A CLINICAL C-ARM BASE-TRACKING SYSTEM USING COMPUTER VISION FOR INTRAOPERATIVE GUIDANCE

by

Luke Haliburton

B.Eng., Dalhousie University, 2015

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF APPLIED SCIENCE

in

The Faculty of Graduate and Postdoctoral Studies (Biomedical Engineering)

THE UNIVERSITY OF BRITISH COLUMBIA (Vancouver)

August 2017

© Luke Haliburton, 2017

Abstract

Mobile C-arm X-ray machines are commonly used for imaging during orthopaedic surgeries to visualize internal anatomy during procedures. However, there is evidence indicating that excess operating time and radiation exposure result from the use of scouting images to aid C-arm positioning during surgery. Additionally, C-arms are currently used primarily as a qualitative tool. Several techniques have been proposed to improve positioning, reduce radiation exposure, and increase quantitative utility, but they require accurate C-arm position tracking. There have been attempts by other research groups to develop C-arm tracking systems, but there are currently no solutions suitable for use in an operating room. The objective of this thesis is therefore to present the development and verification of a real-time C-arm base-tracking system called OPTIX (On-board Position Tracking for Intraoperative X-rays).

The proposed tracking system uses a single floor-facing camera mounted to the base of a C-arm. A computer vision algorithm was developed that tracks motion relative to the operating room floor. This system is capable of relative motion tracking as well as absolute position recovery for previous positions.

The accuracy of the system was evaluated on a real C-arm in a simulated operating room. The experimental results demonstrated that the relative tracking algorithm can measure C-arm translation with errors of less than 0.75% of the total distance travelled, and orientation with errors better than 5% of the cumulative rotation. With the incorporated loop closure step, OPTIX can be used to achieve C-arm repositioning with translation errors of less than 1.10 \pm 0.07 mm and rotation errors of less than 0.17 \pm 0.02°. These results are well within the desired system requirements of 5 mm and 3.1°.

The system has shown promising results for use as a C-arm base-tracking system. The system has clinically acceptable accuracies and should lead to a reduced need for scouting images when reobtaining a previous position. The base-tracking system can be integrated with a C-arm joint tracking system, or implemented on its own for steering guidance. When implemented in an operating room, OPTIX has the potential to lead to a reduction in operating time and harmful radiation exposure to surgical staff.

Lay Summary

Mobile X-ray machines called C-arms are commonly used to provide images of a patient's internal anatomy during surgery. Several image-based tools can be developed if the position of a C-arm is tracked throughout a procedure. These tools could reduce operation times and radiation exposure, and could assist surgeons in performing safe and effective procedures. However, there are currently no C-arm tracking solutions that can be practically implemented in an operating room. We have developed position tracking system called OPTIX that uses a floor-facing camera mounted to a C-arm to monitor movements of the wheeled base.

OPTIX can provide steering guidance on its own, or it can be integrated with a joint tracking system to provide full positional information. We performed tests in a simulated operating room and achieved clinically useful levels of accuracy. We have created a promising tool with the potential to reduce operating time and radiation exposure.

Preface

This thesis is an original piece of work by the author, Luke Haliburton. Algorithm development and experimental evaluations were carried out by the author with guidance from my co-supervisors, Dr. Antony Hodgson and Dr. Carolyn Anglin, who also provided revisions for thesis writing. Dr. Pierre Guy, my clinical co-supervisor, provided guidance regarding clinical applications and experimental scenarios.

The original concept of using a downward-facing camera for C-arm base-tracking was initially proposed by Hooman Esfandiari in his Masters thesis at the University of Calgary in 2014 under the supervision of Dr. Carolyn Anglin and Dr. Derek Lichti.

All figures in this thesis that were obtained from copyrighted sources have been reproduced with written permission from the copyright owners.

The author designed all the parts used to mount the camera and processing unit to the C-arm, as well as the mounting system for the light ring. The mounting parts were fabricated by the author using the 3D printer available in the Centre for Hip Health and Mobility.

Table of Contents

Abstract	ii
Lay Summary	iii
Preface	iv
Table of Contents	v
List of Tables	vii
List of Figures	viii
List of Abbreviations	xi
Acknowledgements	xiii
Chapter 1: Motivation and Background	1
1.1. C-arm Use in Operating Rooms	
1.2. C-arm Fluoroscopy	2
1.3. Ionizing Radiation in the Operating Room	4
1.4. Enhancing C-arm Utility	7
1.5. Existing Solutions and Gap	8
1.5.1. External Optical Tracking	8
1.5.2. Electromagnetic Tracking	9
1.5.3. Radiographic Fiducials	9
1.5.4. Mounted In-line Camera	
1.5.5. TC-arm	11
1.5.6. Gap in the Current Solutions	11
1.6. Observed Need: C-arm Tracking	
1.7. Base-Tracking Methods	13
1.7.1. Encoders	
1.7.2. Optical Flow Sensors	13
1.7.3. Camera-based Tracking	14
1.7.4. Proposed Solution	15
1.8. Computer Vision Background	15
1.9. Previous Work	
1.10. Thesis Objectives	
1.11. Thesis Outline	19
Chapter 2: Design and Development of a C-arm Base-Tracking System	20
2.1. Tracking System Requirements	
2.2. OPTIX Design Stages	
2.2.1. Hardware Selection	
2.2.2. Software Packages	

	2.2.3.	Camera Calibration	38
	2.2.4.	Base-Tracking Algorithm Design	46
2.3.	Exp	erimental Evaluation Methods	62
	2.3.1.	Operating Room Floor Survey	62
	2.3.2.	Development Testing	62
	2.3.3.	Calibration Evaluation	64
	2.3.4.	C-arm Testing in a Simulated Operating Room	66
Chapte	r 3:	Experimental Results	78
3.1.	Ор	erating Room Floor Survey	78
3.2.	Dev	velopment Testing	79
	3.2.1.	Comparing Open-Loop Tracking Methods	79
	3.2.2.	Comparing Closed-Loop Tracking methods	80
	3.2.3.	Comparing Descriptor Matching Algorithms	81
	3.2.4.	Comparing Detector and Descriptor Algorithms	82
3.3.	Cal	ibration Evaluation	83
3.4.	C-a	rm Testing in a Simulated Operating Room	84
	3.4.1.	Summary of Results	85
	3.4.2.	Open-Loop Translation Error	86
	3.4.3.	Open-Loop Rotation Error	86
	3.4.4.	Closed-Loop Repositioning Translation Error	87
	3.4.5.	Closed-Loop Repositioning Rotation Error	89
	3.4.6.	Reacquiring X-rays	90
3.5.	Dis	cussion of Results	93
	3.5.1.	Translation Error in Open-Loop Tracking Algorithm	94
	3.5.2.	Rotation Error in Open-Loop Tracking Algorithm	95
	3.5.3.	Translation Error for Closed-Loop C-arm Repositioning	97
	3.5.4.	Rotation Error for Closed-Loop C-arm Repositioning	98
	3.5.5.	Reacquiring X-rays	99
Chapte	r 4:	General Discussions and Conclusions	. 100
4.1.	The	esis Contributions	. 100
4.2.	Lim	itations	. 101
4.3.	Fut	ure Directions	. 103
4.4.	Cor	ncluding Remarks	. 105
Referer	nces		. 107
Append	dix A – I	ighting Mounting System	. 112
Append	dix B – C	Camera Mounting System	. 117

List of Tables

Table 1	Fluoroscopy time and patient entrance surface doses (ESD) for various orthopaedic procedures (Tsalafoutas, et al. 2008)	6
Table 2	Comparison of technical specifications for previous camera (1), and current camera (2)	25
Table 3	Comparison of GigE and USB3.0 interface standards (FLIR Integrated Imaging Solutions Inc. 2015)	25
Table 4	Technical specifications for the camera (1), and lens (2)	26
Table 5	Comparison of five ARM-based embedded boards recommended by FLIR	27
Table 6	Technical specifications for Odroid-XU4 ARM board	28
Table 7	Technical specifications for Odroid-VU7+ touch screen interface	28
Table 8	Bill of materials for the base-tracking system	37
Table 9	Mean RMSE value of a straight line before and after applying calibration parameters	83
Table 10	Summary of the relative tracking error rates	85

List of Figures

Figure 1	(right) The C-arm fluoroscopy machine used in this thesis, and (left) monitoring workstation where the images are viewed.	3
Figure 2	All available C-arm movements	4
Figure 3	A pattern demonstrating A) a corner, B) an edge, and C) a featureless region	16
Figure 4	C-arm repositioning stats for two orthopaedic surgeons attempting to recreate a given radiograph (Touchette 2017)	20
Figure 5	An aerial view of a C-arm with the locations of wheeled base an imaging center labelled. The angular error that would cause an error in the imaging center of 4.3 cm is identified in the triangle on the right	21
Figure 6	Diagram of the Siemens Arcadis Mobile C-arm showing (left) the mounting location of the camera and user interface	23
Figure 7	Blackfly camera geometry	24
Figure 8	Chameleon camera geometry	24
Figure 9	Odroid-VU7+ touch screen interface.	29
Figure 10	Circuit diagram of the LED ring	30
Figure 11	Lighting system (white) mounted to camera.	30
Figure 12	An annotated drawing of the bottom light holster part with a snap fit component on its rim	31
Figure 13	Design variables for snap-fit hook	31
Figure 14	Relevant design dimensions for snap-fit hook (all dimensions in mm).	32
Figure 15	(left) Wheeled C-arm base, (right) wheeled C-arm base with rubber cover removed to expose the two holes used to mount the camera.	33
Figure 16	Camera mounting system with labelled C-arm attachment points and camera mounting point	34
Figure 17	Annotated drawing of camera mounting system with labelled C-arm attachment points, camera mounting point, and adjustable slots	34
Figure 18	Mounting location of the Odroid-XU4 on the bottom of the mounting platform	35
Figure 19	Two types of radial lens distortions: (left) barrel distortion and (right) pinhole distortion. The distortion in these images are exaggerated for clarity	38
Figure 20	Examples images of the calibration checkerboard pattern.	42
Figure 21	Example image of the scale and perspective calibration target.	43
Figure 22	(Starting top left) Original image, bilateral filter applied, Hough circles detected, mask applied, Canny edges detected, Hough lines detected, intersection points on original image.	45
Figure 23	ORB features in an image of an operating room floor	51
Figure 24	Initial screen when OPTIX program is started.	55
Figure 25	OPTIX GUI showing two saved points; the red point is currently activated for	
-	reacquisition	56

Figure 26	OPTIX GUI showing two saved points, the red point is currently activated for reacquisition. The labelled dark grey circle and light grey triangle represent the 5 mm and 3.1° tolerances respectively.	57
Figure 27	OPTIX GUI showing successful reacquisition of a previously saved position	58
Figure 28	OPTIX GUI calibration tab	58
Figure 29	Flowchart of the base-tracking algorithm with color-coded processes and queues. Purple background is the user interface process, green background is the camera process, and orange background is the tracking process. Queues of matching color are connected	60
Figure 30	Testing setup used during development, including a camera slider, camera, and vinyl flooring.	63
Figure 31	Example images of the line used for calibration evaluation.	65
Figure 32	BioEng Lab setup for C-arm testing	66
Figure 33	An example of C-arm motion for open-loop translation testing	67
Figure 34	Diagram of C-arm motion for open-loop rotation testing.	68
Figure 35	Diagram of side-to-side C-arm motion path for testing.	69
Figure 36	Diagram of in-and-out C-arm motion for testing.	69
Figure 37	Diagram of oblique C-arm motion path for testing	70
Figure 38	Diagram of four-position C-arm motion for testing.	70
Figure 39	Optimal range for Optotrak marker placement (©Northern Digital Inc., with permission).	71
Figure 40	The positions of the three Optotrak markers on the C-arm in the BioEng Lab	72
Figure 41	Optotrak probe used to digitize the floor plane.	73
Figure 42	OR floors from (left to right) the Kelowna General Hospital, Kelowna General Hospital, Vancouver General Hospital, and Vancouver General Hospital	78
Figure 43	(left) The flooring used for C-arm testing in the BioEng lab at CHHM and (right) for development testing with the camera slider.	78
Figure 44	Preliminary absolute accuracy results comparing Lucas Kanade optical flow tracking with ORB descriptor matching.	80
Figure 45	Absolute repositioning error for descriptor matching with and without a loop closure step.	81
Figure 46	Mean loop speed for feature matching algorithms over 3000 loops.	81
Figure 47	Mean loop speed for the detector and descriptor algorithms over 1000 loops	82
Figure 48	Mean number of matches per second for the detector and descriptor algorithms	83
Figure 49	Example result for the evaluation of the straightness of a linear feature before and after applying calibration parameters. RMSE value is shown for each line of best fit	84
Figure 50	Mean tracking error for all C-arm repositioning tasks executed with the absolute position recovery algorithm (blue) compared against the accuracy thresholds for typical clinical applications (orange).	85
Figure 51	Translation error for relative tracking of C-arm motion without loop closure step. The dashed lines show the 0.75% error envelope and the 5 mm bounds for clinical acceptability.	86

Figure 52	Rotation error for relative tracking of C-arm motion without loop closure step. The dashed lines show the 5% error envelope and the 3.1° error bounds.	.87
Figure 53	Mean translation error for each of the four C-arm repositioning tasks	. 88
Figure 54	Closed-loop translation error as a function of distance travelled for all C-arm repositioning tasks. The dashed lines represent the clinically acceptable error	
	envelope	. 88
Figure 55	Mean rotation error for each of the four C-arm repositioning tasks	. 89
Figure 56	Rotation error as a function of distance travelled for all C-arm repositioning tasks. The dashed lines represent the clinically acceptable error envelope	.90
Figure 57 (to	op right) Initial x-ray, (top left) repositioned after side to side motion, (bottom left) after in and out motion, and (bottom right) after oblique motion	.91
Figure 58	The three reacquired x-rays overlaid on top of the initial image with the head to head distance labelled	. 92
Figure 59	The screw centerline and labelled angle for the (top right) initial x-ray, (top left) repositioned after side to side motion, (bottom left) after in and out motion, and (bottom right) after oblique motion.	.93

List of Abbreviations

AKAZE	Accelerated KAZE
ALARA	As low as reasonably achievable
ARM	Advanced RISC machine
BRIEF	Binary robust independent elementary features
CAOS	Computer-assisted orthopaedic surgery
CCD	Charge-coupled device
CDS	Crash dummy symbol
CMOS	Complementary metal oxide semiconductor
СТ	Computed tomography
CPU	Central processing unit
CSV	Comma separated value
Ds	Scattered Dose
ESD	Entrance surface dose
FLANN	Fast library for approximate nearest neighbors
FAST	Features from accelerated segment test
GigE	Gigabit Ethernet
GPU	Graphical processing unit
GUI	Graphical User Interface
Gy	Gray
ICRP	International Council of Radiation Protection
IMU	Inertial Measurement Unit
LED	Light emitting diode
MRT	Medical Radiation Technologist
OPTIX	On-board position tracking for intraoperative x-rays
ORB	Oriented FAST and rotated BRIEF

RAM	Random Access Memory
DICC	
RISC	Reduced instruction set computer
SIFT	Scale invariant feature transform
Sv	Sievert
TKA	Total knee arthroplasty
USB	Universal serial bus

Acknowledgements

I would first like to extend a tremendous thank you to my supervisor, Dr. Antony Hodgson. You have an infectious excitement and an unending supply of wonderful questions. You constantly push me to look at problems in a new light, and I am deeply grateful for everything you have taught me throughout my degree.

I would also like to extend my gratitude to my co-supervisor, Dr. Carolyn Anglin. Thank you for always being interested and taking the time to provide support, often from the other side of the world. Your constant motivation and new perspectives were invaluable additions to my graduate experience.

To Dr. Pierre Guy, my clinical co-supervisor, thank you for providing important insights from a clinical perspective and helping me to become a better engineer. Perhaps more importantly though, thank you for always being genuinely interested.

I am grateful to the Centre for Hip Health and Mobility for housing me during my degree and allowing me opportunities to learn and grow. Thank you especially to all the people who make the Centre feel like a community.

I am abundantly thankful to all my lab members in the Surgical Technologies Lab. Whether it be direct help with my project, inspiration through your own innovative work, or simply being supportive friends. I feel very lucky to have been surrounded by such an inspiring and genuine group of people. Hooman Esfandiari, you've been a great mentor and friend, thank you for all your input and support.

To my friends, on both coasts and in between, thank you for always lifting me up and giving me balance. Because of you, I have never felt a lack of support, laughter, or adventure.

There will never be enough room on a page for me to express the gratitude I have for my family. To my parents, Joan and Terry, you constantly motivate me to be the best I can be in all aspects of my life. You are both tremendous inspirations, and I am incredibly thankful to have you on my team. Chelsea, you are a shining gem in this world, and I am deeply proud to call you my sister.

Janet, my best friend, my adventure buddy, my muse, my love. Your support has been endless, and I lack the words to express how important you are to me. Thank you for keeping me sane, thank you

for challenging me, thank you for making my life musical and fun, and thank you for pushing me to be my best self.

Lastly, thank you to the National Sciences and Engineering Research Council for your financial support through Engineers in Scrubs.

Chapter 1: Motivation and Background

Mobile C-arm fluoroscopy machines have become an increasingly common imaging modality for orthopaedic surgeries. However, there is significant evidence to suggest that excess time and radiation exposure result from positioning C-arms in the operating rooms. Additionally, C-arms are presently used primarily as a qualitative tool. Techniques have been proposed to improve positioning, reduce radiation exposure, and enhance quantitative functionality, but they all rely on accurate tracking of the C-arm position. Several groups have attempted to develop position tracking systems, but there are presently no solutions that can be practically applied to the challenging operating room environment. The goal of this thesis is therefore to develop an accurate position measurement system for the C-arm base that can be effectively used in an operating room. The focus of this research was primarily technology development with the aim of making steps towards clinical implementation.

1.1. C-arm Use in Operating Rooms

C-arms are normally maneuvered and operated by Medical Radiation Technologists (MRT) in the operating room based on verbal instructions from the surgeon. MRTs are trained in the operation of multiple forms of medical imaging, including conventional x-ray, computed tomography (CT), and fluoroscopy. The surgeon must rely on vocal communication with MRTs for C-arm positioning and imaging since the C-arm controls are located outside of the sterile field. Ineffective communication between surgeons and MRTs can lead to increased procedure time and radiation exposure (Pally and Kreder 2013). The C-arm must be manually positioned in the operating room, which is cumbersome due to its large size.

MRTs must often take multiple trial-and-error x-rays, commonly called scouting images, while attempting to achieve a radiographic view requested by the surgeon. A recent study carried out in our lab found that an average of 8 scouting images were required to achieve a satisfactory radiograph of a given anatomical view (Touchette 2017). These scouting images expose the patient and staff to unnecessary radiation and result in increased procedure time (Matthews, et al. 2007).

The typical workflow observed in surgery involves MRTs taking multiple initial images of the patient's anatomy of interest. There are commonly multiple radiographic views that are important to the surgeon. In a tibial intramedullary nailing procedure, for example, the surgeon would desire fluoroscopic images of the site of injury, the ankle, and the knee. The C-arm is maneuvered frequently

throughout procedures, depending on the complexity of the surgery. One study found the number of C-arm movements in orthopaedic trauma procedures to range from approximately 10 to 50, with the majority consisting of gross movements of the wheeled base (Suhm, et al. 2004). We have observed up to 100 C-arm movements in intramedullary nailing procedures, with base movements again being the most frequent. It is common for a surgeon to request multiple repetitions of one or more specific x-ray positions to monitor the progress of tools or implants. These repetitions require that MRTs accurately reobtain previous C-arm positions, which typically requires the use of scouting images. It has been reported that 80% of all C-arm movements in orthopaedic trauma procedures were related to repositioning the C-arm to recreate a previous view (Matthews, et al. 2007) (Suhm, et al. 2004). We have observed MRTs placing lines of tape on the operating room floor as an attempt to mark previous x-ray positions. Placing the tape is time consuming and impractical since the lines must be removed between procedures. Increased operation time is associated with an increased risk of infection (Cizik, et al. 2012) and a longer length of stay in hospital (Tan, et al. 2012).

1.2. C-arm Fluoroscopy

A C-arm fluoroscopy machine is an x-ray imaging modality that allows physicians to obtain real-time radiographic images throughout surgical procedures. C-arms have become ubiquitous in orthopaedic procedures for the surgeon to visualize patient bone structures, implants, and surgical tools. A C-arm, shown in Figure 1, consists of an x-ray emitter and image intensifier attached to a C-shaped arm with multiple degrees of freedom connected to a wheeled base. A separate wheeled podium has two viewing screens that display the radiographic images.



Figure 1 (right) The C-arm fluoroscopy machine used in this thesis, and (left) monitoring workstation where the images are viewed.

C-arm imaging allows surgeons to monitor the position of tools and implants inside the patient, which enables the use of minimally invasive surgical techniques. Minimally-invasive procedures tend to reduce blood loss, pain, likelihood of infection, and time spent in hospital (Grelat, et al. 2016), and have led to decreased patient morbidity and shorter operation times (Mahajan, et al. 2015). C-arms are used in orthopaedic procedures such as fracture fixation using intramedullary nails or plates, pedicle screw insertion, and sacroiliac screw placement in fractured pelvis fixation (Tsalafoutas, et al. 2008). They are also used in surgical procedures outside of orthopaedics such as brachytherapy dosimetry (Moult, et al. 2011), and catheter placement. The primary focus of this thesis will be on orthopaedic applications.

C-arms have achieved widespread surgical use primarily because of their mobility and ability to produce images on demand with relatively low radiation doses. A standard mobile C-arm has 5 intrinsic degrees of freedom, which are clinically referred to as orbit, wig-wag, up-down, in-and-out, and tilt. C-arms can also freely translate or rotate on the floor plane using the wheeled base, allowing for a full 6 degrees of freedom, illustrated in Figure 2. This mobility allows surgeons to obtain radiographic images from virtually any desired viewing angle (within the constraints of the operating table). Multiple images are often taken from different angles to ascertain the 3D position of surgical tools and implants as they are positioned inside the patient.



Figure 2 All available C-arm movements.

Another primary factor motivating widespread C-arm use is the ability to obtain x-ray images in real time. The C-arm can be operated in continuous video mode or single image mode. The images are displayed on the viewing screens in both cases, allowing the surgeon to receive immediate and on-demand visual feedback.

The primary drawback of C-arm fluoroscopy as an intraoperative imaging modality is the fact that it emits harmful ionizing radiation. It is important to consider the effects of this radiation exposure on both patients and surgical staff. The following section will further explore ionizing radiation exposure from C-arm imaging.

