Imperial College London

SHARED CONTROL FOR NATURAL MOTION AND SAFETY IN HANDS-ON ROBOTIC SURGERY

Joshua George Petersen

Department of Mechanical Engineering

Imperial College London

June 2015

This thesis is submitted to Imperial College London in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

Declaration of Originality

The research presented in this thesis is my own work except where appropriately referenced.

Josh Petersen

March 23, 2016

Abstract

Hands-on robotic surgery is where the surgeon controls the tool's motion by applying forces and torques to the robot holding the tool, allowing the robot-environment interaction to be felt though the tool itself. To further improve results, shared control strategies are used to combine the strengths of the surgeon with those of the robot. One such strategy is active constraints, which prevent motion into regions deemed unsafe or unnecessary. While research in active constraints on rigid anatomy has been well-established, limited work on dynamic active constraints (DACs) for deformable soft tissue has been performed, particularly on strategies which handle multiple sensing modalities. In addition, attaching the tool to the robot imposes the end effector dynamics onto the surgeon, reducing dexterity and increasing fatigue. Current control policies on these systems only compensate for gravity, ignoring other dynamic effects. This thesis presents several research contributions to shared control in hands-on robotic surgery, which create a more natural motion for the surgeon and expand the usage of DACs to point clouds. A novel null-space based optimization technique has been developed which minimizes the end effector friction, mass, and inertia of redundant robots, creating a more natural motion, one which is closer to the feeling of the tool unattached to the robot. By operating in the null-space, the surgeon is left in full control of the procedure. A novel DACs approach has also been developed, which operates on point clouds. This allows its application to various sensing technologies, such as 3D cameras or CT scans and, therefore, various surgeries. Experimental validation in point-to-point motion trials and a virtual reality ultrasound scenario demonstrate a reduction in work when maneuvering the tool and improvements in accuracy and speed when performing virtual ultrasound scans. Overall, the results suggest that these techniques could increase the ease of use for the surgeon and improve patient safety.

Acknowledgements

I would first like to thank my supervisor, Dr. Ferdinando Rodriguez y Baena, for his years of advice and direction in my research. His guidance helped me conceptualize my ideas and transform them into publishable work.

I would also like to thank Professor Brian Davies, who first developed hands-on robotic surgery and active constraints, for providing valuable feedback in the early stages of my PhD.

Thanks to the members and associates of the Mechatronics in Medicine Laboratory whom I have spent many hours in the lab and a few less outside the lab. In alphabetical order: Elisa Baretta, William Blyth, Stuart Bowyer, Michele Brugnoli, Chris Burrows, Arnau Casanovas, Hadi El Daou, Antonio Forte, Enrico Franco, Stefano Galvan, Josh Giles, Trevor Hawke, Pia Kaufmann, Alex Leibinger, Fangde Liu, Elnaz Nobari, Matthew Oldfield, Riccardo Secoli, Ryo Takeda, Vani Virdyawan, and Thomas Watts.

I would like to acknowledge the Active Constraints Technologies for Ill-defined or Volatile Environments project (FP7-ICT-2009-6-270460) for funding my research.

Lastly, I would like to thank my partner, Miranda Simon, for her patience and assistance in keeping me sane throughout the PhD process.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation

Approaches to surgical robotics can be subdivided into three main areas: autonomous, where the robot moves independently of the surgeon for at least part of the procedure; teleoperative, where the surgeon controls the robot through a remote console; and hands-on robotic surgery, where the surgical tool attached to the end effector of the robot is maneuvered directly by forces and torques applied by the surgeon to the tool.

Autonomous surgical systems include those in which the robot is a passive tool holder, and those in which it autonomously interacts with the patient. The first robotic assisted surgery, performed by [Kwoh et al., 1988], used an autonomous industrial robot to position a surgical needle guide for brain tumor biopsy, replacing the manually-adjustable stereotactic frame. To minimise potential brain damage, a probe needs to be inserted on a straight-line path, avoiding major blood vessels and other vital parts of the brain. Historically, the surgeon localised the area to be operated on, attached the stereotactic frame to the skull of the patient and manually adjusted the angle of the tool, which was calculated from a computerised tomographic (CT) image of the patient. Like the stereotactic frame, this autonomous surgical robot was a positioning guide -it played no role in tissue removal-, but it offered quicker targeting and improved accuracy. A commercially available autonomous clinical neurosurgical robot of this type is the Neuromate (Renishaw Plc.) [Benabid et al., 1987] [Li et al., 2002].

Autonomous surgical robots which interact with the patient, such as the Probot [Davies et al., 1991] [Harris et al., 1997], which targeted prostatectomies, and the ROBODOC [Paul et al., 1992] [Pransky, 1997] and the CASPAR [Siebert et al., 2002], which targeted lower limb arthroplasty, also aim to improve accuracy in surgical resections. However, autonomous systems which remove tissue from the patient leave many open questions concerning safety, liability, and ethics [Davies et al., 2004]. Removing the surgeon from the procedure can also lead to anxiety for clinical regulators, surgeons, and patients [Davies, 1996].

For these reasons, the ROBODOC is the only autonomous robotic system which removes tissue to have seen widespread commercialisation. Due to their simplicity, passive tool holder robots have been more widely commercialised in systems such as the Pathfinder (Prosurgics Inc.) and Spine Assist (Mazor Robotics Rtd.), however, they are limited to positioning tasks and therefore, cannot be applied to more general surgeries.

Teleoperative robotic surgery, whereby the surgeon controls the robot remotely via a master console, allows for greater surgeon involvement than autonomous systems. Commercial systems, such as the da Vinci Surgical System (Intuitive Surgical Inc.) and the Sensei X (Hanson Medical Inc.), have made robotic laparoscopic and endovascular surgery possible. This minimally invasive technique, where surgery is performed through a small incision, has been demonstrated to reduce pain and loss of function following surgery as compared with conventional procedures [Alexander, 1997].

However, by controlling the robot remotely, the forces and torques of the tool's interaction with the surgical environment are lost. The lack of this haptic feedback has been demonstrated to increase average force application to the tissue by at least 50% and peak force by a factor of 2 [Wagner et al., 2002]. In addition, a lack of force feedback has been shown to result in an increase in gallbladder perforation in laparoscopic cholecystectomy [Joice et al., 1998] and an increase in suture ruptures in anastomosis in porcine intestine [Ruurda et al., 2004].

To restore this haptic feedback in telemanipulation, research systems such as the Neuroarm [Sutherland et al., 2013] use a force/torque sensor on the slave device and replicate the sensed wrench at the master console. However, challenges still remain. The force/torque sensing solution must be able to undergo the harsh sterilisation procedures performed using heat and/or

chemicals, adding to the already steep costs of force/torques sensors [Okamura et al., 2010]. Additionally, research is ongoing in ensuring stability in these systems [Okamura, 2009] and, to date, no commercialised teleoperative system provides haptic feedback.

Hands-on robotic surgery, a term first coined in [Davies et al., 1999], is an alternative approach in which the surgical tool attached to the end effector of the robot is maneuvered directly by forces and torques applied by the surgeon to the tool itself. This type of strategy is advantageous, as the forces and torques of the interaction between the robot and the surgical environment can be felt directly through the instrument, without the need for an additional force/torque sensor at the end effector. This natural feedback can be used by the surgeon to more precisely apply forces and torques to the hard or soft target tissue. This allows hands-on robotic surgery to avoid the above mentioned issues with stability and sterilisability in providing haptic feedback in teleoperative setups.

In addition, the surgeon is more involved in the procedure as compared to autonomous robotic solutions, which can increase its acceptability among surgeons and regulatory bodies [Davies et al., 2004]. The robot only applies control to ensure safety or accuracy through active constraints, which only limit motion when the tools are brought into regions which have been deemed unsafe or unnecessary for the surgery.

Current clinical applications of active constraints in systems such as the Rio (Mako Surgical Corp.) are limited to rigid bone surgeries, as constraint representations are fixed and intra-operative tracking and registration is simple in such scenarios. To expand the benefits of active constraints in accuracy and safety to soft tissue surgery, research into dynamic active constraints, which can protect deforming anatomy, has recently begun.

At present, dynamic active constraints strategies are restricted to particular representations of the anatomy's geometry or place limits on the motion of the deforming region, preventing their applicability to general surgical procedures. Deformation of soft tissue due to tool interactions and patient breathing and heart beat can take arbitrary shape and therefore, achieving a constraint representation that is able to handle unpredictable deformations is one of the main challenges to making hands-on robotic surgery more widely used.

In addition, the deformations of the soft tissue anatomy must be tracked intra-operatively

to ensure precise application of the constraints to aid the surgeon. Research is ongoing into various techniques and sensing modalities for real-time nonrigid registration and tracking of deformable tissue using ultrasound [Weon et al., 2015] and stereo camera imaging [Faria et al., 2014]. Additionally, the ability to use a particular tracking strategy is dependent on the type of surgery to be performed. As a result, the current limitation on the constraint representation in dynamic active constraints strategies further prevents their general applicability to surgery, as they are unable to handle various tracking solutions.

Another problem in hands-on robotic surgery is that, by attaching the surgical instrument to the tool point of the robot, the dynamics of the robot at the end effector are imposed on the surgeon when maneuvering the tool. These dynamics can include the mass and inertia of the robot links, joint friction, gear backlash, and other disturbances, such as interactions with the sterile draping. The imposition of additional dynamics can reduce the surgeon's dexterity and increase fatigue over long procedures which range from 40 to 120 minutes [Lang et al., 2011] for current hands-on surgeries.

1.2 Thesis Aim

The difficult interaction between the robot and the surgeon in hands-on robotic surgery and its confinement to rigid bone procedures, are some of the most important issues limiting the acceptance of this strategy and its expansion to more types of surgery. Hence, the aim of this thesis is to develop robot control strategies that can allow surgeons to interact with the robot in a way that more closely resembles his or her interaction with the tool unattached to the robot, and active constraints that can handle the unpredictable deformations found in soft tissue, and can be applied using various tracking modalities and therefore, to a variety of procedures.

More specifically, this thesis aims to create a more natural motion for the surgeon by reducing the dynamics of the robot at the end effector using the redundant degrees-of-freedom of the many robot kinematic configurations in the literature. Redundant robots are those in which there are more joints than are necessary to achieve an arbitrary pose in space. Using a null-space control strategy, these redundant degrees-of-freedom can be varied without affecting the

tool pose at the end effector. Additionally, the tool point dynamics are a function of the joint configuration of the robot and change based on the position of the redundancy. Therefore, by using a null-space based natural motion control strategy, the end effector dynamics can be reduced in order to decrease fatigue and restore dexterity for the surgeon in hands-on robotic surgery.

Secondly, this thesis aims to create a form of dynamic active constraints for hands-on robotic surgery which does not require a mesh representation and instead can generate constraints directly from streaming point clouds. Such a strategy can be applied to general tracking modalities as many of these techniques can be reduced to point clouds and therefore, the methodology could be applied to many type of soft tissue surgeries. Additionally, by generating a constraint surface from point clouds, the constraint representation can take on arbitrary deformations and motions, where the main bottleneck becomes the refresh rate of point cloud information.

1.3 Thesis Contributions

Research performed towards these aims has resulted in three main contributions;

A novel redundancy optimisation controller has been developed, which lessens the impact of the dynamics of the robot at the end effector on the surgeon by reducing the mass, inertia, friction force, and friction torque at the tool point. By lessening the effect of these dynamics, the feeling of the surgeon's interaction with the robot is made more natural, closer to the surgical instrument unattached to the robot, thereby, restoring dexterity and reducing fatigue. The optimisation is performed in the null-space of the end effector as to ensure the controller does not impact the surgeon's desired tool commands, leaving the surgeon in full control of the surgery at hand. This research is presented in Chapters 3 and 4.

A new dynamic active constraints approach for hands-on robotic surgery has been proposed which creates an implicit constraint surface from a point cloud representation of the region to be protected using the metaball approach from computer graphics. By generating the dynamic active constraints from point clouds, this methodology can be applied to general soft tissue tracking modalities, thereby expanding the applicability of the technique to a variety of surgical

procedures. In addition, the dynamics of the robot have been incorporated into the enforcement of the constraint. These dynamics are necessary in hands-on robotic surgery to ensure an appropriate response to the surgeon's force and torques. This research is presented in Chapter 5.

A virtual reality environment for hands-on robotic surgery has been created, which allows for the testing of shared control strategies for dynamic soft tissue in a more immersive setting, while still preserving the true interaction between the surgeon and the robot. Bringing surgical robotic control research to a clinical setting is a long and arduous process due to necessary safety regulations and phantoms which can exhibit complex soft tissue deformations are difficult to fabricate. Conversely, a virtual setup can provide a safe and versatile way in which to test novel strategies. Registration of the robot with the virtual surgical theater, such that the robot's position in the virtual environment matches that in the real world, allows the complex dynamic interactions between the physical hardware, the control strategy, and the surgeon to be preserved. In addition, the forces and torques can be applied to simulate the interaction of the robot with the virtual soft tissue. This research is presented in Chapter 6.

As a result of this research, currently, two peer-reviewed papers have been published, a workshop presentation has given, and a journal paper is currently under second review:

J. G. Petersen and F. Rodriguez y Baena, "Mass and friction optimization for natural motion in hands-on robotic surgery", IEEE Transactions on Robotics, Under Review.

J. G. Petersen and F. Rodriguez y Baena, "Mass and inertia optimization for natural motion in hands-on robotic surgery", in Proceedings of the IEEE International Conference on Intelligent Robots and Systems, pp. 4284-4289, Sept. 2014.

J. G. Petersen and F. Rodriguez y Baena, "A dynamic active constraints approach for hands-on robotic surgery", in Proceedings of the IEEE International Conference on Intelligent Robots and Systems, pp. 1966-1971, Nov. 2013.

J. G. Petersen and F. Rodriguez y Baena, "Shared control for hands-on robotic surgery", IEEE International Conference on Intelligent Robots and Systems Workshop on Cognitive Surgical Robotics, Nov. 2013.

1.4 Chapter Summaries

In Chapter 2, a comprehensive literature review of the state-of-the art in shared control for general and surgical robotics is presented. Additionally, a categorical structure for classifying various cooperative control strategies is developed and applied throughout the survey.

In Chapter 3, a strategy to reduce the impact of the end effector dynamics on the surgeon to create a more natural motion is investigated. Optimisation criteria for reducing the mass and inertia at the tool point are developed and a control strategy to perform the optimisation in the null-space of the end effector to ensure the surgeon's desired commands are not affected, is shown. Extensive simulation work on a planar robot demonstrates the effectiveness of the proposed methodology in reducing the mass and inertia over the workspace of a planar robot. In addition, compatibility of mass and inertia optimisation joint configurations and the impact of modeling uncertainty on the optimisation results are investigated. Finally, simulated trials demonstrate a reduction in the time required to follow a trajectory as compared with classical redundancy resolution schemes.

In Chapter 4, the natural motion strategy of Chapter 3 is extended to optimise the projection of the joint friction onto the end effector. Simulations of a seven degree-of-freedom lightweight robot demonstrate the efficacy of the proposed mass, inertia, friction force, and friction torque optimisation techniques over the workspace of the robot. Additionally, combined optimisations of mass and friction force are considered. Through experimental user trials, the effectiveness of these strategies in producing a more natural motion for the user over standard surgical redundancy control strategies is shown.

In Chapter 5, a dynamic active constraints approach, which can directly generate a constraint surface from point cloud data and incorporates the kinematics and dynamics of the robot, is developed. Null-space control is used to allow unimpeded motion tangent to the constraint surface and for the simultaneous application of natural motion and dynamics active constraints in Chapter 6. The proposed methodology is demonstrated in simulation on constraint surfaces generated from two different points clouds, one in which the point cloud's motion is unknown and the other in which the point cloud's motion is known (i.e. can be predicted). Lastly, an

experiment using a dynamic brain phantom and a lightweight seven degree-of-freedom robot is used to demonstrate the feasibility of the proposed technique on physical hardware.

In Chapter 6, a virtual reality environment for testing hands-on robotic surgery control techniques is demonstrated, which utilises a physical robot registered to a virtual surgical theater to preserve the dynamic iteration between the robot, the surgeon, and the control strategy, while providing force and torque feedback based on virtual interactions with dynamic soft tissue. Within this setup, the dynamic active constraints and natural motion of Chapters 3, 4, and 5 are demonstrated working simultaneously in a chest ultrasound scenario. The results show efficacy of the dynamic active constraints in improving the user's accuracy and ensure safety for the patient when performing ultrasound scans.

In Chapter 7, conclusions are drawn about the thesis contributions with regards to the stated research aims. Additionally, I present and discuss the limitations of the work presented here, and areas for future research are identified.

Chapter 2

Shared/Cooperative Control in Robotics

This chapter presents a thorough review of the published literature on shared control in robotics. Portions of this chapter are edited versions of the work published in:

> J. G. Petersen and F. Rodriguez y Baena, "Mass and friction optimization for natural motion in hands-on robotic surgery", IEEE Transactions on Robotics, Under Review.
>
> J. G. Petersen and F. Rodriguez y Baena, "Mass and inertia optimization for natural motion in hands-on robotic surgery", in Proceedings of the IEEE International Conference on Intelligent Robots and Systems, pp. 4284-4289, Sept. 2014.
>
> J. G. Petersen and F. Rodriguez y Baena, "A dynamic active constraints approach for hands-on robotic surgery", in Proceedings of the IEEE International Conference on Intelligent Robots and Systems, pp. 1966-1971, Nov. 2013.

2.1 Introduction

Shared control strategies, also known as cooperative or collaborative control strategies, are those in which the user and robot simultaneously enact control policies to achieve a task. The aim of these type of methodologies is to perform the task better than the device or user can alone,

by combining the benefits of both the human and robot. In general, the human has a better overall understanding of the task to be completed and thus, is more trusted to perform the most critical tasks and make high level decisions about the workflow or any plan modifications. On the other hand, robots are more accurate and precise due to better spatial analysis and registration. Additionally, robots are much stronger than humans and do not tire over lengthy tasks. Cooperative control has been demonstrated to improve many aspects of various functions, both in medicine and in general robotics, including reducing the user's mental load and fatigue, as well as the time taken to complete a task, and increasing accuracy and precision.

This literature review aims to provide an overview of the current state-of-the-art research in shared control for robotics, by first presenting a categorical structure to classify various co-operative control strategies and then, examining shared control's usage in general robotics. Finally, present clinical and research applications of cooperative control in surgical systems are detailed.

2.2 Types of Shared Control

Shared control strategies can be divided into several areas, each providing a different method for mixing the control of the human and the robot to achieve better overall results in the desired task. However, these methodologies are not mutually exclusive and can be combined to draw on the individual benefits of each control type. Four categories for shared control have been devised; degree-of-freedom sharing, input modification, space sharing and active guidance.

2.2.1 Degree-of-Freedom Sharing

Degree-of-freedom sharing cooperative control strategies are those in which the control policies of the user and the robot function in orthogonal spaces, such that the control actions of the two do not directly affect one another. There are two main ways in which these strategies are implemented.

The first method uses the robot to constrain particular degrees-of-freedom of the end effector while the user retains control of the remaining motions. In this way, the robot can position

a majority of the degrees-of-freedom of the necessary tool pose and leave the remaining, often critical degrees-of-freedom to the user. Performing a task this way can increase accuracy and/or decrease task completion time as compared with performing the task by hand or using a manually adjusted fixture. Leaving the user in control of the speed and direction of the task has increased acceptance of these types of systems, as compared to autonomous robots, particularly in surgical systems [Davies et al., 2004]. For this form of degree-of-freedom sharing to function precisely, accurate calibration and registration of the robot to the target must be performed to ensure the robot is properly aligned and constraining the correct degrees of freedom to a sufficient level of accuracy for the desired task.

The second implementation for degree-of-freedom sharing is null-space control methodologies on redundant robots. Redundant robots mimic human capabilities in that they have more degrees-of-freedom than is necessary to achieve any pose at the end effector in three-dimensional space. Redundancy is most commonly used to avoid obstacles in the way of the body of the robot without affecting the main task at the tool point [Maciejewski and Klein, 1985] [Minami and Takahara, 2003] [Park et al., 2008].

In this form of degree-of-freedom sharing, the robot reconfigures the redundant degrees-of-freedom without affecting the surgeon's desired motions and forces at the end effector, again leaving the surgeon fully in control of the procedure. These methodologies aim to improve the user's performance in the tool point task by taking advantage of the fact that end effector performance in aspects such as force and velocity transmission and effective mass varies based on the position of the redundant degrees-of-freedom.

In redundant robots, the kinematic model in admittance controlled systems and the kinematic and dynamic models in impedance controlled systems must be as accurate as possible to ensure that the appropriate null-space is computed. An inaccurate null-space computation can result in torques intended for the null-space to affect the end effector, disrupting the user's command of the tool. Additionally, inaccurate modeling can reduce the intended improvements in the user's performance, as the characteristics at the end effector must be computed from the model.

2.2.2 Input Modification

Input modification is used to alter the user's commanded forces, torques, velocities, and/or positions to improve the performance of the user in the task at hand. The simplest form of input modification is used in impedance controlled devices to compensate for gravity's effect on the robot and tool. As these systems are typically back-drivable (a force or torque on a link of the robot can cause a rotation/translation of its driving motor), they require this control as a baseline to ensure the robot does not fall over during use. This strategy can be thought of as a constant additional force and torque input added to the user's desired commands.

More complicated force and torque input modification strategies focus on scaling the user's applied wrench. Scaling up these wrenches can improve the user's ability to manipulate heavy objects. However, careful consideration must be taken to minimise the scaling up of noise to ensure stability. Conversely, methodologies which scale down the user's wrench can increase safety by preventing strong forces and torques from being applied to delicate environments.

Similarly, motion scaling can be used in teleoperative scenarios to modify the user's command input such that larger movements by the user correspond to proportionally smaller movements of the robot. By doing so, small, delicate tasks can be performed at a scale more natural for the user, resulting in higher accuracy [Prasad et al., 2004].

In environments where there is some periodic and predictable component to the motion of the target, this periodicity can be added to the user's desired motion, allowing him or her to act on an effectively stationary target. In addition to potentially increasing the accuracy of the task, this shared control strategy can also reduce the user's mental load, as he or she no longer needs to add this component of motion manually. This technique can additionally enable tasks to be performed on fast, periodic targets which normally would be very difficult or impossible to work on unassisted [Trejos et al., 1999]. Accuracy in predicting the target motion is a key requirement of this type of technique, since poor prediction could result in the robot colliding and damaging the target.

In the frequency domain, input modification is used to filter the commanded position of the user. For instance, by removing the high frequency components of the user's motion, the impact

of hand tremor on the position of the tool point can be reduced. Additionally, in tasks which operate near the order of the amplitude of the noise and tremor in the signal, frequency based techniques can be used to identify intended user motions from involuntary ones, which can increase accuracy in tasks on the micro-scale. Frequency filtering techniques must take care not to remove frequencies which the user is controlling with and, hence, degrade controllability of the tool. In addition, in systems which utilise user intention identification to determine when and where to move the tool, careful application must be used, as a false positive could cause the system to move against the surgeon's wishes and harm the patient.

2.2.3 Space Sharing

Space sharing cooperative control strategies define free spaces where the user can openly move and constrained spaces where the controller directs the user back to the free space region. Space sharing is used to prevent the tool from entering hazardous areas or to assist in training the user to work in a particular area. Since an absolute position constraint is very difficult to execute in practice, the level to which the tool can penetrate into the constrained space varies upon application and the control system used. The boundary of the constrained region is often offset from the damaging region to decrease the possibility of the tool reaching the tissue.

To ensure the free and constrained spaces align with the desired setup, the robot must be properly registered to the task space at sufficient accuracy for the task at hand. Inaccurate registration can result in the user being able to freely move into spaces which are considered dangerous and to be constrained from work areas which are necessary to complete the desired task. Furthermore, in scenarios where the environment is dynamic, such as in soft tissue surgery, the spaces can change in time and must be precisely tracked and updated throughout the procedure.

2.2.4 Active Guidance

Active guidance shared control methodologies are those which apply forces and torques to the user in order to guide them along a corridor, trajectory, or towards a point. Of the shared control categories presented here, this technique is the most similar to autonomous robotics.

As such, it is more commonly used in training scenarios, due to the complicated safety, legal, and ethical concerns which can arise when using these strategies on patients. Within training environments, active guidance is used to lead a trainee in the proper movement and forces necessary for the specific application.

2.3 Shared Control in General Robotics

In this section, I will focus on shared control in general robotics, that is, strategies the primary application of which has been presented as non-clinical. By doing so, I aim to evaluate the suitability of non-medical control strategies for potential surgical use, thereby expanding the range of methods which can be applied to robotic surgery.

2.3.1 Active Constraints in General Robotics

Active constraints, also known as virtual fixtures, were first developed by [Rosenberg, 1992] as a perceptual overlay designed to "reduce (the) mental processing required to perform the task, reduce the work load of certain sensory modalities, and most of all allow precision and performance to exceed natural human abilities" (pp. 3). Analogous to mechanical fixtures, which physically limit tool motion, active constraints use haptic feedback to virtually impose constraints on the motion of the user. Virtual implementation allows for the constraints to be applied in a variety of complex situations where physical fixtures would normally be difficult to implement, including small scale tasks. In addition, virtual fixtures can be changed much more easily than mechanical fixtures, to adapt to new tasks. Recent research into dynamic active constraints also demonstrates fixtures which change shape to accommodate environments which deform and shift during the task [Gibo et al., 2009] [Rydén and Chizeck, 2013].

Active constraints can be subdivided into two types: forbidden region active constraints and guidance active constraints. Forbidden region active constraints are space sharing control strategies which restrict the tool's motion into particular regions of the workspace. Guidance constraints, which can be either space sharing or active guidance cooperative control methodologies, encourage the user to move along a specific surface or path, or towards a point. In space shar-

ing form, guidance constraints function similarly to forbidden region constraints, where the free space is a path or point, and forces are applied to return the user to the path or point [Abbott et al., 2007]. Active guidance methodologies apply forces to direct tool motion along a desired path or surface [Pezzementi et al., 2007] [Castillo-Cruces and Wahrburg, 2010].

However, a recent survey on virtual fixtures in both general and medical applications found active guidance was rarely used [Bowyer et al., 2014]. Particularly in surgery, guidance constraints have open issues about liability and safety. Therefore, this section will focus on some of the recent developments in forbidden region active constraints research in general robotics. A more complete review of active constraints/virtual fixtures can be found in [Bowyer et al., 2014].

Many virtual fixtures implementations use virtual springs to enforce the motion constraint [Ho et al., 1995] [Turro et al., 2001] [Gibo et al., 2009]. These springs apply a force which is proportional to the distance between the tool point and the constraint surface in order to provide a simple yet effective strategy for directing the user back to the free space region. However, these constraints store elastic potential energy as the tool enters the constraint region and releasing the device can cause unintentional and potentially dangerous motions from the release of this energy. In addition, these types of constraints require a penetration into the constrained region before any opposing action is taken.

The work of [Kikuuwe et al., 2008] aims to overcome these issues by creating constraints which use a model of plasticity. In plasticity, the controller provides a stiff collision at the constraint boundary to prevent motion into the constraints up to a certain force level. Past this force threshold, the tool can enter the constraint, however, the controller in the constraint only dissipates energy. While direct application of this type of strategy is discontinuous and, hence, can cause instabilities, a discretised proxy-based method for plasticity was developed in [Kikuuwe et al., 2006] and demonstrated on impedance and admittance controlled devices in [Kikuuwe et al., 2008]. However, a comparison of this plasticity based method was not made with other virtual fixture techniques.

Similar work on reducing constraint penetration using passive means was performed by [Hennekens et al., 2008]. Using a previously developed method [Constantinescu et al., 2005], con-

straint impulse forces based on the current velocity, coupled with a PD controller to reduce drift in the constraint, were applied at the boundary and inside the constraint. Experimental results demonstrated a decrease in the penetration depth compared to a spring-damper constraint and subjects reported sensing the contact as a plastic collision.

The recent widespread availability of low cost three-dimensional (3D) cameras and range finders has enabled large scale point cloud representations of environments to be quickly acquired. Using these streaming point clouds, research into forbidden region virtual fixtures has been performed for haptic rendering in teleoperative scenarios.

Implicit surfaces were used by [Leeper et al., 2012] to generate a smooth representation of the point cloud data for haptic rendering. The surfaces were generated through two weighting methods, metaballs [Blinn, 1982] and point-sets [Adamson and Alexa, 2003], and their surface normals were computed via the gradient. Constraints were imposed based on the distance to the zero set of the implicit surface and in the direction of the gradient.

An alternative proxy-based method for haptic rendering of point clouds was presented in [Rydén et al., 2011]. When the user commanded the haptic device to move through a point cloud, the haptic proxy would not penetrate the point cloud and a force would be applied to the master in the direction of the proxy and proportional to the distance between the two. Directly using the point cloud points allowed for more accurate representations of rough surfaces and corners, but smooth surfaces had increased roughness due to noise in the data. Further developments to this method included the ability to constrain to an offset from the point cloud [Rydén and Chizeck, 2012] and an extension to six degrees-of-freedom constraints [Rydén and Chizeck, 2013].

To date, active constraints in general robotics that can handle dynamic environments have been limited to methodologies for teleoperation and haptic rendering, which allow for the haptic master device and slave robot to be decoupled. This simplifies the issue of forbidden region protection to position controlling the slave. Hands-on scenarios require an approach which considers the kinematics and dynamics of the robot in order to provide an appropriate impedance when constrained and a compliant response otherwise, as the surgeon is directly applying forces to the device.

2.3.2 Power Assist Systems

Power assist systems use input modification shared control strategies to allow the human user to more easily lift and/or carry heavy objects. Due to the limitations in their applicability to surgery, which will be discussed below, a complete review of power assist systems has not been performed and I refer the reader to [Gopura et al., 2011] and [Lee et al., 2012].

Power assist systems typically attempt to detect user intent to move using force sensors [Rahman and Ikeura, 2010] [Aguirre-Ollinger et al., 2012] and some combine this with electromyographic (EMG) sensors attached the skin to estimate user stiffness and damping [Gopura et al., 2009] [Farina and Merletti, 2000]. In both cases, these methods scale the user's applied forces, effectively reducing the mass of the object being lifted or maneuvered.

In [Rahman and Ikeura, 2010], it was shown that, in collaborative lifting scenarios, users would tend to apply too much force to the object. This was because, in the small time between the start of the user's push and the controller acting, they could sense the full weight of the object. These peak forces were reduced by using an exponentially decreasing effective mass impedance model.

The Berkeley Lower Extremity Exoskeleton used positive feedback of the system dynamics to amplify the system's sensitivity to external forces, without the use of force/torque sensors [Kazerooni et al., 2005]. However, this amplification treats all forces equally, such that external disturbances cause the system to respond in the same way as the user's commands. Additionally, the control strategy was shown to be very sensitive to modeling uncertainty and the stability of the system was shown to be less robust using the controller [Zoss et al., 2005].

Amplification of the user's forces through control can have a potentially dangerous effect. When a controller attempts to emulate dynamics that differ significantly from the intrinsic hardware dynamics, an increased risk of coupled or contact instability arises [Colgate and Hogan, 1989] [Buerger and Hogan, 2007]. In addition, as shown in [Kazerooni and Snyder, 1995], a very stiff impedance interacting with such controls would be unstable. In hands-on robotic surgery, the robot interacts with potentially deforming and/or pulsating environments, which can range in stiffness from hard contacts, such as bone, to deformable contact with soft tissue. The

complexity of this environment would make it very difficult to ensure stability. Thus, gaining regulatory approval for a potentially unstable controller in a clinical environment would be taxing.

Additionally, the amplification effect can increase interaction forces with tissue. The desired forces and torques, which the surgeon wishes to apply to the tool and, consequently, the anatomy, are based on the surgeon's medical training and understanding of the goals and decision making processes necessary to complete the surgery safely. The application of scaled forces and torques could potentially cause the surgeon to harm the patient unintentionally.

Lastly, the surgeon's forces can range from delicate (for interacting with soft tissue) to large (for gross positioning). This would require the force sensor to have a large range and still be of sufficient sensitivity to detect changes in more delicate interactions.

