Algorithms for Advancing the Patient-Personalization of Preoperative Planning in Image-Guided Robotic Surgery

By

Michael Allen Siebold

Dissertation
Submitted to the Faculty of the
Graduate School of Vanderbilt University
in partial fulfillment of the requirements
for the degree of

DOCTOR OF PHILOSOPHY
in
Electrical Engineering

March 31, 2021

Nashville, Tennessee

Approved:

Robert J. Webster, III, Ph.D.
J. Michael Fitzpatrick, Ph.D
Robert F. Labadie, M.D., Ph.D.
Benoit M. Dawant, Ph.D.
Jack H. Noble, Ph.D.

ACKNOWLEDGMENTS

First and foremost, I would like to thank my Lord and Savior Jesus Christ who truly has left me, “utterly astonished, saying, ‘He has done all things well.’” (Mark 7:37) All things are made and made possible though Him, including all of the work in this dissertation. Second, I would like to thank my advisor, Bob Webster, for his continuous patience, advice, and mentoring throughout the long and circuitous journey that I have taken on my way to this degree. I have no doubt in my mind that without both his guidance and his patience, I would not have been able to finish this process.

Of course there are many others without whom this work would not exist. Throughout Wendell Berry’s Port William novels he beautifully describes what he calls “membership” and “succession.” “Membership” somewhat oversimplified is the kind of tight community that helps each other through thick and thin without keeping any record of debts and debtors. His description of “succession” is one of the relationship and help passed down from previous generations. I have been incredibly blessed to be a part of several “memberships” so far in my life, and I have been the recipient of a long line of ”succession” on my path to finishing this document and subsequent degree.

The MEDLab and CAOS Labs are certainly a membership and have been full of people past and present that have helped at a moment’s notice with everything from brainstorming research problems, to preparing meals and providing food for my family as we have walked through medical difficulties. I’ll never forget Neal Dillon, Loris Fichera, and Trevor Bruns not even asking if they could help, but just showing up to work one day with bags full of home cooked food for us when they knew it was a particularly stressful time. Then a few years later after mom had passed and dad was undergoing a prolonged set of treatments, everyone from both labs got together to send a massive gift card for grocery delivery since it was in the middle of the COVID-19 lockdown. For these reasons and so many others, I am and will be forever grateful to you all, Alan Kuntz, Andria Remirez, Art Mahoney, Bob Webster, Bryn Pitt, David Comber, Dominick Ropella, Ernar Amanov, Hunter Gilbert, Isuru Godage, James Ferguson, Jason Shrand, Josephine Granna, Joshua Gafford, Katy Riojas, Loris Fichera, Lou Kratchman, Margaret Rox, Max Emerson, Mike Fitzpatrick, Neal Dillon, Patrick Anderson, Patrick Wellborn, Philip Swaney, Raul Wirz, Ray Lathrop, Richard Hendrick, Rob Labadie, Smita De, Tayfun Efe Ertop, and Trevor Bruns.

I would like to thank my dissertation committee, Dr. Robert J. Webster, III, Dr. J. Michael Fitzpatrick, Dr. Robert Labadie, Dr. Benoit Dawant, and Dr. Jack Noble for the time and effort that they all put in with me on this work. I remember working with Dr. Fitzpatrick my first semester here and the patience that he showed explaining registration concepts to me and bringing me into the research projects that have culminated in Chapters 2 and 3. I would also like to thank the National Science Foundation, the National Institutes of Health, IBM, and Vanderbilt for supporting my research here at Vanderbilt. I would especially like to thank the NSF for awarding me their Graduate Research Fellowship which provided me the freedom to pursue several of the research topics in this dissertation in a more in-depth manner.

To start the outside of Vanderbilt gratitude, I would like to especially thank my parents, Bert and Sharon Siebold, for investing so much of themselves into homeschooling me. I think of our lessons constantly and I pray that I someday have the influence on my own children's lives that you both have had on mine. And though my stepmother, Debbie Siebold, has been in the family for a relatively short time, her kind, helpful, and generous spirit have already made an impact on my time here at Vanderbilt.

The years I have spent at Vanderbilt have not been easy personally. Between having four surgeries myself not to mention being hospitalized for pneumonia, both mom's and dad's cancer diagnoses, mom's passing, and the last four years spent in significant chronic pain after a car wreck, there have been many hard days. I've had so much trouble just walking from the car to the lab that my next door neighbor, who works in Medical Center East on the same floor as I did, didn't even recognize me one afternoon when she passed me on the way to the parking garage. I've even been offered a wheelchair ride from a sympathetic hospital transport employee who mistook me for a patient. All that to say, I have needed no small measure of help, and I have been amazed at what people have been willing to do to provide it.

I would like to thank Trevor Bruns for teaching me so much about astrophotography and sharing that hobby with me. There were so many good times designing our first star trackers, and later going out to set up telescopes under the cold night sky trying to figure out how to keep the equipment from fogging up and us from freezing. Both Kristen and Rhonda Miller have been such good friends and have walked with me through some of the darkest days I've faced so far in life. Your encouragement has truly meant the world to me. The entire Cooper family, Nathan, Jennifer, Zoe Grace, Brandon, Ethan, Matt, Alan, Lisa, and Stephen have truly formed a ”membership” and included me in such an amazing way. You all will never know the impact you have had on me over the years. Whether it's 3 am conversations about everything from theology to rewriting the plots of Star Wars movies, one of the many times I've needed a hand, or even just a distraction from pain, you've always been there. I am more grateful than I can express.

I would like to finish this section where I began, with family. Since childhood dad has inspired me to want to follow in his footsteps both in life and academics. Finishing this degree is one of the first steps of my career and I can only hope that mine has even a part of the impact that his has had on the lives of his students. Mom was always my encourager. So many times she reminded me why I was working so hard and helped me maintain perspective. She was not only my mother, but she, along with dad, was also my closest friend. Her diagnosis with pancreatic cancer in 2016 was a blow that none of us saw coming. I remember talking to Bob about the amount of work I was missing to help with her treatments here at Vanderbilt and he told me that I was doing exactly what I should be doing and that we would take care of catching up on work later. His attitude of putting people first allowed me to have so much more time with her in her final two years than I ever could have imagined when the doctor came into the room and said that he had found stage four cancer. This degree has been a lifelong goal of mine. Finishing it is such an incredible feeling. That being said, what I have valued most here at Vanderbilt is the extra time that I was given with mom largely because Bob Webster put people first. That is a gift I can never repay.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Figure Page

Chapter 1

Introduction

1.1 Motivation

Image-guided surgery provides surgeons, either robotic or human, the ability to incorporate patient-specific preoperative surgical plans intraoperatively in ways never before imagined. Virtual fixtures or "no-fly zones" can be created around surgical boundaries to prevent the surgeon from moving the robot into an unintended area, or they may be designed to assist the surgeon in following the preoperative plan. Augmented or virtual reality can give the surgeon visual access to subsurface features while performing complex operations. Autonomous robotic systems can preoperatively plan entire portions of a procedure based on three-dimensional medical images and then perform those plans in the operating room on the patient.

The goal of this dissertation is to advance patient personalization of preoperative surgical planning for image-guided robotic procedures. In particular, a method is proposed for generating patient-specific three-dimensional surgical tool path plans for autonomous robotic bone milling. Additionally, two algorithms are presented that allow a surgeon to mitigate the danger posed by registration error to specific critical structures throughout an image-guided procedure.

1.2 Background

1.2.1 Tool Path Planning in Autonomous Bone Milling

Tool path planning is a well studied problem in CNC machining [7–10]. Unfortunately, these planners are mainly designed for machining homogeneous materials where the design of the part may be modified to facilitate the planned tool path. Many of these planners are designed to provide the highest feed rate possible given known material and cutter properties. These types of path planner are not well suited for autonomous bone milling. Not only are the materials heterogeneous, but their densities vary from patient to patient. Workpiece fixation is an issue that is relatively easily solved in machining, but is not as easily or rigidly accomplished in surgery. One of the more common methods of securing the skull is to use a Mayfield Clamp. A Mayfield clamp, while effective at securing the patient's head, is relatively invasive and can lead to complications [11]. An added difficulty is that bone milling is often performed in close proximity to critical structures that, if damaged, could leave the patient permanently disabled [12]. This is to say that while extensive literature exists in CNC path planning, the application of that literature in surgery has been sparse because the goals of CNC machining and the goals of autonomous surgical bone milling are different.

Tool path planning for autonomous robotic bone milling has evolved significantly since 1994 when Taylor et al. [13] published a paper describing the ROBODOC system. This planner was initially designed to mill pockets that receive surgical implants for total hip arthroplasty (THA) though that would later include total knee arthroplasty (TKA) as well. The cutter (an end mill) is oriented parallel to the long axis of the implant. Bone is removed one layer at a time by moving the cutter along successive linear paths whose separation distance is determined in relationship to the diameter of the cutter. This kind of layer by layer milling is commonly referred to as 2.5D milling because it achieves a 3D pocket by partitioning the volume into a set of approximately planar subvolumes (layers). These subvolumes are then milled in an ordered sequence from highest to lowest.

Several years later the world would see its first force controlled surgical robot for otoneurosurgery, Federspil et al. [1]. While the path plans that were tested on this robot were all 2.5D plans (Figure 1.1), the force exerted by the robot for several different 2.5D strategies was measured and compared and an optimal strategy chosen. This ranking of the path planning strategies by the magnitude of force required was an important step forward

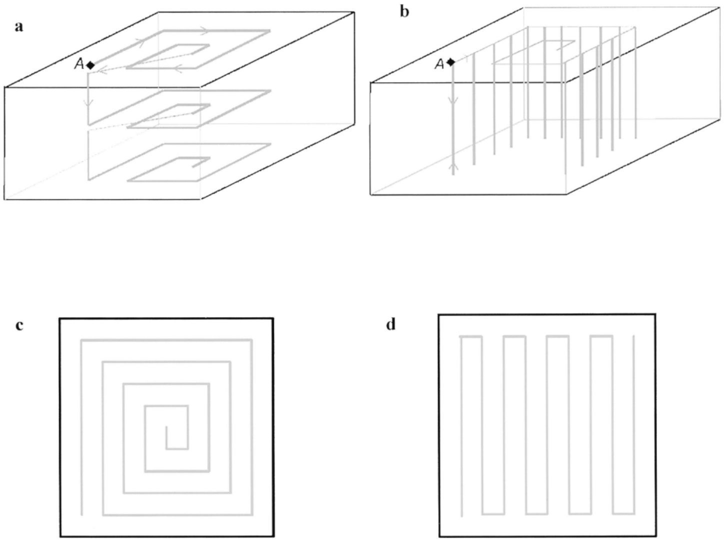

Figure 1.1: Several path planning strategies are seen in the figure. (a) horizontal milling, (b) vertical milling, (c) spiral horizontal milling, and (d) zigzag horizontal milling. Figure from Federspil et al. [1].

in designing tool path plans specific to the force requirements of the bone being milled. It should be noted that for this system the milling was carried out with a surgical spherical cutting burr and the tool shaft approximately perpendicular to the test surface throughout the milling process.

The next autonomous robotic milling path planner to be described is innovative in that it is not actually an image-guided system, Abraham et al. and Wolf et al. [14, 15]. This choice was made to simplify the preoperative planning. The path planner takes advantage of the known geometric properties of the knee implants' design to mill the bone surfaces to match those implants features. Another advance in this planner is that though it is a 2.5D planner, it uses a cell based decomposition to break the volume of bone to be removed into several smaller segments that more easily accommodate the simpler tool paths.

Sugita et al. presented a new innovation for total knee arthroplasty in the form of an image-guided tool path planner [2] whose path is planned under the constraint that the motion of the tool shaft is restricted to pass through a relatively small incision on the front of the knee to perform its 2.5D milling path (Figure 1.2). This purpose of the restriction

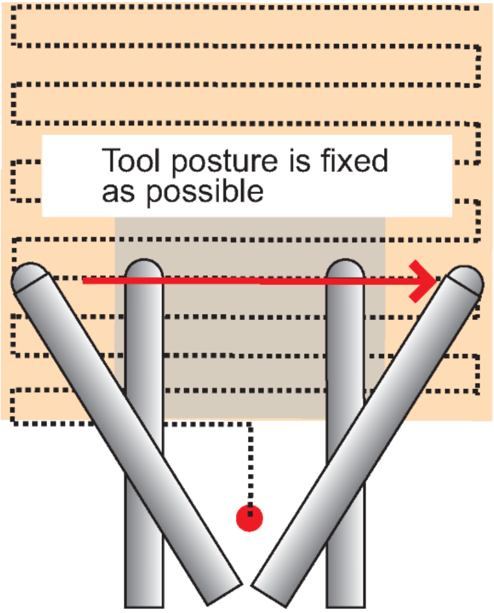

Figure 1.2: Minimally invasive total knee arthroplasty drilling strategy. From Sugita et al. [2].

of the motion of the tool shaft is so that it can be used in minimally invasive total knee arthroplasty. Note that though this planner restricts the movement of the shaft to pass through a small incision, it does not model the drill shaft or interactions between it and the patient.

At the same time as Sugita et al. a tool path planner that takes a different approach to 2.5D milling was created for the CRANIO robot by Cunha-Cruz et al. [3]. This planner was designed to perform an autonomous craniotomy (later the project would move to a cooperatively controlled system). Since the bone to be removed was a thin section of the skull with a typically strong curvature, the standard 2.5D path planner had to be altered. The planner is still a 2.5D planner, but its layers are curves approximating the exterior surface of the skull (Figure 1.3). These curves are projected down into the skull and are deformed to match the curvature of the interior of the skull as the planner approaches the bottom layer. This modification of the typical 2.5D layer by layer algorithm made milling volumes with inherent curvatures far more efficient.

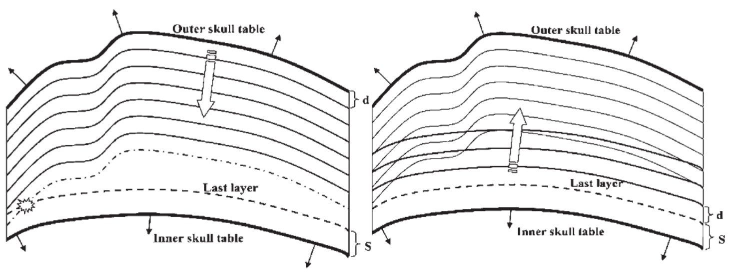

Figure 1.3: The deformation of the curves used by the path planner proposed by Cunha-Cruz et al. for milling a craniotomy. Figure from Cunha-Cruz et al. [3].

The limitation placed on a tool path planner by a 2.5D algorithm is that it must mill all of the bone on the current layer before proceeding further down into the pocket. This layer by layer approach makes checking for interactions between the tool shaft and the unmilled bone fairly simple, because the only bone that needs to be checked for interaction with the tool shaft is the boundary of the finished pocket. If the milled pocket is relatively shallow and the tool shaft stays close to perpendicular to the milled surface such interactions need not be modeled at all. Moving layer by layer also means that when milling arbitrarily shaped pockets, as is often required in robotic bone milling, the tool tip is prevented from following a small segment of pocket boundary that leaves the current plane, and the planner will have to come all the way back to that portion of the next layer down to remove that bone. Another limitation is that a planner may be able to keep more of the cutting burr engaged by leaving the current layer, but again it is prevented from doing so by the basic milling strategy. Revisiting the idea of force control, it is possible to model the surgical forces involved with bone milling as is often done in haptic simulations [16, 17]. If such a simulation was incorporated in a 2.5D tool path planner and a point requiring a lower force to access was available one layer below the current cutting depth, the planner could not choose that point. These are some of the motivations for creating a 3D path planner.

The only true 3D path planner for autonomous robotic bone milling is the inspiration for the work in Chapter 2. This planner was designed by Danilchenko et al. to perform an autonomous mastoidectomy [4]. In this 3D planner, every voxel targeted for removal

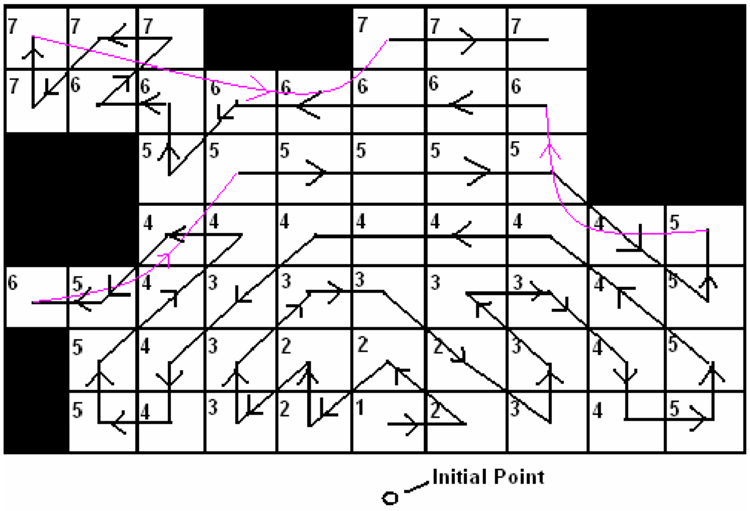

Figure 1.4: A visualization of a 2D slice of the "greedy" path planner's closest neighbor behavior described by Danilchenko et al. From Danilchenko et al. [4].

by the planner is ranked by its distance from the starting voxel. The tool then proceeds in a "greedy" manner from voxel to voxel always choosing to move to the next neighboring targeted voxel that is closest to the tool's starting position (Figure 1.4). An additional innovation is that this planner also relies on preprocessing the segmented medical images to account for the diameter of the spherical cutting burr used by the system. This preprocessing step removes the need for mathematically computing the offset from the edge of the planned pocket associated with the thickness of the cutter being used.

A preliminary version of the path planner presented in Chapter 2 appears in Dillon et al. [18].

1.2.2 Quantifying Risk to Critical Patient Structures in Image-Guided Robotic Surgery in the Presence of Rigid Registration Error

To use an example from the previous section, what if an autonomous bone milling robot needed to work in close proximity to structures that were vital to the patient's quality of life or survival? What assurance would the patient have that the structures would not be damaged during the procedure? These are the types of questions that the algorithms described in this background section seek to answer.

While computing the spatial error associated with various registration modalities has been extensively studied for the last several decades, relatively little work has been done to allow the surgeon to specify the safety level of specific critical structures in the planning stage of a surgery performed under image guidance. In fact, many times safety margins used in image-guided surgery are uniformly thick margins whose dimensions are based on intuition.

There have been two different types of answers given to questions posed at the beginning of this section. The first solution was proposed by Haidegger et al. [5], that models the error associated with an instantaneous position of the tip of a tracked tool in real time via Monte Carlo simulation (Figure 1.5). This error distribution is used to determine the probability that the tool tip is in violation of a virtual fixture put in place around a critical structure (the eighth cranial nerve, the acoustic nerve, during a vascular decompression to treat hemifacial spasm via a suboccipital approach). This result can be used to protect the critical structure by informing the surgeon (robotic or human) when the danger of violating that boundary reaches an unacceptable level. While an important first step toward specifying the safety level of critical structures during a procedure, this method is limited by only supplying instantaneous estimates of danger to the critical structures.

The second solution, extends the idea of a danger probability associated with a critical structure from an instantaneous event to a probability associated with the duration of the procedure. Noble et al. [19] created an algorithm that analyzes a linear bone drilling path in close proximity to critical structures in the middle and inner ear. The probability of the drill colliding with the critical structures is estimated by randomly varying the positions of the start and end of the proposed drilling trajectory. These varied plans are then checked to see if they collided with the critical structures. Probabilities of collision with the critical structures are computed from the collision counts. This is the first time a damage probability for critical structures has been created that was valid for the duration of the procedure.

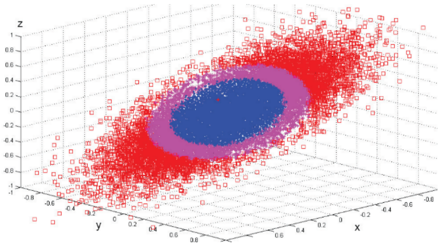

Figure 1.5: This figure from Haidegger et al. highlights the type of instantaneous probability of violating virtual fixtures their algorithm provides [5]. "The POI [Point of Interest] (tooltip) transformed to the coordinate space of the patient. Purple stars show where the overall RMS error is larger than 0.2 mm and red squares mark the region where the error is over 0.4 mm. The exact probability of the POI being beyond the VF [Virtual Fixture] is 0.438 and 0.214 for the 0.2 and 0.4 mm VF, respectively. Red dot shows the theoretical position, black dot represents the effect of the registration errors."

Building on these results, the algorithm created in Chapter 3 uses the combination of a theoretical result from the field of fiducial registration error and Monte Carlo simulation to create safety margins that envelop an arbitrarily shaped critical structure that are spatially varying and that guarantee safety to the critical structure for the duration of the entire surgery. After the publication of Chapter 3, that work was extended by Dillon et al. [6] to include several additional error sources pertinent to the specific application of robotic mastoidectomy performed by the Acoustic Neuroma Surgical Robot [18] (Figure 1.6). The addition of these error sources highlights the value of the algorithm by giving the surgeon the ability to set a statistical guarantee of safety to individual structures in a patient before the surgery takes place.

1.3 Dissertation Contributions

The goal of this dissertation is to advance patient personalization of preoperative surgical planning for image-guided robotic procedures. This advancement is achieved by focusing on two specific problems: (1) tool path planning for autonomous robotic bone milling, and (2) the danger posed to specific critical structures by the registration error present in image-guided robotic procedures over the duration of an invasive procedure.

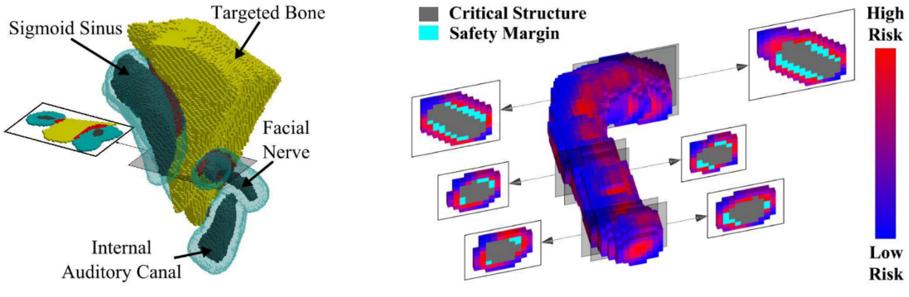

(a) (b)

Figure 1.6: (a) The final safety margins for several critical structures. Portions of the original target volume (in red) are removed where the safety margins overlap the volume originally identified for removal by the surgeon. (b) A schematic of the risk of drilling points surrounding the facial nerve and its current safety margin at one iteration of the algorithm. At each iteration, a percentage of the highest risk voxels are added to the safety margin to bring the probability of preserving the structure closer to the desired threshold. From Dillon et al. [6].

1.3.1 Tool Path Planning Algorithm for Autonomous Robotic Bone Milling

Autonomous robotic bone milling is a common application of surgical robotics because of the inherent rigidity of the bone and the repeatability of robotic manipulators. Up to this point, the primary methods for planning tool path trajectories have been layer by layer 2.5D planners [1–3, 13, 15]. Path planners not limited to removing bone in a layer by layer fashion have the potential to incorporate many additional patient-specific factors in the shape of the planned path. One 3D autonomous bone milling tool path planner has been demonstrated [4]. The implementation of a 3D autonomous bone milling algorithm significantly complicates the interactions between the tool shaft and the unmilled tissue during the procedure.

Chapter 2 of this dissertation is inspired by the path planner presented in [4] and contributes new functionality in two separate areas. The first and primary contribution presented in Chapter 2 is the ability to model and incorporate tool shaft and cutting burr interactions with unmilled bone tissue into the creation of a true 3D path planner. A secondary contribution presented in Chapter 2 is the ability to easily vary the cutting depth of the planned path in a true 3D path planner. Control over the cutting depth is possible because of a novel reapplication of the concept of a “supervoxel” from image and video processing literature [20–22]. The ability to control the cutting depth of such an algorithm decreases the operating time of the robot and increases the speed of the 3D path’s computation such that the plan can be created intraoperatively.

1.3.2 Creating Spatially Varying Margins that Encapsulate Critical Structures Within the Patient via Analysis of Registration Error

Previous work has analyzed the instantaneous danger posed to a critical structure by a tracked tool in an image-guided system in the presence of registration error [5]. Additionally, the danger posed by bone drilling in the presence of registration error to a given critical structure has been analyzed for entire linear bone drilling paths [19]. This dissertation builds on these results by generalizing these safety analyses to extend to a critical patient structure of arbitrary shape, throughout the duration of the procedure. Chapter 3’s primary contribution is to create the first spatially varying safety margins enveloping critical structures based on a statistical analysis of fiducial point-based registration error. If preserved throughout the procedure, these safety margins provide a statistical guarantee that the contained structures will remain undamaged. Chapter 4’s primary contribution is to create statistically guaranteed safety margins with registration modalities other than fiducial point-based registration. The algorithm was simulated with surface based iterative closest point registration. An additional contribution of Chapter 4 is to reconceptualize the safety margins to represent tissue that must be removed. The specific application is a cancerous tumor, and the safety margin defines the volume of tissue that must be fully removed to ensure that the tumor is fully removed, given registration uncertainty.

Chapter 2

A Robust Tool Path Planner for Autonomous Robotic Bone Milling Based on Preoperative Imaging

2.1 Introduction

Bone milling is an excellent example of a class of procedures that benefits greatly from advanced patient personalization. The volumes of bone to be removed in a wide variety of procedures from craniotomy implants, to mastoidectomies, to knee and hip replacements all have aspects of the procedure that can be custom tailored to the individual patient. Of course the volumes of bone to be removed must be planned for the individual patient, and these unique pocket dimensions present their own set of challenges. Beyond the shape of the pocket such considerations as minimizing milling forces around critical structures uniquely placed in each patient, modeling heat buildup around nerves embedded in the bones to avoid damage, choosing a tool path that optimizes cutter orientation to improve the efficiency of the milling, and simply finding a shorter tool path that a standard 2.5D zigzag planner will yield and save the patient time in the operating room. These patient-specific considerations and many other possibilities are able to be accounted for by the algorithm presented in this chapter.

Bone milling is also an ideal target for surgical robotics because of the rigidity of the bone. In less rigid tissue, preoperative plans must be adapted as the tissue deforms. Compensating for this tissue deformation is a relatively difficult task. The rigidity of bone eliminates the need for deformation modeling and greatly simplifies preoperative planning. This rigidity is a primary reason that bone milling was one of the first applications of surgical robotics [23–25].

Note that a distinction is being drawn between bone milling and bone drilling. Bone milling is being defined as the removal of a multi-faceted volume of bone, which is distinct from bone drilling, the removal of bone tissue along a direct linear path. It should also be noted that this bone milling is often performed via a surgical drill; therefore, references will be made to a surgical drill performing bone milling.

Knee [26], hip [27], skull and skull base [28, 29] surgeries have all benefited from the accuracy and repeatability that robotic bone milling provides. For example, in total knee arthroplasty (TKA) Siebert el al. [30] found that the average difference between the planned and achieved tibio-femoral alignment was 0.8° for robotically treated patients and 2.6° in the standard of care group; however, it remains unclear how significant a benefit the additional accuracy is to long term patient outcomes [31]. In total hip arthroplasty (THA), Wu et al. performed a comparison between robotic and manual preparation of the femoral cavity [32]. In this study, they found that only 61% of the implant is in contact with the bone in their manually performed procedures, and that the maximum gap between the implants and the manually broached hole was 2.97 mm with an average gap of 0.77 mm. In comparison, their robotically performed procedures revealed that the average gap between implant and cavity was only 9% of that left by the manual group. In a craniotomy, one benefit of robotic procedures is the ability to model the volume of the removed skull tissue so that an accurate, well-fitting implant may be manufactured before the surgery takes place [33]. In mastoidectomy, a skull base surgery, there are several critical structures embedded in close proximity to the volume of bone that is to be removed. Damage to these structures (e.g. facial nerve, chorda tympani, sigmoid sinus, etc.) could range from permanent facial paralysis to severe bleeding [12]. Robotic bone milling allows the removed volume of bone to be planned such that it does not endanger the patient-specific critical structures in close proximity to the relevant bone tissue [34].

The first commercialized orthopedic surgical robot was ROBODOC (Curexo, Inc., Seoul, South Korea) [13] which ultimately was used clinically for both hip and knee replacement surgeries. ROBODOC performed the bone removal autonomously based on a preoperative plan. Computer Assisted Surgical Planning and Robotics (CASPAR) was a commercial system that also worked in knee and hip surgery performing bone milling autonomously. CASPAR was a part of a German company URS Ortho GmbH and Company KG [35]. MAKO's RIO system is a clinically used cooperatively controlled robot designed for hip and knee surgery. RIO enforces constraints on the surgeon's tool path in real time so that unintended bone is not removed. In this kind of surgeon/robot partnership the surgeon performs the procedure while the robot ensures the safety of the patient and the accurate adherence to the planned removal of bone. Other robotic bone milling systems for hip and knee surgery include Mini Bone-Attached Robotic System for joint arthroplasty (MBARS) [15], MBARS2 [36], and Active Constraints Robot (ACROBOT) [37].

Bone milling robot systems also exist in skull and skull base surgeries such as craniotomy and mastoidectomy. The Robot and Computer-Assisted Craniotomy (CRANIO) system [3] is a system designed to perform a robotically assisted craniotomy. This system started as an autonomous bone milling system and progressed to a cooperatively controlled milling system. The Acoustic Neuroma Surgical Robot (ANSR) robot [18] is a compact, bone-attached, 4 degree of freedom robot that autonomously performs a mastoidectomy. The system implemented by Xia et al. [38] is a cooperatively controlled robot designed for skull base surgery. The systems described by Lim et al. [39] and Danilckenco et al. [40] are autonomous systems for mastoidectomy. Danilckenko's work is a precursor to the ANSR system described by Dillon et al. [18].

There are two distinct approaches to bone milling. In the first approach, the robot autonomously executes a preplanned trajectory based on the preoperative imaging. ROBODOC, CASPAR, MBARS, and ANSR all fall into this category. The second approach utilizes preoperative imaging to create virtual fixtures or "no-fly zones" [41, 42]. The robot then is driven by a surgeon. In this case, the robot does not have a planned trajectory and its added safety is provided by enforcement of spatial boundaries that do not allow the surgeon to remove unintended bone tissue. MAKO RIO, and ACROBOT fall into this category. CRANIO started as an autonomous preplanned robot and has transformed into a cooperative control system.

Many different aspects of bone milling have been explored. Several examples include, methods to determine what type of bone is being milled in real time [43]. There are methods to detect breakthrough into non-bony tissue [44]. Optimal milling angles have been determined from preoperative data [45]. Temperature and force related to bone milling have been modeled from preoperative imaging [46].

Limited detail has been published about planning the trajectory that the mill traverses through the removed bone tissue. The most thorough treatments are given in [1, 2, 15]. In Sugita et al. [2] the path plan of a 5 axis floor standing robot is planned for a minimally invasive knee surgery. Most of the complication involved in the system relates to the system having to operate through a relatively small incision while achieving the desired trajectory. Note that this planner does not model the drill shaft or interactions between it and the patient. Federspil et al. [1] tested several 2.5D approaches with force control on an industrial robot. Of the layer by layer methods tested a "spiral" trajectory minimized the force applied to the cutting burr and was deemed the best tested method. It should be noted that in Federspil et al. the milling was carried out with a spherical drill tip and the tool shaft was approximately perpendicular to the test surface. Wolf et al.'s [15] MBARS robot utilizes knowledge of the implant's design to aid in planning the milling path. A cell decomposition technique is used to partition the volumes of bone to be removed into subvolumes. Then, a layer by layer approach is used to remove the tissue in the partitioned cells.

Planning strategies are briefly described in system overview papers for an early version of ROBODOC [13], CRANIO [3], and ANSR [18]. The description in [13] indicates that it followed a zigzag plan where parallel lines are traversed until no material is left on the first layer. This process is repeated on the next layer down and so forth until the desired volume of bone is removed. It is a reasonable assumption that a more involved tool path planner was implemented as the planning system, Orthodoc [47], matured, but they did not publish the details of an updated tool path planner. It is not a large step to assume that plans were likely developed for each specific implant used by the system since the characteristics of the implants are known and are independent of the patient. The path planner described for the CRANIO robot is a 2.5D projection of a tracing of the surface of the bone that is to be removed. This tracing deforms as the trajectory reaches the interior wall of the skull [3]. The path planner described for the ANSR robot [18] is a preliminary version of the planner detailed in this work.

Tool path planning is a well studied problem in CNC machining [7–10]. CNC planners are generally not well suited for bone milling. A large number of these planners focus on generating the highest feed rate possible given known material and cutter properties. In bone milling, not only are the materials heterogeneous, but they vary from patient to patient. Workpiece fixation is an issue that is relatively easily solved in machining, but is not as easily or rigidly accomplished in surgery. One common and effective method of securing the skull, a Mayfield Clamp, is also relatively invasive and can lead to complications [11]. An added difficulty is that bone milling is often performed in close proximity to critical structures that, if damaged, could leave the patient permanently disabled [12].

This chapter presents a voxelized framework novel to robotic bone milling that enables the simple implementation of a wide range of tool path planning algorithms using either 4 or 5 DOF manipulators. This framework models not only the drill's cutting burr, but also the drill's shaft to avoid potential interactions between the shaft and bone tissue that would impact the ability of the system to accurately follow the preoperative plan. The framework relies heavily on the re-interpretation of a "supervoxel". Supervoxels are a common concept in image segmentation [20–22]. The concept is extended in this chapter to greatly increase the efficiency of the path planning framework. Finally, this chapter details a true 3D surgical bone milling path planner that is easily customized for the procedure or the individual patient's needs.

The algorithm is demonstrated on several mastoidectomy procedures. It was implemented and tested on the 4 DOF Acoustic Neuroma Surgical Robot (ANSR) [18]. The algorithm may be applied to any surgical milling procedure whose planning is based on voxelized preoperative images.

2.2 Developing a Voxelized Framework

The voxelized framework requires a 3D medical image (e.g. a CT scan) that has been segmented to highlight the different volumes that concern the planning process. These segmented regions are as follows: bone that is to be preserved (protected bone), bone that is to be removed (targeted bone), and air. An example of these segmentations can be seen in Figure 2.1. The other pieces of information required by the path planner are the dimensions of the drill shaft, the diameter of the spherical drill tip, and the dimensions of the voxels. The voxelized framework is built around two parallel simulations. One simulation is carried out in a modified version of the segmented 3D medical scan.

The modified simulation is used to plan the sequence of Cartesian positions that the spherical cutting burr will traverse as it mills the targeted bone. The modified simulation takes place in the modified volume because the modified volume has been designed such that the center of the spherical cutting burr may be positioned at the center of any targeted bone voxel in the modified volume without the spherical cutting burr removing any untargeted bone voxels.

