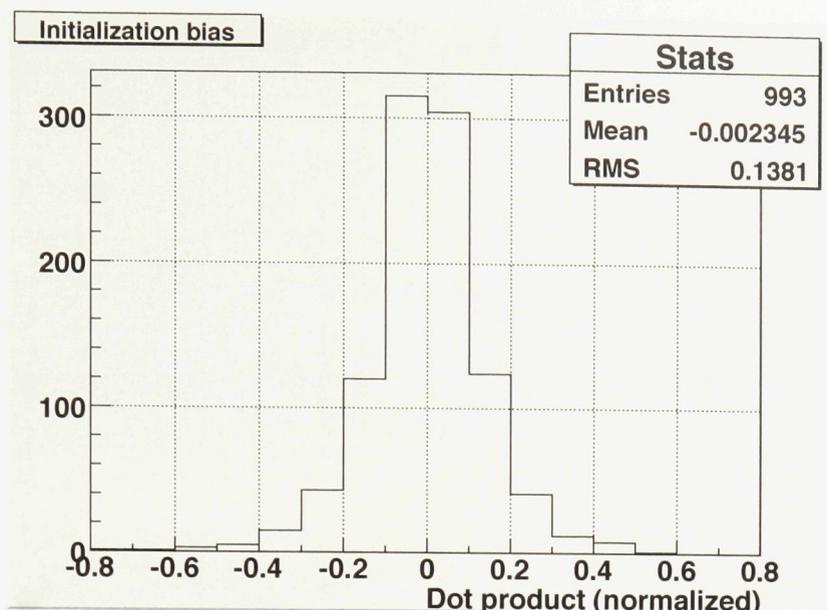


Figure 5.2: Projection of the normalized vector pointing from the localized marker to its true location onto the normalized vector pointing from the initial location to the true location of the marker.

initialization point.

5.1.5 Discussion

Results from the proposed method indicate that a combination of centre of mass initialization and fail-safe re-initialization methods is capable of successfully localizing all markers within a few seconds of the start of the data acquisition. Only an estimate of the distance between the markers is necessary for proper localization. Thus, patient setup errors, changes in the tumour shape, and marker migration should have little effect on the tracking capability of PeTrack.

Figure 5.3 shows the x coordinate as a function of time for one marker of scan #5. The failed localizations occur during the first 1.2 s of the data acquisition. The re-initialization is done very quickly (typically on the order of a few milliseconds). Once a marker's location has been re-initialized, the algorithm waits until the end of the current localization time interval (100 ms, in this case) before localizing all markers again. This waiting period until the end of the current localization time interval and