2.3.3 Manipulability Measures

Manipulability measures are used to find optimal poses in order to improve the performance of manipulators in specific tasks. When optimised in the null-space, these methodologies are degree-of-freedom shared control strategies, however, most manipulability measures have not been demonstrated in this capacity.

Kinematic manipulability ellipsoids were demonstrated in [Yoshikawa, 1985b] for velocity and force transmission, and a control strategy was presented for their maximisation in redundant robots. In [Chiu, 1988], these kinematic manipulability measures were extended to allow for directional force and velocity optimisation.

To account for the dynamics of the manipulator, the dynamic manipulability ellipsoid was presented in [Yoshikawa, 1985a] for optimising acceleration transmission. This measure was extended to account for the effects of gravity [Chiacchio, 2000]. The effects of velocity on dynamic manipulability were shown in [Rosenstein and Grupen, 2002], however, no extension to the measure to account for the effects of velocity was proposed.

An additional measure of the dynamic performance of a manipulator, the dynamic conditioning index, was shown in [Ma and Angeles, 1990] and aimed to quantify the amount of dynamic

coupling and the numerical stability of the inertia matrix. This measure was later applied to the end effector mass matrix by [Ficuciello et al., 2014] to quantify the coupling between the directions of motion at the tool. The authors applied an impedance controller with reduced inertia at the end effector to improve the ease of use for the user, while optimising this measure in the null-space in order to reduce the stability issues which occur when applying inertia reducing controllers. The authors found increased stability when using their controller, however, the force amplification issues discussed in Section 2.3.2, were not addressed.

The avoidance manipulability ellipsoid was developed in [Minami and Takahara, 2003] in order to quantify the ability of a manipulator to reconfigure its redundancy to avoid an obstacle, while still achieving its goals at the end effector. An index was created from this ellipsoid in [Tanaka et al., 2005] for obstacle avoidance during trajectory tracking and manipulator dynamics were included into the measure in [Kobayashi et al., 2013].

It was demonstrated in [Chiacchio, 2000] that the manipulability ellipsoids were an approximation of the actual performance of manipulators and that the abilities of a robot were more accurately represented by a polytope, or n-dimensional polygon. Several methods to optimise, according to the full polytope of various measures, have been developed.

From the acceleration polytope, the acceleration radius was proposed in [Graettinger and Krogh, 1988], which specified the lower bound on the acceleration of the end effector for a particular configuration. The acceleration radius was generalised in [Bowling and Khatib, 1998] to overcome the inhomogeneity problem in mixed manipulability optimisations of rotations and translations described in [Doty et al., 1995]. More recently, the power manipulability ellipsoid was proposed in [Mansouri and Ouali, 2011] as an alternative method for dealing with inhomogeneity. This measure was extended in [Choi et al., 2012] to the convex hull-based power manipulability index, allowing for analysis with respect to the power polytope.

While measures using polytopes have been shown to more accurately characterise manipulator behavior, they have yet to be demonstrated to function in real-time on a redundant system capable of achieving arbitrary poses in Cartesian space. They have been successfully applied to offline robotic design optimisation using the full polytope for manipulability in [San Martin et al., 2007] and [Yu and Liang, 2012], however, the aim of the work presented in this PhD

thesis is to optimise the interaction for a given robot online, as this allows the benefits of the optimisation to be used in current clinical systems without significant system redesign. Additionally, these measures have focused on the ability of the robot to act on its environment. For hands-on surgery, the surgeon is applying forces and torques to the robot and therefore, the measures I develop should optimise the way in which the robot reacts to the surgeon.

2.3.4 Cooperative Object Manipulation

Cooperative object manipulation involves a robot assisting a human in the maneuvering of an object. This area is a natural application for shared control, as the robot can simultaneously apply a control policy to ease the motion for the human. However, it can be difficult to apply a cooperative object manipulation strategy which is both responsive to the user's commands and ensures stability. Typical control in this area falls under the category of input modification, as the robot is often compensating for gravity and the coupling of the robot's impedance control strategy with the human's provides an inherent filtering.

One area of cooperative object manipulation research has aimed to quantify the control strategy employed in a task where two humans are cooperating to move an object. In this task, only one of the humans knows the desired trajectory or position of the object and the other is considered to be passive assistance. By determining the passive human's control strategy, this research aims to create a shared object manipulation strategy for the robot which is intuitive for the user to interact with due to its human-like characteristics.

Early work by [Ikeura et al., 1994] found the passive human's control strategy could be characterised by an impedance model consisting of a spring, mass, and damper. The most significant effect in the model was found to be the damper, whereas the mass and spring were nearly neglected. Using impedance control robots in this manner was shown to improve accuracy results over admittance controlled robots in [Kosuge and Kazamura, 1997].

This work was expanded to utilise a variable impedance model in [Ikeura and Inooka, 1995] and [Rahman et al., 1999], as it was found that humans exhibit a high impedance at low velocities for accurate positioning and a low impedance at high velocities for quick motion. Further, it was demonstrated by [Ikeura et al., 2002] that the human's variable impedance was

an optimal solution to a cost function consisting of a weighting of damping magnitude and change in damping. This solution again resulted in high impedance at low velocities and low impedance at high velocities.

An alternative variable impedance model was presented in [Tsumugiwa et al., 2002], which utilised an estimate of the user's arm stiffness from force and position data at the end effector to adjust the damping coefficient in the robot's impedance control. This control strategy was found to increase stability and accuracy in trajectory tasks as compared to the standard invariable impedance model.

In [Tsumugiwa et al., 2003], it was found that the human-like impedance control strategy for the robot was not precise enough to collaboratively perform a peg-in-hole task. Alternatively, a hybrid control strategy was devised which utilised the impedance model when moving toward the target and a torque control strategy consisting of gravity and friction during fitting.

While the previous methodologies aim to directly reproduce the control strategy a human would utilise in shared manipulation for positioning and trajectory tracking, an alternative strategy intends to modify the natural dynamics of the object through control in order to increase the ease of manipulation for the user. To overcome the high applied torques and unintuitive object motion which occurs when positioning long objects in typical collaborative control, a virtual holonomic constraint at the point held by the robot was proposed in [Takubo et al., 2002]. Modeling the constraint point as a virtual wheel restricted the possible motions of the object, simplifying the necessary steering and allowing for quicker and more accurate positioning.

2.3.5 Summary

Cooperative control has been applied to a variety of problem areas in general robotics and various improvements have been demonstrated which neither the human nor the robot could achieve alone. In this section, I have investigated the general robotics shared control strategies of active constraints, power assist systems, manipulability measures, and cooperative object manipulation. However, many of these strategies were found to be not applicable to clinical scenarios due to safety concerns or did not offer potential improvements for hands-on robotic surgery. In the next section, I will examine the current state-of-the-art in shared control for

surgery.

2.4 Shared Control in Surgery Today

Shared control has seen widespread usage in surgery since the first introduction of robotics for clinical applications. The first usage of robots in surgery was performed by [Kwoh et al., 1988], in which a PUMA 200 industrial robot replaced the manually adjusted stereotactic frame for positioning a surgical needle guide in brain tumor biopsy. After positioning, the robot was turned off and the surgeon inserted the needle through the guide. This degree-of-freedom shared control strategy allowed for quicker targeting and improved accuracy.

Even in the initial stages of surgical robotics, a common theme in shared control for surgery was present: to improve clinical outcomes while still leaving the surgeon in control of the critical portions of the surgery. In [Davies et al., 2004], it was suggested that systems with the surgeon in the loop are more easily accepted in a clinical environment than autonomous systems such as the Probot [Harris et al., 1997] and ROBODOC [Kazanzides et al., 1992] due to uncertainties about who is in charge of the procedure. Indeed, recent medical robotics literature surveys show many more surgeon-operated than autonomous systems [Beasley, 2012] [Bergeles and Yang, 2014].

In addition, locking the robot in place and leaving the surgeon in charge of the crucial task of needle insertion in [Kwoh et al., 1988] highlight the imperative for safety in shared control in surgery. Robotic systems for surgery typically contain many software fallback systems to ensure that the robot does not do anything unexpected and injure the patient. The Passive Arm with Dynamic Constraints (PADyC) aimed to ensure safety by implementing passive degree-of-freedom sharing and space sharing cooperative control strategies via mechanical clutches at the joint level [Troccaz and Delnondedieu, 1996] [Schneider and Troccaz, 2001]. In this way, the device could not move autonomously or exert forces by itself, but could still direct motion and protect the patient. However, this concept was not taken further due to the, sometimes, discontinuous nature of applying constraint at the joint level, as well as rapid wear of the clutched joints.

In this section, I will focus on current commercial systems in use clinically and research systems which are still in development. This is because surgical robotics has a high rate of system turnover due to the difficulties of achieving improvements to clinical outcomes under stringent requirements in the areas of equipment cost, time required for procedures, and safety concerns. We refer the reader to the medical robotics reviews in [Dogangil et al., 2010], [Beasley, 2012], and [Bergeles and Yang, 2014] for a more complete history of robotics in surgery.

Additionally, shared control has been used extensively in rehabilitation robotics, however, the majority of the work in this area has focused on active guidance of trajectories, which is not commonly used in surgery. Therefore, rehabilitation robotics will not be covered in this review and I refer the reader to the literature survey of [Marchal-Crespo and Reinkensmeyer, 2009] for information on this area.

2.4.1 Surgical Guides

"Surgical guides" are degree-of-freedom sharing robots which are used to position surgical tools. After being registered to the patient anatomy, these devices move to the appropriate positions for the surgery and are often fixed in place while the surgeon is performing the main task. The surgeon uses the guide to control the insertion or cutting degrees-of-freedom while the robot maintains the remaining positional coordinates, leading to a more accurate and/or quicker procedure, with the surgeon still in charge of the critical portions of the surgery. Due to their general simplicity and safer operation, as compared to systems which are active during the procedure, surgical guides have seen widespread clinical usage. However, their application remains solely in surgical areas with well-defined rigid anatomy to use for reference.

The Neuromate (Renishaw Plc.) is a stereotactic neurosurgical robot used for electrode implantation, neuro-endoscopy, and biopsy [Varma and Eldridge, 2006]. The system achieves the same accuracy as a conventional stereotactic frame, however, it results in a faster procedure as it can accurately reposition for multiple targets more quickly [Li et al., 2002]. Additionally, a reduction in clinical complication rates has been demonstrated [Cossu et al., 2005].

The Renaissance (Mazor Robotics Ltd.), called the SpineAssist in its first generation, is a surgical guide which is directly mounted to the patient's bone structure [Shoham et al., 2003].

The Renaissance is designed to have a small operating workspace, matching that required by minimally invasive procedures, as opposed to floor mounted general purpose systems like the Neuromate. This smaller design reduces the cost of the robot and increases patient safety due to lower inertia and smaller motors. Additionally, by attaching the robot directly to the patient's anatomy, the patient does not need to be immobilised or tracked in real time throughout the surgery to maintain registration [Shoham et al., 2007].

Another bone-mounted surgical guide is the iBlock (OMNIlife science Inc.), which is used for total-knee arthroplasty [Plaskos et al., 2005]. The device positions a planar bone-cutting guide that the surgeon uses to perform manual cutting. As compared to free-hand methods for bone-cutting, the iBlock has been shown to be more accurate and reduce preparation time due to the actuated positioning [Koulalis et al., 2011].

2.4.2 Motion Filtering and Scaling in Surgery

Motion filtering and motion scaling, which are input modification cooperative control strategies, have been used in teleoperative surgery to allow surgeons to operate more easily at the microscale. At this level, involuntary motions, such as tremor, can impact accuracy more severely [Riviere et al., 1998]. To reduce these inaccuracies, frequency filtering is used to extract the intended master motions. In motion scaling, motion commands of the master device are scaled down, such that the slave device moves a proportionally smaller distance than the master. Motion scaling has been shown to allow for higher accuracy, more so than tremor filtration in teleoperation [Prasad et al., 2004].

The da Vinci (Intuitive Surgical Inc.) is a commercial teleoperative surgical system for minimally invasive surgery which is claimed to have been used on over 2 million patients [Intuitive Surgical Ltd., 2015]. The device uses motion scaling and tremor filtering to make "delicate motions easier to perform" in the space limited surgical environment [Guthart and Salisbury, 2000]. The da Vinci consists of a surgical master console providing the surgeon with a 3D view of the surgery to improve spatial awareness of the surgical environment and a patient-side cart with 3 or 4 slave robotic arms for various tools and the 3D cameras.

While the effectiveness of the da Vinci as compared to laparotomy and conventional laparoscopy

has been established [Health Quality Ontario, 2010], the device does not provide haptic feedback, and this has been demonstrated to increase the chances of tissue damage. In an analysis of laparoscopic cholecystectomy surgeries, it was found that, in 15 out of 20 procedures, gallbladder perforation had occurred, 11 incidences of which were attributed to "too much force/rotation/displacement" [Joice et al., 1998]. The authors suggested that a potential cause was inappropriate levels of force being applied, due to diminished tactile and force feedback. In a simulated blunt dissection procedure, it was discovered that the lack of force feedback increased average force application to the tissue by at least 50%, peak force by a factor of 2, and the number of errors which damaged tissue by a factor of over 3 [Wagner et al., 2002]. As a result, a large amount of research has been performed into incorporating stable force feedback into a laparoscopic teleoperative setup [Westebring-van der Putten et al., 2008] [Puangmali et al., 2008].

The RAVEN II, developed at the University of Washington and University of California Santa Cruz, is an open-source teleoperative surgical platform [Lum et al., 2009]. The system consists of two 7 degree-of-freedom cable driven arms with remote center of motion mechanisms and provides an open API to connect various haptic devices [Hannaford et al., 2013]. Motion scaling has been demonstrated on the device [Jong Yoon et al., 2013] and tactile haptic feedback has been integrated through additionally developed sensors [Wottawa et al., 2013].

While tremor filtering is not effective in teleoperative scenarios, it is more effective in hands-on devices, where motion scaling is limited. The Micron is a hands-on microsurgical tool developed at Carnegie Mellon University designed to increase accuracy through active stabilisation of the tool tip [Riviere and Thakor, 1996]. The device uses both motion scaling and tremor filtering and has demonstrated a significant reduction in position error in experiments with both surgeons and non-surgeons [Maclachlan et al., 2012].

2.4.3 Motion Compensation in Surgery

To expand robotic surgery to dynamic soft tissue procedures such as beating heart surgery, teleoperative systems have used motion compensation, an input modification collaborative control technique. By sensing and predicting periodic organ motion and overlaying this information

on top of the surgeon's master device commands, the surgeon can operate on an effectively static organ. Due to camera data processing delays and servo time delays, research in this area focuses mainly on the prediction of the soft tissue motion in order to ensure the additional motion commanded to the slave matches, as accurately as possible, the current behavior of the dynamic anatomy.

In addition to the increased accuracy which these techniques can provide, they have also been shown to cause less damage to the organs as compared to conventional methods for operating on dynamic soft tissue that use mechanical stoppers or induce cardiac arrest. Procedures such as direct coronary artery bypass require precise hand motions and performing them on beating hearts requires strong mental concentration [Nakamura et al., 2001]. Using motion compensation can help to reduce this mental load on the surgeon, allowing him or her to focus more closely on the required task. Lastly, procedures which were previously very difficult to perform due to the quick nature of the tissue motion, such as mitral valve annuloplasty [Yuen et al., 2009], can be performed more easily due to these techniques. However, ensuring robust and accurate compensation and prediction in spite of time delays and soft tissue deformation variability still remains a research challenge.

The first research into incorporating motion prediction for surgery was, however, performed in a simulated hands-on surgery using a moving platform [Trejos et al., 1999]. The surgeon rested or attached his or her hand holding the tool to the platform and the platform was actuated in synchronisation with the motion of the organ being operated on. The system resulted in a higher accuracy over free-hand, however, task completion time was lengthened, even when a stationary image of the operation was shown to the user.

The first implementation of motion compensation for teleoperative surgery was demonstrated in [Nakamura et al., 2001] for minimally invasive cardiac surgery. An autoregressive model was used to track and predict the heart's motion for compensation, however, a large tracking error was present due to the speed of the camera feedback system.

To improve the robustness to irregularities in the quasi-periodic nature of the heart motion, it was proposed in [Ortmaier et al., 2005] to use a method based on Takens Theorem, which states that a sufficiently long vector derived from previous values in a time series can be used to

reconstruct the underlying system dynamics. ECG signals were utilised to estimate the heart's motion when the view of the tracking cameras was occluded. However, the performance of this method degraded in situations where there was too much variation in the tracked data.

The use of future heart motion prediction to increase accuracy was demonstrated in [Rotella, 2005] and [Ginhoux et al., 2005]. Model predictive controllers provided an estimate of the heart's position several time steps into the future. This method compensates for time delays in camera processing, resulting in a higher accuracy than the previously discussed methods.

The ECG signal from the heart was used in [Bebek and Çavusoglu, 2007] to estimate the period of the heartbeat for a receding horizon model predictive controller. Prediction and tracking were both improved as compared to estimation without ECG.

An extended Kalman filter (EKF) was utilised in [Yuen et al., 2009] to update the parameters of a time-varying Fourier series model which predicted and compensated for the 1D motion of mitral valve annulus. The EKF approach was shown to better handle changing heart rates and to have improved state prediction over a range of measurement noise, as compared to an AR model, as the EKF adjusted "to rapid changes in heart rate through explicit modeling of quasiperiodicity" [Yuen et al., 2009].

2.4.4 Active Constraints for Surgery

Active constraints for surgery were first demonstrated by [Ho et al., 1995] for precise cutting of the femur and tibia for mounting prosthetic knee components. Knee prosthetics require high precision cuts for a good alignment of the prosthesis, which ensures the joint performs well and has a long life. In addition, in this work, the robot was backdrivable and controlled by the surgeon holding a handle mounted at end effector, making this the first implementation of hands-on robotic surgery. Ho *et al.* allowed free motion in the interior of the cut and, as the surgeon approached a predefined constraint boundary, increased the stiffness of the robot, resulting in a very high stiffness at the edge of the forbidden region.

Since then, a significant amount of research has been performed on active constraints for surgical applications due to their potential to increase accuracy and safety. In this section, I will focus on

active constraints strategies which are currently being utilised on clinical devices and large scale research systems still in development. In addition, I will review the current state of dynamic active constraints research. For a full history of active constraints usage in surgery and general robotics, I refer the reader to [Hager et al., 2008] and [Bowyer et al., 2014].

Commercial Systems

Currently, there are three commercially employed surgical systems which use active constraints. The Sculptor, formerly the Acrobot, which has been recently acquired by Mako Surgical Corp, was developed from the first research into active constraints for surgery [Ho et al., 1995]. The system is a hands-on surgical robot which targets bone resection for knee prosthesis implantation. To reduce the mechanical impedance for the surgeon at the end effector, the device consists of a three axis spherical backdrivable manipulator mounted on a six axis large scale gross positioning device [Jakopec et al., 2003]. The system has demonstrated significant results in improving accuracy over conventional techniques using jigs and fixtures [Cobb et al., 2006].

The Rio (Mako Surgical Corp.) was developed as a competitor to the Sculptor and is currently being used in Total Hip Arthroplasty and Unicompartmental and Total Knee Arthroplasty [Lonner and Kerr, 2012]. The hands-on surgical system is backdrivable and utilises a tracking system to continuously align the planned active constraints intra-operatively. Similarly to the Sculptor, the Rio has demonstrated improved accuracy over conventional manual cutting techniques for hip and knee arthroplasty [Nawabi et al., 2013] [Citak et al., 2013].

The Navio Precision Freehand Sculpting system (Blue Belt Technologies Inc.), is a hand-held robotic tool which is used to resect bone in unicondylar knee replacement [Smith et al., 2014]. Alternatively to the other devices discussed in this section, the Navio implements active constraints by physically retracting the cutting burr as the tool enters into the forbidden area [Brisson et al., 2004]. This has the advantage of not requiring large motors or a braking system to limit the user's motion and allows the entire device to be hand-held.

All three of these commercial systems discussed in this section target rigid bone resection due to the simplification in registering the forbidden region, which can be performed via markers rigidly

attached to the bone structure. Additionally, in these surgeries, the constraint region does not need to be tracked and updated intra-operatively due to the lack of soft tissue deformation. These simplifications reduce the applicability of these devices to other surgeries involving soft tissue.

Large scale systems such as the Sculptor and Rio allow for multiple tools to be used on a single system. By doing so, the robot can assist in several tasks, as opposed to the Navio which can only perform one. However, by attaching a full robotic arm to the tool, the surgeon must additionally interact with the dynamics of the manipulator which can include joint friction, backlash, the mass and inertia of the robot links and other disturbances, such as interactions with the sterile draping. These effects can contribute to a sense of bulkiness and drag, which can reduce the surgeon's comfort and ease of use in manipulation tasks, particularly over the length of current hands-on robotic surgical procedures, which can range from 40 to 120 minutes [Lang et al., 2011].

Research Systems

In addition, there are several large scale research systems which utilise active constraints and various other shared control techniques.

The Johns Hopkins University Steady-Hand Robot has seen continued development for over a decade [Taylor et al., 1999] [Mitchell et al., 2007] [He et al., 2012] and several different types of active constraints strategies have been tested on the device [Kumar et al., 1999] [Li and Okamura, 2003] [Bettini et al., 2004] [Li et al., 2007]. The hand held smart tool is intended for microsurgeries, such as vitreoretinal surgery similarly to the Micron [Maclachlan et al., 2012] (discussed in Section 2.4.2). Through input modification shared control strategies, such as tremor filtration [Mitchell et al., 2007] and scaled motion [Kumar et al., 2000], the Steady-Hand aims to improve the accuracy of the surgeon operating at such a small scale. However, as the system consists of a small scale admittance controlled robot to which the surgical tool is attached, active constraints and scaled force feedback [He et al., 2012] can also be implemented.

The Neuroarm (University of Calgary) is a MR-compatible teleoperative robotic platform in-

tended for neurosurgery [Sutherland et al., 2003]. By providing real-time images to the surgeons, they "can correct for brain shift and can ensure complete removal of a tumor during operation" [Sutherland et al., 2013]. Motion scaling has been demonstrated on the system and, due to the presence of haptic feedback, other forms of shared control such as active constraints are possible.

The Mirosurge (DLR, German Aerospace Center), developed from the Kinemedic, is a multi-robot surgical platform which can operate in both hands-on and teleoperative modes [Hagn et al., 2010]. The system consists of several MIRO robots [Hagn et al., 2008] which are directly attached to the operating table and a surgical console made up of two Omega.7 haptic devices (Force Dimension Ltd.) and a Miracube 3D display (Miracube Inc.). The system exhibits several forms of shared control, including guidance active constraints [Ortmaier et al., 2006], degree-of-freedom sharing in hands-on mode [Konietschke et al., 2006] and motion scaling in teleoperative mode [Konietschke et al., 2009].

The Active Constraint Technologies for Ill-Defined or Volatile Environments (ACTIVE) FP7 project is an "integrated redundant robotic platform for neurosurgery" [ACTIVE, 2012]. This multi-university collaborative project incorporates two seven degree-of-freedom LWR 4+ robots (KUKA GmbH.), two 7 degree-of-freedom Sigma.7 haptic devices (Force Dimension Ltd.), and the Active Headframe, a novel actuated device for head support. The project incorporates active constraints and can operate in both teleoperative and hands-on modes. The author of this thesis was involved in the software implementation of the active constraints within this project and their integration with the overall system.

Similarly to the commercial devices discussed in the previous section, which utilise active constraints, hands-on systems using large scale robots impose additional dynamics on the surgical tool and, hence, the surgeon. The Mirosurge and the Kuka LWR 4+ used in the ACTIVE project aim to reduce the effect of the end effector dynamics on the surgeon through lightweight system designs. However, neither of these systems utilise control to make the interaction of the surgeon with the robot closer to the experience of the tool unattached to the robot.

Dynamic Active Constraints

Recently, techniques to expand the capabilities of active constraints to handle the dynamic nature of soft tissue surgery have been explored. To ensure the protection of delicate tissue structures, these dynamic active constraints must be able to handle the arbitrary shapes and deformations which can occur from tool interactions and patient heartbeat and breathing.

Dynamic virtual fixtures using potential fields were created in [Ren et al., 2008] for beating heart surgery. The constraints were precomputed from preoperative MR and CT images and the authors proposed synchronising the constraints with the patient using ECG signals and intra-operative ultrasound. Due to their precomputed nature, the constraints were unable to handle unpredicted motion, such as deformation due to contact with tools or abnormal changes in cardiac behavior.

A one degree-of-freedom dynamic regional constraint was demonstrated by [Gibo et al., 2009] for assisting in a teleoperative positioning task in which the user attempted to depress an actuated soft tissue phantom to a particular depth. Two methods for enforcing the constraint were tested; one based on the current position of the tissue and one based on a prediction of its position. Experimental results found that the methodology was an improvement over static constraints and unconstrained teleoperation, however, the two constraint methods tested by the authors had comparable results.

Real-time MR imaging was used in [Navkar et al., 2012] to generate dynamic active constraints for assistance in beating heart surgery. The constraints aimed to prevent the tool from contacting the inner walls of the heart by applying forces towards the center of constrained pathway. Experimental testing in a simulated scenario demonstrated an improvement in off-path error as compared to both visual-only guidance and no guidance.

A form of active constraints referred to as dynamic frictional constraints, distinguish between when the user has actively moved the tool into the forbidden region and when the dynamic constraint has moved onto where the tool is positioned [Bowyer and Rodriguez y Baena, 2014]. These constraints only act when the former situation has occurred as, in the latter case, the tool “is no longer reactive to the actions of the surgeons, instead, it is actively generating tool

motion of its own volition" [Bowyer, 2014]. To prevent this, a fully dissipative friction-based control policy was presented which ensures the active constraint does not take control of the tool's motion, leaving the surgeon fully in control at all times.

To date, no work has been presented on dynamic active constraints which can handle arbitrary shapes and deformations. Early work considered single point constraints and the work of [Navkar et al., 2012] was specific to beating heart surgery. While the constraints presented in [Bowyer, 2014] provide a unique strategy for handling safety in dynamic soft tissue surgery, the control methodology requires a meshed representation of the constraint geometry. Real-time non-rigid registration of soft tissue is an ongoing research problem [Faria et al., 2014] [Weon et al., 2015] and therefore, real-time updating of constraint meshes to ensure patient protection is difficult. Alternatively, a method which can utilise streaming point cloud data from 3D cameras and range finders, in addition to mesh data, would be able to handle arbitrary and deforming constraint geometries in real-time.

2.5 Summary

This chapter has presented a literature review on the current state-of-the-art in shared control for both general and surgery specific robotics. Current hands-on robotic devices are bulky, making them physically hard to maneuver. However, hand-held smart tools are typically surgery specific and therefore, not easily adaptable to new tasks like large scale robotic systems. A control strategy which can improve the transparency of the human-robot interaction to make it more natural — that is, closer to the tool unattached to the robot, without affecting the surgeon's usability — could reduce fatigue and improve the surgeon's ease of use, leading to increased acceptance of such devices.

To this end, in Chapters 3 and 4, a novel hands-on degree-of-freedom sharing cooperative control strategy is developed for optimising the mass, inertia, frictional force, and frictional torque at the end effector of a redundant robot, without affecting the pose of the tool. The efficacy of the proposed optimisation technique is demonstrated through extensive simulations and a user-based experimental setup built around a commercial robotic platform.

In addition, current dynamic active constraint approaches require structured representations of the underlying geometry and can not handle arbitrary deformations. A dynamic active constraints approach which does not require *a priori* knowledge of the structure to be constrained could bypass the need for non-rigid registration and feature mapping, and directly constrain a streaming point cloud representation of the anatomy.

In Chapter 5, a novel shared control technique for hands-on robotic surgery is presented for generating meshless dynamic active constraints directly from point clouds using implicit surfaces. The combined strategies are validated in a virtual reality ultrasound scenario in Chapter 6.

Chapter 3

Redundant Mass and Inertia Optimisation for Natural Motion

This chapter presents a novel hands-on control methodology for optimising the mass and inertia at the end effector of a redundant robot, without affecting the pose of the tool. By reducing these effects at the tool point, the surgeon can move the surgical instrument more naturally, closer to the tool unattached to robot and therefore, dexterity is increased and fatigue is reduced. The research presented here is an edited version of the work published in:

> J. G. Petersen and F. Rodriguez y Baena, "Mass and inertia optimization for natural motion in hands-on robotic surgery", in Proceedings of the IEEE International Conference on Intelligent Robots and Systems, pp. 4284-4289, Sept. 2014.

3.1 Introduction

As discussed in Chapter 1, hands-on robotic surgery is advantageous because it enables the forces and torques of the interaction between the robot and the environment to be felt directly through the tool, without the need for an additional force sensor at the end effector. The surgeon can utilise this natural feedback to more precisely apply forces and torques to the hard or soft target tissue, a valuable feature when performing surgery with robotic assistance. In addition, the surgeon is in control of the speed and direction of the surgical task, which has lead

to increased acceptance of these types of systems, as compared to autonomous robots [Davies et al., 2004].

However, by attaching the surgical tool to the end effector of the robot, the surgeon must also interact with the dynamics of the robot at the tool point. These dynamics can include joint friction, gear backlash, the mass and inertia of the robot links, and other disturbances, such as interactions with the sterile draping. Each of these factors can reduce the dexterity of the surgeon and increase fatigue, particularly over the time frame of current hands-on procedures, which can range from 40 to 120 minutes [Lang et al., 2011]. The aim of the work presented in this chapter is to reduce these end effector dynamics. This will achieve a more natural feeling for the surgeon when performing hands-on robotic surgery, by providing an experience which is closer to the feeling of the tool unattached to the robot.

The effects of the system dynamics on the user are more apparent in large scale systems such as the Sculptor (formerly Stanmore Implants Worldwide Ltd., recently acquired by Mako Surgical Corp.) [Cobb et al., 2006] and the RIO (Mako Surgical Corp.) [Hagag et al., 2011], making them bulky and physically harder to maneuver. Hand-held smart tools such as the Steady-Hand Robot [He et al., 2012] do not have a large dynamic effect on the surgeon, however, they are typically surgery specific and, as a result, not easily adaptable to new tasks like large scale robotic systems. Therefore, a natural motion solution for large scale systems can, in principle, provide the largest possible benefits for the surgeon and these improvements could affect a wider range of surgeries.

One strategy for modifying the end effector dynamics is to mechanically design the robot to have low mass and inertia at the tool point. Early hands-on robotic systems, such as the Acrobot [Jakopec et al., 2003], employed specially designed mechanisms with a limited number of degrees-of-freedom to reduce these effects. In [Khatib and Burdick, 1987], a methodology for creating a robot with uniform and isotropic bounds on the end effector acceleration through optimisation of the link masses, inertias, and lengths was presented. The German Aerospace Center (DLR) has developed two robots which aim to have low inertia and mass in order to enable safe human-robot collaboration: the Light-Weight Robot (LWR) [Albu-Schäffer et al., 2007] and the MIRO [Hagn et al., 2008].

Robotic surgical systems are quite costly to develop as they have special requirements for safety and sterilisability and can take years of testing for clinical certification. A significant redesign of a current surgical system would also be required to undergo recertification. Additionally, for some systems, reducing the dynamic effects on the surgeon may not be possible without compromising performance goals. Furthermore, the tool point dynamics can still be varied in lightweight redundant systems for a given pose using control and so, further improvements upon design considerations can be made by optimising in the null-space. Therefore, a control solution to minimise the effects of the robot's end effector dynamics on the surgeon can offer improvements to current clinical systems, without the need for a costly system redesign.

Control strategies exist in the general robotics literature which aim to ease the user's ability to maneuver objects in collaborative tasks. Power assist systems, for which literature reviews can be found in [Gopura et al., 2011] and [Lee et al., 2012], amplify the user's forces and torques such that the effective mass of the object the human is interacting with is reduced, allowing the user to more easily lift and/or maneuver heavy objects. The user's forces are typically detected with force sensors as in [Aguirre-Ollinger et al., 2012] and [Ammar et al., 2010], and some methods utilise EMG sensors attached the user's skin to provide an estimate of the subject's arm stiffness and damping, as in [Gopura et al., 2009] and [Farina and Merletti, 2000].