The unmodified simulation is performed in parallel with the modified simulation. It takes place in the unmodified volume, that is to say the original segmented 3D medical image. The unmodified simulation is used to determine the appropriate drill shaft orientations, if any exist, that do not collide with unmilled bone tissue for a cutting burr position chosen in the modified simulation. This separate simulation is necessary because the diameter of the spherical cutting burr will typically be much larger than the diameter of the drill shaft. The modified volume accounts for the diameter of the spherical cutting burr. Therefore the much smaller targeted voxel segmentation present in the modified volume

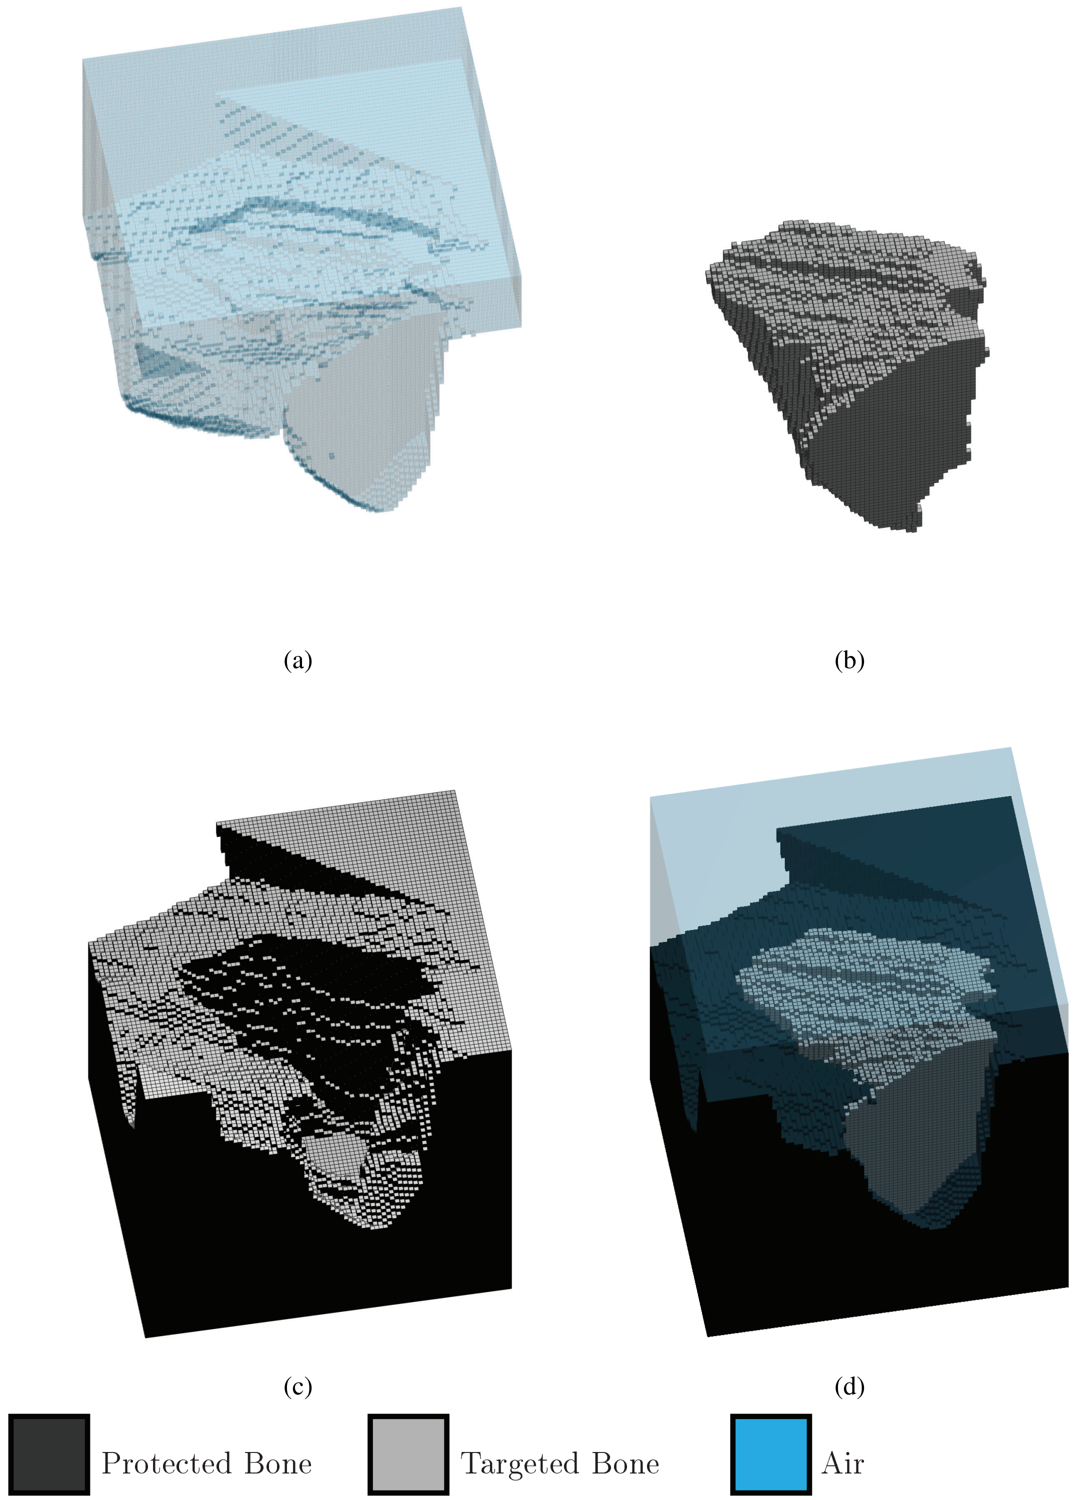
(a) (b)

(c) (d)

Figure 2.1: 3D visualizations of a segmented CT scan are shown in the figure. Part (a) is air. Part (b) is targeted bone. Part (c) is protected bone that remains unmilled throughout the procedure. Part (d) combines all of three of the segmentations into a single visualization.

does not allow accurate calculation of the angles by which the targeted bone voxels may be accessed. Table 2.1 contains a glossary of frequently used terms.

Term	Definition
Immediately adjacent	All voxels in one segmentation (e.g. Targeted Bone, Protected Bone, etc.) that share a vertex with any voxel in the other segmentation (26-connected) are considered to be immediately adjacent to the other segmentation.
Supervoxel	A supervoxel as used in this algorithm is a cubic structuring element with odd dimensions (e.g. 3 x 3 x 3 or 5 x 5 x 5) whose center is coincident with the center of the cutting burr. During path planning as the supervoxel moves through the unmodified volume, targeted bone voxels that overlap the non-central voxels within it are counted as milled. Targeted bone boundary voxels in the modified volume are counted as milled only when they are overlapped by the central voxel of the supervoxel. In the modified simulation the center voxel of the supervoxel may not occupy a protected bone voxel, but non-central voxels within the supervoxel may occupy voxels labeled as protected bone.
Protected bone	This is bone that is to be preserved throughout the milling procedure. The protected bone contains critical structures whose damage would cripple the patient.
Targeted bone	This is bone that is to be removed by the milling procedure. In the modified volume, this bone is marked as removed after contact with any portion of the supervoxel. In the unmodified volume, this bone is marked as removed after contact with any portion of the cutting burr's structuring element.

Targeted bone boundary	This is targeted bone that is immediately adjacent to the protected bone. Voxels assigned this label are only marked as removed by the center of the supervoxel. This segmentation does not exist in the unmodified volume.
Unmodified volume	This 3D array is the original segmentation of the 3D medical scan. This voxelized region has been segmented with values corresponding to protected bone, targeted bone, and air.
Modified volume	This 3D array is a copy of the unmodified region that has been resegmented. The volume of the targeted bone segmentation has been reduced by a morphological erosion using a structuring element representing the spherical cutting burr. The voxels removed from the targeted bone segmentation have been added to protected bone segmentation in the modified volume. The modified volume also contains a new label differentiating between targeted bone and targeted bone that is immediately adjacent to the protected bone. The generation of this volume is described in Section 2.2.1.
Modified simulation	This simulation is used to determine the positions of the path as it is developed. This simulation takes place in the modified volume.
Unmodified simulation	This simulation is used to determine the permissible orientations of the drill shaft. In the algorithm, this simulation is also useful for developing a more detailed fitness metric as described in Section 2.3.5.

Table 2.1: A glossary of terms that are frequently used in this chapter.

2.2.1 Preprocessing the Segmented 3D Medical Image: Creating the Modified Volume

This framework is designed to perform the milling with surgical drills already in use in the operating room. These surgical drills use spherical cutting burrs to remove bone tissue. The diameter of the chosen burr impacts the path planner. To illustrate this situation, consider a plan created via a CT scan whose voxels represent 0.4 mm cubes. If this milling plan is created for a 5 mm diameter spherical cutting burr, each time the cutting burr is centered on a single voxel, it will cover many additional voxels. This situation can be seen in Figure 2.2.

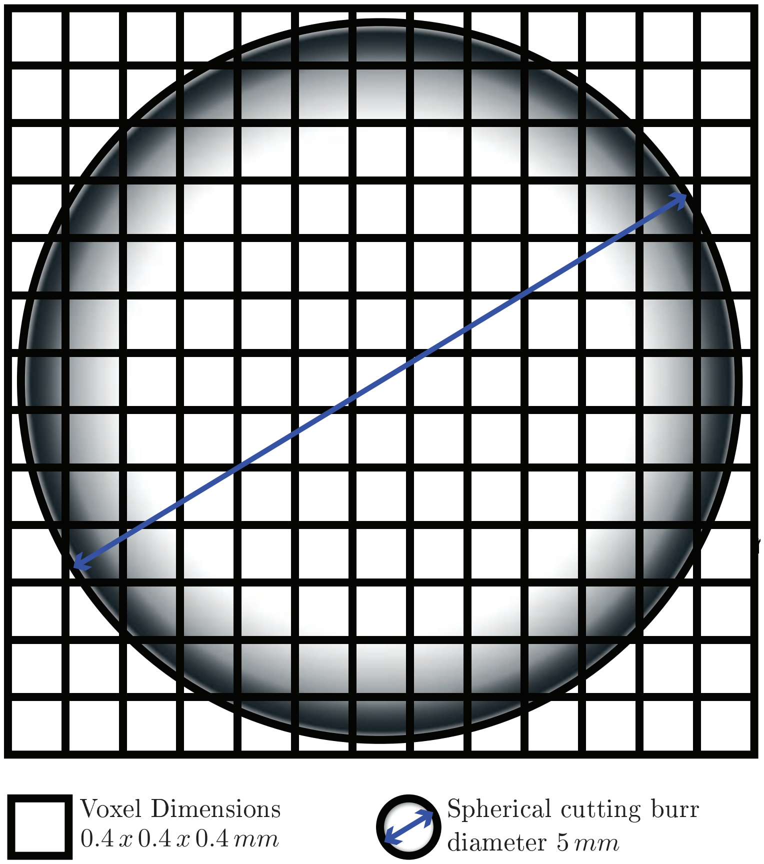

Figure 2.2: A two dimensional comparison between the size of an example spherical cutting burr and voxel size is shown in the figure. The cutting burr shown is 5 mm in diameter. The voxels shown are 0.4 mm cubes. The large number of voxels covered by the cutting burr in this example illustrates why the size of the cutting burr needs additional modeling, both for the preparation of the modified region and in the simulated milling in the modified simulation.

This relationship proves problematic when the cutting burr is commanded to a location near the edge of the volume of bone that is segmented to be removed. If the center of the cutting burr is commanded to a position centered on a voxel near the border of the protected bone, the radius of the spherical cutting burr will reach into the protected bone even though the center of the spherical cutting burr has not left the targeted bone segmentation. This removal of protected bone happens because the diameter of the cutting burr is larger than the voxels that partition the volume. To account for this relationship, a voxelized structuring element that encompasses the entire burr is created. Such a structuring element would contain all voxels overlapped by the circle in Figure 2.3a. A morphological erosion operation is performed on the targeted bone (Figure 2.3b). After this erosion is complete, the center of the cutting burr may be safely commanded to the centers of the voxels remaining following this erosion. Note that there is an additional region in Figure 2.3b, Eroded Target: Unreachable Bone. These voxels represent bone that cannot be removed by a burr of the chosen diameter without removing bone that is to be preserved. From this point on, the eroded target and the eroded target: unreachable bone are considered protected bone, and the remaining smaller targeted bone is the region containing voxels that are to be visited by the center of the spherical cutting burr.

After the targeted bone has been reduced to compensate for the cutting burr's diameter, the modified targeted bone is dilated into the air using the same structuring element as before. The purpose of this dilation is to compensate for the cutting burr's radius when it is milling the top surface of the bone. Without this dilation the first voxel targeted would require the cutting burr to be cutting bone with its entire radius for every targeted bone voxel immediately adjacent to the air. This is illustrated in Figure 2.4. The new volume created by this preprocessing is the modified volume and will form the basis of the modified simulation that determines the cutting burr's positions throughout the path. The modified volume is shown in Figure 2.5.

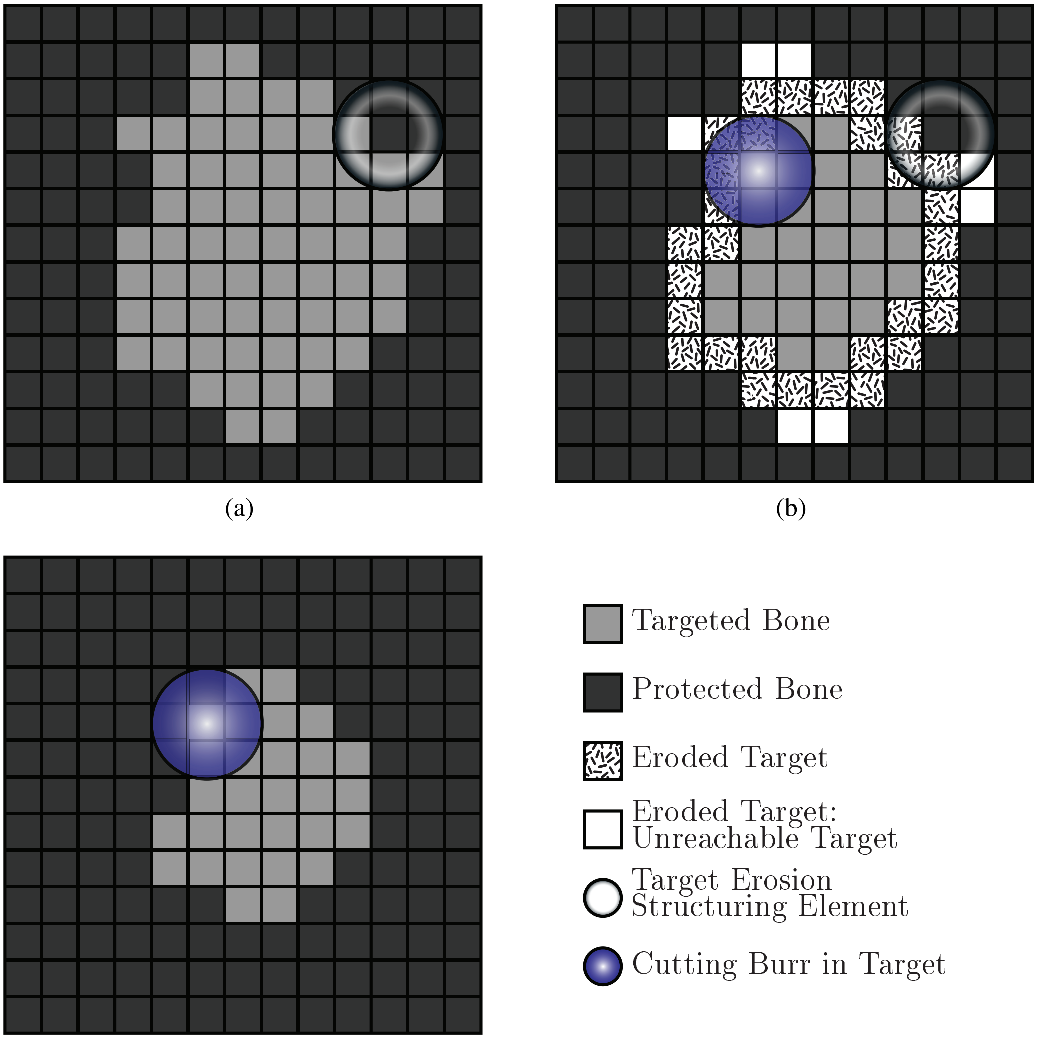

(a) (b)

(c) (d)

Figure 2.3: Part (a) shows the initial segmented medical scan. Part (b) shows the process of modifying the segmented scan so that each targeted voxel may be visited by the drill's center. Part (c) shows the final modified volume. Part (d) is a legend for the different regions highlighted in parts (a) - (c). It should be noted that Eroded Target and Eroded Target: Unreachable Bone are removed from the modified target the same way. The distinction is that the Eroded Target: Unreachable Target region represents voxels that could never have been reached without removing protected bone because of the dimensions of the cutting burr.

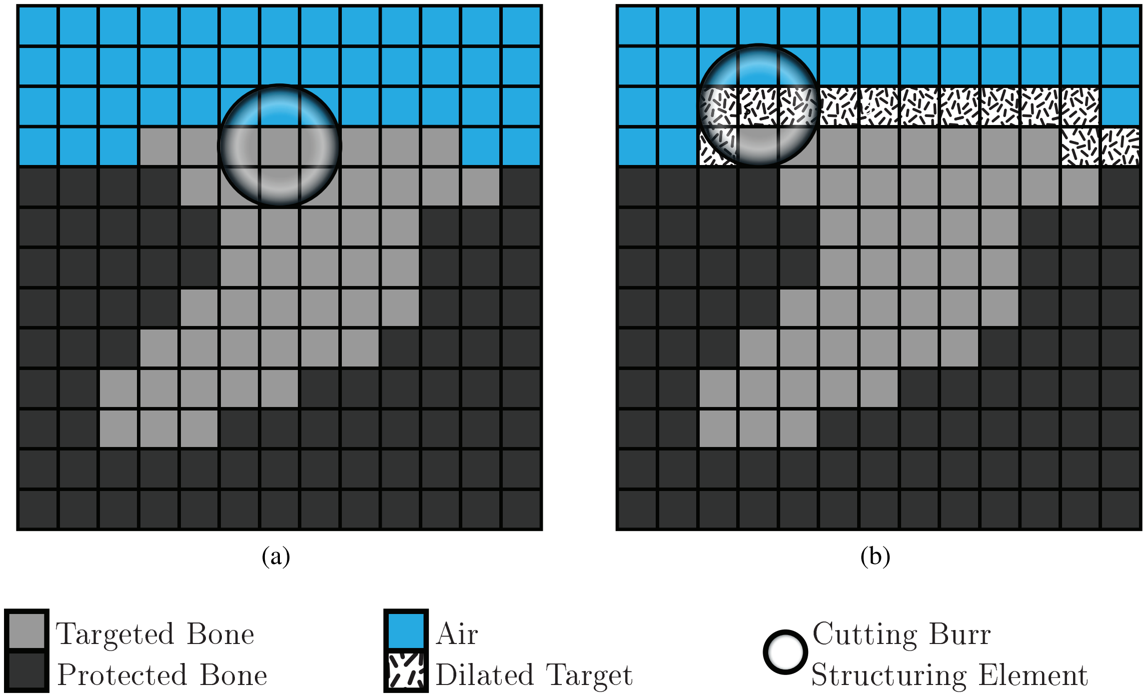

Figure 2.4: Part (a) shows a two dimensional vertical slice of the modified region. Part (b) shows the same slice after the target has been dilated into the air region.

2.2.2 Modeling the Cutting Burr in the Unmodified Simulation: Structuring Element

This portion of the framework takes place in the unmodified volume, and is a part of the unmodified simulation. The model of the cutting burr in the unmodified volume is simply the structuring element that was used to create the modified volume. This structuring element may be placed at any voxel position within the unmodified volume and any voxel that the structuring element covers is considered to be milled.

2.2.3 Modeling the Drill Shaft and Collision Detection

This portion of the framework takes place in the unmodified volume, and is a part of the unmodified simulation. The drill shaft of the surgical drill chosen for use in the Acoustic Neuroma Surgical Robot has three distinct diameters. There is the barrel of the drill, the chuck or gripper, and the shaft which leads up to the cutting burr. The modeling of the

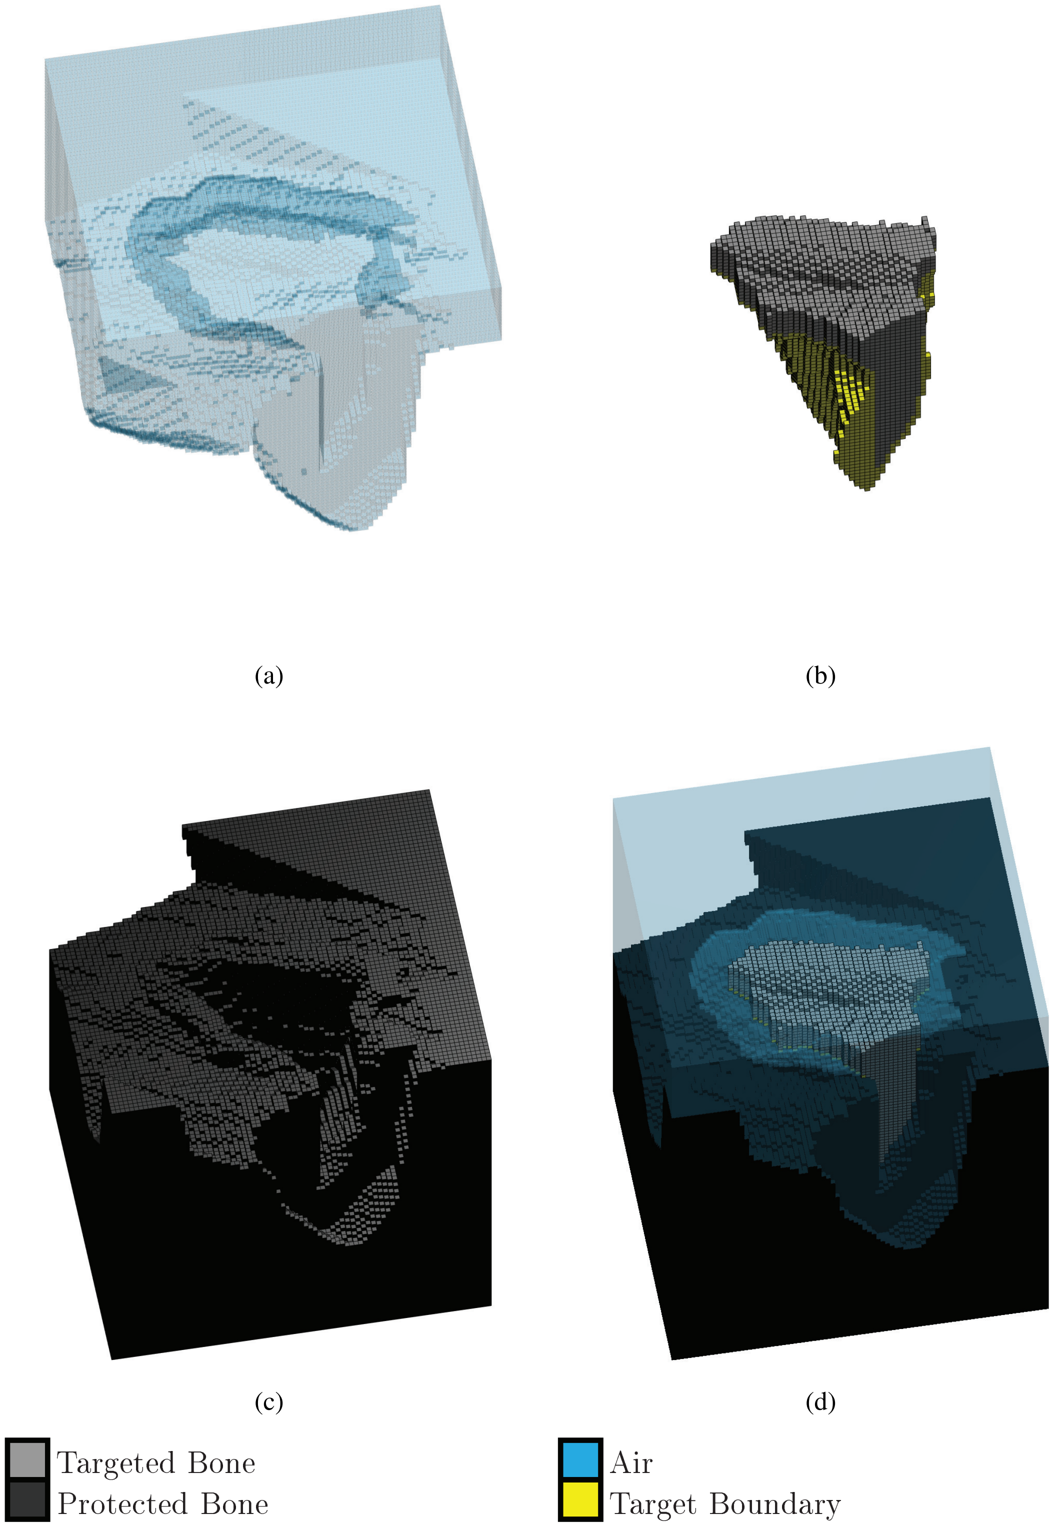

Figure 2.5: 3D visualizations of a preprocessed segmented CT scan are shown in the figure. Part (a) is air. Part (b) is targeted bone including the targeted bone boundary. Part (c) is protected bone that remains unmilled throughout the procedure. Part (d) combines all three of the segmentations into a single visualization.

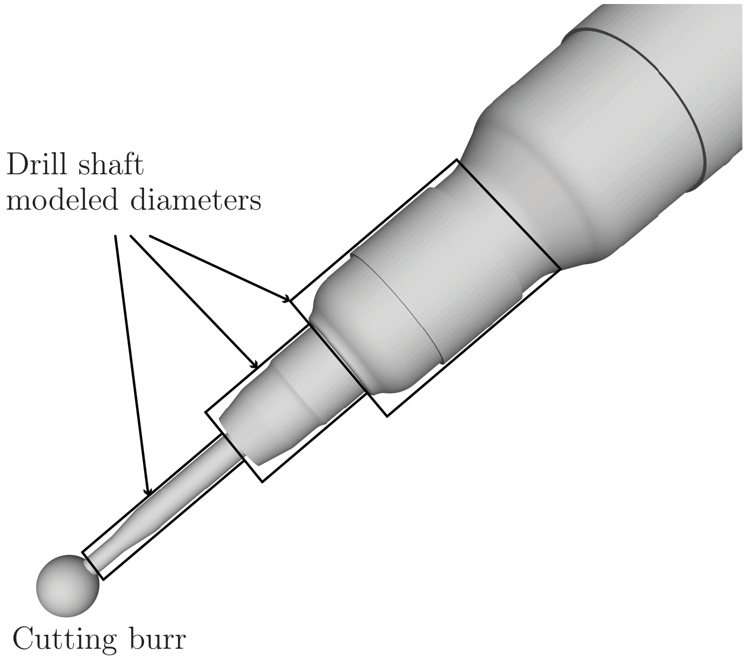

Figure 2.6: The figure shows the cutting burr at the tip of the drill shaft mounted in the drill. The combination of these three sections of the drill will be referred to as the drill shaft. Because of the restrictions of the robot used for experimental validation, additional portions of the shaft do not need to be modeled because they cannot physically interfere with the procedure.

drill and the collision detection methods outlined in this section are the only differences between the voxelized framework's 4 DOF and 5 DOF variations.

In the 4 DOF version of the framework, a simplifying approximation can be made. The rotational axis is aligned parallel to one of the scan's axes. This orientation choice means that the drill axis is always parallel to one of the voxel planes. The information stored in the several-voxel-thick region containing the drill shaft can then be collapsed into a single voxel thickness.

This collapsing of the slices is accomplished by comparing the voxels along the dimension being reduced. If any such voxel is unmilled bone (either protected or targeted), the voxel in the collapsed slice is counted as unmilled bone, if all voxels along the dimension being reduced are segmented either as air, or as having been previously milled, those voxels are counted as air in the collapsed slice. This method of collapsing the unmilled or protected bone voxels and air or milled bone voxels from several slices of the modified volume into a single slice is illustrated in Figure 2.7.

Line segments that correspond to sides of the drill that exist within the collapsed slice are then generated. These line segments are generated for each of the drill's three radii. These line segments that represent the sides of the drill for each of the three drill shaft's radii are shown in Figure 2.6. The resulting line segments are checked for intersections with unmilled bone tissue in the collapsed slice. The line segments are swept through the possible angles that the robot could achieve and a range of accessible angles for the tested voxel location is stored in the path plan. A sample of a modeled drill shaft being checked against a collapsed slice of the CT scan is illustrated in Figure 2.7.

In the 5 DOF version of the framework, it is necessary to use a much more computationally expensive simplification of the unmodified volume. A tree based spatial partitioning method was chosen because the boundaries of the volume of partially milled bone being checked for interactions with the shaft of the drill are quite complex. Since the unmodified volume is already partitioned into uniform cubes (voxels), an octree was chosen to partition the unmodified volume (Figure 2.8). An octree may be used as a method for partitioning $\mathbf{R}^3$ space for efficient searches. In principal, an octree starts out with the smallest box that contains all points of interest. This box is divided into eight equally sized "child" boxes, hence the name "octree". These subdivisions continue recursively until the dimensions of the smallest child boxes correspond to the desired level of precision. A modification to the standard octree was necessary. The original parent box needed to be sized and positioned such that its finest resolution child boxes would be voxel sized cubes concentric with the segmented voxels. An additional consideration was made because of the size of the octree. Even a partitioned search for collisions with the number of voxels relevant to the milling was impractical. Instead of modeling the entire volume of both protected and targeted bone tissue in the octree, a one voxel thick surface of bone (both protected and targeted) is modeled. In this way, every interface that the drill shaft could collide with is modeled and many

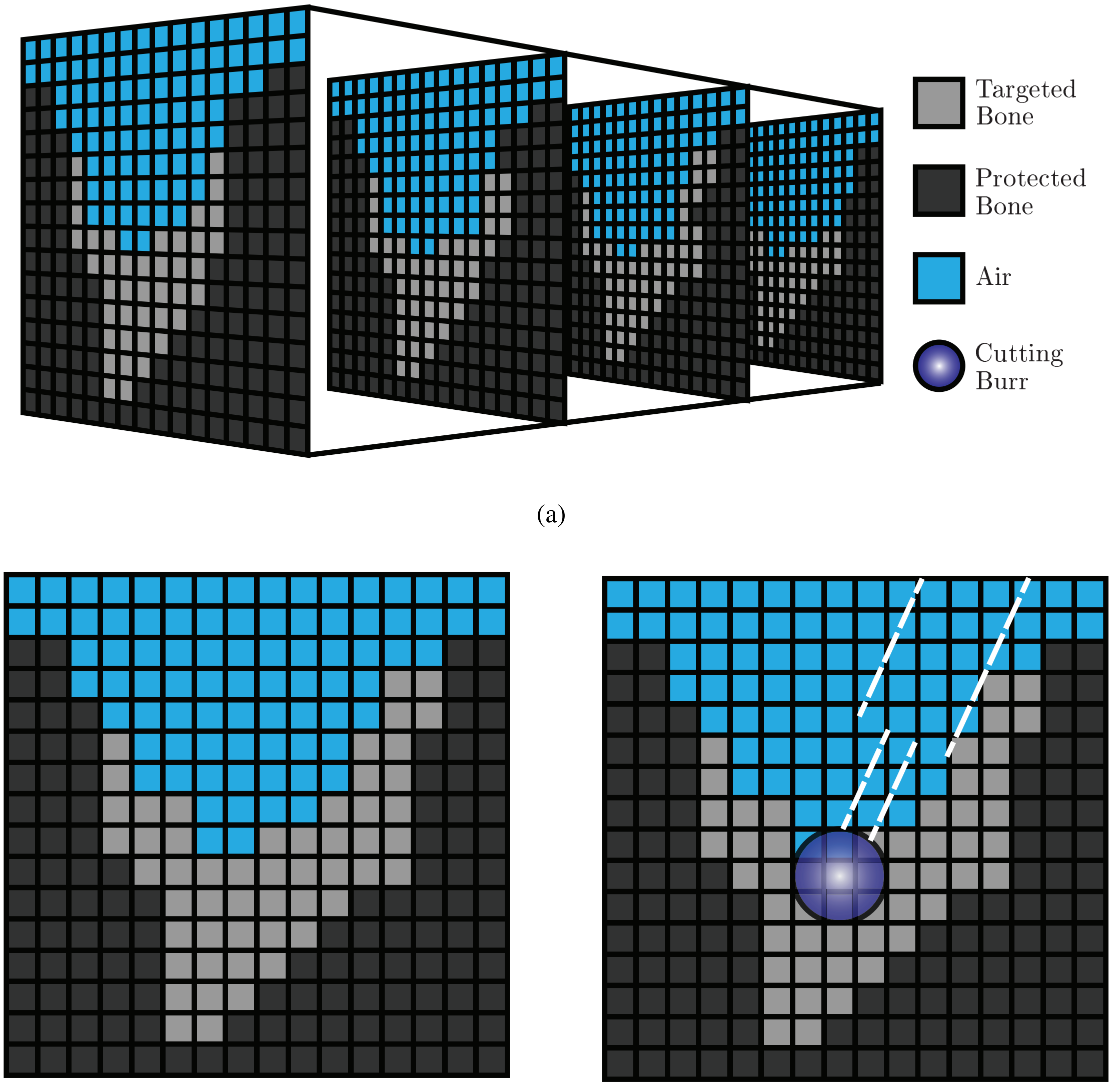

(a)

(b)

(c)

Figure 2.7: This figure demonstrates the simplification that allows for rapid collision detection in the 4 DOF voxelized framework. Part (a) shows several vertical 2D slices of a sample drilling. Part (b) shows the information in all of these slices combined into a single voxel thick slice. The combination is carried out such that if a voxel was unmilled bone (targeted or protected) in any of the slices, then in the combined slice it is represented as unmilled bone. Part (c) shows an example of checking a drill shaft's orientation for collision with the single combined slice from Part (b).

unnecessary calculations are avoided.

In the 5 DOF voxelized framework, the drill is modeled as three approximate cylinders, one for each of the three diameters shown in figure 2.6. Any drill pose may be checked for collision with unmilled bone by comparing the three cylinders that contain the drill with the octree representing the top layer of bone tissue via the Gilbert–Johnson–Keerthi (GJK) algorithm [48]. Sample shafts being checked for interactions with the bone's surface are shown in Figure 2.9.

2.2.4 Modeling the Cutting Burr in the Modified Simulation: Defining the Supervoxel

This portion of the voxelized framework relates to the modified volume and the modified simulation. Given the previous sections, the cutting burr's center may be safely commanded to all target voxels in the modified region. Interactions between the drill shaft and the unmilled bone may also be anticipated and avoided. Another issue arises if the cutting burr's center is commanded to pass through each voxel in the modified volume's targeted voxels before considering that voxel to be removed. Using the previous example, if the cutting burr's diameter is 5 mm and the voxels are 0.4 mm cubes, the cutting burr will be commanded to make cuts of at most 0.4 mm deep. This is inefficient, and leads to long paths and extra computation time.