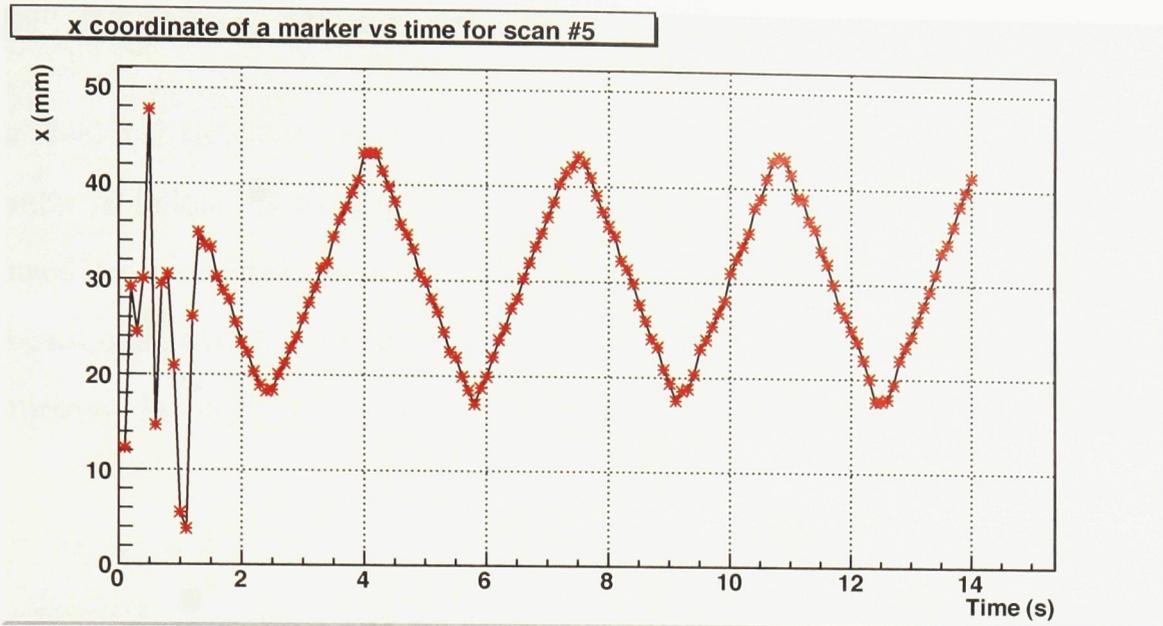


Figure 5.3: x coordinate of a marker as a function of time for scan#5. From $t = 0$ to $t = 1.2$ s, the marker location fluctuates randomly. This corresponds to the initialization phase of the algorithm. The tracking recovers at approximately $t = 1.3$ s.

the randomness of the initialization are the main limiting factors in the time of the initialization phase.

Once the markers are correctly localized at the end of the initialization phase, the algorithm effectively locks-on to the markers. This is seen after $t = 1.3$ s on Figure 5.3. Over all data sets studied, including both simulation and experimental data, the algorithm never lost track of any marker past the initialization phase.

In a clinical setting, this means that the initialization will eliminate patient setup error. If the patient is already in the field of view when the tracking algorithm is launched, it will take approximately 750 ms of data acquisition on average to lock-on to the markers. Even if a marker gets lost during the treatment, it will be identified by the conditions listed in Section 5.1.2. The delivery of treatment will be interrupted and an alert will be sent to the operator of the system. The fail-safe methods will allow the tracking to recover and the treatment to resume quickly.

5.2 Motion Prediction

In an ideal real-time tumour tracking system, the therapy beam would be repositioned instantly to follow the tumour at every localization. This would effectively eliminate the need for a tumour motion margin from the treatment planning. Unfortunately, the beam can never be perfectly synchronized to the tracking since every system has an inherent latency. These delays are due to the computational time required to localize the tumour and to the mechanical reaction time of the hardware. Litzenberg *et al.* report a delay of 55 ms in their dynamic multi-leaf collimator system. [30] Seppenwoolde *et al.* have a 90 ms delay in their beam-gated radiotherapy system. [24] Mechanical systems that need to reposition the beam will have even longer delays, up to 200 ms or more. [4]

These delays justify the need for motion prediction. Given that a system with latency ΔT has localized the tumour at time t , the tumour location at time $t + \Delta T$ needs to be predicted. The system can then be repositioned according to the predicted tumour location. The system latency should be kept to a minimum since prediction error has been shown to increase with latency. [31]

This section presents two simple prediction methods. I tested these methods on the data sets from the second experiment conducted at the Ottawa Hospital Cancer Centre (see Chapter 4).

5.2.1 Linear Extrapolation

The first method I implemented is a very simple linear extrapolation scheme described by Sharp *et al.*. [31] I use the previous two locations of a marker to calculate its velocity and predict its next location. I assume the marker maintains constant velocity. This

prediction is given by the following equation:

$$\tilde{x}_{t_n} = 2x_{t_{n-1}} - x_{t_{n-2}} \quad (5.1)$$

where t_n is the time of the n^{th} localization in milliseconds, \tilde{x}_{t_n} is the predicted location of the marker at the next localization, and $x_{t_{n-1}}$ and $x_{t_{n-2}}$ are the two previous locations of the marker, as determined by the algorithm.