1.3. Ionizing Radiation in the Operating Room

Ionizing radiation can affect humans through two mechanisms, one causing deterministic effects and the other creating stochastic effects. Deterministic effects result from direct exposure to a high-dose beam of radiation. Stochastic effects are caused by an accumulation of low-dose radiation exposure over time. Radiation exposure is commonly expressed in Grays (Gy), where 1 Gy = 1 joule/kilogram, and Sieverts (Sv), where 1 Sv is the equivalent biological effect of absorbing 1 Gy into a body tissue.

Deterministic effects depend on radiation dosage and location on the body, and can range in severity from skin reddening to sterility. Typical x-ray and fluoroscopy machines operate well below the level

of radiation exposure required to cause deterministic effects, and they are therefore not typically a concern. Deterministic effects are not expected to manifest in the average person at entrance surface doses (ESD) lower than 6000 mGy (6 Gy), which would require 1500 minutes (25 hours) of direct C-arm exposure (Geleijns and Wondergem 2005). Table 1 shows fluoroscopy exposure times and patient ESD for various orthopaedic procedures, which are all significantly lower than the amount required to cause deterministic effects. The dosage levels for hospital staff are significantly lower since their primary exposure is through scattered x-rays that reflect off the patient and table. However, precautions are taken to monitor and minimize exposure for all patients and staff members according to the As Low As Reasonably Achievable (ALARA) principle (Mahajan, et al. 2015) (Hiniker and Donaldson 2014).

Surgical Procedure	Number of Patients	Fluoroscopy time (min)	ESD (mGy)
Intramedullary nailing of peritrochanteric fracture	113	3.2 ± 1.7	183 ± 138
Open reduction and internal fixation of malleolar fracture	16	1.5 ± 1.2	21 ± 27
Intramedullary nailing of diaphyseal femoral fracture	15	6.3 ± 2.7	331 ± 21
Arthroscopy for anterior cruciate ligament reconstruction with artificial ligament	13	0.9 ± 0.7	19 ± 20
Tibial plateau (plate and medullary screws)	13	1.2 ± 1.0	35 ± 36
Bilateral pedicle screw placement in the lumbar spine	11	0.8 ± 0.6	46 ± 32
Tibial intramedullary nailing	8	5.7 ± 3.5	137 ± 111
Fracture of the distal radius fixed with plate	7	1.8 ± 0.9	17 ± 10
Bilateral pedicle screw placement in the cervical spine	4	4.2 ± 3.0	173 ± 162
Vertebroplasty	4	5.1 ± 1.3	323 ± 51
All procedures	204	3.0 ± 2.3	149 ± 152

Table 1Fluoroscopy time and patient entrance surface doses (ESD) for various orthopaedic
procedures (Tsalafoutas, et al. 2008)

Current research tends to focus on stochastic effects, which result from accumulated radiation exposure over time. The primary mode of radiation exposure for hospital staff is through repeated daily low-dose scatter over the course of their careers. Radiation damage accumulates in tissues, and even low doses can increase the likelihood of stochastic effects when repeated over time (Grelat, et al. 2016). Long-term exposure can lead to serious negative consequences such as the development of cancer, cataracts, and other disorders (Srinivasan, et al. 2014).

A recent study reported that female orthopaedic surgeons in the United States have an increased prevalence of breast cancer (2.90X) and any cancer (1.85X) relative to the general population (Chou, et al. 2015). Male spine surgeon members of the Scoliosis Research Society have been found to have

a higher incidence of thyroid cancer (25X) and other cancers than the general male population (Wagner, Lai and Asher 2006). Additionally, interventional cardiologists present an increased incidence (52%) of radiation-associated cataracts relative to controls (9%) (Ciraj-Bjelac, et al. 2010).

Although there are serious potential side-effects involved with exposure to ionizing radiation, the resulting x-ray images are vital to many surgical procedures. Fortunately, most of these effects can be prevented or minimized through safe work practices and safety-focused engineering design. A large focus of research has therefore been aimed at safe radiation protocols and protective measures. The following section discusses methods of mitigating radiation risk, which provides context for this thesis work as a means of reducing radiation exposure.

1.4. Enhancing C-arm Utility

In addition to reducing surgical radiation exposure, this thesis is motivated by a desire to enhance the functionality of the C-arm as an intraoperative tool. The C-arm is currently used primarily in a qualitative manner to provide surgeons with visual guidance throughout procedures. However, its role as a quantitative tool could be expanded if we had a practical way to accurately calculate and track the position and orientation of the C-arm. Several functions that could be developed for a tracked C-arm include an expanded field of view for alignment measurements, visual limb measurement, and computer assisted orthopaedic surgery (CAOS).

An expanded field of view is useful for intraoperatively evaluating alignment for several procedures that require long-form radiographs. Spinal alignment in scoliosis correction presently requires the use of plain radiographs to perform alignment measurements. These radiographs are sometimes not large enough to visualize enough adult spinal anatomy (Vidal, 2015) (Vidal, et al. 2016). Additionally, since the C-arm is already present in the operating room to perform these procedures, using it to acquire long form radiographs could reduce procedure time and improve accuracy. The C-arm field of view can be expanded by stitching multiple radiographic images together to create a panoramic view (Amiri, et al. 2014). The C-arm pose must be tracked to reconstruct an accurate panoramic view of the anatomy, which would provide the surgeon with a powerful intraoperative tool to evaluate alignment metrics.

Intraoperative limb length measurement is performed by placing a radio-opaque ruler next to the limb and taking C-arm images at both ends of the desired bone. This allows the surgeon to measure the patient's bones and determine appropriate implant size. There are two primary issues with this

method however, related to accuracy and radiation exposure. Radio-opaque rulers have been found to result in incorrect intramedullary nail sizes in 6% of cases (Galbraith, et al. 2012). These errors are likely due to the ruler being misaligned with the bone, parallax from non-parallel C-arm images, or the ruler changing position between images. Additionally, surgeons sometimes need to hold the ruler in place to get a useful measurement, which causes their hands to be directly in the x-ray beam. A tracked C-arm could be calibrated to provide quantitative limb measurements with high precision, while allowing surgeons to stand away from the table and receive smaller doses of radiation.

Computer-assisted orthopaedic surgery involves tracking surgical tools and patient anatomy to provide intraoperative guidance. Computer-assisted navigation has been shown to increase accuracy in pedicle screw placement (Tian, et al. 2017), and leads to improved implant alignment in TKA (Todesca, et al. 2017). A tracked C-arm is required for many CAOS systems to register fluoroscopic images with preoperative CT scans. A practical method of C-arm tracking could increase the usability of navigated procedures, leading to increased surgical accuracy for a greater number of patients.

1.5. Existing Solutions and Gap

Several research groups have attempted to develop C-arm tracking systems, and while some have achieved satisfactory levels of accuracy, their limitations often lie in the practicality of implementing the proposed systems in the operating room. Clinical operating rooms are a challenging environment, with cables, blood splatter, many personnel, and sterility issues. Notable proposed solutions in the existing literature include external infrared cameras, electromagnetic trackers, radiographic fiducials, and in-line cameras. The following section will briefly outline each technology and identify the primary limiting factors to demonstrate the gap that this research aims to fill.

1.5.1. External Optical Tracking

External optical tracking with infrared markers is a well-established method for C-arm positional tracking. A stereo infrared camera is positioned in the operating room such that it is facing the C-arm, and infrared markers are rigidly attached to the gantry. These markers can either be active infrared emitters or passive markers designed to be detected by the camera. There are several commercial-grade optical tracking systems available, such as Optotrak (Northern Digital Inc., Waterloo, ON, Canada) or Polaris (Northern Digital Inc, Waterloo, ON, Canada), which can track all 6 degrees of freedom of the C-arm with clinically useful levels of accuracy. One group at the University of Bern in Switzerland compared the accuracy of two Optotrak 3020s, a Polaris P4, and a Polaris Spectra and found tracking errors ranging from 0.86 mm to 1.15 mm (Rudolph, Ebert and Kowal 2010). These

optical tracking systems are also capable of tracking surgical tools and the patient by attaching additional infrared markers.

The primary limitation for this C-arm tracking method is the reliance on line of sight. The infrared markers must constantly be within the field of view of the tracking camera, which requires alterations in surgical workflow and operating room setup. The Optotrak 3020 has a pyramid-shaped effective workspace that has a cross section of 4.2 m by 3 m at the large end, reducing to 0.5 m by 0.9 m over a depth of 5.5 m (Northern Digital Inc. 2015). The Polaris sensors also have pyramid-shaped effective volumes, with cross sections of 1.57 m by 1.31 m, reducing to an unspecified smaller cross section over a depth of 1.45 m (Northern Digital Inc. 2017). This requirement can be challenging in a busy operating room with many personnel and other equipment. In addition, these commercial tracking systems are expensive to purchase, costing tens of thousands of dollars.

1.5.2. Electromagnetic Tracking

One group from Johns Hopkins University in Maryland has developed a C-arm tracking setup based on electromagnetic tracking (Yoo, et al. 2013). An electromagnetic field generator is mounted to the operating table, which tracks the motion of a sensor attached to the C-arm. The system can track all 6 degrees of freedom with a mean tracking error of 1.90 ± 1.38 mm within a cylindrical volume with a diameter of 460 mm and a depth of 650 mm (Yoo, et al. 2013).

Several practical limitations exist for this tracking system. The accuracy degrades near the edge of the field of view of the tracker, which means that the field generator must be moved along the table to accurately track the C-arm over the full length of a patient. This movable mounting setup requires a significant modification to the operating table. Additionally, electromagnetic tracking is known to be susceptible to errors in the presence of metallic objects, such as surgical tools and other equipment (Moult, et al. 2011).

1.5.3. Radiographic Fiducials

Researchers at Queens University in Ontario have developed a pose estimation method using radiographic fiducials that allows the position of the C-arm to be estimated following a fluoroscopic shot (Moult, et al. 2011). A pattern of radio-opaque beads is placed on the operating table so that it appears in the C-arm images. The appearance of the bead pattern is used to calculate the pose of the C-arm throughout the procedure. The system can accurately track all 6 degrees of freedom with translational errors of 0.64 ± 0.24 mm and rotational errors of $0.68 \pm 0.06^{\circ}$ (Moult, et al. 2011).

The limitations associated with this tracking approach are primarily associated with the use of a physical tracking pattern on the operating table. The beads must be in the field of view to determine the pose, which limits the effective range of the C-arm. The bead pattern must also be physically present on the operating table, and cannot change position, which could interfere with the surgical procedure. Additionally, the beads appear in each radiographic image, which could interfere with the appearance of anatomical features and decreases the quality of the image. Finally, the algorithm requires an average of 31 seconds to estimate the C-arm pose, which may not be not practical for clinical use (Moult, et al. 2011).

1.5.4. Mounted In-line Camera

One research group at Johns Hopkins University Maryland has developed an on-board, optical tracking system for C-arms (S. Reaungamornrat, et al. 2011). The system makes use of a stereo camera mounted on the C-arm gantry near the image intensifier. A reference pattern is placed on the patient table, which is used to calculate the pose of the C-arm. Accuracy testing has demonstrated a tracking error of 0.87 ± 0.25 mm (S. Reaungamornrat, et al. 2014). The system has several notable benefits: it can track all 6 degrees of freedom, it makes use of an established, accurate, commercial-grade video tracker, (MicronTracker, Claron Technology Inc., Toronto, ON, Canada), and surgical tools can be tracked by attaching reference patterns.

A group from Technische Universitat Munchen in Germany has created a similar setup with a single monocular camera mounted on the C-arm (Wang, Traub and Heining, et al. 2008). The camera is mounted near the image intensifier and is aligned to be parallel to the x-ray beam. A reference pattern is placed near the patient, and the appearance of the reference pattern is used to calculate the C-arm pose. The system has been shown to generate stitched x-ray panoramas with errors of less than 1% of the true image position (Wang, Traub and Heining, et al. 2008) (Wang, Traub and Weidert, et al. 2010).

Some practicality issues exist for both mounted-camera tracking systems, primarily concerning the use of a reference pattern for tracking. The pattern must constantly be in the field of view of the camera to calculate the pose, which can limit the effective range of the C-arm. Also, the pattern must be physically present on or near the operating table, which may be inconvenient for surgical staff during an operation. Finally, tracking capabilities are lost if the pattern changes position during the operation.

1.5.5. TC-arm

A low-cost on-board C-arm tracking system (called the TC-arm) was developed by other members of our research group at the University of British Columbia and the University of Calgary (Amiri, et al. 2014). Two inertial measurement units (IMUs) are attached to the C-arm gantry to measure orbit, tilt, and wig-wag motions, and two laser distance sensors are used to measure up-down and in-and-out motions. This system can track the C-arm with errors of 1.5 ± 1.2 mm (Amiri, et al. 2014). The TC-arm is low-cost relative to conventional tracking systems, and has no line of sight requirements, making it a practical solution for clinical use. The primary limitation of the TC-arm is the fact that it only tracks 5 degrees of freedom, while assuming the C-arm base remains stationary. Since the base is frequently moved during surgery, this limitation significantly reduces the effective range of tracking the C-arm. The desire to remove this limitation is one of the factors that motivates the current work.

1.5.6. Gap in the Current Solutions

The primary limitations of conventional C-arm tracking systems are line-of-sight requirements and comparatively high costs. Several research groups have proposed alternative tracking methods, but each of the solutions discussed have practical limitations preventing their clinical implementation in an operating room. One common limitation of several tracker designs is a lack of effective range, which is particularly hindering for some orthopaedic procedures that require the C-arm to take images over a large working area (e.g. imaging the full length of an adult spine). Some of the proposed solutions have a limited range because they rely on fiducials or reference patterns, or a limited effective sensor range. The TC-arm system measures the position of the C-arm joints relative to the base and therefore requires an additional system to measure base movement (the creators have implemented a fiducial-based system for this purpose).

Tracking the base of the C-arm is useful for integrating with the TC-arm, or a similar joint tracking system, to achieve full, 6 degrees of freedom tracking. Base-tracking is also a useful development on its own since base motion is the most common C-arm movement. Base-tracking can be used to provide positioning guidance to MRTs when reobtaining a previous x-ray position in scenarios where base motion is the primary movement. Planar panoramic images, such as those desired for scoliosis correction, can be generated solely through base translation without adjusting the C-arm joints. A base-tracking system could also be integrated into a robotic base to provide closed loop control of base positioning maneuvers.

The desired output of a base-tracking system would be global planar translation and rotation measurements relative to any arbitrary starting position. The system must be designed to perform in the challenging operating room environment. Additional details and specifications can be found in Chapter 2.

1.6. Observed Need: C-arm Tracking

We have identified a need for a practical intraoperative C-arm tracking system to reduce unnecessary radiation exposure, decrease procedure time, and expand the quantitative functionality of the C-arm.

A system capable of tracking the position and orientation of the C-arm throughout a surgical procedure has many potential applications, including repeated positioning guidance, artificial x-ray generation (Touchette 2017), panoramic stitching (Amiri, et al. 2014) (Amini 2016), and limb length measurement.

These applications have three primary potential benefits: 1) reduction in radiation exposure for patients and operating room staff, 2) reduction in procedure times, and 3) more accurate information for the surgeon.

Radiation exposure could be reduced by providing positional guidance to MRTs when reacquiring a previous image position to reduce the need for scouting images. Artificial x-rays can provide visual feedback to the MRT and surgeon to ensure that the C-arm is in the correct position before acquiring a real radiograph. Allowing the surgeon to stand further away from the x-ray beam while measuring limbs by reducing their reliance on radio-opaque rulers could also contribute to a reduction in radiation exposure.

Operation times could be decreased as a by-product of reducing the number of scouting images taken during a procedure. Enabling the use of C-arms to obtain long-form radiographs through panoramic stitching could also reduce overall procedure times compared to the use of plain radiographs.

Expanding the field of view of the C-arm could enable surgeons to make informed intraoperative decisions about alignment in spinal correction procedures. Providing accurate positioning guidance could lead to more accurate imaging and smoother communication between surgeons and MRTs.

1.7. Base-Tracking Methods

There are several viable sensor options that must be considered when selecting the most appropriate solution for C-arm base-tracking. The previous sections have established that a practical tracking system for an operating room would ideally be mounted on-board the C-arm to eliminate line of sight issues. Three possible approaches include: wheel encoders, optical flow sensors, and camera-based tracking, all of which are presented in the following sections. Ultimately, camera tracking was chosen as the most suitable method for this application, and our implementation is further detailed in Chapter 2.

1.7.1. Encoders

Rotary encoders are sensors typically used to measure relative displacements of a wheeled device based on wheel rotation and circumference. Wheel-based odometry has well-known, inherent, systematic errors that result from calibration error, imperfect wheel geometry, and unequal wheel sizes (Borenstein, et al. 1997). Encoder systems are also subject to unpredictable random errors caused by bumps, floor irregularities, and wheel slippage (Bonarini, Matteucci and Restelli 2005). Additionally, installing a shaft encoder would require alterations of the current C-arm wheel setup, making it difficult to quickly and inexpensively retrofit the system.

Rotary encoder systems are also inherently open-loop relative tracking systems, which means that measurement errors will accumulate over time.

1.7.2. Optical Flow Sensors

An optical flow sensor is a common form of planar displacement measurement that is perhaps most commonly found in optical computer mice. Light from a light emitting diode (LED) is directed to the ground surface using a system of mirrors. The reflected light then passes through a lens and is captured by a charge-coupled device (CCD) or a complementary metal oxide semiconductor (CMOS) sensor. The sensor is typically mounted directly onto a processing chip, which calculates displacement as the device is moved across a surface. Optical flow sensors overcome many of the issues associated with encoders since they are non-contact, and therefore are unaffected by wheel geometry imperfections or ground surface irregularities.

This sensor option was tested in an initial iteration of the C-arm base-tracking project by our research group (Esfandiari 2014), but the accuracy was not adequate. The measurement error ranged from

2.1% over 39.9 mm of travel to 5.3% over 59.1 mm of travel. These results indicate that the error in this measurement technique will increase to unacceptable levels over displacements likely to be encountered in practice, so it was determined that the accuracy is insufficient for this application.

In addition to demonstrating an unacceptable level of measurement error, the optical flow sensor system, like the encoder approach, is an open-loop relative movement sensor. It could be possible to create a closed loop system by adding reference markings to the floor of the operating room, but this is an undesirable approach.

1.7.3. Camera-based Tracking

Camera-based tracking is frequently employed in fields such as robotics and autonomous vehicle navigation. Camera-based navigation systems can be developed using relatively inexpensive off-theshelf components, and are able to extract a large amount of information from the surrounding scene. Computer vision is the name of the technical field that uses cameras as primary sensors to make measurements and interpret the surroundings. Camera tracking techniques are non-contact and independent of wheel geometry, which enables them to bypass many of the limitations faced by encoders. Another distinct advantage of camera-based tracking is the ability to perform absolute loop closure. An image acquired at a position of interest can be saved and later used to reset any accumulation of error without requiring external markers.

Camera-based navigation techniques can be broadly divided into two categories: forward-facing and downward-facing. Forward-facing cameras are commonly used in robotics and vehicle navigation, and have been widely explored for both outdoor (Nister, Naroditsky and Bergen 2005), and indoor applications (Rodriguez-Telles, Torres-Mendez and Martinez Garcia 2013) (Gadd and Newman 2015). Given the wide variety of moving equipment and the number of people in an operating room, a forward-facing camera is likely not a practical solution for this application.

Downward-facing cameras calculate camera motion relative to the ground surface using the appearance of visual texture. This sensor configuration has been explored by several research groups for robot localization (Kelly 2000) (Dille, Grocholsky and Singh 2009) (Killpack, et al. 2010), and vehicle navigation (Aqel, et al. 2015). The most comparable application has reported promising results with errors of less than 0.52% of the total distance travelled (Killpack, et al. 2010) To our knowledge, a downward-facing camera-based tracking system has not been developed for any equipment in a hospital operating room. The presence of visual texture in operating room floors is

further explored in Section 3.1, and previous work by our research group has shown this to be a viable technique (Esfandiari 2014). A floor-facing camera could be mounted to the base of the C-arm so that the measurements will not be interrupted by surgical staff.

1.7.4. Proposed Solution

We therefore propose a C-arm base-tracking system based on a floor-facing, on-board camera. The system will calculate incremental translation and rotation information between each video frame, and will have the ability to reset accumulated errors through absolute loop closure. The natural visual texture present in operating room floors will be used for tracking purposes. The camera will be mounted beneath the base of the C-arm so that the lighting can be controlled and the surgical staff will not accidentally interfere with the measurements during a procedure.

1.8. Computer Vision Background

Camera-based motion tracking involves calculating the change in position of a camera from one video frame to the next. There are several computer vision techniques that can be employed for motion tracking, which can broadly be divided into dense and sparse methods.

Dense algorithms make use of every pixel in each video frame to determine the change in position. These dense tracking algorithms are commonly used for image registration, and function by iteratively searching for the optimum geometric transform between consecutive video frames (Evangelidis and Psarakis 2008). The geometric transform can then be used to extract incremental displacement and rotation information between each frame. These techniques tend to be very accurate, since they make use of all available information in each image. However, dense tracking algorithms are computationally expensive, often making them too slow for real-time applications.

Sparse tracking algorithms are designed to leverage features in images that can be matched in consecutive frames to determine the geometric transform between the frames. Features are small regions of pixels that are well localized within the image. Figure 3 shows examples of a corner, an edge and a featureless region. Corners are characterized as areas where the image gradient changes in more than one direction (Forsyth and Ponce 2012). A small window located at a corner would appear different if it was translated in any direction. Conversely, a small window placed on an edge could be moved to a new position along the edge and appear the same, indicating that it is not well localized. Edges are useful in other computer vision applications, but corners are necessary for feature-based motion tracking.



Figure 3 A pattern demonstrating A) a corner, B) an edge, and C) a featureless region

Feature-based tracking techniques tend to be less accurate than dense methods, but they are considerably less computationally expensive. Since the features used in sparse tracking methods are high information regions within an image, they are still able to achieve low levels of error. There have been many feature detection algorithms developed that can be harnessed for motion tracking.

Two main categories of methods to identify corresponding features in consecutive frames are optical flow and descriptor matching.

Optical flow methods, such as the well-known Lucas Kanade algorithm (Lucas and Kanade 1981), detect features in the first frame and then use information about this set of features to search for corresponding features in the following frame. The features are detected using an algorithm such as Harris Corners (Harris and Stephens 1988), or Features from Accelerated Segment Test (FAST) (Rosten and Drummond 2006). These algorithms define a small window around the location of a feature point in the first image, and use this window to search for the corresponding feature in the next image. This technique assumes that feature points will only move a small amount from one image to the next, which renders them ineffective in cases of large frame-to-frame displacements. Features will be lost from the field of view as camera motion progresses, requiring that a new set of features be re-detected when a lower threshold amount is reached. Before re-detection, the feature points end up clustered to one side of the image. This means that the information in a portion of the image is not being exploited, which decreases the robustness of the tracking algorithm. Frequent re-detection can alleviate this issue, but this opposes the primary speed advantage of optical flow schemes. Optical flow algorithms are fast because they do not need to find feature points in every

video frame. For these reasons, optical flow algorithms tend to be faster than descriptor matching techniques, but they also tend to be less accurate.