The presented solution is to extend the idea of a supervoxel. A typical supervoxel is simply the expansion of a voxel to include some set of its neighbors sharing similar characteristics. This supervoxel's dimensions are defined based on the voxels' dimensions and the desired cutting depth. For instance, if each voxel represents a 0.4 x 0.4 *x* 0.4 mm cube and an approximate cutting depth of 1 mm is desired, then the supervoxel will consist of a volume of 5 x 5 x 5 voxels centered on the cutting burr's center. Given these dimensions, the supervoxel will reach out 2.5 voxels, or 1 mm, in each direction from the center of the supervoxel. Thus in the modified simulation all voxels segmented as targeted bone within 1 mm of the center of the supervoxel will be counted as milled. The

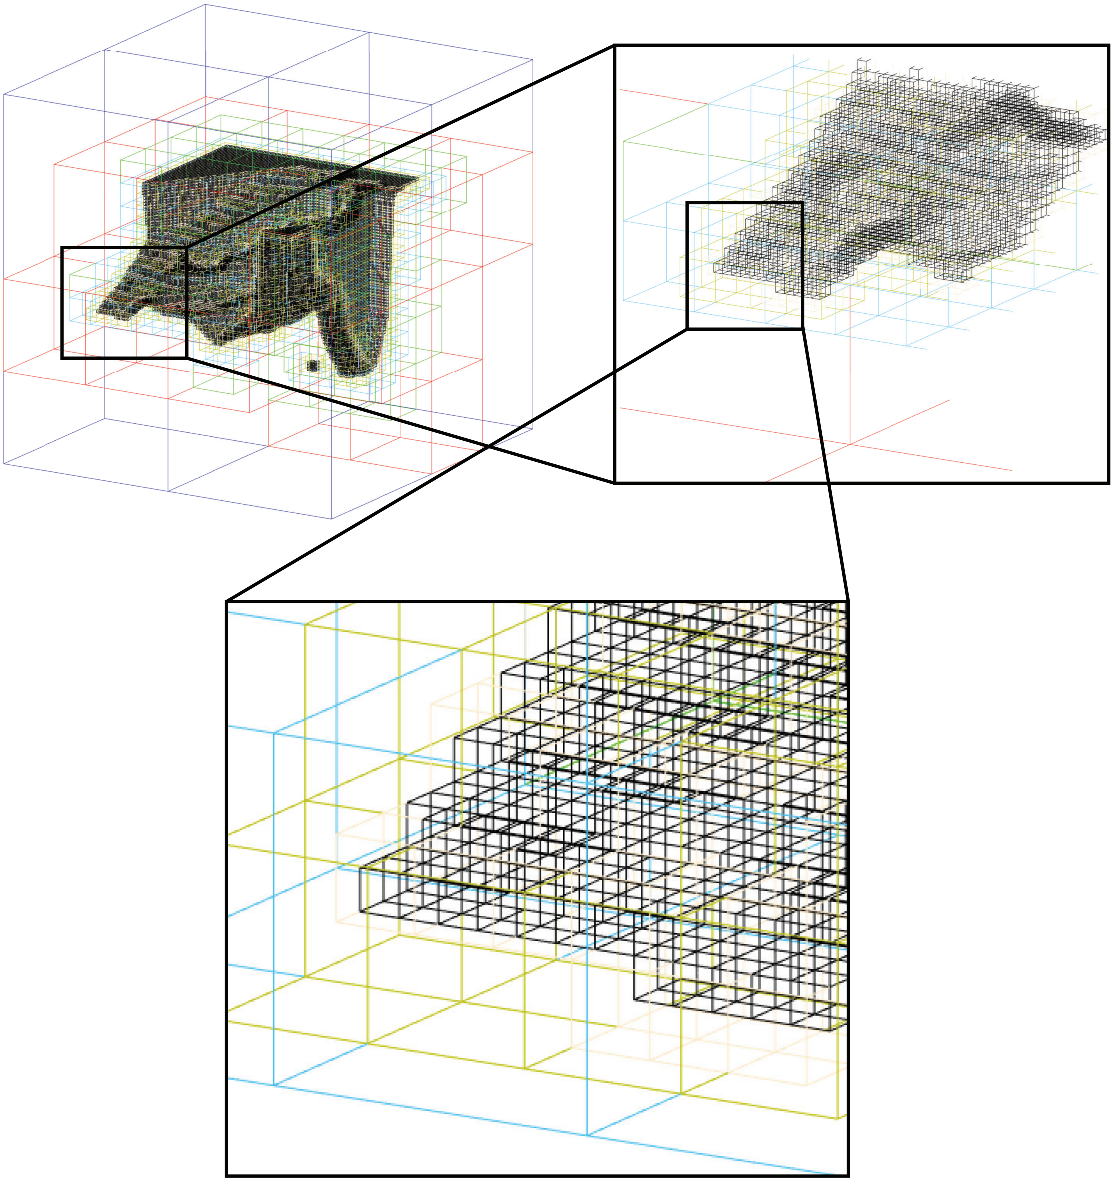

Figure 2.8: A sample octree of the unmilled surface of bone during path planning is shown in the figure. Each color box represents a different level of the octree. The highest octree level is the single box that surrounds the whole volume. The lowest level is composed of the black boxes. These black boxes are chosen to correspond to the unmilled bone surface at a given point in the simulated milling (unmodified simulation). To test for collision between a drill shaft with the unmilled bone surface (black boxes) the drill shaft is first compared with the highest level (largest box pictured). If this highest level is in collision with the drill shaft, the shaft is then checked against the 8 blue boxes (the next level in the octree). This proceeds hierarchically until the drill shaft is checked against only the unmilled surface voxels (black boxes) that are nearest the drill shaft. This octree representation allows for fast and accurate collision detection between the modeled drill shaft and the unmilled surface of bone (the black cubes).

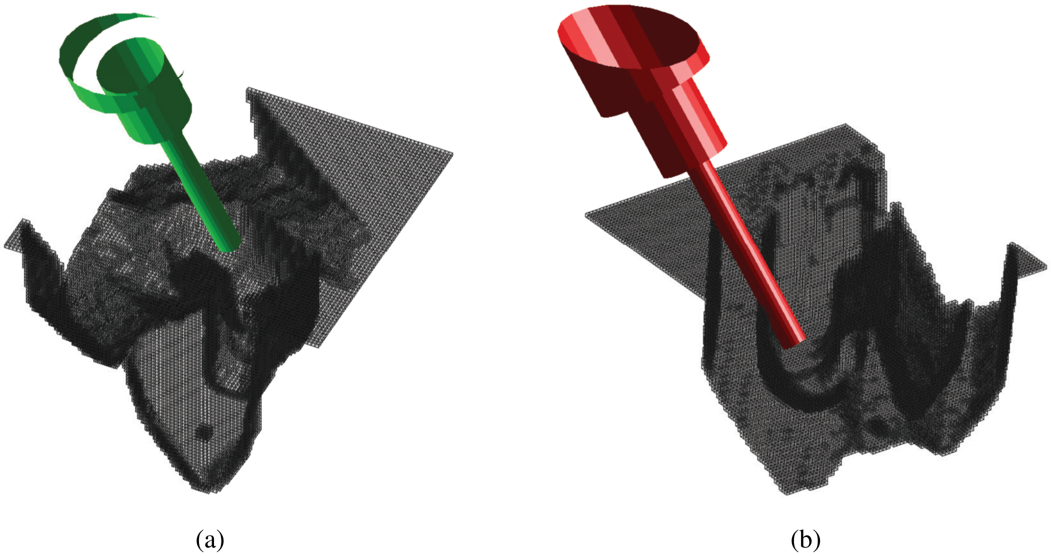
(a) (b)

Figure 2.9: This figure shows an example of a drill shaft being checked for collision with an unmilled surface of bone. Part (a) shows a drill shaft not in collision with the unmilled bone surface. Part (b) show a drill shaft that is in collision with the unmilled bone surface.

supervoxel's dimensions will vary depending on the medical image being used, the density of the bone being milled, and the diameter of the cutting burr. An example supervoxel can be seen in Figure 2.10.

The introduction of the supervoxel requires the addition of two new features to the voxelized framework. These features are outlined in Sections 2.2.5 and 2.2.6.

2.2.5 Defining the Target Boundary Segmentation

This portion of the voxelized framework relates to the modified volume and the modified simulation. The first feature required by the introduction of the supervoxel is related to ensuring the entire volume of targeted bone is removed. The targeted bone voxels in the modified volume were created by an erosion operation where the cutting burr structuring element was applied to the target voxels in the unmodified volume (Section 2.2.1). This erosion was performed so that the center of the cutting burr may be commanded to the boundary of the targeted bone in the modified simulation without the edge of the cutting

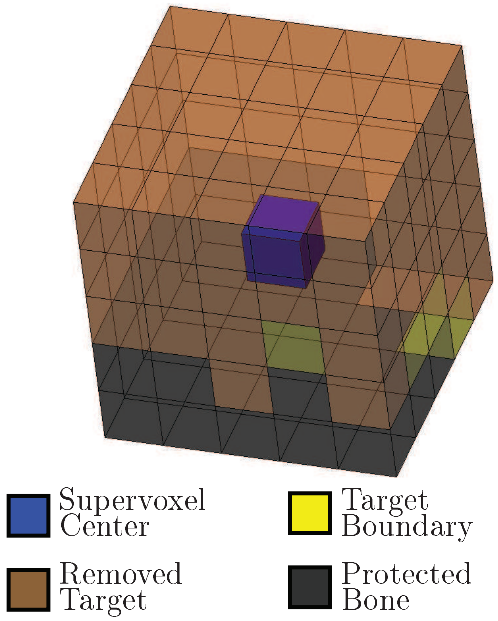

Figure 2.10: A example of a supervoxel is shown in this figure. Note that the type of the voxels contained within the supervoxel are also shown in the legend.

burr removing any protected bone. By defining a supervoxel that counts voxels not covered by the center of the cutting burr as milled, a situation has been created where a target voxel in the modified volume that is situated on the boundary with the protected bone may be marked as milled without the center of the supervoxel reaching that voxel. Figure 2.11 was created to illustrate this situation. Figure 2.11 contains a simulation of a 2D path chosen purely for example. If the $5x5$ supervoxel traverses the white path in Figure 2.11b, then every targeted bone voxel present in the modified simulation will have been marked as milled. Figure 2.11c and 2.11d represent the unmodified simulation running in parallel with the modified simulation shown in Figure 2.11a and 2.11b. There is a 1-1 correspondence between voxel positions in the modified and unmodified simulations. Note that in the figure the modified simulation has been cropped to highlight its details. This has been done because all voxels outside the shown area would be labeled as protected bone.

The band of unmilled voxels present in Figure 2.11d demonstrates that if voxels adjacent to protected bone in the modified simulation are marked as milled by any part of the

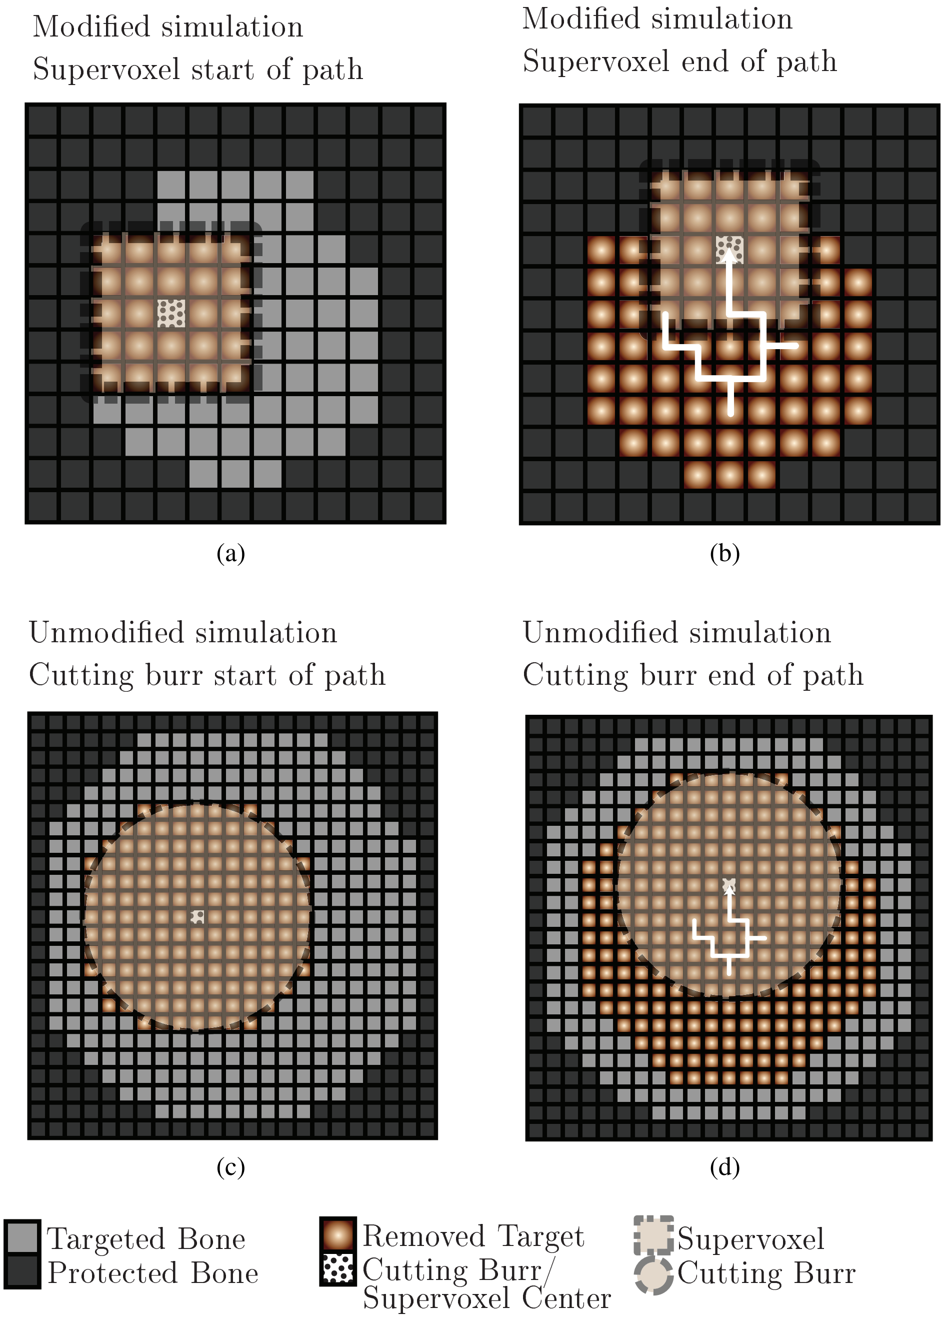

(a) (b)

(c) (d)

Figure 2.11: Parts (a) and (c) of the figure show 2D slices of a point in the modified simulation in parallel with the unmodified simulation. Note the distance between the upper left portion of the "removed" region and the targeted bone in (c). Parts (b) and (d) also represent 2D slices of the simultaneous modified and unmodified simulations and contain a sample path. Note that in part (b) all targeted voxels have been removed, and in part (d) a band of targeted bone was not removed. Ensuring the removal of this band of unmilled targeted bone requires the introduction of an additional voxel label.

supervoxel other than the center voxel, targeted bone in the unmodified simulation (and subsequently the physical space) may be unmilled at the end of the planned path. Therefore, all of the targeted bone voxels adjacent to the protected bone in the modified region must be visited by the supervoxel's center. To overcome this difficulty, an additional label is added to the modified volume. The new segmentation consists of each of the voxels that share faces with (6-connected) the protected bone segmentation in the modified volume. The voxels in this new, targeted bone boundary, segmentation are not counted as removed unless the they are visited by the single voxel representing the center of the cutting burr. A sample trajectory utilizing this new target boundary label is shown in Figure 2.12.

2.2.6 Defining a Simple Objective Planner

This portion of the voxelized framework relates to the modified volume and the modified simulation. The second feature required by the introduction of the supervoxel is related to safety and the size of the supervoxel. By expanding the supervoxel to include a volume of additional voxels, it is now possible that some of those voxels within the supervoxel will be protected bone. Since any targeted bone voxel within the supervoxel has been marked as milled, any path planning algorithm using this supervoxel will be forced to consider voxels exterior to the supervoxel as candidates for the next tool position. Since voxels not immediately adjacent to the voxel representing the current center of the spherical cutting burr (also the center of the supervoxel) are being considered for the next tool position, and since collisions between the drill shaft and unmilled bone are calculated for voxels that represent candidate next tool positions, the distance that the drill will travel between checks for collision between unmilled bone tissue and the drill shaft has been increased.

These two considerations led us to develop a simple planner that would generate a path through consecutive voxels leading from the supervoxel's current center to its desired next position. This simple planner works by analyzing a subvolume of voxels where the targeted voxel and the current center of the supervoxel occupy two opposite corners. The voxels in

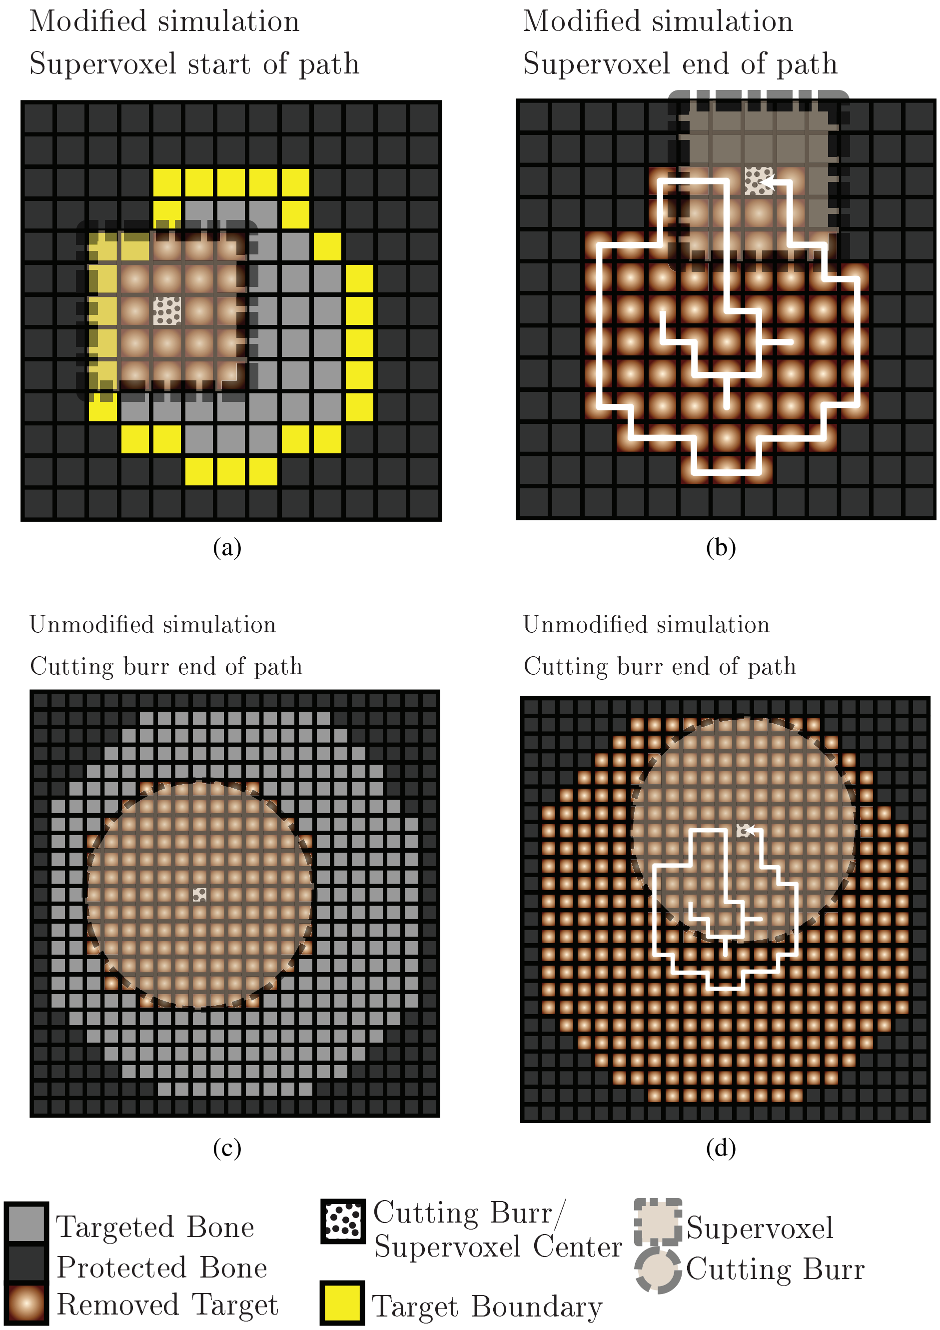

(a) (b)

(c) (d)

Figure 2.12: Parts (a) and (c) of the figure show 2D slices of a point in the modified simulation in parallel with the unmodified simulation. Note the new segmented target boundary in Part (a). Parts (b) and (d) also represent 2D slices of the simultaneous modified and unmodified simulations and contain a sample path. Comparing Part (d) to Figure 2.11 (d), it is observed that the entire region of targeted bone has been removed now that the target boundary has been included in the modified simulation.

the subvolume are labeled with integers corresponding to their Manhattan distance from the current voxel center. Since it is possible for protected bone voxels to be within the supervoxel, and it is impossible for the supervoxel's center to be commanded to a protected bone voxel, it is possible that a desired position exterior to the supervoxel is walled off from the center of the supervoxel by protected bone voxels. Because of this possibility, a check is made that each integer step between the current center voxel and the target voxel exists in the subvolume. If this check is confirmed, the algorithm proceeds to step to lower labeled voxels while checking for shaft collisions with unmilled target and backtracking when necessary. Ties between equal voxel labels are broken randomly. This labeling and a sample path are seen in Figure 2.13.

In this way it is ensured that the largest step taken without checking for collisions between the drill shaft and the unmilled bone corresponds directly to the voxel size. Note that if this voxel size is too large, the segmented volumes may be upsampled to ensure patient safety. The simple objective planner algorithm is outlined in Figure 2.14.

2.3 Description of 3D Path Planning Algorithm

This section presents the true 3D path planning algorithm that has been created to run in the voxelized framework defined by the previous section. This algorithm requires the definition of both the modified and unmodified simulations. These simulations are run in parallel in the modified and unmodified volumes respectively. The modified simulation will be used to evaluate the candidate cutting burr positions at each step of the algorithm. The unmodified simulation is necessary to determine the accessibility of of the candidate positions at each step of the algorithm. The position of the center of the supervoxel in the modified simulation, and the position of the cutting burr's structuring element in the unmodified simulation are synchronized between the two parallel simulations.

The path planner proceeds using this pattern:

1. Preprocessing: segmented CT scans

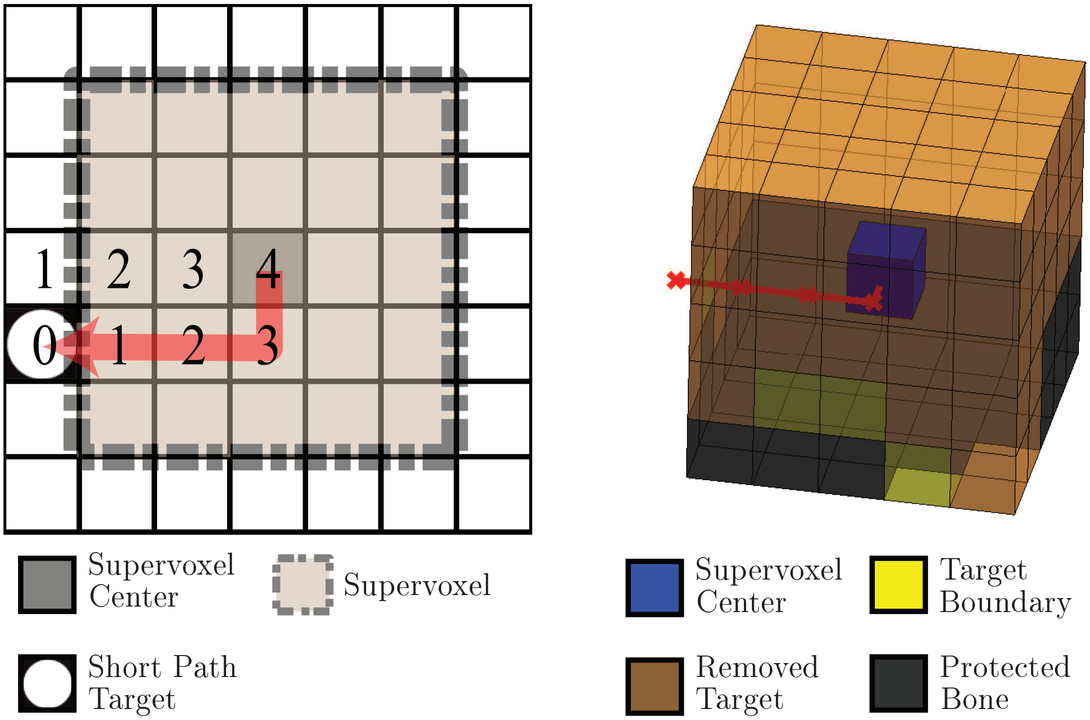

(a) (b)

Figure 2.13: This figure shows examples of the short path mechanism used by the voxelized framework to plan paths to the targets within or immediately adjacent to the supervoxel. Part (a) shows a 2D representation of a supervoxel that needs to move to a target adjacent to the supervoxel. The short path algorithm chooses neighbors with values that are smaller than the current voxel's value for its next step. Part (b) shows a 3D representation of the supervoxel and the path that is chosen through the example shown in 2D in Part (a).

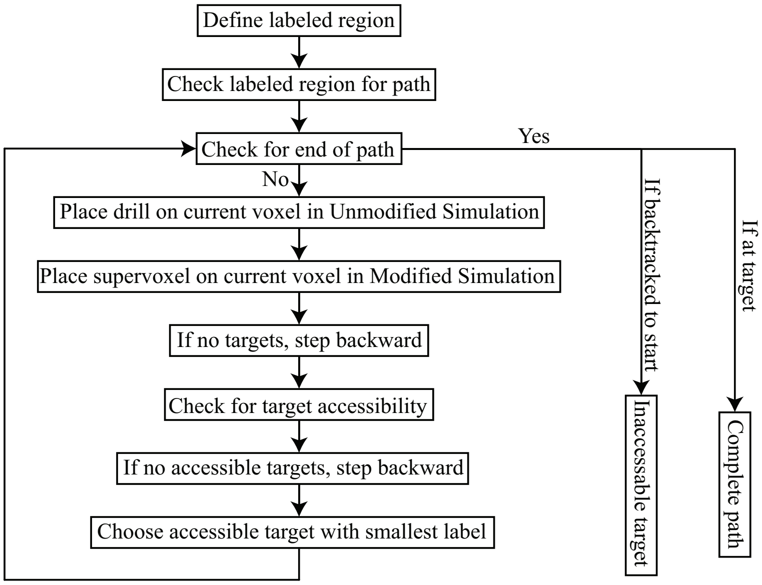

Figure 2.14: The figure shows the flowchart associated with the simple objective planner outlined in Section 2.2.6.

Create the modified and unmodified segmented volumes as described in Section 2.2.1. These volumes will form the basis for the **Modified Simulation** and the **Unmodified Simulation** in the following steps.

2. Preprocessing: modeling the spherical cutting burr

   Modified simulation: Define a supervoxel (Section 2.2.4) to represent the cutting burr in the modified simulation. The dimensions of this supervoxel are chosen to limit the cutting depth of the milling strategy.

   Unmodified simulation: Define a structuring element to represent the physical dimensions of the spherical cutting burr in the modified simulation.

3. Taking the first step

   Modified simulation: Place the center of the supervoxel created in Step 2 on a voxel on the surface of the targeted bone segmentation. Mark the voxels within the supervoxel as milled according to Section 2.2.4.

Unmodified simulation: Place the center of a structuring element also defined in Step 2 on the voxel that corresponds to the initial voxel chosen in the modified volume. Mark the voxels contained within the structuring element as milled. Check for the angles from which the voxel is accessible using the methods defined in Section 2.2.3.

4. Identify candidate next steps

   Modified simulation: Identify unmilled target voxels that are candidates for the next drill pose from the set of voxels within and immediately adjacent to the supervoxel.

5. Check candidate next steps for accessibility

   Unmodified simulation: Evaluate the accessibility of those candidates by comparing a model of the drill shaft with the surface of unmilled bone (either protected or targeted) as described in Section 2.2.3.

6. Possibility: no accessible candidates

   Modified simulation: If there are no accessible candidates, plan a direct path to the nearest accessible target voxel using an extension of the simple objective planner presented in Section 2.2.6.

7. Possibility: there are accessible candidates. Rank them by fitness

   Modified simulation: Evaluate and rank the accessible candidate poses by a predefined fitness metric.

8. Plan the length of the step

   Modified simulation: If there are no protected bone voxels present within the supervoxel, a single step is taken directly to the chosen candidate voxel position. Otherwise, a short path is planned to the candidate avoiding any preserved bone or inaccessible voxels present in the supervoxel via the methods described in Section 2.2.6.

9. Move the cutting burr

   Modified simulation: Place the supervoxel in its new position. Mark the center voxel of the supervoxel as removed. Mark all target voxels within the supervoxel that are not directly adjacent to the protected bone as removed.

   Unmodified simulation: Place the cutting burr's structuring element centered on the new position in the unmodified volume. Mark all voxels within the cutting burr's structuring element as removed.

10. Repeat steps 4-9 until no more valid candidate voxels remain.

11. **Modified simulation:** Use the simple objective planner described in Section 2.2.6 to plan a path back to the starting position.

It should be noted that, depending on the physical application, it may be desirable to mill the volume multiple times with different diameter cutting burrs to increase the efficiency of the milling process. This algorithm can easily accommodate such successive millings by being applied first to the unmilled volume with the larger cutting burr. Then, the results of that first plan are used as the initial state of the second plan for the smaller cutting burr. This can be repeated as many times as the individual situation requires.

2.3.1 Preprocessing the Segmented Medical Image

In addition to the steps described in Section 2.2, each targeted bone voxel undergoes an initial accessibility check. This is done to identify targeted bone voxels that cannot be reached by the spherical cutting burr because of severe overhangs of unmilled bone. These identified voxels are then marked as inaccessible. This preprocessing accessibility check is simply an application of the methods presented in Section 2.2.3 to every voxel in the targeted bone segmentation. This accessibility check compares only drill shaft and protected bone, since it is designed to remove from consideration voxels that could never be reached by the cutting burr because of the pocket's shape.

2.3.2 Identify Candidates

Modified simulation: The supervoxel is placed in the segmented medical image at the cutting burr's current position. The segmentation values of the voxels within and immediately adjacent to the supervoxel are identified. All of these voxels representing unmilled targeted bone or targeted bone boundary are analyzed to determine if they are accessible to the drill.

2.3.3 Evaluate Accessibility

Unmodified simulation: This accessibility check is the only difference between the 4 and 5 DOF versions of this algorithm. Both checks are performed as described in Section 2.2.3. It is necessary not only to find the accessible voxels, but also to rank the various candidate voxels based on the angles by which they can be accessed. This desire is motivated by earlier studies performed by Dillon et al. relating to the efficiency of cutting with a spherical burr [45]. By determining many different accessible angles, the algorithm may now modify the planned path based on the projected efficiency of the cutter at the candidate positions.

For the 4 DOF case, the first and last angles representing the range of orientations from which the voxel may be approached by the drill shaft are stored. This range is then checked for sufficient overlap with the current position's range of accessible angles. If a candidate voxel's accessibility window has a sufficiently large range of accessible angles, and a sufficient overlap with current voxel's accessibility window, then the voxel is marked accessible and the range of accessible angles are stored (Figure 2.15).

For the 5 DOF case, a discretized range of poses is analyzed as a concession to the computational intensity of the collision detection. These drill shaft poses are then checked against the octree representing the surface of the unmilled and protected bone as described in Section 2.2.3. If a threshold percentage of the evaluated poses are accessible, then the

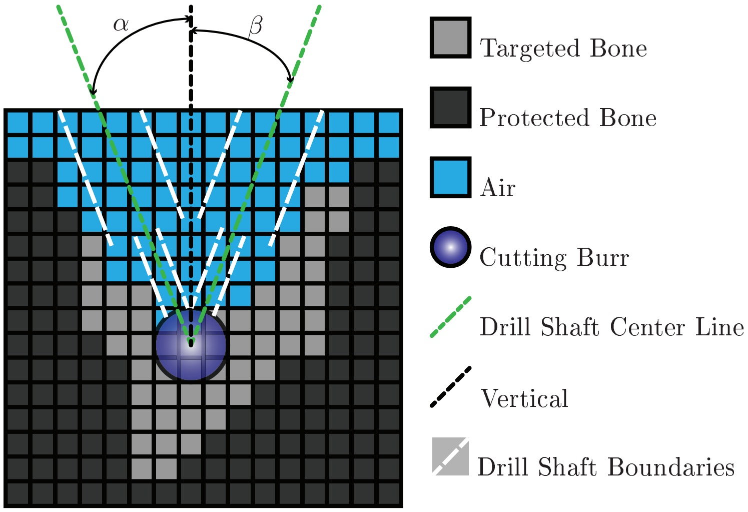

Figure 2.15: The 2D slice shown in the figure is the collapsed slice created in Figure 2.7. α and β are the angles stored by the algorithm to give an range of accessible angles for that particular voxel.

pose is marked accessible and the angles that define the position of the shaft whose cutting burr is best aligned with the targeted bone's surface is saved corresponding to that voxel's position.

2.3.4 Plan Short Path

Modified simulation: If no accessible candidate voxels exist, then a search of all unmilled target left in the modified volume is performed. The unmilled targeted bone is sorted based on its Euclidean distance from the cutting burr's current position. These voxels are checked for accessibility and a path is planned to the closest accessible voxel (Section 2.2.6). The process then restarts at Section 2.3.2.

2.3.5 Rank Candidate Voxels

Modified simulation: If accessible candidate voxels do exist, then they are subjected to a fitness metric and ranked accordingly. Using this fitness metric to guide the drill tip is the source of the flexibility of the algorithm presented in this chapter. Any number of

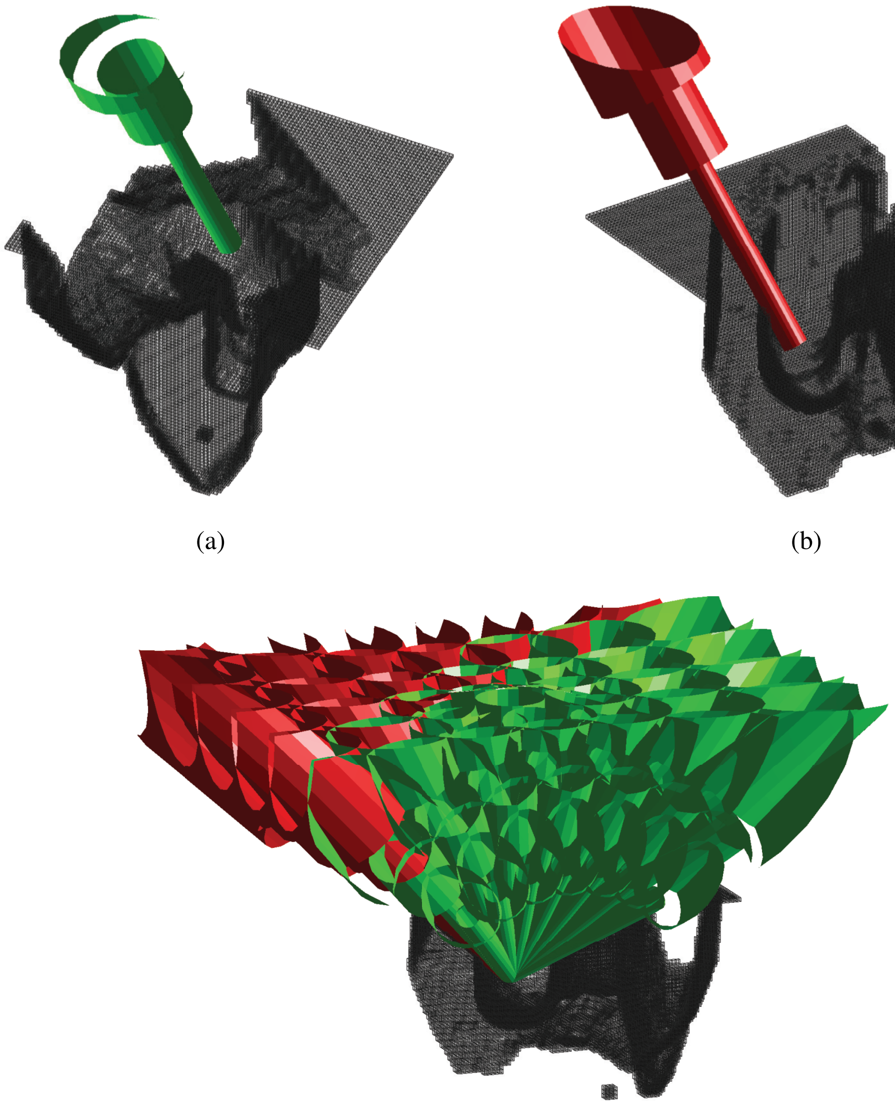

(a)

(b)

(c)

Figure 2.16: This figure shows an example of a drill shaft being checked for collision with an unmilled surface of bone. Part (a) shows a drill shaft not in collision with the unmilled bone surface. Part (b) show a drill shaft that is in collision with the unmilled bone surface. Part (c) shows a range of shafts being tested for collision with the unmilled bone surface.

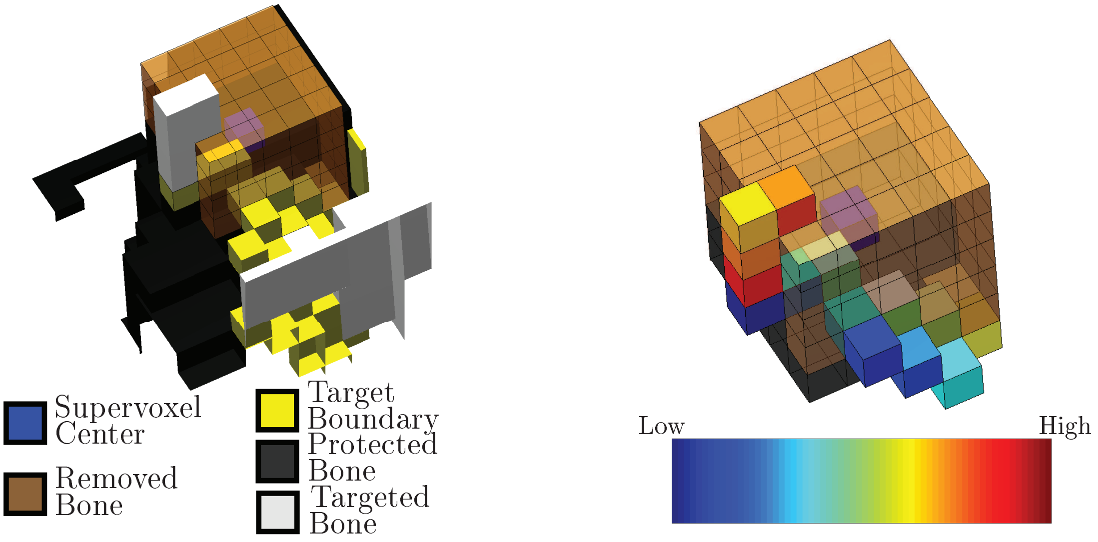

(a) (b)

Figure 2.17: This figure shows the ranking of candidate voxels that the algorithm will choose between for the next step in the planned path. Part (a) shows a 3D visualization of the volume surrounding the supervoxel with the targeted voxels highlighted. Part (b) shows a 3D visualization of the supervoxel surrounded by the candidate steps and the candidates are ranked by fitness from low to high.

conditions may be scaled and added to the normalized metric. For example, because of the geometry of the spherical cutting burr, it does not perform well when plunging along the drill shaft's axis to make deep cuts. This poor performance is because the robot uses a spherical cutting burr that is far less efficient at removing material with the bottom or near the top of the cutting burr. This cutting efficiency is of course dependent on the angle of the drill shaft as well, though the simplifying assumption that the cutting burr is aligned with the z axis has been made. To consider this physical limitation of the spherical cutting burr in the unmodified simulation, the number of target voxels present in the top half of the cutting burr is calculated for every candidate voxel. This number is then scaled and added to the fitness metric used in the modified simulation. Thus a preference for candidates with fewer voxels in the top half of the cutting burr is built into the implementation of this algorithm. The fitness metric used in this chapter's experiments will be included in the simulations Section 2.4.1. A visual example of this ranking can be seen in Figure 2.17.