However, amplification of the user's forces through control can have potentially dangerous effects. An increased risk of contact or coupling instability arises when an impedance controller attempts to emulate dynamics that differ significantly from the intrinsic hardware dynamics [Colgate and Hogan, 1989] [Buerger and Hogan, 2007]. Additionally, it was shown in [Kazerooni and Snyder, 1995] that a very stiff impedance interacting with such controls would be unstable. The environment in hands-on robotic surgery is quite complex, as the target tissue may pulsate and deform and the stiffness of the tissue can range from hard contacts, such as bone, to soft contacts, such as heart, liver, or brain. Due to the complexity of this environment, ensuring stability would be difficult and, therefore, such a controller could not be applied in a clinical scenario.

In addition, amplifying the user's wrench can increase the surgeon's interaction forces and torques with the tissue. Surgeons use their medical training and understanding of the goals and

decision making processes necessary to complete the surgery safely, to apply appropriate forces and torques to the patient through the surgical tool. Scaling the surgeon's applied forces and torques could potentially cause the surgeon to unintentionally harm the patient through the robot. Similarly, the force and torque scaling in these type of impedance controllers affects the surgeon's ability to accurately feel the forces and torques of the interaction between the surgical tool and the tissue through the end effector. Thus, such strategies remove an essential benefit of hands-on robotic surgery.

Lastly, these types of impedance controllers require force/torque sensors at the tool point in order to measure the user's desired input. However, for surgery, these sensors are required to undergo harsh sterilisation procedures, typically performed through heat and/or chemical sterilisation [Okamura et al., 2010]. Thus, specially designed force/torques sensors are necessary to handle such sterilisation processes, driving up the already high cost of these devices.

In [Ficuciello et al., 2014], an impedance controller with reduced inertia was applied at the end effector of the robot in a collaborative drawing task in order to improve the subject's ease of use. In addition, the dynamic conditioning index [Ma and Angeles, 1990] was optimised in the redundancy to minimise the coupling between the Cartesian degrees-of-freedom at the end effector in order to reduce the above mentioned stability issues, which occur when applying inertia reducing controllers. The region of stability for the controller was found to be increased using the controller, however, the effects of force amplification on the tool's interaction with the surface were not addressed.

Therefore, as an alternative to using an impedance controller at the tool point, a control strategy which minimises the end effector dynamics using the redundant degrees-of-freedom would be advantageous, as it avoids the stability and safety problems with tool point impedance control strategies and ensures that the surgeon can accurately sense the tool interactions with the environment directly through the surgical instrument.

For these reasons, a null-space based optimisation control methodology has been developed, which reduces the effect of the end effector dynamics on the surgeon. By creating a control strategy rather than a design solution, expensive design or redesign and certification work is avoided and the methodology could potentially be applied to current clinical hands-on robotic systems.

In addition, by reducing the end effector dynamics in the redundant degrees-of-freedom, stability and safety issues with impedance controllers applied at the end effector are avoided.

This chapter will first derive model based representations of the mass, inertia, frictional force, and frictional torques at the end effector of a general n-degree-of-freedom manipulator. Following this, measures for mass and inertia are developed for directional and overall optimisation, and combined mass and inertia optimisation methodologies are demonstrated. Friction force and torque optimisation will be performed in Chapter 4. A null-space optimisation control strategy is presented to reduce the effect of these tool point dynamics on the surgeon, without affecting his or her desired commands. Simulation results show the extent of the reduction in mass and inertia, which the optimisation strategy provides over the workspace of a four degree-of-freedom planar robot, and the methodology's sensitivity to uncertainties in the dynamic model parameters is investigated. Combined mass and inertia optimisations are performed to examine the relative reductions in these quantities as compared to single optimisations. Lastly, simulated trajectory trial results demonstrate a reduction in the time required to move the tool through a trajectory, as compared with conventional redundancy resolution techniques in hands-on robotic surgery.

3.2 End Effector Dynamics

When maneuvering the tool attached to the robot, the surgeon is interacting with the end effector dynamics of the robot. In this section, I will derive these configuration dependent dynamics from the joint space dynamics for later optimisation.

The joint space dynamics of an n degree-of-freedom serial manipulator can be written in the Lagrangian formalism as,

$$A(q)\ddot{q} + b(q, \dot{q}) + g(q) + \tau_f = \tau_m + \tau_j \tag{3.1}$$

where $A(q)$ is the $n \times n$ inertia matrix, $b(q, \dot{q})$ is the $n \times 1$ vector of Coriolis and centrifugal torques, $g(q)$ is the $n \times 1$ vector of gravity torques, τ_f is the $n \times 1$ vector of joint friction torques,

and τ_m is the $n \times 1$ vector of joint motor torques. τ_j is the projection of the user's end effector forces and torques onto the joints, that is, τ_j is the torques at the joints resulting from the user's forces and torques at the end effector.

$$\tau_j \triangleq J^T \begin{bmatrix} F_u \\ \tau_u \end{bmatrix} \tag{3.2}$$

where F_u and τ_u are the forces and torques applied by the user at the end effector.

The dynamics of the end effector position and orientation can be computed from the joint space dynamics by multiplication with the transpose of the dynamically consistent inverses of the linear velocity Jacobian, J_v, and the angular velocity Jacobian, J_ω, respectively,

$$\overline{J}_v = A^{-1} J_v^T \Lambda_v \text{ and } \overline{J}_\omega = A^{-1} J_\omega^T \Lambda_\omega \tag{3.3}$$

where $\Lambda_v = \left(J_v A^{-1} J_v^T\right)^{-1}$ is the pseudo kinematic energy matrix associated with mass and $\Lambda_\omega = \left(J_\omega A^{-1} J_\omega^T\right)^{-1}$ is the pseudo kinematic energy matrix associated with inertia.

The resulting equations of motion provide a mathematical description of how the surgeon and robot interact at the tool point. The equations of motion associated with linear motion can be expressed as,

$$\Lambda_v \ddot{x} + v_v + p_v + \overline{J}_v^T \tau_f = F_m^{EE} + F_u \tag{3.4}$$

where v_v is the vector of Coriolis-centrifugal forces at the end effector, p_v is the vector of end effector gravity forces, $\overline{J}_v^T \tau_f$ is the projection of the frictional torques onto the linear end effector dynamics and $F_m^{EE} \triangleq \overline{J}_v^T \tau_m$ is the vector of forces at the end effector resulting from motor torques. The equations of motion associated with angular motion can be written as,

$$\Lambda_\omega \dot{\omega} + v_\omega + p_\omega + \overline{J}_\omega^T \tau_f = \tau_m^{EE} + \tau_u \tag{3.5}$$

where v_ω is the vector of Coriolis-centrifugal torques at the end effector, p_ω is the vector of end

effector gravity torques, $\overline{J}_{\omega}^{T}\tau_f$ is the projection of the frictional torques onto the angular end effector dynamics and $\tau_m^{EE} \triangleq \overline{J}_{\omega}^{T}\tau_m$ is the vector of torques at the end effector resulting from motor torques.

The projections of the frictional torques onto the end effector, $\overline{J}_{v}^{T}\tau_f$ and $\overline{J}_{\omega}^{T}\tau_f$, are the frictional forces and torques which the surgeon feels when moving the tool as a result of the friction on the joints. We will define the effective end effector friction force, F_f^{EE} and the effective end effector friction torque, τ_f^{EE} as these projections.

$$F_f^{EE} \triangleq \overline{J}_{v}^{T}\tau_f \text{ and } \tau_f^{EE} \triangleq \overline{J}_{\omega}^{T}\tau_f \tag{3.6}$$

Assuming the motor torques applied to the space of the end effector are used only to compensate for gravity and that Coriolis-centrifugal forces are small, the resulting dynamic systems which the surgeon interacts with consist of the operational space mass and inertia and the effective end effector frictional forces and torques.

$$\Lambda_v\ddot{x} + F_f^{EE} = F_u \text{ and } \Lambda_\omega\dot{\omega} + \tau_f^{EE} = \tau_u \tag{3.7}$$

From these equations, two factors can be identified which affect the way in which the robot interacts with the surgeon. First, the resulting accelerations from the surgeon's forces and torques are scaled by the mass and inertia of the manipulator. The surgeon performs surgery with a set of tools, each of which has specific inertial properties that the surgeon is likely accustomed to. Attaching these tools to the robot modifies the mass/inertia which the surgeon experiences when using them. To provide a more natural response of the tool to the surgeon's forces (i.e a closer experience to the tool being handled while not attached to the robot), I will minimise the end effector inertial properties without affecting the surgeon's desired tool pose.

Second, the effective frictional forces and torques reduce the impact of the surgeon's forces and torques, requiring the surgeon to apply more forces to achieve the same acceleration for the same mass/inertia. Clearly, the tool alone does not exhibit friction of this type. To reduce this

effect, I will minimise the friction forces and torques without changing the surgeon's desired end effector position and orientation.

This chapter will develop a null-space optimisation strategy to reduce the mass and inertia at the end effector of the robot, without affecting the surgeon's commands at the tool point. Friction force and torque optimisation and combined mass and friction optimisation will instead be addressed in Chapter 4.

3.3 End Effector Mass and Inertia Optimality Measures

3.3.1 Overall Mass and Inertia Optimality Measures

As shown in Section 3.2, the end effector mass and inertia with which the surgeon is interacting can be characterised by the pseudo kinetic energy matrices for mass, Λ_v, and inertia, Λ_ω, respectively. As the inverses of these matrices are guaranteed to exist and are more robust to singular or near singular configurations [Hogan, 1984], I will use the inverses in creating our optimisation criteria.

The effective mass and inertia at the end effector were shown to be representable geometrically as belted ellipsoids [Khatib, 1995], in the form depicted in Figure 3.1 for the 2D case. The shape of this ellipsoid depends on the configuration of the robot and, if the robot is redundant, the configuration of the redundancy. This means that I can reconfigure the redundancy to achieve a different effective mass/inertia for the robot, without affecting the surgeon's desired end effector pose.

In order to minimise or maximise either the mass or the inertia, I will examine the Singular Value Decomposition (SVD) of the symmetric inverse pseudo kinetic energy matrix.

$$\Lambda^{-1} = U\Sigma U^T \tag{3.8}$$

where U is the orthogonal matrix consisting of the singular vectors of Λ^{-1} and Σ is the diagonal matrix of singular values. The singular vectors correspond to the principal axes of the

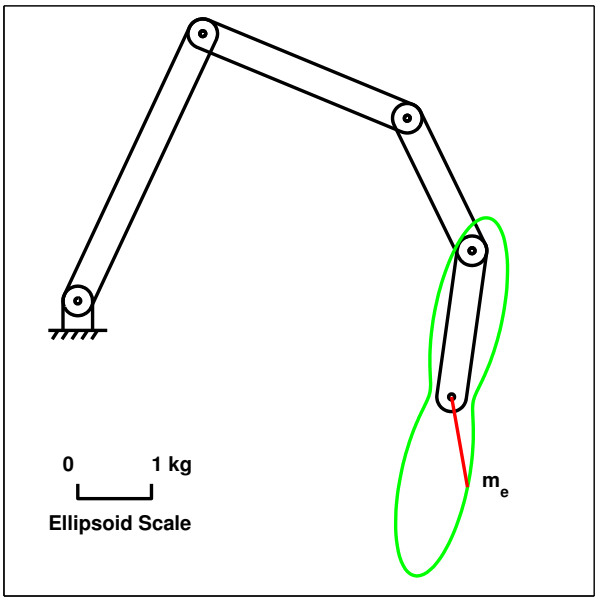

Figure 3.1: The belted mass ellipsoid [Khatib, 1995] overlaid on a simulated four degree-of-freedom planar robot with the effective mass, m_e, for a particular direction highlighted.

mass/inertia ellipsoid and the singular values correspond to inverse of the mass/inertia in the principal directions.

As the determinant of an orthogonal matrix is ± 1, the determinant of the inverse mass/inertia matrix is equal to the product of its singular values.

$$\det(\Lambda^{-1}) = \prod_{i=1}^{m} \sigma_i \tag{3.9}$$

Therefore, to minimise the mass or inertia, the negative of the determinant of the inverse pseudo kinetic energy matrices should be minimised.

$$c_{det} = -\det(\Lambda^{-1}) \tag{3.10}$$

So far, I have considered mass and inertia separately. This is due to the inhomogeneity resulting from an attempt to combine the two [Doty et al., 1995]. To simultaneously optimise mass and inertia, a weighted sum can be used, where the weights reflect the relative importance to be given to rotations, as opposed to translations.

$$c_m^c = w_v \det(\Lambda_v^{-1}) + w_\omega \det(\Lambda_\omega^{-1}) \tag{3.11}$$

where w_v and w_ω are the weights of the mass and inertia optimisations respectively.

3.3.2 Directional Mass and Inertia Optimality Measures

In addition to optimising the overall mass and inertia at the end effector, specific directions can also be optimised to create a preferred direction of motion that could be used in cases where there is frequent motion along a particular pathway or where the intended direction is known *a priori*. The mass or inertia felt when applying a force or torque in a specific direction is referred to as the effective mass/inertia [Khatib, 1995]. The effective mass (m_e) for a particular direction is shown in Figure 3.1.

The effective mass or inertia in a specified direction can be computed by,

$$\frac{1}{m_u} = u^T \Lambda_v^{-1} u \text{ and } \frac{1}{I_u} = u^T \Lambda_\omega^{-1} u \tag{3.12}$$

where u is the unit vector in the desired direction, and m_u and I_u are the effective mass and inertia respectively, in the direction of u. As these are the quantities that I am directly interested in modifying, I use these as our measure for optimisation.

$$c_m^d = -u^T \Lambda_v^{-1} u \text{ and } c_i^d = -u^T \Lambda_\omega^{-1} u \tag{3.13}$$

Additionally, multiple directions can simultaneously be optimised by a weighted sum of the effective mass or inertia in the specified directions.

$$c_w^d = -\sum_{i=1}^{n} w_i u_i^T \Lambda_v^{-1} u_i - \sum_{j=1}^{m} w_j u_j^T \Lambda_\omega^{-1} u_j \tag{3.14}$$

where w_i and w_j are the weights associated with i-th effective mass and j-th effective inertia, respectively.

3.4 Redundant Optimisation of End Effector Mass

For each of the natural motion metrics discussed above, the optimisation problem I wish to solve is to find the joint angles which minimise the specified metric, subject to the constraint that the joint angles must result in a pose equal to the surgeon's current pose on the entire path between the initial configuration and optimum.

$$\begin{aligned} \underset{q}{\text{minimize}} \quad & c(q) \\ \text{subject to} \quad & \int_{t_0}^{t_f} \|x(q_0) - f(r(t))\| \, \mathrm{d}t = 0 \end{aligned} \tag{3.15}$$

where $r(t)$ is a parametrisation of the path between the initial joint angles and optimum such that $r(t_0) = q_0$ and $r(t_f) = q_{opt}$ and $f(q)$ is the forward kinematics of the manipulator. Note that this constraint is stronger than simply requiring the end effector pose corresponding to the optimum joint angles to be equal to the surgeon's initial pose. As the null space manifold can be discontinuous, joint angles can exist which satisfy the desired pose, but cannot be reached from the current joint angles without violating the desired end effector pose in the transition. Therefore, only solutions which are reachable from the current joint position are valid.

To simplify this problem and ensure it remains solvable in real-time, I use a local optimiser which controls the redundancy to the local end effector dynamics measure optima, relative to the current set of joint angles. The control necessary to move towards the local optima is computed and projected into the null space of the current end effector pose. The local optima to the current pose can be reached by simply computing the gradient of the optimality measure and performing gradient descent. A velocity saturation method for gradient descent was presented in [Sentis, 2007] which I will use here,

$$\begin{aligned} \omega_{des} &= -\frac{k_p \nabla c}{k_d} \\ v_d &= \min\left(1, \frac{\omega_{max}}{|\omega_{des}|}\right) \\ F_p &= -k_d(\dot{q} - v_d \omega_{des}) \end{aligned} \tag{3.16}$$

where k_p is a spring gain, k_d is a damping gain, and ω_{max} is the maximum allowed velocity.

Null-space controllers are used to simultaneously achieve a hierarchy of tasks such that the tasks do not affect each other and higher priority tasks take precedence over lower priority tasks in the event of a loss of degrees of freedom. We use the dynamic recursive null-space formulation presented in [Sentis, 2007] in order to allow for concurrent dynamic active constraints application [Petersen and Rodriguez Baena, 2013] and towards our goal of developing a complete hands-on robotic control system. Equation (3.17) shows the general form of the recursive null-space controller which allows N hierarchical tasks to operate concurrently.

$$\Gamma = \sum_{k=1}^{N} J_{k|prec(k)}^{T} F_{k|prec(k)} \tag{3.17}$$

where $J_{k|prec(k)} = J_k N_{prec(k)}$ is the Jacobian of the k-th task operating in the null space of the previous $k-1$ tasks, $N_{prec(k)} = I - \sum_{k=1}^{N} \overline{J}_{k|prec(k)} J_{k|prec(k)}$ is the combined null space of the higher order tasks, and $F_{k|prec(k)}$ represents the forces of the k-th task acting in the null space of the previous $k-1$ tasks.

As our goal is to optimise while not affecting the desired tool position and orientation of the surgeon, the work in this chapter uses two tasks: dynamic compensation of the point in contact with the surgeon and the null space optimisation. The two-task version of (3.17) used here is as follows:

$$\Gamma = J_t^T \left(u_t + p_t\right) + N_t^T J_p^T \left(F_p + u_p + p_p\right) \tag{3.18}$$

where J_t is the linear and angular velocity Jacobian of the contact point with the surgeon, u_t and p_t are the Coriolis-centrifugal and gravity forces respectively at the contact point with the surgeon, J_p is the Jacobian of the posture, F_p represents the optimisation forces (3.16), u_p and p_p are the Coriolis-centrifugal and gravity forces respectively of the posture, and $N_t = I - \overline{J}_t J_t$ is the dynamically consistent null space of J_t.

Using this strategy, the robot can reconfigure without affecting the surgeon, to achieve the optimal local configuration in the null-space.

3.4.1 Additional Optimality Measures

In addition to mass/inertia optimisation, this formulation allows for the addition of further terms to cater for other criteria using a weighted summation. Safety can be improved using gradient based methods such as joint limit avoidance [Liegeois, 1977] [Klein and Huang, 1983]; obstacle avoidance to prevent collisions with the patient, surgeon, or equipment [Maciejewski and Klein, 1985] [Choi and Kim, 1999]; and self collision avoidance [Sugiura et al., 2006] [Stasse et al., 2008].

3.5 Simulations

To verify our optimisation strategy, a four degree-of-freedom planar manipulator was implemented in MATLAB 2013a (Mathworks Inc.) and its dynamics were simulated using MATLAB's ODE functions. Four degrees-of-freedom were chosen as they provide a single degree of redundancy when constraining the end effector's position and orientation in a planar scenario. The parameters of this model are shown in Table (3.1) and the center of mass for each link was located at half the link's length. These values have been chosen to reflect the link length and inertial distribution found in general robots with decreasing length and inertial properties tending away from the base degree-of-freedom. Further natural motion simulations and experiments using the parameters of a real system will be performed in Chapter 4.

Link	**Length (m)**	**Mass (kg)**	**Inertia ($kg*m^2$)**
1	1.00	2.00	0.010
2	0.75	1.50	0.005
3	0.50	0.75	0.005
4	0.50	0.50	0.001

Table 3.1: The kinematic and dynamic parameters used in the simulation of the four degree-of-freedom planar robot.

3.5.1 Effectiveness of Mass Optimisation Over the Robot's Workspace

To demonstrate the general effectiveness of the mass optimisation strategy presented here, it must be shown that for a majority of end effector poses over the robot's workspace, there is

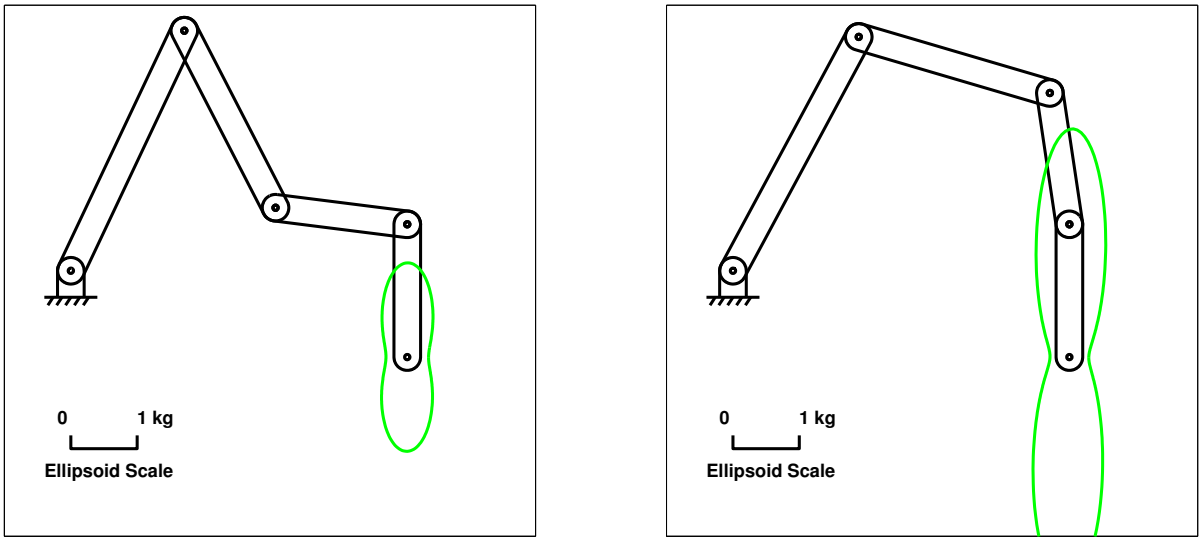

(a) Minimum Mass Determinant (b) Maximum Mass Determinant

Figure 3.2: The local minimum and maximum belted mass ellipsoids for a single end effector position and orientation overlaid on their corresponding joint configurations.

a substantial improvement in the mass properties at the joint pose corresponding to the local minimum of the mass metric as compared to the joint point corresponding to the local maximum. The local minimum and maximum poses are the joint configurations which correspond to the local minimum and maximum of the optimality measures, relative to the initial joint angles. Figure 3.2 depicts the minimal and maximal mass configurations for a single end effector pose with the belted mass ellipsoids overlaid.

Therefore, the mass determinant of 1000 random, initially nonsingular configurations of the robot were maximised and minimised. 1000 trials were chosen as I believe this covers a majority of the workspace of the robot. Additionally, this number of trials was chosen as it allowed them to be run within a reasonable amount time, approximately 12 hours for the 2000 optimizations. Initial joint configurations were found by drawing from a uniform distribution over $[0, 2\pi)$ for each joint and checking for singularities by computing the singular value decomposition of JJ^T. Keeping the initial position and orientation of the end effector fixed, the robot used (3.18) to optimise (3.10) from the initial joint parameters to find the local minimum. A proportional-derivative (PD) controller with gains $k_{p,ee} = 30$ and $k_{d,ee} = 5$ was applied to the end effector to eliminate integration drift and the parameters of (3.16) were set as $k_p = 50$, $k_d = 15$, and $\omega_{max} = 3rad/s$. As (3.16) is used to minimise the optimisation criterion, a modified version

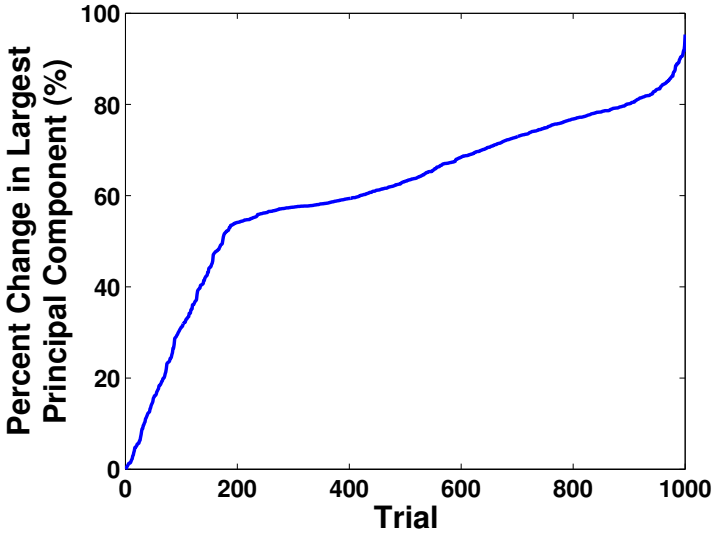

Figure 3.3: The percentage change in the largest principal mass component between the minimum and maximum joint configurations for randomly generated end effector poses.

was used to obtain the local maximum:

$$\begin{aligned} \omega_{des} &= \frac{k_p \nabla c}{k_d} \\ v_d &= \min\left(1, \frac{\omega_{max}}{|\omega_{des}|}\right) \\ F_p &= -k_d(\dot{q} - v_d \omega_{des}) \end{aligned} \tag{3.19}$$

Figure 3.3 depicts the percentage change in the largest principal mass component between the maximum and minimum mass determinants for each of the random trials, arranged in ascending order. For this specific robot, the minimum configuration represents, on average, a 60.85% reduction in the largest principal component of the mass ellipsoid relative to the maximum configurations with a standard deviation of 19.35%. This difference in the end effector mass properties of the local minimum and local maximum demonstrates a reduction in mass over the robot's workspace when using our optimisation technique.

3.5.2 Effectiveness of Inertia Optimisation Over the Robot's Workspace

Additionally, to demonstrate the effectiveness of the inertia optimisation scheme over the robot's workspace, 1000 random, initially nonsingular configurations of the robot were minimised and

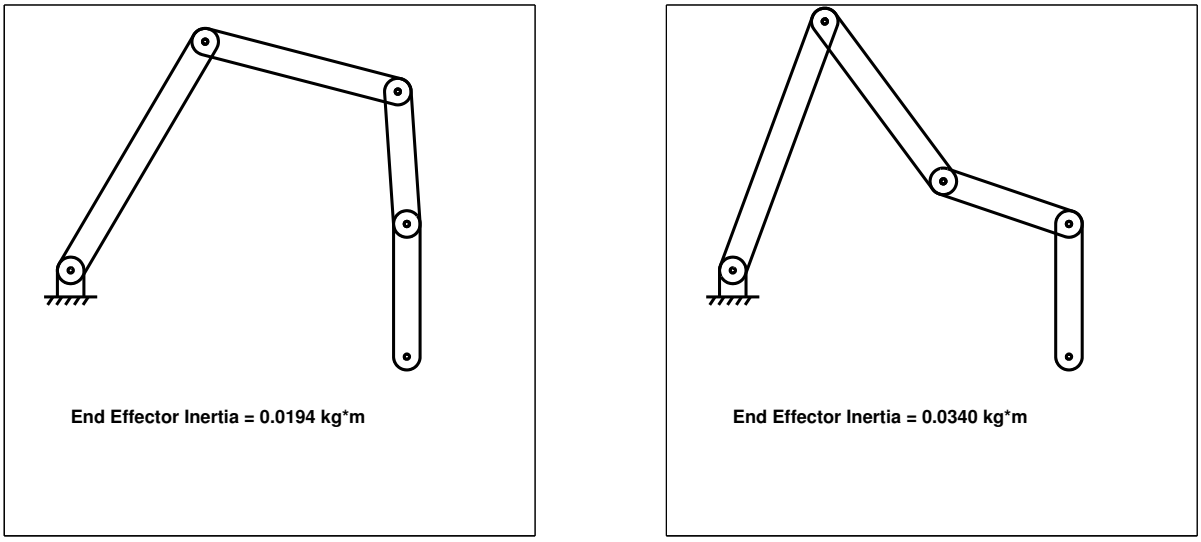

(a) Minimum Inertia Determinant (b) Maximum Inertia Determinant

Figure 3.4: The local minimum and maximum inertia joint configurations for a single end effector position and orientation.

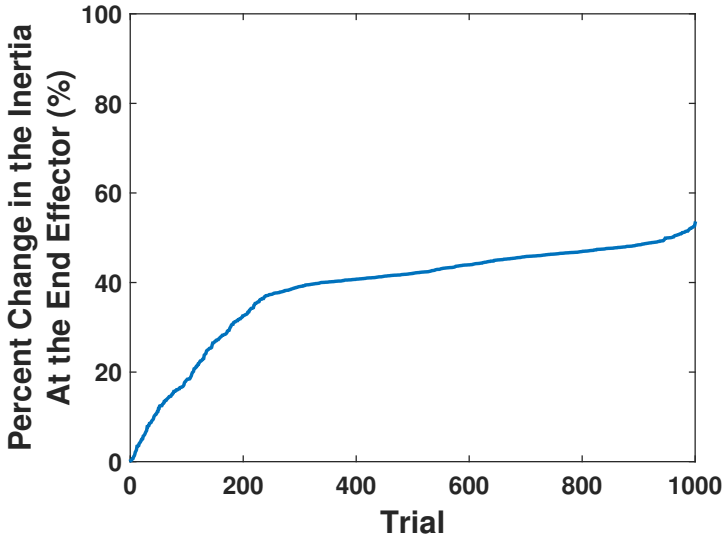

Figure 3.5: The percentage change in the inertia between the minimum and maximum joint configurations for randomly generated end effector poses.

maximised. Similarly to the previous section, initial joint configurations were computed by drawing from a uniform distribution over $[0, 2\pi)$ for each joint and checking for singularities by computing the singular value decomposition of JJ^T. A PD controller with gains $k_{p,ee} = 30$ and $k_{d,ee} = 5$ was applied to the end effector in order to avoid integration drift and ensure the initial position and orientation of the end effector was fixed. The robot was commanded with (3.18) to optimise (3.10) from the initial joint configuration and the parameters of (3.16) were set as $k_p = 50$, $k_d = 15$, and $\omega_{max} = 3rad/s$. Figure 3.4 depicts the minimal and maximal inertia configurations for a single example end effector pose.

Figure 3.5 depicts the percentage chance in the inertia at the end effector between the maximum and minimum joint poses. The minimum joint configuration has a 38.70% average reduction in the inertia as compared with the maximum inertia joint configuration, with a standard deviation of 11.57%. This represents a large reduction in the inertia the surgeon would feel in general when rotating the robot about the tool point.

3.5.3 Compatibility of Mass and Inertia Optimal Joint Configurations

Comparing the minimum and maximum joint configurations for mass (Figure 3.2) and for inertia (Figure 3.4), it appears that mass and inertia optimisation are not compatible for this specific robot, as the optimal mass configuration is close to the least optimal inertia pose and *vice-versa*. To further examine this issue, 1000 random end effector poses over the workspace of the robot were optimised for mass and inertia, using the same methodology as in the previous sections. Figure 3.6a depicts the resulting distance between the optimal mass and inertia joint poses for each end effector position and orientation, demonstrating a large difference in general between these configurations. Further, Figure 3.6b depicts the impact that optimising for mass has on the inertia at the tool. The inertia at the end effector at the mass optima is, in general, substantially larger than at the optimal inertia configuration.