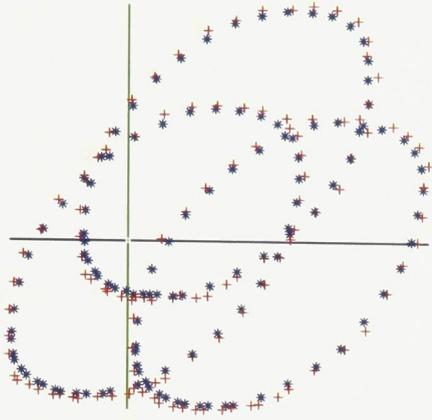
5.2.2 Circular Approximation

I implemented a second prediction method. Since tumour motion presents hysteresis in half of patients, I assume that for a relatively short time interval, the markers move on the arc of a circle. [24] I calculate the average angular velocity ω of the marker during this interval by using the previous three locations of the marker. I predict the marker's next location by extending the arc of the circle by $\omega \times t$, where t is set to the localization interval used while tracking the markers, i.e. $t = 100$ ms. The circle on which to extend the arc is the circle that passes through the last three locations of the marker.

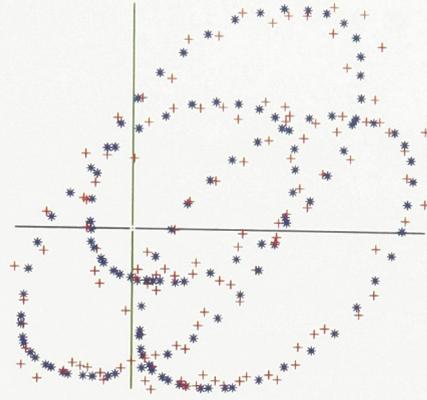
5.2.3 Results

Since the true location of the markers in the frame of the scanner is not known, I compared the predicted location to the localized track of the markers.

For linear extrapolation, the average error between the predicted location and the localized markers was 1.1 ± 0.7 mm. The prediction error was 2.3 mm at a 95% confidence level. This means that 95% of the time, the prediction error will be less than 2.3 mm. In the case of the circular approximation, the average prediction error was 3.6 ± 2.3 mm. The prediction error was 7.7 mm at a 95% confidence level. Figure 5.4 shows the predicted path of the markers for scan #1 using linear extrapolation and

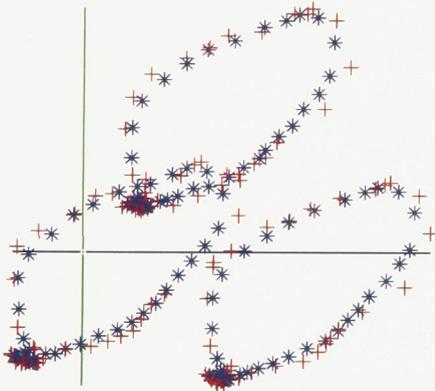


(a) Linear prediction

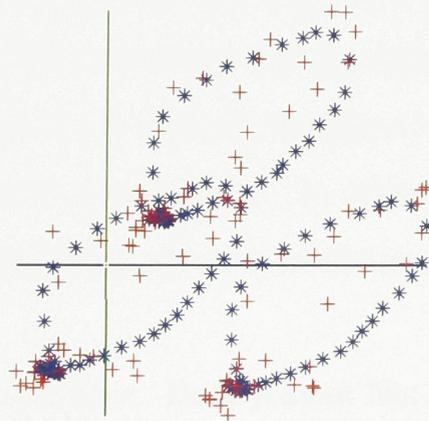


(b) Circular approximation

Figure 5.4: Comparison of predicted marker location (red crosses) and tracked path of the markers (blue stars) for scan #1 (projected onto the $x - y$ plane).



(a) Linear prediction



(b) Circular approximation

Figure 5.5: Comparison of predicted marker location (red crosses) and tracked path of the markers (blue stars) for scan #8 (projected onto the $x - y$ plane).

circular approximation compared to the localized track of the markers. Figure 5.5 shows the results for scan #8.

5.2.4 Discussion

Linear extrapolation outperformed circular approximation in all scans. The range of prediction error observed with circular approximation is very large and makes this technique unsuitable for clinical use. Linear extrapolation shows more promise, as the average error was 1.1 ± 0.7 mm. This is very promising. Margins to account for organ motion and setup errors for lung cancer tumours can vary between 5 mm when using tumour tracking and 10–15 mm when no tumour tracking is used. [23,32] With an average prediction error of 1.1 ± 0.7 mm and a maximum error of 2.3 mm at a 95% confidence level for linear extrapolation, treatment margins could be reduced to only a few millimetres by using PeTrack.