Descriptor matching is a more robust method of finding corresponding features between video frames. There are multiple detection algorithms available, such as Scale Invariant Feature Transform (SIFT) (Lowe 2004), Accelerated KAZE (AKAZE) (Alcantarilla, Nuevo and Bartoli 2013), Binary Robust Independent Elementary Features (BRIEF) (Calonder, et al. 2010), or Oriented FAST and Rotated BRIEF (ORB) (Rublee, et al. 2011), with varying levels of performance. The overall method involves locating and describing features in two images using keypoints and descriptors. The algorithm first searches for keypoint locations, then calculates a histogram (in the case of SIFT and its variants) or a binary matrix (for BRIEF and its variants) that describes information such as the feature strength and orientation. This histogram or matrix is referred to as a descriptor, and is a useful way of identifying features. Once the features are detected in both images, an algorithm compares the keypoint descriptors to find matches between the images. There are multiple matching algorithms available, but the two most common are Brute Force and Fast Library for Approximate Nearest Neighbors (FLANN) (Muja and Lowe 2009). Descriptor matching techniques are robust due to several important characteristics. Descriptors uniquely describe feature points, allowing for consistent matching. Descriptor-based methods can handle larger frame-to-frame displacements than optical flow because there are no location assumptions made during matching. Keypoints are detected and matched between each video frame, rather than carrying a single set of feature points through consecutive frames as in optical flow. Finally, certain descriptor algorithms are designed to be insensitive to changes in lighting, orientation, and scale. The performance of multiple feature detection, description, and matching algorithms is explored in Section 3.2.3.

1.9. Previous Work

The initial research for this project was carried out by another member of our research group in 2014 (Esfandiari 2014). A camera calibration procedure was established and an offline tracking algorithm was developed to calculate the displacement and rotation that occurred in a video recorded by a floor-facing camera.

The tracking algorithm made use of the Harris Corner detector (Harris and Stephens 1988) and Lucas Kanade optical flow algorithm (Lucas and Kanade 1981), to find corresponding feature points to calculate the frame-to-frame homography. The incremental translations and rotations between each

frame were extracted from the homography matrix, which were then summed to determine the relative camera motion.

The system achieved translational and rotational errors of less than 2-4% of the total one-way distance travelled and 2.7% of the total one-way rotation, which is clinically acceptable for some applications (Esfandiari 2014). The effects of varying camera distances, velocities, and floor textures were explored, with the results indicating that this is an appropriate method for C-arm base-tracking. Part of this research also indicated the usefulness of absolute loop closure to eliminate the accumulation of incremental errors from the relative tracking scheme.

While this initial work demonstrated the potential for a floor-facing camera-based tracking system for C-arms, there were several aspects that required further development. The most prominent issue was the fact that the system operated offline, which is not practical for use in a clinical environment. A real-time implementation is required to meet the demanding needs of a surgical operating room. Due to the nature of the optical flow algorithm, features clustered to one side of the image as the camera moved. Until enough features were lost to warrant re-detection, this clustering effect meant that not all the available information was being used in each frame. A feature tracking scheme with more evenly distributed keypoints should increase the robustness of the system. This system was primarily a relative tracking method, although the use of a reference pattern on the floor was explored for the purposes of absolute loop closure. A marker-free loop closure system would be more suitable for clinical implementation.

This thesis focuses on continuing and extending the initial research carried out by Esfandiari. Updated methods are used to address each of the limitations identified in the original work, and multiple hardware changes are implemented to better suit the clinical environment.

1.10. Thesis Objectives

The goals of this thesis are to:

- 1. Design and implement a real-time base-tracking algorithm using a single, floor-facing camera mounted on a C-arm with the ability to accurately re-obtain a previous position; and
- 2. Verify the performance of the tracking system in a simulated operating room setting.

1.11. Thesis Outline

The thesis is laid out as follows:

Chapter 1: Information on C-arm usage and radiation exposure in the operating room, constituting the motivation for this research. A review of relevant studies available in the literature is presented along with relevant background information.

Chapter 2: The design and development of all aspects of the base-tracking system is detailed, including hardware selection, calibration, and the tracking algorithm. Experimental methods are also outlined in this chapter.

Chapter 3: The results of experimental testing are presented and interpreted. A discussion of the implications of these results is also included.

Chapter 4: A discussion of thesis contributions and limitations in the developed system is presented. Suggestions are included for the future directions of this research.

Chapter 2: Design and Development of a C-arm Base-Tracking System

This chapter provides a detailed overview of the design and development of a C-arm base-tracking system called OPTIX (On-Board Position Tracking for Intraoperative X-rays). The design requirements are outlined, along with a full treatment of the system development including hardware selection, a robust calibration procedure, algorithm design, and finally a description of the experimental methods.

2.1. Tracking System Requirements

The required level of accuracy for a C-arm base-tracking system depends on the application. One of the focuses of this research is on repositioning tasks, which require that the repeated radiographic image be clinically equivalent to the original. Another member of our lab group recently carried out a study that involved evaluating orthopaedic surgeons performing C-arm repositioning tasks (Touchette 2017). The position and rotation errors measured in this study are shown in Figure 4, and can be used to define a ceiling for allowable tracking system error.



Figure 4 C-arm repositioning stats for two orthopaedic surgeons attempting to recreate a given radiograph (Touchette 2017).

These data show that two x-ray images can be considered clinically equivalent with a C-arm lateral error of approximately 4.3 cm and angular error of 3.8°.

To further define the acceptable angular error, we consider the task of repositioning the C-arm so that an anatomical view of interest is within the field of view. The field of view is a circle approximately 20 cm in diameter, and the distance between the wheeled base and the C-arm imaging

center is approximately 80 cm. To ensure that the center point of the previously viewed image does not move more than the desired 4.3 cm towards the edge of the image, the C-arm base must be tracked with an angular error of less than $\tan^{-1}(\frac{4.3}{80}) = 3.1^{\circ}$, as depicted in Figure 5.



Figure 5 An aerial view of a C-arm with the locations of wheeled base an imaging center labelled. The angular error that would cause an error in the imaging center of 4.3 cm is identified in the triangle on the right.

For other base-tracking applications, such as panoramic image stitching or limb measurement, more stringent error bounds are required. A recent study by a member of our research group defined 5 mm as a clinically acceptable accuracy for panoramic stitching to evaluate scoliosis alignment intraoperatively (Amini 2016). This value was based on sagittal vertical axis (SVA) misalignment. SVA is defined as the horizontal distance between a vertical line drawn from the center of the C7 vertebra and the posterior superior corner of the S1 vertebra. A study has shown that an SVA below 50 mm results in a higher health-related quality of life (Schwab, et al. 2010), which led to the definition of 5 mm panoramic stitching error (10% of the proposed clinical margin) (Amini 2016). Additionally, limb lengths are considered clinically equivalent within a tolerance of 5 mm (Clave, et al. 2015). A base-tracking system with an accuracy of 5 mm could therefore be used to evaluate scoliosis alignment and measure limb lengths with clinically useful results.

Based on these various application criteria, we have decided to set the requirements for this system to be 5 mm of translational error and 3.1° of rotational error.

A base-tracking system must be able to function in an operating room to be clinically relevant. This requirement is less easily quantified than accuracy, but is of equal importance and has a large impact on design decisions. A practical base-tracking system must perform reliably without interfering with the surgical staff or contributing to a disruption in the operating procedure. Cables, tools, and other equipment in the operating room must not cause the performance of the system to drop to unacceptable levels.

2.2. OPTIX Design Stages

There are three principal design stages for the OPTIX system: (1) hardware selection, (2) camera calibration, and (3) base-tracking algorithm development.

2.2.1. Hardware Selection

The tracking system makes use of a floor-facing camera mounted onto the base of a C-arm. The system consists of several hardware components, including a camera, processing unit, user interface display, lighting, mounting, and power supply. The first design consideration was the camera mounting location, which subsequently influenced other hardware design choices.

Practical implementation in the operating room is one of the primary design goals guiding this process. It is therefore required that the camera be mounted such that the surgical staff will not enter the field of view without altering their standard workflow. There is a hollow space behind the front wheels of the Siemens Arcadis Mobile C-arm (Siemens, Erlangen, Germany) that provides an advantageous mounting location, as depicted in Figure 6. A similar hollow space is present on C-arms from major manufacturers including Siemens, General Electric, Ziehm, and Philips. It is highly unlikely that surgical staff will interfere with camera measurements in this location, and the dark enclosure provides the opportunity to easily control the lighting conditions.


Figure 6 Diagram of the Siemens Arcadis Mobile C-arm showing (left) the mounting location of the camera and user interface

The limitation of this mounting location is that it requires the camera to be placed very close to the floor, effectively limiting its field of view. The total distance from the floor to the C-arm wheel strut in this location ranges from 16 cm to 18.5 cm. Despite this limitation, this is the only practical mounting location on the C-arm that is reliably protected from accidental interference by personnel in the operating room.

2.2.1.1. Camera and Lens Assembly

A Blackfly Gigabit Ethernet (GigE) Vision grayscale camera (BFLY-PGE-09S2M-CS, FLIR Integrated Imaging Solutions Inc., Richmond, Canada) was used in the initial iteration of this project (Esfandiari 2014). Although the imaging performance was satisfactory, the camera geometry, Figure 7, was not compatible with the limited vertical space underneath the C-arm base. The Blackfly camera is designed with the GigE network connection port opposite to the imaging sensor, which requires the camera to be mounted closer to the floor to make space for the cable. We therefore opted to use a Chameleon Universal Serial Bus (USB) 3.0 Vision grayscale camera (CM3-U3-13S2M-CS, FLIR Integrated Imaging Solutions Inc., Richmond, Canada) instead.



Figure 7 Blackfly camera geometry.

The Chameleon camera geometry, Figure 8, is better suited to the limited space available beneath the C-arm base. The USB3.0 port is on a face that is perpendicular to the imaging sensor, which enables the camera to be mounted as far from the floor as possible. This cable configuration results in an estimated 75 mm of extra space based on measurements of the locking GigE cable required for the Blackfly camera. The overall length of the Chameleon along the optical axis is 29 mm, compared to 37 mm for the Blackfly. The total floor-to-camera height advantage for the Chameleon is approximately 83 mm. Given that the total available height ranges from 16 cm to 18.5 cm as mentioned in Section 2.2.1, this corresponds to a significant 45%-52% space savings.



Figure 8 Chameleon camera geometry.

The Chameleon and Blackfly cameras are similar in terms of imaging performance. A comparison of the technical specifications for the two cameras is presented in Table 2.

Make & Model	Imaging Sensor	Imaging Architecture	Resolution	Frame Rate	Machine Vision Standard
(1) FLIR BFLY-PGE-09S2M-CS	1/3", 4.08 μm pixels	Global shutter CCD	1288 x 728	30 fps	GigE Vision v1.2
(2) FLIR CM3-U3-13S2M-CS	1/3", 3.75 μm pixels	Global shutter CCD	1288 x 964	30 fps	USB3 Vision v1.0

 Table 2
 Comparison of technical specifications for previous camera (1), and current camera (2)

The table shows that primary differentiating factor between the two cameras is the machine vision standard. The machine vision standard refers to the global communication protocol standard used to interface with the camera. The Blackfly uses a GigE cable for communication while the Chameleon transmits data through a USB3.0 cable. A comparison of the two interface standards can be found in Table 3.

Table 3	Comparison Inc. 2015)	of GigE and USB3.0 interface standards (FLIR Integrated Imaging So				
		GigE	USB3.0			

	GigE	USB3.0
Bandwidth	125 MB/s	400 MB/s
Maximum cable length	100 m	3 m
CPU usage	Medium	Low
System cost (single camera)	Medium (\$70 USD)*	Low (\$50 USD)*

*Based on peripherals required to implement a single camera.

The USB3.0 interface provides advantages in terms of data transfer speed, CPU usage, and cost, all of which are important factors for this application. The primary disadvantage of USB3.0 is the maximum cable length of 3 m. In this project, however, the camera and processing unit are both mounted to the C-arm, so a long communication cable is not necessary. USB3.0 is therefore clearly the best interface option for this application, further supporting the use of the Chameleon camera.

The lens chosen for this project has a variable focal length that allows it to focus on objects at short distances. This is crucial since the camera is mounted close to the floor, and must bring the visual texture present on the flooring into focus. The technical specifications of the camera and lens assembly are summarized in Table 4.

Make & Model	Imaging Sensor	Imaging Architecture	Resolution	Frame Rate	Machine Vision Standard	Focal Length
(1) FLIR						
CM3-U3-13S2M-CS	1/3", 3.75	Global	1288 x 964	30 fps	USB3	2.8–8 mm
(2) Fujinon YV2.8x2.8SA-2	mm pixels	shutter CCD	1200 / 301	00.00	Vision v1.0	210 0 11111

Table 4	Technical specifications for the camera (1), and lens (2)
---------	---

The costs for all camera and lens components are outlined in Section 2.2.1.7.

2.2.1.2. Processing Unit

An on-board processing unit was determined to be the most appropriate configuration due to the mobile nature of the C-arm. This setup simplifies the data communication process since the camera and processing unit are both mounted directly to the C-arm, as opposed to an external processing unit requiring lengthy cables or wireless communication.

There are several requirements that a processing unit must satisfy to be appropriate for this application:

- A USB3.0 port to communicate with the Chameleon camera
- A multi-core CPU to enable parallel processing
- Enough processing power (RAM, CPU, GPU) to handle computationally demanding computer vision tasks

Advanced Reduced instruction set computer Machine (ARM) development boards are specifically designed to be low-power mobile processing units with advanced computing capabilities. FLIR, the manufacturer of the Chameleon camera, has published a Technical Note on five commonly used ARM-based embedded processing boards (FLIR Integrated Imaging Solutions 2016). According to the

results, all five boards can support computer vision applications with high-speed demands. Technical specifications and reported costs for each of the five boards are provided in Table 5.

	Reported Cost	CPU Cores	RAM	USB3.0 Port
Odroid-XU4	\$120 (Diigiit Robotics 2017)	Samsung Exynos5422 Cortex [™] -A15 2 GHz and Cortex [™] -A7 Octa core	2 GB	Yes
Samsung Exynos 5250 Arndale	\$324 (CNX Software 2017)	Exynos5 Octa Cortex [™] - A15 – 1.6 GHz quad core	2 GB	Yes
NVIDIA Jetson TK1	\$262.23 (NVIDIA 2017)	Tegra K1 SOC – quad core	2 GB	Yes
NVIDIA Jetson TX1	\$394.56 (Arrow 2017)	Cortex A57 – quad core	4 GB	Yes
NVIDIA DRIVE PX	\$1,500 (The Motley Fool 2016)	Togra V1 guad core		Voc
	\$15,000 (Forbes 2016)	regra VI – duan cole	4 00	TES

 Table 5
 Comparison of five ARM-based embedded boards recommended by FLIR

One of the primary focuses of this research is the development of a low-cost C-arm tracking solution. Given the recommendation from FLIR that any of these embedded boards can adequately support demanding computer vision tasks, the most economical board was chosen. The NVIDIA Drive is prohibitively expensive, and is sold to automobile manufacturers rather than directly to consumers, so this board was immediately ruled out. The Samsung Exynos was eliminated because it has a less powerful CPU than the Odroid-XU4, while costing more than double the price. The NVIDIA Jetson TK1 was tested by researchers at the University of Leuven in Belgium, and was found to have unsatisfactory speed performances for real-time computer vision applications (Hulens, Verbeke and Goedeme 2015). The NVIDIA Jetson TX1 has more computing power and would provide satisfactory performance, but it is the most expensive of the consumer grade boards. Ultimately, the Odroid-XU4 was chosen for this application since it has the required ports, sufficient processing power, and is less than half the price of the nearest competitor. More detailed technical specifications for the Odroid-XU4 are included in Table 6.

Make & Model	Operating System	СРU	GPU	RAM	Ports
		Samsung			2x – USB 3.0
Odroid-XU4 1	Linux Ubuntu 16.04	Exynos5422 Cortex [™] - A15 2 GHz and Cortex [™] -A7 Octa core	ARM Mali [™] - T628 MP6		1x – USB 2.0
				2 GB LPDDR3 RAM PoP stacked	1x – Ethernet RJ45
					1x – HDMI
					1x – eMMC5.0 HS400
					1x – microSD

Table 6 Technical specifications for Odroid-XU4 ARM board

The Odroid-XU4 is an open source board that runs the Linux operating system. This allows for a large degree of development control, which was useful in implementing the base-tracking system.

2.2.1.3. User Interface

A touch screen user interface display was chosen to provide visual information, encourage intuitive interaction with the tracking system, and avoid unnecessary components such as a computer mouse. The ODROID-VU7+ is designed specifically to interact with Odroid ARM boards, the specifications for the touch screen interface are detailed in Table 7.

 Table 7
 Technical specifications for Odroid-VU7+ touch screen interface

Make & Model	Screen Size Resolution		Touch Input Cost		
Odroid-VU7+	7 inch TFT-LCD	1024 x 600	5 finger capacitive touch input	\$120 CAD	

The software component of the user interface is detailed in Section 2.2.4.4. Figure 9 shows the Odroid-VU7+ and its communication cables.



Figure 9 Odroid-VU7+ touch screen interface.

The Odroid-VU7+ receives visual information from the Odroid-XU4 through an HDMI cable and transmits touch screen responses through a USB cable. The USB cable also provides power to the screen. The two cables are contained in a single wrap of braided shielding to reduce clutter and protect the wiring.

2.2.1.4. Lighting

The camera is mounted beneath the front C-arm wheel strut, which means that most of the floor beneath the camera is consistently in shadow. This provides an opportunity to control the lighting conditions through a self-contained lighting system.

A circular ring of white LEDs (Lee's Electronics, Vancouver, BC) was purchased to surround the camera lens. The ring is rated at 12 VDC to illuminate the 30 individual LEDs. A variable resistor was used to determine the amount of resistance required to achieve appropriate brightness beneath the C-arm strut. Satisfactory brightness was achieved at approximately $1.2 \text{ k}\Omega$, which caused the floor to be illuminated without noticeable reflections or loss of contrast. A simplified circuit diagram is shown in Figure 10; note that only 7 of 30 LEDs are represented in the diagram for clarity.



Figure 10 Circuit diagram of the LED ring.

A custom mounting system was designed to connect the LED ring to the camera to ensure equal and consistent lighting throughout tracking. The LED ring was positioned facing upwards to reflect off the underside of the C-arm wheel strut. This configuration prevented reflection in the floor and created a diffused lighting condition. Figure 11 shows the assembled lighting system mounted to the camera.



Figure 11 Lighting system (white) mounted to camera.

The custom lighting mounting system was 3D printed using Polylactic acid (PLA) filament on an Ultimaker 2 3D printer. The system was designed to be mounted directly to the Chameleon camera using four tapped screw holes available on the camera. The mounting system rigidly suspends the LED ring around the camera lens without impeding on the field of view.

The individual parts were designed to snap-fit together to minimize the use of fasteners. There are several general design constraints for the cantilevered lock component of snap-fitting parts to ensure that all bending strain remains in the elastic region. The bottom light holster part is shown in Figure 12 with a snap-fit lock component on its rim.



Figure 12 An annotated drawing of the bottom light holster part with a snap fit component on its rim.

A generic lock component is shown in Figure 13 to illustrate the important design variables.



Figure 13 Design variables for snap-fit hook.

The important design constraints for reusable snap-fit locks are as follows (Bonenberger 2016):

• Beam length L_b and beam thickness T_b are constrained by:

$$5T_b \le L \le 10T_b$$

• Insertion face angle α is bound by:

$$25^o \leq \alpha \leq 30^o$$

• Retention face depth *Y* is related to the wall thickness T_w by:

$$Y \leq T_w$$

• Finally, the retention face angle β is related to the coefficient of friction μ of the material by:

$$\beta \leq \tan^{-1}\left(\frac{1}{\mu}\right)$$

The coefficient of friction for PLA is highly variable depending on the 3D printer properties, layer direction and other factors. Several values have been reported by research groups: 0.385 (Mathurosemontri, et al. 2017), 1.11 (Wu, et al. 2016), and a range from 0.277 to 0.888 (Bajpai, Singh and Madaan 2012). This corresponds to a range of β threshold values from 42.1° to 74.5° for good locking performance. The dimensions of the cantilevered lock constructed according to these constraints are shown in Figure 14.



Figure 14 Relevant design dimensions for snap-fit hook (all dimensions in mm).

Fully dimensioned drawings of all 3D printed lighting system mounting parts are included in Appendix A.

2.2.1.5. Mounting System

The camera must be rigidly attached to the C-arm to accurately track the base. The chosen mounting location on the Siemens Arcadis Orbic Mobile C-arm, beneath the front wheel strut, has two convenient holes that can be used to attach the camera system. Figure 15 shows the location of the mounting holes beneath the rubber cover.



Figure 15 (left) Wheeled C-arm base, (right) wheeled C-arm base with rubber cover removed to expose the two holes used to mount the camera.

A rigid mounting platform was designed to attach to the C-arm strut through the two holes using two bolts. Figure 16 shows the assembled mounting system with the C-arm mounting holes and camera mounting location labelled.



Figure 16 Camera mounting system with labelled C-arm attachment points and camera mounting point.

Figure 17 shows four views of the mounting system with important dimensions and attachment points labelled.



Figure 17 Annotated drawing of camera mounting system with labelled C-arm attachment points, camera mounting point, and adjustable slots.

The strut is not perpendicular to the floor, so some adjustability was built into the mounting platform. The two white components have small slots where they are bolted to the central black platform. These slots enable the angle of the platform to be adjusted to ensure proper alignment.

Space was also included to attach the Odroid-XU4 to the mounting platform, which allows the full system to be concealed beneath the C-arm and provides protection for the processor board. Figure 18 shows the Odroid-XU4 mounted to the platform. Compression springs are employed in tandem with the Odroid-XU4 mounting bolts to alleviate tension from the board and prevent bolt loosening.



Figure 18 Mounting location of the Odroid-XU4 on the bottom of the mounting platform.

A mounting calibration process was established to ensure perpendicularity between the camera and the floor despite the slanted C-arm strut. It is difficult to adjust the system alignment once it is mounted beneath the C-arm due to the limited space and inability to see the parts. Therefore, the mounting platform is first connected to the top of the C-arm strut, then adjusted until a tubular level indicates that it is horizontal. Once the mounting platform adjustments have been made, the system is detached from the top of the C-arm strut and mounted to the proper location beneath the strut. This calibration process is only required when attaching the system to a new C-arm.

Fully dimensioned drawings of all the mounting system parts can be found in Appendix B.