2.3.6 Move Cutting Burr

Modified simulation: The supervoxel is moved to the voxel location corresponding to the highest fitness metric. The center voxel of the supervoxel is marked as removed. All unmilled targeted bone voxels within the supervoxel that are not adjacent to protected bone voxels are marked as removed. The process then restarts at Section 2.3.2 until all accessible targeted bone has been removed.

Unmodified simulation: The cutting burr's structuring element is moved to the same coordinates as the supervoxel, but in the unmodified volume. All voxels contained within the cutting burr's structuring element in the unmodified volume are marked as removed. The parallel movement of the supervoxel and cutting burr through the modified and unmodified simulations respectively can be seen in Figure 2.18.

2.3.7 Plan Path to Start Voxel

Modified simulation: When all accessible targeted voxels have been milled, use the simple objective planner described in Section 2.2.6 to plan a path back to the starting position. Note that there may still be a small volume of targeted bone left unmilled. This volume represents voxels that cannot be approached with enough overlap between the accessibility windows described in Section 2.3.3.

2.4 Results

2.4.1 Simulations

The algorithm was evaluated by applying both its 4 DOF and 5 DOF variations to 36 cadaver CT scans that had been segmented for robotic mastoidectomy. These scans were performed on one of two xCAT ENT portable CT scanners (Xoran Technologies, Ann Arbor, MI, USA). The voxel size of these scanners is set to 0.4 mm cube. The scan volume is either 640 x 640 x 355 or 548 x 548 x 312 voxels. The scans were segmented

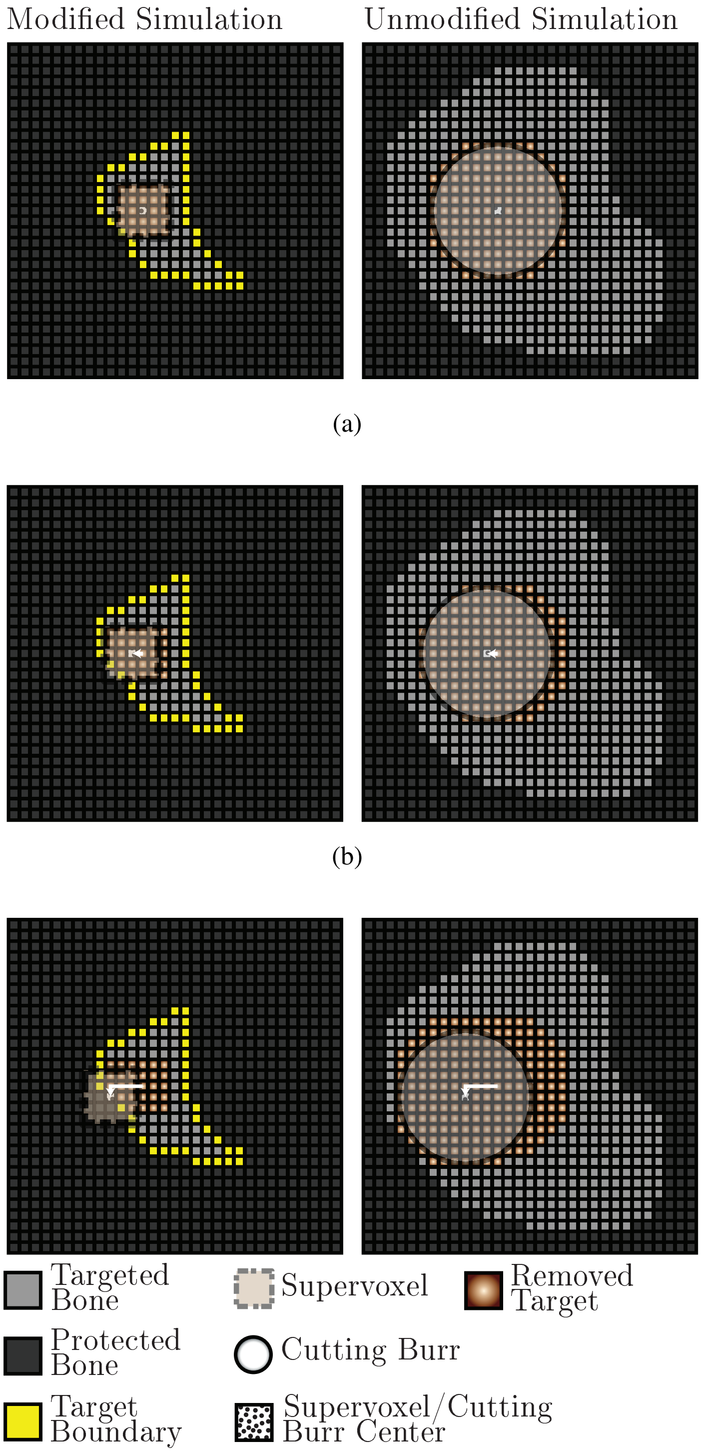

(a)

(b)

(c)

Figure 2.18: This figure shows an example of the simultaneous modified and unmodified simulations. Part (a) shows the first step of simulated milling on a given 2D slice. Part (b) shows step two in the same simulation. Part (c) shows the same simulation after several steps have been performed.

into three different areas: targeted bone, protected bone, and air. The preprocessing steps described in Section 2.2 were applied to each of these scans. A 5 mm diameter cutting burr was modeled. A supervoxel of size 5 x 5 x 5 was used to correspond to an approximate cutting depth of 1 mm. The three modeled shaft diameters for the Midas Rex Legend surgical drill (Medtronic, Minneapolis, MN) used for physical experiments were 2.4 mm, 6.8 mm, and 11.6 mm. After the preprocessing, the algorithm described in Section 2.3 was applied to the segmented regions. The metric, F_m for determining which steps were chosen next at each point along the planned path as described in Section 7 is a normalized combination of five different considerations: (1) preserving the current planar direction of the cutting burr's movement, (2) covering as many voxels as possible on the next step, (3) placing the center of the supervoxel on a target boundary voxel, (4) limiting the volume of targeted bone within the top half of the cutting burr, and (5) adding a bias toward upward movements of the cutting burr. These metrics are individually normalized and combined with the following weights before the final metric is normalized, $F_m = (1) + (2) + 10(3) + 5(4) + (5)$.

Additionally, using the voxelized framework described in Section 2.2, a simple, 2.5D zigzag, path planner was developed for comparison with the 3D algorithm. The necessarily arbitrary shape of the targeted bone volumes is difficult for such a simple planner to cover, so two small modifications were made: (1) The planner was allowed to identify additional pockets of unmilled targeted bone and move to them if possible. (2) After the zigzag portion of the drilling was completed for a layer, the algorithm would make a circuit around any reachable undrilled target present on the boundaries of the slice. This mimics traditionally separate rough and finish cuts because the 3D algorithm presented in this chapter also combines those into one step. The results of these 4 and 5 DOF simulated 2.5D and 3D path plans are compared in the following section.

For the 4 DOF version of the 3D planner, a joint trajectory plan for the Acoustic Neuroma Surgical Robot (ANSR) previously developed by Dillon et al [18] was also simulated.

Simulated joint trajectories were also generated for the baseline 4 DOF zigzag planner. This will be the same robot and joint trajectory planner used in the experiments described in Section 2.4.3.

The robot joint trajectory planner required a minimum and maximum cutting burr speed because the planner analyzes the CT scan of the milling to estimate the bone density being milled at each step of the planned cutting burr trajectory. This approximated bone volume removed per robot step is used to modify the cutting burr's linear speed throughout the milling process. This minimum and maximum cutting burr speed were set to 0.5 mm/s and 5 mm/s respectively. These speeds are comparable to cutting burr speeds chosen by other bone milling path planners [2, 3, 45]. Note that the speed is lower than Sugita et al. This choice was made because the robot described in that paper mills on the knee and does not work in close proximity to critical structures such as those embedded in the temporal bone. (e.g. facial nerve, chorda tympani, sigmoid sinus, etc.)

2.4.2 Simulation Results

A path plan is shown within the segmented targeted bone in Figure 2.19. Also shown is the simulated volume of bone that has been removed by the drill traversing that path. Note that it has not drilled any unintended volume, and that some volume has been left behind as inaccessible to the drill bit, either because a cutting burr of that size could not fit into the space required to remove some targeted bone without removing some protected bone, or because the targeted bone voxels could not be reached by the shaft of the drill bit without the shaft touching undrilled bone.

These algorithms were both implemented in MATLAB utilizing the parallel computing toolbox and executed on a custom built machine designed around an Intel Corei9 9900k processor with 8 cores running at 4.8 GHz with 64 GB of memory. Thirty-six separate cadaver scans were used for simulation of both the 3D and the baseline zigzag path planner for both the 4 DOF and 5 DOF algorithms. Complete simulation results are shown in Tables

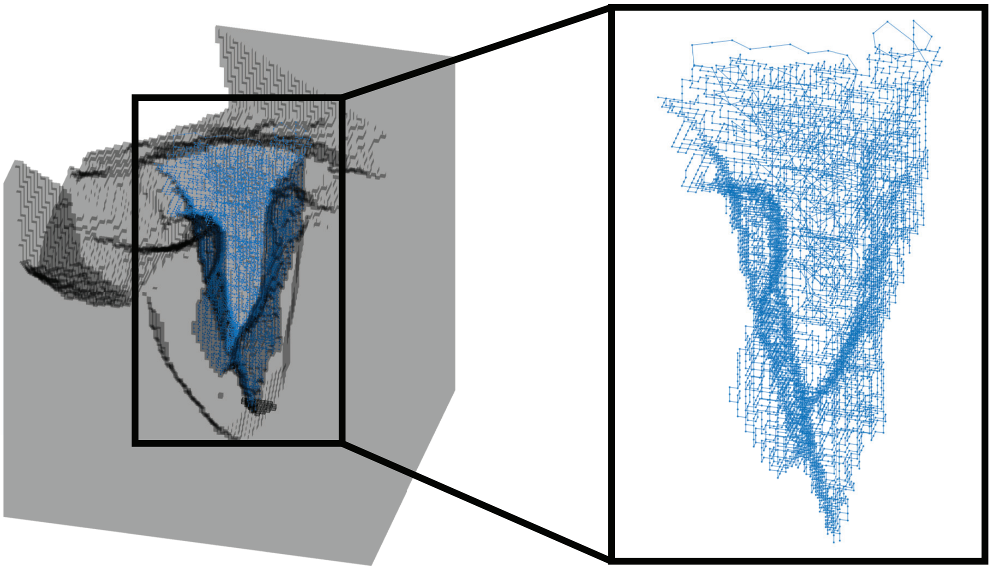

Figure 2.19: The left side of the figure shows a planned path (blue line) placed within the modified region. The call out on the right hand side isolates the path for better visualization. Note the how the final planned path freely travels along the vertical axis.

A.1, A.2, A.3, and A.4. A visualization of a planned path in a modified region appears in Figure 2.19.

2.4.3 Experiments

A set of experiments to validate the simulation results was designed. These experiments were conducted using the Acoustic Neuroma Surgical Robot (ANSR) presented in Dillon et al. [18]. This robot was designed to autonomously mill a mastoidectomy as a component of a translabyrinthine approach for acoustic neuroma (a.k.a. vestibular schwannoma) removal. The ability of the ANSR system to accurately remove segmented pockets of bone has been previously established by Dillon et al. [49]. ANSR is a 4 DOF robot, so only the 4 DOF simulations will be validated by this set of experiments; though, the 4 DOF and 5 DOF algorithms are identical except for their collision detection methods (Section 2.2.3).

	Predicted path length (mm)	**Milling path length (mm)**	**Modified final targets / initial**	**Unmodified final targets / initial**	**Modified initial target (cm^3)**	**Simulation time (s)**	**Predicted milling time (s)**	**Average milling speed (mm / s)**	**Steps in contact / path length**	**Average voxels milled per step in contact**	**Ratio predicted path length (mm / cm^3)**
Zigzag 4 DOF Simulation Average Results											
means	722	719	1.3e-02	1.2e-01	3.15	133	1937	3.7	0.66	8.0	238
±	236	235	4.2e-03	3.9e-02	1.03	44	633	1.2	0.22	2.6	78
3D 4 DOF Simulation Average Results											
means	656	600	1.8e-04	1.0e-01	3.15	653	1342	4.6	0.08	32.9	227
±	214	196	5.9e-05	3.4e-02	1.03	213	438	1.5	0.03	10.7	74
Zigzag 5 DOF Simulation Average Results											
means	718	-	2.4e-02	1.0e-01	3.10	8125	-	-	0.7	7.96	241.0
±	234	-	7.9e-03	3.3e-02	1.01	2654	-	-	0.2	2.60	78.7
3D 5 DOF Simulation Average Results											
means	641	-	1.7e-02	9.3e-02	3.10	60592	-	-	0.1	32.34	223.0
±	209	-	5.7e-03	3.1e-02	1.01	19794	-	-	0.0	10.56	72.8

Table 2.2: Simulation results averaged over all 36 segmentations based on cadaver CT scans that were evaluated for this study. These results are from both the zigzag and the 3D path planning algorithms and both 4 and 5 DOF versions of each algorithm. All means are reported with 95% confidence intervals. The full simulation results are found in Appendix A, Tables A.1, A.2, A.3, and A.4. Note that several columns of the 5 DOF table are empty. These are related to the milling simulations and physical millings that were conducted with the ANSR robot. Since the ANSR robot is a 4 DOF robot, this data was not generated for 5 DOF.

The 3D milling algorithm presented in Section 2.3 and the zigzag algorithm designed for Section 2.4.1 were applied to a set of 15 cadaveric temporal bones fixed in resin. This fixation is necessary to provide enough surface area to attach the robot to the temporal bones. These temporal bones had been cut off at the level of the facial nerve. This will limit the physical millings to mastoidectomies instead of the much deeper mastoidectomies combined with the translabyrinthine approach for acoustic neuroma removal, such as those simulated in all of the cadaver head scans in the previous section.

The CT scans were obtained previously for the simulations performed in Section 2.4.1. The minimum and maximum linear speeds for the cutting burr were set to 0.5 mm/s to 5 mm/s to agree with the simulations reported in the previous section. The segmentation and algorithm preprocessing steps were identical to the corresponding simulations. Of the 15 temporal bones that were milled, seven of them were milled using the zigzag milling strategy and eight of them were milled using the 3D milling strategy. The results of these 2.5D and 3D milling strategies are reported in the following section.

2.4.4 Experimental Results

The algorithms were implemented in MATLAB utilizing the parallel computing toolbox and executed on a custom built Intel Corei9-9900k 8-core machine water cooled and set at 4.8 GHz across all cores. The machine has 64 GB of ram and an m.2 SSD hard drive for fast access to large data sets. Experimental results appear in Table 2.3.

2.5 Discussion and Future Work

2.5.1 Simulations

The average simulation results found in Table 2.2, show that on average the true 3D algorithm presented in this chapter generated a path plan that was linearly 9.1% shorter than that the simple zigzag planner that was implemented for comparison. Beyond that when

Specimen	Modified initial target (cm^3)	Predicted milling time 3D (s)	Predicted milling time zigzag (s)	True milling time (s)	% diff predicted & actual milling times	Specimen	Modified initial target (cm^3)	Predicted milling time zigzag (s)	Predicted milling time 3d (s)	True milling time (s)	% diff predicted & actual milling times
Milled with Zigzag						**Milled with 3D**					
TB01	4.3	1607	2782	2910	4	TB03	2.4	1479	1045	1159	10
TB05	1.3	634	884	934	6	TB04	1.7	1092	807	884	9
TB09	2.2	1046	1471	1559	6	TB06	1.9	1146	860	937	9
TB10	1.1	678	908	971	7	TB07	2.1	1372	978	1061	8
TB12	4.6	1774	2820	2953	5	TB11	4.0	2720	1729	1861	7
TB13	4.8	1970	3009	3159	5	TB14	1.0	590	656	724	10
TB16	6.0	2277	3441	3626	5	TB15	2.6	1649	1126	1219	8
						TB17	3.5	2342	1555	1679	8
means	3.5	1427	2188	2302	5	means	2.4	1549	1094	1191	9
±	0.3	110	168	177	0.4	±	0.2	119	84	92	0.7

Table 2.3: Results from milling 15 temporal bones with both the 3D and zigzag algorithms. Means are reported with 95% confidence intervals.

the average predicted milling time for the 4 DOF implementations of the 3D and zigzag algorithms applied to the exact same cadaver scans is examined, a 30.7% reduction in predicted milling time is achieved. It is necessary to note that the ANSR joint trajectory path planner internally simplifies the planned paths which accounts for an approximately 8.5% reduction in the 3D planner's path length and a much smaller 0.4% reduction of the zigzag planner's path length. This reduction in length is caused by the ANSR joint trajectory planner rounding off corners of the Manhattan distance based path plans generated by the 3D path planner. Since the zigzag path planner's trajectory consists mainly of long straight lines, there are far fewer such corners to round off. The remainder of the reduction in time is created by the ANSR joint trajectory planner analyzing every step of the planned path and comparing it to the pre-operative CT scan to determine what density of bone is being removed during each step of the trajectory. This data is used to inform the instantaneous linear speed of the cutting burr throughout the planned path. From the data in Table 2.2, the average percentage of the steps of the zigzag path planner and the 3D path planner that are in contact with unmilled bone are 66% and 8% respectively. Additionally, when only considering path steps that are in contact with the cutting burr, the zigzag planner and the 3D planner on average remove 8.0 and 32.9 voxels per step in contact with unmilled bone respectively. The disparity between these numbers leads to the conclusion that the zigzag planner is in far more constant contact with the surface of the bone than the 3D planner. Fewer steps in contact with the bone under the 3D planner allows the robot to move with a far greater linear speed through the portions of the path where bone is not being milled. For plans generated by the 3D planner, the ANSR joint trajectory planner does have to slow down to mill the bone more aggressively once it is in contact with bone. Comparing the average linear speed of the cutting burr over the exact same set of mastoidectomy target volumes, the 3D planner and the zigzag planner yield 4.65 mm/s and 3.7 mm/s respectively.

Analyzing the individual predicted milling times for the 4 DOF zigzag and 3D path planners reported in Tables A.1 and A.2 shows not only large variation in the predicted gains, but that for temporal bone 14 the zigzag milling path is predicted to be $66s$ faster than the 3D milling path. Temporal bone 14 is also the smallest mastoidectomy volume planned in these simulations at 0.99 cm^3. The next smallest is temporal bone 10 at an initial modified volume of 1.12 cm^3. The predicted simulation time for the 3D tool path planner is 230 s faster than the zigzag planner. The relationship between predicted milling time and initial target volume is analyzed further in Figure 2.20. The plot also contains the predicted simulation times for each 4 DOF milling strategy vs the initial modified target volume the planned tool path traverses.

The modified target volume is used for analysis because, as discussed in Section 2.2.1, there are voxels in the unmodified volume that could never be milled by the cutting burr without removing unintended bone. These voxels that cannot be milled by either the zigzag or 3D path planners should not be part of the efficiency calculations presented in this section. Therefore, the modified volume is used for these comparisons.

It should be noted that while both the 3D and zigzag planners removed the vast majority of the bone segmented as target in the CT scans, the 3D planner consistently removed more of the intended target volume than the zigzag planner. It is believed this is because the 3D planner can access a voxel from any of several directions, while the zigzag planner is often limited by the row that it is moving along.

The advantages gained by the 3D path planner to come at a computational cost. The average simulation time for the 4 DOF 3D planner is $653s$ while on average the zigzag planner only took $133s$. It is noteworthy that both planners are implemented in Matlab for ease of implementation and modification. It is reasonable to assume that the time required for either algorithm to execute would be clinically acceptable if these algorithms were implemented in a faster, compiled language such as C++.

Given the 0.5% and 2.3% percent differences between the mean predicted path lengths of the 4 DOF and 5 DOF zigzag and 3D path planners respectively, it is reasonable to expect the results from the 5 DOF algorithm simulations found in Table 2.2 to be fairly close to

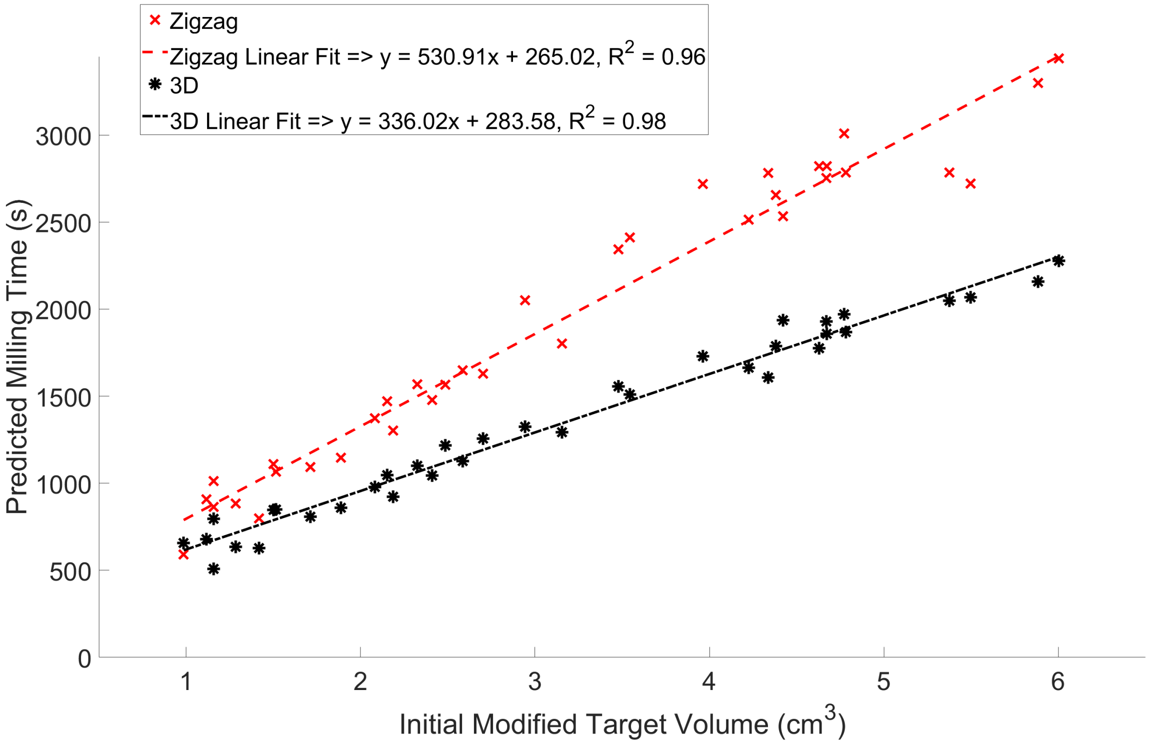

Figure 2.20: For each algorithm the relationship between predicted milling times and initial modified target volumes is shown. Since the slope of the line fitted to the 3D data set is lower than the slope of the line fitted to the zigzag data set, it is clear that the advantages of the 3D algorithm over the zigzag algorithm become more pronounced as the volume of the removed bone increases.

the data collected from the 4 DOF algorithm, and that is largely the case. There are two exceptions that require discussion.

First, the average ratio of initial to final modified volume for the 3D planner is $1.8e^{-4}$ cm^3 and $1.7e^{-2}$ cm^3 in the 4 DOF and 5 DOF 3D planners respectively. In one sense, these two volumes of unmilled bone are so negligible compared to the average of all the modified initial target volumes, 3.15 cm^3, as reported in table A.1 and A.2 that the point is moot. Still, it would be expected that the 5 DOF 3D algorithm would be able to consistently reach more volume than the 4 DOF path planner. The likely cause for this counter intuitive result is the limited number of 5 DOF shafts being sampled in the collision detection (Section 2.2.3). Since the 5 DOF collision detection was very computationally expensive, a limited number of shafts were tested for each step (324 shafts). These shafts were sampled from shaft orientations within 2.5° of the shaft's current location. This relatively course discretization may cause accessible voxels to be marked as inaccessible. As a result, those voxels would remain unmilled.

Second, the simulation time is significantly higher for the 5 DOF path planner than for the 4 DOF. The average simulation time for the 3D 5 DOF planner is so high, 60592 s, that it is impractical when programmed in this un-compiled language on this desktop computer. There are modifications that could reasonably be expected to increase the speed of the 5 DOF algorithm. The only difference between the 4 DOF and 5 DOF algorithms is their means of collision detection to determine what drill shaft angles are achievable by the robot and do not collide with unmilled bone (Section 2.2.3). Therefore, the computational complexity of the 5 DOF collision detection should be reduced. The 5 DOF algorithm currently samples a grid of shafts centered on the current drill shaft orientation and extending radially 2.5° from the current shaft location in both rotational DOFs. This 5° x 5° window is populated with 18 shafts along each rotational DOF ($18^2 = 324$). It is very likely that as the cutting burr descends into the target volume that a large portion of the 5° x 5° window that is being sampled is obscured by unmilled bone and therefore inaccessible. It may be possible to perform the simplified estimate used for the 4 DOF algorithm for each rotational axis. The result of such a preprocessing collision detection step would be to create new limits for the possible angle windows in both rotational DOFs. These new limits would restrict the number of shafts tested significantly by shrinking the currently set 5° x 5° sampling window for a large percentage of the targeted points in the modified volume. Testing that many fewer drill shafts for collision with the unmilled bone would greatly reduce the computational time required for the 5 DOF algorithm.

2.5.2 Experiments

Analysing the results from the physical experiments is far more straightforward than those from the previous section. The parameter being tracked in this set of experiments was the actual time the mastoidectomy millings required. Examining the results from Table 2.3 reveals that the percent difference between the predicted and actual milling times were $5 \pm 0.4\%$ and $9 \pm 0.7\%$ for the zigzag and 3D planners respectively. The consistently longer than predicted milling times are intuitive since the ANSR trajectory planner does not account for resistance at the cutting burr when creating the estimated milling duration. Also, the longer milling times for the 3D planner are easily explained by the fundamental differences between the zigzag and 3D milling strategies. The zigzag planner consists largely of long segments of path that are in contact with a small volume of bone. The 3D planner spends far less time in contact with the bone. When it does contact the bone, it is in contact with several times the volume of bone as discussed in the previous subsection. This higher volume removed per step implies a higher force is applied to the cutting burr. Since the time estimate does not account for such forces, it is intuitive that the percent difference between the zigzag planned simulations and experiments would be lower than that of the 3D planned experiments and simulations. It should be noted that the underestimations are very consistent, which implies that the predicted milling times are still a useful metric for comparison of the larger data set simulated in the previous section. A visualization of

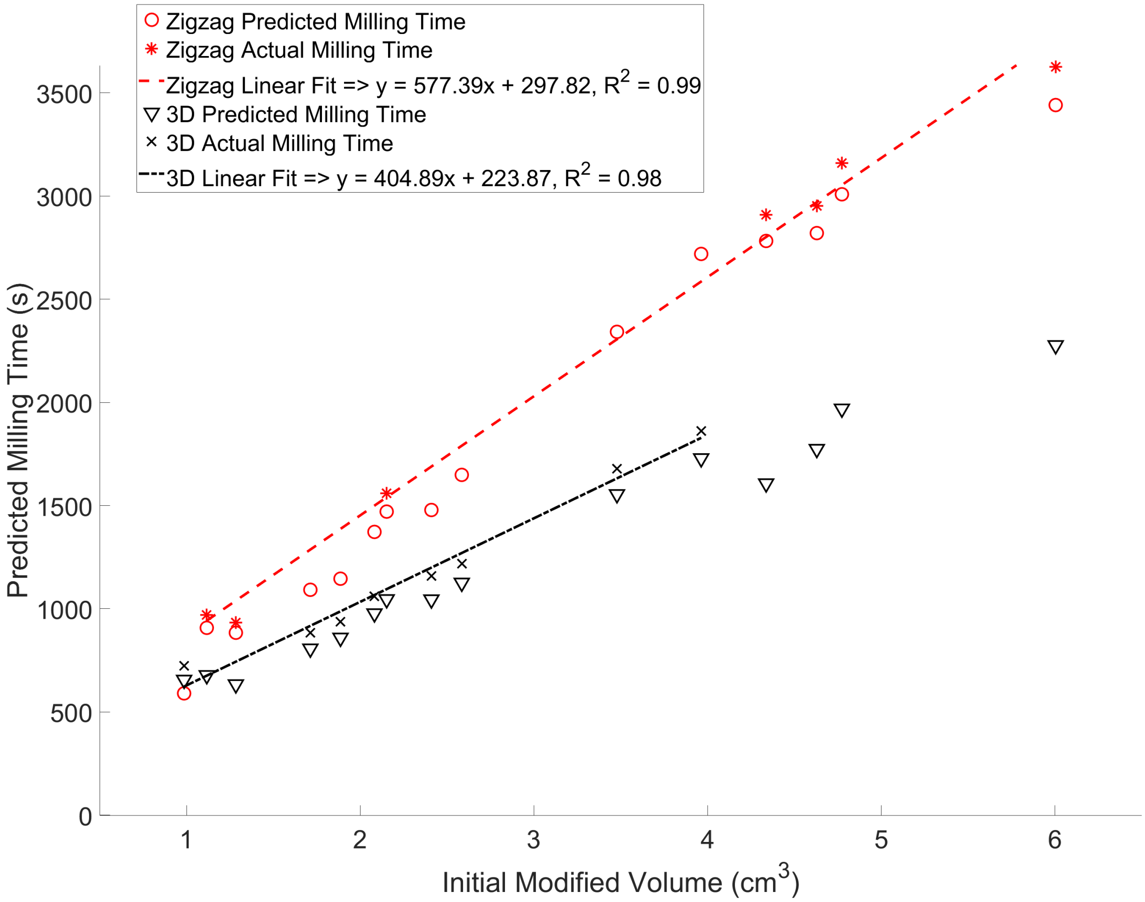

Figure 2.21: Similarly to Figure 2.20, for each algorithm the relationship between experimental milling times and initial modified target volumes is shown. Note that the fitted curves are only interpolations of the experimental data.

the data reported in Table 2.3 that relates the predicted milling times to the experimental milling times and compares them to the initial modified target volume similarly to Figure 2.20 can be seen in Figure 2.21.

From the data reported in these two figures, a comparison may be made between the simulated versus experimental relationship between initial modified target volume and the time taken to mill the pocket. Both slopes are larger than the simulations predicted. The milling times for the zigzag planner are increasing at a rate of approximately 577 s/cm^3 for the experimental data and the 3D planner's are increasing at a rate of approximately 405 s/cm^3. The percent difference between the zigzag and 3D planners simulated and experimental rates of increase are 8.4% and 18.6% respectively. The percent difference for the 3D path planner is approximately double that of the zigzag planner which is similar to the discrepancy between the percent difference of the average predicted and experimental milling times reported in Table 2.3.

2.5.3 Conclusions and Future Work

This chapter has shown that while zigzag planners are common in bone milling applications, they are not well suited to the irregular pocket shapes present in the application of a mastoidectomy. The computational time was higher for the 3D path planner. Given that the average predicted milling time of the true 3D 4 DOF planed paths was approximately 30.7% shorter than those generated the zigzag path planner for the exact same data set, the extra computation seems justifiable.

This chapter has presented a robust voxelized framework that allows a wide variety of tool path planning algorithms to implemented in image-guided surgery with the same level of accuracy as the medical scan that is used for guidance. A true 3D tool path planner for bone milling using this framework has also been implemented and evaluated. The algorithm's efficacy in removing the desired tissue via a shorter path than a simple 2.5D zigzag planner has been demonstrated. The benefits of such an algorithm go beyond straightforward path length and milling time comparisons. Using this path planner, it is now possible to incorporate any modeled quantity relevant to the procedure into the milling plan. Additional factors determining the planned path could include modeling the force applied to the robot for each point based on the CT scan and choosing locations that minimize the force applied to the cutting burr on each step of the plan. These factors could also involve modeling regions around especially critical structures embedded in the bone to mill them in a manner that reduces the danger of violating those structures. The work of Dillon et al. [45] to determine optimal cutting angles for bone milling with surgical cutting burrs could be incorporated into the fitness metric to use the cutting burr more efficiently. What has been created in this voxelized framework and 3D algorithm is a novel and efficient tool path planning algorithm designed specifically for bone milling. This novel 3D path planner has made customizing the tool path planning algorithm for specific surgical applications far easier than it has been.

Chapter 3

Safety Margins in Robotic Bone Milling: From Registration Uncertainty to Statistically Safe Surgeries

This chapter is adapted from ”Safety margins in robotic bone milling: from registration uncertainty to statistically safe surgeries” published in The International Journal of Medical Robotics and Computer Assisted Surgery and has been reproduced with the permission of the publisher and my co-authors Neal P. Dillon, Loris Fichera, Robert F. Labadie, Robert J. Webster, III, and J. Michael Fitzpatrick.

3.1 Introduction

In this chapter, the focus will be on creating safety margins specific not only to the individual patient's anatomy, but also specific to the registration being used to guide the robotic bone milling procedure. The major benefit of this level of patient personalization of the preoperative planning is an increased confidence, both for the surgeon and the patient, that critical structures will be preserved throughout the duration of the procedure.