Both methods used only a few previous locations of a marker to predict its next location (2 previous locations in the case of linear extrapolation, and 3 in the case of circular approximation). The random error associated with the position of a marker makes it very difficult to estimate its velocity. The case of circular approximation is even worse because two more pieces of information are needed: the plane of motion and the centre of the circle, each with associated uncertainties. In addition, both methods assume constant velocity of the markers. This makes the prediction methods only applicable to systems with low latency.

The problem of motion prediction is rather complex. Each patient has a unique breathing pattern. Though there is a periodic component to the breathing motion, each cycle varies randomly in period and amplitude. [33] Furthermore, the prediction is based on the tumour localization which has uncertainties associated with it. Mathematical models have been used to predict breathing motion, but the intrinsic random fluctuations of breathing can make these models inaccurate. [24]

Another approach to tumour motion prediction is based on linear prediction and seems to be promising. [8,31,34] Linear prediction assumes that the next position of the tumour will be a linear combination of its past locations:

$$\tilde{x}_{t_n} = a_0 + a_1x_{t_{n-1}} + \dots + a_nx_{t_0} \quad (5.2)$$

Using several past histories smoothes out the random fluctuations of the motion. Furthermore, adaptive filters can be used with this technique. An adaptive filter adjusts the parameters of the linear combination in real time after each new prediction in order to minimize the prediction error. [34] This could potentially improve the prediction accuracy compared to linear extrapolation and circular approximation.

5.3 Conclusion

This chapter presented work done to improve the PeTrack algorithm. The centre of mass method was developed to initialize the algorithm. Self-correction methods were implemented to recover from incorrectly localized markers. Two methods were tested to predict tumour motion. Linear extrapolation performed better than circular approximation in all data sets. I plan to continue work on the prediction part of the algorithm. Future work will explore the possibility of using linear prediction and adaptive filters to predict tumour motion.

Chapter 6

General Discussion

In this chapter, I will discuss the principal sources of errors that affect the performance of PeTrack and I will address some of the limitations of my work.

6.1 Sources of Errors

The accuracy and precision of PeTrack is limited by several factors. If the intrinsic spatial resolution of the PeTrack detector is σ_d , then the accuracy of the localization algorithm is approximately given by σ_d/\sqrt{N} , where N is the total number of coincidence lines collected during a localization time interval. [13] The performance of the tracking could be improved by increasing the number of coincidence lines used per localization. However, this entails increasing the activity of the markers, which in turn would deliver a higher dose to the patient. The correct balance between desired accuracy of the tracking and dose delivered to the patient should be determined carefully during treatment planning.

The curvature of the path of the markers introduces a systematic error in the localization. I suspect that the markers are being localized near the centre of mass of their track during a localization time interval. In the simulation study, the markers were moving along a circular path. The minimum distance between a circular arc

and its centre of mass is given by:

$$\delta_{com} = \rho \left(1 - \frac{\cos \alpha/2}{\alpha/2} \right), \quad (6.1)$$

where ρ is the radius of the circle and α is the angle subtended by the arc at the centre of the circle. In the simulation study, the markers were moving at an angular velocity of $120^\circ/\text{s}$, which translates to $\alpha = 12^\circ$ for a localization time interval of 100 ms. Using Equation 6.1 and the radius of rotation used in the simulations ($\rho = 5$ mm to $\rho = 31$ mm), I find that the distance between the middle of the track and its centre of mass varies between $\delta_{com} = 0.01$ mm and $\delta_{com} = 0.06$ mm. This is much lower than the values I have determined for the systematic error in Section 3.9, which ranged from 0.10 to 0.37 mm in absolute value. This indicates that the markers are not being localized exactly at the centre of mass of the track. This effect should be studied further to determine its exact causes and to find ways to compensate for it. One method of remedying to this effect could be to shorten the track length of the markers. This would be achieved by increasing the sampling rate of the tracking algorithm. However, there is a limit to the minimum computational time required for the algorithm to converge. This limit should be determined experimentally. In addition, enough coincidence lines need to be collected during each localization time interval to achieve the required accuracy.