2.2.1.6. Power

Wall-power was chosen as the preferred power method for this research due to the demanding requirements of the processing unit and lighting system. The Odroid-XU4 has a relatively high peak

current at start-up, with a 5 VCD/4 A power supply recommended by the manufacturer (Hardkernal 2017). The LED light ring requires a 12 VDC power supply to illuminate. It is difficult to source a battery system capable of meeting these requirements for long surgical procedures. However, development of a sufficient battery system in the future could be useful to increase the portability of the system and reduce cable clutter.

All wall-powered devices in Canadian hospitals are required to uses hospital grade power cords to minimize risk of damage or electrical shock. The British Columbia Safety Authority provides a list of approved certification marks for electrical products, which was adhered to when purchasing all power components for this research (British Columbia Safety Authority 2016).

A 20-foot NEMA 5-15p hospital grade cord (Canadian Standards Association certified) was purchased to provide power to all the required electrical components. The cord is of comparable length to the C-arm power cord, so overall mobility will not be affected.

The Odroid-VU7+ touch screen interface and the Chameleon camera are both directly powered by the Odroid-XU4 through USB cables. The Odroid-XU4 is powered by a 5 VDC/4 A power adaptor (Underwriters Laboratories certified) and the LED ring is powered by a 12 VDC/1.2 A adaptor (Canadian Standards Association certified). Both AC/DC adaptors are connected to the main cord through a splitter (Underwriters Laboratories certified).

2.2.1.7. System Cost

One of the motivations for using computer vision for C-arm tracking is its relatively low implementation cost. The system is assembled mostly from off the shelf components aside from the mounting platform and lighting mounting, which were 3D printed. Table 8 summarizes the hardware costs associated with the base-tracking system.

Item	Detail	Quantity	Unit Cost (CAD)	Cost (CAD)
Camera	FLIR CM3-U3-13S2M-CS	1	\$295.00	\$295.00
Camera cable	USB 3.0 cable	1	\$10.00	\$10.00
Lens	Fujinon YV2.8x2.8SA-2	1	\$80.00	\$80.00
	Odroid-XU4	1	\$120.00	\$120.00
	eMMC module with Linux	1	\$75.00	\$75.00
User interface	Odroid-VU7+	1	\$130.00	\$130.00
Communication	HDMI cable	1	\$22.95	\$22.95
	USB cable	1	\$5.00	\$5.00
Lighting	LED ring	1	\$16.50	\$16.50
LED power adaptor	12 VDC adaptor	1	\$12.50	\$12.50
Lighting mounting	3D printer material	5.26 m	\$0.80/m	\$4.21
Camera mounting platform	3D printer material	6.89 m	\$0.80/m	\$5.51
	Hospital grade power cord	1	\$22.85	\$22.85
Power	Splitter	1	\$4.75	\$4.75
Total				\$804.27

Table 8	Bill of materials for the base-tracking system
---------	--

This cost total is considerably lower than conventional C-arm tracking systems, which can cost tens of thousands of dollars.

2.2.2. Software Packages

All programs developed for the calibration procedure and base-tracking algorithm were written in the Python 2.7 programming language. An open source computer vision library, OpenCV 3.1, was used to provide a number of well-established algorithms involved in calibration and tracking. Most OpenCV algorithms are implemented in C, which enables fast computation regardless of the language used to write the overall program. The graphical user interface was developed using wxPython, a

Python-specific user interface library. Experimental data was processed using MATLAB R2017a (Mathworks®, Natick, Massachusetts, US).

2.2.3. Camera Calibration

Camera calibration is a necessary step to extract meaningful 3-dimensional measurements from 2dimensional images. Calibration is used to calculate the intrinsic and extrinsic camera parameters, correct for the effects of lens distortions, and to establish the relationship between image coordinates and world coordinates.

The camera intrinsic parameters are used to relate the pixel space to the image plane space. These parameters include the focal length, pixel spacing, and optical centre. The extrinsic parameters relate the image plane space to the world coordinate system. These parameters are related to the position and orientation of the camera in space (Forsyth and Ponce 2012). Both intrinsic and extrinsic parameters are determined through the calibration procedure described in Section 2.2.3.1.

There are several forms of lens distortion that can influence image quality. The most common effect in modern camera lenses is radial distortion, which causes straight lines to appear curved (Szeliski 2011). Figure 19 shows examples of the two common presentations of radial distortion, barrel and pinhole, applied to a checkerboard pattern. Other lens distortion mechanisms include tangential distortion, decentering, and chromatic aberration, although these are reported to have significantly less effect in practice (Zhang 2000). Lens distortions must be corrected before accurate measurements can be made, which is addressed in Section 2.2.3.1. The effect of calibration on the expected measurement error is quantified in Section 3.3 based on experimentation.



Figure 19 Two types of radial lens distortions: (left) barrel distortion and (right) pinhole distortion. The distortion in these images are exaggerated for clarity. In this research, camera motion is calculated relative to the planar floor surface. It is therefore necessary to determine the scale of the floor plane once the camera has been mounted to the C-arm. This information enables meaningful measurements to be calculated in world coordinates. A custom calibration procedure was previously developed by our research group to determine the scale factor, which is presented and expanded upon in Section 2.2.3.2.

2.2.3.1. Undistortion Calibration

To correct for the effects of distortion, we use a closed-form calibration technique developed at Microsoft Research in Washington, USA (Zhang 2000), which is outlined in this section. The technique is based on the conventional pinhole camera model. A 3D point is denoted by $P = \begin{bmatrix} X & Y & Z \end{bmatrix}^T$, and its corresponding 2D image projection is given by $p = \begin{bmatrix} u & v \end{bmatrix}^T$. Augmented forms of these matrices are formed by adding a 1, such that $\tilde{P} = \begin{bmatrix} X & Y & Z \end{bmatrix}^T$ and $\tilde{p} = \begin{bmatrix} u & v \end{bmatrix}^T$. A 3D point is related to its 2D image projection by Equation 1 (Szeliski 2011).

$$\tilde{p} = \lambda K [R_{3x3} \quad t_{3x1}] \tilde{P} \tag{1}$$

An arbitrary scale factor is represented by λ , the selection of which will be addressed in Section 2.2.3.2. The extrinsic parameters (*R*, *t*) represent the rotation and translation relating the world coordinates to the camera coordinates. The intrinsic camera matrix *K*, denoted by Equation 2 (Szeliski 2011), is defined by the optical centre point (u_0 , v_0) and the pixel scale factors along the *u* and *v* axis, α and β . The amount of skew between the two axes is denoted by γ .

$$K = \begin{bmatrix} \alpha & \gamma & u_0 \\ 0 & \beta & v_0 \\ 0 & 0 & 1 \end{bmatrix}$$
(2)

The closed-form solution to calibration draws constraints from the symmetric matrix *B*, which is defined as (Zhang 2000):

$$B = K^{-T}K^{-1} = \begin{bmatrix} \frac{1}{\alpha^2} & -\frac{\gamma}{\alpha^2\beta} & \frac{v_0\gamma - u_0\beta}{\alpha^2\beta} \\ -\frac{\gamma}{\alpha^2\beta} & \frac{\gamma^2}{\alpha^2\beta^2} + \frac{1}{\beta^2} & -\frac{\gamma(v_0\gamma - u_0\beta)}{\alpha^2\beta^2} - \frac{v_0}{\beta^2} \\ \frac{v_0\gamma - u_0\beta}{\alpha^2\beta} & -\frac{\gamma(v_0\gamma - u_0\beta)}{\alpha^2\beta^2} - \frac{v_0}{\beta^2} & \frac{(v_0\gamma - u_0\beta)^2}{\alpha^2\beta^2} + \frac{v_0^2}{\beta^2} + 1 \end{bmatrix} = \begin{bmatrix} B_{11} & B_{12} & B_{13} \\ B_{12} & B_{22} & B_{23} \\ B_{13} & B_{23} & B_{33} \end{bmatrix}.$$
(3)

The symmetric matrix *B* can be represented by the 6-vector $b = \begin{bmatrix} B_{11} & B_{12} & B_{13} & B_{22} & B_{23} & B_{33} \end{bmatrix}$.

Further constraints are provided by considering planar homography. The model plane is assumed to lie on Z = 0 in world coordinates, which modifies the rotation matrix to $R_{2x2} = [r_1 \ r_2]$ (where r_i is a column vector) and the model point $\tilde{P} = [X \ Y \ 1]^T$. The planar homography H is therefore defined, up to a scale factor λ , as the relationship between the model point and its image projection (Zhang 2000):

$$\tilde{p} = H\tilde{P} \tag{4}$$

$$H = [h_1 \quad h_2 \quad h_3] = \lambda K [r_1 \quad r_2 \quad t].$$
(5)

In the above equation, h_k is a column vector representing the kth column of the homography matrix. This provides two constraining equations:

$$h_1^T B h_2 = 0 \tag{6}$$

$$h_1^T B h_1 = h_2^T B h_2 \tag{7}$$

Defining the column vectors of *H* to be $h_i = \begin{bmatrix} h_{i1} & h_{i2} & h_{i3} \end{bmatrix}^T$ gives:

$$h_i^T B h_j = \mathbf{v}_{ij}^T. \tag{8}$$

The vector $\mathbb{v}_{ij} = [h_{i1}h_{j1} \quad h_{i1}h_{j2} + h_{i2}h_{j1} \quad h_{i2}h_{j2} \quad h_{i3}h_{j1} + h_{i1}h_{j3} \quad h_{i3}h_{j2} + h_{i2}h_{j3} \quad h_{i3}h_{j3}]^T$. This allows equations 6 and 7 to be rewritten as:

$$\begin{bmatrix} \mathbf{v}_{12}^T \\ (\mathbf{v}_{11} - \mathbf{v}_{22})^T \end{bmatrix} b = 0$$
(9)

Given *n* calibration images, Equation 9 is stacked *n* times to give:

$$\mathbb{V}b = 0 \tag{10}$$

The solution to 10 is the eigenvector of $\mathbb{V}^T \mathbb{V}$ with the smallest eigenvalue. This provides the solution to *b*, which is used to calculate the values of α , β , γ , u_0 , and v_0 according to Equation 3. These variables fully define the intrinsic camera matrix *K* (Zhang 2000).

The intrinsic and extrinsic camera parameters are initially calculated with an assumption of zero distortion. Once these parameters have been estimated, radial lens distortion is corrected by assuming that the center of the distortion is at the principal point and the skew $\gamma = 0$. The two

equations that relate the observed pixel coordinates (x, y) and the corrected image coordinates (\check{x}, \check{y}) are given by (OpenCV Dev Team 2014):

$$\check{x} = x + x[k_1(x^2 + y^2) + k_2(x^2 + y^2)^2 + k_3(x^2 + y^2)^3]$$
(11)

$$\check{y} = y + y[k_1(x^2 + y^2) + k_2(x^2 + y^2)^2 + k_3(x^2 + y^2)^3].$$
(12)

The radial distortion coefficients k_1 , k_2 and k_3 are solved using the Levenberg-Marquardt algorithm to minimize (Zhang 2000):

$$\sum_{i=1}^{n} \sum_{j=1}^{m} \left\| p - \breve{p} \left(K, k_1, k_2, k_3, R_i, t_i, P_j \right) \right\|^2.$$
(13)

The function $\breve{p}(K, k_1, k_2, k_3, R_i, t_i, P_j)$ is the projection of the world coordinate point P_j onto the image coordinates according to Equation 5 and then distorted by Equations 11 and 12.

The OpenCV implantation of the camera calibration algorithm also calculates parameters for tangential distortion (OpenCV Dev Team 2014). The tangential distortion parameters p_1 and p_2 are related to the image coordinates by:

$$\tilde{x} = x + \left[2p_1xy + p_2\left((x^2 + y^2) + 2x^2\right)\right]$$
(14)

$$\breve{y} = y + \left[2p_1\left((x^2 + y^2) + 2x^2\right) + 2p_2xy\right].$$
(15)

In practice, this calibration scheme is implemented by taking multiple photos of a checkerboard pattern from varying angles. Knowledge of the planar checkerboard pattern is used to solve the above calibration problem. This technique is practical to implement since the pattern can be printed using a standard printer and mounted to a flat object such as a book or table. Figure 20 shows three examples of calibration images taken for this calibration procedure. The effects of radial distortion are clearly visible in the images.



Figure 20 Examples images of the calibration checkerboard pattern.

The calibration procedure described in this section is implemented as a built-in method in OpenCV. Images of checkerboard patterns are provided as inputs to the algorithm, and the outputs are the camera intrinsic matrix *K* and distortion coefficient matrix *D*. These outputs are applied to each video frame throughout the base-tracking algorithm, which is detailed in Section 2.2.4.

The intrinsic camera matrix resulting from this calibration process that was used throughout this research is:

$$K = \begin{bmatrix} \alpha & \gamma & u_0 \\ 0 & \beta & v_0 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 722.473 & 0 & 666.673 \\ 0 & 724.568 & 439.090 \\ 0 & 0 & 1 \end{bmatrix} px$$

The distortion parameters calculated by the calibration procedure are:

 $D = [k_1 \ k_2 \ p_1 \ p_2 \ k_3]^T = [-0.385 \ 0.227 \ 0.00110 \ -0.000491 \ -0.0848]^T \text{ px}$

Note that the tangential coefficients are very small relative to the other distortion parameters. This indicates that the effect of tangential distortion is relatively small compared to radial distortion, as expected.

2.2.3.2. Scale Calibration

A custom target was created for the purposes of scale and perspective calibration, as depicted in Figure 21. The procedure used in this research is a modified version of the procedure during the initial iteration of this project by another member of our lab group (Esfandiari 2014). The target is square with crash dummy symbols (CDS) on each corner for easy detection.



Figure 21 Example image of the scale and perspective calibration target.

The calibration procedure consists of the following steps:

- 1. A bilateral filter (Tomasi and Manduchi 1998) is applied to remove noise while preserving edges. Effective parameter values were determined through trial and error* to be: pixel neighborhood d = 4 and sigma space $\sigma = 75$.
- 2. A circular Hough transform (Hough and Arbor 1962) is used to detect the four circular features. The parameter values were chosen through trial and error**, with resulting values of: accumulator ratio dp = 1.9, minimum circle spacing $D_{min} = 50 px$, and maximum radius $R_{max} = 100 px$.
- 3. Mask frames are created for each detected circle with all pixel values set to 0 outside the detected circle.
- 4. Each generated mask is applied to the original image and saved separately.
- 5. A Canny edge detector (Canny 1986) is applied to each masked frame, followed by a linear Hough transform (Hough and Arbor 1962) to detect the line segments in each circle. Effective Canny edge detection parameters were determined to be: lower threshold $t_1 = 5$, and upper threshold $t_2 = 200$. Linear Hough parameters were set to: distance resolution $\rho = 1 px$, angular resolution $\theta = 1^o$, and minimum vote threshold $t_v = 65$. The parameters for both algorithms were found through trial and error***.
- 6. The intersection points of the line segments in each circle are calculated

7. The mean side length of the identified square is calculated using the detected corner points

*The values were adjusted through trial and error based on visual inspection. Success was defined as values that led to a reduction in image noise without removing the appearance of edges in the target. For the pixel neighborhood, we tested a range of values from 1-10 with a step size of 1, and for the sigma space we tested a range of values from 30 to 100 with a step size of 5.

**An initial guess for the parameters was based on the size of the circles in the target. The values were then adjusted through trial and error until the correct number of circles was detected in multiple images of the target. For the accumulator ratio, we tested a range of values from 1 to 2 with a step size of 0.1. The minimum circle spacing and maximum radius values were chosen to be smaller and larger than the actual measurements found in sample images of the target, respectively.

***The parameters were adjusted through trial and error until the correct line segments were detected in multiple images of the target. For the Canny edge detection lower threshold, a range from 0 to 20 was tested with a step size of 5, and for the upper threshold a range from 150 to 250 was tested with a step size of 5. For the Linear Hough transform, the distance resolution was chosen to be the smallest available value of 1 pixel to provide the highest resolution. A range from 0.5° to 3° with a step size of 0.5° was investigated for the angular resolution. Finally, for the minimum vote threshold, a range of values from 30 to 100 was tested with a step size of 5.

Figure 22 shows an image at each step of the calibration algorithm.



Figure 22 (Starting top left) Original image, bilateral filter applied, Hough circles detected, mask applied, Canny edges detected, Hough lines detected, intersection points on original image.

The side length of the square target was measured using digital calipers, which allows a scale factor to be calculated in units of millimeters per pixel. This scale factor directly relates transformations measured in the image coordinate space to world coordinate space. It is equal to λ in section 2.2.3.1. This calibration procedure is an attempt to provide robust feature extraction that is insensitive to changes in target size, lighting conditions, and image noise.

2.2.4. Base-Tracking Algorithm Design

The following sections will detail the frame-to-frame tracking algorithm, absolute loop closure, and the user interface. Details on several important existing concepts are also provided to enhance understanding of the overall algorithm design.

2.2.4.1. Important Concepts

The developed base-tracking algorithm leverages several existing computer vision concepts. The following sections provide an overview of the ORB feature detector and descriptor, the Brute Force matching algorithm, the ratio test, and geometric image transformations to enable a better understanding of the overall algorithm, and to illustrate the novel aspects of the base-tracking system.

2.2.4.1.1. Oriented FAST and Rotated BRIEF Feature Detector and Descriptor

Oriented FAST and Rotated BRIEF (ORB) is a binary feature description algorithm that uses modified FAST keypoint detection in combination with a modified BRIEF descriptor (Rublee, et al. 2011). ORB was chosen for this application due to its speed and accuracy relative to other feature detectors. It was specifically designed with low power, real-time applications in mind and is resistant to rotation and changes in lighting conditions.

FAST is a well known keypoint detection algorithm that is designed to be fast enough for use in realtime applications. The FAST detector examines a ring of 16 pixels surrounding a candidate pixel and compares their intensities. The pixel test *F*, shown in Equation 16, calculates the intensities of each ring pixel I_p^i and compares them to the intensity of the candidate pixel I_c plus or minus a threshold value *t*.

$$F(I_p;t;i) = \sum_{i=0}^{16} \begin{cases} 1, \ I_p^i < I_c - t \\ 1, \ I_p^i > I_c + t \\ 0, \ otherwise \end{cases}$$
(16)

The center pixel is identified as a keypoint if the test sums to 9 or higher (Rosten and Drummond 2006). The ORB detector augments the original FAST algorithm by adding a scale pyramid scheme, a measure of corner strength, and a measure of orientation. The pyramid scheme allows features to be detected at multiple scales by running the algorithm on multiple scaled versions of each image. A Harris corner score is used to determine the strength of a detected keypoint and reject points that lie on edges. The Harris score identifies points where the image gradient changes in more than one direction, which increases the robustness of the detected keypoints (Harris and Stephens 1988). Finally, the intensity centroid of a small window surrounding the keypoint is calculated, and the vector between this centroid and the center of the keypoint as a measure of orientation. This modified FAST algorithm is called Oriented FAST (oFAST) (Rublee, et al. 2011).

BRIEF (Calonder, et al. 2010) is a keypoint descriptor algorithm that was developed as a less computationally intensive alternative to SIFT (Lowe, 2004) and its variants. BRIEF uses a binary vector to describe an image patch surrounding a detected keypoint rather than a difference-of-Gaussian histogram (Calonder, et al. 2010). The descriptor performs a binary test τ on a small image patch p surrounding the keypoint that is defined as:

$$\tau(p; x; y) = \begin{cases} 1, \ p(x) < p(y) \\ 0, \ otherwise \end{cases}$$
(17)

where p(x) is the intensity of the pixel at the point x and p(y) is the intensity of the pixel at point y. The set of (x,y) pairs are selected in a random pattern on a small image patch that has been smoothed using a Gaussian filter (Calonder, et al. 2010). The result is a 256-bit binary vector that describes the area surrounding a keypoint.

Binary descriptors allow for fast computation speeds and efficient storage. BRIEF is robust to changes in lighting, image blur, and perspective changes; however, it performs poorly under planar rotation (Rublee, et al. 2011). ORB makes use of a modified BRIEF descriptor called rBRIEF, which was developed using the orientation measure from oFAST and a principal component analysis learning step. The rBRIEF descriptor is robust to planar rotation and is very computationally efficient to compute and match (Rublee, et al. 2011).

ORB is the combination of the oFAST detector and rBRIEF descriptor, which means that it is robust to changes in lighting, perspective, planar rotation, and image blur. These characteristics, along with the fact that it is computationally inexpensive, are the primary motivations for using ORB for mobile applications. Experimental results for multiple detector and descriptor algorithms are presented in Section 3.1 to further support the use of ORB in this application.

2.2.4.1.2. Brute Force Matching Algorithm

The Brute Force descriptor matching algorithm is a straightforward method of finding corresponding keypoints. The Hamming distance, which calculates the number of positions where one binary string differs from another, is used to compare the binary ORB descriptors from two images. The Brute Force algorithm performs an exhaustive search to find the smallest Hamming distance, and therefore best match, for each keypoint. In this application, the smallest two Hamming distances were recorded to perform a ratio test to filter the results.

2.2.4.1.3. Ratio Test

The ratio test is a fast method of filtering keypoint matches proposed by Lowe (Lowe 2004) to remove features that do not have a good match in the reference image. The ratio between the smallest match distance d_1 and the second smallest distance d_2 for each keypoint is calculated:

$$Ratio = d_1/d_2. \tag{18}$$

A small ratio indicates that the closest match is much stronger than the second closest match, which tends to indicate that a correct match has been found. Rejecting all matches with a ratio greater than 0.8 has been shown to eliminate 90% of false matches while incorrectly eliminating only 5% of real matches (Lowe 2004).

2.2.4.1.4. Geometric Image Transformations

Geometric transformations can be applied to images in the form of matrix operations to translate, rotate, and distort image domains. Geometric image transforms are defined by the number of degrees of freedom. A Homography, also called a projective transform, allows up to 8 degrees of freedom as shown in Equation 19.

$$\begin{bmatrix} h_1 & h_2 & h_3 \\ h_4 & h_5 & h_6 \\ h_7 & h_8 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} u \\ v \\ w \end{bmatrix}$$
(19)

These 8 degrees of freedom can be decomposed into two translations, planar rotation, scale, aspect ratio, shear, and two projective distortions (i.e. out of plane tilting) (Szeliski 2011). Note that the

augmented 2-dimensional image coordinates $[x \ y \ 1]^T$ of the original image can be warped into full, 3-dimensional coordinates $[u \ v \ w]^T$.