Bone milling was one of the first applications considered in surgical robotics [23–25] due to the similarity to computer-assisted manufacturing processes. Examples of systems that have been commercialized and used clinically include the ROBODOC (Curexo, Inc., Seoul, South Korea), RIO Robotic Arm Interactive (Mako Surgical Inc., Ft. Lauderdale, FL, USA), and the Computer Assisted Surgical Planning and Robotics (CASPAR) system (URS Ortho GMBH & Co. KG, Rastatt, Germany). These platforms are examples of two different ways that robotic bone milling can be accomplished. ROBODOC is a fully autonomous system that plans the cutting burr's path preoperatively before it is carried out. In contrast, RIO is a cooperatively controlled robotic arm that uses preoperative planning to enforce “no fly zones” called virtual fixtures or active constraints [41, 42], ensuring that the surgeon removes only the desired volume, while leaving the surgeon in control of the exact path the robot traverses. Both approaches rely on accurate registration between the preoperative images to the patient in the operating room. Any errors in registration make it so that locations in the surgeon's preoperative plan, which is made in CT image space, will not perfectly align with corresponding locations in the patient. As imaging resolution and registration techniques have improved, it has become possible to apply robotic bone drilling and milling to applications that require higher accuracy than the orthopedic surgeries addressed by the commercial systems mentioned above. Skull base surgeries are particularly well suited for the use of image guidance and robotics because of their high accuracy requirements and the frequency with which the procedures are performed and this has motivated several research groups to pursue such solutions [1, 38–40, 50–55]. One example procedure is mastoidectomy, i.e. bone removal in the mastoid portion of the temporal bone, to gain access to the underlying anatomy of the ear. Mastoidectomy is performed approximately 120,000 times annually in the United States alone (extrapolating the results of [56] to the current population and to include both inpatient and outpatient procedures). It is performed as a preliminary step in more complex procedures of the middle and inner ear and to remove abnormal bone tissue arising from mastoiditis, cholesteatoma or other diseases. Critical structures are present within or near the surgical field during mastoidectomy that must be preserved during milling, including the facial nerve (damage causes facial paralysis), the chorda tympani (damage causes impaired sense of taste), the sigmoid sinus (damage causes bleeding), and the internal auditory canal (contains the auditory, vestibular, and facial nerves which may be irreparably damaged leading to hearing loss, balance disorders, and/or facial paralysis). Avoiding damage to these critical structures is challenging. The surgeon must first locate them using visual and tactile feedback while milling nearby bone then carefully remove the necessary bone. An example of a complex otologic procedure that requires a mastoidectomy is acoustic neuroma tumor removal surgery. Acoustic neuromas, more properly called vestibular schwannomas, are benign tumors which arise within or in close proximity to the internal auditory canal. While benign, their growth causes local compressive damage including hearing loss, balance disturbances, facial paralysis, and – in rare cases when they grow to large sizes – death from intracranial complications (e.g. intracranial bleeding, compression of the brainstem suppressing respiratory drive). Treatment options include radiation therapy and surgical removal with the most common surgical approach being a translabyrinthine approach which consists of a mastoidectomy followed by deeper bone milling to remove the labyrinth, including the semicircular canals, allowing access to the internal auditory canal. Figure 3.1 shows a preoperative segmentation of a translabyrinthine approach with critical structures annotated. The translabyrinthine approach is a strong candidate for image-guided robotic assistance for two reasons. First, from a surgical perspective, mastoidectomy and labyrinthectomy involve bulk removal of bone and are currently done by a human surgeon over several hours to reach the internal auditory canal following which the surgeon must change tasks to delicately handle exposed neural tissue in removing the tumor. Second, from a technological standpoint, since the anatomy is rigid and thus does not deform relative to the preoperative image, a robot that is guided by a path planned in a preoperative image could be programmed to remove the necessary bone while avoiding the critical anatomy. This chapter focuses on improving safety in robotic bone milling for otologic surgery with a specific emphasis on mastoidectomy and acoustic neuroma tumor removal surgery. No method currently exists to provide statistical confidence that the inevitable small registration errors will not lead to accidental damage to critical structures. Thus, the purpose of this work is to describe an algorithm that can provide such statistical confidence for systems that make use of point-based registration. The algorithm provides this confidence by establishing safety margins of bone around each critical structure that will not be targeted in the planned path. With a stiff robot and tool, milling up to the boundary of these safety margins will enable the cutting burr to come as close to the critical structure as possible, based on registration error statistics and a surgeon-specified safety level for each structure.

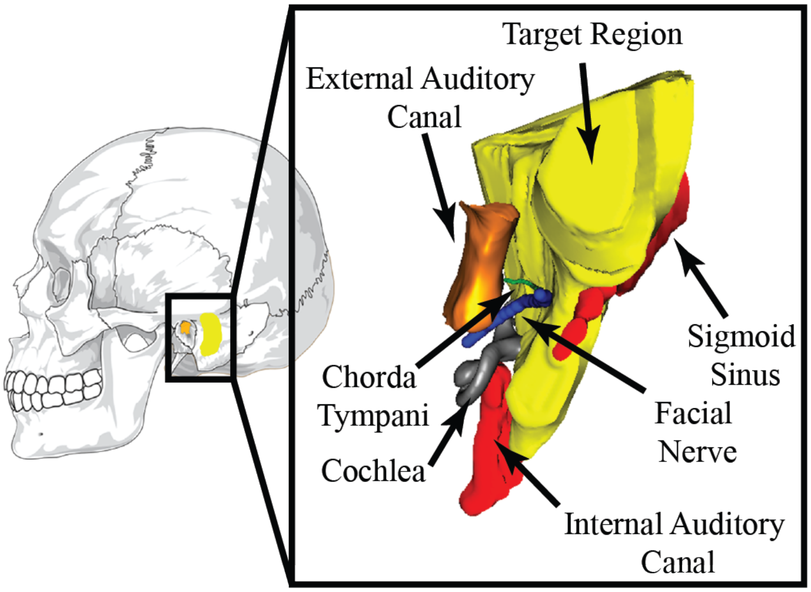

Figure 3.1: Mastoidectomy involves the removal of the bone volume highlighted in yellow. Several critical structures lie near the volume to be milled, including the facial nerve, chorda tympani, internal auditory canal, sigmoid sinus, etc. Damage to these critical structures causes complications for the patient. Note: the segmentation on the right has been rotated for ease of visualization.

For example, the surgeon might wish to be 99.9% sure that the system does not accidentally contact the facial nerve. These safety levels are ensured by choosing an appropriately sized and shaped safety margin, which is determined using numerical simulation and target registration error (TRE) theory [57]. Separate safety margins are obtained for each critical structure so that individual safety levels can be specified.

Examples of prior work that are most similar to that presented here are that of Haidegger et al. [5], and Noble et al. [19]. Haidegger et al. estimated the instantaneous level of danger to critical structures for an optically tracked tool [5]. In contrast, this chapter seeks a global approach that analyzes the entire procedure before any milling commences. Noble et al. estimated the danger to critical structures posed by a linear drill path for minimally invasive cochlea access [19]. In this chapter, the danger estimates are generalized to nonlinear tool paths. A preliminary version of this work was presented in [58]. The present work encompasses those results, which have not previously been published in archival form, and also extends them by providing a more accurate approach to collision detection and a more extensive numerical evaluation that includes simulations on five cadaver specimens, additional critical structures to avoid, and a comparison against constant thickness safety margins.

3.2 Materials and Methods

The workflow for robotic mastoidectomy previously developed by Dillon et al. [18] constitutes the general framework of this study. Briefly, the target volume is defined by the surgeon in a preoperative computed tomography (CT) scan and then used to generate a milling path for the robot. In the operating room, the preoperative plan is registered to the patient using point-based registration of bone-implanted fiducial markers that are localized in an intraoperative CT scan. This chapter's goal is to calculate safety margins around vital patient anatomy to limit the risk of accidental collisions with the robotic milling tool caused by registration error. Critical structures involved in a typical mastoidectomy procedure include the facial nerve, the chorda tympani, the sigmoid sinus and the internal auditory canal. Safety margins will be iteratively grown around these structures to define a region in which the robot will not be allowed to operate (Figure 3.2).

3.2.1 Algorithm Overview

The proposed algorithm takes as input a three-dimensional voxelized representation of the patient's anatomy wherein critical structures have been segmented. An example is shown in Figure 3.1. For each critical structure, individual spatially varying safety margins of minimal volume are determined using the four-step iterative process illustrated in Figure 3.3. The safety margin is initialized to include zero voxels, and is progressively expanded by including voxels that surround the critical structure. In the first step of the algorithm, a high number of imperfect registrations between the preoperative plan and the patient's anatomy in the operating room are simulated. These simulations are used to determine

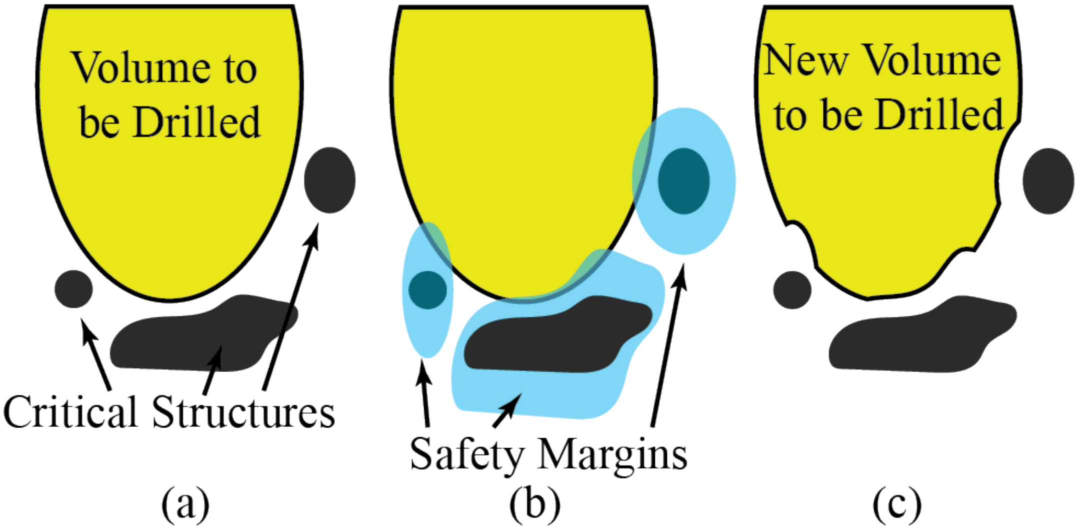

Figure 3.2: (a) An illustration of a segmented volume, with critical structures and a volume to be milled. (b) The same volume superimposed with safety margins. (c) The final result after the intersection of the safety margins and the volume to be milled has been removed from the volume to be milled. This reduced volume to be milled may now be milled with statistical assurance that the protected structures are safe from registration error.

the "overall damage risk" which is defined as the risk of accidental overlap between the critical structure and the voxels that surround the critical structure which, if milled, would cause damage to the critical structure. Registrations are simulated by generating fiducial locations from repeated sampling of the distribution of Fiducial Localization Error (FLE), which is a property of the imaging system and fiducial markers. If the Overall Damage Risk does not meet the tolerance level specified by the surgeon, the safety margin is grown by including neighboring voxels. Priority is given to those voxels that, if reached during the milling procedure, would pose a higher threat to the critical structure. Such a threat is quantified using a metric called "point damage risk". This process is iterated until the overall damage risk falls below the threshold specified by the surgeon. Table 3.1 contains a glossary of frequently used terms.

Term	**Definition**
Image Space	The coordinate system associated with the preoperative scan
Physical Space	The coordinate system associated with patient in the operating room
Fiducial Localization Error (FLE)	Positional error in the fiducial markers' locations caused by the inability to consistently determine their locations in both the physical and image spaces.
Target Registration Error (TRE)	The difference in the positions of an arbitrary point (not a fiducial location) in physical space and image space after registration.
Critical Structure	An anatomical structure in the patient that is also represented by a set of voxels in a medical image.
Safety Margin	A set of voxels surrounding the critical structure. This set is enlarged iteratively until the specified value of the overall damage risk is reached.
Overall Damage Risk	The probability that the critical structure in the patient will be damaged if all voxels surrounding the union of the critical structure and the safety margin were to be milled.
Neighboring Voxels	Voxels that share at least one common vertex. Each voxel has 26 neighbors.
Analyzed Shell	Voxels that share a voxel vertex with the union of the critical structure and the safety margin.
Point Damage Risk	The probability that the critical structure in the patient will be damaged if an individual voxel neighboring the analyzed shell were to be milled. This value is calculated for each voxel in the analyzed shell.
Transfer Percentage	The percentage of the analyzed shell that is transferred into the safety margin during a given iteration.

Table 3.1: A glossary of terms that are frequently used in this chapter.

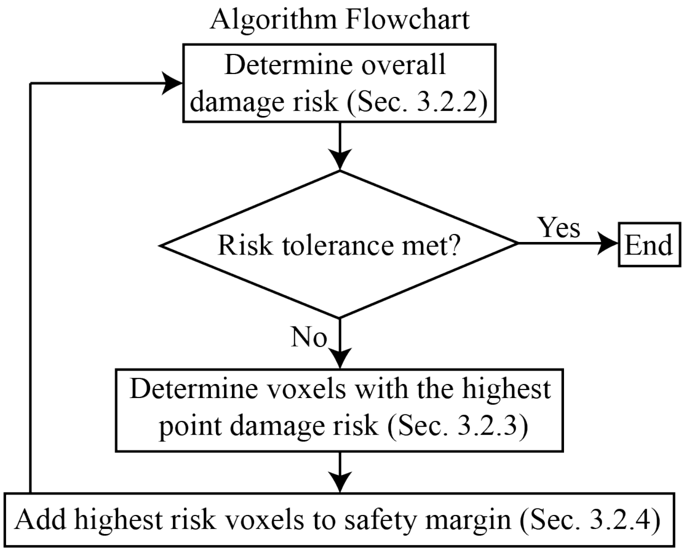

Figure 3.3: A flow chart is shown outlining the method employed to generate the safety margins surrounding the critical structures.

3.2.2 Overall Damage Risk

The overall damage risk is determined via numerical simulation. To perform this simulation, covariance matrices are estimated to describe the FLE distributions for each of the fiducial markers used during registration. These error distributions are then sampled to simulate imperfect fiducial locations that are used in a rigid registration to obtain a transformation from image space to physical space. This transformation is used to transform the voxels exterior to the safety margin from image space to physical space, where a check is performed for overlap with the critical structure. Note that because of the computational intensity of this approach, only those voxels directly neighboring the outer boundary of the safety margin (voxels are considered as neighboring if they share a vertex, and as a result each voxel has 26 neighbors) are considered. This outer shell of voxels is called the "analyzed shell" and is updated as the safety margin is "grown" through algorithm iterations. The overlap check is performed by discretizing faces of the voxels on the interior of the analyzed shell into a set of points (voxel corners). If, after registration, any point from this set lies within the critical structure, a collision has occurred. Figure 3.4 shows an example of such a registration. In Figure 3.4(a), the true critical structure is black and the safety margin is cyan. Figure 3.4(b) shows the set of points associated with the interior faces of the voxels in the analyzed shell. Figure 3.4(c) shows a 3D rendering and 2D slices taken from the rendering of Figure 3.4(a) and (b) registered together. Here, the green points are on the face of the analyzed shell and are external to the critical structure. The red points are on the face of the analyzed shell and are within the critical structure after registration. This simulation is repeated N_r (typically thousands) times. N_r is chosen by running the algorithm on a sample scan many times for various N_r values and recording the resulting damage probabilities. An acceptable N_r value is found when the variation of the trials falls below a given threshold (i.e. a value is selected that yields consistent results but is not overly computationally intensive). The Overall Damage Risk is finally calculated as the fraction of those registrations containing overlap between the critical structure and the analyzed shell.

3.2.3 Point Damage Risk

If the overall damage risk is above the safety level originally specified by the surgeon, additional voxels will be included in the safety margin. One simple approach would be to include the voxels closest to the critical structure. However, because of the anisotropic nature of the target registration error (TRE) this would produce a sub-optimal safety margin, i.e. a margin with uniform thickness but highly varying risk at the margin boundary. Here a different approach is used, in which each voxel is individually evaluated based on the probability that, after registration, the critical structure would be damaged if the voxel was targeted by the robotic mill. This probability is the “point damage risk”. Given the FLE covariances described in Section 3.2.3, the approach of Danilchenko and Fitzpatrick is used to determine the covariance of the TRE distribution at any point [57]. This covariance can be used to determine point damage risk. This risk is determined by considering the smallest Mahalanobis distance between the point under consideration and a set of points

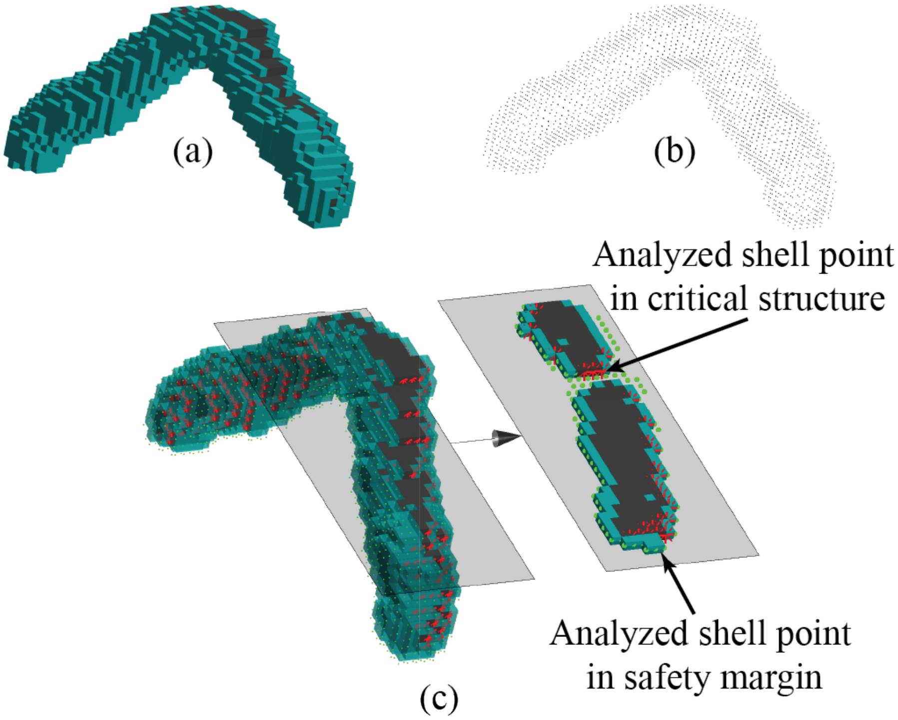

Figure 3.4: The effect of simulated registration error. (a) The true location of the critical structure (dark gray) and the safety margin (cyan) in physical space. (b) A set of points located on the faces of the analyzed shell voxels that are neighbors of the union of the critical structure and the safety margin in image space. (c) Superposition of (a) and (b) after they have been registered together. The error in registration from image space to physical space is due to fiducial localization error (FLE) that is added to the true fiducial positions in both spaces before registration. Red stars represent points on the analyzed shell that lie within the critical structure, and green dots represent points on the analyzed shell that are external to the critical structure. Computing many such simulated registrations and tabulating the fraction of registrations that contain at least one red point yields the "overall damage risk."

on the surface of the critical structure. Mahalanobis distance is a multidimensional generalization of the standard deviation and can be related to Euclidean distance by applying a "whitening transformation" (so-called because of an analogy with "white" noise in visible images), which is an affine transformation dependent on the TRE covariance matrix. The TRE distribution in the whitened space is isotropic so the Mahalanobis distance and Euclidean distances are equivalent. The probability that the point under consideration actually lies outside a sphere due to registration error (illustrated as a circle in Figure 3.5(c)) can be calculated by evaluating a three-degree-of-freedom Chi-Squared cumulative distribution function at the Mahalanobis distance squared. This computation yields the probability that despite registration error the point will remain within the sphere in any direction, not necessarily the direction that would cause damage to the critical structure. Thus, the complement of this probability is a conservative estimate of the point damage risk. Note that because of the anisotropic nature of the TRE distribution, the point damage risk of the point under consideration in Figure 3.5 is higher than would have been assumed looking at only the Euclidean distance (red line in Figure 3.5 (b)). As explained above, the point damage risk is inversely related to the shortest Mahalanobis distance between the point and the critical structure. This distance is calculated for each voxel in the analyzed shell and is used to rank the voxels based on their point damage risk. A visualization of the risk level of the voxels within an analyzed shell surrounding a critical structure and its safety margin is shown in Figure 3.6. It is important to note that even though these individual probability estimates are conservative, they will not lead to an oversized safety margin because the overall damage risk, calculated via the simulation method described in Section 3.2.2, is the final arbiter on how many voxels are included.

3.2.4 Growing the Safety Margin

The safety margin is expanded by transferring a percentage of the voxels in the analyzed shell associated with high risk to the critical structure (as identified in Section 3.2.3) into the

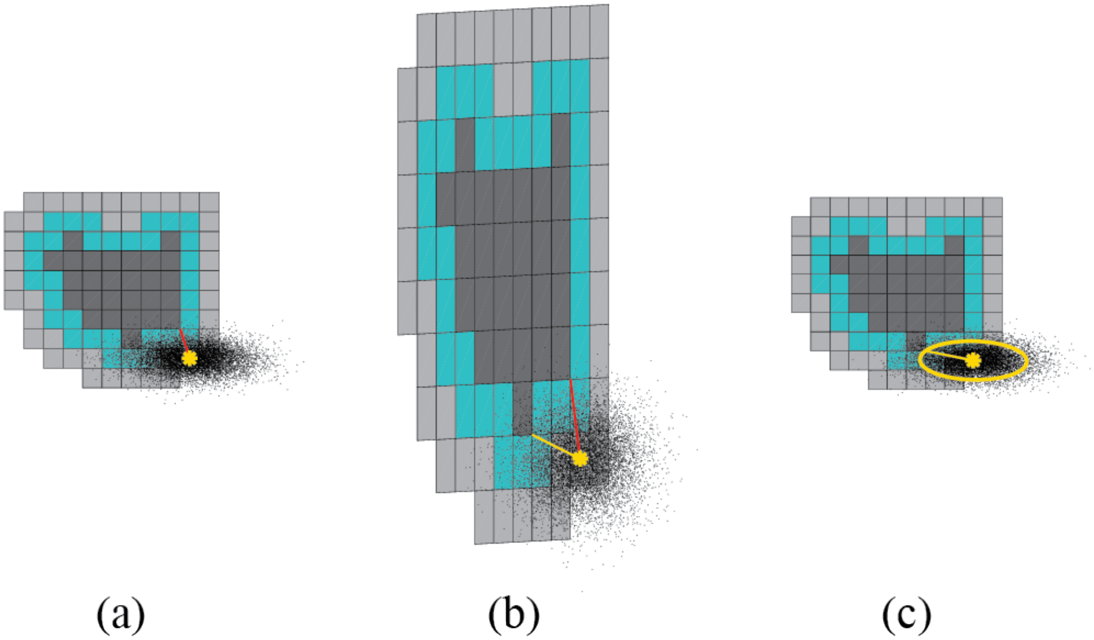

Figure 3.5: An illustration of the Mahalanobis distance calculation. (a) A critical structure (black), safety margin (cyan), and analyzed shell (gray) are shown. The risk to the critical structure of milling a specific point (yellow star) is examined. The shortest distance between it and the critical structure is shown in red. (b) "Whitened" space, formed by applying the whitening transformation. The shortest distance is shown in yellow. (c) In the whitened space, the probability that the point is within the yellow circle is easily obtained. The complement of this probability is a conservative estimate of the probability that the registration error could cause the center of a mill bit at the center of the circle to fall within the critical structure.

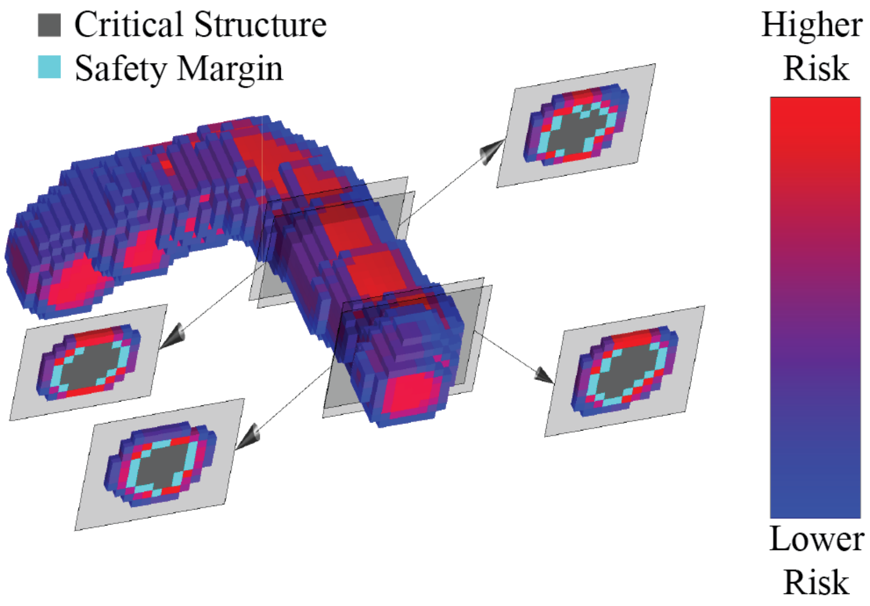

Figure 3.6: All the voxels surrounding a portion of a critical structure and its safety margin are shown. These voxels are color coded by the relative risk posed by each voxel to the critical structure in the patient, if that voxel location were to be milled. Several 2D cross sections of the region of interest can be seen on either side of the 3D rendering.

current safety margin. This transfer is repeated as illustrated in Figure 3.7. The percentage of high-risk voxels moved into the safety margin at each iteration (transfer rate) is the key parameter of this process: a low transfer rate results in a high number of algorithm iterations, thereby increasing computation time; by contrast, a high transfer rate ensures quick convergence of the algorithm, but may result in unnecessarily thick margins. To address this tradeoff, the transfer rate is adapted as the algorithm progresses: it is initially set to a fixed value and then linearly decreases as the overall damage risk approaches the value specified by the surgeon (as illustrated in Figure 3.8). It should be noted that the final result will always be conservative, since, while it is possible for slightly too many voxels to be transferred through this process, it is not possible for too few to be transferred to match the surgeon's desired safety threshold. Numerical simulations in Section 3.3.2 will show that the overshoot is small, when using the transfer percentage function illustrated in Figure 3.8.

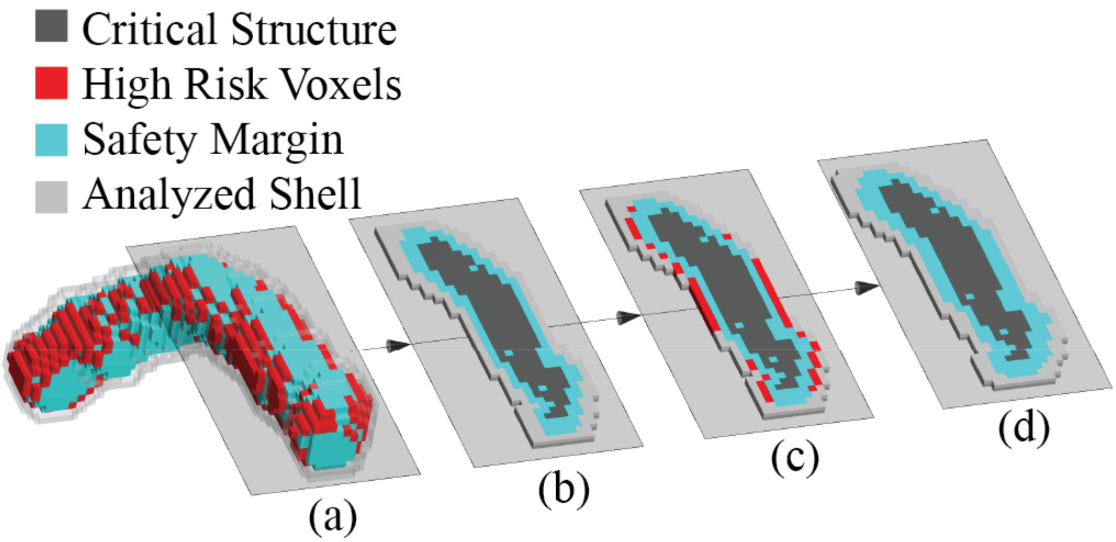

Figure 3.7: One iteration in the process of growing a safety margin around a critical structure is shown. The 2D slices are taken from the indicated plane of the 3D rendering (a). (b) 2D slice at the beginning of the iteration. (c) High-risk voxels in red surrounding the critical structure + current safety margin. (d) The high-risk voxels have been transferred to the safety margin, and the analyzed shell surrounding the union of the critical structure + safety margin has been updated. This result is the start of the next iteration.

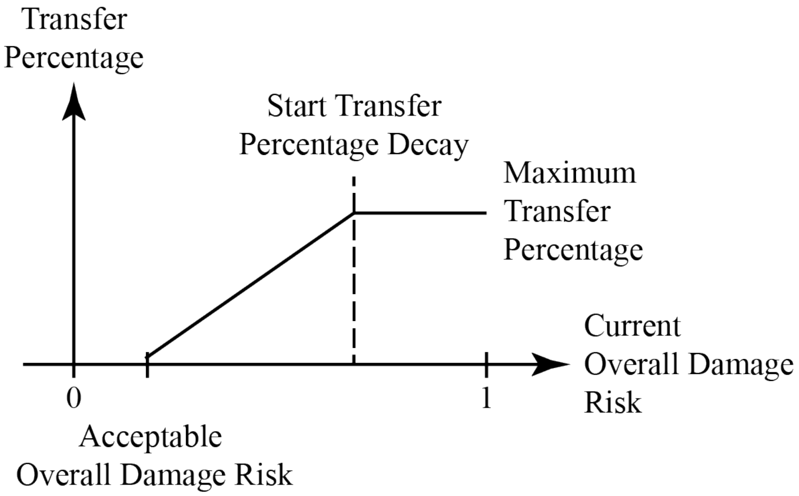

Figure 3.8: The linear interpolation by which transfer percentage is reduced is illustrated. Note that the overall damage risk is being reduced each time the algorithm iterates; therefore, the current overall damage risk moves from right to left along the figure's horizontal axis.

3.3 Results

3.3.1 Experiments

The algorithm was evaluated by applying it to the preoperative planning of mastoidectomy. The facial nerve chorda tympani, internal auditory canal, external auditory canal, and the sigmoid sinus are the critical structures for which safety margins are generated in this example. The procedure summarized in Figure 3.2.1 has been applied to five cadaver specimens. The scans were obtained using a xCAT ENT portable CT scanner (Xoran Technologies, Ann Arbor, MI, USA) whose voxel is a 0.4 mm cube with a scan volume of 640 x 640 x 355 voxels. For comparison, a second set of scans was created by upsampling each of the five scans such that the voxels became 0.2 mm cubes. The volume of bone to be removed, the internal auditory canal, and the sigmoid sinus were manually segmented, and the chorda tympani, facial nerve, and external auditory canal were automatically segmented via methods described by Noble et al. [59]. A configuration of six bone-implanted fiducials unique to each specimen was located roughly 20 mm above the volume of bone to be removed (Figure 3.9). These fiducials were localized in the image [60], and also serve as the attachment points for the robot to the patient [18, 52]. Fiducial localization errors were generated for the computer simulations by selecting error displacements from true fiducial positions from an isotropic distribution. The distribution was normal with zero mean and standard deviation equal to $0.176/\sqrt{3} = 0.1016$ mm, which produces a root-mean-square three-dimensional error length, FLE, of 0.176 mm. This value is the average of the results of a detailed analysis of FLE for several scanners and localization methods [61]. The localization error was also “homogeneous”, meaning that the same distribution was used for each fiducial. For each registration in Section 3.2.2, a random value was selected from this distribution for each component of each fiducial in each space. The 3D renderings of the initial segmentations can be seen in Figure 3.9. The parameters associated with two of the critical structures and their acceptable overall damage risks were taken from Noble et al., [19], as follows: facial nerve = 0.001, and the external auditory canal = 0.05. For the other structures, the acceptable overall damage risks were chosen with input from a surgeon as follows: chorda tympani = 0.05, internal auditory canal = 0.01, and the sigmoid sinus = 0.01. It was determined via simulation that the number of simulated registrations, N_r, required in order for these probabilities to be achieved to three decimal places was $N_r = 25000$. The initial transfer percentage (Figure 3.8) was 20%. The overall damage risk threshold values after which the transfer percentage begins to decay (Figure 3.8) were 0.3 for the facial nerve, 0.4 for the chorda tympani, 0.5 for the internal auditory canal, 1.0 for the sigmoid sinus, and 0.5 for the external auditory canal. To further validate the final safety margins generated by the algorithm, a shell of voxels that shared voxel faces with the union of the final safety margin and the critical structure was generated. The average and standard deviation of the minimum Mahalanobis distances between the centers of the voxels in this final shell and the same set of discretized points on exterior of the critical structure used in Section 3.2.3 were recorded. The algorithm was also run with the Mahalanobis distance in the point damage risk calculation in Section 3.2.3. replaced with a Euclidean distance. This caused the algorithm to generate safety margins that were uniformly distributed around their respective critical structures. These uniformly thick safety margins are then compared to the spatially varying safety margins developed by spatially varying algorithm.

3.3.2 Results

The results for each specimen can be found in Table 3.2. Figure 3.9 shows scan 4 before the algorithm is applied (solid black critical structures), and after the TRE safety margins have been generated (transparent cyan safety margins). The simulation was written in MATLAB (with the Mahalanobis Distance calculation in a mex file) and the workload was distributed in parallel among six CPUs using MATLAB's parallel loop processing facility. The simulations were run on a Dell Precision 5810 with a six-core 3.5 GHz Intel Xenon

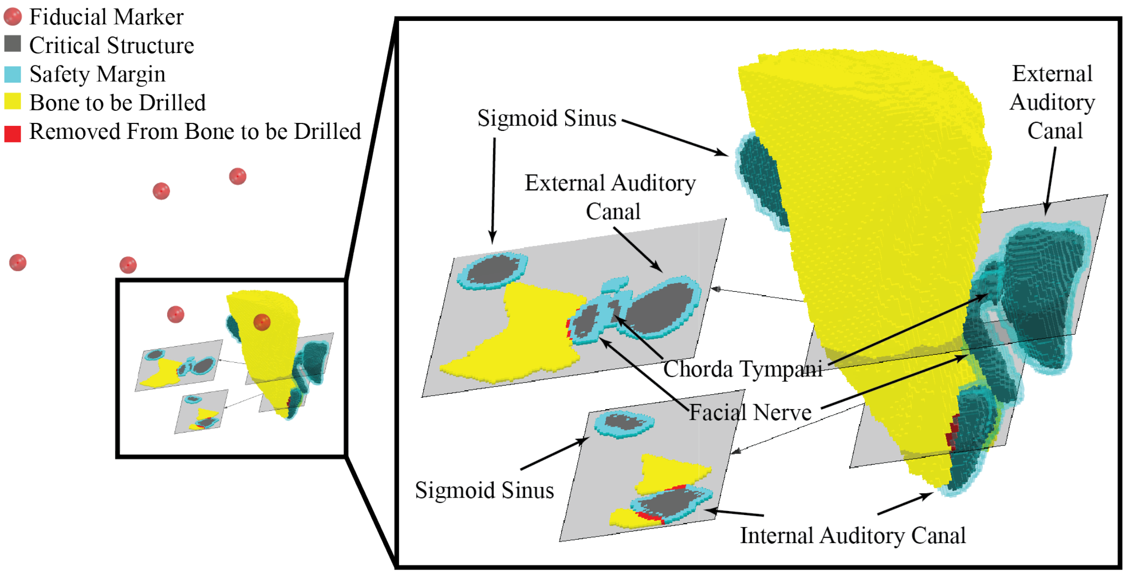

Figure 3.9: A cadaver scan is shown with six fiducial locations (red spheres), target region, critical structures, and the safety margins that the spatially varying algorithm generated. After the removal of the intersection of the safety margins and the target region, the resulting reduced target region can be milled with statistical assurance of safety from registration error.

processor and 16 GB of ram.

	Voxel Size 0.4^3 mm^3 Spatially Varying Safety Margins	Voxel Size 0.4^3 mm^3 Uniformly Thick Safety Margins	Voxel Size 0.2^3 mm^3 Spatially Varying Safety Margins	Voxel Size 0.2^3 mm^3 Uniformly Thick Safety Margins
Facial Nerve				
Volume of Safety Margin (mm^3)	229.13 ± 19.55	262.46 ± 21.45	140.18 ± 24.53	162.66 ± 15.25
Final Damage Probability (Acceptable = 0.001)	$6.56E-04 \pm 3.81E-04$	$4.64E-04 \pm 4.06E-04$	$6.88E-04 \pm 4.88E-04$	$6.16E-04 \pm 2.99E-04$

Average Mahalanobis Distance From Final Shell Voxels to Critical Structure	7.70 ± 0.49	8.95 ± 0.55	5.59 ± 0.82	6.71 ± 0.48
Standard Deviation of Mahalanobis Distances From Final Shell Voxels to Critical Structure	1.07 ± 0.14	2.40 ± 0.24	0.52 ± 0.04	1.77 ± 0.18
Chorda Tympani				
Volume of Safety Margin (mm^3)	69.17 ± 9.87	47.80 ± 7.21	24.94 ± 2.20	26.71 ± 3.53
Final Damage Probability (Acceptable = 0.05)	0.0105 ± 0.0269	0.0068 ± 0.0119	0.0486 ± 0.0019	0.0368 ± 0.0131
Average Mahalanobis Distance From Final Shell Voxels to Critical Structure	7.77 ± 0.61	6.52 ± 0.28	4.32 ± 0.09	4.89 ± 0.12
Standard Deviation of Mahalanobis Distances From Final Shell Voxels to Critical Structure	1.17 ± 0.20	1.94 ± 0.41	0.51 ± 0.11	1.42 ± 0.24

Internal Auditory Canal				
Volume of Safety Margin (mm^3)	227.39 ± 56.00	222.96 ± 55.25	142.39 ± 26.46	158.28 ± 27.24
Final Damage Probability	0.0012 ± 0.0014	0.0059 ± 0.0053	0.0078 ± 0.0054	0.0071 ± 0.005
Average Mahalanobis Distance From Final Shell Voxels to Critical Structure	7.34 ± 0.19	7.57 ± 1.55	4.84 ± 0.48	5.57 ± 0.59
Standard Deviation of Mahalanobis Distances From Final Shell Voxels to Critical Structure	1.15 ± 0.15	1.93 ± 0.57	0.55 ± 0.13	1.32 ± 0.35
Sigmoid Sinus				
Volume of Safety Margin (mm^3)	646.57 ± 184.37	661.26 ± 265.14	470.96 ± 154.14	514.44 ± 174.98
Final Damage Probability (Acceptable = 0.01)	0.0024 ± 0.0043	0.0051 ± 0.0054	0.0075 ± 0.0052	0.0034 ± 0.0038
Average Mahalanobis Distance From Final Shell Voxels to Critical Structure	7.35 ± 0.19	7.56 ± 1.29	5.24 ± 0.67	5.76 ± 0.38

Standard Deviation of Mahalanobis Distances From Final Shell Voxels to Critical Structure	1.03 ± 0.09	1.73 ± 0.45	0.56 ± 0.05	1.04 ± 0.14
External Auditory Canal				
Volume of Safety Margin (mm^3)	426.98 ± 65.60	347.70 ± 39.15	241.88 ± 29.13	262.82 ± 17.77
Final Overall Damage Risk (Acceptable = 0.05)	0.0203 ± 0.0283	0.0105 ± 0.0162	0.0477 ± 0.0022	0.0407 ± 0.0064
Average Mahalanobis Distance From Final Shell Voxels to Critical Structure	7.39 ± 0.46	6.56 ± 0.49	4.44 ± 0.30	4.99 ± 0.20
Standard Deviation of Mahalanobis Distances From Final Shell Voxels to Critical Structure	1.16 ± 0.18	1.93 ± 0.66	0.51 ± 0.12	1.22 ± 0.17
Time (min)	3.85 ± 1.79	14.87 ± 11.46	108.70 ± 40.76	137.49 ± 82.23

Table 3.2: The results (with 95% confidence intervals) from running the TRE compensation algorithm on five cadaver scans.