However, even when the distance between optimal poses is large, a combined optimisation solution can still simultaneously reduce the mass and inertia of the robot at the end effector. 1000 random end effector poses were optimised using a 50%-50% mass-inertia ratio and the resulting mass and inertia were compared with the least optimal mass and inertia joint configurations. As

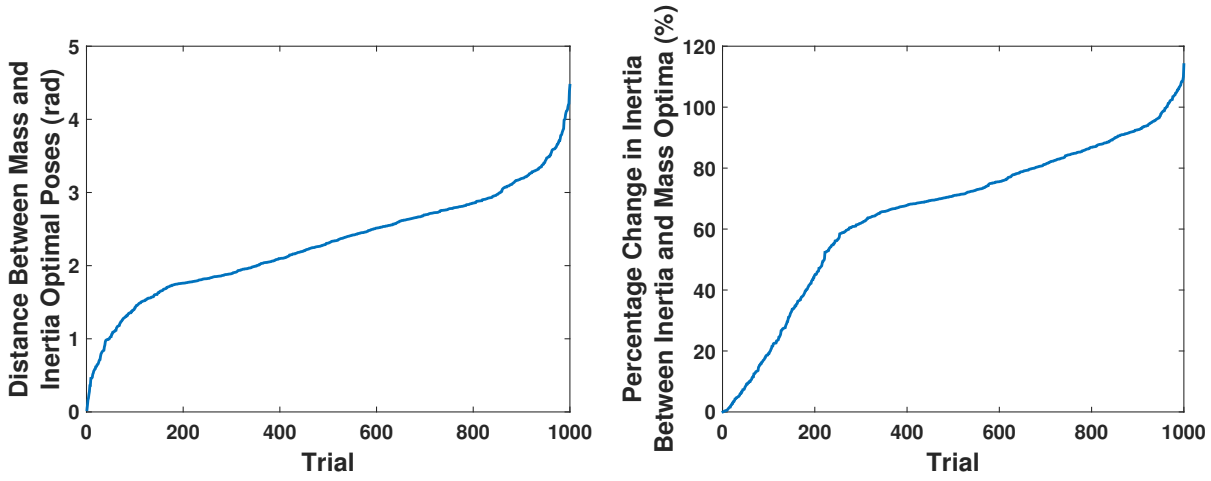

(a) (b)

Figure 3.6: The L^2 distance between the optimal mass and optimal inertia joint configurations (a) and the percentage change in the inertia between the optimal mass and optimal inertia poses (b) for 1000 trials.

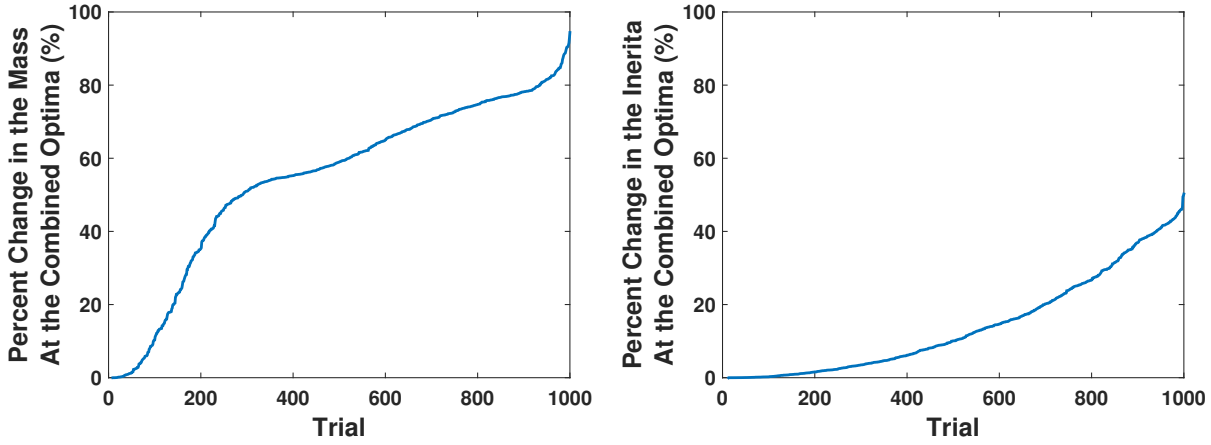

(a) (b)

Figure 3.7: The percentage improvement in mass (a) and inertia (b) at the combined optima relative to the least optimal mass and inertia poses.

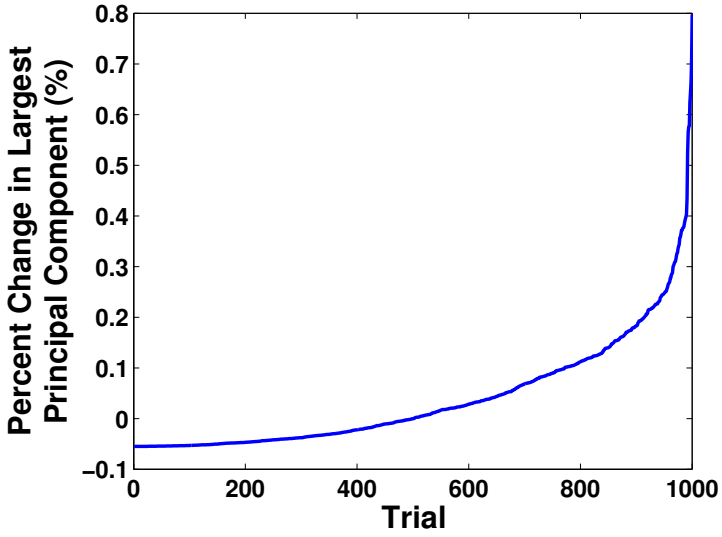

Figure 3.8: The percentage change in the largest principal mass component between an optimisation using the true dynamic parameters and an optimisation using normally distributed dynamic parameters.

shown in Figure 3.7, the combined optimisation, in general, reduced both the mass and inertia at the tool point, although not to the level seen in the single optimisations.

There are several ways in which the choice of using a single mass or inertia optimisation or using a combined strategy can be made. The decision can be made based on the most significant quantity affecting the robot, the quantity that affects the most common movement performed when using the robot, or based on the quantity which is most reduced by the null-space strategy in general. On the specific robot used in the simulations here, mass is the more significant quantity affecting the end effector motions and, in addition, is more substantially reduced by the proposed simulations. Further investigation of the combined optimisation effects will be performed on a robot capable of achieve general six dimensional end effector poses in Chapter 4.

3.5.4 Effectiveness Under Modeling Uncertainty

In a real system, the parameters of the dynamic model are only estimates of the true values. In order to test the impact of modeling uncertainty on the optimisation procedure, 1000 random trials of a constant end effector pose were performed, in which the mass and inertia properties were drawn from a normal distribution, with the mean equal to the true value of the parameter

and the standard deviation equal to 10% of its value ($m_{i,e} \sim \mathcal{N}(m_i, (0.1m_i)^2)$ and $I_{i,e} \sim \mathcal{N}(I_i, (0.1I_i)^2)$). Using the proposed control strategy (3.18), these trials optimised the effective mass at the end effector which was computed from the estimated parameters in order to compare the resulting joint configuration and actual effective mass with an optimisation using the true parameter values. A PD controller with gains $k_{p,ee} = 30$ and $k_{d,ee} = 5$ was applied to the end effector to eliminate integration drift and the parameters of (3.16) were set as $k_p = 50$, $k_d = 15$, and $\omega_{max} = 3rad/s$.

Figure 3.8 depicts the percentage difference between the optimisation results using the true and estimated mass and inertia parameters on the largest principal mass component, in ascending order. The results demonstrate that the change in the principal mass component while optimising using mass and inertia parameter estimates is, in the worst case, 0.7976% less optimal that using the true values. This suggests that this level of parameter variation does not have a large effect on the optimisation results. Note that a negative change in the largest principal component does not reflect a more optimal pose, as optimality is defined as the product of the principal components.

3.5.5 Effectiveness of Directional Mass Optimisation Over the Robot's Workspace

To investigate the effectiveness of the proposed directional mass optimisation strategy over the workspace of the robot, the directional mass optimality measure (3.13) was minimised and maximised using (3.18) for 1000 random, initially nonsingular configurations. Initial joint configurations were found by drawing from a uniform distribution over $[0, 2\pi)$ for each joint and checking for singularities by computing the singular value decomposition of JJ^T. The optimisations were performed for end effector positions and orientations fixed to the initial pose and the direction optimised for all cases was arbitrarily chosen to be the negative vertical direction. The choice of the optimisation direction does not affect the overall results as the orientation angle is random due to the way in which the joint angles are chosen, preventing any biasing with respect directions which exhibit larger changes in mass. A PD controller with gains $k_{p,ee} = 30$ and $k_{d,ee} = 5$ was applied to the end effector to eliminate integration drift and the parameters of (3.16) were set as $k_p = 50$, $k_d = 15$, and $\omega_{max} = 3rad/s$.

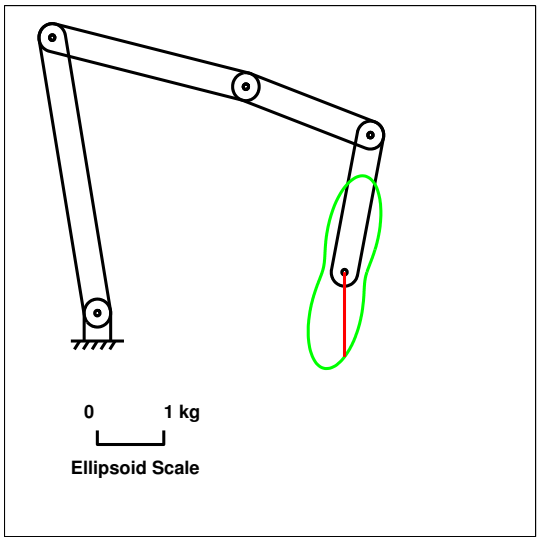

(a) Minimum Effective Mass in the Vertical Direction

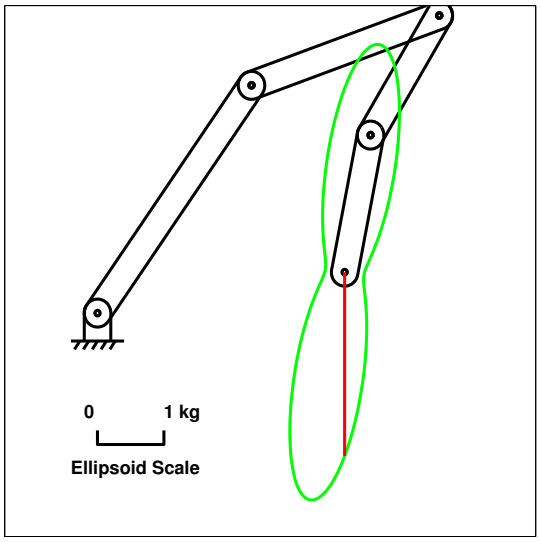

(b) Maximum Effective Mass in the Vertical Direction

Figure 3.9: The local minimum and maximum effective mass in the vertical direction and belted mass ellipsoids for a single end effector position and orientation overlaid on their corresponding joint configurations.

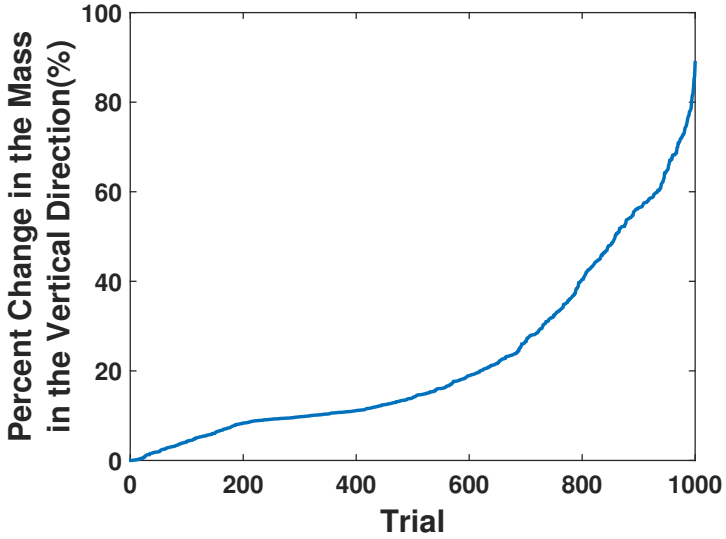

Figure 3.10: The percentage change in the effective mass in the vertical direction between the minimum and maximum joint configurations for randomly generated end effector poses.

Such an optimisation could be used to ensure that the user's interaction with the robot feels "lightweight" when moving up and down. This would help reduce fatigue and improve ease of use when moving towards a target, or if motion in the vertical direction is much more frequent than any other in a particular task. Figure 3.9a shows the configuration with locally optimal posture for minimising the mass in the specified direction. Additionally, Figure 3.9b depicts the configuration for the maximum effective mass in the negative y-direction.

Figure 3.10 depicts the percentage change in the effective mass in the vertical direction between the minimum and maximum directional mass optimisation metrics for each of the random trials, arranged in ascending order. The minimum configuration on this specific robot represents, on average, a 22.64% reduction in the effective mass relative to the maximum configurations with a standard deviation of 19.91%. The reduction in effective mass over the workspace of this specific robot is not as substantial as the overall mass metric results, however, does demonstrate a reduction in effective mass for most poses. Figure 3.11 illustrates a case in which the reduction in mass is not as substantial. In configurations for which the desired direction to be optimised aligns with the rotation axis of the last joint such that an instantaneous translation in the specified direction can be performed only by a rotation of the final joint, a substantial change in the mass properties is not possible because the effects of the last joint cannot be reduced. Therefore, this strategy is most effective in configurations where moving the robot in the specified direction requires multiple links of the robot to move in coordination. This reduction in mass may be more substantial for more poses on a robot capable of six dimensional end effector poses in space and will be investigated in Chapter 4. Directional inertia optimisation was not performed in this section since the planar nature of the robot allows for only a single rotation direction.

3.5.6 Simulation Trajectory Trials

In order to compare the performance of these measures against current surgical methodologies for redundancy resolution, a test scenario was created. A translational trajectory task requiring acceleration and deceleration was designed to test the effects of mass optimisation and its effect on motion due to external forces from the surgeon as mass is more substantially reduced by

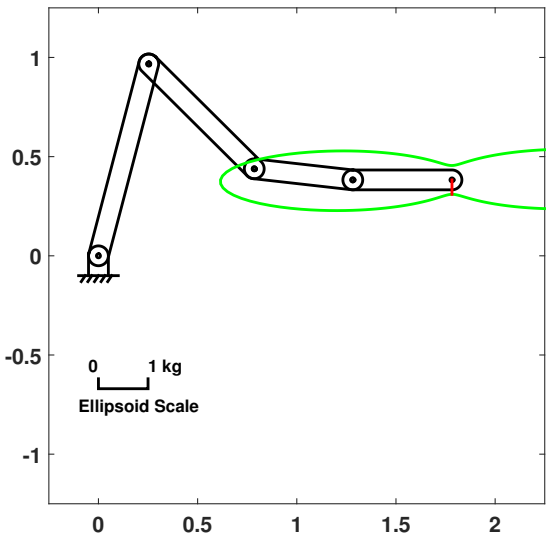

(a) Minimum Effective Mass in the Vertical Direction

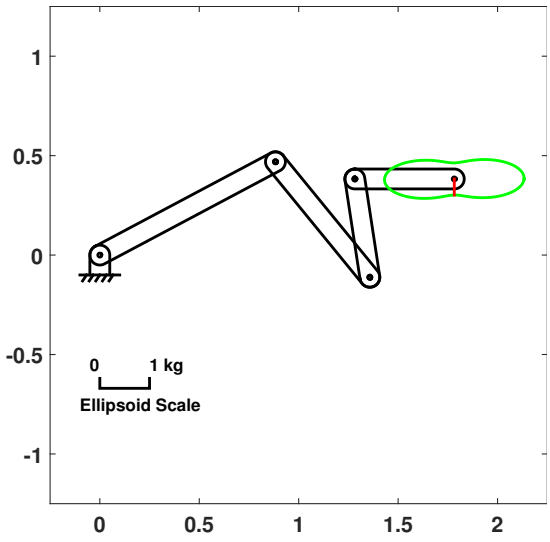

(b) Maximum Effective Mass in the Vertical Direction

Figure 3.11: The local minimum and maximum effective mass in the vertical direction and belted mass ellipsoids for a single end effector position and orientation demonstrating the limited reduction in mass when optimising in a direction corresponding to the rotation axis of the last joint.

our proposed methods as discussed in Section 3.5.3. Figure 3.12 shows the trajectory that the robot followed in each trial. The corner points were used as references for the robot with the desired target switched when the robot reached the current point. To simulate the surgeon following the trajectory, external forces were applied to the system using a force saturated PD controller.

$$\begin{aligned} v &= k_p(x_{des} - x) - k_d\dot{x} \\ F &= \min(v, F_{max}) * \frac{v}{\|v\|} \end{aligned} \tag{3.20}$$

where F_{max} is the maximum force allowable and k_p and k_d are the position and damping gains respectively. For the trajectory trials presented here, the gains were set empirically to allow for minimal overshoot when combined with the robot controllers as follows: $k_p = 30$, $k_d = 10$, and $F_{max} = 1N$.

Four control methodologies were tested for this pattern; active local mass determinant minimisation, initial minimum posture, “elbow up” posture, and damped posture. The gains of the active mass minimisation controller were empirically set to $k_p = 50$, $k_d = 15$, and $\omega_{max} = 3rad/s$

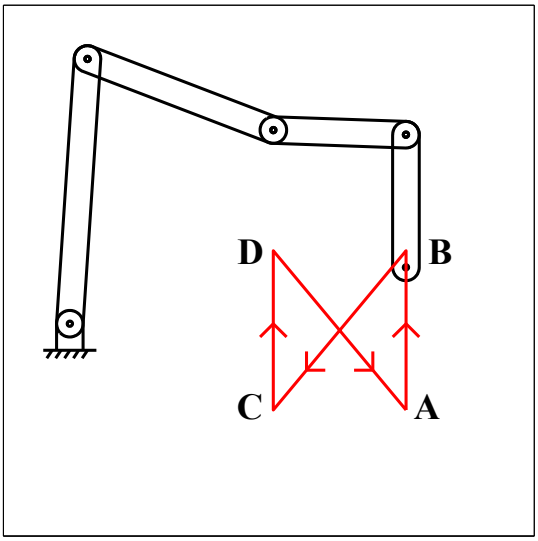

Figure 3.12: The test trajectory used for comparing the performance of the proposed null-space based optimisation strategy with alternative redundancy resolution strategies.

such that the null-space exhibited approximately critical damped response to the optimal posture.

The initial minimum pre-optimised posture case controlled the redundancy to the position corresponding to the minimum mass determinant of the starting end effector pose using a PD controller with gains who were empirically chosen to exhibit close to critically damped response: $k_{p,pos} = 30$ and $k_{d,pos} = 15$. The null-space strategy in this case minimises the sum of squared errors to the desired joint configuration when the end effector pose is not consistent with the chosen set of joint angles. This strategy was included to compare a pre-optimised configuration with active optimisation during motion, in order to determine if such a control technique could provide natural motion in a more constrained surgical theater.

In the "elbow up" posture controller, the robot's redundancy was controlled to a constant "elbow up" position using a PD controller with gains who were empirically chosen to approximate a critically damped response: $k_{p,pos} = 30$ and $k_{d,pos} = 15$. This strategy is generally used to keep the robot in a position which avoids the general work area. In the scenario here and in general for many surgical applications, the main body was controlled in the redundancy to be above the end effector in order to stay out of the way of the motion.

In the damped posture controller, the joint configuration of the robot in the redundancy was

controlled to have dynamic compensation and damping with the damping gain empirically set to $k_{d,pos} = 15$ to allow for the degree of null-space motion commonly seen in this type of control. This control mode is commonly used in redundant systems to allow for reconfiguration by the user by hand. The redundancy remains in approximately the same position when moving the end effector and, by applying a force on the redundancy, the user can reconfigure the robot manually for the current task.

Control Strategy	**Time (s)**
Determinant Minimisation	5.32
Initial Min Posture	6.16
Elbow Up Posture	8.26
Damped Posture	10.30

Table 3.2: The time taken for each control strategy to perform the desired trajectory.

When applying the same force to the end effector, the resulting acceleration will be a function of the mass of the robot at the tool point. Therefore, the time the robot takes to complete the trajectory is indicative of the mass of the robot throughout the trajectory. The time taken by each strategy to perform the trials is shown in Table (3.2).

3.6 Discussion

While the "feel" of the response of the robot to the surgeon depends on several factors, the focus of the work in this chapter is on the apparent mass/inertia at the end effector. The results presented are limited by using a 2D robot and a simulated surgeon however, they demonstrate the impact of mass and inertia on the performance of the robot. User trials on a robot capable of achieving 6D poses at the end effector will be demonstrated in Chapter 4. Section 3.5.1 and 3.5.2 demonstrate how optimising the mass and inertia, on average, makes a difference over the workspace, i.e. it can produce a substantial change in the apparent mass/inertia during a hands-on interaction. Additionally, although limited to a fixed pose, the approach is shown to still be effective in the presence of model uncertainty (Figure 3.8).

Section 3.5.6 shows that methods which consider the mass of the device are able to follow the trajectory in less time than the "elbow up" posture and damped posture cases. Figure 3.13 shows why this is the case by plotting the maximal principal component of the belted

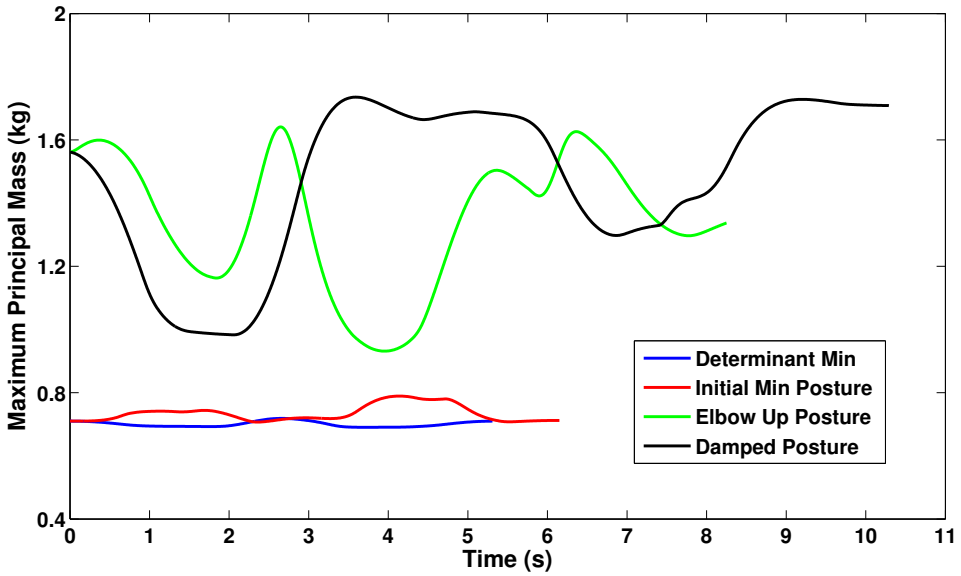

Figure 3.13: The principal mass component over time for each controller during the trajectory trials.

mass ellipsoid over time for each control strategy. The "elbow up" posture and damped posture control strategies show a consistently higher mass throughout the simulations, causing the robot to be more difficult to accelerate/decelerate and thus guide along the course. Therefore, the optimisation strategy provides a more natural motion for the user.

In the case of the damped posture presented here, the maximum mass component varies by up to 1.5 times its smallest value, as shown in Figure 3.13 with corresponding joint poses and mass ellipsoids shown in 3.14. From the viewpoint of the user, this would require a different control strategy to achieve good performance in all areas of the workspace. By removing the need for this motor skill, the surgeon can focus more on the surgery at hand and less on determining the right amount of force to achieve the desired motion.

In some surgeries, the workspace is heavily constrained and using an active optimisation strategy may not be possible. The pre-optimised constant posture strategy demonstrates that many of the benefits of considering mass can still be gained without using the active solution. In these situations, the optimal joint configuration for the area where the surgery will take place can be found pre-operatively for constant redundancy position control and the surgical equipment can be arranged around the already optimised robot.

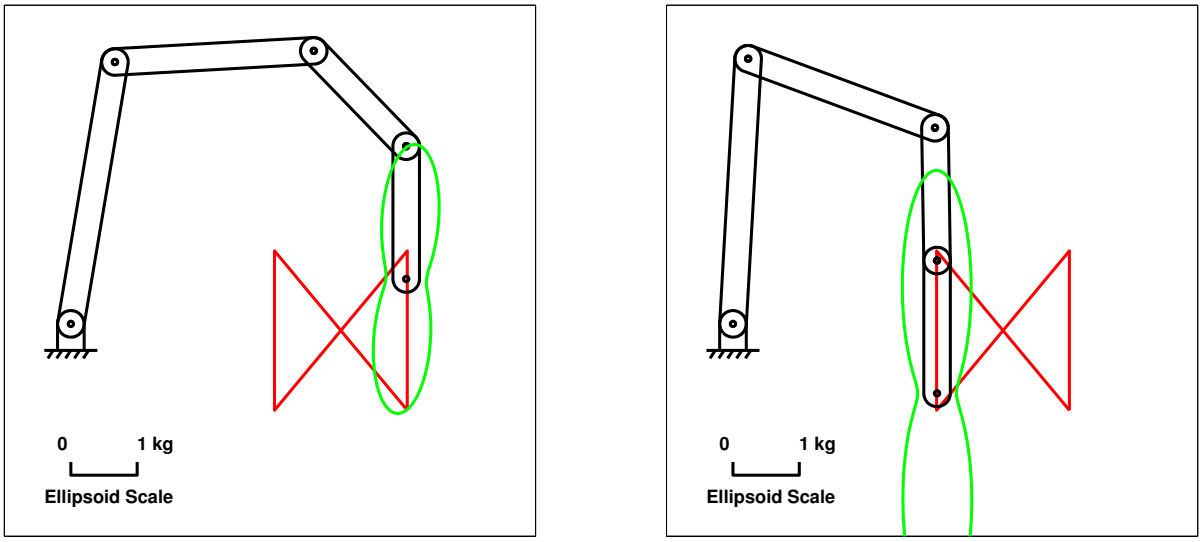

Figure 3.14: A comparison of the mass ellipsoid at two times during the trajectory while using the damped posture controller demonstrating the large change in the mass properties which occurs when using this controller.

3.7 Summary

This chapter has demonstrated a null-space based optimisation strategy which reduces the mass and inertia at the end effector of the robot in order to create a more natural motion, closer to the tool unattached to the robot. By optimising in the null-space of the tool pose, the desired commands of the surgeon at the tool point are unaffected and safety concerns regarding the stability and force amplification in impedance controllers applied to the end effector is avoided. By providing a control strategy rather than a design solution, the approach presented here can be applied to current large scale hands-on robotic surgery devices without the need for a costly and lengthy system redesign.

In this chapter, the dynamics of the robot at the end effector have been derived and measures to quantify the overall and directional tool point mass and inertia have been developed. Weighted solutions to simultaneous mass and inertia optimisations have been discussed and a redundant optimisation strategy has been demonstrated, which minimises the measures in the null-space of the tool pose such that the optimisation does not affect the surgeon's desired commands at the end effector. Simulation work on a four degree-of-freedom planar robot has demonstrated the reduction in mass and inertia which the optimisation strategy provides over the workspace of

the robot, and an investigation into the effect of uncertainty in the parameters of the dynamic model has revealed that a reasonably small amount of parameter variation does not substantially affect the optimisation results. Combined optimisation simulation results demonstrate that simultaneous optimisation of mass and inertia can reduce both these quantities even when their individual optimal poses are conflicting. Finally, simulated trajectory trial results show a reduction in the time required to complete a trajectory when using the optimisation controller, as compared to conventional redundancy resolution controllers, damped posture and elbow up.

In the next chapter, frictional force and torque optimisation measures will be developed in order to reduce the frictional effects at the end effector, on the surgeon. In addition, combined mass and friction optimisation will be presented. Simulation results on a seven degree-of-freedom system will show the efficacy of the proposed methods in reducing the end effector dynamics over the workspace of the robot. Lastly, experimental results on the same robot used in simulation will demonstrate a reduction in the work required to move the device in point-to-point motions, as compared to conventional surgical redundancy resolution schemes, creating a more natural motion for the surgeon.

Chapter 4

Redundant Mass and Friction Optimisation for Natural Motion

This chapter presents a novel hands-on control methodology for optimising the friction force and torque at the end effector of a redundant robot without affecting the pose of the tool. The efficacy of the overall and directional mass and friction optimisation methodologies over the workspace of a robot is demonstrated through simulation of a lightweight seven degree-of-freedom robot. Additionally, these metrics are shown in combined optimisations with mass to simultaneously reduce both of these quantities at the end effector. Lastly, experimental user trials validate the proposed strategies by demonstrating a reduction in the work required to move the tool, thereby providing a more natural motion for the surgeon.The research presented here is an edited version of the work under review in:

> J. G. Petersen and F. Rodriguez y Baena, "Mass and friction optimisation for natural motion in hands-on robotic surgery", IEEE Transactions on Robotics, Under Review.

4.1 Introduction

Natural motion intends to reduce the effects of the robot dynamics at the end effector on the surgeon which result from attaching the surgical tool to the device. By providing an experience which is closer to the tool unattached to the robot, the surgeon's ease of use can be improved

and fatigue from interacting with these dynamics over long surgeries can be lessened. While the previous chapter addressed the impact of the mass and inertia of the robot at the end effector on the surgeon, other dynamic effects of the robot at the tool point can reduce the ability of the surgeon to maneuver the surgical instrument. This chapter aims to reduce the effect of the friction forces and torques at the end effector which result from the joint level friction in order to provide a more natural motion for the surgeon.

There exist several models which aim to accurately capture the nonlinear and complex effects of friction [Olsson et al., 1998] [Bona and Indri, 2005]. Recent models for friction based on the work of [Dahl, 1968] have addressed the Stribeck effect [Wit et al., 1995] and incorporated elastoplastic effects into the presliding region [Dupont et al., 2002] [Hayward et al., 2009].

However, performing a null-space optimisation using a model-based estimate of the current friction on the robot's joints can result in motions which are undesirable for hands-on robotic surgery. As opposed to the effective mass and inertia at the end effector which depend only on the joint angles and the direction of motion, the friction at the tool point exhibits an additional dependence on the velocity of the tool, outside of the static friction regime [Olsson et al., 1998] [Bona and Indri, 2005]. As a result, changes in the direction and magnitude of the velocity of the end effector cause the direction and magnitude of the friction forces and torques at the tool point to change. Therefore, in an optimisation scheme which utilised an accurate model-based estimate of the joint friction, quick changes in the velocity of the tool would result in rapid reconfiguration of the redundant degrees-of-freedom in order to maintain an optimal pose with respect to these changing values. Clinical application of such a strategy would likely be difficult, as such rapid changes would be disconcerting for the surgeon. Additionally, these changes in direction and corresponding reconfigurations may cause the system to spend more time in non-optimal poses as it transitions between the optimal joint configurations, lessening the reduction in friction which the controller would have.

Such a control strategy would also only optimise when the robot was in motion, as the frictional forces and torques at the end effector are zero when the robot is stopped. An optimisation technique which could continue to reconfigure the redundancy to ensure low friction for the next motion of the tool would be able to reduce the frictional effects of that motion.

As the surgeon is in control of the motion of the tool and, hence, the speed and direction of the surgical procedure, the motion of the tool is unknown to the robot controller. Research has been performed into understanding and predicting the surgeon's intent, however, this research has been limited to general task descriptions (e.g. place or position) [Hundtofte et al., 2002] [Li and Okamura, 2003] or has focused on high level workflow decisions [Rosen et al., 2006] [Holden et al., 2014] due the complexity of such a task. Alternatively, by providing a joint configuration which reduces the overall frictional forces and torques at the end effector, these effects could be lessened regardless of the subsequently surgical instrument motion.

As a result of these concerns, optimality measures for the friction force and torque at the tool point have been developed, which are independent of the velocity of the robot's joints, in order to prevent rapid reconfigurations of the redundancy to accommodate changes in the frictional forces and torques resulting from quick tool velocity changes. In addition, this methodology configures the robot to ensure a general improvement in these quantities regardless of the direction of motion as the intended motion of the surgical instrument is unknown to the controller.

First, this chapter will develop joint velocity independent optimality metrics for the frictional force and torque at the end effector using the novel concept of the "potential velocity" of the end effector, which is used to create a theoretical friction force and torque at the tool point for which to optimise. Simulation results on a seven degree-of-freedom lightweight robot will demonstrate the effectiveness of the overall and directional mass and friction optimisations over the robot's workspace. Next, the trade-offs when performing a combined mass and friction optimisation are examined and the impact of modeling uncertainty on end effector friction is investigated. Finally, the results of an experimental point-to-point user trial demonstrate a reduction in the work required to move the tool, as compared with standard surgical redundancy resolution techniques, indicating a more natural motion when using the proposed optimisation strategies.