In this research, the camera is rigidly fixed to the C-arm so that it maintains a constant distance and angle from the planar floor surface. These constraints indicate that projective distortions, scale, aspect ratio, and shear should not change from one video frame to the next. Limiting the degrees of freedom to translation and rotation along a constant plane z = 0 defines a rigid geometric transform:

$$\begin{bmatrix} R_{1,1} & R_{1,2} & T_1 \\ R_{2,1} & R_{2,2} & T_2 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix}.$$
 (20)

Note that in this case, the original 2-dimensional image coordinates $\begin{bmatrix} x & y & 1 \end{bmatrix}^T$ can only be warped to new 2-dimensional coordinates $\begin{bmatrix} x' & y' & 1 \end{bmatrix}^T$, since the distortion is constrained to a single plane. It is also important that the rigid transformation matrix can be easily decomposed into rotation and translation, which is necessary for extracting physical camera movements from the matrix. A rigid transform is often represented in the form shown in Equation 21.

$$Rigid = [R|t] = \begin{bmatrix} R_{1,1} & R_{1,2} & T_1 \\ R_{2,1} & R_{2,2} & T_2 \end{bmatrix}$$
(21)

A rigid transform is calculated for each set of consecutive video frames to determine the incremental translation and rotation. The incorporation of this step into the overall algorithm is further detailed in Section 2.2.4.2.

2.2.4.2. Relative Motion Tracking

Frame-to-frame relative tracking is the primary tracking method for this system. The algorithm calculates the incremental change in position and orientation that occurs between each video frame to continually update the C-arm pose in real-time.

The global coordinate system is always established with its origin at the location where C-arm tracking was initiated. The x-axis is directed horizontally to the right from the perspective of a user standing at the back of the C-arm (where the controls are located), and the y-axis is directed away from a user in the same position. A yaw rotation that is counter-clockwise when viewed from above the C-arm is considered positive.

The relative motion tracking algorithm consists of the following steps:

- 1. A pixel mapping function is calculated based on the calibration parameters
- 2. The mapping function is applied to the first video frame to remove distortion
- 3. Initial feature points are calculated using the ORB detection and description algorithm (Rublee, et al. 2011)
- 4. Current feature points are saved as reference points
- 5. Calibration mapping function is applied to the next video frame
- 6. Feature points are calculated using the ORB detection and description algorithm (Rublee, et al. 2011)
- 7. A set of matching features is determined using the Brute Force matching algorithm
- 8. The Ratio Test is applied to filter the best matches (Lowe 2004)
- 9. A rigid transform is calculated that best fits the transformation between the reference points and the current points
- 10. Displacement and rotation measurements are extracted from the rigid transformation matrix
- 11. Repeat steps 3 through 10

An example of 950 ORB features detected in an image of the floor in the BioEngineering Lab in the Centre for Hip Health and Mobility is shown in Figure 23.



Figure 23 ORB features in an image of an operating room floor.

The calibration mapping function that is generated in the first iteration creates a transformation map for each pixel based on the calibration parameters. Applying this mapping scheme is significantly faster than using the OpenCV undistort function on each loop, with identical results. The undistort function calculates the transformation map each time, which is unnecessary since the calibration parameters do not change. The algorithm efficiency was therefore improved by separating the calculation and application of the transformation mapping.

The incremental displacement and rotation measurements are calculated from the rigid transformation matrix between each video frame. These incremental measurements are accumulated throughout tracking to continuously determine the real-time pose of the C-arm. The rigid transformation matrix has the following form:

$$Rigid = \begin{bmatrix} R_{1,1} & R_{1,2} & T_1 \\ R_{2,1} & R_{2,2} & T_2 \end{bmatrix} = \begin{bmatrix} \cos(\Delta yaw) & -\sin(\Delta yaw) & T_x \\ \sin(\Delta yaw) & \cos(\Delta yaw) & T_y \end{bmatrix},$$
 (22)

where Δyaw is the incremental yaw between the previous frame and the current frame. T_x and T_y are the incremental displacements between the previous frame and the current frame in terms of a temporary local coordinate system established at the previous frame.

The incremental yaw measurement is extracted from the transformation matrix according to Equation 23.

$$\Delta yaw = \tan^{-1} \left(\frac{R_{2,1}}{R_{1,1}} \right) = \tan^{-1} \left(\frac{\sin(\Delta yaw)}{\cos(\Delta yaw)} \right)$$
(23)

The total yaw between the starting point and the current position is calculated by simply accumulating all the incremental yaw measurements, as shown in Equation 24. This summation is carried out continuously to maintain an updated orientation measurement θ throughout tracking.

$$\theta = \sum \Delta yaw \tag{24}$$

The incremental position changes in the rigid transformation matrix, Equation 22, must be transformed to the global coordinate system. The calibration scale factor λ and the global yaw measurement θ are used in Equation 25 to convert to global coordinates.

$$\begin{bmatrix} \Delta X \\ \Delta Y \end{bmatrix} = \lambda \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} T_x \\ T_y \end{bmatrix}$$
(26)

In non-matrix form:

$$\Delta X = \lambda \left(T_{\rm x} \cos \theta - T_{\rm y} \sin \theta \right) \tag{27}$$

$$\Delta Y = \lambda (T_x \sin \theta + T_y \cos \theta). \tag{28}$$

The total displacement from starting position, in global coordinates, is simply the summation of incremental position changes. These summations, shown in Equations 29 and 30, are carried out on each loop of the algorithm to maintain a continuous measurement of the current C-arm position.

$$X = \sum \Delta X \tag{29}$$

$$Y = \sum \Delta Y \tag{30}$$

The pose of the C-arm base in global coordinates (X, Y, θ) is therefore continuously determined in real-time through the accumulation of incremental measurements. The relative tracking algorithm operates at approximately 8 fps on the Odroid-XU4.

2.2.4.3. Absolute Position Recovery

There are two primary motivations for implementing an absolute position recovery method in the tracking algorithm. The first is that it allows users to return to previously obtained C-arm positions with increased accuracy, and the second is the ability to reset the accumulation of errors that occurs during relative tracking.

Absolute position recovery is enabled using saved points. The user can save the current C-arm position at any time through the user interface, which will be further detailed in Section 2.2.4.4. The user can choose any saved point that they wish to reacquire, and the system will provide guiding cues when the C-arm is near the correct location. The absolute position recovery aspect of the tracking algorithm has been programmed to automatically initiate when the user is within 3 cm from a previously saved point that they are attempting to reacquire. A radius of 3 cm is approximately half of the camera field of view, and was chosen to provide sufficient overlap between the current image and the saved point.

When a user saves a position, the global translation and orientation information is stored along with the set of reference keypoints and descriptors associated with that position. The global reference frame is defined using the initial C-arm position as the origin. When the C-arm is near the selected save point, the algorithm uses these stored feature points as reference points to determine the displacement and rotation motion required to estimate the current position relative to the saved position and thereby update the system's estimate of its current location. This process enables the user to reposition the C-arm with an increased level of accuracy, and resets any measurement errors that may have accumulated during the relative tracking process. The absolute tracking mode operates continuously while the C-arm remains within the specified 3 cm radius from the saved point. If the C-arm is moved away from the saved point, the algorithm will automatically switch back to the relative tracking method. Both methods operate at approximately the same frame rate, so there is no loss in performance incurred through the addition of the loop closure step.

The process of calculating position measurements in the absolute tracking mode is slightly different than in relative tracking. The global position and orientation of the target point are known, and the

rigid transformation matrix provides a transform between the current position and the target position. Therefore, the current position can be determined in global coordinates directly, rather than through the accumulation of incremental measurements.

The rigid transformation matrix in absolute tracking is defined in Equation 31. The rotational aspect is described by yaw_d , which represents the angular difference between the current frame A and the target frame. T_x and T_y are the translations between frame A and the target frame, defined in a temporary coordinate system established at frame A.

$$Rigid = \begin{bmatrix} R_{1,1} & R_{1,2} & T_1 \\ R_{2,1} & R_{2,2} & T_2 \end{bmatrix} = \begin{bmatrix} \cos(yaw_d) & -\sin(yaw_d) & T_x \\ \sin(yaw_d) & \cos(yaw_d) & T_y \end{bmatrix}$$
(31)

The angle between the current position and target position is calculated by Equation 32.

$$yaw_{d} = \tan^{-1} {\binom{R_{2,1}}{R_{1,1}}} = \tan^{-1} {\binom{\sin(yaw_{d})}{\cos(yaw_{d})}}$$
(32)

The global orientation θ can be directly calculated given the angular difference between the current position and the target position, as well as the known global angle associated with the target position yaw_{target} :

$$\theta = yaw_{target} + yaw_d. \tag{33}$$

The displacement between current position and target position in global coordinates X_d and Y_d can be calculated using the local displacements from the rigid transform in Equation 31. The calibration scale factor λ and the global orientation measurement θ are used to convert to global coordinates, as described in Equations 34 and 35.

$$X_d = \lambda \left(T_{\rm x} \cos \theta - T_{\rm y} \sin \theta \right) \tag{34}$$

$$Y_d = \lambda \left(T_x \sin \theta + T_y \cos \theta \right) \tag{35}$$

Finally, the global position can be directly calculated from the target position (X_{target}, Y_{target}) and the displacement between the current position and target position:

$$X = X_{target} + X_d \tag{36}$$

$$Y = Y_{target} + Y_d. \tag{37}$$

In this manner, the global position and orientation of the C-arm base (X, Y, θ) are determined based on the location of the saved point. The previous position and orientation of the C-arm at the saved point $(X_{target}, Y_{target}, yaw_{target})$ can also be reacquired with a high level of accuracy. The absolute tracking algorithm operates at approximately the same frame rate (8 fps) as the relative tracking algorithm on the Odroid-XU4.

2.2.4.4. Graphical User Interface

A simple and functional graphical user interface (GUI) was designed to allow the user to interact with the system through an Odroid-VU7+ touch screen interface. The interface was constructed using wxPython, a common GUI library for Python. The button sizes and positions were designed to accommodate touch screen interactions based on guidelines developed by a group at the University of Maryland (Anthony, et al. 2014).

The GUI has two main tabs, one for tracking and one for calibration. The tracking tab, shown in Figure 24, provides the user with the option to initiate tracking using the 'Start' button.

OPTIX C-Arm Tracking	008
Tracking Calibration	
Start	Save Current Location

Ready to Track

Figure 24 Initial screen when OPTIX program is started.

Once the 'Start' button has been pressed, the interface changes to provide the user with the option to end tracking using the 'Stop' button, or save the current position using the 'Save Current Location' button. When a position is saved it appears as a circular target on the screen with crosshairs to aid alignment and a line indicating the forward position of the C-arm. The user may save multiple locations around the operating room by pressing the 'Save Current Location' button. The current C-arm position is represented as a square target in the centre of the screen, with crosshairs used to align the C-arm with a previously saved position. Figure 25 shows the GUI with two points saved, one of which has been activated for reacquisition.



Figure 25 OPTIX GUI showing two saved points; the red point is currently activated for reacquisition.

The user can touch a previously saved point that they would like to repeat. This activates the return function in the algorithm and allows for absolute loop closure when the C-arm is close to the saved position. When a user touches a saved point it will turn red, indicating that it is active for return guidance. Once the C-arm is within 5 mm of displacement and 3.1° of rotation from the saved point, the circle will turn green. Figure 26 shows a dark grey region representing the translation target area with a radius of 5 mm and a light grey region representing the $\pm 3.1^{\circ}$ target angle bounds.



Figure 26 OPTIX GUI showing two saved points, the red point is currently activated for reacquisition. The labelled dark grey circle and light grey triangle represent the 5 mm and 3.1° tolerances respectively.

These limits are based on the tracking system requirements established in Section 2.1. Figure 27 shows a successfully reacquired position.



Figure 27 OPTIX GUI showing successful reacquisition of a previously saved position.

The second tab in the GUI is the calibration tab, which is only used during the initial system setup. The calibration tab is shown in Figure 28.

• OPTIX C-Arm Tracking	(908
Tracking Calibration		
	Run Camera Calibration	
	Run Scale Calibration	
Ø Default	⊖ Manual	
Ready to Calibrate		

Figure 28 OPTIX GUI calibration tab.
The user has the option to initiate distortion calibration by pressing 'Run Camera Calibration', or to initiate scale calibration by pressing 'Run Scale Calibration'. The user also has the option to manually select the calibration files or use the default files. The calibration images are acquired separately from the OPTIX GUI, and this tab is used to initiate the required functions to read the selected calibration images and determine the resulting parameters. The calculated parameters are saved for use in the tracking algorithm.

2.2.4.5. Program Architecture

The overall architecture of the tracking algorithm was designed to take advantage of the multi-core CPU on the Odroid-XU4 through parallel processing. A colour-coded flow chart representation of the overall program is depicted in Figure 29. The program is split into three processes that run simultaneously, which increases the speed and allows for a responsive user interface. The three separated aspects are: (1) the user interface, (2) the camera interface, and (3) the tracking algorithm. The flowchart contains all three parallel processes with labels and coloured backgrounds. Square blocks represent actions, rhomboid blocks represent decisions, and circles indicate the initiation of a parallel process. Independent loops exist within each of the three processes, with connections established through message queues. The queue messages are color matched to show the six possible messages that can be transferred through queue channels across processes.



Figure 29 Flowchart of the base-tracking algorithm with color-coded processes and queues. Purple background is the user interface process, green background is the camera process, and orange background is the tracking process. Queues of matching color are connected.

The user interface process controls the GUI for the user to interact with, and is initially the only running process when the program starts. The camera interface process is initiated from within the user interface process when the 'Start' button is pressed. Placing the GUI in its own process enables a responsive interface that updates in real-time. The processes communicate with one another through queue messages. Queues are a method of relaying messages between parallel processes, which enables the real-time C-arm position information to be displayed on the GUI. The user interface process sends information about button presses to the other two processes, and receives messages about the C-arm position and saved points. The user interface process is represented on the left side of Figure 29.

The camera interface process establishes the connection between the Odroid-XU4 and the Chameleon camera when tracking is initiated using the 'Start' button. The tracking process is initiated from within the camera process once the camera connection has been established. The camera process reads images from the camera and applies the calibration coefficients to each frame before sending them to the tracking process. The camera process sends a new image to the tracking process any time the queue between them is empty, indicating that the tracking algorithm has completed a loop and is now processing a new image. This strategy ensures that there is no build up of images, which would create a measurement lag in the system unless the camera frame rate and algorithm rate were exactly synchronized. Applying the calibration distortion parameters to each image is one of the most computationally expensive aspects of the tracking system, so separating this process from the main tracking algorithm increases the overall loop speed. The camera interface process constitutes the center portion of the flowchart in Figure 29.

Finally, the third process contains the relative and absolute tracking algorithms. This process is initiated from within the camera process and receives calibrated images through the queue. The tracking process also communicates with the user interface process through the queue by sending the current C-arm position and save point information. The bulk of the calculations are carried out in the tracking process, including feature point detection and matching, transform estimation, and pose calculation. The tracking process constitutes the right side of Figure 29.

Three separate, independent queue connections are established to streamline communications between processes. The image queue is dedicated solely to transmitting images from the camera process to the tracking process. A second queue is dedicated to transferring messages from the user interface process to the tracking process. The user interface sends messages through this queue to save the current position, initiate reacquisition of a previous position, or to stop tracking. Finally, the third queue channel is responsible for communicating messages from the tracking process to the user interface process. This queue is used to transmit the current C-arm position on each loop and when a position has been saved.

2.3. Experimental Evaluation Methods

Several experimental procedures were carried out to determine the accuracy of the overall tracking system. A short survey was conducted to obtain photographs of operating room floors from multiple operating rooms. A setup was assembled for use throughout algorithm development, which was used to carry out preliminary accuracy tests to guide the selection of various algorithms included in the final tracking system. An experimental evaluation was also developed to measure the accuracy of the calibration procedure. Finally, the primary accuracy evaluation for the OPTIX system was carried out using the C-arm in a simulated operating room.

2.3.1. Operating Room Floor Survey

The base-tracking algorithm measures C-arm motion relative to the operating room floor based on the appearance of visual texture. It is therefore necessary to ensure that there is a satisfactory level of visual texture present in a typical operating room floor.

Photographs were voluntarily collected by various medical students and surgical staff from multiple hospitals. The pictures were taken using smart phone cameras positioned roughly parallel to the floor at an approximate height 10 cm. No patients, hospital personnel, or any other details from the operating room were included in any photos. The results of the survey are shown in Section 3.1.

2.3.2. Development Testing

A desktop testing rig was constructed for use throughout the development of the tracking algorithm. This setup allowed rapid, continuous performance testing as small adjustments were implemented in the tracking system.

The testing rig consists of a camera slider (StudioFX) with a pivoting ball mount (DMKfoto) capable of directing the camera towards the floor. A section of vinyl flooring with a visual texture similar to that found in operating rooms was placed beneath the camera slider. The experimental setup is shown in Figure 30.



Figure 30 Testing setup used during development, including a camera slider, camera, and vinyl flooring.

This experimental setup ensured that the camera could be moved smoothly at a consistent distance from the flooring to allow for rapid preliminary algorithm testing. The camera slider track provides 1113 mm of travel distance in a single direction. The total distance travelled can be increased by traversing back and forth across the slider multiple times. This preliminary testing was used to guide several design choices, outlined in Sections 2.3.2.1 and 2.3.2.2, when developing the algorithm.

2.3.2.1. Comparing Tracking Methods

The development testing rig was used to compare the accuracy of optical flow based tracking algorithms with descriptor matching. These measurements were primarily carried out using straight line trajectories, with the knowledge that more complex paths would be more easily evaluated with the camera mounted on the C-arm in later testing.

The camera was initially positioned firmly against one end of the camera slider to ensure a consistent return position and physically measurable amount of displacement. The camera was then propelled along the slider by hand while the algorithm tracked the motion. The camera was moved through one or more laps of the camera slider and eventually returned to its initial position.

The final displacement and rotation measured by the tracking algorithm were recorded, as well as the total distance travelled. The absolute repositioning error was determined directly, since the true final location was equal to the initial location. The relative repositioning error was calculated by dividing the absolute error by the total distance travelled. The ORB descriptor and detector algorithm (Rublee, et al. 2011) was used to represent descriptor matching, and Lucas Kanade (Lucas and Kanade 1981) was used as the optical flow technique. The value of adding an absolute loop closure step was also explored in this preliminary accuracy study. The absolute loop closure step follows the algorithm outlined in Section 2.2.4.3. The results for these tests can be found in Section 3.2.1.

2.3.2.2. Comparing Descriptor Matching Algorithms

There are multiple keypoint detector and descriptor algorithms available in the OpenCV library. It was important to evaluate the performance of the available algorithms to determine which was the most appropriate for this application.

AKAZE (Alcantarilla, Nuevo and Bartoli 2013), BRISK (Leutenegger, Chli and Siegwart 2011), and ORB (Rublee, et al. 2011), are non-patented, open source feature detector and descriptor algorithms in the OpenCV library that are aimed at real-time applications. FLANN (Muja and Lowe 2009) and Brute Force are the two most commonly used matching algorithms available in OpenCV. The three detector and descriptor algorithms were tested with each of the two matching algorithms for speed and number of successful matches. Two images of the visually textured floor were taken with a large amount of overlap (i.e. a small displacement) for testing purposes.

A testing program was written that recorded the time required to detect, describe and match keypoints in the two images. The program iterates through the detection and matching sequence multiple times and calculates the average loop speed, along with the number of successful matches. The timeit Python library was used to measure the time required for each loop.

The two matching algorithms were tested with each of the three detectors and descriptors. The loop speed was averaged over 1000 iterations for each iteration. Once a matching algorithm was selected, the loop times for each of the detector and descriptor algorithms were compared. The results are summarized in Section 3.2.3.

2.3.3. Calibration Evaluation

The OpenCV implementation of the flexible calibration algorithm (Zhang 2000) calculates and reports an internal root mean square error (RMSE) for each calibration dataset. It is common practice in computer vision applications to iteratively calibrate a camera until an RMSE of less than 0.5 pixels

is reported. When we performed this standard calibration procedure prior to conducting the evaluation study, we found an internal RMSE value of 0.48 pixels for the calibration image set.

However, since this RMSE is calculated based on a randomly selected image from within the calibration dataset, rather than an external test image, it is important to also externally verify the success of the calibration parameters with images that are not included in the training dataset. To externally validate the calibration parameters, we evaluated the straightness of a linear feature. Ten images were taken of a straight line at multiple orientations, and the calibration parameters were applied to remove lens distortion. Figure 31 shows an example of a linear feature before and after the distortion parameters are applied.



Figure 31 Example images of the line used for calibration evaluation.

The straightness of each feature was determined by calculating a line of best fit for both the original and corrected images. The pixel locations (x, y) for a sample of n points were recorded for each image. The RMSE was then calculated using the best-fit line values \hat{y}_i and the actual pixel locations y_i according to Equation 38.

$$RMSE = \sqrt{\frac{\sum_{i=1}^{n} (\hat{y}_i - y_i)^2}{n}}$$
(38)

A higher RMSE value corresponds to more distortion in the image. The results of this evaluation can be found in Section 3.3.

2.3.4. C-arm Testing in a Simulated Operating Room

A Siemens Arcadis Orbic Mobile C-arm (Siemens, Erlangen, Germany), located in the Bioengineering Laboratory (BioEng Lab) at the Centre for Hip Health and Mobility (CHHM), was used to evaluate the accuracy of the OPTIX system. The simulated surgical suite in the BioEng Lab (see Figure 32) enabled test conditions resembling a real operating room. An Optotrak Certus (Northern Digital Inc., Waterloo, ON, Canada) optical tracking system was used to measure the position of the C-arm and compare with OPTIX measurements.



Figure 32 BioEng Lab setup for C-arm testing.

The OPTIX camera and processing unit were rigidly attached to the front wheel strut of the C-arm with the camera facing the floor. The LED ring was illuminated to remove the effects of shadows and specular reflections. The GUI described in Section 2.2.4.4 was used throughout the C-arm evaluation procedures to save positions and receive visual guidance from OPTIX. The current translation and rotation measurements calculated by OPTIX were saved on each loop in comma separated value (CSV) format for later evaluation. Separate testing schemes were designed to evaluate the accuracies of the relative tracking algorithm and the loop closure algorithm.

2.3.4.1. Relative Tracking Algorithm Evaluation Methods

The open-loop relative tracking tests were separated into translation and rotation tests. The purpose of these studies was to assess the relative tracking error as a function of the cumulative distance or rotation. For the relative tracking translation test, depicted in Figure 33, the C-arm was propelled through arbitrary distances up to approximately 3.5 m, which is approximately twice the length of a conventional operating room table. The C-arm was stopped at small intervals for 10 seconds to record measurements with Optotrak. The position recorded by OPTIX was stored continuously. The 10 second pauses allowed the Optotrak and OPTIX data to be synchronized during processing and enabled a mean to be calculated at each position from multiple measurements, which helps account for some noise. Multiple repetitions of this test were repeated in varying directions.



Figure 33 An example of C-arm motion for open-loop translation testing.