3.4 Discussion

These results indicate that, using the method described in this chapter, critical structures may now be given statistical safety guarantees with respect to registration uncertainty throughout robotic bone milling. This work is the first to generate safety margins that compensate for registration error by preserving critical structures (e.g. vasculature, nerves, etc.) to a specified safety level throughout the duration of a robotic surgical procedure. The safety levels were selected based on estimated complication rates in the conventional approaches as well as prior work in planning image-guided otologic surgery near vital anatomy. However, it is important to note that the values can be selected by the surgeon according to a variety of factors, including weighing the importance of the structure with the importance of removing nearby bone. Every application of the algorithm was successful in that for each structure in each cadaver scan the acceptable overall damage risk values were satisfied. The similarity of the final overall damage risk values to their predefined acceptable values demonstrates the effectiveness of varying the analyzed shell transfer percentage. The calculation of instantaneous tool damage probabilities is an important first step toward increasing patient safety, as so aptly pointed out by Haidegger et al [5]. However, it is noteworthy that the average point damage risks of the shell of voxels that share at least one face with the union of the final safety margin and critical structure are on the order of 10^{-7} or smaller. This value is several orders of magnitude smaller than any of the final overall damage probabilities. Such a large disparity indicates that determining the safety of a critical structure based solely on instantaneous tool damage probabilities may greatly underestimate the danger presented to the critical structure. This observation underscores the need to generate damage probabilities for the entire path traversed by the cutting burr, rather than relying on damage probability estimates that account only for the burr's instantaneous positions. The final overall damage risks of the upsampled regions are much closer to the acceptable overall damage risks. The standard deviation of the Mahalanobis distances from the final shell voxels to the critical structure from the upsampled scans are smaller than those calculated for the original scans. These observations imply that the higher resolution of the upsampled volumes enables the shape of the optimal safety margins to be more closely approximated than they are in the original scans. The significantly smaller standard deviations of the Mahalanobis distances in the spatially varying safety margins coupled with their smaller safety margin volumes implies that the spatially varying safety margins are more optimally shaped than the safety margins of uniform thickness. Therefore, upsampling the input medical scan is an effective method to improve the volumetric efficiency of the safety margins generated by this algorithm. In this work, the guarantees are based on the assumption that rigid point-based registration is the sole source of error. While this assumption is imperfect, other error sources whose statistics are known could be incorporated as well. These sources could include robot-specific physical errors such as calibration errors, joint positioning errors, system compliance, etc. Initial work toward modeling and incorporating these additional system errors into the safety margin algorithm is presented in [6]. The relative contribution of registration and other error sources is specific to a given system. Like registration error, other system error sources are typically spatially varying and anisotropic; thus, simulations like those described in this chapter are necessary to account for patient-specific conditions related to each error source. Note that all of these errors represent positional uncertainty and potential damage to the underlying structure via direct contact with the cutting burr. Damage to sensitive anatomy such as nerves can also occur as a result of excessive heat [62]. In mastoidectomy, the open cavity can be irrigated throughout the procedure, minimizing the risk of thermal damage. However, in other procedures, such as minimally invasive drilling for cochlear implantation, irrigation is difficult and thermal damage to nerves needs to be considered. In this case, the heat rise can be modeled in a manner similar to that of Feldmann et al. [63] and included in the generation of safety margins. While the use of this algorithm has now been demonstrated on a system for robotic mastoidectomy previously developed by Dillon et al. (Figure 3.10h), there are several other robotic systems designed for skull base surgery that could benefit from this approach as well. Examples include freestanding or table mounted systems that optically track either the entire procedure or a tool that is used to register the patient to the robot (Figure 3.10a-e) [1, 38–40, 51] and bone-attached robots (Figure 3.10f-h) [18, 50, 52]. These systems are designed to perform skull base bone milling for otologic surgery (Figure 3.10a,b,d,h) and neurosurgery (Figure 3.10), and linear drilling for minimally invasive cochlear implantation (Figure 3.10e-g). The accuracy and technical requirements of skull base bone milling and minimally-invasive cochlear implantation are similar since both involve removing bone in close proximity to critical structures to provide access to anatomy deep beneath the surface of the skull. These requirements suggest that any of these systems would benefit from the application of the algorithm presented in this work. All that is required to use the algorithm is segmentations of the vital anatomy in the medical image, knowledge of the locations of the fiducial markers relative to the anatomy, and estimates of the FLE, which can be used to generated estimates of registration error distributions near the anatomy. FLE of a given marker imaged with a particular scanner can be determined experimentally using geometric precise phantoms (see [61, 64]).

Additionally, the algorithm described in this chapter has the potential to be applied to other surgical procedures substantially different from those previously discussed. The algorithm can be applied to any procedure that is reliant on a registration method whose TRE may be estimated statistically, and it provides a way to convert surgeon-specified safety thresholds into non-uniform margin thicknesses that provide a statistical assurance of safety. For example, in the placement of a deep-brain stimulation electrode, registration is typically accomplished via either bone-implanted fiducial markers or surface registration. The algorithm presented in this chapter could be used in this application to ensure the preservation of critical structures such as those listed by Bériault: “surface veins, arteries running within the sulci, ventricles, critical motor and sensory cortices, and deep nuclei such as the caudate nucleus” [65]. Thermal ablation of cancerous tumors and needle placement for biopsy are also procedures that might benefit from this algorithm. It is also

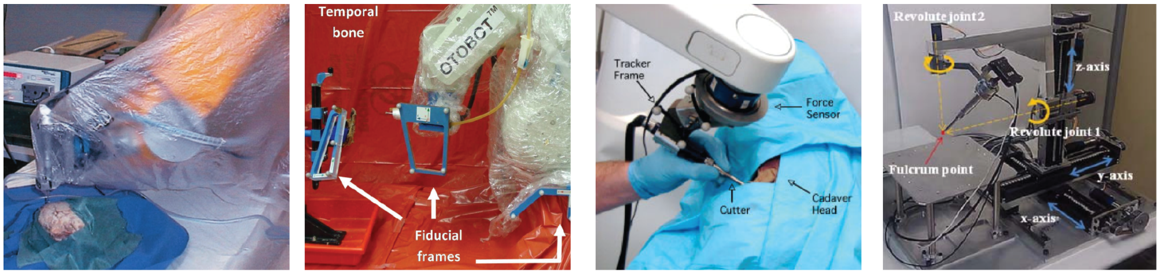

(a) Federspil et al. (b) Danilchenko et al. (c) Xia et al. (d) Lim et al.

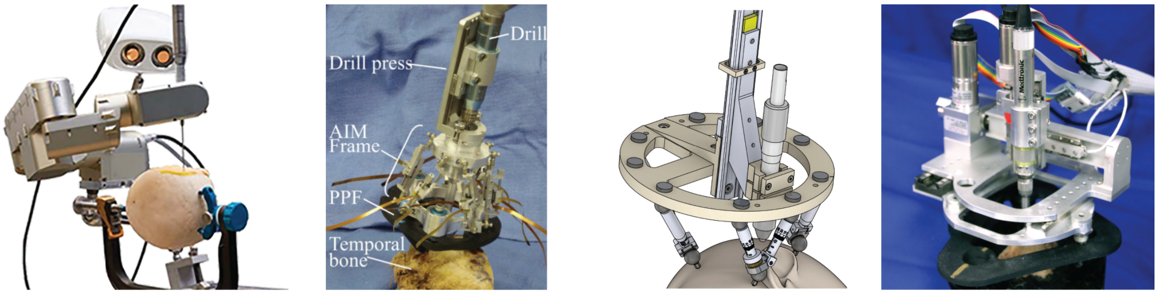

(e) Bell et al. (f) Kratchman et al. (g) Kobler et al. (h) Dillon et al.

Figure 3.10: Several robotic systems developed for skull base surgery. Systems (a-d) are free standing robots used to mill portions of the skull base under guidance of an external tracking system. Systems (a) and (b) are autonomous robots while systems (c) and (d) are cooperatively controlled by the surgeon. Systems (e)-(g) are robots that drill a tunnel through the mastoid for minimally invasive cochlear implantation. System (e) mounts to the patient's bed and is guided by an external tracking system while systems (f) and (g) are attached directly to the patient. Finally, system (h) is a bone-attached milling robot for mastoidectomy. All of these systems served as inspiration for the present work.

possible that this algorithm will enable the aforementioned procedures to be targeted in a safer manner by methods utilizing non-linear trajectories such as steerable needles. By developing this algorithm for generally shaped regions of interest, safety margins based only on intuition have been replaced with a statistically sound approach that has the potential to increase the safety in a broad range of surgeries.

Chapter 4

Registration Based Safety Margins in Robotic Partial Nephrectomy: From Intuition Based Margins to Statistically Safe Surgeries

4.1 Introduction

This chapter takes the safety margins that were personalized to the individual patient and surgery in the last chapter and alters them in two important areas. First, the margins developed in the last chapter were only valid under fiducial point-based registrations. The algorithm presented in this chapter removes that restriction and will generate statistically valid safety margins for any registration modality. Second, this chapter reconceptualizes the purpose of the safety margin. In Chapter 3, the safety margins were put in place to prevent damage to critical structures throughout the duration of a bone milling procedure. In this chapter, the safety margin concept will be extended to include a volume of tissue that must be entirely removed. Specifically, the margins in this chapter will be grown surrounding an individual patient's kidney tumor during a partial nephrectomy. The margins generated by the new algorithm will represent all of the tissue that must be removed for the surgeon to be confident that no positive margin was left in the patient.

The number of reported cases of renal cancer is dramatically rising. A 23% increase in diagnoses was reported globally between 1990 and 2013 [66]. The United States is estimated to have 73,820 new cases and 14,770 deaths in 2019 alone [67]. One of the several factors thought to contribute to this increase is more prevalent screening. This includes detection of renal masses while imaging the patient for other reasons. These incidental renal cancer diagnoses account for up to 50 percent kidney cancer cases in western countries [66, 68]. These masses that are often found inadvertently and by increased access to medical imaging, they are often found earlier than if they were detected after the patient had become symptomatic. This early detection has led to an increase in incidence of small renal masses, where small renal masses are defined as having a largest cross sectional distance of 4 cm or less [68]. This increase in incidence of small renal masses has led to an increase in the number of cases eligible for partial nephrectomy as opposed to radical nephrectomy.

Surgery is an effective treatment for renal cancer. Radical nephrectomy, or the removal of the entire kidney, is the most common treatment [69]. Unfortunately, radical nephrectomy leaves patients with an increased risk of chronic kidney disease, elevated cardiovascular risk, and increased mortality [70, 71]. Partial nephrectomy, or the removal of a portion of the kidney including and surrounding the renal mass, mitigates these long term risks by sparing as much renal tissue as possible. Partial nephrectomy is an underutilized procedure because it is more difficult to perform in a minimally invasive setting, and because it is associated with a higher rate of such complications as urine leakage or postoperative hemorrhage [69, 72].

A major contributor to the complexity of laparoscopic partial nephrectomy is the difficulty of performing tasks that require complex manipulation laparoscopically (e.g. suturing). This technical difficulty has limited the adoption of this technique to relatively few highly skilled laparoscopic surgeons [72]. The da Vinci Surgical System's (Intuitive Surgical, Inc. Sunnyvale, CA) wristed tools and stereoscopic vision provide additional dexterity and visualization that prove useful in performing the complex tasks associated with laparoscopic partial nephrectomy [73]. This assistance in performing complex intracorporeal tasks is valuable because the need for extreme skill and specialization is significantly reduced. Link et al. reported procedure time continuing to decrease for a single surgeon after 178 cases of laparoscopic partial nephrectomy [72]. The authors suppose their result to be an underestimate of the actual learning curve because the complexity of the cases increased along with the volume. This laparoscopic learning curve is in striking contrast to Deane et al.'s study of one surgeon's initial 11 robotically assisted partial nephrectomy cases. In this study, both the procedure time and the hospital stay were reduced in the surgeon's initial 11 robot assisted cases when compared to an expert surgeon's 11 laparoscopic cases [74]. It should be noted that the surgeon performing the robotically assisted partial nephrectomies had performed more than 200 robot-assisted laparoscopic radical prostatectomies.

Robotic assistance also benefits the expert laparoscopic surgeon. Benway et al. performed a retrospective study of three expert surgeons at high volume centers, 118 consecutive laparoscopic partial nephrectomies, and 129 consecutive robotically assisted partial nephrectomies [75]. All robotically assisted partial nephrectomies were performed with the da Vinci robot. Even with expert surgeons performing both the traditional laparoscopic and robot assisted surgeries several advantages of robot assistance emerged: (1) Hospital stays were shorter for the robotically assisted cases (2.4 vs 2.7 days). (2) Tumor complexity had no effect on the overall operative time for the robotically assisted cases, while tumor complexity had a significant impact on the overall operative time for the laparoscopic cases. (3) Warm ischemic time was significantly lower (30%) for the robotically assisted cases. Perhaps the result reported by Benway et al. most pertinent to this study is that, though not statistically significant, 5 of the 129 (3.9%) robotically assisted cases yielded positive margins as opposed to 1 of 118 (0.8%) of the laparoscopic cases. The authors suppose that the higher precision available to the surgeon during a robotic procedure might lead the surgeon to attempt to resect a smaller margin than the surgeon would in a laparoscopic surgery.

Assuming that Benway et al.'s supposition is correct, this is an excellent example of the dangers of an increase in precision without an associated increase in accuracy. The most intuitive way to increase the accuracy of the procedure is to provide the surgeon with accurate visualization of the preoperative surgical plan during the procedure. Utilizing preoperative data is difficult in robotic assisted partial nephrectomy, requiring the surgeon to visualize 2D slices of preoperative images mentally in the operative field. Currently a robotically assisted surgeon examines 2D slices of preoperative imagery. During the procedure the surgeon attempts to mentally map the relevant information from the slices onto the surgical field with nothing more than occasionally looking at the slices to verify details of the renal mass and critical structures' subsurface locations.

The need for surgeon access to accurate real-time visualization of preoperative imaging has driven significant research interest in this area. Goals of these systems include increasing the safety of robotically assisted partial nephrectomy, and increasing the accessibility of the procedure to patients by reducing the burden on the surgeon [76–80]. Additionally, Medical Engineering and Discovery Lab (MEDLab) research collaborators at Vanderbilt University and Vanderbilt University Medical Center have done extensive examination of system components including registration requirements [81–83] and the accuracy of the da Vinci Surgical System [84, 85]. Methods exist to compensate for the relative inaccuracy of the encoders on the da Vinci robot's setup arms where necessary [84–86]. One component common to each of these systems is a registration from preoperative imaging to the surgical space. Registration error is inherent to registration.

This chapter presents a method to create statistically valid spatially varying safety margins that compensate for the error inherent to registration. These safety margins are patient specific and the acceptable level of risk that a structure protrudes beyond the boundary of its safety margin is specified by the surgeon prior to the generation of the safety margins. This chapter uses the partial nephrectomy procedure and iterative closest point (ICP) surface registration [87] for an example of the generation of these safety margins. However, the methods created here are applicable to any procedure requiring accurate spatial localization of critical structures defined in a preoperative plan. Additionally, this new algorithm may utilize any registration modality that can be simulated many times to generate the error distributions required by the algorithm. Practically, the limit placed on which registration modalities are available to this algorithm is set by the execution time of the required simulations.

Prior works similar to the methods presented here are Haidegger et al. [5], Noble et al. [19], Siebold et al. [34], and Dillon et al. [6]. Haidegger et al. estimated the instantaneous level of danger to critical structures for an optically tracked tool [5]. Nobel et al. estimated the probability of damage that a surgical drill traversing a linear path would damage critical

Term	**Definition**
Critical Structure	An anatomical structure in the patient that is also represented by a set of voxels in a medical image.
Safety Margin	A set of voxels surrounding the critical structure. This set is enlarged iteratively until the specified value of the overall damage risk is reached.
Overall Damage Risk	The probability that the critical structure in the patient will be damaged if all voxels surrounding the union of the critical structure and the safety margin were to be removed.
Analyzed Shell	Voxels that are exterior to the union of the safety margin and critical structure that have been overlapped by the simulated tumor registration process outlined in Section 4.2.4.
Point Damage Risk	The probability that the critical structure in the patient will be damaged if an individual voxel neighboring the union of the critical structure and the safety margin were to be removed. This value is calculated for each voxel in the analyzed shell.
Transfer Percentage	The percentage of the analyzed shell that is transferred into the safety margin during a given iteration.

Table 4.1: A glossary of terms that are frequently used in this chapter.

structures in close proximity to that path. Most similar to this work are Siebold et al. and Dillon et al. [6, 34]. Siebold et al. provides the inspiration for this work and outlines the development of a similar algorithm for a procedure based on point based fiducial registration in the context of mastoidectomy. Dillon et al. was an extension of Siebold et al. that considered additional error sources inherent to the robotic mastoidectomy robot designed by Dillon et al. [18]. It should be noted that a similar extension is possible for the algorithm presented in this chapter. Such an extension designed for a specific robot and procedure would allow all error sources that can be modeled to be included in the generation of the safety margins. A specific efficiency improvement to Siebold et al.'s method is also included in the presented algorithm. A glossary of terms used in this chapter is seen in Table 4.1.

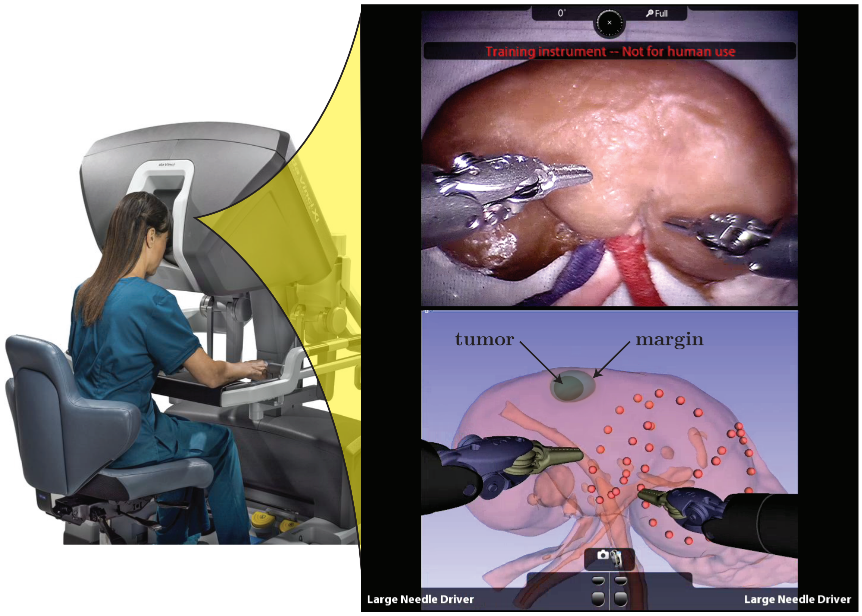

Figure 4.1: The upper half of the figure shows the da Vinci robot arms in relation to a phantom kidney with an implanted tumor. The lower half of the figure shows the display of the robot arms in relationship to the tumor and kidney segmentations from preoperative imaging.

4.2 Materials and Methods

Ferguson et al. in the Medical Engineering and Discovery Lab is developing a method for intraoperative visualization of the tumor and vasculature in a partial nephrectomy [83]. This project seeks to place a 3D visualization of the surgeon's tools interacting with the segmented preoperative 3D medical images (e.g. Computed Tomography (CT), Magnetic Resonance Imaging (MRI), etc.) in the surgeon's field of view as a real time simulation displayed in the da Vinci's surgical console via Tile-Pro.

The goal of this real time simulation is to enable the surgeon to quickly and accurately localize and resect the tumor. The algorithm described here generates safety margins surrounding the tumor and additional desired vasculature. These safety margins allow the surgeon to be confident that, in a procedure guided by these visualizations, the tumor will be removed with no positive margins and that all possible vasculature will be spared. These safety margins are illustrated in Figure 4.6. The portion of the basic workflow of the da

Vinci visualization project relevant to this algorithm is as follows: (1) Segment a medical image to show the tumor and basic vasculature. (2) Using the da Vinci manipulator, swab the exposed surface of the kidney. The points in $\mathbf{R}^3$ that the manipulator arm's end effector passes through during the surface swab are stored. This set of points is used to register the preoperative plan to the patient in the operating room. (3) The representation of the current location of the surgeon's instruments relative to the preoperatively segmented medical image is displayed in the da Vinci surgeon console via Tile-Pro.

4.2.1 Algorithm Overview

The algorithm presented here fits between steps (2) and (3) and takes as input: (1) the segmented 3D medical image, (2) the surface tracing of the kidney in the operating room, and (3) the acceptable risk levels associated with each critical structure (e.g. tumor, vasculature, etc.). The output of the algorithm is a medical image where the critical structures have been augmented with safety margins that ensure the risk posed to those structures does not rise above acceptable levels. These augmented critical structures are displayed to the surgeon in the da Vinci console via Tile-Pro. The algorithm proceeds as follows:

1. Perform a large set of registrations and store the resulting transformations (Section 4.2.2).
2. Determine overall damage risk (Section 4.2.3).
3. Check overall damage risk against surgeon defined risk tolerance for the critical structure.
4. Determine the voxels with the highest point damage risk (Section 4.2.4).
5. Add highest risk voxels to the safety margin (Section 4.2.5).
6. Repeat steps 2 - 5 until the overall damage risk has reached the acceptable level chosen by the surgeon.

The structure of this algorithm is similar to that presented by Siebold et al. [34]. Two major differences in the algorithm's structure are: (1) Moving the performance of registrations outside of the loop (Section 4.2.2) and storing the resulting homogeneous transformations for later use in Sections 4.2.3 and 4.2.4. This modification increases the efficiency of the algorithm. (2) A new method is required for estimating the risk level of the voxels surrounding the union of the safety margin and the critical structure (Section 4.2.4). This modification is required by the change in registration modality. Previously, the algorithm relied on a closed form solution to estimate the error covariance matrices of the individual voxels surrounding the critical structure and safety margin. This formula is valid only for rigid point-based fiducial registration. Extending the algorithm to utilize additional registration modalities requires an alternative method to estimate these risk levels.

4.2.2 Perform Registrations

Before the algorithm starts its iterative process, N_r registrations are performed and the resulting homogeneous transformations, $H_i \; \forall \; i \in [1,...,N_R]$, are stored for later use. These registrations are between the surface swab of the kidney, and the surface of the segmented kidney. The segmentation of the kidney is accomplished in 3D Slicer (www.slicer.org) [88, 89], a medical imaging software suite developed and maintained specifically for medical and medical robotics research. In 3D Slicer, the segmentation is converted into an STL, which is a file containing the surface mapping of the segmented volume consisting of a set of vertices on the segmentation's surface which are connected into triangles to represent faces of the surface. The vertices are extracted and used as the bases for an iterative closest point (ICP) registration with the surface swab of the kidney. This surface swab of the kidney is registered to the kidney in the medical image. The registered surface swab is then perturbed randomly in both position and orientation to simulate different initializations of the ICP algorithm. Each of these N_r perturbed swabs is registered to the point cloud obtained by randomly downsampling the vertices of the STL representing the surface of the kidney. Each of these registrations yields a homogeneous transformation, H_i. This collection of transformations will be used in Sections 4.2.3 and 4.2.4.

4.2.3 Overall Damage Risk

The overall damage risk is the risk of damage to the critical structure if the tissue contained in every voxel neighboring that structure were to be removed. This is a worst case damage metric. To calculate the overall damage risk, the critical structure is perturbed many times by simulating registration error. This perturbation is accomplished by applying one of the homogeneous transformations, H_i, calculated in Section 4.2.2 to a set of points lying on the surface of the critical structure. A check is then made to determine if any portion of this registered shell extends outside the critical structure and safety margin after the registration. This check for a critical structure lying outside of the safety margin after registration is illustrated in Figure 4.2. This process is repeated for all N_r of the H_i's that were stored in Section 4.2.2. A simple ratio is calculated of perturbed critical structures that protrude from the safety margin over the number of registrations. This ratio is the overall damage risk.

4.2.4 Point Damage Risk

The point damage risk is the risk that an individual voxel outside the critical structure and safety margin will overlap the critical structure after registration. The point damage risk is determined by performing a set of registrations similar to that used to determine the overall damage risk in Section 4.2.3. A set of discrete points from the exterior surface of the critical structure is sampled. The registrations' transformations are then applied to this set of points representing the surface of the critical structure. If any of these registered surface points then overlap voxels exterior to the safety margin and critical structure, then those overlapped voxels' coordinates are recorded and one violation is recorded for each overlapped voxel. We define N_p to be the number of registrations preformed to determine

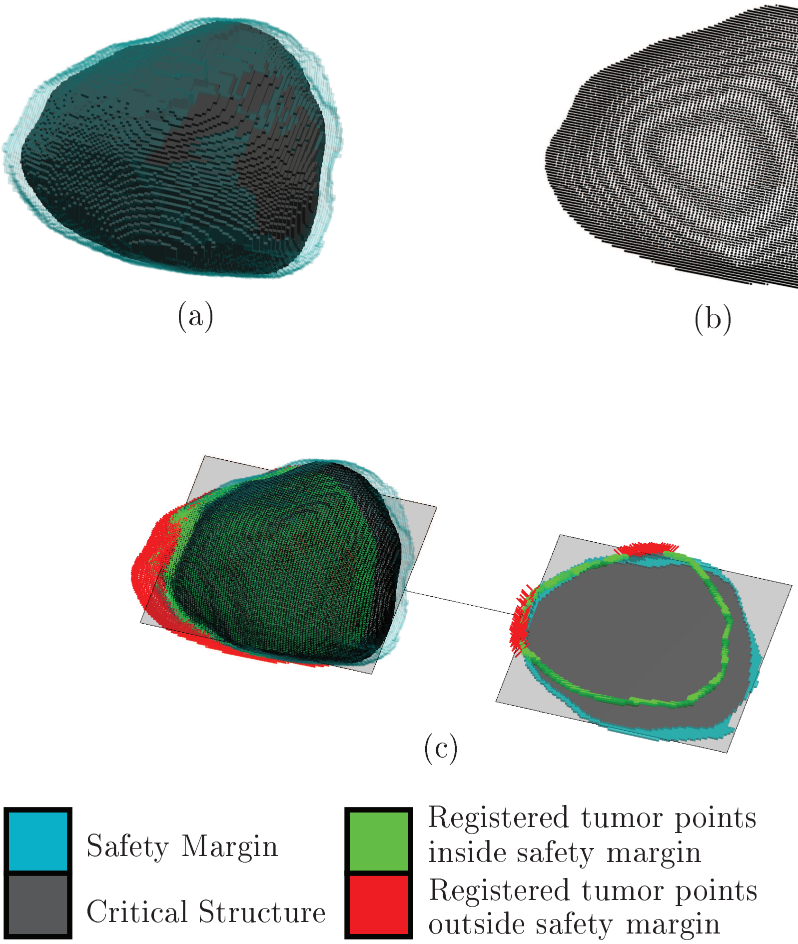

Figure 4.2: Part (a) shows a critical structure and its associated safety margin. Part (b) shows the collection of points on the surface of the critical structure. The left side of part (c) shows the two registered together and the protrusion of points on the registered surface of the critical structure outside of the safety margin. The right side of part (c) shows a one voxel thick layer pulled from the 3D portion on the right to better illustrate the registration.

the point damage risk. If the number of violated voxels does not outnumber the number of voxels to be transferred by a factor of R_s, the process is repeated for another N_p registrations until the number of violated voxels does outnumber the number of voxels to be transferred into the safety margin. Since the number of violations has been counted for each voxel, a ranking of voxels most dangerous to the critical structure has been established. By performing only a sufficient number of registrations to determine the voxels to be transferred into the safety margin one iteration at a time, voxels violated by outlying registrations could become part of the safety margin. Voxels covered by such outlying registrations would not pose a high enough point damage risk to be included in the safety margin were all the registrations required by the point damage risk calculation performed at once. Performing all of the registrations required by the point damage risk calculation is not possible since the total number of registrations required by the point damage risk calculation is unknown until the final result of the algorithm. Therefore, to prevent gross misregistrations from impacting which voxels are added to the safety margin, a threshold value, T_r is defined. This T_r value represents the number of simulated registrations in which the registered tumor must overlap a voxel before it can be included in the safety margin. Also, it is possible that by discretizing only the exterior of the critical structure some voxels immediately adjacent to the critical structure may have been overlooked. These voxels will be accounted for when the safety margin is grown in Section 4.2.5.

Note that this process results in an approximation of the point damage risks associated with voxels exterior to the critical structure and the safety margin. The overall damage risk is the metric that determines the ending of the growth of the safety margin. The point damage risk is only used to determine which voxels are moved into the safety margin next, so it has no direct affect on the final safety level of the safety margin. An example of the resulting ranked voxels surrounding the critical structure and safety margin is shown in Figure 4.3.

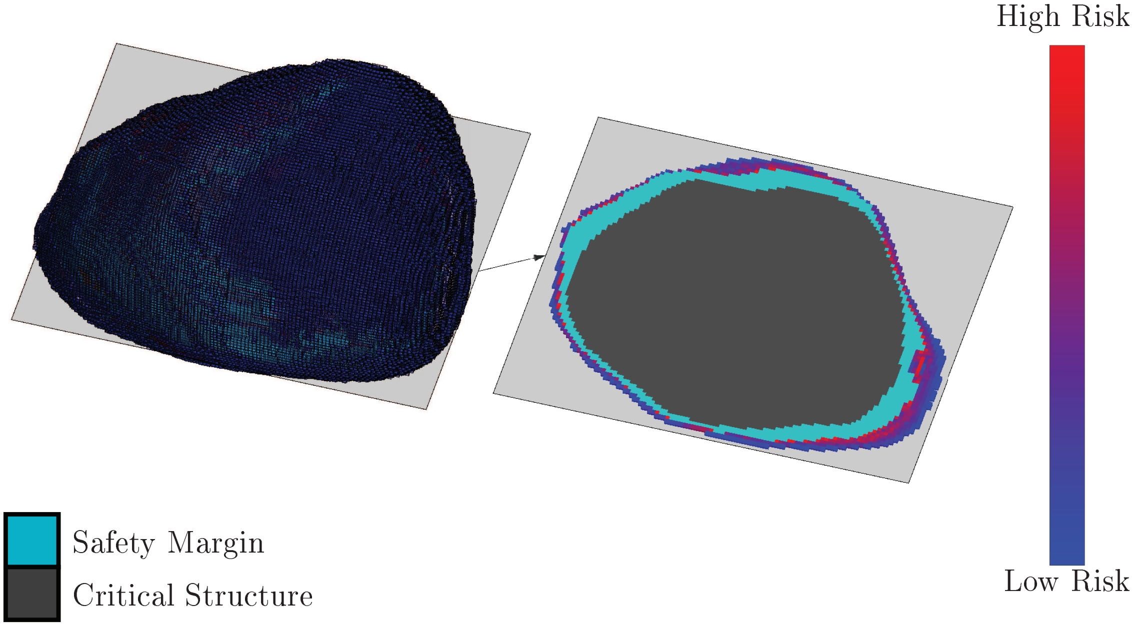

Figure 4.3: The figure shows an analyzed shell ranked by its point damage risk. The analyzed shell surrounds the safety margin which surrounds the critical structure.

4.2.5 Growing the Safety Margin

To grow the safety margin, a percentage of voxels ranked in Section 4.2.4 are transferred into the safety margin. The number of voxels transferred into the safety margin is a percentage of the total number of voxels immediately adjacent (26-connected) to the safety margin and critical structure. Prior to transfer, the voxels in the analyzed shell are sorted by point damage risk and only the voxels associated with the highest point damage risk are transferred in to the safety margin. A morphological close operation is performed on the union of the safety margin and critical structure to ensure that any voxels between the new safety margin voxels and the existing critical structure and safety margin are included in the updated safety margin. This process is illustrated in Figure 4.4.

To reduce the number of iterations that the algorithm must execute, a variable transfer percentage is used. This transfer percentage starts out relatively high and is reduced linearly as the overall damage risk approaches the acceptable damage risk. This linear decay of the

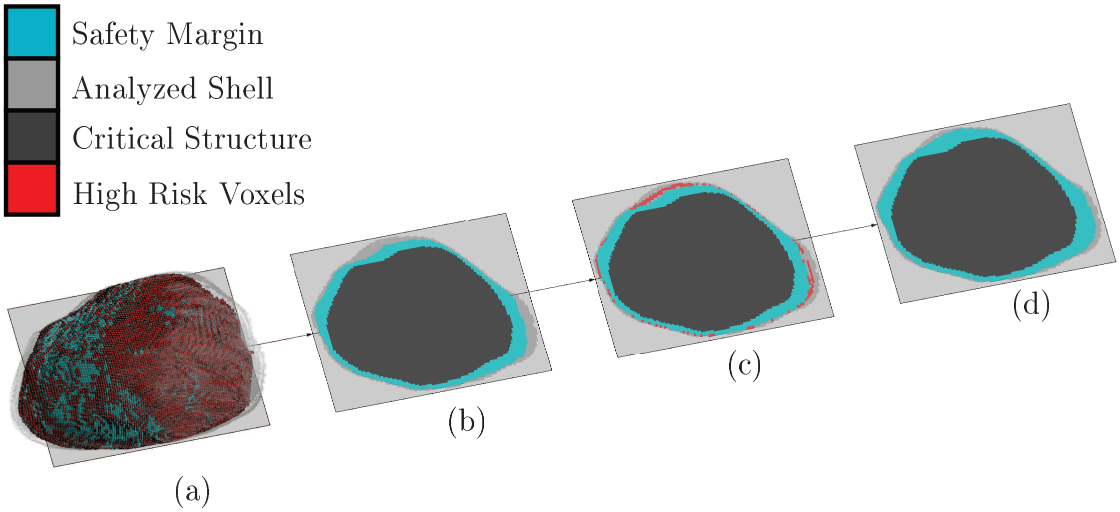

Figure 4.4: Part (a) shows a critical structure and safety margin with the highest risk voxels in the analyzed shell highlighted in red. Part (b) shows the state of the selected one voxel thick slice of Part (a) at the beginning of the algorithm's current iteration. Part (C) shows a slice of Part (a). Part (d) show the result of the current iteration of the algorithm where the high risk voxels from the analyzed shell have been transferred into the safety margin.

transfer percentage is shown in Figure 4.5.