The intrinsic spatial resolution of the PeTrack system also factors into the accuracy of the localization. The spatial resolution of the system is limited by three main factors: the range of the positrons, the width of the crystals used in the detector, and the non-collinearity of the annihilation photons. I discussed the positron range in Section 3.9. For a ^{124}I source confined in a gold, tungsten or titanium metal capsule, the positron CSDA range at the average energy of 604 keV is between 0.22 mm (in tungsten) and 0.70 mm (in titanium). This indicates that with the appropriate choice

of capsule material and isotope, the positron range effect should be small.

The width of the crystals used in the detector also affects resolution. The FWHM of this effect is $w/2$, where w is the width of the crystal. [35] Reducing the crystal size improves this factor. In the case of PeTrack, the crystal width is $w = 4$ mm (FWHM= 2 mm). Using a crystal twice as small could bring this factor down to a FWHM of 1 mm. However, this increases the cost of the system.

The non-collinearity of annihilation photons occurs when the total momentum of the positron and electron is non-zero when the pair annihilates. This causes the gamma rays to deviate from a 180° emission angle. The distribution of the deviation angle has a FWHM of $0.0044R$, where R is the radius of the detector. [12] For a distance of 60 cm between opposite detectors ($R = 30$ cm), as used with the prototype PeTrack detector, this translates into a FWHM of 1.32 mm. However, there is a limit to how close the detector modules can be to each other since they need to be mounted on a Linac gantry. Current data suggest that a distance of 1 m ($R = 50$ cm) between opposite detector modules might be more realistic for use of PeTrack in a clinical setting. This translates to a FWHM of 2.2 mm. This is higher than all other factors described above and indicates that non-collinearity has the highest contribution to the performance of the localization of markers. Since there is a limit to how close the detectors can be to each other, the other sources of errors need to be minimized.

6.2 Limitations

I identify a few limitations in my research. In all my experiments, the motion of the markers was constrained to a two-dimensional plane. This was done to reduce the complexity of the simulations and experimental protocols. In real life, the markers might move in all three dimensions. Future experiments should include three-dimensional motion of the markers. In addition, the amplitude of motion of the

markers was slightly higher than the average reported amplitude of motion of lung tumours. Shirato *et al.* reported an average amplitude of motion of 8.2 ± 6.5 mm in the right-left direction, 10.7 ± 8.6 mm in the superior-inferior direction, and 8.8 ± 7.0 mm in the anterior-posterior direction. [21] The smallest amplitude I used was 20.2 mm, which is on the high end of the range of average amplitude of motion for lung tumours. The reported size of lung tumours (from 9 to 38 mm) would also motivate reducing the distance between the markers. [23] This would also present an opportunity to determine what is the minimum distance needed between two markers for PeTrack to resolve them.

Chapter 7

Conclusion

7.1 Conclusions

I presented PeTrack, a technique that can track tumours in real-time using markers that emit positrons. I have evaluated the accuracy and precision of the tracking component of PeTrack using simulation and experimental studies. In both cases, the technique was shown to achieve submillimetre accuracy and precision. I conclude this performance is sufficient for real-time tumour tracking.

I developed a method to initialize the tracking algorithm around the centre of mass of midpoints. Furthermore, the tracking algorithm can accurately identify failed localizations and correct the situation. During my testing, the initialization phase of the tracking never lasted more than 1.6 seconds from the start of data acquisition and the markers never got lost beyond the initialization phase. These features will reduce the complexity of use of PeTrack in a clinical setting since initialization and recovery of lost markers will be done with little to no human input.

I experimented with two simple methods to predict tumour motion. The average prediction error of linear extrapolation and circular approximation was 1.1 ± 0.7 mm and 3.6 ± 2.3 mm, respectively. Linear extrapolation might be able to reduce treatment margins to only a few millimetres.

7.2 Future Work

Some elements need to be addressed in the development of PeTrack. In future experiments, the path of the markers should not be constrained to only one plane of motion. This might be achievable with a respiratory motion phantom. The markers should also be brought closer together to better reflect tumour size. The minimum distance required between two markers in order to be resolved by PeTrack should also be determined experimentally. The systematic error causing the markers to be localized inside their curved track should be studied in more depth. A method of compensating for this error should be developed. Finally, I recommend more work into tumour motion prediction. Linear prediction and adaptive filters might be able to predict tumour motion more accurately than linear extrapolation.