The relative tracking rotation test, shown in Figure 34, involved rotating the C-arm through angles up to approximately 70°. The largest rotations we observed in orthopaedic procedures was approximately 30° in a single direction, which could lead to an expected maximum back-and-forth rotation of approximately 60°. The C-arm was stopped at small intervals to record the position in Optotrak and allow synchronization with OPTIX data during processing. The Optotrak measurements and OPTIX measurements were compared to determine the tracking error for both relative tracking studies.



Figure 34 Diagram of C-arm motion for open-loop rotation testing.

The results for the relative tracking algorithm evaluation are presented in Sections 3.4.2 and 3.4.3.

2.3.4.2. Loop Closure Algorithm Evaluation Methods

The loop closure algorithm was evaluated through a series of repositioning tasks. Four predefined patterns of C-arm motion were established to evaluate the repositioning accuracy of the OPTIX system. When reacquiring a saved position, the C-arm was maneuvered until the OPTIX GUI target turned green. This testing method was designed to evaluate the ability of the OPTIX system to track C-arm positioning near previously stored locations after moving around the operating room. The motion schemes were chosen to represent C-arm movements that occur in the operating room while ensuring that a wide range of possible motions was evaluated. Common C-arm motions were determined through observation of orthopaedic surgeries and consultation with our clinical collaborator, Dr. Pierre Guy.

The first motion scheme involved translating the C-arm side-to-side in a straight line along the operating table, as shown in Figure 35. This scheme evaluated OPTIX accuracy in a single translation dimension while closely resembling C-arm translation observed during lower limb or spine surgery. Arbitrary initial and final positions were saved in OPTIX and the C-arm was moved between the two points multiple times. The total amount of translation was varied arbitrarily up to the full length of the patient table, ensuring that the full range of potential motions was represented.



Figure 35 Diagram of side-to-side C-arm motion path for testing.

The second motion pattern also consisted of a single translation dimension. The C-arm was moved fully away from the table and then repositioned multiple times, as depicted in Figure 36. This scenario is commonly observed in pelvis fixation and spinal surgery to provide the surgeon with more room to operate. The initial and final positions were saved in OPTIX before repetition, and the total distance away from the table was varied arbitrarily.



Figure 36 Diagram of in-and-out C-arm motion for testing.

The third motion pattern involves both translation and rotation to test both aspects of OPTIX. Two positions along the patient table are saved in OPTIX with the C-arm placed at different angles to the table, shown in Figure 37. This maneuver, named oblique motion, is representative of motions observed in some lower limb and spinal surgeries.



Figure 37 Diagram of oblique C-arm motion path for testing.

Finally, a series of four arbitrary save points was tested to evaluate the effect of more complex motions. The C-arm was moved to four positions, represented in Figure 38, and a save point was recorded in OPTIX at each position. The C-arm was maneuvered to the four points multiple times to evaluate the ability for OPTIX to measure C-arm position near multiple save points in one session.





An Optotrak Certus (Northern Digital Inc., Waterloo, ON, Canada) optical tracking system was used as the gold standard against which the OPTIX measurements were compared. The C-arm body was equipped with three active infrared markers to track its position in space. Each marker emits an infrared signal with a unique frequency that is detected by the Optotrak optical sensor. The Optotrak shutter is synchronized with the infrared signals and it transmits the position data for each marker to the Certus System Control Unit where it is analyzed by the Optotrak interface software. Two Optotrak software packages, 1st Principles and 3D Architect, were used for marker calibration and measurement for this study. The Optotrak has a resolution of 0.01 mm and an error of ± 0.1 mm when the infrared markers are within the optimal ranges defined in Figure 39 (Northern Digital Inc. 2015).





The infrared markers were taped rigidly to the C-arm body approximately 20 cm apart, as shown in Figure 40. Based on the expected Optotrak measurement error, this configuration corresponds to an estimated angular error of approximately $\tan^{-1}(0.1 \text{ mm}/_{200 \text{ mm}}) = 0.03^{\circ}$.



Figure 40 The positions of the three Optotrak markers on the C-arm in the BioEng Lab.

Optotrak provides 3-dimensional position information for each marker in its own local coordinate system. The following section details how these data were interpreted for comparison with the planar translation and rotation measurements from OPTIX.

2.3.4.3. Interpretation of Optotrak Measurements

Measurements made by Optotrak are oriented in a local coordinate system in the form of 3dimensional position data $(x_i, y_i, z_i)_j$ for each infrared marker. However, OPTIX provides translation and rotation measurements relative to the C-arm starting position in a single plane, assumed to be parallel to the floor. The Optotrak measurements must therefore be projected onto the floor plane and expressed in terms of translation and rotation relative to the initial position to compare with the OPTIX data. This was accomplished by measuring multiple points on the floor at the beginning of each test session to establish the floor plane. All measured C-arm points for each test session could then be projected onto the corresponding floor plane for direct comparison with OPTIX.

A digitization probe (Northern Digital Inc., Waterloo, ON, Canada), depicted in Figure 41, was used to measure points on the floor. The digitization probe has four active infrared markers, and it is calibrated using the 3D Architect software. A pivot test was carried out to establish the relationship

between each of the infrared markers and the tip of the probe. This process produces a calibration file that is uploaded into the 1st Principles software before conducting measurements with the probe.



Figure 41 Optotrak probe used to digitize the floor plane.

The probe tip can be used to digitize a point by placing the probe tip in direct contact. In this case, the probe tip was placed on the floor in 6 or more locations to collect the points that were used to define the floor plane. The collected points were spaced over an area of approximately 2 m x 2 m, which corresponds to a maximum estimated angular error in the floor plane of approximately 0.003° .

The marker positions were collected at 100 Hz using 1st Principles and processed in Matlab. Each probe location was averaged over the full 10 seconds of collection time. A planar surface was then fit to the set of points using Matlab's built-in pcfitplane function with a maximum tolerance of 0.1 mm to correspond with the Optotrak measurement error. This process was repeated for each new test session to account for the position and orientation of the Optotrak.

All subsequent Optotrak measurements of C-arm motion were projected on to the established floor plane. This effectively constrained the Optotrak data to planar motion. An infrared marker point, defined in Optotrak coordinates as $P_{opt} = [x_{opt} \quad y_{opt} \quad z_{opt}]^T$, was projected onto a floor plane defined by a normal vector $N_{floor} = [u \quad v \quad w]$ and a point on the plane $Q_{floor} = [x_{floor} \quad y_{floor} \quad z_{floor}]^T$ according to Equation 39. I_{3x3} is the 3x3 identity matrix and P_{proj} is the resulting projected point.

$$P_{proj} = P_{opt} \left(I_{3x3} - N_{floor}^T N_{floor} \right) + Q_{floor} \left(N_{floor}^T N_{floor} \right)$$
(39)

The next step is to use the infrared markers to establish a local coordinate system to determine the translation and rotation at each C-arm location relative to the initial position. The relationship between the three infrared markers was constant since they were rigidly attached to the C-arm. A local coordinate system was defined at each position using the vectors between the infrared markers. The markers were arbitrarily assigned labels 1 through 3, which were maintained throughout the calculations. The vector X_1 between the initial position of the first marker P_1^1 and the initial position of the second marker P_1^2 (arbitrarily assigned) was defined as:

$$X_1 = P_1^1 - P_1^2 = \begin{bmatrix} x_{1,2} & y_{1,2} & z_{1,2} \end{bmatrix}^T.$$
 (40)

The vector Y'_1 between the initial position of the first marker and the initial position of the second marker P_1^3 was:

$$Y'_{1} = P_{1}^{1} - P_{1}^{3} = \begin{bmatrix} x_{1,3} & y_{1,3} & z_{1,3} \end{bmatrix}^{T}.$$
(41)

An orthogonal vector \mathbb{Z}_1 was calculated by taking the cross product of the two marker vectors:

$$Z_1 = X_1 \times Y'_1 \tag{42}$$

Since the marker locations were projected onto the floor plane, this vector will always be normal to the floor plane. A second orthogonal vector Y_1 was calculated by taking the cross product of the first orthogonal vector and one of the marker vectors.

$$Y_1 = Z_1 \times X_1 \tag{43}$$

Since the vector X_1 was used in both cross products, it was guaranteed to be orthogonal to both Y_1 and Z_1 . These three vectors were then normalized to form the axes of the local coordinate system $(i, j, k)_1$:

$$i_1 = \frac{X_1}{norm(X_1)}$$
 (44)

$$j_1 = {Y_1 / norm(Y_1)}$$
 (45)

$$k_1 = \frac{Z_1}{norm(Z_1)}$$
(46)

The axes of the local coordinate system at the initial position were combined to form a rotation matrix R_1 that defined the initial C-arm orientation by:

$$R_{1} = \begin{bmatrix} i_{1}^{1} & j_{1}^{1} & k_{1}^{1} \\ i_{1}^{2} & j_{1}^{2} & k_{1}^{2} \\ i_{1}^{3} & j_{1}^{3} & k_{1}^{3} \end{bmatrix}$$
(47)

The rotation information can be combined with the position of the mounted camera to fully define the initial pose of the C-arm in a global coordinate frame. Using the camera location as the initial position allows the OPTIX and Optotrak measurements to be compared.

Before attaching the camera mounting system to the C-arm, the probe was used to record the location of the camera and the mounting bolts. This allowed the camera location to be calculated relative to the bolts when it was mounted beneath the wheel strut. After mounting the camera, the probe was used to record the location of the two bolts and one point on the wheel strut approximately above the camera. The point approximately above the camera was used to ensure that the rigid relationship between the bolts and the camera was applied in the appropriate direction. The resulting camera location (x_c , y_c , z_c) was taken as the reference marker, so the initial position and orientation of the local coordinate system 0 relative to the Optotrak coordinate system G was given by the transformation matrix T_{G0} (Zatsiorsky 1998):

$$T_{G0} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ x_c & i_1^1 & j_1^1 & k_1^1 \\ y_c & i_1^2 & j_1^2 & k_1^2 \\ z_c & i_1^3 & j_1^3 & k_1^3 \end{bmatrix}$$
(48)

In general, the transformation T_{12} between two local coordinate systems T_{01} and T_{02} is defined by:

$$T_{12} = T_{01}^{-1} T_{02} \tag{49}$$

The initial C-arm position was taken as the origin for Optotrak measurements since OPTIX data is also measured relative to the starting position. The transformation between the initial position and any other C-arm position was therefore used to determine the translation and rotation of the C-arm. The local coordinate system transformation matrices T_{Gi} were calculated for each C-arm position in the test set. We then calculated each C-arm position T_{0i} relative to the starting location T_{G0} by:

$$T_{0i} = T_{G0}^{-1} T_{Gi} \tag{50}$$

Since the motion was constrained to the floor plane and the Z axis was defined to be normal to the floor plane, the resulting global transformation matrices will always have the following form:

$$T_{0i} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ x_i & \cos \theta_i & -\sin \theta_i & 0 \\ y_i & \sin \theta_i & \cos \theta_i & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$
(51)

The C-arm orientation relative to the starting point at each position θ_i was extracted from the transformation matrix by:

$$\theta_i = \operatorname{atan2}\left(T_{0i_{3,2}}, T_{0i_{2,2}}\right) = \operatorname{atan2}(\sin(\theta_i), \cos(\theta_i))$$

The C-arm displacement was calculated at each position relative to the starting location for both the Optotrak and OPTIX measurements for convenient comparison by:

$$d_i = \sqrt{x_i^2 + y_i^2}$$
(52)

When the excursions were carried out, the C-arm was held stationary at each save point for 10 seconds to collect multiple measurements at each location. The mean position and rotation values were compared for each. The results of these motion studies are presented in Section 3.4.

2.3.4.4. Reacquiring X-ray Images from a Saved Location

Optotrak measurements are being used as the primary method of verification for this research. As an additional confirmation of the repositioning accuracy, x-rays were taken of a screw in a phantom pelvis model.

A radiographic image of the screw was taken and the position was recorded in OPTIX as the initial position. The C-arm was then maneuvered along the table away from the initial position (side to side motion), then repositioned until the OPTIX target turned green. A second radiograph was taken at the repositioned location. The C-arm was then moved away from the table (in and out motion) before repositioning and acquiring another x-ray. Finally, the C-arm was moved along the table and rotated (oblique motion), before reobtaining the initial position and acquiring a third x-ray.

This reacquisition test was largely user independent because the C-arm position when repeating the x-ray was fully determined by the OPTIX target. The target turns green when the C-arm is within a

translation and rotation of 5 mm and 3.1° respectively from the saved point. The purpose of this study was to verify the accuracy of OPTIX near a previously-saved location.

The reacquired x-rays were overlaid on to the initial image at a 50% transparency. The diameter of the head of the screw was measured to be 7 mm using digital calipers. This measurement allowed the distance between the leftmost edges of the screw heads between each image to be measured relative to the scale of the screw head diameter. This evaluation provided a measure of the translation error, which should be within 5 mm when the OPTIX target is green.

The angle of the screw axis was also evaluated to determine the repositioning rotation error. The screw centerline was determined manually by drawing a line from the tip through the center of the head. The angle of this line relative to the horizontal was taken to be the screw angle, and the difference was calculated between the reacquired images and the initial x-ray.

Chapter 3: Experimental Results

The results of the experimental procedures used to evaluate multiple aspects of the OPTIX system are presented in this chapter. The experimental methods used to obtain these results were outlined earlier in Section 2.3. The implications of these experimental results are discussed in Section 3.5.

3.1. Operating Room Floor Survey

A sample of operating room floors at two hospitals in British Columbia was conducted to visualize the appearance of a typical floor. Examples of the collected images are shown in Figure 42.



Figure 42 OR floors from (left to right) the Kelowna General Hospital, Kelowna General Hospital, Vancouver General Hospital, and Vancouver General Hospital.

For comparison, the floorings used for experimental evaluations for this research are shown in Figure 43. In all observed cases, the flooring consists of irregularly-patterned visual texture that can be used for computer vision based tracking.



Figure 43 (left) The flooring used for C-arm testing in the BioEng lab at CHHM and (right) for development testing with the camera slider.

The products from several flooring manufacturers, (Florock 2017) (Silikal America 2017) (Mats Inc 2010), were also examined to further define a typical operating room floor. Significant visual texture was present in all cases, suggesting that our proposed tracking method is likely applicable to a wide range of hospitals.

Although this survey is not intended to be exhaustive, we have seen no evidence so far to suggest that solid, untextured flooring is commonplace.

3.2. Development Testing

Testing was carried out on the camera slider setup described in Section 2.3.2 to guide several algorithm design decisions. The following two sections present the results and consequences of these preliminary studies.

3.2.1. Comparing Open-Loop Tracking Methods

A preliminary accuracy study was carried out to compare the performance of descriptor matching with optical flow for open-loop tracking. As described in Section 1.9, an optical flow technique was employed in the previous iteration of this research (Esfandiari 2014). A new descriptor matching approach was designed to improve robustness and accuracy, and the purpose of this testing was to ensure that this expectation held true for our application. The ORB descriptor and detector algorithm (Rublee, et al. 2011) was used to represent descriptor matching, and the Lucas Kanade method (Lucas and Kanade 1981) was used as the optical flow technique. The results shown in Figure 44 illustrate the absolute and relative errors from 10 translational movement trials of motions of approximately 2 m. There is over an order of magnitude reduction in translational error when using descriptor matching compared to the prior optical flow algorithm (from $\sim 1.75\%$ down to 0.14%).



Figure 44 Preliminary absolute accuracy results comparing Lucas Kanade optical flow tracking with ORB descriptor matching.

These results are consistent with the expectation that descriptor matching should be a more robust and accurate tracking technique than optical flow.

3.2.2. Comparing Closed-Loop Tracking methods

Further preliminary testing was carried out to explore the value of incorporating an absolute loop closure step. This step involves fine tuning the repositioning measurements when the camera is close to the initial position by matching with an image of the initial position. The results are shown in Figure 45 in terms of absolute repositioning error. Relative positioning error is not presented for this test since the loop closure step is independent of the distance travelled. The closed-loop system can provide sub-millimetric position measurement accuracy, which is an order of magnitude better than the open-loop descriptor matching approach following an excursion of approximately 2 m (typical intraoperative C-arm movement distance).



Figure 45 Absolute repositioning error for descriptor matching with and without a loop closure step.

3.2.3. Comparing Descriptor Matching Algorithms

The two descriptor matching algorithms available on OpenCV were tested for compute speed with each of the descriptor and detector algorithms. The results are presented in Figure 46 in terms of loop speed. A higher loop speed value indicates a faster algorithm.



Figure 46 Mean loop speed for feature matching algorithms over 3000 loops.

The algorithms performed identically in terms of the number of identified matches, but the Brute Force algorithm had a slight advantage in terms of computation time. FLANN (Muja and Lowe 2009) is designed to be efficient for large data sets by finding approximate nearest neighbors. This means that FLANN will find good matches, but not always the best match, while Brute Force checks every combination to find the best possible match. Since the number of features in this application is relatively small (hundreds) and the test results provide evidence of a slight speed advantage, we conclude that the Brute Force matching algorithm is the most appropriate for this project.

3.2.4. Comparing Detector and Descriptor Algorithms

Each of the three detector and descriptor algorithms was analyzed separately using the Brute Force matcher to determine the most suitable for this research. The results are presented in terms of loop speed in Figure 47, with a larger value indicating a faster algorithm.



Figure 47 Mean loop speed for the detector and descriptor algorithms over 1000 loops.

The results show that AKAZE (Alcantarilla, Nuevo and Bartoli 2013) was significantly slower than ORB (Rublee, et al. 2011) and BRISK (Leutenegger, Chli and Siegwart 2011). ORB and BRISK had very similar performance in terms of loop speed.

We then calculated the number of successful matches identified per second to further differentiate the algorithms. The results are shown in Figure 48, with a larger value indicating a faster algorithm with more successful matches.



Figure 48 Mean number of matches per second for the detector and descriptor algorithms.

These results indicate that ORB identifies the greatest number of matching features in the least amount of time. This should, in principle, correlate to more robust tracking at a greater frame rate.

Based on the evidence gathered from these preliminary studies, the ORB detector and descriptor algorithm was chosen along with the Brute Force matching algorithm.

3.3. Calibration Evaluation

The mean results of the calibration evaluation conducted on 10 images of a straight line at multiple orientations is shown in Table 9.

Table 9 Intern Niel Value of a straight line before and after apprying campration parame	Table 9	Mean RMSE value of a straight line before	and after applying calibration parameter
--	---------	---	--

	Before Correction	After Correction
Mean RMSE (px)	9.30 ± 6.46	0.851 ± 0.542

The results demonstrate an order of magnitude decrease in RMSE for the straight line after applying the calibration parameters. The large RMSE value before correction indicates significant radial lens distortion. Given the scale value defined in Section 2.2.3.2 of $\lambda = 0.137$ mm/px, this corresponds to an expected measurement error due to calibration of 1.27 ± 0.89 mm.

The mean RMSE after correction is less than a single pixel. Based on the scale value above, this corresponds to an expected measurement error due to calibration of 0.117 ± 0.074 mm.



An example result for one calibration correction is shown in Figure 49.

Figure 49 Example result for the evaluation of the straightness of a linear feature before and after applying calibration parameters. RMSE value is shown for each line of best fit.

Note that the best fit line after correction is not visible because it so closely matches the data. The effect of the calibration parameters is clearly apparent, which indicates that the calibration procedure was successful.

3.4. C-arm Testing in a Simulated Operating Room

The accuracy of the OPTIX system was evaluated using several C-arm movement patterns, as described in Section 2.3.4.2. Displacement and rotational motion was carried out to determine the rate of error accumulation in the relative tracking algorithm. Repositioning tasks were also conducted to quantify the error in the loop closure step. The motions ranged from single-direction translation to more complex patterns of translation and rotation to fully evaluate the performance of the system. A summary of the evaluated errors is presented followed by a separate assessment of the relative and absolute tracking algorithms.

3.4.1. Summary of Results

The relative tracking errors are summarized in Table 10.

Table 10	Summary	of the	relative	tracking	error rates.

Translation Error (% of cumulative distance travelled)	Rotation Error (% of cumulative rotation)		
0.75	5		

More details on the relative tracking error results are provided in Sections 3.4.2 and 3.4.3.

Figure 50 presents a summary of the mean rotation and translation errors for all C-arm repositioning tasks using the absolute position recovery algorithm for loop closure.



Figure 50 Mean tracking error for all C-arm repositioning tasks executed with the absolute position recovery algorithm (blue) compared against the accuracy thresholds for typical clinical applications (orange).

The translation and rotation errors were approximately 20% and 5%, respectively, of the clinically acceptable values established in Section 2.1. More detailed analyses of these results are presented in Sections 3.4.4 and 3.4.5.

3.4.2. Open-Loop Translation Error

The translation error was determined for open-loop relative tracking through C-arm excursions without the loop closure step. Relative motion tracking inherently results in an accumulation of small errors from each frame to frame measurement. The translation error is therefore presented in Figure 51 as a function of the distance travelled for several relative motion tests.



Figure 51 Translation error for relative tracking of C-arm motion without loop closure step. The dashed lines show the 0.75% error envelope and the 5 mm bounds for clinical acceptability.

A 0.75% error bound is included in Figure 51 as an approximate envelope of the recorded errors, as well as the clinically acceptable translation error bound of 5 mm. A positive error in the Figure indicates that OPTIX is overestimating the true displacement measurement, while a negative error indicates an underestimation. All the negative errors in the figure are from a single trial, while the rest are positive. This figure shows that problematic errors would accumulate after about 0.7 m of travel in relative tracking mode

3.4.3. Open-Loop Rotation Error

The C-arm was maneuvered without the loop closure step to determine the rotation error in the open-loop relative tracking algorithm. The results for multiple relative rotation tests are shown in Figure 52.



Figure 52 Rotation error for relative tracking of C-arm motion without loop closure step. The dashed lines show the 5% error envelope and the 3.1° error bounds.

A positive error in Figure 52 indicates that OPTIX is overestimating the C-arm angle, while a negative error indicates an underestimation. A 5% error envelope is shown in the Figure along with the desired angular accuracy of 3.1°. This figure shows that problematic errors would accumulate after a rotation of about 65° in relative tracking mode.

3.4.4. Closed-Loop Repositioning Translation Error

The translation error was analyzed separately for each of the four repositioning tasks introduced in Section 2.3.4. Save-points were recorded in OPTIX to trigger the loop closure step each time a position was repeated. Figure 53 shows the mean closed-loop translation error assessed by comparison with Optotrak for each of the four tasks.



Figure 53 Mean translation error for each of the four C-arm repositioning tasks.

The translation error was also evaluated with respect to the cumulative distance travelled by the Carm during repetitions of each of the repositioning tasks, shown in Figure 54. This presentation of the data should allow an assessment of the effectiveness of the loop closure step, which is designed to reset the accumulated relative tracking error at each save point.