4.3 Results

4.3.1 Experiments

To evaluate the algorithm, it was applied to the preoperative planning of four partial nephrectomies. The algorithm developed a safety margin surrounding the scanned tumor present in each kidney or phantom. This safety margin indicated minimum volume of tissue that must be removed to maintain the given confidence level that the tumor boundary will not be violated during the procedures. The tumors were manually segmented, and the patches used for registration were obtained by tracing the surface of the kidney with one of the arms of the da Vinci and recording the tool tip position throughout the tracing process. A portion covering roughly half of the top surface of each kidney was chosen as suggested by [81]. Initial swab poses were generated for the series of N_r registrations by

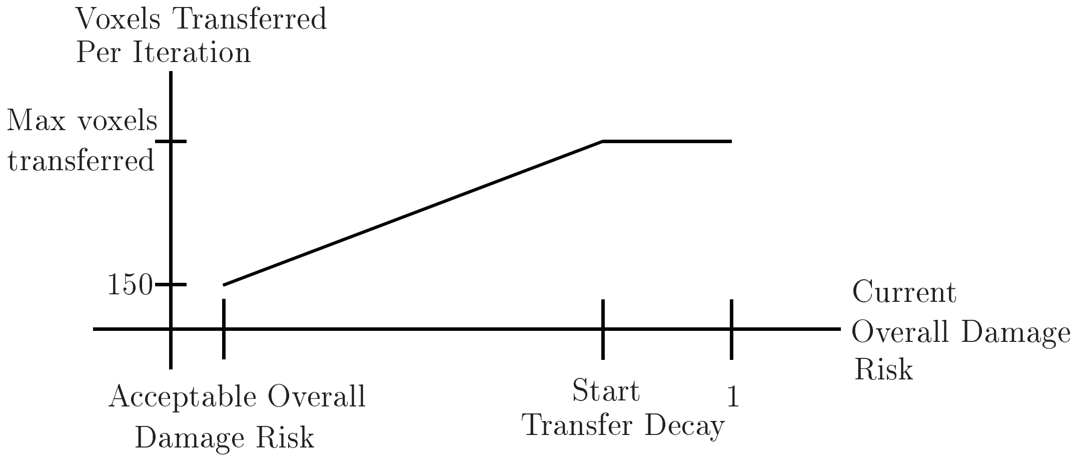

Figure 4.5: The figure shows the relationship between the number of voxels transferred into the safety margin each iteration and the current overall damage risk.

registering the point clouds obtained from the da Vinci's tool tip via GoICP [90]. GoICP is a variant of ICP that optimally registers a point cloud to a surface without needing a relatively accurate initial pose. These initial poses were perturbed by inserting normally distributed positional noise with a standard deviation of 5 mm in x, y, and z. Rotational error was introduced by estimating a normal vector for the center of the patch. Normally distributed rotations away from that vector were generated with a standard deviation of 5°. This value is chosen such that approximately 99.7% of the angular variation would fall with ±30° of the approximated normal. Additionally, safety margins were generated for a range of positional and rotational errors to test the sensitivity of the volumes output to the size of the error input. The range of standard deviation values tested was all combinations of 5 mm, 3.33 mm, and 2.5 mm of positional error and 5°, 3.33°, and 2.5° of rotational error. The results of these tests are reported in Appendix B. A final acceptable overall damage risk was chosen to be 0.01, N_p was chosen to be 5000, and the number of simulated registrations was

5000. The final overall damage risk was chosen to roughly agree with the rates observed by Benway et al. [75]. The number of simulations was chosen by generating a safety margin for each of the four specimens using $N_r = 35000$. The overall damage risk of this safety margin was tested for N_r values varying from 1000 to 10000 in steps of 1000. Each N_r was tested 50 times and 95% confidence intervals were calculated for the 50 overall damage risks generated for each N_r. An N_r equal to 5000 yielded a confidence interval of below $0.5e-3$ for each of the four specimen. The maximum transfer of voxels from the analyzed shell into the safety margin was 0.5 multiplied by the number of voxels that share a vertex (26-connected) with the exterior of the safety margin. The transfer rate was set to decay starting at an overall damage risk of 0.1. A minimum transfer rate of 150 voxels was used.

The results of the spatially varying safety margins are compared to constant thickness safety margins generated for the same kidney and tumor with the same surface swab. These safety margins of constant thickness are generated by applying the algorithm, but calculating the point damage risk as simply the closest Euclidean distance between each voxel surrounding the critical structure and the surface of the critical structure.

4.3.2 Results

The simulation results are reported in Table 4.2 and the simulated safety margin associated with tumor 2 can be seen in Figure 4.6. The bulk of the simulation was written in MATLAB while a non-parallelizable portion of the point damage risk calculation was written into a mex file for faster computation. A custom built 44-core dual socket Xeon E5-2696v4 machine with 128 GB of ram was used to compute the reported results. MATLAB's parallel toolbox was used to distribute the workload of the simulation onto all 44 available cores.

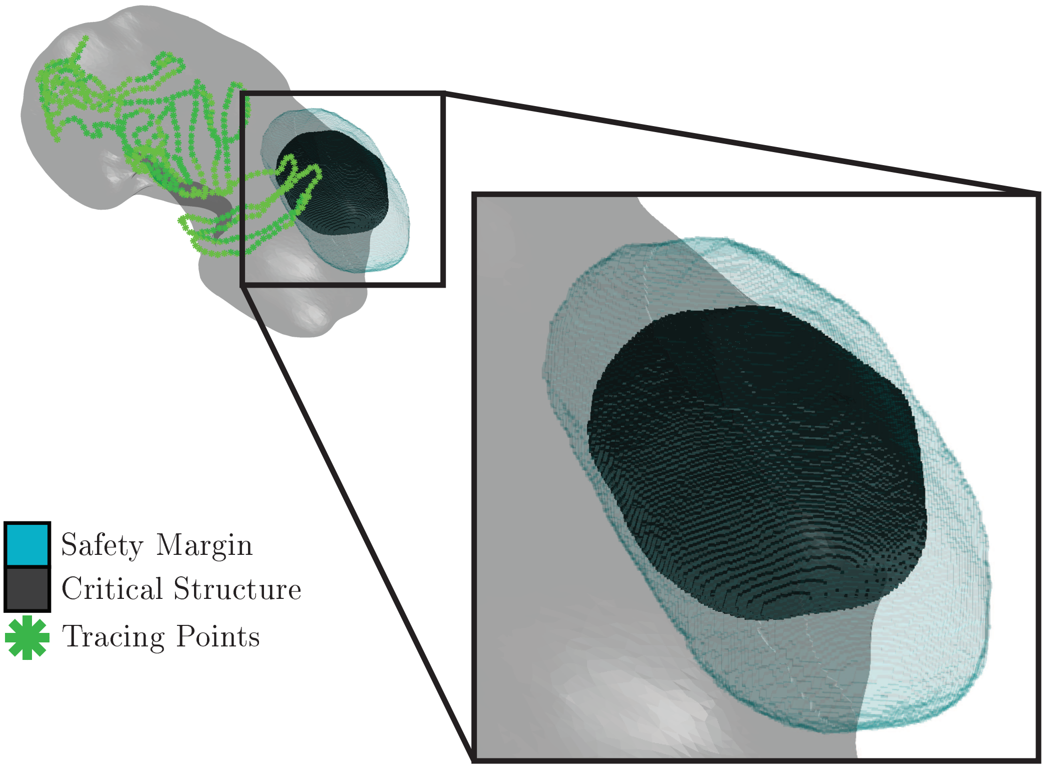

Figure 4.6: The figure shows a segmented preoperative scan of a kidney with a tumor placed at one end. The light blue region surrounding the tumor is the safety margin created by this algorithm. The green points are from the da Vinci tool tip tracing of the kidney's surface. These are the points used for the N_r registrations.

	Tumor 1		**Tumor 2**		**Tumor 3**		**Tumor 4**	
	SV	CT	SV	CT	SV	CT	SV	CT
Margin Volume (cm^3)	2.2e+4	6.3e+4	8.0e+3	1.3e+4	7.5e+4	1.3e+5	2.0e+4	3.9e+4
Final Damage Probability	1.1e-2	9.0e-3	9.4e-3	8.0e-3	1.0e-2	1.0e-2	1.0e-2	1.0e-2
Mean M-Dist to Surrounding Voxels	1.87	4.37	2.12	2.69	2.05	3.07	2.39	3.36
STDev of M-Dists to Surrounding Voxels	0.40	1.89	0.54	0.66	0.33	1.28	0.95	1.33
time (min)	35.4	61.3	17.5	4.8	100	110	28.6	53.9

Table 4.2: Numeric results for the spatially varying (SV) margins and the constant thickness (CT) margins created surrounding a tumor placed in a kidney. M-Dist refers to the Mahalanobis distance.

4.4 Discussion and Future Work

Assuming that the point damage risks are normally distributed, the results seen in Table 4.2 show that the standard deviations of the Mahalanobis distances are lower for the spatially varying safety margin. This implies that the point damage risks of the voxels immediately adjacent to the final safety margin are more consistent with each other for the margin created by the spatially varying algorithm as opposed to the uniform thickness margins.

As expected, all of the final spatially varying safety margins were smaller than their constant thickness counterparts. This is also true for the larger set of tests reported in Appendix B. It is interesting to note from the larger set of results reported in Appendix B, that the positional error has a greater impact on the volume of the safety margins than the rotational error. Beyond that observation the results from the additional testing were wholly predictable. All spatially varying margins were smaller than the corresponding uniformly thick margins. Also, all of the standard deviations of the Mahalanobis distances between the exterior of the tumor and the points exterior to the margin were smaller for the spatially varying margins. This implies that the level of danger posed by the voxels exterior to the spatially varying safety margins is more uniform than for the uniformly thick safety margins.

Safety margins were successfully generated surrounding both natural and simulated tumors in four CT scans of both human and phantom kidneys. This is the first time such spatially varying safety margins have been generated using surface registration. An important observation is that the algorithm easily incorporates any registration modality. The level of safety provided by this margin can be set as desired by the surgeon for the specific critical structure (e.g. lesion, vasculature, etc.). It should be noted that these safety margins were created under the assumption that registration error was the only source of error in the system. This assumption is a first step for generating general safety margins. Other error sources that can be modeled may be incorporated into the safety margins created by this algorithm. Dillon, et al. performed a similar extension of the fiducial based safety margins created in [34, 58]. Several additional error sources were modeled by the extended algorithm including error contributions from the CT scanner and physical robot. These error contributions were specific to the individual application of that paper. Similar extensions to this work are possible for any specific physical system with error sources that can be modeled.

Beyond simulation, Ferguson et al. in the Medical Engineering and Discovery Lab has created a platform for image guidance using the da Vinci robot [83]. A set of experiments where an expert user resects several phantoms with injected "lesions" has been planned. One set is to be performed with image guidance and no safety margins, the other set will be performed using both image guidance and the safety margins created by this algorithm. The results of this experiment are expected to demonstrate the utility of the margins created by this algorithm.

This chapter presents a robust algorithm to generate statistically guaranteed safety margins for any registration modality where repeated registrations may be assumed to be independent events. This method may be extended to incorporate other error sources specific to the systems to which it is applied. This algorithm is the first to provide the ability to create a system where the surgeon may specify the level of safety for specific structures unique to individual patients where the source of error is any registration modality. These safety guarantees are valid throughout the duration of the procedure.

Chapter 5

Future Work and Conclusions

The novel algorithms presented in this dissertation enhance the patient personalization of preoperative surgical planning. The specific areas of preoperative planning that are addressed here are: tool path planning for autonomous robotic bone milling, and the creation of optimal safety margins surrounding critical structure for a wide range of surgical applications. These algorithms were shown to be effective by applying them to the tool path planning of a robotically performed mastoidectomy, the creation of safety margins to inform the preoperative planning of the volume of bone to be removed in a mastoidectomy, and the creation of safety margins surrounding cancerous tumors to inform the surgeon's resection of healthy tissue in partial nephrectomy. In each case the algorithms performed the desired tasks yielding shorter tool paths and smaller safety margins when compared to standard approaches, such as a 2.5D tool path planning algorithm or a uniformly thick safety margin. The major contributions of this dissertation are found in autonomous robotic tool path planning and in automatic safety margin creation and are as follows:

Tool path planning: Up to this point, the primary methods for planning tool path trajectories have been layer by layer 2.5D planners [1–3, 13, 15]. The implementation of a 3D autonomous bone milling algorithm significantly complicates the interactions between the tool shaft and the unmilled tissue during the procedure. This dissertation has extended the 3D path planner presented by Danilchenko in [4] in two separate areas.

- The modeling and incorporation of drill shaft and cutting burr interactions with unmilled bone tissue at every point along a true 3D path (Chapter 2).
- The creation of a method to control the cutting depth of the true 3D planned path via a reapplication of the "supervoxel" concept. Controlling this cutting depth also shortens the generated paths such that they are able to be calculated quickly enough

to be generated intraoperatively (Chapter 2).

Safety margins: Both instantaneous danger posed to a critical structure by a tracked tool in an image-guided system in the presence of registration error [5], and the danger posed to a given critical structure by drilling a linear trajectory through bone in the presence of registration error [19] have been analyzed. This dissertation builds on these results by generalizing these safety analyses to extend to a critical patient structure of arbitrary shape for the duration of the procedure. First, safety margins are created for rigid point-based registration (Chapter 3, [34]). Second, safety margins are created for Iterative Closest Point surface registration (Chapter 4). This new method for building margins under surface registration error can incorporate any registration modality. Specific contributions are as follows.

- The creation of the first spatially varying safety margins enveloping critical structures based on a statistical analysis of fiducial point-based registration error. If preserved throughout the procedure, these safety margins provide a statistical guarantee that the contained structures will remain undamaged (Chapter 3, [34]).
- The creation of statistically guaranteed spatially varying safety margins via registration modalities other than fiducial point-based registration. The algorithm was simulated with surface based iterative closest point registration (Chapter 4).
- The reconceptualization of the safety margins to represent tissue that must be removed as opposed to protected. The specific application is a cancerous tumor that must be fully removed to prevent regrowth of the tumor.

5.1 Future Work in Path Planning for Autonomous Bone Milling

The new 3D tool path planning algorithm presented in this dissertation creates the opportunity to design many different approaches to the preoperative planning of image-guided robotic bone milling. These possibilities include planning a milling trajectory based on the ability of the cutting burr to perform more efficiently on that path. For more delicate procedures, it is possible to predict force required to pass the cutter through the unmilled bone such as is often simulated in haptic bone milling simulation [16, 17]. Such a prediction of required forces could be used to inform a trajectory that minimizes the force required for each step of the cutting burr's trajectory. For procedures that require work near easily damaged critical structures, a gradient of labels could be created that leads away from those structures. By utilizing that gradient in the planner algorithm's fitness metric the bone surrounding those critical structures could be milled last so that all forces exerted by the bone tissue on the cutter pointed away from the critical structure. Another possible use of the fitness metric would be to incorporate Dillon et al.'s work analyzing optimal milling angles in bone milling [45]. This work could be used in the fitness metric to mill the voxels in an order such that at every step the optimal milling angle was used. This would naturally lead to lower force throughout the milled path. It should also be possible to include Abdul et al.'s work in temperature modeling of bone milling [46] to create tool paths that optimize the temperature generated while milling the volume. The flexible metric that prioritizes each next step in the path can be used to incorporate any quantity that can be modeled into the tool path plan. This metric then opens up any number of future explorations that have been unachievable to this point.

5.2 Future Work in Spatially Varying Safety Margins

The novel algorithm to generate statistically guaranteed spatially varying safety margins in Chapter 3 was presented using only registration error and with only fiducial point-based registration. With the new algorithm presented in Chapter 4 the possibilities for additional applications increase many times over. Though this algorithm was evaluated with iterative closest point surface registration, safety margins can now be generated for any registration modality. Applying the algorithm to these additional modalities will greatly increase its impact on the safety of image-guided surgical procedures. Additionally, the interpretation of the safety margins may change. In this dissertation, safety margins have been used to prevent a drill from entering critical structures and to prevent a surgeon from crossing the boundary of a tumor during a partial nephrectomy. Another possible application for a safety margin grown surrounding a critical structure could be ensuring coverage of planned ablation zones during cancer treatment. The safety margin concept could be altered so that the margins are grown into the critical structure ensuring that any position within the margin was within the critical structure. This type of shrunken safety margin could ensure: a needle hits its desired target for tumor biopsy, accurate placement of radioactive seeds in the treatment of prostate cancer, or accurate needle placement in a Deep Brain Stimulation (DBS) surgery.

Experimental validation of this algorithm is planned to be performed using the image-guided partial nephrectomy system developed by Ferguson et al. [83]. These experiments would include several "lesions" that would be injected into phantoms. These phantoms would then have the simulated lesions removed under image guidance. Half of the simulated lesions would have safety margins grown around them via the methods described in Chapter 4. The other half would be encased in constant thickness margins. From preop and postop CT scans, the volume of removed tissue would be determined for each simulated lesion. This data would allow the determination of number of tumors with positive margins and the volume removed in each surgery. These results should validate the safety and volumetric savings presented in the simulation results in Section 4.3.2.

Additional error sources may be included in this algorithm. Dillon et al. [6] has already applied the algorithm presented in Chapter 3 to include 10 additional error sources in the generation of safety margins under rigid point-based registration for an autonomous mastoidectomy. Similar extensions can be created for a wide variety of systems. Now that such safety margins can be generated for any registration modality, the number of specific systems that could benefit from the safety margins generated in Chapter 4 increases dramatically. Any number of these systems will have additional error sources that could be included in the generation of the safety margins in a manner similar to Dillon et al. [6]. For example, given the ICP based registration used as an example in Chapter 4, it is possible to easily incorporate the error of the digitizer that generates the point cloud used in the registration. All one would need is a characterization of the error associated with the digitizer and sample that error distribution many times for each point. These sampled points could be used to run the algorithm in Chapter 4 many times, once for each set of sampled points. The union of the resulting many safety margins would constitute a safety margin that accounted not only for registration error, but for the physical error associated with tracing the kidney's surface.

Additionally, all applications of the algorithm attempted to this point have been under the assumption of a rigid surgical field. Future work with this algorithm could include modeled tissue deformation in the creation of the safety margins to account for the effect that deformation has on image guidance.

5.3 Conclusion

The main contributions of this dissertation are: In Chapter 2, the ability to model and incorporate tool shaft and cutting burr interaction with unmilled bone tissue in the creation of a true 3D tool path planner, and the ability to easily vary the cutting depth in a 3D path plan via a reapplication of the "supervoxel" from image and video processing. In Chapter 3, the creation of the first spatially varying safety margins that allow a surgeon to specify the acceptable level of risk a critical structure may be exposed to for the duration of a procedure. In Chapter 4, the creation of the first statistically guaranteed safety margins with registration modalities other than point based registration. Additionally, the safety margin was altered such that it represented a volume of tissue that must be removed rather than as a volume that must be persevered.

Each of these contributions directly advances the state of the art in preoperative patient specific image-guided robot surgical planning. Robotic image-guided surgery is an area that offers incredible benefits both to the surgeon and to the patient. It will continue to redefine what is possible in surgery by making what were once prohibitively complex procedures more widely available to patients. This availability is accomplished by easing the burden on the surgeon performing the procedure either by robotically performing delicate and physically taxing portions of the procedure or by offering easy and accurate intraoperative access to preoperative scan data in ways that have never before been possible.

BIBLIOGRAPHY

[1] Philipp A. Federspil, Urban W. Geisthoff, Dominik Henrich, and Peter K. Plinkert. Development of the First Force-Controlled Robot for Otoneurosurgery. *The Laryngoscope*, 113(3):465–471, 3 2003.

[2] Naohiko Sugita, Taiga Nakano, Takeharu Kato, Yoshikazu Nakajima, and Mamoru Mitsuishi. Tool path generator for bone machining in minimally invasive orthopedic surgery. *IEEE/ASME Transactions on Mechatronics*, 15(3):471–479, 2010.

[3] V. Cunha-Cruz, A. Follmann, A. Popovic, P. Bast, T. Wu, S. Heger, M. Engelhardt, K. Schmieder, and K. Radermacher. Robot- and computer-assisted craniotomy (CRANIO): From active systems to synergistic man-machine interaction. *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine*, 224(3):441–452, 2010.

[4] Andrei Danilchenko. *Fiducial-based registration with anisotropic error*. PhD thesis, Vanderbilt University, 2011.

[5] Tamas Haidegger, Sandor Gyori, Balazs Benyo, and Zoltan Benyo. Stochastic approach to error estimation for image-guided robotic systems. In *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology*, pages 984–987. IEEE, 8 2010.

[6] Neal P Dillon, Michael A Siebold, Jason E Mitchell, Gregoire S Blachon, Ramya Balachandran, J. Michael Fitzpatrick, and Robert J. Webster. Increasing safety of a robotic system for inner ear surgery using probabilistic error modeling near vital anatomy. In Robert J. Webster and Ziv R. Yaniv, editors, *Medical Imaging 2016: Image-Guided Procedures, Robotic Interventions, and Modeling*, volume 9786. SPIE, 3 2016.

[7] S. Marshall and J. G. Griffiths. A survey of cutter path construction techniques for milling machines. *International Journal of Production Research*, 32(12):2861–2877, 12 1994.

[8] Ali Lasemi, Deyi Xue, and Peihua Gu. Recent development in CNC machining of freeform surfaces: A state-of-the-art review. *CAD Computer Aided Design*, 42(7):641–654, 2010.

[9] Ramy F. Harik, Hu Gong, and Alain Bernard. 5-axis flank milling: A state-of-the-art review. *CAD Computer Aided Design*, 45(3):796–808, 2013.

[10] L W Kariuki, B W Ikua, and G N Nyakoe. Generation and Optimization of Pocket Milling Tool Paths - A Review. *International Conference on Sustainable Research and Innovation*, 5(May):129–133, 2014.

[11] Michael J. Lee and Eric L. Lin. The use of the three-pronged Mayfield head clamp resulting in an intracranial epidural hematoma in an adult patient. *European Spine Journal*, 19(SUPPL.2):187–189, 2010.

[12] A. Julianna Gulya. Anatomy of the Temporal Bone with Surgical Implications. *Otolaryngology- Head and Neck Surgery*, 1995.

[13] Russell H. Taylor, Brent D. Mittelstadt, Howard A. Paul, William Hanson, Peter Kazanzides, Joel F. Zuhars, Bill Williamson, Bela L. Musits, Edward Glassman, and William L. Bargar. An image-directed robotic system for precise orthopaedic surgery. *IEEE Transactions on Robotics and Automation*, 10(3):261–275, 6 1994.

[14] Nathan Abraham, Alon Wolf, and Howie Choset. A potential function approach to surface coverage for a surgical robot. *Computer Aided Surgery*, 11(1):1–9, 1 2006.

[15] Alon Wolf and Branislav Jaramaz. MBARS: Mini bone attached robotic system for joint arthroplasty. *Proceedings of the First IEEE/RAS-EMBS Interna-*

tional Conference on Biomedical Robotics and Biomechatronics, 2006, BioRob 2006, 2006(2):1053–1058, 2006.

[16] M. Arbabtafti, M. Moghaddam, A. Nahvi, M. Mahvash, and A. Rahimi. Haptic and visual rendering of virtual bone surgery: A physically realistic voxel-based approach. In *2008 IEEE International Workshop on Haptic Audio visual Environments and Games*, number October, pages 30–35. IEEE, 10 2008.

[17] Mohammadreza Arbabtafti, Majid Moghaddam, Ali Nahvi, Mohsen Mahvash, Barry Richardson, and Bijan Shirinzadeh. Physics-Based Haptic Simulation of Bone Machining. *IEEE Transactions on Haptics*, 4(1):39–50, 1 2011.

[18] Neal P. Dillon, Ramya Balachandran, J. Michael Fitzpatrick, Michael A. Siebold, Robert F. Labadie, George B. Wanna, Thomas J. Withrow, and Robert J. Webster. A Compact, Bone-Attached Robot for Mastoidectomy. *Journal of Medical Devices*, 9(3):031003, 9 2015.

[19] Jack H. Noble, Omid Majdani, Robert F. Labadie, Benoit Dawant, and J. Michael Fitzpatrick. Automatic determination of optimal linear drilling trajectories for cochlear access accounting for drill-positioning error. *The International Journal of Medical Robotics and Computer Assisted Surgery*, 6(3):281–290, 9 2010.

[20] Chenliang Xu and Jason J. Corso. Evaluation of super-voxel methods for early video processing. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1202–1209, 2012.

[21] Aurĺien Lucchi, Kevin Smith, Radhakrishna Achanta, Graham Knott, and Pascal Fua. Supervoxel-based segmentation of mitochondria in em image stacks with learned shape features. *IEEE Transactions on Medical Imaging*, 31(2):474–486, 2012.

[22] Suyog Dutt Jain and Kristen Grauman. Supervoxel-consistent foreground propagation in video. *Lecture Notes in Computer Science (including subseries Lecture Notes*

in Artificial Intelligence and Lecture Notes in Bioinformatics), 8692 LNCS(PART 4):656–671, 2014.

[23] Howard A. Paul, William L. Bargar, Brent Mittlestadt, Bela L. Musits, Russell H. Taylor, Peter Kazanzides, Joel Zuhars, Bill Williamson, and William Hanson. Development of a Surgical Robot for Cementless Total Hip Arthroplasty. *Clinical Orthopaedics and Related Research*, &NA;(285):57???66, 12 1992.

[24] S.C. Ho, R.D. Hibberd, and B.L. Davies. Robot assisted knee surgery. *IEEE Engineering in Medicine and Biology Magazine*, 14(3):292–300, 1995.

[25] Peter Kazanzides, Brent D. Mittelstadt, Bela L. Musits, William L. Bargar, Joel F. Zuhars, Bill Williamson, Phillip W. Cain, and Emily J. Carbone. An integrated system for cementless hip replacement. *IEEE Engineering in Medicine and Biology Magazine*, 14(3):307–313, 1995.

[26] Maria Bautista, Jorge Manrique, and William Hozack. Robotics in Total Knee Arthroplasty. *The Journal of Knee Surgery*, pages 1203–1211, 3 2019.

[27] Samik Banerjee, Jeffery J. Cherian, Randa K. Elmallah, Todd P. Pierce, Julio J. Jauregui, and Michael A. Mont. Robot-assisted total hip arthroplasty. *Expert Review of Medical Devices*, 13(1):47–56, 2016.

[28] M. Matinfar, C. Baird, A. Batouli, R. Clatterbuck, and Peter Kazanzides. Robot-assisted skull base surgery. *IEEE International Conference on Intelligent Robots and Systems*, pages 865–870, 2007.

[29] Nirmeen Zagzoog and Victor X.D. Yang. State of Robotic Mastoidectomy: Literature Review. *World Neurosurgery*, 116:347–351, 2018.

[30] Werner Siebert, Sabine Mai, Rudolf Kober, and Peter F Heeckt. Technique and first

clinical results of robot-assisted total knee replacement. *The Knee*, 9(3):173–180, 9 2002.

[31] Samik Banerjee, Jeffrey J. Cherian, Randa K. Elmallah, Julio J. Jauregui, Todd P. Pierce, and Michael A. Mont. Robotic-assisted knee arthroplasty. *Expert review of medical devices*, 12(6):727–35, 2015.

[32] Li-dong Wu, H. J. Hahne, and J. Hassenpflug. The dimensional accuracy of preparation of femoral cavity in cementless total hip arthroplasty. *Journal of Zhejiang University-SCIENCE A*, 5(10):1270–1278, 10 2004.

[33] P. Bast, A. Popovic, T. Wu, S. Heger, M. Engelhardt, W. Lauer, K. Radermacher, and K. Schmieder. Robot- and computer-assisted craniotomy: resection planning, implant modelling and robot safety. *The International Journal of Medical Robotics and Computer Assisted Surgery*, 2(2):168–178, 6 2006.

[34] Michael A. Siebold, Neal P. Dillon, Loris Fichera, Robert F. Labadie, Robert J. Webster, and J. Michael Fitzpatrick. Safety margins in robotic bone milling: from registration uncertainty to statistically safe surgeries. *The International Journal of Medical Robotics and Computer Assisted Surgery*, 13(3):e1773, 9 2017.

[35] B. L. Davies, F. M. Rodriguez y Baena, A. R.W. Barrett, M. P.S.F. Gomes, S. J. Harris, M. Jakopec, and J. P. Cobb. Robotic control in knee joint replacement surgery. *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine*, 221(1):71–80, 2007.

[36] A Wolf, S Amir, and A B Mor. MBARS2: A New Design for a 3- Degrees of Freedom Parallel Structure for Joint Arthroplasty. In *Volume 4: Fatigue and Fracture; Fluids Engineering; Heat Transfer; Mechatronics; Micro and Nano Technology; Optical Engineering; Robotics; Systems Engineering; Industrial Applications*, pages 445–451. ASME, 2008.

[37] B.L. Davies, K.L. Fan, R.D. Hibberd, M. Jakopec, and S.J. Harris. ACROBOT - using robots and surgeons synergistically in knee surgery. In *1997 8th International Conference on Advanced Robotics. Proceedings. ICAR'97*, pages 173–178. IEEE, 1997.

[38] Tian Xia, Clint Baird, George Jallo, Kathryn Hayes, Nobuyuki Nakajima, Nobuhiko Hata, and Peter Kazanzides. An integrated system for planning, navigation and robotic assistance for skull base surgery. *The International Journal of Medical Robotics and Computer Assisted Surgery*, 4(4):321–330, 12 2008.

[39] Hoon Lim, Jung-Min Han, Jaesung Hong, Byung-Ju Yi, Seung Hwan Lee, Jin Hyeok Jeong, Nozomu Matsumoto, Masamichi Oka, Shizuo Komune, and Makoto Hashizume. Image-guided robotic mastoidectomy using human-robot collaboration control. In *2011 IEEE International Conference on Mechatronics and Automation*, pages 549–554. IEEE, 8 2011.

[40] Andrei Danilchenko, Ramya Balachandran, Jenna L. Toennies, Stephan Baron, Benjamin Munske, J. Michael Fitzpatrick, Thomas J. Withrow, Robert J. Webster, and Robert F. Labadie. Robotic Mastoidectomy. *Otology & Neurotology*, 32(1):11–16, 1 2011.

[41] Stuart A. Bowyer, Brian L. Davies, and Ferdinando Rodriguez y Baena. Active Constraints/Virtual Fixtures: A Survey. *IEEE Transactions on Robotics*, 30(1):138–157, 2 2014.

[42] Jake J. Abbott, Panadda Marayong, and Allison M. Okamura. Haptic Virtual Fixtures for Robot-Assisted Manipulation. In *Robotics Research*, pages 49–64. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007.

[43] Yu Dai, Yuan Xue, and Jianxun Zhang. Milling State Identification Based on Vibration Sense of a Robotic Surgical System. *IEEE Transactions on Industrial Electronics*, 63(10):6184–6193, 2016.

[44] Wen Yo Lee and Ching Long Shih. Control and breakthrough detection of a three-axis robotic bone drilling system. *Mechatronics*, 16(2):73–84, 2006.

[45] Neal P. Dillon, Loris Fichera, Patrick S. Wellborn, Robert F. Labadie, and Robert J. Webster. Making robots mill bone more like human surgeons: Using bone density and anatomic information to mill safely and efficiently. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, volume 2016-Novem, pages 1837–1843. IEEE, 10 2016.

[46] Kais I.Abdul lateef Al-Abdullah, Hamid Abdi, Chee Peng Lim, and Wisam Yassin. Force and temperature modelling of bone milling using artificial neural networks. *Measurement: Journal of the International Measurement Confederation*, 116(August 2017):25–37, 2018.

[47] William A Hanson, Howard A. Paul, William Williamson, and Brent Mittlestadt. AN IMAGE DRIVEN. In *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, volume 12, pages 1931–1932, 1990.

[48] Elmer G. Gilbert, Daniel W. Johnson, and S. Sathiya Keerthi. A Fast Procedure for Computing the Distance Between Complex Objects in Three-Dimensional Space. *IEEE Journal on Robotics and Automation*, 1988.

[49] Neal P. Dillon, Ramya Balachandran, Michael A. Siebold, Robert J. Webster, George B. Wanna, and Robert F. Labadie. Cadaveric Testing of Robot-Assisted Access to the Internal Auditory Canal for Vestibular Schwannoma Removal. *Otology & Neurotology*, 38(3):441–447, 3 2017.

[50] Jan-Philipp Kobler, Kathrin Nuelle, G. Jakob Lexow, Thomas S. Rau, Omid Majdani, Lueder A. Kahrs, Jens Kotlarski, and Tobias Ortmaier. Configuration optimization and experimental accuracy evaluation of a bone-attached, parallel robot for

skull surgery. *International Journal of Computer Assisted Radiology and Surgery*, 11(3):421–436, 3 2016.

[51] Brett Bell, Christof Stieger, Nicolas Gerber, Andreas Arnold, Claude Nauer, Volkmar Hamacher, Martin Kompis, Lutz P. Nolte, Marco Caversaccio, and Stefan Weber. A self-developed and constructed robot for minimally invasive cochlear implantation. *Acta Oto-Laryngologica*, 132(4):355–360, 4 2012.

[52] Louis B. Kratchman, Grgoire S. Blachon, Thomas J. Withrow, Ramya Balachandran, Robert F. Labadie, and Robert J. Webster. Design of a Bone-Attached Parallel Robot for Percutaneous Cochlear Implantation. *IEEE Transactions on Biomedical Engineering*, 58(10):2904–2910, 10 2011.

[53] Jan-Philipp Kobler, Jens Kotlarski, Julian Öltjen, Stephan Baron, and Tobias Ortmaier. Design and analysis of a head-mounted parallel kinematic device for skull surgery. *International Journal of Computer Assisted Radiology and Surgery*, 7(1):137–149, 1 2012.

[54] Christof Stieger, Marco Caversaccio, Andreas Arnold, Guoyan Zheng, Jonas Salzmann, Daniel Widmer, Nicolas Gerber, M. Thurner, Claude Nauer, Yves Mussard, Martin Kompis, Lutz P. Nolte, Rolf Häusler, and Stefan Weber. Development of an auditory implant manipulator for minimally invasive surgical insertion of implantable hearing devices. *The Journal of Laryngology & Otology*, 125(3):262–270, 3 2011.

[55] Brett Bell, Nicolas Gerber, Tom Williamson, Kate Gavaghan, Wilhelm Wimmer, Marco Caversaccio, and Stefan Weber. In Vitro Accuracy Evaluation of Image-Guided Robot System for Direct Cochlear Access. *Otology & Neurotology*, 34(7):1284–1290, 9 2013.

[56] Lesley C. French, Mary S. Dietrich, and Robert F. Labadie. An estimate of the number

of mastoidectomy procedures performed annually in the United States. *Ear, Nose and Throat Journal*, 87(5), 2008.

[57] Andrei Danilchenko and J. Michael Fitzpatrick. General Approach to First-Order Error Prediction in Rigid Point Registration. *IEEE Transactions on Medical Imaging*, 30(3):679–693, 3 2011.

[58] Michael A Siebold, Neal P Dillon, Robert J. Webster, and J. Michael Fitzpatrick. Incorporating target registration error into robotic bone milling. In Robert J. Webster and Ziv R. Yaniv, editors, *Medical Imaging 2015: Image-Guided Procedures, Robotic Interventions, and Modeling*, volume 9415, Orlando, Florida, United States, 3 2015. SPIE.

[59] Jack H. Noble, Frank M. Warren, Robert F. Labadie, and Benoit M. Dawant. Automatic segmentation of the facial nerve and chorda tympani in CT images using spatially dependent feature values. *Medical Physics*, 35(12):5375–5384, 11 2008.

[60] Xiaofeng Liu, Hakan Cevikalp, and J. Michael Fitzpatrick. Marker orientation in fiducial registration. In Milan Sonka and J. Michael Fitzpatrick, editors, *Medical Imaging 2003: Image Processing*, volume 5032, San Diego, California, United States, 5 2003. SPIE.

[61] Jan-philipp Kobler, Jesus Díaz Díaz, J. Michael Fitzpatrick, G. Jakob Lexow, Omid Majdani, and Tobias Ortmaier. Localization accuracy of sphere fiducials in computed tomography images. In Ziv R. Yaniv and David R. Holmes, editors, *Medical Imaging 2014: Image-Guided Procedures, Robotic Interventions, and Modeling*, volume 9036. SPIE, 3 2014.

[62] Yi-chu Lin, Gianlorenzo Dionigi, Gregory W Randolph, I-cheng Lu, Pi-ying Chang, Shan-yin Tsai, Hoon Yub Kim, Hye Yoon Lee, Ralph P. Tufano, Hui Sun, Xiaoli Liu, Feng-Yu Chiang, and Che-Wei Wu. Electrophysiologic monitoring correlates of

recurrent laryngeal nerve heat thermal injury in a porcine model. *The Laryngoscope*, 125(8):E283–E290, 8 2015.