4.2 End Effector Friction Force and Torque Optimisation

When the surgeon applies forces and torques to the end effector, he or she not only feels the mass and inertia of the manipulator, but also the joint friction torques. As shown previously in Section 3.2, the joint friction torques can be projected onto the end effector using the dynamically consistent inverses of the Jacobian (3.3) to find the effective friction force and torque which the surgeon feels at the tool point.

$$F_f^{EE} \triangleq \overline{J}_v^T \tau_f \text{ and } \tau_f^{EE} \triangleq \overline{J}_\omega^T \tau_f \tag{4.1}$$

where τ_f is the vector of joint friction torques, $\overline{J}_v^T$ and $\overline{J}_\omega^T$ are the dynamically consistent inverses of the linear and angular Jacobians, respectively, and F_m^{EE} and τ_m^{EE} are the vectors of forces and torques, respectively, at the end effector resulting from the motor torques.

4.2.1 Directional End Effector Friction Force and Torque Optimality Measures

Friction models such as [Wit et al., 1995] and [Dupont et al., 2002] aim to accurately model the friction forces present at the joint by incorporating complex effects such as stiction and the Stribeck effect. These models typically exhibit a dependence on velocity outside of the static friction regime [Olsson et al., 1998] [Bona and Indri, 2005].

However, using an accurate real-time estimation of the current joint friction to optimise the pose of the robot in hands-on robotic surgery can lead to behavior that is undesired. In such a set-up, rapid changes in the velocity of the end effector may cause the direction and magnitude of the frictional forces and torques to quickly change and, as a result, cause the system to reconfigure rapidly to optimise these changing values. Such rapid changes are likely to be disconcerting for the surgeon and dissuade clinical usage. Secondly, with an accurate friction estimate, if the surgeon was not moving the tool, the robot would no longer change configuration. A strategy which is able to continuously optimise when the robot is at rest would be able to reduce the frictional effects for the next motion. Lastly, I assume that I do not know the surgeon's intended

motion. Therefore, a general optimisation strategy, which can reduce the friction that would occur regardless of the direction of the next motion, would be advantageous.

For these reasons, I have developed a methodology which sacrifices the precision of the true frictional forces to allow for values which are more consistent with the joint pose rather than the velocity in order to simplify the optimisation and to avoid the aforementioned issues. To do this, I propose the concept of the "potential velocity" of the end effector, which is to be used in combination with a standard joint friction model. At the current pose, the joint velocities required for the end effector to move at a particular velocity at a specified pose can be computed using the dynamically consistent inverse of the linear and angular Jacobians.

$$\dot{q}_v^P \triangleq \overline{J}_v \dot{x}_p \text{ and } \dot{q}_\omega^P \triangleq \overline{J}_\omega \omega_p \tag{4.2}$$

$\dot{q}_v^P$ and $\dot{q}_\omega^P$ are defined as the vectors of potential joint velocities for a potential end effector linear velocity, $\dot{x}_p$, or angular velocity, ω_p.

The potential joint velocities for a given direction and speed can be used to compute the potential frictional force and torque at the end effector, $F_f^{EE}(\dot{q}_v^P)$ and $\tau_f^{EE}(\dot{q}_\omega^P)$. The potential friction here does not represent the true friction occurring at the end effector at the current time. Rather, it represents what the friction would be if the device were moving at the potential velocity at the current pose. In this way, the friction being optimised is more closely tied to the current pose, as opposed to the current velocity. To minimise the potential frictional force or torque for a given direction and speed, it suffices to minimise the squared magnitude of the friction vector.

$$c_f^d = \|F_f^{EE}(\dot{q}_v^P)\|^2 \text{ and } c_\tau^d = \|\tau_f^{EE}(\dot{q}_\omega^P)\|^2 \tag{4.3}$$

By optimising friction when moving at the potential joint velocities, the robot will continue to optimise even at rest and avoid rapid reconfiguration due to changes in tool motion direction and speed.

4.2.2 End Effector Friction Force and Torque Optimality Measures

To achieve general optimisation with respect to frictional force and torque when moving the robot in any direction, I utilise a summation of the directional optimisation measure in the orthogonal Cartesian directions. The optimality measure associated with frictional force can be expressed as,

$$c_f = \sum_{i=1}^{3} \|F_f^{EE}(\dot{q}_v^{P,i})\|^2 \tag{4.4}$$

where $\dot{q}_v^{P,1} = \overline{J}_v\dot{x}$, $\dot{q}_v^{P,2} = \overline{J}_v\dot{y}$, and $\dot{q}_v^{P,3} = \overline{J}_v\dot{z}$ are vectors of potential joint velocities associated with unit linear velocities in the three Cartesian directions, $\dot{x}$, $\dot{y}$, and $\dot{z}$ using (4.2).

Similarly, the optimality measure associated with frictional torques can be expressed as,

$$c_\tau = \sum_{i=1}^{3} \|\tau_f^{EE}(\dot{q}_\omega^{P,i})\|^2 \tag{4.5}$$

where $\dot{q}_\omega^{P,1} = \overline{J}_\omega\omega_x$, $\dot{q}_\omega^{P,2} = \overline{J}_\omega\omega_y$, and $\dot{q}_\omega^{P,3} = \overline{J}_\omega\omega_z$ are the potential joint velocities associated with unit angular velocities about the three Cartesian directions, ω_x, ω_y, and ω_z using (4.2).

Multiple end effector velocity directions and combined optimisations of frictional forces and torques can be accomplished using a weighted sum. Similarly to (3.11) and (3.14), a weighted sum is required for combined optimisations, since forces and torques can not be directly added.

$$c_f^c = w_f \sum_{i=1}^{n} \|F_f^{EE}(\dot{q}_v^{P,i})\|^2 + w_\tau \sum_{j=1}^{m} \|\tau_f^{EE}(\dot{q}_\omega^{P,j})\|^2 \tag{4.6}$$

4.2.3 Combined Optimisation

Finally, the optimality measures developed above can be combined into a single metric, in which each measure must be weighted according to, for example, relative importance or expected

effect,

$$c = w_m c_m + w_i c_i + w_f c_f + w_\tau c_\tau \tag{4.7}$$

where w_m, w_i, w_f, and w_τ are the weights for the mass, inertia, frictional force, and frictional torque metric respectively.

4.2.4 Kuka LWR 4+ Friction Model

The robot I will be using for testing these measures in simulation and in an experimental setup will be the Kuka LWR 4+ (Kuka Robotics GmbH), which has seven degrees-of-freedom and is able to be controlled in impedance mode. To perform friction optimisation, an appropriate joint friction model for the robot must be chosen.

The LWR 4+ is intended for compliant assembly in industrial tasks and, as such, aims for a higher payload ($7kg$) than traditional impedance systems. To accomplish this, the system is inherently non-back-drivable and, instead, provides a virtual impedance through modeling. This, however, results in a limitation of the system whereby each joint on the Kuka in impedance mode requires $1.5Nm$ of torque before it will move. This limitation can be viewed as a virtual friction which I can use in optimisation.

To prevent discontinuities in optimising, a smooth approximation to the friction present in the LWR is used,

$$\tau_f = \tau_{max} \left(\frac{2}{\pi} \right) \arctan\left(\beta \dot{q}\right) \tag{4.8}$$

where τ_{max} is the maximum frictional torque and β is a parameter which affects the slope of the friction torque near zero. As β tends to infinity, τ_f approaches $\tau_{max}\text{sgn}(\dot{q})$.

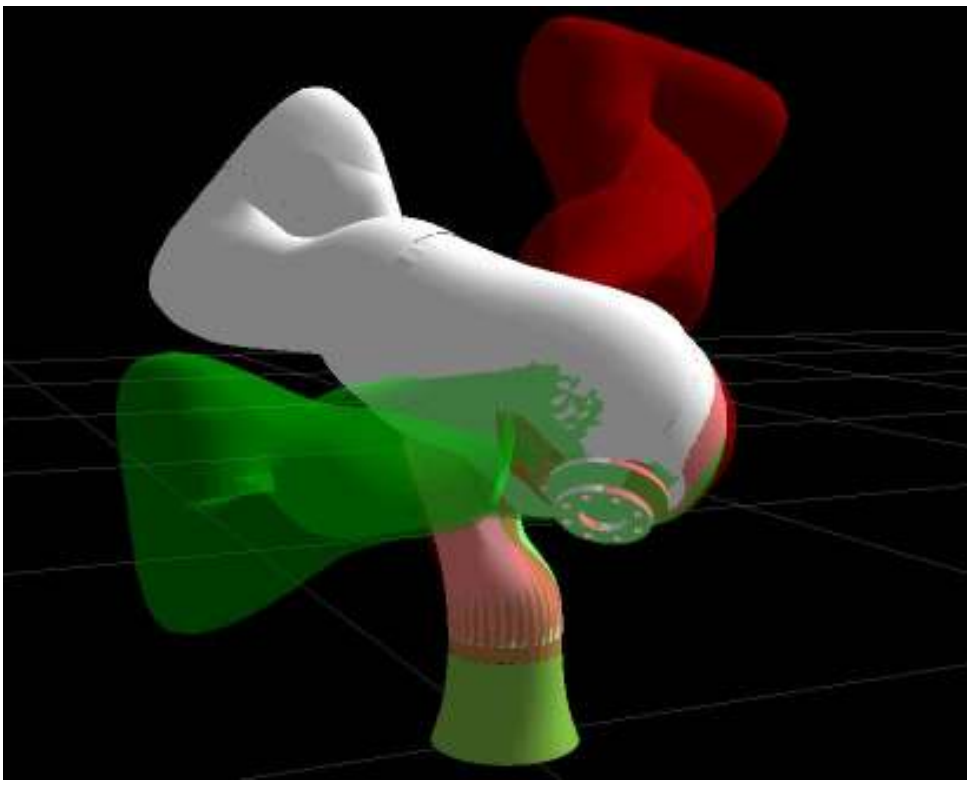

Figure 4.1: A visualisation of the local overall minimum mass pose (green) and the local overall maximum mass pose (red) overlaid on the initial joint posture (white).

4.3 Simulation Results and Discussion

Simulations were performed on a PC (i7-2600 @ 3.4 GHz) running Ubuntu 12.04, implemented using the ROS (www.ros.org) and OROCOS (www.orocos.org) frameworks. The kinematics and dynamics of the robot were provided in a library by Kuka and these equations of motion of the robot were integrated using the Dormand Prince 5 method [Dormand and Prince, 1980] provided by the odeint library in boost::numeric (www.boost.org).

4.3.1 Effectiveness of Overall Optimisations Over the Robot's Workspace

To demonstrate the effectiveness of considering mass, inertia, friction force, and friction torque over the workspace of the LWR, I aim to show that, for a large sample of arbitrary, non-singular, initial joint angles and corresponding end effector poses, the difference in these quantities between the local minimum posture and the local maximum posture is substantial in a majority of cases. The local minimum and maximum postures are defined as the sets of joint angles

corresponding to the local minimum and maximum, respectively, of the optimisation measure, relative to the initial pose. Figure 4.1 depicts an example of the local minimum and maximum postures overlaid on the initial joint posture.

1000 random, non-singular, initial joint configurations were found by drawing from a uniform distribution over the robot's joint limits for each link and testing for singularities by computing the singular value decomposition of JJ^T. Keeping the initial end effector position and orientation constant, the robot was then optimised using (3.18) from the initial joint angles with the parameters of (3.16) set to $k_p = 50$, $k_d = 10$, and $\omega_{max} = 10 rad/s$. c was set to the respective version of (3.10) for mass and inertia and the respective version of (4.4) for friction force and torque. As (3.16) is used to minimise the optimisation criterion, a modified version was used to obtain the local maximum:

$$\begin{aligned} \omega_{des} &= \frac{k_p \nabla c}{k_d} \\ v_d &= \min\left(1, \frac{\omega_{max}}{|\omega_{des}|}\right) \\ F_p &= -k_d(\dot{q} - v_d \omega_{des}) \end{aligned} \tag{4.9}$$

Figure 4.2 depicts the resulting percentage change in the optimisation quantities across all the trials between the local minimum and maximum configurations, arranged in ascending order. Additionally, Table 4.1 summarises the means and standard deviations of the percentage improvements that the local minimum configurations represent, relative to the maximum configurations for each criteria over the 1000 trials. For each of the optimisation measures, the minimum configurations, on average, represent a substantial reduction in their respective optimisation quantity, relative to the maximum configurations. Therefore, the overall mass, inertia, friction force, and friction torque are reduced over the workspace of the robot using our proposed methods and, as a result, a more natural motion will occur when using these techniques.

4.3.2 Effectiveness of Directional Optimisations Over the Robot's Workspace

Similarly, the effects of the directional mass, inertia, friction force and friction torque optimisation over the workspace of the robot were examined by optimising on a set of random joint

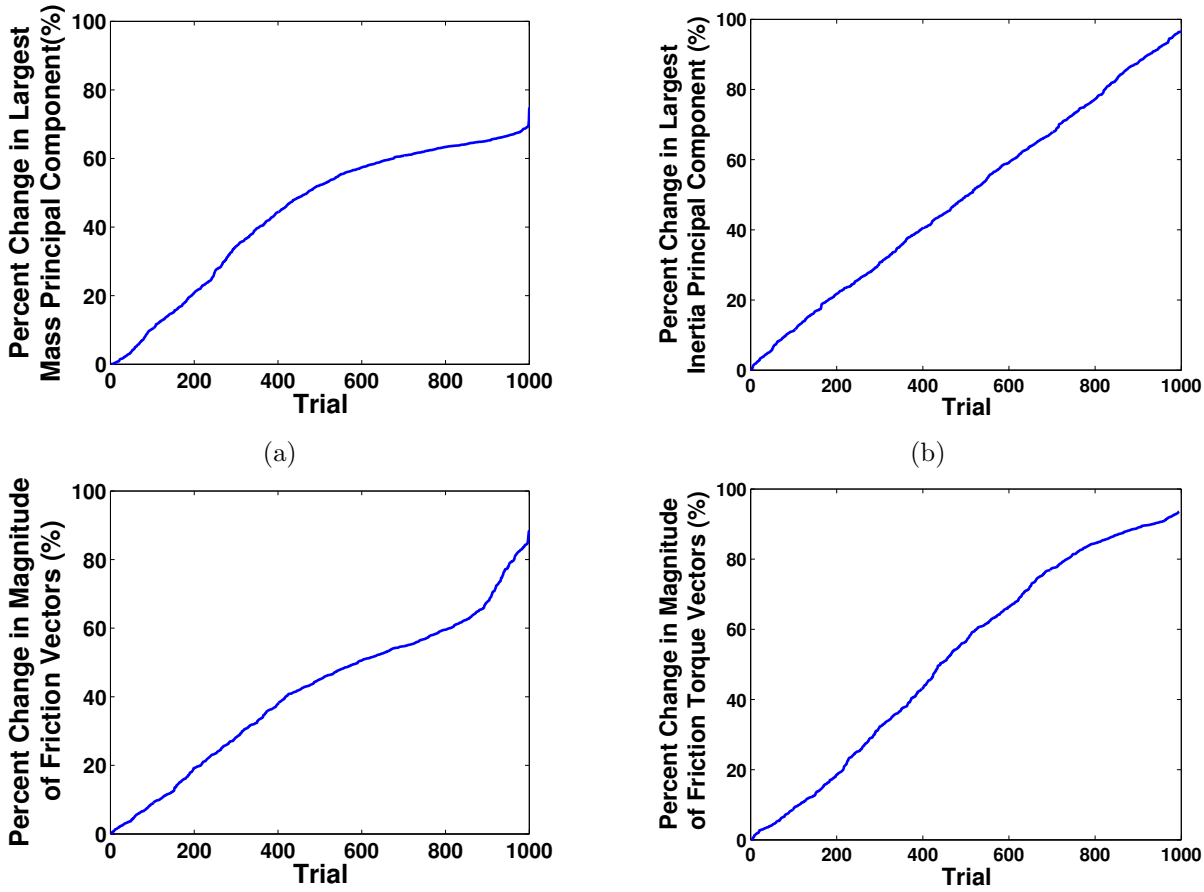

Figure 4.2: The percent change in the local minimum and maximum mass postures (a), the local minimum and maximum inertia postures (b), local minimum and maximum friction force postures (c), and the local minimum and maximum friction torque postures (d) for 1000 random poses.

	Mass	**Inertia**	**Friction Force**	**Friction Torque**
Mean (%)	44.23	49.99	41.44	52.26
Standard Deviation (%)	20.97	27.67	21.93	29.67

Table 4.1: Mean and standard deviations of the percentage reduction at the optimal pose relative to the maximum pose for the overall mass, inertia, friction force, and friction torque optimisations.

configurations, sampled over the robot's workspace. 1000 random, non-singular, initial joint configurations were found by drawing from a uniform distribution over the robot's joint limits for each link and testing for singularities by computing the singular value decomposition of JJ^T. The robot was optimised in the downward direction using (3.18) from the initial joint angles, while keeping the initial tool pose constant with the parameters of (3.16) set to $k_p = 50$, $k_d = 10$, and $\omega_{max} = 10rad/s$. c was set to the respective versions of (3.13) for mass and inertia and the respective versions of (4.3) for friction force and torque. Similarly to the previous section, a modified version of (3.16) was used to find the local maximum (4.9).

The results of these simulations can be found in Figure 4.3, which shows the percentage change in the effective optimisation criteria across all the trials between the local minimum and maximum configurations, arranged in ascending order. The means and standard deviations over the 1000 trials for each metric of the percentage improvements that the local minimum configurations represent relative to the maximum configurations are shown in Table 4.2. The effective mass and friction force in the negative vertical direction and the effective friction torque about the downward direction are substantially reduced on average across the trials at the local minimum poses, as compared to the local maximum poses. These quantities are lessened over the workspace of the robot, demonstrating the efficacy of our methods in creating a more natural motion in the preferred direction.

There is a less substantial change in the directional inertia about the vertical direction, on average, than in the other quantities for the Kuka LWR, however, the optimisation does reduce the effective inertia for most poses. Figure 4.4 depicts a case in which there is very little change in the inertia about negative vertical direction at the tool point by overlaying the optimal pose and least optimal pose on top of the initial joint configuration. While the optimal and least optimal joint poses are far apart, there is less than a 1% change in the directional inertia for this end effector pose.

This reduced impact is due to the rotation axis of the final joint aligning with the desired optimisation axis. The degree to which this alignment occurs affects the amount to which the effective inertia can be reduced, since this joint must always be rotated regardless of the null-space posture, when it is not orthogonal to the desired rotation axis. The null-space based

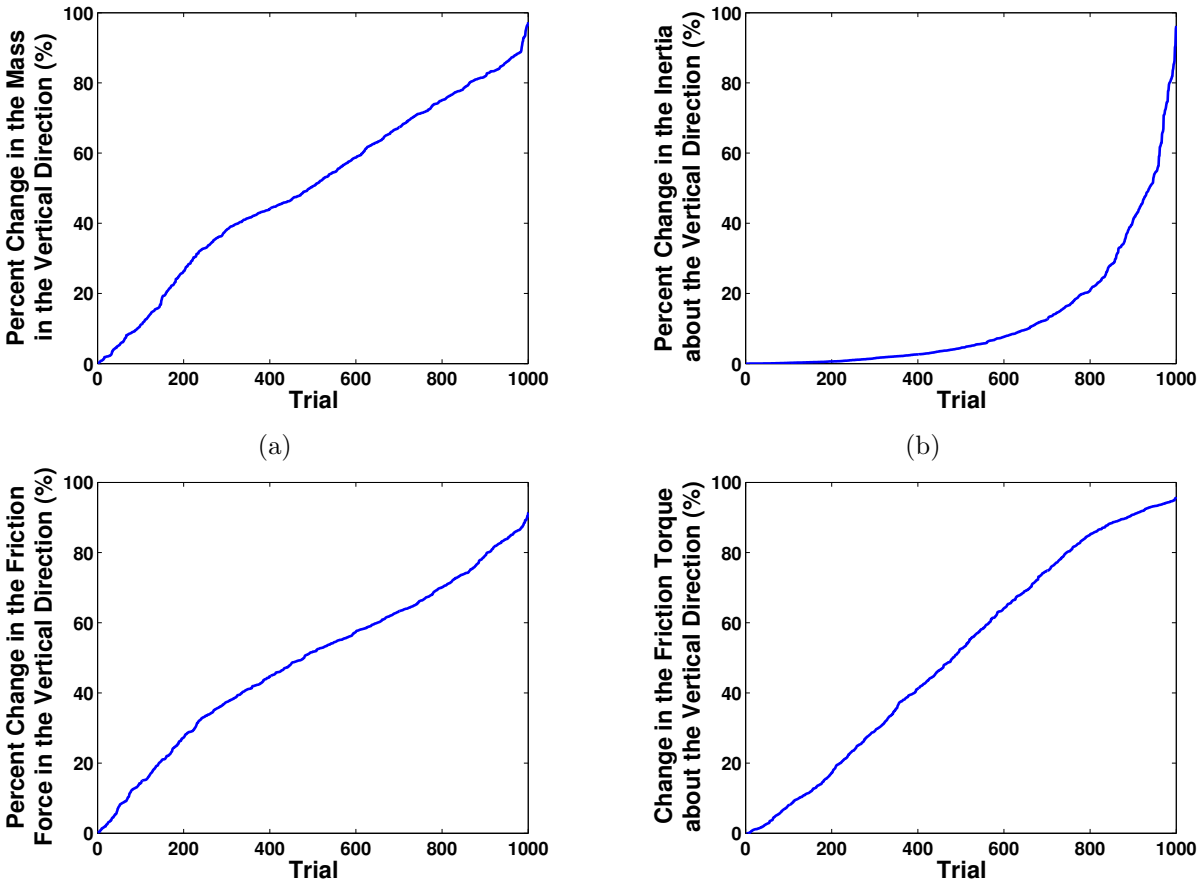

Figure 4.3: The percent change in the effective mass (a), effective inertia (b), effective frictional force (c), and effective frictional torque (d) in the vertical direction between the corresponding local minimum and maximum postures for 1000 random poses.

	Mass	**Inertia**	**Friction Force**	**Friction Torque**
Mean (%)	49.99	12.86	49.15	50.83
Standard Deviation (%)	25.05	18.61	23.14	31.48

Table 4.2: Mean and standard deviations of the percentage reduction at the optimal pose relative to the maximum pose for the directional mass, inertia, friction force, and friction torque optimisations.

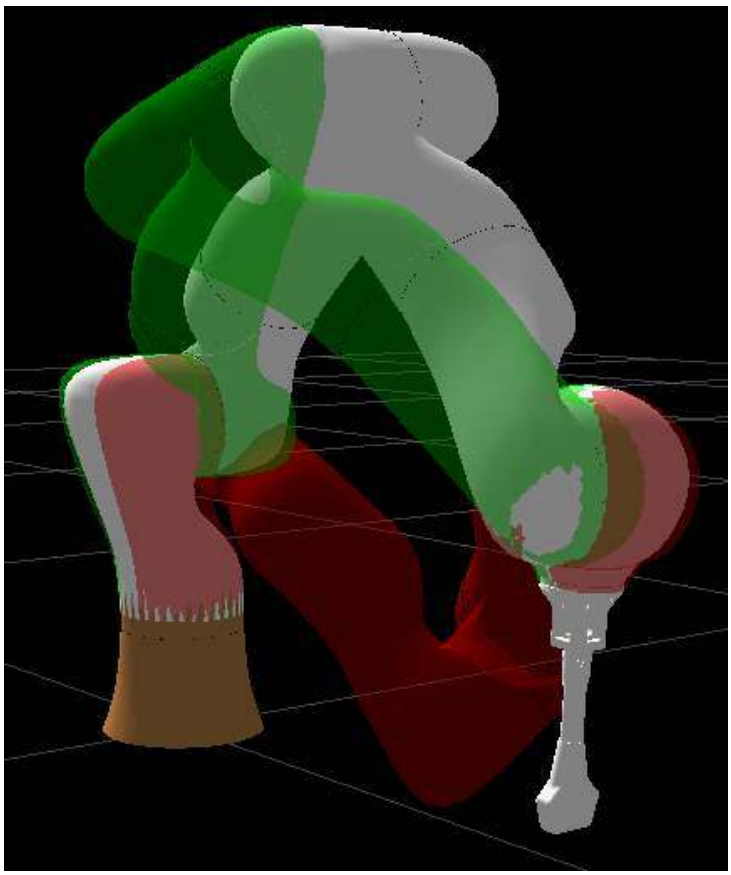

Figure 4.4: A visualisation of the local minimum inertia pose (green) and the local maximum inertia pose (red) overlaid on the initial joint posture (white), for a case in which there is limited improvement in the inertia about the vertical direction.

optimisation scheme cannot affect the inertia about the final joint axis, however, off-axis inertia is affected by the positions of the other joints, which is why the overall inertia optimisation controller more substantially decreases the inertial in the principal directions.

4.3.3 Combined Mass and Friction Optimisation Trade-off

As mass, inertia, friction force, and friction torque have different units of measurement, there must be a weighted trade-off to simultaneously optimise any of the two quantties. Figure 4.5 depicts an example of this trade-off using RVIZ (wiki.ros.org/rviz), by overlaying the optimal mass pose, optimal frictional force pose and an equally weighted trade-off solution.

The weights in a multivariate optimisation determine the relative importance of the quantities being optimised. However, if the optimal poses are too far apart, which can occur when

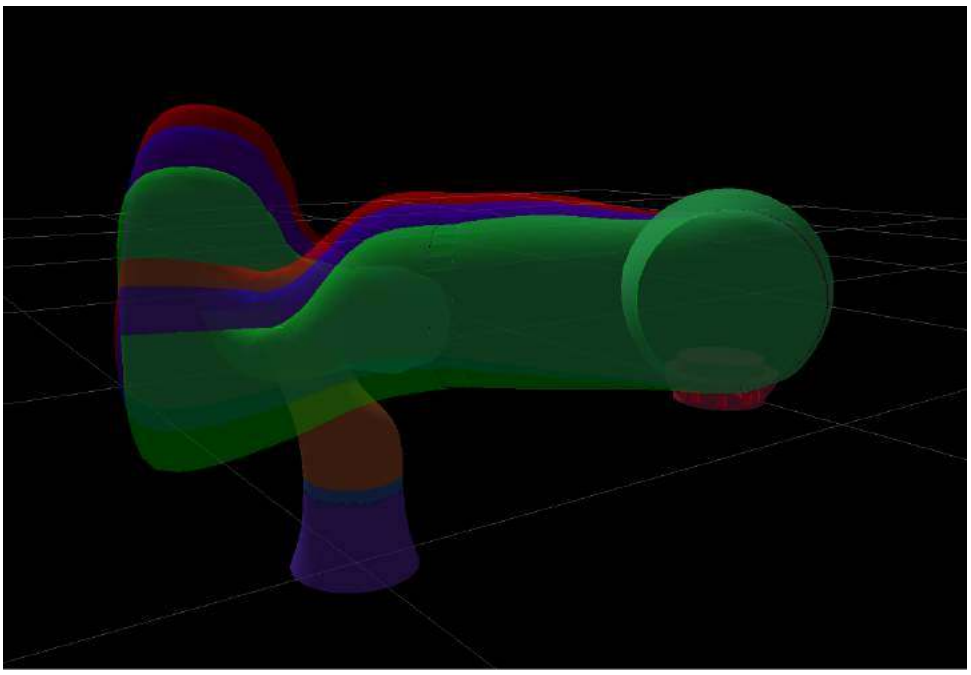

Figure 4.5: A comparison of the optimal mass pose (red), optimal frictional force pose (green), and an equally weighted trade-off pose (blue).

optimisation goals are conflicting, a trade-off solution may not work well in reducing either quantity and only one of the criteria must be chosen. This section aims to determine, on the Kuka LWR 4+, how much of an optimisation trade-off occurs in general and to find reasonable weight parameters to be used in the experiments.

To examine the effect of weighted optimisation, weighted solutions for 1000 random configurations were simulated and the resulting optimality measures were compared with the optimal mass poses and optimal friction force poses. Weighted simulations for 90%-10%, 50%-50%, and 10%-90% mass-friction ratios were performed.

Figure 4.6a shows the percentage trade-off for the principal mass component for the weighted solutions, compared with the optimal mass configurations found in Section 4.3.1. Similarly, Figure 4.6b depicts the percentage difference for the magnitude of the friction force vectors of the weighted solutions compared with the optimal frictional force pose found in Section 4.3.1.

In the majority of cases, the weighted optimisation does not have a substantial effect on the mass and frictional force values, as compared to the single optimisation results. Indeed, the

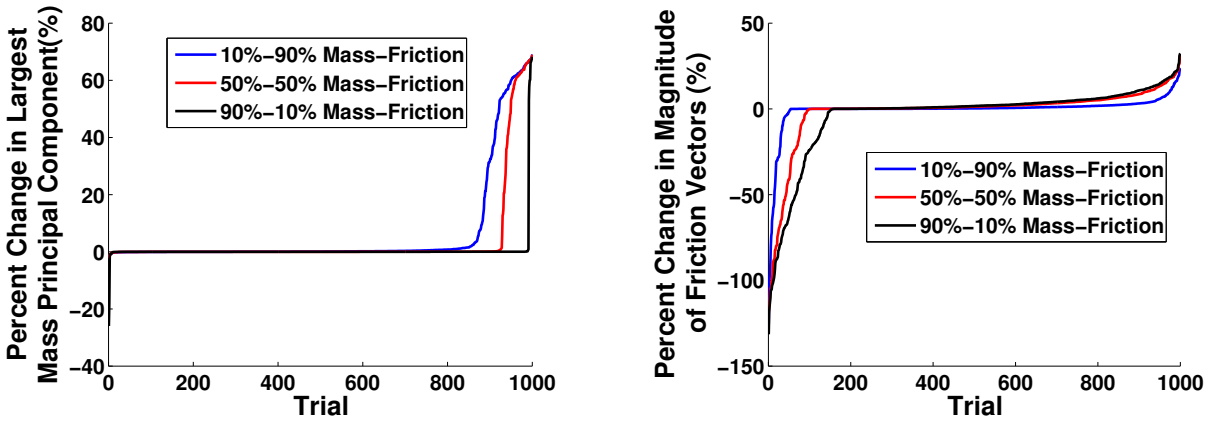

(a) (b)

Figure 4.6: Percentage change in the largest principal mass component between the optimal and weighted trade-off solutions (a) and percentage change in the magnitude of frictional force vectors between the optimal and weighted trade-off solutions (b).

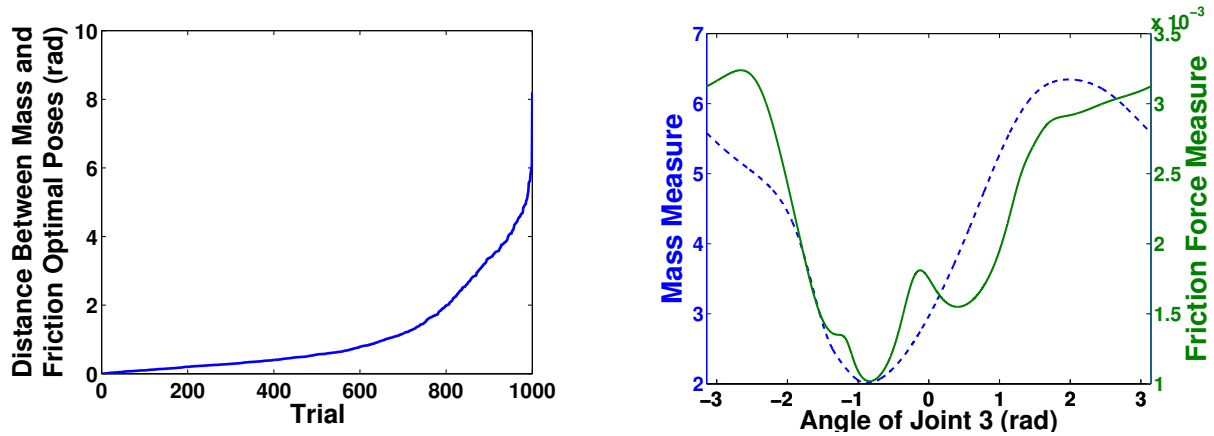

(a) (b)

Figure 4.7: The L^2 distance between the optimal mass and friction force poses for 1000 random trials (a) and the mass and frictional force measures for a complete rotation of one of the joints for a single pose, demonstrating the local optima that can be found in the frictional force measure.

change is close to zero because the distance between the optimal frictional force and mass poses is small in general, as shown in Figure 4.7a, which depicts the L^2 joint distance between the local optimal mass and local optimal frictional force postures.