Figure 54 Closed-loop translation error as a function of distance travelled for all C-arm repositioning tasks. The dashed lines represent the clinically acceptable error envelope.

A positive error in Figure 54 indicates that OPTIX is overestimating the C-arm translation measurement. There are no obvious trends apparent in the error measurements with respect to the cumulative distance travelled, which indicates that the error reset functionality is working well.

3.4.5. Closed-Loop Repositioning Rotation Error

The rotation error was also analyzed separately for each of the four C-arm repositioning tasks. The mean rotation error for each is shown in Figure 55.



Figure 55 Mean rotation error for each of the four C-arm repositioning tasks.

The rotation error was analyzed with respect to the cumulative distance travelled by the C-arm. This enables an evaluation of whether the loop closure step effectively eliminates the accumulation of relative tracking errors. Figure 56 shows rotation error as a function of the distance travelled through multiple repetitions of the repositioning tasks.



Figure 56 Rotation error as a function of distance travelled for all C-arm repositioning tasks. The dashed lines represent the clinically acceptable error envelope.

A positive rotation error indicates that OPTIX is overestimating the C-arm angle. There does not appear to be any correlation between the rotation error and the cumulative distance travelled. Which again indicates that the error reset process is working well.

3.4.6. Reacquiring X-rays

Radiographic images were taken using the C-arm to further verify the repositioning guidance provided by OPTIX. An image was initially taken of a screw in a phantom pelvis model, then the C-arm was maneuvered away before repositioning it multiple times based on the OPTIX user interface. The results are shown in Figure 57.



Figure 57 (top right) Initial x-ray, (top left) repositioned after side to side motion, (bottom left) after in and out motion, and (bottom right) after oblique motion.

These results clearly demonstrate the ability for OPTIX to successfully reobtain a previous x-ray position with little visually perceptible difference in positioning. Figure 58 shows the three repeated x-rays overlaid on top of the initial image with the head to head distance labelled.





Figure 58 The three reacquired x-rays overlaid on top of the initial image with the head to head distance labelled.

The tip to tip distances are all within the set algorithm tolerance of 5 mm. Figure 59 shows the manually determined centerline and labelled angle for the screw in each of the four images.



Figure 59 The screw centerline and labelled angle for the (top right) initial x-ray, (top left) repositioned after side to side motion, (bottom left) after in and out motion, and (bottom right) after oblique motion.

The angles of the screws in the reacquired images are all within 1° of the initial image. The screw centerlines were determined manually, which could be a source of error. However, the screw angles in the repeated images are well within the clinically acceptable 3.1° bound.

3.5. Discussion of Results

The experimental procedures in this study were designed to evaluate the position measurement errors in the OPTIX C-arm base-tracking system. Preliminary testing was carried out on a camera slider to guide several algorithm design choices. The purpose of this testing was to empirically determine which tracking methods would be best suited to this application. An evaluation of the camera calibration procedure was conducted using the appearance of a straight line before and after the application of the calibration parameters. This study was intended to externally verify the calibration procedure using images that were not included in the training set. Finally, multiple C-arm motion studies were carried out to evaluate the accuracy of OPTIX on a real C-arm in a simulated operating room. The relative tracking algorithm was assessed by translating and rotating the C-arm and tracking the accumulation of errors relative to a reference provided by a high resolution external tracking system (Optotrak). The absolute tracking algorithm was evaluated through several repositioning tasks. The motion patterns for the repositioning tasks were chosen to be representative of typical C-arm maneuvers during surgery. The translation and rotation errors were evaluated relative to the tracking requirements established in Section 2.1. A demonstration of the repositioning accuracy was conducted by using OPTIX to repeat x-rays of a screw after maneuvering the C-arm away from the initial position multiple times.

3.5.1. Translation Error in Open-Loop Tracking Algorithm

The C-arm was maneuvered through distances up to 3.5 m during the open-loop translation evaluation, which is longer than the full length of an operating room table. The open-loop algorithm is expected to accumulate incremental errors as the C-arm travels. This accumulation of error is an inherent drawback of any relative tracking method.

Nearly all the recorded translation errors were positive, which indicates that OPTIX overestimated the translation relative to the Optotrak measurements. This positive bias appears to be relatively consistent, so it could potentially be characterized and corrected for in future iterations. The errors are approximately bounded by lines drawn at 0.75% of the distance travelled, which is a notable two-fold improvement over the 2% error reported in the previous phase of this research (Esfandiari 2014). This reduction in relative tracking error is likely a result of the improvements made to the feature detection and matching techniques. ORB is among the state-of-the-art feature detection and mapping) has gained traction as an popular robot navigation method in recent years because of its high accuracy and ability to be applied in real-time (Mur-Artal, Montiel and Tardos 2015).
For repositioning tasks, a 0.75% error means the C-arm can be maneuvered approximately 5.7 m and expect to achieve the same level of accuracy (4.3 cm error) as that reported for orthopaedic surgeons in a study carried out by a member of our research group (Touchette 2017).

A translation error of 0.14% was measured for open-loop ORB tracking during development testing, which is notably smaller than the 0.75% reported for C-arm testing. An important potential source of reduced error in the development testing setup is the fact that the camera was moved back and forth across the slider. This was done to extend the effective travel distance of the camera, but moving backwards could lead to an artificial reduction in the measurement error. This was carried out for all development testing, so the results for different tracking algorithms tested on the slider can be compared to one another. However, the slider result may have an unfair advantage compared to the one-way C-arm translation.

To satisfy the more stringent requirements of panoramic image generation (\leq 5 mm error) (Amini 2016), 0.75% error corresponds to a maximum C-arm excursion distance of 0.67 m. This means that the error from one image to the next in the panorama should satisfy the requirements. However, the error relative to the first image could accumulate to unacceptable levels over the full length of an adult spine. A study should be carried out in the future that specifically focuses on evaluating the use of this base-tracking method for panoramic stitching to determine if this error rate is acceptable.

A potential source of error for the translation measurement is the scale value calculated by the scale calibration step. However, the scale calibration procedure was repeated on several images of the same target, and the scale value was consistent within ±0.01 mm/px. Therefore, the error is likely not due to any variability in our calibration procedure. A potential method of reducing any possible scale errors would be to design a different method of determining the scale. Tracking a C-arm movement with OPTIX without any scale value and comparing the result to an Optotrak measurement could be one method of determining the scale. Expanding the size of the current scale target, moving the C-arm a specified distance using physical constraints, or placing a ruler-like target on the floor could be other potential methods.

3.5.2. Rotation Error in Open-Loop Tracking Algorithm

The C-arm was rotated up to approximately 70° for the open-loop evaluation tasks, a greater angle than what would usually be expected in a typical surgical setting. The open-loop tracking method

inherently leads to an accumulation of error, so it was expected that the angular error would increase as the C-arm is rotated.

There was a clear negative bias apparent in the data, which indicates that OPTIX underestimated the C-arm rotation. The negative bias appears to be consistent, and could therefore potentially be characterized and corrected for in future iterations. It is notable that the accumulations of errors occur in the opposite direction for translation and rotation. However, the calculation of rotation is independent of the scale value, while the translation calculation is not. Any error introduced by the scale value would therefore have an effect on the translation measurement, while the rotation errors must be from a different mechanism.

The errors appear to be approximately bounded by a line drawn at 5% of the C-arm rotation, which is consistent with the results obtained in the previous phase of this research using an optical flow method, even though the translation errors improved by more than a factor of two (Esfandiari 2014). A possible reason for this discrepancy is the fact that the rotation error was reported for a much larger camera height in the previous research (70 cm) compared to this system (~10 cm). This significant reduction in camera height leads to a smaller field of view. It is possible that the decreased field of view has a stronger effect on rotation errors than translation error, which would explain our results.

The rotation error could be improved by increasing the field of view using a wide-angle lens or a system of mirrors, or by adding a second camera to significantly increase the moment arm when calculating orientation. We expect that a second camera could be added to the system without reducing the framerate, since there are currently 3 parallel processes operating on 8 cores. A new process could be introduced for the second camera to operate in parallel, which theoretically should not reduce the loop rate. The current moment arm is constrained by the camera field of view, which is approximately 6 cm. The second camera could be added to the rear of the C-arm, which would significantly increase the moment arm to 80 cm. We expect that the rate of accumulation of rotation errors would notably decrease with this modification. Using an inertial measurement unit or magnetometer to measure the C-arm orientation could also improve the rotation error through sensor fusion.

A desired rotation error of better than 3.1° was established in Section 2.1 to satisfy envisioned surgical applications. A 5% error bound means that OPTIX is expected to have better than 3.1° of error until approximately 62° of cumulative rotation in a single direction. This is an encouraging

result since 60° is the maximum expected rotation in the operating room (moving from 30° in one direction to 30° in the other direction) based on observations. These results indicate that OPTIX has clinically relevant rotation accuracy even without the use of the loop closure step.

3.5.3. Translation Error for Closed-Loop C-arm Repositioning

The mean translation error for all closed-loop repositioning tasks was 1.10 ± 0.07 mm. This result is well within the established error requirement of 5 mm, indicating that OPTIX can be used for repositioning tasks with clinically relevant accuracy. The accuracy of systems proposed by other research groups range from 0.64 mm (Moult, et al. 2011) to 1.9 mm (Yoo, et al. 2013), as outlined in Section 1.5, which is comparable to our result. It is important to note that the suitability for use in the operating room of the systems proposed by other groups was often the issue, despite successful accuracy results.

The translation errors were analyzed separately for each of the four repositioning tasks. There were no important differences in the errors between the tasks, as all four motion patterns had errors well within the required 5 mm bound. The mean error value for each task was within 0.5 mm of the mean for all tasks. The side to side task had the smallest reported translation error, which could be because the wheels on the C-arm can be turned horizontally. This configuration limits the C-arm motion and allows the user to move it side to side along a consistent path with minimal rotation.

The repositioning translation error was also examined as a function of the cumulative distance travelled. The loop closure step in the repositioning algorithm was shown to successfully reset the accumulation of relative tracking errors between save points for all tasks. The repositioning errors were spread approximately evenly between overestimating and underestimating the C-arm position, and the measured errors lie within the required 5 mm bound regardless of the distance travelled. There were several errors in this evaluation that were close to the 5 mm bound, which is larger than expected. Notably, 6 of the largest errors were from just 2 trials, indicating that an issue may have occurred during those trials. Other potential explanations for these large errors could be the fact that the Optotrak markers were placed far from the C-arm base. This presents an opportunity for error to be introduced when transforming the Optotrak data into the floor plane. We expect that a measurement error due to the setup would be more likely than an actual error of 5 mm in the loop closure algorithm.

These results demonstrate that OPTIX provides consistent, accurate repositioning measurements regardless of the distance travelled or the number of times a save-point is repeated. This study also demonstrated that our system can accurately track the C-arm over clinically useful distances. This is in contrast to most of the tracking systems proposed by other research groups that limited the tracked range of the C-arm such as the mounted in-line camera (S. Reaungamornrat, et al. 2014) (Wang, Traub and Heining, et al. 2008) and the TC-arm (Amiri, et al. 2014).

3.5.4. Rotation Error for Closed-Loop C-arm Repositioning

The mean rotation error for all closed-loop repositioning tasks was $0.17 \pm 0.02^{\circ}$. This result is significantly smaller than the established error requirement of 3.1° , which indicates that OPTIX has clinically useful rotation accuracy for repositioning tasks.

Only one other group identified in Section 1.5 reported a rotation accuracy. The radiographic fiducial tracking system reported an average rotation error 0.68° (Moult, et al. 2011). The rotation error for our system is a fraction of this value for repositioning tasks.

The mean rotation errors were also examined separately for each of the four repositioning tasks. The oblique and four position tasks had smaller rotation errors than the side to side and in and out tasks despite the former being more complex motion patterns. However, all four tasks had errors well below the requirement and were within 0.08° of the mean for all tasks, so these discrepancies do not appear to be significant.

The repositioning rotation error was also assessed as a function of the cumulative distance travelled to determine the effectiveness of the loop closure step in eliminating any accumulation of relative tracking errors. The results indicate that the loop closure step was successful in eliminating accumulated tracking errors at each save point. The errors appear to be approximately evenly distributed between overestimating and underestimating the C-arm orientation and all the measured errors lie well within the required 3.1° bounds regardless of the distance travelled.

These results are promising and demonstrate that OPTIX can be used for C-arm repositioning tasks with rotation errors that are comfortably within the clinical requirements and smaller than those reported by other research groups. The loop closure step effectively eliminates the accumulation of relative tracking errors and enables accurate repositioning regardless of the distance travelled.

3.5.5. Reacquiring X-rays

The x-ray reacquisition task demonstrated the ability for OPTIX to accurately track the C-arm position near a previously saved point. The C-arm was maneuvered until the OPTIX target turned green, so the resulting images were largely user-independent within the set algorithm translation and rotation tolerances of 5 mm and 3.1° respectively. The reacquired images showed little visually perceptible differences relative to the initial position. The head to head distances and centerline angles were all within the set tolerances, which indicates that the repositioning algorithm was functioning well.

The accuracy of the reacquired images could be improved by tightening the set tolerances in the OPTIX algorithm. Currently the target will turn green when the C-arm is within 5 mm and 3.1° of the previous position. However, tightening these tolerances would likely increase the amount of time required to successfully position the C-arm. The time was not recorded in this study, but would be a useful metric in a future user study. There were no scouting images required for any of the images, and all reacquisition tasks were successful on the first try. A recent study by a member of our lab found that an average of 8 scouting images were required to obtain a given radiograph (Touchette 2017), so this represents a marked improvement. A formal user study could be carried out in the future to further explore the utility and inter-user performance of the OPTIX guidance system.

Chapter 4: General Discussions and Conclusions

This chapter provides a summary of the thesis contributions, discusses limitations in the developed system, and identifies several directions for future work.

4.1. Thesis Contributions

A real-time, on-board base-tracking system using a single floor-facing camera was developed for tracking C-arm fluoroscopes across unaltered operating room floors. The system is based on previous work carried out in our lab group (Esfandiari 2014), which saw the development of an off-line tracking algorithm for relative motion using an optical flow algorithm. This thesis has expanded on the previous work in several important ways. The current system incorporates new and more robust feature detection and matching algorithms, leading to a demonstrated increase in relative tracking accuracy. The system now functions in real-time, providing continuously updated position measurements through a graphical user interface. The algorithm architecture exploits multiple CPU cores to parallelize several processes, enabling accurate tracking at 8 fps and a responsive user interface. The algorithm can estimate C-arm position with translation errors better than 0.75% of the one-way distance and rotation errors better than 5% of the cumulative one-way rotation. A loop closure step was developed that has been shown to eliminate the accumulation of errors during relative tracking when returning to previously-visited locations and to enable highly accurate C-arm repositioning. The loop closure algorithm achieved average position and rotation errors of $1.10 \pm$ 0.07 mm and 0.17 \pm 0.02° respectively. The developed user interface features an application to save multiple C-arm positions for later reacquisition, which automatically triggers the loop closure step. The user interface has been implemented on a touch screen device, which provides an intuitive platform to interact with the tracking system. Multiple hardware systems were changed or designed to advance the base-tracking system towards clinical relevance. A custom mounting system was constructed to rigidly fix the camera and processing unit beneath the base of a C-arm. The mounting configuration provides adequate height for tracking despite the limited available space. A selfcontained lighting system was also constructed to provide consistent illumination and adequate contrast for the camera.

The developed tracking system cost approximately \$800, which is inexpensive relative to conventional tracking systems, and uses a single monochrome camera as the lone sensor. The entire system can be easily retrofitted to existing C-arms in a matter of minutes.

Individual aspects of the developed algorithm incorporate multiple well-established computer vision techniques. However, the overall algorithm and implementation are novel, to our knowledge. Tracking systems based on floor-facing cameras have been developed for some robotics (Kelly 2000) (Dille, Grocholsky and Singh 2009) (Killpack, et al. 2010) and vehicle (Aqel, et al. 2015) (Farraj, et al. 2013) applications, but the configuration has not been applied to the medical field to our knowledge. The systems developed by Kelly (Kelly 2000) and Aqel (Aqel, et al. 2015) make use of full image correlation, which is more computationally expensive than feature-based matching. Additionally, Kelly's system requires that the user first moves the robot along the desired path manually to generate a mosaic-based map before the robot can travel on its own. This map-based tracking system is inherently an absolute tracking method. The systems developed by Dille (Dille, Grocholsky and Singh 2009) and Killpack (Killpack, et al. 2010) both use optical flow algorithms for tracking, which we have shown to be less accurate than descriptor matching. The experimental setup used by Killpack (Killpack, et al. 2010) used checkerboard patterns on the floor rather than standard, unaltered floor textures. All the other systems proposed in the literature either have some form of absolute tracking or loop closure, altered floor appearances, or lower accuracy. To the best of our knowledge, the open-loop tracking aspect of our system is the most accurate real-time open-loop tracking system based on a single floor-facing camera developed to date.

This thesis has made several important steps towards clinical implementation of the C-arm basetracking system. The electrical components were chosen according to requirements established by the British Columbia Safety Authority to ensure the system can be used in an operating room. The sensor and processing unit are mounted fully on-board the C-arm, which eliminates the line of sight limitation associated with conventional tracking systems. The system has been implemented and tested on a real C-arm in a simulated operating room. Several motion patterns, representative of real intraoperative C-arm movements, were incorporated in the experimental testing to fully quantify the measurement error in the system. Experimental evaluations of the developed system have demonstrated clinically useful levels of accuracy.

4.2. Limitations

The maximum speed that the C-arm can be maneuvered while tracking is limited by the camera frame rate. The tracking system has been verified up to approximately 0.04 $m/_S$ during experimental testing, which is representative of typical C-arm speeds we have observed during orthopaedic procedures. The camera used in the system was chosen to maximize the distance from the floor, and the tracking algorithm was designed with a focus on maximizing the frame rate. These two factors

directly effect the maximum C-arm speed that can be accurately tracked. However, if the C-arm is moved too quickly, which unfortunately is not uncommon, there will not be enough overlap between consecutive images to determine the geometric transformation. This scenario can result in an underestimate of the total C-arm displacement or loss of tracking since some motion is not accounted for. There are at least three methods of compensating for this speed limitation. First, a mosaic-based map method similar to the approach used by Kelly (Kelly 2000) could be added to the tracking algorithm. This would allow the system to correct its position when the C-arm was in any previouslytravelled location. The current system has shown that position correction at explicitly saved points is an effective method of resetting any error accumulation, so an expanded implementation of this method could further improve the performance of the system. Any time a loss of tracking or major error occurs, the system could indicate to the user that they should maneuver the C-arm to any previously-travelled area to re-orient the tracking system without losing any saved points. A large reduction in storage space relative to Kelly's implementation could be achieved by storing feature points and descriptions rather than full images. However, based on some preliminary research, this solution would likely require a more powerful processor to operate at an acceptable frame rate since the algorithm would need to search a relatively large dataset. This leads to the second potential method of compensating for the speed limitation: the processing system could be upgraded to provide a faster frame rate. The current camera can provide images at 30 fps, but the tracking algorithm processes images at approximately 8 fps. A more powerful processing unit could increase the algorithm loop speed. As an example, the NVIDIA Jetson TX1 mentioned in Section 2.2.1.2 has more RAM and a more powerful GPU. This board costs over 3 times the price of the Odroid-XU4, but the increased CPU power and potential GPU acceleration could lead to a two-fold increase in frame rate. The third approach to handling this limitation is simply to maintain reasonable C-arm speeds during tracking. Base-tracking performance should not be affected if the MRT takes care not to maneuver the C-arm at greater than normal speeds.

The translation error in the relative tracking algorithm is expected to accumulate at a rate of approximately 0.75% of the total distance travelled. Although this represents a marked improvement over the previous iteration of this system, this rate can potentially lead to clinically problematic levels of error over long excursion distances (e.g. an error of 7.5 mm over 1 m of travel). Similarly, the rotation error accumulates at approximately 5% of the cumulative rotation. The errors for panoramic image generation and other potential applications need to be explored in future research phases to determine whether this error rate is acceptable for specific applications. The most straightforward method to improve this error rate would be to increase the field of view of the camera. This would

allow the camera to view a larger area of the floor and detect a greater number of features. An increased field of view would also provide a larger moment arm that could increase the rotation accuracy. The vertical space beneath the C-arm base is limited, but the field of view could be increased using a wide-angle lens or a series of mirrors. A wide-angle lens introduces a large amount of distortion to images of close objects, which could counteract any accuracy improvement from the increased field of view. A series of mirrors could also be used to increase the field of view by effectively extending the distance between the camera and the floor. These and other methods could be explored in the future.

Another limitation in the current system is the fact that the relative positions of the operating table and the patient are not tracked. The tracking algorithm calculates motion relative to the initial C-arm position, while assuming the patient will remain stationary. In practice, however, a patient's limbs are often maneuvered by the operating staff throughout the procedure. This motion may be minimal for procedures focused on the spine or pelvis, but can be significant for the extremities. A complete system for intraoperative guidance would also need elements for tracking the other degrees of freedom of the C-arm as well as motions of the patient; such elements are beyond the scope of the current work.

Finally, the current system has not yet been implemented in a real operating room. There are additional challenges in surgical scenarios that have not yet been explicitly tested and are difficult to predict. In particular, blood splatter and other debris could change the appearance of the floor between passes of the C-arm base. We have not yet fully tested these conditions, although the descriptor matching algorithm we chose is theoretically capable of accommodating some percentage of missing features between image pairs without significantly affecting the registration accuracy. This capability was explored during development testing by occluding part of the floor with a sheet of paper while reacquiring a previous position using the loop closure algorithm. This testing showed that the algorithm can identify the appropriate previous position with some level of occlusion, but should be explored more formally in the future.

4.3. Future Directions

The next step for this research is to test the performance and utility of OPTIX in an intraoperative setting. This would require approval from the Biomedical Engineering department at Vancouver General Hospital. Steps have been taken throughout the design process to use approved electrical power components that meet the Biomedical Engineering department standards. Additionally,

approval from the Vancouver Coastal Health Research Institute and the University of British Columbia research ethics board. The primary goals of this study would be to assess the utility of the repositioning guidance through feedback from MRTs as well as usage data. Useful data would include the time spent moving the C-arm and the number of scouting images taken during surgical procedures with and without the use of OPTIX.

Next, we would like to integrate the base-tracking system with a system to track the rest of the C-arm joints. An on-board system such as the TC-arm (Amiri, et al. 2014) could be incorporated to enable full 6 degree of freedom tracking. The addition of joint tracking will necessitate further development of the user interface to provide information on the full C-arm pose.

Another member of our lab group is developing a powered omni-wheel base to facilitate C-arm maneuvering. The OPTIX system could be integrated with the powered base to provide a closed-loop propulsion and tracking system. The integration of the powered base and tracking of the other C-arm joints would result in a fully tracked C-arm that could serve as a platform for enabling a wide range of quantitative functions and associated surgical applications.