[63] Arne Feldmann, Juan Anso, Brett Bell, Tom Williamson, Kate Gavaghan, Nicolas Gerber, Helene Rohrbach, Stefan Weber, and Philippe Zysset. Temperature Prediction Model for Bone Drilling Based on Density Distribution and In Vivo Experiments for Minimally Invasive Robotic Cochlear Implantation. *Annals of Biomedical Engineering*, 44(5):1576–1586, 5 2016.

[64] A.D. Wiles and T.M. Peters. Real-Time Estimation of FLE Statistics for 3-D Tracking With Point-Based Registration. *IEEE Transactions on Medical Imaging*, 28(9):1384–1398, 9 2009.

[65] Silvain Bériault, Fahd Al Subaie, D. Louis Collins, Abbas F Sadikot, and G Bruce Pike. A multi-modal approach to computer-assisted deep brain stimulation trajectory planning. *International Journal of Computer Assisted Radiology and Surgery*, 7(5):687–704, 9 2012.

[66] Geolani W. Dy, John L. Gore, Mohammad H. Forouzanfar, Mohsen Naghavi, and Christina Fitzmaurice. Global Burden of Urologic Cancers, 1990–2013, 2017.

[67] Rebecca L Siegel, Kimberly D Miller, and Ahmedin Jemal. Cancer statistics, 2019. *CA: A Cancer Journal for Clinicians*, 69(1):7–34, 1 2019.

[68] Borje Ljungberg, Karim Bensalah, Steven Canfield, Saeed Dabestani, Fabian Hofmann, Milan Hora, Markus A. Kuczyk, Thomas Lam, Lorenzo Marconi, Axel S. Merseburger, Peter Mulders, Thomas Powles, Michael Staehler, Alessandro Volpe, and Axel Bex. EAU Guidelines on Renal Cell Carcinoma: 2014 Update. *European Urology*, 67(5):913–924, 5 2015.

[69] Steven Campbell, Robert G. Uzzo, Mohamad E. Allaf, Eric B. Bass, Jeffrey A. Cadeddu, Anthony Chang, Peter E. Clark, Brian J. Davis, Ithaar H. Derweesh, Leo

Giambarresi, Debra A. Gervais, Susie L. Hu, Brian R. Lane, Bradley C. Leibovich, and Philip M. Pierorazio. Renal Mass and Localized Renal Cancer: AUA Guideline. *Journal of Urology*, 198(3):520–529, 2017.

[70] R. Houston Thompson, Stephen A. Boorjian, Christine M. Lohse, Bradley C. Leibovich, Eugene D. Kwon, John C. Cheville, and Michael L. Blute. Radical Nephrectomy for pT1a Renal Masses May be Associated With Decreased Overall Survival Compared With Partial Nephrectomy. *Journal of Urology*, 179(2):468–473, 2 2008.

[71] James McKiernan, Rachel Simmons, Jared Katz, and Paul Russo. Natural history of chronic renal insufficiency after partial and radical nephrectomy. *Urology*, 59(6):816–820, 6 2002.

[72] Richard E. Link, Sam B. Bhayani, Mohammed E. Allaf, Ioannis Varkarakis, Takeshi Inagaki, Craig Rogers, Li-Ming Su, Thomas W. Jarrett, and Louis R. Kavoussi. Exploring the learning curve, pathological outcomes and perioperative morbidity of laparoscopic partial nephrectomy performed for renal mass. *Journal of Urology*, 173(5):1690–1694, 5 2005.

[73] Sanjeev Kaul, Rajesh Laungani, Richard Sarle, Hans Stricker, James Peabody, Ray Littleton, and Mani Menon. Da Vinci-Assisted Robotic Partial Nephrectomy: Technique and Results at a Mean of 15 Months of Follow-Up. *European Urology*, 51(1):186–192, 1 2007.

[74] Leslie A Deane, Hak J Lee, Geoffrey N Box, Ori Melamud, David S Yee, Jose Benito A Abraham, David S Finley, James F Borin, Elspeth M. McDougall, Ralph V Clayman, and David K Ornstein. Robotic versus Standard Laparoscopic Partial/Wedge Nephrectomy: A Comparison of Intraoperative and Perioperative Results from a Single Institution. *Journal of Endourology*, 22(5):947–952, 5 2008.

[75] Brian M. Benway, Sam B. Bhayani, Craig G. Rogers, Lori M. Dulabon, Manish N. Pa-

tel, Michael Lipkin, Agnes J. Wang, and Michael D. Stifelman. Robot Assisted Partial Nephrectomy Versus Laparoscopic Partial Nephrectomy for Renal Tumors: A Multi-Institutional Analysis of Perioperative Outcomes. *Journal of Urology*, 182(3):866–873, 9 2009.

[76] Asaki Hattori, Naoki Suzuki, Makoto Hashizume, Tomohiko Akahoshi, Kozo Konishi, Shohei Yamaguchi, Mitsuo Shimada, and Mitsuhiro Hayashibe. A robotic surgery system (da Vinci) with image guided function–system architecture and cholecystectomy application. *Studies in Health Technology and Informatics*, 94:110–116, 2003.

[77] Sam B. Bhayani and Devon C. Snow. Novel dynamic information integration during da Vinci robotic partial nephrectomy and radical nephrectomy. *Journal of Robotic Surgery*, 2(2):67–69, 7 2008.

[78] Li-Ming Su, Balazs P. Vagvolgyi, Rahul Agarwal, Carol E. Reiley, Russell H. Taylor, and Gregory D. Hager. Augmented Reality During Robot-assisted Laparoscopic Partial Nephrectomy: Toward Real-Time 3D-CT to Stereoscopic Video Registration. *Urology*, 73(4):896–900, 4 2009.

[79] S. Duke Herrell, David M. Kwartowitz, Paul M. Milhoua, and Robert L. Galloway. Toward Image Guided Robotic Surgery: System Validation. *Journal of Urology*, 181(2):783–790, 2 2009.

[80] Francesco Volonté, Nicolas C. Buchs, François Pugin, Joël Spaltenstein, Boris Schiltz, Minoa Jung, Monika Hagen, Osman Ratib, and Philippe Morel. Augmented reality to the rescue of the minimally invasive surgeon. The usefulness of the interposition of stereoscopic images in the Da Vinci™ robotic console. *The International Journal of Medical Robotics and Computer Assisted Surgery*, 9(3):e34–e38, 9 2013.

[81] Anne B. Benincasa, Logan W. Clements, S. Duke Herrell, Sam S. Chang, Michael S.

Cookson, and Robert L. Galloway. Feasibility study for image guided kidney surgery: assessment of required intraoperative surface for accurate image to physical space registrations. In Kevin R. Cleary and Robert L. Galloway, Jr., editors, *Medical Imaging 2006: Visualization, Image-Guided Procedures, and Display*, volume 6141. SPIE, 3 2006.

[82] Rowena E. Ong, Courtenay Glisson, Hernan Altamar, Davis Viprakasit, Peter Clark, S. Duke Herrell, and Robert L. Galloway. Intraprocedural Registration for Image-Guided Kidney Surgery. *IEEE/ASME Transactions on Mechatronics*, 15(6):847–852, 12 2010.

[83] James M. Ferguson, Leon Y. Cai, Alexander Reed, Michael Siebold, Smita De, S. Duke Herrell, and Robert J. Webster. Toward image-guided partial nephrectomy with the da Vinci robot: exploring surface acquisition methods for intraoperative re-registration. *Medical Imaging 2018: Image-Guided Procedures, Robotic Interventions, and Modeling*, 10576(615):8, 3 2018.

[84] David M. Kwartowitz, S. Duke Herrell, and Robert L. Galloway. Toward image-guided robotic surgery: determining intrinsic accuracy of the da Vinci robot. *International Journal of Computer Assisted Radiology and Surgery*, 1(3):157–165, 11 2006.

[85] David M. Kwartowitz, S. Duke Herrell, and Robert L. Galloway. Update: Toward image-guided robotic surgery: determining the intrinsic accuracy of the daVinci-S robot. *International Journal of Computer Assisted Radiology and Surgery*, 1(5):301–304, 3 2007.

[86] David M. Kwartowitz, Michael I. Miga, S. Duke Herrell, and Robert L. Galloway. Towards image guided robotic surgery: multi-arm tracking through hybrid localization. *International Journal of Computer Assisted Radiology and Surgery*, 4(3):281–286, 5 2009.

[87] Paul J. Besl and Neil D. McKay. Method for registration of 3-D shapes. In Paul S. Schenker, editor, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 1611, pages 586–606, 4 1992.

[88] Ron Kikinis, Steve D. Pieper, and Kirby G. Vosburgh. 3D Slicer: A Platform for Subject-Specific Image Analysis, Visualization, and Clinical Support. In *Intraoperative Imaging and Image-Guided Therapy*, pages 277–289. Springer New York, New York, NY, 2014.

[89] Andriy Fedorov, Reinhard Beichel, Jayashree Kalpathy-Cramer, Julien Finet, Jean-Christophe Fillion-Robin, Sonia Pujol, Christian Bauer, Dominique Jennings, Fiona Fennessy, Milan Sonka, John Buatti, Stephen Aylward, James V. Miller, Steve Pieper, and Ron Kikinis. 3D Slicer as an image computing platform for the Quantitative Imaging Network. *Magnetic Resonance Imaging*, 30(9):1323–1341, 11 2012.

[90] Jiaolong Yang, Hongdong Li, Dylan Campbell, and Yunde Jia. Go-ICP: A Globally Optimal Solution to 3D ICP Point-Set Registration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(11):2241–2254, 11 2016.

Appendices

Appendix A

Path Planning Full Results

Specimen	Predicted path length (mm)	Milling path length (mm)	Modified final targets / initial	Unmodified final targets / initial	Modified initial target (cm^3)	Simulation time (s)	Predicted milling time (s)	Average milling speed (mm / s)	Steps in contact / path length	Average voxels milled per step in contact	Ratio predicted path length (mm / cm^3)
CH01	491	491	7.8e-3	2.4e-1	2.19	93	1303	3.8	0.67	8.2	224
CH02	626	626	4.5e-3	1.3e-1	2.70	97	1628	3.8	0.72	8.1	232
CH03	451	451	3.5e-3	1.7e-1	1.50	67	1110	4.1	0.61	8.5	300
CH04	1072	1072	6.8e-3	1.5e-1	5.38	200	2783	3.9	0.63	7.8	199
CH05	1001	1001	4.6e-3	7.4e-2	4.42	209	2533	4.0	0.67	7.9	226
CH06	634	634	1.8e-3	1.1e-1	2.33	124	1568	4.0	0.65	8.1	273
CH07	1058	1058	1.8e-3	8.8e-2	4.67	228	2752	3.8	0.67	7.8	227
CH08	1049	1049	4.1e-3	9.2e-2	4.67	254	2821	3.7	0.66	7.6	224
CH09	705	705	2.7e-3	4.7e-2	3.15	110	1801	3.9	0.71	7.5	224
CH10	289	289	1.5e-3	1.9e-1	1.16	42	864	3.3	0.60	8.4	249
CH11	381	379	3.8e-3	1.4e-1	1.16	106	1013	3.7	0.62	8.9	329
CH12	1181	1181	3.2e-3	6.3e-2	5.88	231	3300	3.6	0.65	7.4	201
CH13	1112	1112	3.5e-3	5.0e-2	5.50	253	2721	4.1	0.66	7.2	202

CH14	1023	1023	2.4e-3	3.2e-2	4.78	188	2784	3.7	0.70	7.1	214
CH15	889	889	2.7e-3	3.7e-2	3.54	166	2412	3.7	0.62	7.3	251
CH16	949	949	1.5e-3	3.6e-2	4.23	181	2514	3.8	0.65	7.3	225
CH17	1000	1000	2.6e-3	4.2e-2	4.38	183	2655	3.8	0.66	7.4	228
CH18	663	619	6.7e-4	4.7e-2	2.49	105	1565	4.0	0.65	7.5	267
CH19	407	407	2.8e-2	1.3e-1	1.51	65	1066	3.8	0.65	8.2	269
TB01	949	949	3.7e-4	8.3e-2	4.34	144	2782	3.4	0.66	7.6	219
TB02	294	294	2.5e-2	3.5e-1	1.42	61	798	3.7	0.65	10.0	208
TB03	530	530	4.2e-3	6.1e-2	2.41	71	1479	3.6	0.71	8.7	220
TB04	407	407	5.9e-3	1.5e-1	1.71	56	1092	3.7	0.71	8.5	238
TB05	317	317	4.0e-3	2.4e-1	1.28	45	884	3.6	0.67	8.8	247
TB06	434	434	3.7e-3	1.6e-1	1.89	65	1146	3.8	0.70	8.3	230
TB07	517	508	5.7e-3	8.2e-2	2.08	71	1372	3.7	0.66	8.1	248
TB08	748	747	4.7e-3	9.7e-2	2.94	134	2051	3.6	0.66	7.9	254
TB09	552	552	9.3e-3	9.6e-2	2.15	77	1471	3.8	0.67	8.1	257
TB10	345	345	1.8e-2	1.2e-1	1.12	51	908	3.8	0.63	8.8	309
TB11	946	946	6.5e-3	9.0e-2	3.96	182	2720	3.5	0.65	8.1	239
TB12	983	983	1.9e-3	5.2e-2	4.63	187	2820	3.5	0.69	7.6	212
TB13	1069	1047	4.6e-3	7.9e-2	4.77	191	3009	3.5	0.65	7.9	224
TB14	217	217	2.5e-1	4.9e-1	0.99	33	590	3.7	0.61	9.0	220
TB15	596	596	2.4e-2	1.2e-1	2.58	99	1649	3.6	0.67	8.6	231
TB16	1252	1221	6.1e-3	7.3e-2	6.00	276	3441	3.5	0.65	7.6	209
TB17	850	850	1.4e-3	5.5e-2	3.48	153	2342	3.6	0.66	7.7	244
means	722	719	1.3e-02	1.2e-01	3.15	133	1937	3.7	0.66	8.0	238
’±	236	235	4.2e-03	3.9e-02	1.03	44	633	1.2	0.22	2.6	78

Table A.1: Simulation results for the baseline 4 DOF zigzag algorithm. Means are reported with 95% confidence intervals.

Specimen	Predicted path length (mm)	Milling path length (mm)	Modified final targets / initial	Unmodified final targets / initial	Modified initial target (cm^3)	Simulation time (s)	Predicted milling time (s)	Average milling speed (mm / s)	Steps in contact / path length	Average voxels milled per step in contact	Ratio predicted path length (mm / cm^3)
CH01	476	426	5.9e-4	2.3e-1	2.19	542	922	4.6	0.08	31.3	218
CH02	601	561	0	1.2e-1	2.70	457	1256	4.5	0.08	36.8	222
CH03	429	413	0	1.7e-1	1.50	287	848	4.9	0.08	28.0	286
CH04	951	884	2.4e-5	1.3e-1	5.38	954	2048	4.3	0.09	33.9	177
CH05	912	823	7.2e-5	7.2e-2	4.42	1025	1937	4.2	0.08	36.0	206
CH06	561	500	0	1.0e-1	2.33	460	1100	4.5	0.08	33.5	241
CH07	916	847	0	8.5e-2	4.67	820	1929	4.4	0.08	40.2	196
CH08	928	829	1.4e-4	8.9e-2	4.67	883	1858	4.5	0.08	33.2	199
CH09	658	586	0	4.3e-2	3.15	408	1293	4.5	0.08	32.4	209
CH10	271	238	0	1.9e-1	1.16	157	507	4.7	0.09	25.4	234
CH11	401	387	5.5e-5	1.4e-1	1.16	254	796	4.9	0.09	25.4	347
CH12	933	854	0	5.7e-2	5.88	1064	2159	4.0	0.10	36.0	158
CH13	935	878	0	4.6e-2	5.50	1277	2068	4.2	0.09	35.7	170
CH14	886	801	0	2.6e-2	4.78	2479	1869	4.3	0.09	32.7	185
CH15	716	664	0	3.4e-2	3.54	495	1510	4.4	0.08	34.9	202
CH16	798	721	0	3.2e-2	4.23	562	1662	4.3	0.08	36.2	189
CH17	850	756	0	3.9e-2	4.38	668	1787	4.2	0.08	35.7	194
CH18	608	553	0	4.4e-2	2.49	407	1218	4.5	0.08	27.7	244

CH19	433	405	2.8e-3	1.1e-1	1.51	537	850	4.8	0.07	28.2	286
TB01	839	748	3.0e-5	8.2e-2	4.34	584	1607	4.7	0.08	33.5	194
TB02	323	306	1.1e-3	3.5e-1	1.42	298	626	4.9	0.08	31.0	228
TB03	544	500	2.4e-4	5.8e-2	2.41	347	1045	4.8	0.10	28.9	226
TB04	422	380	3.7e-5	1.4e-1	1.71	287	807	4.7	0.08	31.0	246
TB05	328	309	1.5e-4	2.3e-1	1.28	206	634	4.9	0.06	35.1	255
TB06	419	402	0	1.5e-1	1.89	271	860	4.7	0.08	34.6	222
TB07	486	457	3.1e-5	7.3e-2	2.08	377	978	4.7	0.07	32.1	233
TB08	667	603	0	9.1e-2	2.94	537	1326	4.5	0.07	32.8	227
TB09	550	494	0	7.8e-2	2.15	330	1046	4.7	0.07	33.2	256
TB10	352	332	0	1.1e-1	1.12	203	678	4.9	0.07	31.7	315
TB11	876	773	0	8.3e-2	3.96	713	1729	4.5	0.08	34.6	221
TB12	854	793	0	4.9e-2	4.63	771	1774	4.5	0.08	39.2	184
TB13	935	864	5.4e-5	7.7e-2	4.77	763	1970	4.4	0.08	35.9	196
TB14	366	320	6.5e-5	1.8e-1	0.99	243	656	4.9	0.08	28.4	371
TB15	586	517	1.2e-3	1.1e-1	2.58	638	1126	4.6	0.07	36.2	227
TB16	1041	962	0	7.0e-2	6.00	2610	2277	4.2	0.10	33.4	173
TB17	766	724	1.8e-5	5.1e-2	3.48	576	1555	4.7	0.08	29.7	220
means	656	600	1.8e-04	1.0e-01	3.15	653	1342	4.6	0.08	32.9	227
±	214	196	5.9e-05	3.4e-02	1.03	213	438	1.5	0.03	10.7	74

Table A.2: Simulation results for the 4 DOF 3D milling algorithm. Means are reported with 95% confidence intervals.

Specimen	**Predicted path length (mm)**	**Modified final targets / initial**	**Unmodified final targets / initial**	**Modified initial target (cm^3)**	**Simulation time (s)**	**Steps in contact / path length**	**Average voxels milled per step in contact**	**Ratio predicted path length (mm / cm^3)**
CH01	518	3.2e-2	1.9e-1	2.19	4689	0.7	8.26	236.9
CH02	600	2.1e-2	1.0e-1	2.67	5315	0.7	8.13	225.1
CH03	447	7.8e-3	9.4e-2	1.45	3833	0.6	8.29	307.9
CH04	1010	6.5e-3	5.9e-2	4.65	9867	0.6	7.48	217.0
CH05	993	1.6e-2	7.3e-2	4.42	9563	0.7	7.93	224.7
CH06	624	9.3e-3	1.3e-1	2.33	6860	0.6	7.97	268.2
CH07	1016	2.0e-2	8.8e-2	4.59	10423	0.7	7.59	221.4
CH08	927	5.9e-4	5.4e-2	3.88	11670	0.6	7.37	239.1
CH09	705	2.7e-3	4.7e-2	3.15	6090	0.7	7.47	223.6
CH10	236	3.1e-4	7.0e-2	0.83	1960	0.6	8.01	283.0
CH11	311	1.3e-1	2.4e-1	1.11	2906	0.6	8.92	280.5
CH12	1205	2.3e-3	3.3e-2	5.88	11961	0.7	7.37	204.8
CH13	1113	3.6e-3	3.9e-2	5.50	15280	0.7	7.21	202.4
CH14	1024	1.6e-3	2.8e-2	4.78	9690	0.7	7.13	214.2
CH15	889	2.6e-3	3.7e-2	3.54	10146	0.6	7.32	250.9
CH16	949	1.5e-3	3.6e-2	4.23	10995	0.7	7.33	224.5
CH17	840	1.7e-2	5.0e-2	3.84	6494	0.6	7.37	218.8
CH18	581	2.4e-3	5.1e-2	2.06	5732	0.6	7.73	282.2

CH19	1000	2.5e-3	4.2e-2	4.38	8724	0.7	7.44	228.3
TB01	797	4.1e-2	1.3e-1	3.96	15568	0.7	8.23	201.5
TB02	341	1.3e-1	3.0e-1	1.42	3621	0.7	8.96	240.1
TB03	494	5.8e-2	1.4e-1	2.41	4947	0.7	8.68	205.1
TB04	401	3.5e-2	1.7e-1	1.71	13477	0.7	8.49	233.9
TB05	363	3.4e-2	1.3e-1	1.28	2986	0.7	8.61	282.6
TB06	447	2.1e-2	1.4e-1	1.89	20863	0.7	8.53	236.4
TB07	522	5.9e-3	7.2e-2	2.08	4091	0.7	8.10	250.9
TB08	749	4.6e-3	9.7e-2	2.93	7423	0.7	7.85	256.0
TB09	469	2.2e-2	1.1e-1	1.91	3726	0.7	8.07	245.7
TB10	326	2.2e-2	1.2e-1	1.07	3120	0.6	8.66	306.1
TB11	964	1.3e-2	9.0e-2	3.96	9995	0.7	7.96	243.3
TB12	996	1.0e-3	4.7e-2	4.59	10496	0.7	7.56	217.0
TB13	1033	1.8e-3	6.4e-2	4.39	9354	0.7	7.68	235.3
TB14	257	1.7e-1	3.7e-1	0.98	2414	0.6	9.68	261.6
TB15	648	1.5e-2	1.2e-1	2.58	6730	0.7	7.98	250.6
TB16	1190	6.5e-3	4.6e-2	5.59	12917	0.6	7.56	213.1
TB17	851	1.3e-3	5.2e-2	3.48	8560	0.7	7.71	244.7
means	718	2.4e-02	1.0e-01	3.10	8125	0.7	7.96	241.0
±	234	7.9e-03	3.3e-02	1.01	2654	0.2	2.60	78.7

Table A.3: Simulation results for the baseline 5 DOF zigzag algorithm. Means are reported with 95% confidence intervals.

Specimen	Predicted path length (mm)	Modified final targets / initial	Unmodified final targets / initial	Modified initial target (cm^3)	Simulation time (s)	Steps in contact / path length	Average voxels milled per step in contact	Ratio predicted path length (mm / cm^3)
CH01	515	2.6e-2	1.9e-1	2.19	41572	0.1	28.46	235.4
CH02	585	1.2e-2	8.7e-2	2.67	49642	0.1	36.16	219.5
CH03	426	3.8e-3	8.7e-2	1.45	32218	0.1	27.60	293.4
CH04	832	1.3e-3	5.4e-2	4.65	97969	0.1	30.35	178.8
CH05	902	1.1e-2	6.9e-2	4.42	91587	0.1	33.50	204.0
CH06	539	7.7e-3	1.3e-1	2.33	47022	0.1	33.70	231.7
CH07	875	1.8e-2	8.4e-2	4.59	82716	0.1	39.20	190.8
CH08	789	0	5.1e-2	3.88	92719	0.1	29.62	203.6
CH09	660	0	4.3e-2	3.15	54543	0.1	33.24	209.1
CH10	204	0	6.9e-2	0.83	12507	0.1	23.06	244.2
CH11	350	1.1e-1	2.2e-1	1.11	25519	0.1	24.34	316.1
CH12	976	0	2.8e-2	5.88	104354	0.1	37.97	165.9
CH13	928	0	3.6e-2	5.50	103136	0.1	35.99	168.8
CH14	894	0	2.6e-2	4.78	91496	0.1	31.47	187.0
CH15	711	0	3.4e-2	3.54	72174	0.1	29.69	200.6
CH16	795	0	3.2e-2	4.23	87075	0.1	32.63	188.1
CH17	691	0	3.2e-2	3.84	62710	0.1	33.94	180.0
CH18	469	0	5.0e-2	2.06	40618	0.1	31.18	227.8

CH19	861	4.4e-5	3.9e-2	4.38	85391	0.1	32.32	196.5
TB01	782	3.0e-2	1.1e-1	3.96	85682	0.1	34.31	197.7
TB02	364	1.2e-1	2.8e-1	1.42	30468	0.1	31.26	256.5
TB03	521	5.2e-2	1.2e-1	2.41	36843	0.1	30.77	216.3
TB04	414	3.1e-2	1.7e-1	1.71	31792	0.1	28.70	241.9
TB05	379	3.0e-2	1.2e-1	1.28	26371	0.1	28.29	295.2
TB06	439	1.4e-2	1.3e-1	1.89	34665	0.1	36.16	232.1
TB07	488	5.2e-4	6.3e-2	2.08	30871	0.1	33.99	234.2
TB08	644	1.0e-3	9.0e-2	2.93	59941	0.1	32.19	220.3
TB09	476	0	8.0e-2	1.91	35420	0.1	33.84	249.3
TB10	327	2.8e-3	1.1e-1	1.07	21660	0.1	33.48	306.6
TB11	890	9.9e-3	7.7e-2	3.96	94036	0.1	33.04	224.7
TB12	851	0	4.5e-2	4.59	84410	0.1	39.69	185.4
TB13	845	0	6.0e-2	4.39	79904	0.1	37.42	192.4
TB14	296	1.3e-1	3.4e-1	0.98	19799	0.1	30.76	301.1
TB15	600	1.1e-2	1.2e-1	2.58	55806	0.1	36.18	232.3
TB16	977	0	4.2e-2	5.59	100483	0.1	31.40	174.8
TB17	782	7.4e-5	4.8e-2	3.48	78208	0.1	28.36	224.8
means	641	1.7e-02	9.3e-02	3.10	60592	0.1	32.34	223.0
±	209	5.7e-03	3.1e-02	1.01	19794	0.0	10.56	72.8

Table A.4: Simulation results for the 5 DOF 3D algorithm. Means are reported with 95% confidence intervals.

Appendix B

Partial Nephrectomy Safety Margins: Full Results

	Tumor 1		**Tumor 2**		**Tumor 3**		**Tumor 4**	
	SV	CT	SV	CT	SV	CT	SV	CT
	Initialization Error STD 5 mm and 5°							
Margin Volume (cm^3)	2.2e+4	6.3e+4	8.0e+3	1.3e+4	7.5e+4	1.3e+5	2.0e+4	3.9e+4
Final Damage Probability	1.1e-2	9.0e-3	9.4e-3	8.0e-3	1.0e-2	1.0e-2	1.0e-2	1.0e-2
Mean M-Dist to Surrounding Voxels	1.87	4.37	2.12	2.69	2.05	3.07	2.39	3.36
STDev of M-Dists to Surrounding Voxels	0.40	1.89	0.54	0.66	0.33	1.28	0.95	1.33
time (min)	35.4	61.3	17.5	4.8	100	110	28.6	53.9

Table B.1: Numeric results for the spatially varying (SV) margins and the constant thickness (CT) margins created surrounding a tumor placed in a kidney. M-Dist refers to the Mahalanobis distance. This data was gathered by setting the randomly generated position and rotational error standard deviations described in Section 4.3.1 to 5 mm and 5°.

	Tumor 1		**Tumor 2**		**Tumor 3**		**Tumor 4**	
	SV	CT	SV	CT	SV	CT	SV	CT
	Initialization Error STD 5 mm and 3.33°							
Margin Volume (cm^3)	2.1e+4	6.1e+4	8.1e+3	1.5e+4	7.6e+4	1.2e+5	2.0e+4	4.5e+4
Final Damage Probability	6.5e-3	1.0e-2	9.6e-3	6.0e-3	9.8e-3	1.0e-2	9.0e-3	1.1e-2
Mean M-Dist to Surrounding Voxels	1.99	4.30	2.19	2.81	2.06	3.04	2.41	3.58
STDev of M-Dists to Surrounding Voxels	0.49	1.72	0.74	0.82	0.36	1.29	0.92	1.42
time (min)	23.4	32.7	12.7	6.3	116.7	77.7	33.8	91.8

Table B.2: Numeric results for the spatially varying (SV) margins and the constant thickness (CT) margins created surrounding a tumor placed in a kidney. M-Dist refers to the Mahalanobis distance. This data was gathered by setting the randomly generated position and rotational error standard deviations described in Section 4.3.1 to 5 mm and 3.33°.

	Tumor 1		**Tumor 2**		**Tumor 3**		**Tumor 4**	
	SV	CT	SV	CT	SV	CT	SV	CT
	Initialization Error STD 5 mm and 2.5°							
Margin Volume (cm^3)	2.1e+4	7.1e+4	8.3e+3	1.7e+4	7.2e+4	1.2e+5	1.9e+4	4.1e+4
Final Damage Probability	1.0e-2	5.5e-3	9.0e-3	1.0e-2	1.0e-2	1.0e-2	1.0e-2	1.0e-2
Mean M-Dist to Surrounding Voxels	1.86	4.68	1.87	3.30	2.02	3.07	2.39	3.44
STDev of M-Dists to Surrounding Voxels	0.39	1.92	0.40	0.98	0.30	1.31	0.95	1.38
time (min)	25.6	99.5	12.8	128.9	88.0	71.3	29.5	40.2

Table B.3: Numeric results for the spatially varying (SV) margins and the constant thickness (CT) margins created surrounding a tumor placed in a kidney. M-Dist refers to the Mahalanobis distance. This data was gathered by setting the randomly generated position and rotational error standard deviations described in Section 4.3.1 to 5 mm and 2.5°.

	Tumor 1		**Tumor 2**		**Tumor 3**		**Tumor 4**	
	SV	CT	SV	CT	SV	CT	SV	CT
	Initialization Error STD 3.33 mm and 5°							
Margin Volume (cm^3)	1.8e+4	3.8e+4	7.1e+3	1.4e+4	5.3e+4	8.5e+4	1.1e+4	2.0e+4
Final Damage Probability	7.5e-3	8.0e-3	8.0e-3	1.0e-3	1.0e-2	1.0e-2	1.0e-2	1.0e-2
Mean M-Dist to Surrounding Voxels	1.92	3.56	2.13	3.17	2.05	2.99	2.22	3.29
STDev of M-Dists to Surrounding Voxels	0.26	1.39	0.40	1.08	0.19	1.24	0.57	1.35
time (min)	23.7	19.5	14.8	8.9	72.3	81.5	22.6	49

Table B.4: Numeric results for the spatially varying (SV) margins and the constant thickness (CT) margins created surrounding a tumor placed in a kidney. M-Dist refers to the Mahalanobis distance. This data was gathered by setting the randomly generated position and rotational error standard deviations described in Section 4.3.1 to 3.33 mm and 5°.

	Tumor 1		**Tumor 2**		**Tumor 3**		**Tumor 4**	
	SV	CT	SV	CT	SV	CT	SV	CT
	Initialization Error STD 3.33 mm and 3.33°							
Margin Volume (cm^3)	1.6e+4	4.1e+4	6.9e+3	1.1e+4	5.4e+4	7.7e+4	1.1e+4	2.1e+4
Final Damage Probability	1.0e-2	8.5e-3	1.1e-2	5.0e-3	1.0e-2	1.1e-2	1.1e-2	1.0e-2
Mean M-Dist to Surrounding Voxels	1.89	3.82	2.05	2.80	2.08	2.94	2.17	3.29
STDev of M-Dists to Surrounding Voxels	0.30	1.47	0.53	0.93	0.20	1.31	0.53	1.32
time (min)	22.9	30.6	11.7	13.2	67.4	56.4	21.2	52.4

Table B.5: Numeric results for the spatially varying (SV) margins and the constant thickness (CT) margins created surrounding a tumor placed in a kidney. M-Dist refers to the Mahalanobis distance. This data was gathered by setting the randomly generated position and rotational error standard deviations described in Section 4.3.1 to 3.33 mm and 3.33°.

	Tumor 1		**Tumor 2**		**Tumor 3**		**Tumor 4**	
	SV	CT	SV	CT	SV	CT	SV	CT
	Initialization Error STD 3.33 mm and 2.5°							
Margin Volume (cm^3)	1.8e+4	4.6e+4	7.2e+3	1.3e+4	5.1e+4	7.6e+4	1.0e+4	1.9e+4
Final Damage Probability	8.8e-3	1.0e-2	7.0e-3	5.0e-4	1.0e-2	8.5e-3	1.0e-2	8.0e-3
Mean M-Dist to Surrounding Voxels	1.96	4.07	2.11	2.75	2.08	2.89	2.17	3.22
STDev of M-Dists to Surrounding Voxels	0.31	1.52	0.62	1.01	0.22	1.18	0.53	1.31
time (min)	25.3	107.5	12.8	4.3	54.9	66.8	14.2	34.1

Table B.6: Numeric results for the spatially varying (SV) margins and the constant thickness (CT) margins created surrounding a tumor placed in a kidney. M-Dist refers to the Mahalanobis distance. This data was gathered by setting the randomly generated position and rotational error standard deviations described in Section 4.3.1 to 3.33 mm and 2.5°.

	Tumor 1		**Tumor 2**		**Tumor 3**		**Tumor 4**	
	SV	CT	SV	CT	SV	CT	SV	CT
	Initialization Error STD 2.5 mm and 5°							
Margin Volume (cm^3)	1.5e+4	3.2e+4	6.9e+3	1.2e+4	4.1e+4	5.7e+4	8.5e+3	1.2e+4
Final Damage Probability	8.5e-3	8.5e-3	9.6e-3	8.5e-3	1.0e-2	9.5e-3	9.6e-3	1.0e-2
Mean M-Dist to Surrounding Voxels	1.95	3.44	2.29	2.87	2.04	2.74	2.24	2.98
STDev of M-Dists to Surrounding Voxels	0.23	1.21	0.57	0.93	0.16	1.01	0.40	1.17
time (min)	19.4	50.0	12.9	6.1	42.6	59.3	15.3	18.4

Table B.7: Numeric results for the spatially varying (SV) margins and the constant thickness (CT) margins created surrounding a tumor placed in a kidney. M-Dist refers to the Mahalanobis distance. This data was gathered by setting the randomly generated position and rotational error standard deviations described in Section 4.3.1 to 2.5 mm and 5°.

	Tumor 1		**Tumor 2**		**Tumor 3**		**Tumor 4**	
	SV	CT	SV	CT	SV	CT	SV	CT
	Initialization Error STD 2.5 mm and 3.33°							
Margin Volume (cm^3)	1.5e+4	3.0e+4	6.5e+3	1.2e+4	4.0e+4	5.8e+4	7.8e+3	1.3e+4
Final Damage Probability	1.0e-2	1.0e-2	1.1e-2	7.5e-3	1.3e-2	1.0e-2	9.8e-3	1.5e-2
Mean M-Dist to Surrounding Voxels	1.98	3.54	2.12	3.08	2.02	2.77	2.14	3.14
STDev of M-Dists to Surrounding Voxels	0.22	1.29	0.46	1.04	0.15	1.03	0.41	1.30
time (min)	22.8	11.3	12.8	16.0	54.7	29.5	18.9	17.8

Table B.8: Numeric results for the spatially varying (SV) margins and the constant thickness (CT) margins created surrounding a tumor placed in a kidney. M-Dist refers to the Mahalanobis distance. This data was gathered by setting the randomly generated position and rotational error standard deviations described in Section 4.3.1 to 2.5 mm and 3.33°.

	Tumor 1		**Tumor 2**		**Tumor 3**		**Tumor 4**	
	SV	CT	SV	CT	SV	CT	SV	CT
	Initialization Error STD 2.5 mm and 2.5°							
Margin Volume (cm^3)	1.4e+4	2.9e+4	6.9e+3	1.3e+4	3.5e+4	5.9e+4	7.5e+3	1.1e+4
Final Damage Probability	6.5e-3	1.1e-2	8.0e-3	1.1e-2	1.0e-2	7.0e-3	9.6e-3	1.0e-2
Mean M-Dist to Surrounding Voxels	1.93	3.42	2.01	2.98	1.97	2.97	2.26	2.91
STDev of M-Dists to Surrounding Voxels	0.22	1.26	0.47	0.93	0.17	1.23	0.46	1.12
time (min)	19.8	21.4	14.0	11.1	36.5	31.3	18.2	21.0

Table B.9: Numeric results for the spatially varying (SV) margins and the constant thickness (CT) margins created surrounding a tumor placed in a kidney. M-Dist refers to the Mahalanobis distance. This data was gathered by setting the randomly generated position and rotational error standard deviations described in Section 4.3.1 to 2.5 mm and 2.5°.