List of References

- [1] T. Bortfeld. “IMRT: a review and preview.” *Phys. Med. Biol.* **56**, R363–R379 (2006).
- [2] ICRU Report 24. “Determination of absorbed dose in a patient irradiated by beams of x or gamma rays in radiotherapy procedures.” Technical report, International Commission on Radiation Units and Measurements, Washington DC (1976).
- [3] I. Jacobs, J. Vanregemorter, and P. Scalliet. “Influence of respiration on calculation and delivery of the prescribed dose in external radiotherapy.” *Radiother. Oncol.* **39**(2), 123–128 (1996).
- [4] P. Keall, G. Mageras, J. Balter, R. Emery, K. Forster, S. Jiang, J. Kapatoes, D. Low, M. Murphy, B. Murray, C. R. Ramsey, M. van Herk, S. S. Vedam, J. Wong, and E. Yorke. “The management of respiratory motion in radiation oncology report of AAPM Task Group 76.” *Med. Phys.* **33**(10), 3874–3900 (2006).
- [5] J. Hoisak, K. Sixel, R. Tirona, and P. Cheung. “Correlation of lung tumor motion with external surrogate indicators of respiration.” *Int. J. Radiat. Oncol. Biol. Phys.* **60**(4), 1298–1306 (2004).
- [6] 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 (2000).
- [7] H. Shirato, M. Oita, K. Fujita, Y. Watanabe, and K. Miyasaka. “Feasibility of synchronization of real-time tumor-tracking radiotherapy and intensity-modulated radiotherapy from viewpoint of excessive dose from fluoroscopy.” *Int. J. Radiat. Oncol. Biol. Phys.* **60**(1), 335–341 (2004).

- [8] X. Tang, G. Sharp, and S. Jiang. “Fluoroscopic tracking of multiple implanted fiducial markers using multiple object tracking.” *Phys. Med. Biol.* **52**, 4081–4098 (2007).
- [9] J. Balter, J. N. Wright, L. Newell, and B. Friemel. “Accuracy of a wireless localization system for radiotherapy.” *Int. J. Radiat. Oncol. Biol. Phys.* **61**(3), 933–937 (2005).
- [10] T. Willoughby, P. Kupelian, J. Pouliot, and K. Shinohara. “Target localization and real-time tracking using the Calypso 4D localization system in patients with localized prostate cancer.” *Int. J. Radiat. Oncol. Biol. Phys.* **65**(2), 528–534 (2006).
- [11] J. Kindblom, A.-M. Ekelund-Olvenmark, and H. Syren. “High precision transponder localization using a novel electromagnetic positioning system in patients with localized prostate cancer.” *Radiother. Oncol.* **90**(3), 307–311 (2009).
- [12] C. Levin and E. Hoffman. “Calculation of positron range and its effect on the fundamental limit of positron emission tomography system spatial resolution.” *Phys. Med. Biol.* **44**(3), 781–99 (1999).
- [13] T. Xu, J. Wong, P. Shikhaliev, J. Ducote, and M. Al-Ghazi. “Real-time tumor tracking using implanted positron emission markers: Concept and simulation study.” *Med. Phys.* **33**(7), 2598–2609 (2006).
- [14] P. Zanzonico. “Positron emission tomography: a review of basic principles, scanner design and performance, and current systems.” *Semin. Nucl. Med.* **34**(2), 87–111 (2004).
- [15] S. Jan, G. Santin, D. Strul, S. Staelens, K. Assié, D. Autret, S. Avner, R. Barbier, M. Bardiès, and P. Bloomfield. “GATE: a simulation toolkit for PET and SPECT.” *Phys. Med. Biol.* **49**, 4543–4561 (2004).
- [16] S. Agostinelli, J. Allison, and K. Amako. “GEANT4—a simulation toolkit.” *Nucl. Instrum. Methods Phys. Res. A* **506**(3), 250–303 (2003).
- [17] F. Lamare, A. Turzo, Y. Bizais, and C. C. L. Rest. “Validation of a Monte Carlo simulation of the Philips Allegro/GEMINI PET systems using GATE.” *Phys. Med. Biol.* **51**, 943–962 (2006).
- [18] S. Jan, C. Comtat, D. Strul, G. Santin, and R. Trébossen. “Monte Carlo simulation for the ECAT EXACT HR+ system using GATE.” *IEEE Trans. Nucl. Sci.* **52**(3), 627–633 (2005).