A communication channel between the OPTIX system and the C-arm imaging computer could be established to enable further application development based on analysis of images from the C-arm. Implementation methods are currently being explored to transmit radiographic images to the OPTIX processing unit. For example, access to the radiographs would allow the user to select an anatomical point of interest in an image through the user interface. OPTIX could then provide guidance to the user to reposition the C-arm so that the point of interest is centered in the radiograph. Additionally, as mentioned in Section 4.2, a patient-tracking method could be developed to account for any movement of the patient anatomy relative to the C-arm. A configuration similar to the mounted inline camera systems (S. Reaungamornrat, et al. 2014) (Wang, Traub and Heining, et al. 2008) outlined earlier in Section 1.5.4 could be integrated with our base-tracking system as one potential method to track patient motion.

There are several other useful applications that could be developed in the future. Panoramic image stitching could be implemented once the communication channel has been established to obtain radiographic images from the C-arm. A member of our lab (Amini 2016) has developed a stitching process that can be integrated with the developed base-tracking system. Another member of our research group recently developed an artificial x-ray generation system called AXIS (Touchette 2017). Integrating the artificial x-ray system with this on-board tracking system would contribute to

enabling increase the ability for AXIS to be implemented in the operating room. A needs assessment could be conducted to determine the relative priorities for further applications such as limb measurement and computer assisted navigation.

Several user studies could be conducted once the full tracking system and powered base have been integrated. The purpose of these user studies will be to assess the utility of the C-arm tracking system regarding the developed applications. The primary assessment metrics would be focused on usability, accuracy, and whether the system reduces radiation exposure and operation time. The desired participant demographic will depend on the application that is being studied. Working MRTs should be included in all cases since they are the primary operators of C-arms in clinical practice. For applications such as panoramic stitching, orthopaedic surgeons should also be involved in the evaluation. The initial user studies should be carried out in a simulated operating room setting, such as the one used for this research. The next step following these assessments would be to implement the system in a real operating room.

4.4. Concluding Remarks

This thesis has detailed the design, development, and verification of a real-time C-arm base-tracking system. The developed system includes both hardware and software elements, and has been designed to work sufficiently quickly to enable intraoperative use. The camera configuration eliminates the line of sight requirement associated with conventional tracking methods, making the system unobtrusive and therefore highly suitable for use in a busy operating room.

Our proposed tracking system is a novel implementation of well-established computer vision techniques, and is a unique approach to C-arm tracking to our knowledge. Many of the tracking systems that have been proposed by other research groups have limited effective ranges and are not practical for use in the operating room. Our system was designed to measure C-arm motion over the full range of expected base movement. Despite the high frequency of base translation, we are not aware of any other group that has focused on tracking the C-arm base. Our system could be combined with other proposed tracking methods such as the TC-arm (Amiri, et al. 2014) or mounted in-line camera systems (S. Reaungamornrat, et al. 2014) (Wang, Traub and Heining, et al. 2008) to expand their tracking range and lead to a practical, fully tracked C-arm. This research has made important steps towards achieving a practical and accurate intraoperative tracking system for C-arm fluoroscopes.

The accuracy of the system was evaluated on a real C-arm in a simulated operating room. The experimental results demonstrated that the feature-based relative tracking algorithm can track the C-arm position with errors of less than 0.75% of the total distance travelled and orientation with errors better than 5% of the cumulative rotation. With the incorporated loop closure step, OPTIX can be used to achieve C-arm repositioning with average position errors of less than 1.10 ± 0.07 mm and rotation errors of less than $0.17 \pm 0.02^{\circ}$. These results are well within the requirements for clinical utility in a variety of envisioned surgical applications.

The performance of the system is promising: C-arm base-tracking alone could contribute to a decreased reliance on scouting images by providing accurate repositioning guidance, and further development could lead to a fully tracked C-arm with multiple potential clinical applications. In an operating room, OPTIX has the potential to lead to a reduction in operating time and harmful radiation exposure to surgical staff.

References

- Alcantarilla, P F, J Nuevo, and Adrien Bartoli. 2013. "Fast explicit diffusion for accelerated features in nonlinear scale spaces." *British Machine Vision Conference*. Bristol.
- Amini, Mohammed. 2016. "A fluoroscopy-based intraoperative tool for measuring alignments in spinal deformity correction surgery (Masters thesis)." Vancouver.
- Amiri, Shahram, David R. Wilson, Bassam A. Masri, and Carolyn Anglin. 2014. "A low-cost tracked C-arm (TC-arm) upgrade system for versatile quantitative intraoperative imaging." *International Journal of Computer Assisted Radiology and Surgery* 9 (4): 695-711.
- Anthony, Lisa, Quincy Brown, Berthel Tate, Jaye Nias, Robin Brewer, and Germaine Irwin. 2014. "Designing smarter touch-based interfaces for educational contexts." *Personal Ubiquitous Computing* 18: 1471-1483.
- Aqel, Mohammad O. A., Mohammad H. Marhaban, M Iqbal Saripan, Napsiah Bt. Ismail, and Asem Khmag. 2015. "Optimal configuration of a downward-facing monocular camera for visual odometry." *Indian Journal of Science and Technology* 8 (32): 1-8.
- Arrow. 2017. Accessed June 29, 2017. https://www.arrow.com/en/products/900-82180-0001-000/nvidia.
- Bajpai, Pramendra Kumar, Inderdeep Singh, and Jitendra Madaan. 2012. "Tribological behavior of natural fiber and reinforced PLA composites." *Wear* 829-840.
- Bonarini, A, M Matteucci, and M Restelli. 2005. "Automatic error detection and reduction for an odometric sensor based on two optical mice." *International Conference on Robotics and Automation.* Barcelona. 1675-1680.
- Bonenberger, Paul R. 2016. "Lock feature development: rules of thumb." In *The first snap-fit handbook:* creating and managing attachments for plastic parts, 251-268. Hanser Publishers.
- Borenstein, J, H R Everett, L Feng, and D Wehe. 1997. "Mobile robot positioning: Sensors and techniques." *Journal of Robotic Systems* 14 (4): 231-249.
- British Columbia Safety Authority. 2016. "Approved certification marks for electrical products." February 29.
- Calonder, Michael, Vincent Lepetit, Christoph Strecha, and Pascal Fua. 2010. "BRIEF: binary robust independent elementary features." *Proceedings of the 11th European Conference on Computer Vision.* 778-792.
- Canny, John. 1986. "A computational approach to edge detection." *IEEE Transactions on Pattern Analysis and Machine Intelligence* PAMI-8 (6): 679-698.
- Chou, Loretta B, Lori B Lerner, Alex H.S. Harris, Ashley J Brandon, Sabine Girod, and Lesley M Butler. 2015. "Cancer prevalence among a cross-sectional survey of female orthopedic, urology, and plastic surgeons in the United States." *Women's Health Issues* 25 (5): 476-481.
- Ciraj-Bjelac, Olivera, Maden M Rehani, Kui Hian Sim, Houng Bang Liew, Eliseo Vano, and Norman J Kleiman. 2010. "Risk for radiation induced cataract for staff in interventional cardiology: Is there reason for concern?" *Catheterization and Cardiovascular Interventions* 76: 826-834.
- Cizik, Amy M., Michael J. Lee, Brook I. Martin, Richard J. Bransford, Carlo Bellabarba, Jens R. Chapman, and Sohail K. Mirza. 2012. "Using the spine surgical invasiveness index to identify risk of surgical site infection." *The Journal of Bone and Joint Surgery* 94 (4): 335-342.
- Clave, A., V. Sauleau, D. Cheval, T. Williams, C. Lefevre, and E. Stindel. 2015. "Can computer-assisted surgery help restore leg length and offset during THA? A continuous series of 321 cases." *Orthopaedics & Traumatology: Surgery & Research* 101: 791-795.
- CNX Software. 2017. Accessed June 29, 2017. http://www.cnx-software.com/2012/10/28/249-samsung-exynos-5-cortex-a15-arndale-development-board/.

Diigiit Robotics. 2017. Accessed June 29, 2017. http://www.ca.diigiit.com/odroidxu4?search=odroid%20xu4.

- Dille, Michael, Benjamin P. Grocholsky, and Sanjiv Singh. 2009. "Outdoor downward-facing optical flow odometry with commodity sensors." *Proceedings in Field & Service Robotics (FSR '09).* Cambridge.
- Esfandiari, Hooman. 2014. "Photogrammetric advances to C-arm use in surgery (Masters thesis)." Calgary.
- Evangelidis, Georgios D, and Emmanouil Z Psarakis. 2008. "Parametric image alignment using enhanced correlation coefficient maximization." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 30 (10): 1858-1865.
- Farraj, Firas Abi, Daniel Asmar, Elie Shammas, and Imad Elhajj. 2013. "Non-iterative planar visual odometry using a monocular camera." *16th International Conference on Advanced Robotics.* Montevideo: IEEE.
- FLIR Integrated Imaging Solutions Inc. 2015. "FLIR White Papers: Machine vision interface comparison and evolution."
- FLIR Integrated Imaging Solutions. 2016. "Technical Application Note (TAN2014009): Streaming cameras on embedded systems."
- Florock. 2017. Operating & Surgical Room Flooring. Accessed July 20, 2017. http://www.florock.net/industrial-flooring-systems/heathcare-medical-industry-epoxy-flooringsolutions/operating-surgical-room-flooring/.
- Forbes. 2016. Accessed June 29, 2017. https://www.forbes.com/sites/greatspeculations/2016/11/23/what-is-driving-the-surge-innvidias-automotive-revenues/#4ead234a186c.
- Forsyth, David, and Jean Ponce. 2012. Computer Vision: A Modern Approach. New Jersey: Pearson.
- Gadd, Matthew, and Paul Newman. 2015. "A framework for infrastructure-free warehouse navigation." *IEEE International Conference on Robotics and Automation.* Seattle.
- Galbraith, J. G., D. P. O'Leary, H. L. Dailey, A. Mitra, and J. A. Harty. 2012. "Preoperative estimation of tibial nail length Because size does matter." *Injury* 43: 1962-1968.
- Geleijns, Jacob, and Jan Wondergem. 2005. "X-ray imagine and the skin: radiation biology, patient dosimetry and observed effects." *Radiation Protection Dosimetry* 114 (1): 121-125.
- Grelat, Michael, Joel Greffier, Pascal Sabatier, Cyril Dauzac, Guillaume Lonjon, Bertrand Debono, Julien Le Roy, Pascal Kouyoumdjian, and Nicolas Lonjon. 2016. "Assessment of the radiation exposure of surgeons and patients during a lumbar microdiskectomy and a cervical microdiskectomy: a french prospective multicenter study." *World Neurosurgery* 89: 329-336.
- Hardkernal. 2017. *Odroid-XU4*. February 19. Accessed August 10, 2017. http://www.hardkernel.com/main/products/prdt_info.php.
- Harris, Chris, and Mike Stephens. 1988. "A combined corner and edge detector." *Proceedings of Fourth Alvey Vision Conference*. 147-151.
- Hiniker, Susan M, and Sarah S Donaldson. 2014. "ALARA: in radiation oncology and diagnostic imaging alike." *Oncology* 247.
- Hough, P. V. C., and A. Arbor. 1962. Method and means for recognizing complex patterns. United States of America Patent 3069654.
- Hulens, Dries, Jon Verbeke, and Toon Goedeme. 2015. "How to choose the best embedded pocessing platform for on-board UAV image processing." *IEEE International Joint Conference on Computer Vision, Imaging, and Computer Graphics Theory and Applications.*
- Kelly, Alonzo. 2000. "Mobile robot localization from large-scale appearance mosaics." *The International Journal of Robotics Research* 19 (11): 1104-1125.

- Killpack, Marc, Travis Deyle, Cressel Anderson, and Charles C Kemp. 2010. "Visual odometry and control for an omnidirectional mobile robot with a downward-facing camera." *IEEE International Conference on Intelligent Robots and Systems.* Taipei.
- Leutenegger, Stefan, Margarita Chli, and Roland Y Siegwart. 2011. "BRISK: Binary robust invariant scalable keypoints." *IEEE International Conference on Computer Vision.* Barcelona.
- Lowe, David G. 2004. "Distinctive Image Features from Scale-Invariant Keypoints." *International Journal of Computer Vision* 60 (2): 91-110.
- Lucas, Bruce D, and Takeo Kanade. 1981. "An iterative image registration technique with an application to stereo vision." *International Joint Conference on Artificial Inteligence*. Vancouver. 674-679.
- Mahajan, Anupam, Sumant Samuel, Atul K Saran, M K Mahajan, and M K Mam. 2015. "Occupational radiation exposure from C arm fluoroscopy during common orthopaedic surgical procedures and its prevention." *Journal of Clinical and Diagnostic Research* 9 (3): 1-4.
- Mathurosemontri, Suchalinee, Supaphorn Thumsorn, Satoshi Nagai, and Hiroyuki Hamada. 2017. "Investigation of friction and wear behavior of polyoxymethylene/poly(lactic acid) blends." *Key Engineering Materials* 728: 229-234.
- Mats Inc. 2010. *Operating Room Flooring*. Accessed July 21, 2017. http://matsinc.com/blog/architect-designer-news/operating-room-flooring.
- Matthews, Felix, Dominik J. Hoigne, Manfred Weiser, Guido A. Wanner, Pietro Regazzoni, Norbert Suhm, and Peter Messmer. 2007. "Navigating the fluoroscope's C-arm back into position: an accurate and practicable solution to cut radiation and optimize intraoperative workflow." *Journal of Orthopaedic Trauma* 21 (10): 687-692.
- Mechlenburg, Inger, Henrik Daugaard, and Kjeld Soballe. 2009. "Radiation exposure to the orthopaedic surgeon during periacetabular osteotomy." *International Orthopaedics* 33: 1747-1751.
- Moult, E., E. C. Burdette, D. Y. Song, P. Abolmaesumi, G. Fichtinger, and P. Fallavollita. 2011. "Automatic C-arm pose estimation via 2D/3D hybrid registration of a radiographic fiducial." *Proc. SPIE Medical Imaging 2011: Visualization, Image-Guided Procedures, and Modeling.* Lake Buena Vista.
- Muja, Marius, and David G Lowe. 2009. "Fast approximate nearest neightbors with automatic algorithm configuration." *International Conference on Computer Vision Theory and Applications.* 331-340.
- Mur-Artal, Raul, J M. M. Montiel, and Juan D Tardos. 2015. "ORB-SLAM: A versatile and accurate monocular SLAM system." *IEEE Transactions on Robotics* 31 (5): 1147-1163.
- Nister, David, Oleg Naroditsky, and James Bergen. 2005. "Visual odometry for ground vehicle applications." *Journal of Field Robotics* 23 (1): 3-20.
- Northern Digital Inc. 2015. "Optotrak Certus: Technical Specifications." Accessed July 3, 2017. https://www.ndigital.com/msci/products/optotrak-certus/optotrak-certus-technicalspecifications/.

 2017. Polaris Measurement Volume. Accessed August 11, 2017. https://www.ndigital.com/medical/products/polaris-family/features/measurement-volume/.

- NVIDIA. 2017. Accessed June 29, 2017. https://store.nvidia.com/store?Action=DisplayPage&Locale=en_US&SiteID=nvidia&id=QuickBuy CartPage.
- OpenCV Dev Team. 2014. "OpenCV Docs: Camera Calibration with OpenCV." November 10. Accessed July 12, 2017. http://docs.opencv.org/3.0-

beta/doc/tutorials/calib3d/camera_calibration/camera_calibration.html.

Pally, Elliott, and Hans J Kreder. 2013. "Survey of terminology used for the intraoperative direction of Carm fluoroscopy." *Canadian Journal of Surgery* 56 (2): 109-112. Reaungamornrat, S, Y Otake, A Uneri, J W Stayman, W Zbijewski, D J Mirota, J Yoo, et al. 2011. "Trackeron-C: A novel tracker configuration for image-guided therapy using a mobile C-arm." *Computer Assisted Radiology and Surgery.* Berlin.

Reaungamornrat, S., Y. Otake, A. Uneri, S. Schafer, D. J. Mirota, S. Nithiananthan, J. W. Stayman, et al.
2014. "An on-board surgical tracking and video augmentation system for C-arm image guidance." *International Journal of Computer Assisted Radiology and Surgery* 7 (5): 647-665.

Reaungamornrat, Sureerat, Y Otake, A Uneri, S Schafer, D J. Mirota, S Nithiananthan, J W Stayman, et al. 2014. "An on-board surgical tracking and video augmentation system for C-arm image guidance." *International Journal of Computer Assisted Radiology and Surgery* 7 (5): 647-665.

Rodriguez-Telles, Francisco Geovani, L Abril Torres-Mendez, and Edgar A Martinez Garcia. 2013. "A fast floor segmentation algorithm for visual based robot navigation." *International Conference on Computer and Robot Vision.*

Rosten, Edward, and Tom Drummond. 2006. "Machine learning for high-speed corner detection." In *Computer Vision-ECCV 2006*, 430-443. Springer Berlin Heidelberg.

Rublee, Ethan, Vincent Rabaud, Kurt Konolige, and Gary Bradski. 2011. "ORB: an efficient alternative to SIFT or SURF." *Proceedings of the 2011 International Conference on Computer Vision.* Barcelona. 2564-2571.

Rudolph, Tobias, Lars Ebert, and Jens Kowal. 2010. "Comparison of three optical tracking systems in a complex navigation scenario." *Computer Aided Surgery* 15 (4): 104-109.

Schwab, F J, A Patel, J Ungar, Pierre Farcy, and V Lafage. 2010. "Adult spinal deformity - postoperative standing imbalance: how much can you tolerate? An overview of key parameters in assessing alignment and planning corrective surgery." *Spine* 35 (2): 219-226.

Silikal America. 2017. *Operating Room Flooring*. Accessed July 20, 2017. http://www.silikalamerica.com/operating-room-flooring.html.

Singer, Gordon. 2005. "Occupational radiation exposure to the surgeon." *Journal of the American Academy of Orthopaedic Surgeons* 13 (1): 69-76.

Srinivasan, Dushyanth, Khoi D Than, Anthony C Wang, Frank La Marca, Page I Wang, and Thomas C Schermerhorn. 2014. "Radiation safety and spine surgery: systematic review of exposure limits and methods to minimize radiation exposure." *World Neurosurgery* 82 (6): 1337-1343.

Suhm, Norbert, Paul Mueller, Urs Bopp, Peter Messmer, and Pietro Regazzoni. 2004. "The MEPUC concept adapts the C-arm fluoroscope to image-guided surgery." *Injury: International Journal of the Care of the Injured* 35 (1): 120-123.

Szeliski, Richard. 2011. Computer Vision: Algorithms and applications. London: Springer.

Tan, Tze-Woei, Jeffrey A Kalish, Naomi M Hamburg, Denis Rybin, Gheorghe Doros, Robert T Eberhardt, and Alik Farber. 2012. "Shorter duration of femeral-popliteal bypass is associated with decreased surgical site infection and shorter hospital length of stay." *Journal of the Americal College of Surgeons* 215 (4): 512-518.

The Motley Fool. 2016. Accessed June 29, 2017. https://www.fool.com/investing/2016/10/27/teslamotors-inc-is-using-nvidia-corporations-driv.aspx.

Tian, Wei, Cheng Zeng, Yan An, Chao Wang, Yajun Liu, and Li Jianing. 2017. "Accuracy and postoperative assessment of pedicle screw placement during scoliosis surgery with computer-assisted navigation: a meta-analysis." *The International Journal of Medical Robotics and Computer Assisted Surgery* 13 (1).

Todesca, Alessandro, Luca Garro, Massimo Penna, and Jacques Bejui-Hugues. 2017. "Conventional versus computer-navigated TKA: a prospective randomized study." *Knee Surgery, Sports Traumatology, Arthroscopy* 25 (6): 1778-1783.

Tomasi, C, and R Manduchi. 1998. "Bilateral filtering for gray and color images." *Proceedings of the 1998 IEEE International Conference on Computer Vision.* Bombay.

- Touchette, Michele. 2017. "Artificial x-ray imaging system (AXIS) Design and evaluation on c-arm performance in operating room and educational settings (Masters thesis)." Vancouver.
- Tsalafoutas, Ioannis A., Virginia Tsapaki, Alkiviadis Kaliakmanis, Spiridon Pneumaticos, Fotis Tsoronis, Elias D. Koulentianos, and George Papachristou. 2008. "Estimation of radiation doses to patients and surgeons from various fluoroscopically guided orthopaedic surgeries." *Radiation Protection Dosimetry* 128 (1): 112-119.
- Vidal, Christophe, Brice Ilharreborde, Steffen Queinnec, and Keyvan Mazda. 2016. "Role of intraoperative radiographs in the surgical tratment of adolescent idopathic scoliosis." *Journal of Pediatric Orthopaedics* 36 (2): 178-186.
- Wagner, T A, S M Lai, and M A Asher. 2006. "SRS surgeon members' risk for thyroid cancer: is it increased?" *41st Annual Meeting of the Scoliosis Research Society.* Monterey.
- Wang, Lejing, Joerg Traub, Sandro Michael Heining, Selim Benhimane, Ekkehard Euler, Rainer Graumann, and Nassir Navab. 2008. "Long bone X-ray image stitching using camera augmented mobile c-arm." In *Medical Image Computing and Computer-Assisted Intervention*, 578-586. Berlin: Springer.
- Wang, Lejing, Joerg Traub, Simon Weidert, Sandro Michael Heining, Ekkehard Euler, and Nassir Navab. 2010. "Parallax-free intraoperative X-ray image stitching." *Medical Image Analysis* 14: 674-686.
- Wu, Gaihong, Shuqiang Liu, Husheng Jia, and Jinming Dai. 2016. "Preparation and properties of heat resistant polylactic acid (PLA)/Nano-SiO2 composite filament." *Journal of Wuhan University of Technology - Material Science Edition* 31 (1): 164-171.
- Yoo, J., S. Schafer, A. Uneri, Y. Otake, A. J. Khanna, and J. H. Siewerdsen. 2013. "An electromagnetic "Tracker-in-Table" configuration for x-ray fluoroscopy and cone-beam CT-guided surgery." *International Journal of Computer Assisted Radiology and Surgery* 8: 1-13.
- Zatsiorsky, Vladimir M. 1998. Kinematics of Human Motion. Windsor: Human Kinetics.
- Zhang, Zhengyou. 2000. "A flexible new technique for camera calibration." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22 (11): 1330-1334.

Appendix A – Lighting Mounting System

The following annotated drawings of the lighting mounting system are included in this section:

- 1. Camera mount
- 2. Connector
- 3. Top light holster
- 4. Bottom light holster









Appendix B – Camera Mounting System

The following annotated drawings of the camera mounting system are included in this section:

- 5. Vertical plate
- 6. Mounting plate
- 7. Attachment wedge