For a small subset of cases, the frictional force is lower in the combined optimisation than at the local friction measure minimum. This is due to the greater degree of nonlinearity present in the friction model, which leads to having more local optima. To illustrate this, Figure 4.7b depicts the mass and friction force measures for a full null-space rotation of a single pose in which the friction force measure was better at the mass optimum. This figure is parameterized by joint 3 as this joint completes a full 2π rotation during the null-space motion. A local friction force optimum exists at $q_3 = 0.4557rad$ while the global friction force optimum is at $-0.8445rad$, which is very close to the mass optimum at $-0.8599rad$. As a result, in select cases, the combined optimisation leads to a better result in friction.

Overall, the results demonstrate that the trade-off on the Kuka LWR 4+ is minimal in general and therefore, a weighted trade-off solution is valid for the device. To implement the strategy on the experimental system, the weighting coefficients must be selected. One way that these can be chosen is on the basis of the relative magnitudes of the friction and mass (i.e. a system which is more affected by friction is more heavily friction weighted). However, attempting to compare the relative worth of mass in kilograms and friction in Newtons can be difficult. Alternatively, the choice can be made based on the relative improvements made by the strategy (i.e. a system with high mass but very little null-space improvement is weighted more towards friction). We suggest performing simulations optimising mass and friction over the workspace as done in Section 4.3.1 and weighting based on the averages of the resulting percentage improvements. As these percentage improvements are approximately the same on the LWR (44.23% for mass as compared to 41.44% for friction), a 50%-50% mass-friction ratio has been chosen for the experiments.

4.3.4 Impact of Modeling Uncertainty on Optimisation

The dynamic parameters and frictional torques on the joints of a real system can be estimated using techniques such as [Sousa and Cortesão, 2014], [Gautier et al., 2013], and [Kermani et al.,

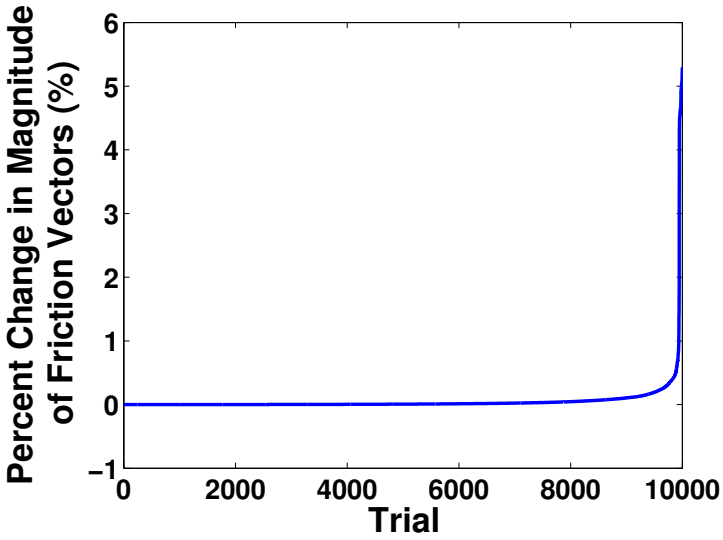

Figure 4.8: Model uncertainty optimisation results for frictional force using normally distributed joint friction estimates.

2007]. However, there can still be a mismatch between the real values and the estimates, which may have an effect on the optimisation with these parameters. In Section 3.5.4, it was shown that reasonably small variations in the inertial parameters do not have a substantial effect on the outcome of the optimisation for mass through 1000 random trials of a single pose, where the inertial parameters were drawn from a normal distribution with mean equal to the true value of the parameter and standard deviation arbitrarily set equal to 10% of its value.

To investigate the effects of frictional torque estimation on optimisation, a series of trials were performed to compare the optimisation results using the true joint friction values and the results found using estimated joint friction values, which were drawn from a normal distribution, $\tau_{f,e} \sim \mathcal{N}(\tau_f, (0.1\tau_f)^2)$. 100 trials were performed for each of 100 random poses (10,000 trials in total).

Figure 4.8 depicts the percentage increase in the magnitude of the friction force for the estimated joint friction optimisation, as compared to the friction force found using the true friction values, arranged in ascending order. The results demonstrate that the resulting friction vector, when optimising using the joint friction estimates is, in the worst case, 5.298% larger in magnitude than the optimal vector found using the true friction values. The average increase was 0.0631% with a standard deviation of 0.361%. As compared to the optimisations of Section 4.3.1, in which the average decrease in the friction vector using the optimisation strategy was found to

be 41.44%, the results suggest that this level of parameter variation does not have a large effect on optimisation.

4.4 Experimental Results and Discussion

So far, I have demonstrated the efficacy of the proposed methods over the workspace of the Kuka LWR 4+ on static end effector poses in simulation. This section aims to validate our technique in a experimental setup, in which the end effector's position is changing over time. The point-to-point user trials performed here demonstrate the method results in a reduction in the work required to move the tool over conventional surgical null-space controllers and, hence, a more natural motion for the user.

4.4.1 Kuka LWR 4+ Implementation

To implement our methods on the Kuka LWR 4+, the control strategy presented previously for simulation needed to be adapted. Kuka provides three controllers for the LWR 4+; joint position, joint impedance, and Cartesian impedance. However, the joint impedance mode focuses on setting desired positions along with stiffness and damping parameters, rather than directly setting torques or current, which the null-space gradient methodology requires. Additional joint torques can be added in joint impedance mode, but, due to the filtering performed by the Kuka controller and the $1.5Nm$ friction threshold discussed in Section 4.2.4, I found the system quite difficult to control using only joint torques, particularly for null-space motions that require more precise combinations of torques to ensure that they do not affect the main task. Cartesian impedance mode, which can be used to set end effector stiffness, damping, position and orientation, has additional null-space parameters which can be to used set the position, stiffness and damping of a joint positioning task projected into the null-space of the Cartesian pose. However, this again does not allow us to use the original formulation of our optimisation, as the null-space based gradient descent directly computes torques as opposed to positions.

To allow us to implement our controller on the robot, I utilised the Cartesian impedance mode

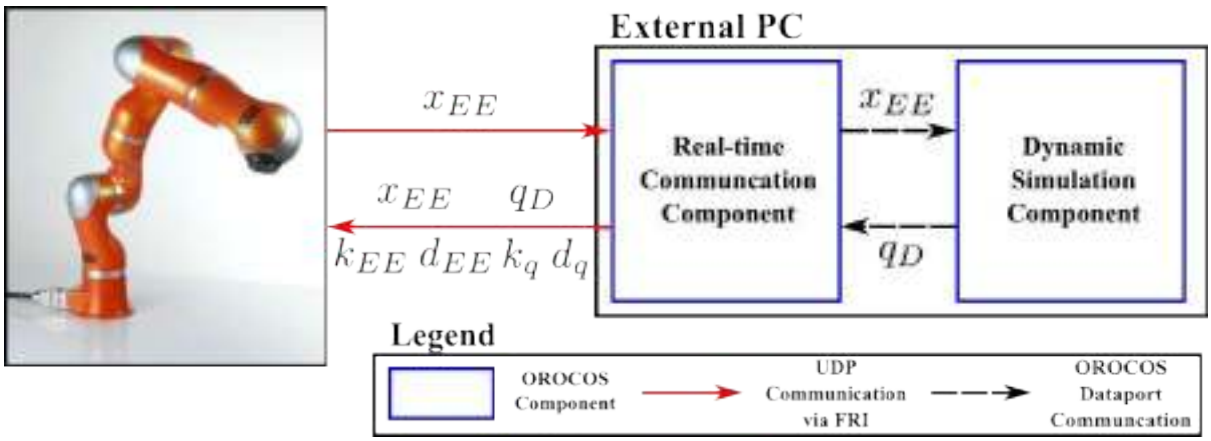

Figure 4.9: Data flow diagram of optimisation controller with Kuka LWR where x_{EE} is the current Cartesian pose of the robot, k_{EE} is the desired Cartesian stiffness, d_{EE} is the desired Cartesian damping, q_D is the desired null-space joint angles computed by the dynamic simulation, k_q is the desired null-space joint stiffness and d_q is the desired null-space joint damping.

with the additional null-space parameters, in combination with a forward simulation of the dynamics using the controller from the simulations (3.18). Figure 4.9 shows a diagram of the controller's data flow. The current Cartesian pose of the robot, x_{EE}, is received through the Fast Research Interface (FRI) [Schreiber et al., 2010] of the robot controller via User Datagram Protocol (UDP). This pose is set to the desired end effector pose in a dynamic simulation of the system, which uses the proprietary rigid body dynamic library of the LWR provided by Kuka and the control torques computed using (3.18) to integrate the system one time step into the future using the Dormand-Prince 5 method [Dormand and Prince, 1980], provided by the odeint library in boost::numeric.

The robot is then commanded through the FRI using the current Cartesian pose of the Kuka, with the Cartesian stiffness and damping, k_{EE} and d_{EE} respectively, set to zero. The joint angles on the forward simulated system, q_D, are commanded as the joint angles in the null-space controller of the Kuka, with an appropriate stiffness and damping, k_q and d_q. To minimise any errors in the null-space control affecting the user's commands, lower gains of $100Nm/rad$ for the stiffness and 0.1 for damping were used.

1 kHz communication was achieved through the FRI on a PC (i7-3770 @ 3.40 GHz) running Ubuntu 12.04 (www.ubuntu.com), with the Xenomai real-time kernel patch (www.xenomai.org), RTnet for real-time networking (www.rtnet.org), and OROCOS for real-time processing compo-

nents (www.orocos.org). On this system, the forward simulation of the dynamics for a time step of $1ms$ was computed in approximately $0.8ms$ for the worst case of simultaneously computing all of the optimisation gradients.

The resulting behavior of this controller mimics the simulations above and takes advantage of the tuned, low level controllers of the Kuka. It assumes that the flexible joint dynamics are handled by the Kuka controllers, i.e. the Kuka is a rigid body robot, and that the rigid body dynamic model provided by Kuka is accurate. In 3.5.4, mass optimisation simulations were performed using estimated model parameters, which were drawn from a normal distribution with mean equal to the true value of the parameter and the standard deviation equal to 10% of its value. It was found that small variations from the true values of the inertial parameters do not affect the results of the optimisation substantially. Section 4.3.4 demonstrates similar results for frictional optimisation.

4.4.2 Kuka LWR 4+ Experimental Setup

Figure 4.10 depicts the experimental setup which includes the Kuka LWR 4+, ATI Gamma force/torque sensor (www.ati.ia.com) for measuring the subject's applied forces and torques, 3D printed tool, and the mock surgical target. The Gamma force/torque sensor was calibrated with a range of $32N$ and resolution of $\frac{1}{160}N$ for forces in the x- and y-direction, a range of $100N$ and resolution of $\frac{1}{80}N$ for forces in the z-direction and a range of $2.5Nm$ and resolution of $\frac{1}{2000}Nm$ for torques.

The experiment aims to demonstrate that the methodology presented here allows for a more natural motion in that it requires less work to perform the same task, when compared to conventional surgical controllers. As I am not applying guiding forces to the end effector, the user's strategy to maneuver the tool can affect the results since work is path dependent. To reduce the degree of this variation and to more clearly see the effect of the optimisation, the experiments focused on optimising mass and frictional force only and utilised a mainly translational task — a series of point-to-point motions. The experiments also solely focused on translation due to the complexity of designing a 6D target and to reduce the time required to perform the experiments, already approximately 30-40 minutes per subject.

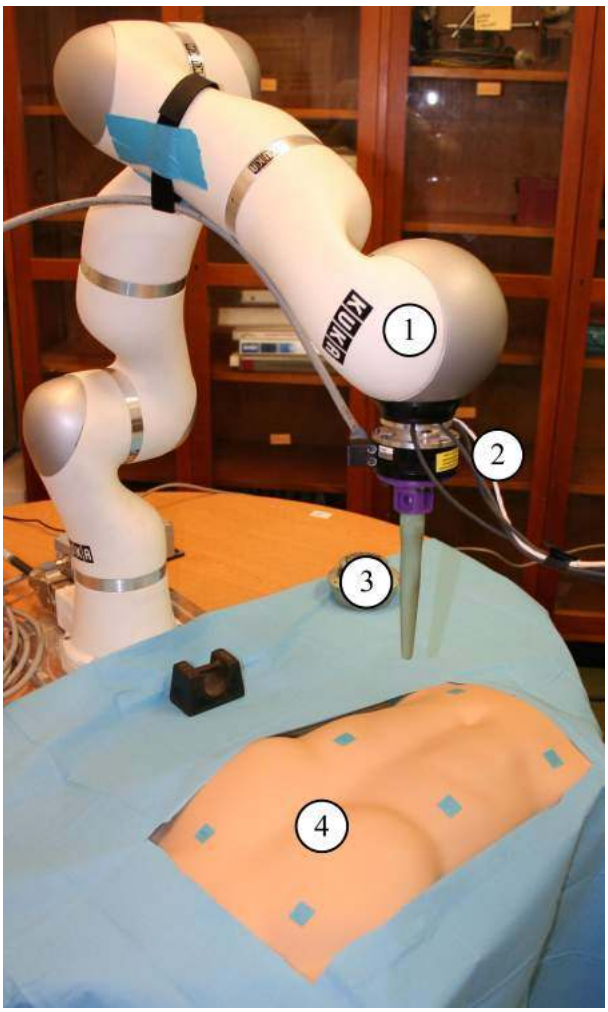

Figure 4.10: Experimental setup consisting of Kuka LWR 4+ (1), ATI Gamma force/torque sensor (2), 3D printed tool (3) and surgical target (4).

Additionally, while the Kuka LWR 4+ provides the required redundant, compliant interface, it is still a research system and does not have the capability to operate on the small scale required for a minimally invasive surgery. Current hands-on surgical robots are capable of aligning orthopedic joint replacement implants within $2mm$ and 5 ° of the preplanned implant position [Citak et al., 2013], a level of precision the LWR cannot achieve. Therefore, the task in which I will evaluate the effectiveness of our controller will function on a larger scale than is typical in conventional, hands-on robotic surgery.

A total of 8 engineering students and post-docs participated in the trials, 5 men and 3 women, with ages ranging from 20-34. 50% of the subjects had not worked on or with a robot previously. The subjects were asked to grip the 3D printed tool and not the robot's body. As the focus

of this experiment was on the optimisation of mass and frictional force, both translational quantities, the subjects were asked to keep the tool's orientation approximately normal to the surgical target. To reduce stress on the subjects' wrist, however, rotations about the normal to the surgical target were allowed. Prior to performing the trials, the subjects were given some time to familiarise themselves with the controllers, but were not given details about the aim of the controllers or the experiment in general.

Six reference points were chosen on the body of the surgical target. Subjects moved the tool to touch these points in five different orders, which were generated randomly and ordered randomly to avoid selection and training bias. The subjects maneuvered through these five sets of points three times for each of four controllers. For each subject, the order of the controllers was also random to avoid training bias. The controllers used in this experiment were active optimisation of mass and frictional force, fixed optimal posture and two other controllers, which I believe to represent the current state-of-the-art in hands-on robotic surgery: elbow up control and damped posture control. Due to the absence of published control strategies for redundant, hands-on surgical robots, I drew this conclusion based on direct experience with the only two commercial systems of this kind, the Mako Rio (now of Stryker Corp.) and the Stanmore Implants Worldwide Sculptor (now also of Stryker Corp.).

Active optimisation was performed using a ratio of $50\% - 50\%$ mass measure to frictional force measure, as the system demonstrates an equal reduction in these quantities under our optimisation technique, as discussed in Section 4.2.3. The parameters of (3.16) were set to $k_p = 50$, $k_d = 10$, and $\omega_{max} = 10 rad/s$, the potential velocity was set to $5 cm/s$ for the reasons explained in Section 4.3.3, and the null-space parameters on the Kuka were set to $k_p^{null} = 100$ and $k_d^{null} = 0.1$.

In 3.5.6, a fixed optimal posture solution was shown in simulation to perform nearly as well as active optimisation. This strategy aims to allow for the benefits of mass and friction optimisation to still be taken advantage of in situations where the surgical environment is heavily constrained. By pre-operatively finding the optimal solution for the region in which the majority of the surgery will take place, a constant redundancy position control strategy can be used to minimise the squared joint position to this optimum and the surgical equipment can be arranged around

an already optimised robot. By including this controller here, I aim to verify that this type of strategy works in an experimental setup.

The elbow up controller maneuvers the redundancy to keep the robot out of the general work area of the end effector. This type of control aims to keep the elbow above the end effector and out of the way of the surgeon. The damped posture controller applies gravity compensation and damping to the redundancy of the robot. This controller is typically used as a simple way for the user to reconfigure the robot manually during tasks. The redundancy remains approximately constant while moving the end effector and a force can be applied to the redundancy to adjust it accordingly.

The elbow up and fixed optimal posture strategies were implemented using the Cartesian Impedance mode in the FRI by setting zero stiffness and damping in the Cartesian task and commanding the necessary constant joint posture in the null-space, with stiffness equal to $100Nm/rad$ and damping ratio equal to 0.1. The damped posture controller was the default gravity compensation provided by Kuka, which exhibits the damped postural characteristics required.

4.4.3 Kuka LWR 4+ Experimental Results and Discussion

Natural motion focuses on reducing the amount of effort it takes for the surgeon to move the tool attached to the robot. Therefore, the metric used here to compare the control methodologies is the sum of the absolute value of work over each trial (i.e. the amount of energy the users transferred to the end effector). The work was computed using the difference in the Cartesian position calculated from the joint position encoders and the wrench from the force/torque sensor.

For a single trial of one of the subjects, the cumulative work over time for all four controllers is shown in Figure 4.11. The damped posture and elbow-up controllers result in more work over time as compared with the fixed and active optimisation trials. Additionally, the separation between the standard and optimal controllers appears to be increasing over time, suggesting that the effect would increase in longer tasks, such as those that take place in surgery.

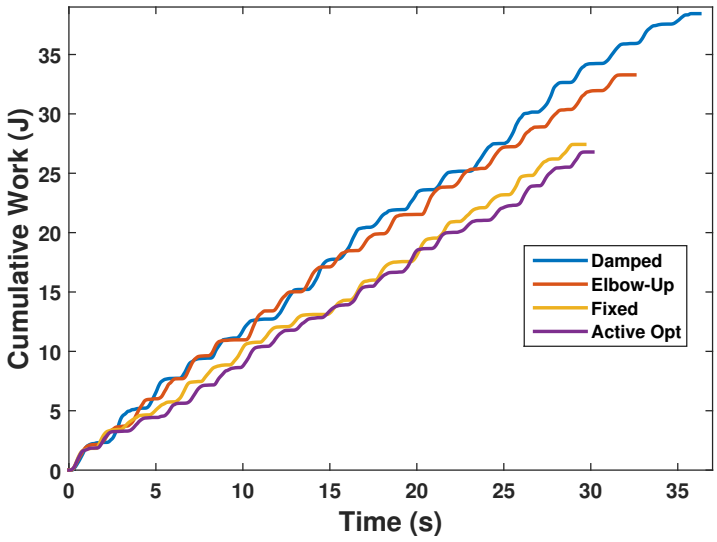

Figure 4.11: Comparison of the cumulative work for one user's trial over time for each experimental test controller.

For the results for all subjects, the work of each trial was normalised by the length of the path over the trial to account for variations in the lengths of the randomised paths. Additionally, as I am not applying guiding forces to the end effector, the precise path cannot be prescribed. General contact-to-contact point style motions, such as those in the trials, exhibit a parabolic shape, however, there exist user specific variations to these motions including the height of the trajectory and the amount of deviation from the vertical plane passing through sequential points [Ziherl and Munih, 2009]. These variations and the amount of rotation applied to the tool must be taken into consideration as work is path-dependent.

To account for these factors and compare the results between subjects, the length normalised work of each trial was additionally divided by the mean of the control method with the highest length normalised work per person. This allows us to examine the relative improvement between users for the various control methods, taking into account user control variability. The normalisation can be summarised as follows:

$$W_n^{i,j,k} = \frac{\frac{W_i^{j,k}}{L_j}}{\max_m \frac{1}{N_{trials}} \sum_{l=1}^{N_{trials}} \frac{W_i^{l,m}}{L_l}} \tag{4.10}$$

$W_i^{j,k}$, the work of subject i for trial j of path k, is divided by the length of the trial, L_j, and then

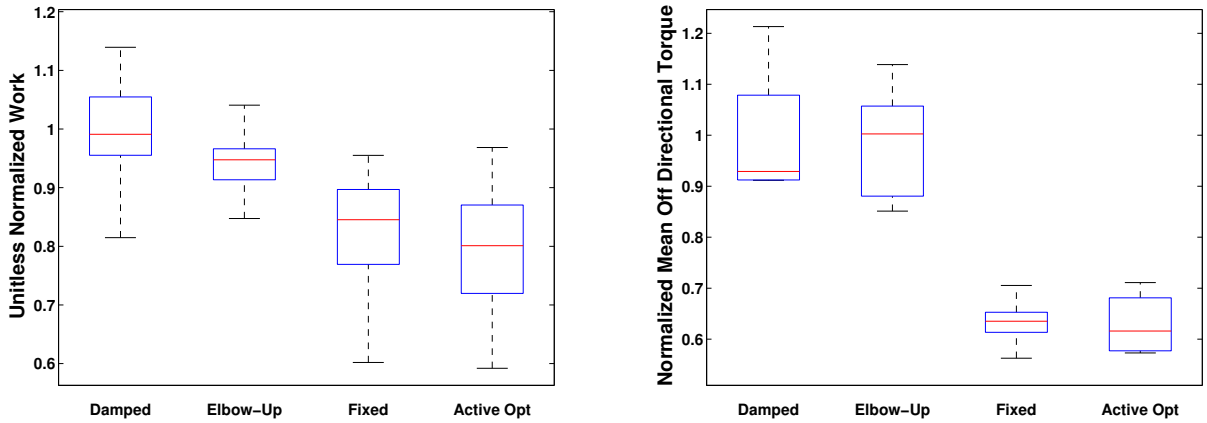

(a) (b)

Figure 4.12: The median and IQR for the cumulative normalised work value (a) and mean torque orthogonal to the direction of motion normalised by the control method with the highest torque (b) across subjects for the four controllers.

by the mean of the path with the highest average length normalised work for subject i.

Figure 4.12a depicts the normalised absolute value of work for all subjects across the experimental trials for each of the four controllers. The Komogorov-Smirnov test for normality was used to determine that at least one of the work distributions was not normal. Therefore, the non-parametric Kruskal-Wallis one-way analysis of variance test was employed to demonstrate that at least one of the distribution's medians was different from at least one other group ($p < 0.01$). Pairwise comparisons were made using the Dunn-Sidak test to analyse the specific pairs for dominance. Table 4.3 depicts the results of this test, where a 1 indicates a significant difference at the 0.01 level and a 0 indicates the opposite. The results demonstrate that the fixed, pre-optimised solution and active optimisation strategies differ significantly from the damped and elbow-up strategies. This shows that the optimisation strategies developed here require less work to perform the same task as compared to standard hands-on surgical redundancy control strategies and therefore, they create a more natural motion.

	Damped	**Elbow-Up**	**Fixed**	**Active Opt**
Damped	X	0	1	1
Elbow-Up	0	X	1	1
Fixed	1	1	X	0
Active Opt	1	1	0	X

Table 4.3: Dunn-Sidak results for the experimental trials where a 1 indicates a significant difference at the 0.01 level and a 0 indicates the contrary.

Additionally, Figure 4.12b depicts the mean magnitude of the torque orthogonal to the direction of rotation, normalised per user to account for within subject variation. The Kruskal-Wallis test determined that at least one of the distribution's medians was different from at least one other group ($p < 0.01$). The Dunn-Sidak test was used to determine that the damped and elbow-up differ significantly from the fixed and active optimisation strategies at the 0.01 significance level. As the trials focused on translational motion, these torques can be interpreted as those which the users applied to ensure a consistent orientation of the end effector. The higher value in the damped and elbow-up trials would suggest that the users found it more difficult to deal with the tool, when compared to the fixed and active optimisation cases. Indeed, all subjects informally commented that the natural motion controllers were much easier to perform the task than the damped and elbow-up modes.

Lastly, the results indicate that the fixed, pre-optimised control strategy can perform nearly as well as the active optimisation in situations where the set of poses that the surgeon will use to perform surgery do not differ significantly from the optimal pose, which is generally the case in minimally invasive surgery. This demonstrates that, even in heavily constrained surgical environments, where changes in redundancy are limited, considering end effector mass and friction can reduce the impact of the end effector dynamics on the surgeon. In less constrained environments, with more significant changes in the end effector poses, the active optimisation controller would be expected to improve results over the fixed optimisations.

4.5 Summary

This chapter has presented a null-space based optimisation methodology which reduces the projection of the joint friction onto the end effector in order to provide the surgeon with an experience that is closer to the tool unattached to the robot. Performing a null-space optimisation of the friction at the end effector directly using the joint friction values estimated from a model would result in rapid changes in the robot configuration when quickly changing direction of the end effector. In addition, no optimisation would occur when the system was at rest since the friction would be zero in such a case. A strategy which continues to optimise when the surgical instrument is not moving is able to ensure a low friction when the next motion takes

place. Lastly, as the surgeon is in complete control of the tool, I do not have knowledge about the direction and speed of the next instrument motion. A methodology which can configure the robot such that the surgeon feels lower friction regardless of the next direction of motion is advantageous. Therefore, I have developed the "potential velocity" concept to create a friction optimisation criterion which is joint pose consistent and can provide a general optimisation with respect to arbitrary motions for our natural motion strategy.

This chapter has developed overall and directional optimality measures for the friction force and torque at the tool point using the "potential velocity" of the manipulator. Simulation work has demonstrated the effectiveness of the proposed mass and friction metrics in reducing their respective quantities over the workspace of a seven degree-of-freedom lightweight serial manipulator. Additionally, a small trade-off was found when performing a combined mass and friction optimisation for this particular robot. An investigation into the effect of optimising friction using estimates of the true friction parameters on the robot demonstrated that, for a given level of parameter variation, there was not a large effect on the resulting optimal pose and corresponding friction. Lastly, an experimental point-to-point motion user trial demonstrated a reduction in work when moving the tool, as compared to classic redundancy resolution techniques, resulting in a more natural motion for the subjects.

In the following chapter, a dynamic active constraints approach will be developed, which will prevent surgical tool motion into deforming soft tissue regions which have been deemed unsafe or unnecessary for the surgery. This strategy will generate an implicit constraint surface directly from point clouds, as this allows the strategy to be applied to various sensing modalities for intra-operative soft tissue deformation tracking, expanding the methodology's application to a variety of surgeries. Simultaneous control of the dynamic active constraints strategy in Chapter 5 and the mass and friction optimisation presented in this chapter will be performed in Chapter 6, through a virtual reality ultrasound scenario.

Chapter 5

Dynamic Active Constraints through Implicit Surfaces

This chapter presents a novel hands-on control methodology for generating meshless dynamic active constraints directly from point clouds using implicit surfaces. By creating the dynamic active constraints from point clouds, this method can be applied to general soft tissue tracking strategies, such as CT scans, MR imaging, and stereo cameras and therefore, to a variety of surgical procedures. Additionally, the control technique presented here utilises the dynamics of the robot to ensure an appropriate constraint impedance, a necessary requirement for hands-on robotic surgery. The research presented here is an edited version of the work published in:

> J. G. Petersen and F. Rodriguez y Baena, "A dynamic active constraints approach for hands-on robotic surgery", in Proceedings of the IEEE International Conference on Intelligent Robots and Systems, pp. 1966-1971, Nov. 2013.

5.1 Introduction

Chapters 3 and 4 demonstrated a natural motion control strategy which reduced the effects of the end effector dynamics on the surgeon. By decreasing the effective mass and friction at the tool point in the null-space, the surgeon's ease of use was improved and fatigue over long surgeries could be lessened, without the surgeon's desired tool pose, velocity, or forces being

affected. In this chapter, I develop a dynamic active constraints control strategy for hands-on robotic surgery to aid in ensuring patient safety. The combined strategy of natural motion and dynamic active constraints will be investigated in Chapter 6 in a virtual reality ultrasound scenario.

Forbidden region active constraints in surgery are space sharing cooperative control strategies which aim to prevent surgical tool motion into regions considered dangerous, or unnecessary for the surgery. They have been successfully applied to procedures taking place on rigid anatomy in commercial hands-on robotic systems such as the Rio (Mako Surgical Corp.) and the Navio Precision Freehand Sculpting system (Blue Belt Technologies Inc.). Dynamic active constraints aim to extend the applicability of active constraints to soft tissue surgeries and thereby, increase accuracy and safety in these deforming environments.

Figure 5.1 depicts a simplified mock-up of a scenario for dynamic active constraints in hands-on robotic surgery. As opposed to static active constraints which only require a rigid registration with the bone structure being operated on, dynamic active constraints require an intra-operative tracking system to precisely measure the deformations of the delicate tissue structures in order to ensure the correct area is protected throughout the surgery. In addition, the motion and deformation of the target and the forbidden regions is, in general, unknown *a priori*. Therefore, the constraint representation must be able to handle arbitrary shapes and deformations resulting from interactions between the surgical tool and soft tissue, patient breathing and pulsation, and other effects such as brain shift [Slotty et al., 2012].

Current dynamic active constraint strategies place restrictions on the motion of the deforming region or are restricted to particular representations of the underlying geometry, preventing their applicability to general surgical procedures and constraint region tracking methodologies. The dynamic virtual fixtures of [Ren et al., 2008] for beating heart surgery were precomputed from MR and CT images and, therefore, were unable to handle unpredicted motion of the heart, such as deformation due to tool contact or abnormal changes in the cardiac motion. The teleoperative dynamic active constraints of [Gibo et al., 2009] were only demonstrated in one dimension. [Navkar et al., 2012] provided a 3D strategy which was specific to constraining tool motion inside blood vessels for heart surgery. Dynamic frictional constraints were demonstrated

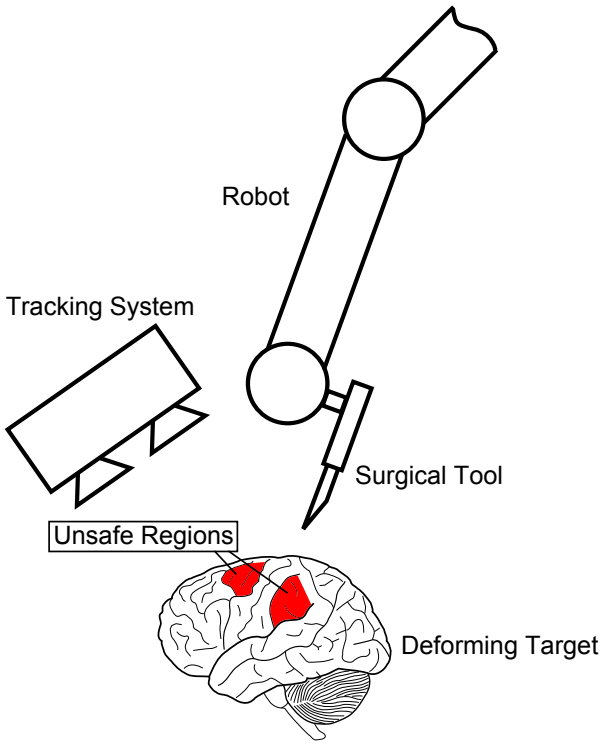

Figure 5.1: A mock-up of a scenario for dynamic active constraints in hands-on robotic surgery consisting of a robot with attached tool, a deforming soft tissue target, and a tracking system.

in position and orientation by [Bowyer and Rodriguez y Baena, 2014], however, the method required a mesh representation of the constraint geometry, thereby restricting its application to sensing modalities which can accurately track and deform mesh vertices.