- [19] S. Nehmeh, Y. Erdi, T. Pan, and E. Yorke. “Quantitation of respiratory motion during 4D-PET/CT acquisition.” *Med. Phys.* **31**(6), 1333–1338 (2004).
- [20] Q.-S. Chen, M. Weinhaus, F. C. Deibel, and J. Ciezki. “Fluoroscopic study of tumor motion due to breathing: facilitating precise radiation therapy for lung cancer patients.” *Med. Phys.* **28**(9), 1850–1856 (2001).
- [21] H. Shirato, K. Suzuki, G. Sharp, K. Fujita, and R. Onimaru. “Speed and amplitude of lung tumor motion precisely detected in four-dimensional setup and in real-time tumor-tracking radiotherapy.” *Int. J. Radiat. Oncol. Biol. Phys.* **64**(4), 1229–1236 (2006).
- [22] S. Surti and J. Karp. “Imaging characteristics of a 3-dimensional GSO whole-body PET camera.” *J. Nucl. Med.* **45**(6), 1040–1049 (2004).
- [23] T. Harada, H. Shirato, S. Ogura, S. Oizumi, and K. Yamazaki. “Real-time tumor-tracking radiation therapy for lung carcinoma by the aid of insertion of a gold marker using bronchofiberscopy.” *Cancer* **95**(8), 1720–1727 (2002).
- [24] Y. Seppenwoolde, H. Shirato, K. Kitamura, S. Shimizu, M. van Herk, J. 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).
- [25] G. McLachlan. “On the choice of starting values for the EM algorithm in fitting mixture models.” *The Statistician* **37**, 417–425 (1988).
- [26] C. Biernacki. “Initializing EM using the properties of its trajectories in Gaussian mixtures.” *Stat. Comput.* **14**(3), 267–279 (2004).
- [27] C. Biernacki, G. Celeux, and G. Govaert. “Choosing starting values for the EM algorithm for getting the highest likelihood in multivariate Gaussian mixture models.” *Comput. Stat. Data Anal.* **41**(3-4), 561–575 (2003).
- [28] D. Coleman and D. Woodruff. “Cluster analysis for large datasets: An effective algorithm for maximizing the mixture likelihood.” *J. Comput. Graph. Stat.* **9**(4), 672–688 (2000).
- [29] O. Gundogdu. “Positron emission tomography particle tracking using cluster analysis.” *Nucl. Instrum. Methods Phys. Res. A* **534**(3), 562–576 (2004).

- [30] D. Litzenberg, J. Moran, and B. Fraass. “Incorporation of realistic delivery limitations into dynamic MLC treatment delivery.” *Med. Phys.* **29**(5), 810–820 (2002).
- [31] G. Sharp, S. Jiang, S. Shimizu, and H. Shirato. “Prediction of respiratory tumour motion for real-time image-guided radiotherapy.” *Phys. Med. Biol.* **49**(3), 425–440 (2004).
- [32] L. Ekberg, O. Holmberg, and L. Wittgren. “What margins should be added to the clinical target volume in radiotherapy treatment planning for lung cancer?” *Radiother. Oncol.* **48**(1), 71–77 (1998).
- [33] G. Benchetrit. “Breathing pattern in humans: diversity and individuality.” *Resp. Physiol.* **122**(2-3), 123–129 (2000).
- [34] S. S. Vedam, P. Keall, A. Docef, D. Todor, and V. Kini. “Predicting respiratory motion for four-dimensional radiotherapy.” *Med. Phys.* **31**(8), 2274–2283 (2004).
- [35] T. F. Budinger. “PET instrumentation: What are the limits?” *Semin. Nucl. Med.* **28**(3), 247–267 (1998).