Dynamic active constraints require real-time updating of the deforming tissue with sufficient procedure dependent accuracy to ensure the delicate anatomy is protected at all times. Registration of deforming soft tissue is an ongoing research problem, with various different methods having been demonstrated in magnetic resonance imaging (MRI) [Ferrant et al., 2002] [Kerkhof et al., 2009], ultrasound (US) [Wein et al., 2013] [Weon et al., 2015], computerised tomography (CT) [Lawson et al., 2007] [Wein et al., 2008], stereo camera imaging [Ji et al., 2014] [Faria et al., 2014], and laser range scanning [Cao et al., 2008] [Ding et al., 2011]. Many of these strategies have not been shown to work in real-time or can only function at low accuracy in real-time. In addition, the applicability of the sensing strategies is highly dependent on the type of surgery being performed. A strategy which can generate the forbidden region directly from point clouds is advantageous, as each of the above methods can be reduced to this data form and, therefore, such a dynamic active constraints approach would not be dependent on

the sensing strategy utilised and could be applied to more types of surgeries. Additionally, as various sensing strategies improve in the areas of accuracy and speed or new sensing modalities are developed, such a strategy could easily be transfered.

The methodology used in this research to generate the dynamic constraint surface from point clouds is similar to work on constraint based haptic rendering for dynamic environments [Leeper et al., 2012] [Rydén and Chizeck, 2013], in that these works use implicit surfaces to define the constraint surface from streaming point clouds. However, the control strategy developed here additionally incorporates the dynamics of the robot into the enforcement of the constraint. In haptic rendering for teleoperation, safety issues for protecting the patient can be vastly simplified by decoupling the position of the master and slave device when the master penetrates the constraint boundary. The slave can be position controlled to remain on the surface of the forbidden region, ensuring the tool does not reach the delicate tissue, while forces are commanded to the user to return the master to free space. In hands-on robotic surgery, such a decoupling strategy is not possible, as the surgeon and the robot jointly control the surgical instrument and, as such, penetration into the constraint region will occur in impedance controlled devices. Additionally, the physical contact between the robot and the surgeon couples the dynamics of the two. Therefore, the dynamics of the robot must be incorporated into the enforcement of the constraint to ensure an appropriate impedance to the surgeon's forces and torques.

An alternative strategy to ensure the tool does not penetrate into the constraint region is to use a non-backdrivable system with admittance control, such as that employed in [Taylor et al., 1999] and [Bettini et al., 2004]. Non-backdrivable motors do not allow external torques to rotate the motor shaft. Therefore, the robot cannot move without a command to the motors. Admittance control systems translate the user's applied forces and torques into linear and angular velocities for the robot, typically in a proportional fashion. These so-called hard active constraints can be implemented in such a control setup by completely removing the commanded forces and torques which would maneuver the tool into the forbidden region when computing the desired velocity for the robot.

However, the use of admittance controlled robots in hands-on robotic surgery removes a critical enhancement that impedance control provides, that is, to allow the surgeon to directly feel

the interaction forces between the robot and the surgical environment through the tool. The surgeon can utilise this natural feedback to more precisely apply forces and torques to the target tissue. A lack of this type of feedback has been linked to increased tissue damage [Joice et al., 1998] [Wagner et al., 2002].

Joint level braking systems such as [Schneider and Troccaz, 2001] and [Tenzer et al., 2012] have been developed in an attempt to implement hard active constraints for impedance controlled systems. In addition to preventing motion into the forbidden region, these types of strategies are energetically passive in that they can filter the user's motion, but cannot apply forces to move the tool. These types of devices are intrinsically safer than autonomous or actuated systems, however, the application of braking torques at the joint level to impose constraints in Cartesian space can be discontinuous. Additionally, implementing this type of strategy requires a robot redesign to accommodate a brake at each joint. A software based solution for dynamic active constraints in hands-on robotic surgery is much more easily applied to general devices, including current clinically applied robots.

For all of the above reasons, a software based dynamic active constraints approach has been developed, which has the potential to be applied in the future to current commercial hands-on robotic surgical systems. The strategy can generate constraints directly from point clouds, allowing the application of this method to general tracking methodologies, and the dynamics of the robot have been incorporated into the enforcement of the constraint to provide an appropriate impedance to the surgeon's wrench.

This chapter will first demonstrate the metaball approach for generation of an implicit surface from a point cloud. Next, the Implicit Function Jacobian will be derived, which defines the normal to an implicit surface's level sets. Dynamic active constraints using the Implicit Function Jacobian will then be demonstrated, including a methodology which can incorporate predicted motion. A multi-priority null-space task hierarchy will be applied to allow for simultaneous application of the dynamic active constraints, unimpeded motion along the surface and posture control for null-space optimisations. Simulations demonstrate the proposed control methodology constraining a planar robot to constraint isosurfaces generated from two different points clouds, one in which the point cloud's motion is unknown and the other in which the point cloud's motion

is known. Finally, an experiment using a dynamic brain phantom and a lightweight seven degree-of-freedom robot demonstrates the feasibility of the proposed method on a physical setup. The dynamic active constraints strategy demonstrated in this chapter will be applied concurrently in Chapter 6 with the natural motion strategy developed in Chapters 3 and 4.

5.2 Implicit Surface Generation from Point Clouds

As previously discussed, there exist a variety of methods to track deforming soft tissue in order to accurately apply active constraints to protect the patient. The forbidden region can be registered using MRI, US, CT, stereo camera imaging or laser range scanning. These sensing modalities are dependent on the surgery type and many are still in development towards providing real-time updates of the soft tissue structure with high accuracy. An active constraints strategy which can function using each of these data types would simplify its application and allow for easy transference between sensing schemes as current strategies are improved in speed and precision and novel techniques are developed. Therefore, the methodology presented here generates the constraints surface from point clouds, as each of the above mentioned methods can be reduced to this constraint representation.

Point clouds are a generic representation of 3D objects which can handle arbitrary deformations and motions. In this section, an implicit surface will be produced from this point cloud to surround the area chosen to be protected. The surface is smooth and continuous in order to allow for unimpeded motion of the robot along the constraint surface. The methodology used in this section to generate the implicit surface is a modified version of that used in [Leeper et al., 2012]. In their work, implicit surfaces were used for haptic rendering of arbitrary point clouds and some of the effects of parameter choice, cloud density, and noise on the formation of the surface were analysed. In this work, implicit surfaces are used to generate a three dimensional constraint to which the robot will be controlled.

Compactly supported radial basis functions have been used to wrap the desired constraint point cloud in order to produce the implicit surface, as these functions provide several useful properties. A compactly supported function is one which is zero outside of some bounded set.

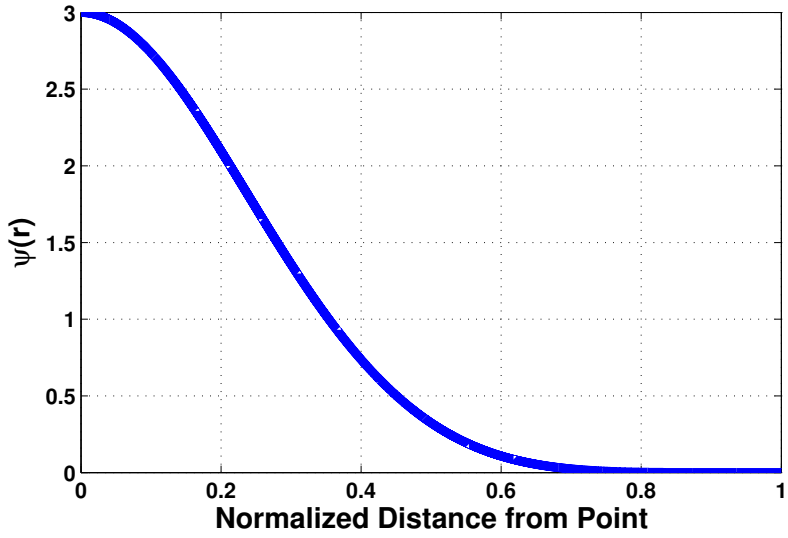

Figure 5.2: A plot of the value of the Wendland function 4,2 as a function of the normalised distance from the origin.

Compactly supported radial basis functions are only non-zero within a particular radius, $[0, R]$, around the point they are centered on. When evaluating where the tool point is in relation to the implicit surface, constraint points which are further than a distance R from the end effector will not affect the computation. This is valuable for enacting active constraints on large point cloud data sets as these data sets can be partitioned to decrease processing time and allow for real-time application. Additionally, the property of compact support can be used to detect that the end effector is approaching the constraint surface and allow control strategy switching. Lastly, the derivatives of compactly supported radial basis functions are zero at the edges of the compact set, $f'(0) = f'(R) = 0$, such that they provide a continuous scalar field. This property prevents discontinuities from forming in the constraint surface when the tool is approaching or moving away from a new point, and helps to ensure smooth control.

The particular radial basis function used in our experiments is the Wendland function $\psi_{4,2}(r)$ which is depicted in Figure 5.2 [Wendland, 1995].

$$\psi_{4,2}(r) = \begin{cases} \left(1 - \frac{r}{R}\right)^6 \left(35\left(\frac{r}{R}\right)^2 + 18\frac{r}{R} + 3\right) & r \leq R \\ 0 & r > R \end{cases} \tag{5.1}$$

where r is the current radius and R is the radius of influence.

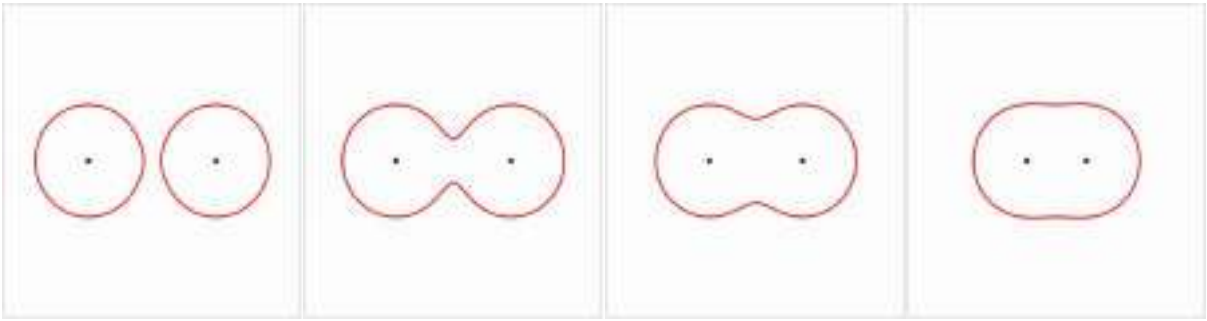

(a) (b) (c) (d)

Figure 5.3: The change in the constraint isosurface (red) which results as the constraint points converge (black crosses).

An implicit function is generated using radial basis functions by finding the weighted contribution of the distance between each sensed constraint point and the tool point. In computer graphics, this is commonly referred to as a metaball surface representation [Blinn, 1982].

$$f(x,t) = T - \sum_i \psi_{4,2}(\|x_{tp}(t) - x_i(t)\|) \tag{5.2}$$

where $f(x,t) = 0$ is the implicit surface which will be the isosurface that defines the boundary of the forbidden region, T is the threshold value, $x_{tp}(t) \in \mathbb{R}^p$ is the position on the robot's tool point at time t and $x_i(t) \in \mathbb{R}^p$ is the i-th sensed constraint point at time t.

To demonstrate how the metaball approach generates the isosurface and how multiple points are handled, Figure 5.3 plots the resulting constraint surface for two constraint points for varying inter-point distances. At far enough a distance, the points do not interact with one another and spherical isosurfaces are generated around each point. As the points come together, the surface interpolates to provide a single smooth isosuface around both the points.

Further, the radius of influence, R, and the threshold, T, impact the resulting constraint surface which is generated from the point clouds. To illustrate this, Figure 5.4 depicts the effect of these properties on the resulting isosurface, shown in red, when constructing the constraint surface from two constraint points, shown as black crosses. Increasing the radius of effect, which is the maximum distance a point can have an effect on the computation of the isosurface, increases the distance between the surface and the point, whereas increasing the threshold value reduces

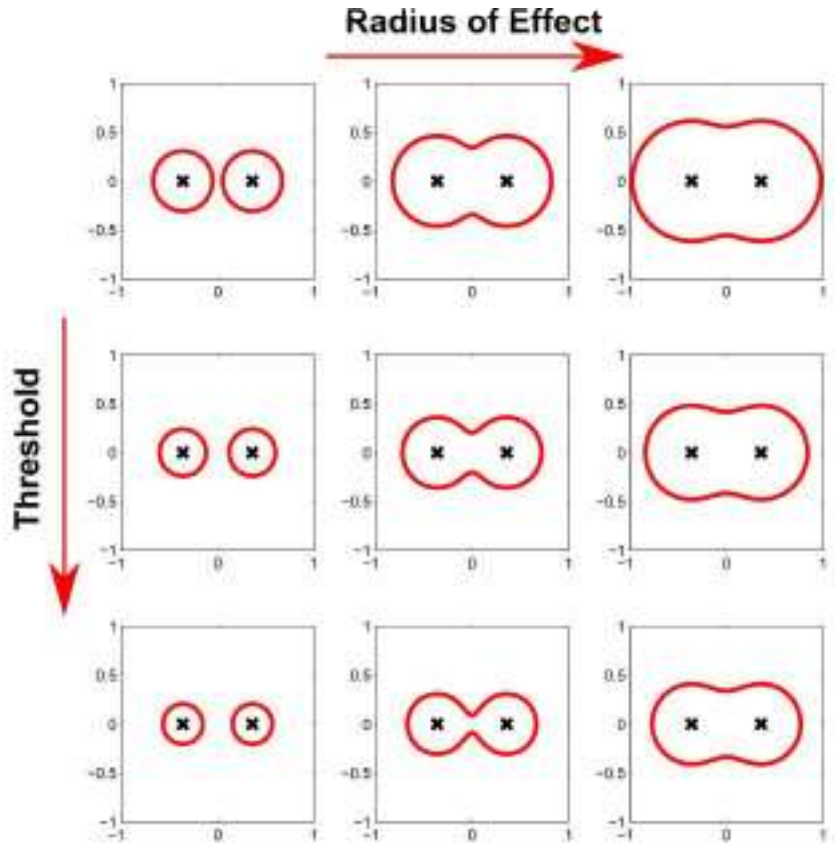

Figure 5.4: The effect of the radius of influence and threshold on the constraint isosurfaces (red) which result from a set of constraint points (black crosses).

the distance between the surface and the point. [Leeper et al., 2012] showed that the choice of the radius of influence and threshold to ensure a smooth isosurface should be based on the density of the point cloud which they defined as the average separation between points in the cloud. Results in the above two graphs shown are in 2D for simplicity, however, metaballs are easily applied in 3D.

Methods exist which can fit an implicit surface to a point cloud such that the zero set directly passes through the points [Savchenko et al., 1995] [Morse et al., 2005]. Alternatively, the method described above generates a surface which surrounds the point cloud defining the delicate tissue. In using an offset distance, a factor of safety is directly included into the representation, ensuring there is some distance for constraint forces to be applied before the surgical instrument reaches the vulnerable anatomy. This is vital since, as discussed in Section 5.1, hard active constraints can not be applied in hands-on robotic surgery. The minimum protection distance of r_{min} can

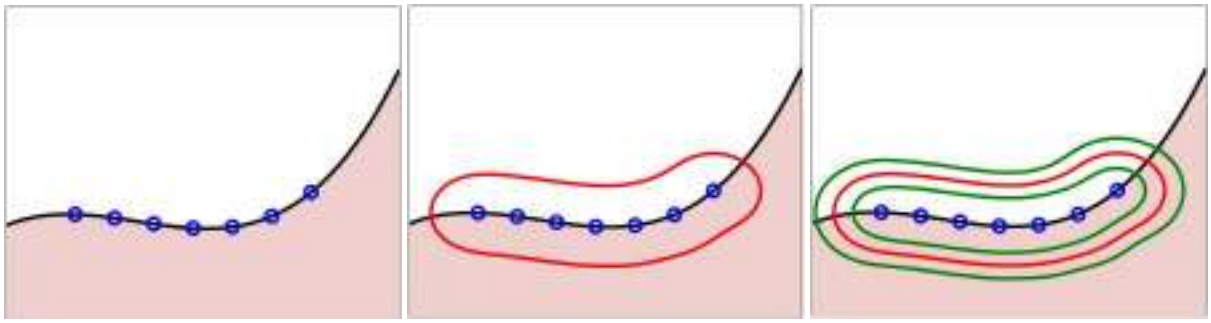

(a) Sensed points (blue circles) on the surface of an area of anatomy which is to be constrained.

(b) Constraint isosurface (red) constructed from the sensed points using (5.2).

(c) The level sets (green) of the constraints isosurface which will be used for control.

Figure 5.5: A demonstration of how the implicit surface is generated from the sensed anatomy points and the level sets which will be used for control.

be computed for a chosen radius of effect, R, by setting the threshold equal to,

$$T = \psi_{4,2}(r_{min}) \tag{5.3}$$

Figure 5.5 summarises the dynamic active constraints strategy thus far. Points are sensed on the anatomy to be constrained, using the surgery appropriate sensing modality in Figure 5.5a. The constraint isosurface is then generated from these sensed points using (5.2), as demonstrated in Figure 5.5b. In the next sections, the level sets to the constraint isosurface shown in Figure 5.5c will be used to control the tool point to the constraint surface.

5.3 The Implicit Function Jacobian

Equation (5.2) defines a time-deforming implicit constraint surface which is to serve as the motion constraint for the robot. To control the motion of the robot in relation to the surface, I will now derive a Jacobian which defines the normal of the current level set of the implicit constraint surface which the robot tool point currently lies on. We have named this the Implicit Function Jacobian. The following derivation requires the constraint isosurface, f, to be continuous and $\frac{df}{dt}$ and $\frac{d^2f}{dt^2}$ to exist and be continuous. This condition is satisfied by $\psi_{4,2}$ and this control methodology will work for any such constraint surface which satisfies these conditions.

To compute the Implicit Function Jacobian, the time derivative of the function which defines

the constraint surface is first found.

$$\frac{df}{dt} = \frac{\partial f}{\partial t} + J_f \frac{dx}{dt} = 0 \tag{5.4}$$

where $\frac{\partial f}{\partial t} \in \mathbb{R}^p$ is the partial derivative of f explicitly with respect to time,

$$J_f = \begin{bmatrix} \frac{\partial f}{\partial x_1} & \cdots & \frac{\partial f}{\partial x_p} \end{bmatrix} \tag{5.5}$$

is the Jacobian of f with respect to the spatial coordinates, and

$$\frac{dx}{dt} = \dot{x} = \begin{bmatrix} \dot{x}_1 \\ \vdots \\ \dot{x}_p \end{bmatrix} \tag{5.6}$$

is the vector of spatial variable velocities.

Since, in general, the spatial variables cannot be controlled directly, the Jacobian matrix which relates the joint velocities to the spatial variable velocities will be used.

$$\dot{x} = J\dot{q} \tag{5.7}$$

Substituting (5.7) into (5.4), rearranging, and normalising yields,

$$\frac{J_f J}{\|J_f J\|}\dot{q} = -\frac{1}{\|J_f J\|}\frac{\partial f}{\partial t} \tag{5.8}$$

$\frac{J_f J}{\|J_f J\|}$ represents a relation between joint velocities and the velocity along the normal to the level set of the constraint function that the robot end effector currently lies on. That is, if the current position of the tool point is x_c such that $f(x_c, t) = c$ for some constant c, then $\frac{J_f J}{\|J_f J\|}\dot{q}$ will be equal to the velocity of the constraint function in the direction of the gradient at the end effector. When the tool point lies in the zero set of the constraint function $\{x | f(x,t) = 0\}$ and, as such, lies on the desired constraint surface, the velocities produced will be along the

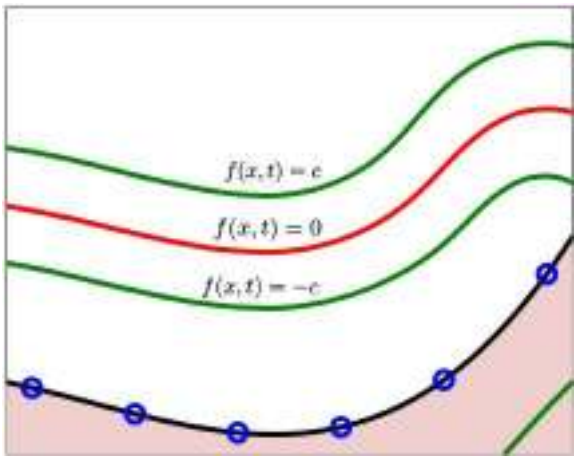

(a) The sensed points (blue), generated constraint isosurface (red), and level sets (green) overlaid on the anatomy which is to be protected from the tool.

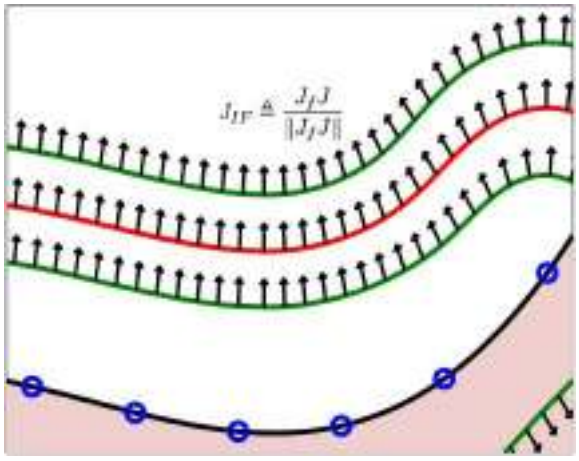

(b) The normals to the level sets and zero set of the constraint function overlaid on the previous diagram, which are used to control the tool point to the zero set.

Figure 5.6: An explanatory diagram of the level sets of the constraint isosurface and the constraint level set normals which are used to guide the tool to the zero set.

normal to the desired constraint surface. The Jacobian that represents this relationship, I will define to be the Implicit Function Jacobian.

$$J_{IF} \triangleq \frac{J_f J}{\|J_f J\|} \tag{5.9}$$

To more clearly illustrate these concepts, the level sets to the constraint isosurface which correspond to where the surface equals a constant are shown in Figure 5.6a. The normals to the level sets of the forbidden region boundary in joint space which are defined by the Implicit Function Jacobian are shown Figure 5.6b.

$\dot{x}_{IF}$ is defined to be the velocity of the constraint function in the direction of the normal to the current level set.

$$\dot{x}_{IF} \triangleq J_{IF}\dot{q} \tag{5.10}$$

The right hand side of (5.8) is the actual velocity of the surface in the direction of its normal. Perfect tracking of the time variations of the surface requires controlling the velocity of the robot's end effector to equal this value in the direction of the normal to the constraint surface.

Therefore, the desired constraint surface normal velocity, $\dot{x}_{IF}^{des}$, is defined to be equal to this value.

$$\dot{x}_{IF}^{des} \triangleq -\frac{1}{\|J_f J\|}\frac{\partial f}{\partial t} \tag{5.11}$$

5.4 Dynamic Active Constraints Using the Implicit Function Jacobian

In order to control the robot to lie on the desired constraint surface, the Implicit Function Jacobian has been incorporated into the operational space approach [Khatib, 1987] and its recursive null-space extension [Sentis, 2007]. The operational space approach transforms the joint dynamics into the dynamics of task space in order to create controllers which are linear in the accelerations of the chosen task.

The dynamics of a robot with n joints can be expressed in the Lagrangian formalism as

$$A(q)\ddot{q} + b(q,\dot{q}) + g(q) = \Gamma \tag{5.12}$$

where A is the inertia matrix, b is the Coriolis-centrifugal vector, g is the gravity vector, and Γ is the generalised force vector.

For some arbitrary task $x_t = T(q)$ and its instantaneous kinematics $\dot{x}_t = J_t\dot{q}$ (where J_t is the task Jacobian), the task space dynamics can be computed by

$$\overline{J}_t^T(A\ddot{q} + b + g = \Gamma) \Rightarrow \Lambda_t\ddot{x}_t + \mu_t + p_t = F_t \tag{5.13}$$

where $\overline{J}_t = A^{-1}J_t^T\Lambda_t$ is the dynamically consistent inverse of J_t, $\Lambda_t = (J_tA^{-1}J_t^T)^{-1}$ is the operational space inertia matrix, $\mu_t = \overline{J}_t^T b - \Lambda_t\dot{J}_t\dot{q}$ is the operational space Coriolis-centrifugal vector, $p_t = \overline{J}_t^T g$ is the operational space gravity vector, and F_t is the vector of operational space forces.

By appropriately choosing $F_t = \Lambda_t \ddot{x}_t^{ref} + \mu_t + p_t$, where $\ddot{x}_t^{ref}$ is the reference acceleration control input, the dynamics of the robot can be linearised in the task space. These operational space forces are mapped back to actuation torques through the transpose of the Jacobian.

$$\Gamma = J_t^T F_t \tag{5.14}$$

Using the operational space technique, I can transform the joint space dynamics into the space of the constraint using J_{IF}.

$$\Gamma = J_{IF}^T F_{IF} = J_{IF}^T (\Lambda_{IF} \ddot{x}_{IF}^{ref} + \mu_{IF} + p_{IF}) \tag{5.15}$$

where Λ_{IF} is the implicit function space inertia matrix, μ_{IF} is the implicit function space Coriolis-centrifugal vector, and p_{IF} is the implicit function space gravity vector.

Next, suitable low level acceleration controllers which will control the robot's end effector to lie on the desired surface, must be defined. Two cases arise when constraining the robot; constraint surfaces the motion of which is known or predicted, and constraint surfaces which are static or the motion of which is unknown.

The motion of known time-varying surfaces can be tracked by using a feed-forward controller which compensates for the normal velocity and acceleration of the surface.

$$\ddot{x}_{IF}^{ref} = \ddot{x}_{IF}^{des} + K_p \left(f_{set} - f(x_c, t_c) \right) + K_d (\dot{x}_{IF}^{des} - J_{IF} \dot{q}) \tag{5.16}$$

where

$$\ddot{x}_{IF}^{des} \triangleq -\frac{d}{dt} \left(-\frac{1}{\|J_f J\|} \frac{\partial f}{\partial t} \right) \tag{5.17}$$

is the acceleration of the surface in the direction of its normal, x_c is the current tool position, t_c is the current time, f_{set} is the desired level set to control to, K_p is a diagonal matrix of position gains, and K_d is a diagonal matrix of velocity gains.

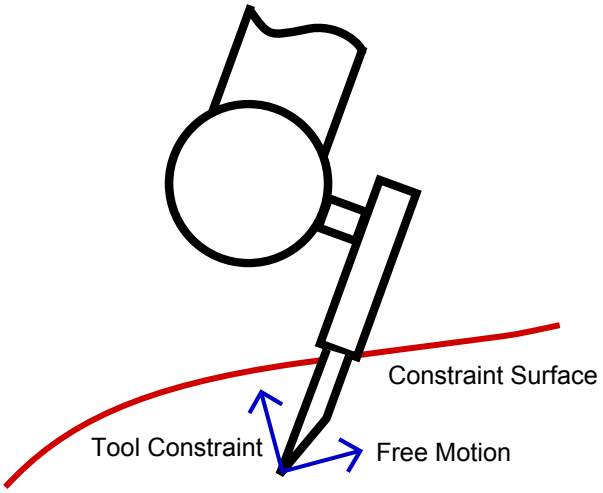

Figure 5.7: Diagram demonstrating the desired behavior of the tool point constraint; to apply forces to move the user and robot out of the constraint while allowing free motion tangent to the constraint isosurface.

If the time deformations of the isosurface are unknown, the current estimate of the constraint can be controlled to using the following controller.

$$\ddot{x}_{IF}^{ref} = K_p \left(f_{set} - f(x_c, t_c) \right) - K_d J_{IF} \dot{q} \tag{5.18}$$

In addition to controlling the robot to lie on the constraint surface, the hands-on controller should provide unimpeded motion along the surface, which is illustrated in Figure 5.7, and provide a posture for the remaining degrees of freedom. This is accomplished with the recursive extension to the operational space approach [Sentis, 2007], which allows for a hierarchy of an arbitrary number of tasks to simultaneously operate. Each task in the hierarchy operates in the degrees of freedom left over from higher order tasks by operating in the null space of the previous Jacobians. This controller takes the form,

$$\Gamma = \sum_{k=1}^{N} J_{k|prec(k)}^{T} F_{k|prec(k)} \tag{5.19}$$

where $J_{k|prec(k)} = J_k N_{prec(k)}$ is the Jacobian of the k-th task operating in the null space of the previous $k-1$ tasks, $N_{prec(k)} = I - \sum_{k=1}^{N} \overline{J}_{k|prec(k)} J_{k|prec(k)}$ is the combined null space of the higher order tasks, and $F_{k|prec(k)}$ is the forces of the k-th task acting in the null space of the previous

$k-1$ tasks.

The three-task prioritised controller which hierarchically constrains the robot, allows motion along the surface of the constraint, and controls the residual degrees of freedom is

$$\Gamma = J_{IF}^T F_{IF} + J_{pt|IF}^T F_{pt|IF} + J_{p|pt|IF}^T F_{p|pt|IF} \tag{5.20}$$

where J_{pt} is a two degree of freedom Cartesian velocity Jacobian and J_p is the Jacobian of the posture. $F_{pt|IF}$ controls the forces in the Cartesian directions along the surface and $F_{p|pt|IF}$ controls the forces in the residual posture.

While the above controller directs the user to the constraint isosurface, outside the constraint, the user should be unimpeded in moving the tool. Therefore, when away from a constraint surface, the following controlle,r which consisted of a Cartesian position controller and a joint controller in the null space, should be used.

$$\Gamma = J_{pt}^T F_{pt} + J_{p|pt}^T F_{p|pt} \tag{5.21}$$

where J_{pt} is the Cartesian velocity Jacobian, F_{pt} is the Cartesian force, $J_{p|pt}$ is the joint or postural Jacobian acting in the null space of the Cartesian Jacobian, and $F_{p|pt}$ is the force of the posture acting in the null space of the Cartesian velocity Jacobian.

When the user approaches the surface with the end effector and reaches the zero set, the controller is switched to the surface constraint controller (5.20). Similarly, when the user moves the tool away from the surface, the controller is switched back to (5.21).

$$\Gamma = \begin{cases} J_{IF}^T F_{IF} + J_{pt|IF}^T F_{pt|IF} + J_{p|pt|IF}^T & f(x_c) \leq 0 \\ J_{pt}^T F_{pt} + J_{p|pt}^T F_{p|pt} & f(x_c) > 0 \end{cases} \tag{5.22}$$

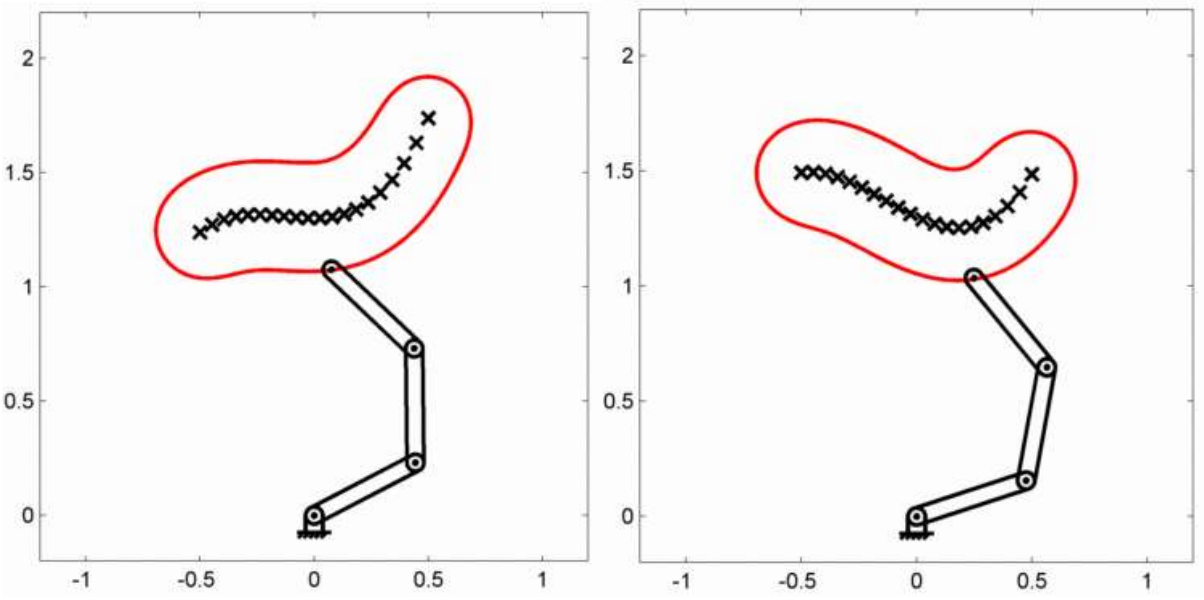

(a) (b)

Figure 5.8: A simulated four degree-of-freedom planar robot being constrained to an implicit surface (red), which was generated from unpredicted constraint points (black crosses) while moving sinusoidally in the horizontal direction. Two points in time are depicted: (a) at $t = 2.5s$ and (b) $t = 12.5s$

5.5 Results

5.5.1 Surface Tracking Simulation

A 2D dynamic simulation was implemented in Matlab R2012a (Mathworks Inc.) to test our control algorithm. To mimic a scenario in which the soft tissue deformation was not periodic or predictable, the controller was first tested using constraint points for which the motion was unknown a priori and the position of the points was updated at 40 Hz to simulate stereo camera data rates. In addition to constraining the end effector to the isosurface, to simulate the surgeon moving the end effector, a feed forward controller was used to control the horizontal direction of the position of the end effector to follow a sinusoid. Lastly, the residual posture was controlled to the zero position.

Figure 5.8 depicts the simulated robot kinematic chain, constraint points and constraint surface at two time intervals during the simulation. In Figure 5.9a, the error in the level set of the constraint over time is shown. Due to the discrete nature of the range data, at every update of the camera, the error in the position constraint jumps to a non-zero value and the controller

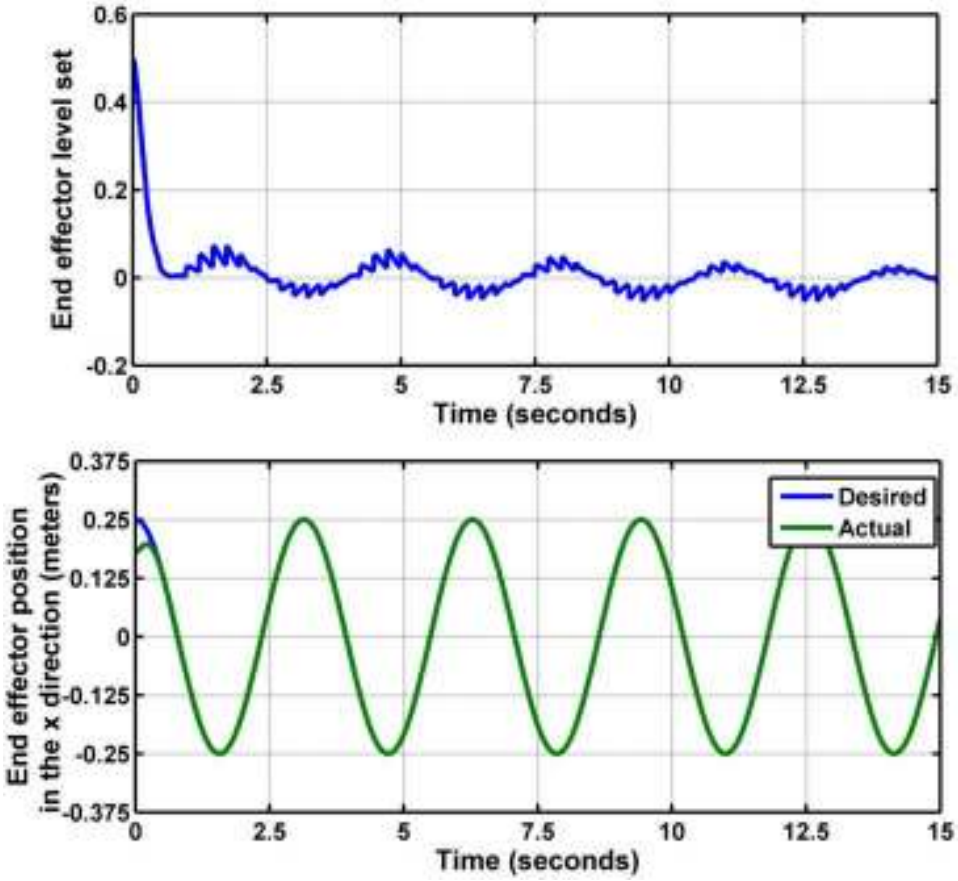

Figure 5.9: Simulation results for the unpredicted surface demonstrating the controller performing gradient descent to the zero set of the end effector after each constraint point update (a) and the horizontal tracking of a sinusoid by the end effector (b).

performs gradient descent along the level sets of the implicit constraint function. As the positioning of the end effector in the horizontal direction functions in the null space of the constraint task, this gradient descent does not affect the horizontal trajectory and the secondary task controller is still able to track the desired trajectory as shown Figure 5.9b. This demonstrates that the primary constraint task does not affect the surgeon's ability to maneuver the tool along the surface, even in the presence of errors in the primary task due to point cloud updating.

To evaluate the proposed strategy in a scenario where the motion of the constraint points is known or predicted, a 2D simulation was created for testing the controller operating on a constraint surface in which the motion of the points was known. Each constraint point's x and y position were defined to be known sinusoids. Similarly to the previous simulation, the motion of the surgeon was mimicked by a sinusoid trajectory commanded on the horizontal direction of the end effector, and the posture was commanded to the zero position.

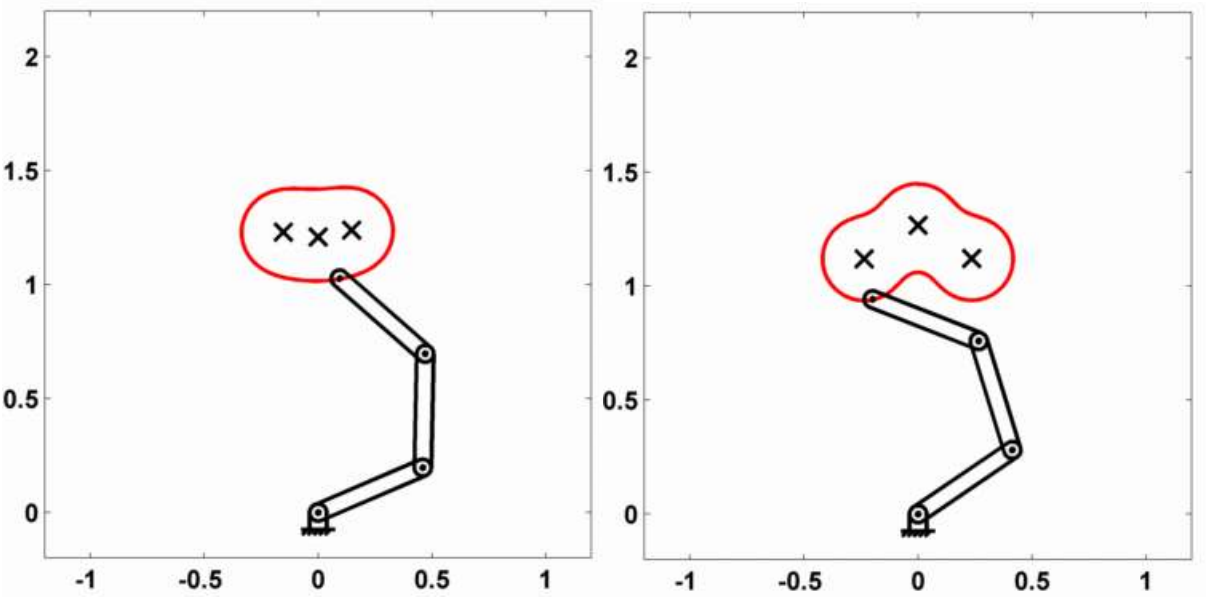

(a) (b)

Figure 5.10: A simulated four degree-of-freedom planar robot being constrained to an implicit surface (red) which was generated from constraint points whose motion was known (black crosses), while moving sinusoidally in the horizontal direction. Two points in time are depicted: (a) at $t = 2.5s$ and (b) $t = 7.5s$

Figure 5.10 shows the robot and constraint at two different times during the simulation. Figure 5.11a demonstrates that the controller performs an approximately critically damped response to a level set of zero while Figure 5.11b shows that the controller simultaneously tracks the desired position trajectory. The prediction of the point cloud motion allows the velocity and acceleration in the normal direction of the level sets of the implicit surface to be computed. Utilising this information in the controller allows the robot to perfectly track the motion of the constraint surface. Additionally, as in the previous simulation, the control actions of the primary constraint task do not affect the simulated surgeon's motions in the null-space. These results are dependent on the accuracy of the prediction.

5.5.2 Experimental Setup

The experimental setup used to verify the proposed control strategy in a physical setup is shown in Figure 5.12. A Kuka LWR 4+ was controlled through the Fast Research Interface (FRI, Kuka Robotics GmbH) with a PC running Ubuntu 10.04, OROCOS (orocos.org) and ROS Electric (ros.org). Torques were commanded to the FRI operating in the axis-specific impedance control

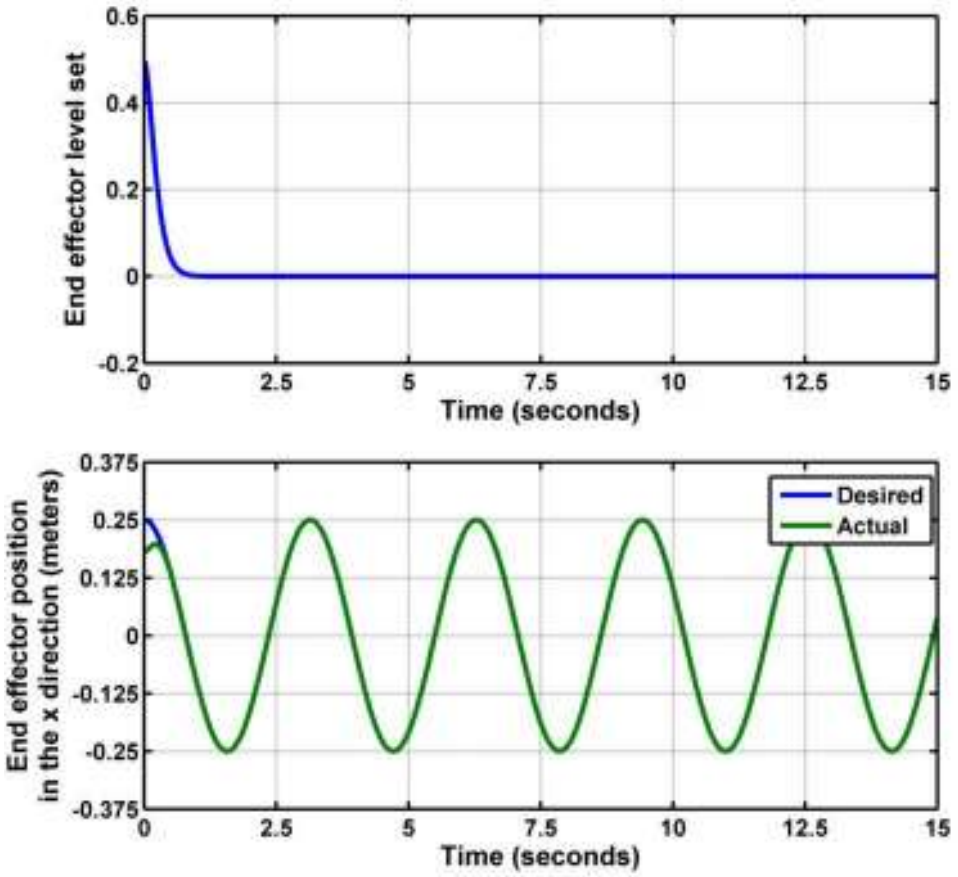

Figure 5.11: Simulation results for the predicted surface demonstrating the controller performing gradient descent to the zero set of the end effector after each constraint point update (a) and the horizontal tracking of a sinusoid by the end effector (b).

mode with the virtual spring stiffness and damping parameters set to zero. We modified the open source Whole-Body Control Software to create our controllers [Philippsen et al., 2011]. A dynamic soft tissue brain phantom made of Platsil Gel was created and actuated using a linear actuator controlled by an Arduino Uno. An Optotrak Certus (Northern Digital Inc.) was used to track two markers attached to the outer surface of the brain phantom and the generated point clouds were sent to the Controller via the User Datagram Packet over Internet Protocol (UDP/IP). The surface variations were approximately 1cm in amplitude and had a frequency of about 0.45 Hz.

In both controllers of equation (5.22), the Cartesian gains were set to zero to allow for the user to freely control the tool in Cartesian space and along the surface. The acceleration level postural controllers were set to damp out any motion in the posture, allowing the user to maneuver the postural degrees of freedom if necessary, while ensuring they approximately held their position

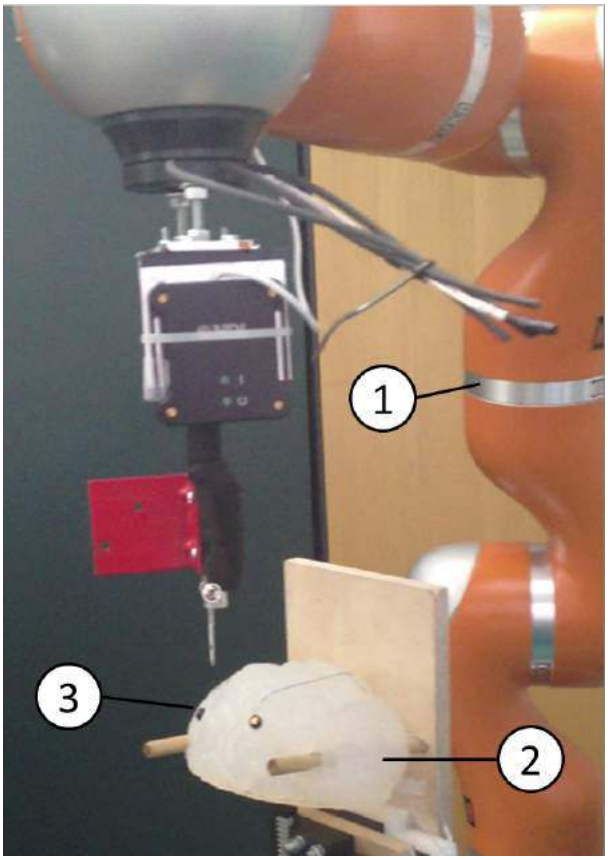

Figure 5.12: The experimental setup consisting of a Kuka LWR 4+ (1), a dynamic soft tissue phantom (2) and Optotrak markers (3)

while the user was controlling the end effector. Additionally, as opposed to the simulations, the Coriolis-centrifugal forces were not compensated for on the LWR. However, due to the low velocity nature of these experiments, this did not affect the results significantly. To more clearly demonstrate the effect of constraining to surface, the level set threshold for switching between the constrained and unconstrained controllers was raised to two from the zero set. This means that for a small region outside of the constraint boundary, the controller will direct the motion of the tool to the constraint surface.

The experiment aimed to demonstrated the ability of the constraint to actively limit the motion of the tool when near the constrained region. The tool began approximately 10cm away from the brain surface and the user attempted to contact the tool with one of the markers on the brain while the dynamic active constraints strategy was applied to prevent this from occurring.

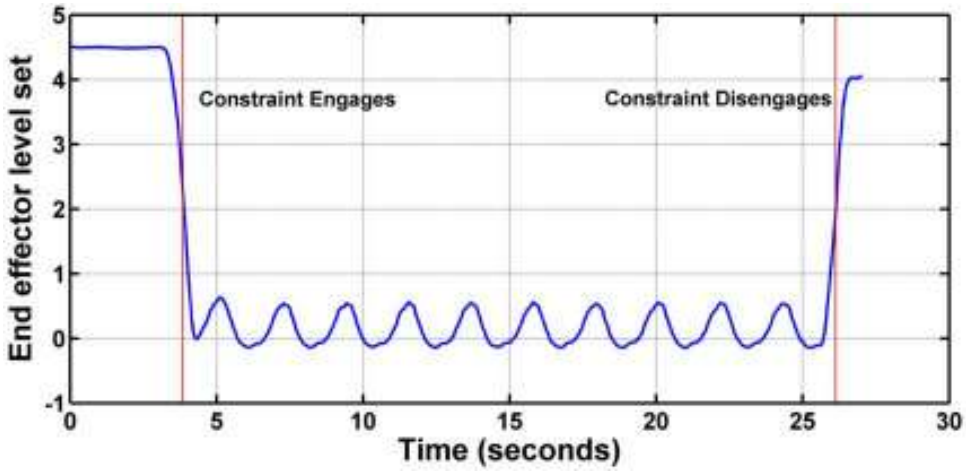

Figure 5.13: Dynamic phantom results depicting level set of the end effector over the trial.

5.5.3 Dynamic Brain Phantom Experimental Results

Figure 5.13 shows the error to the constraint surface during the experiment utilising the phantom and Figure 5.14 shows the magnitude of force the user applied to the robot during the test. At the start of the experiment, the tool was away from the surface and no torques were commanded to constrain the tool. The user moved the tool towards one of the markers on the moving constraint brain and at time $t = 3.85s$, the tool passed the level set limit, engaging the constraint and applying forces to limit the user's motion. At approximately 25.5 seconds, the user exerted a force on the tool away from the constraint and at $t = 26.15s$ the constraint disengaged, allowing unimpeded movement of the end effector. More detailed experimental results will be presented in the following chapter, however, the results here demonstrate the feasibility of the proposed controller to run in real-time on a seven degree-of-freedom manipulator. In addition, the controller successfully prevented the user from colliding with the dynamic brain phantom.

5.6 Summary

In this chapter, a dynamic active constraints approach for hands-on robotic surgery has been demonstrated, which utilises point clouds to generate an implicit surface and a multi-priority null-space controller to constrain the tool point. The ability to apply the technique to point

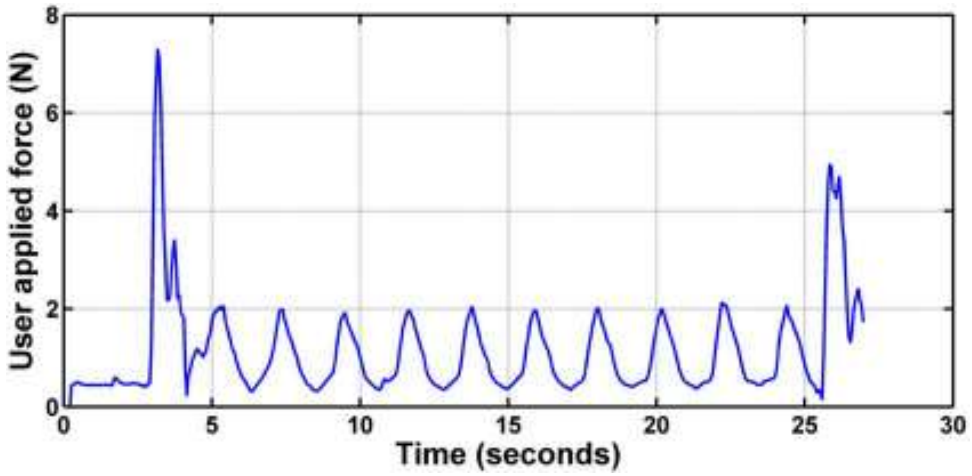

Figure 5.14: Magnitude of user applied forces during dynamic phantom tests.

clouds allows the strategy to be utilised with various sensing modalities which can, in general, be reduced to such a representation, thereby expanding the potential usage of the technique to various surgeries. The null-space controller incorporates the dynamics of the robot into the enforcement of the constraint in order to provide an appropriate impedance to the surgeon's wrench. Additionally, the controller allows for free motion tangent to the constraint and Chapter 6 will demonstrate optimisations in the remaining postural degrees-of-freedom.

Simulation work demonstrated the accuracy of the proposed strategy in constraining the robot to the surface of the constraint and allowing free motion tangent to the constraint isosurface for point clouds the motion of which was either known or unknown. Experiments with a lightweight seven degree-of-freedom redundant robot on a dynamic brain phantom showed the feasibility of the approach in a physical setup. In the next chapter, more thorough experiments will establish the effectiveness of the proposed dynamic active constraints technique in a virtual reality ultrasound scenario. Additionally, the dynamic active constraints methodology developed here will be shown working concurrently with the natural motion controller developed in Chapters 3 and 4 for minimising the effects of the end effector dynamics on the surgeon.

Chapter 6

Virtual Reality for Testing Shared Control

This chapter presents an immersive virtual reality setup for testing shared control strategies for hands-on robotic surgery on dynamic soft tissue. Preliminary controls research on soft tissue requires a suitable testing environment, and developing accurate phantoms which mimic the deformation of real tissue can be difficult. We propose a setup in which the robot is registered with a virtual environment such that the robot's position in the real world aligns with that of the virtual environment, preserving the complex dynamic interaction between the surgeon, the control strategy, and the robot. The motion of soft tissue can be simulated and interaction forces and torques can be applied to the robot. We develop a chest ultrasound scenario for diagnosing pleural effusion and pneumothroax within this virtual environment and use it to demonstrate the simultaneous application of the dynamic active constraints strategy developed in Chapter 5 and the natural motion technique of Chapters 3 and 4. In addition, I show the efficacy of the dynamic active constraints in improving the user's accuracy and ensuring safety for the patient when performing ultrasound scans.

6.1 Introduction

The previous chapters developed shared control strategies to ensure a more natural motion for the surgeon and to protect the patient from tool movement into dangerous deforming regions during surgery. This chapter aims to show the concurrent application of these strategies and to demonstrate the efficacy of the dynamic active constraints approach developed here in a scenario which closer to clinical application. To do so, an environment which allows for the testing of the complex dynamic interactions between the surgeon, the robot, the controller, and the soft tissue is required. *In vivo* clinical trials of surgical control strategies provide direct evidence as to whether a shared control strategy is effective. However, there are safety concerns, particularly in the preliminary stages of the controller's development, where the controllers have not demonstrated stability when interacting with complex deforming soft tissue. To allow for testing surgical control strategies in a safer manner, soft tissue phantoms have been developed. These phantoms aim to replicate the material properties of various organs so that the synthetic soft tissue responds to tool interactions such as indentation and cutting in a similar way to the real tissue and/or results in a comparable image to sensing technologies as scanning real tissue [D'Souza et al., 2001] [Chen et al., 2010] [Faulkner et al., 2015]. Additionally, some phantoms mimic the deformations of the soft tissue in response to heart beat and breathing [Cattilino et al., 2014], and gravity [Reinertsen and Collins, 2006] [Forte et al., 2014].

However, it is difficult to design a phantom which can simultaneously accomplish multiple properties (i.e. appropriate indentation response and cutting response) [Lamouche et al., 2012]. Therefore, soft tissue phantoms still tend to be limited to simpler scenarios. In addition, soft tissue phantoms must be replaced when excessive indentation or cutting has occurred to allow for additional testing or training.

As an alternative to soft tissue phantoms, a virtual reality environment has been developed for testing shared control strategies in hands-on robotic surgery. Like soft tissue phantom solutions, it provides a safe alternative to preliminary testing control strategies *in vivo*, however, the virtual environment is more flexible, as different soft tissues and surgeries can be more easily adapted and tested. Additionally, the virtual environment can be easily reset, which allows for repeatable testing and training, whereas phantoms must be recreated.

VR environments have been used previously to test and train surgeons in laparoscopic surgery [Sun et al., 2007] [Nagendran et al., 2013] [Liu and Curet, 2015]. Surgical trainees who practice in virtual reality have been shown to have improved abilities over those who trained using standard programmatic training methods in laparoscopy [Jordan et al., 2001] [Seymour et al., 2002] demonstrating the value of VR training environments. However, to our knowledge, this is the first application of VR technology to hands-on robotic surgery. Additionally, the VR concept of co-location [Swapp et al., 2006], where sensory cues such as haptics and audio, are aligned in both the virtual and real space, has been used to give the user a better sense of presence. The physical robot is registered and aligned with the virtual scene such that the user can grasp the real robot and see the corresponding virtual copy in the same location. This preserves the interaction between the robot, controller, and user while the soft tissue interaction can be simulated. To our knowledge, this environment represents the first VR surgical theater to co-locate the full robot kinematics within the environment, thereby improving immersion.

Immersion in VR is the sense of being physically present within a virtual environment [Pausch et al., 1997]. The level of immersion experienced by the user is dependent on both the capabilities of the hardware (resolution, field of view, refresh rate, etc.) and the detail of the virtual environment itself [Mahalil et al., 2014]. Training in a more immersive environment, one which is closer to the clinical scenario, "can help surgeons develop high-level skills while also reducing stress and improving their confidence for carrying out those skills in the real environment, with real patients." [Dargar et al., 2015].

A similar concept to virtual reality is augmented reality (AR), which aims to provide additional information to the user by overlaying registered, real-time 3D images onto the real world. This type of strategy has been used in surgery to display sensing information to the surgeon about the underlying anatomy, such as vasculature, and other information, such as pre-operative planning [Alaraj et al., 2011] [Zhu et al., 2014]. AR can be used in conjunction with robotic assistance, such as active constraints, or alone, when assisting the surgeon. Commercially viable AR solutions which are accurate and compact in design remain in development due to issues in both software and hardware. The AR device must be registered in real-time with the environment at sufficient precision and occlusion, where real world objects pass in front of virtual objects, must be handled appropriately. AR hardware must be able to present clear

images at a high refresh rate and have the computational capacity to perform the registration and handle occlusion, yet be ergonomic and light enough to be worn by the user comfortably. At present, however, a VR environment can be used to determine the viability of complex AR strategies. By testing AR strategies in virtual reality, the primary development issues in AR are removed and AR's efficacy in providing useful information to the user can be more directly assessed, towards future implementation in real surgery.

This chapter will first discuss the proposed scenario, ultrasound (US) for diagnosing pleural effusion and pneumothorax, which will be used to validate our shared control strategies. The technical details of the implementation of the virtual surgical environment will then be discussed, followed by the experimental design, which compares performing the US scans with no assistance, augmented reality assistance, and robotic assistance through dynamic active constraints with augmented reality. The results demonstrate an improvement in safety and accuracy when using robotic assistance through dynamic active constraints.

6.2 Simulated Ultrasound for Diagnosing Pleural Effusion and Pneumothorax

Limitations in the accuracy of the Kuka, in addition to limited resources for creating a high fidelity surgical simulation, made finding a suitable clinical scenario difficult to identify. The medical scenario I will use to validate our shared control strategies within the virtual surgical environment is pleural effusion and pneumothorax diagnosis with US. Pleural effusion is an accumulation of excess fluid in the pleural cavity, the fluid-filled region surrounding the lungs. Similarly, pneumothorax is an accumulation of air within this space. These conditions interfere with breathing by restricting the expansion of the lungs and can be caused by physical trauma to the chest or a variety of medical conditions.

Ultrasound is used in determining if pleural effusion or pneumothorax have occurred in the patient and the locations in which fluid has accumulated. It is also used in guiding the needle for thoracentesis, which is the procedure to remove the excess fluid or air resulting from these conditions. Ultrasound is able to better detect pleural effusion than X-ray radiology, including

in cases where there is a very small amount of pleural effusion [Gryminski et al., 1976]. The performance of US in detecting pleural effusion is comparable to that of computerised tomography (CT) [Yu et al., 1993], with the benefit of being radiation free and portable, as well as providing real-time imaging [Beckh et al., 2002]. Performing ultrasound at the bedside is faster and safer than transporting the patient to a CT scanner, particularly in patients in potentially critical condition [Feller-Kopman, 2007]. Lastly, US reduces complications by improving accuracy in needle placement when performing thoracentesis as compared to no US guidance and, consequently, is lower cost due to the lessened length of stay resulting from fewer complications [Patel et al., 2012] [Mercaldi and Lanes, 2013].

However, US is a very operator dependent procedure as image acquisition and interpretation occur simultaneously, requiring a sufficiently high level of skill and concentration [Havelock et al., 2010]. To achieve a good scan for pleural effusion or pheumothorax, the US probe must be positioned between the ribs and orientated such that the scanning plane is directed between ribs and the probe is perpendicular to the skin [Doelken and Mayo, 2008]. Additionally, pressure must be applied in the direction perpendicular to the skin and the magnitude of this force depends on the patient's body mass index [Dhyani et al., 2014]. Studies have shown that if force application is too high, musculoskeletal injuries can occur [Mirk et al., 1999].

Due to these requirements, ultrasound requires a training program supervised by an experienced professional and a program for maintaining these skills. The minimum standards for chest ultrasound training in the UK have been established by the Royal College of Radiologists [The Royal College of Radiologists, 2012]. To achieve the most basic level of competency, Level 1, trainees require one session of at least 5 ultrasound scans per week for 3 months supervised by a doctor with Level 2 competency or 2 years of experience at Level 1. In addition, to maintain Level 1 competency, 20 exams must be performed per year, with no more than 3 months of time lapsing between scans.

This chapter aims to demonstrate that robotic assistance can be used to improve the accuracy and speed of performing ultrasound scans, leaving the surgeon able to better focus on interpreting the ultrasound scans. By applying dynamic active constraints at the target points, the user can feel when he or she is applying too much force, and constraint forces can be applied to

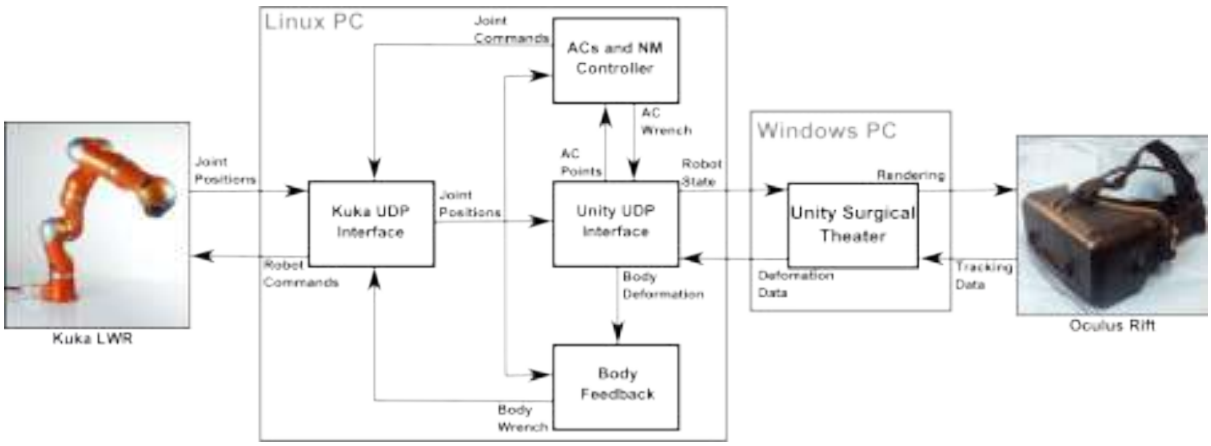

Figure 6.1: System diagram of the hands-on robotic surgery virtual environment.

ensure the user's forces do not reach levels that can harm the patient. Additionally, constraining the orientation of the tool simplifies the operator's control of these degrees-of-freedom, thereby improving accuracy.

6.3 Virtual Surgical Environment

6.3.1 System Overview

The virtual reality surgical environment setup used for the ultrasound scenario is shown in Figure 6.1. In this system, the physical robot, a Kuka LWR 4+ controlled by a Linux PC, is registered and aligned with a virtual surgical environment running on a Windows PC and rendered on an Oculus Rift (Oculus VR, LLC.) head-mounted display, such that the physical robot's position in the real world matches that in the virtual one. In doing so, the user can interact with the real robot, allowing for the complex dynamic interaction between the user, the control strategy, and the robot to take place, while the user is immersed in the virtual surgical theater. In addition, the motion of the soft tissue can be simulated and the interaction forces and torques between the robot and this tissue can be applied to the robot.

The Linux PC (i7-3770 @ 3.40 GHz), which communicated with the Kuka LWR 4+ and computed the tissue interaction wrenches and the DACs and natural motion control commands, ran Ubuntu 12.04 (www.ubuntu.com), with the Xenomai real-time kernel patch (www.xenomai.org), RTnet for real-time networking (www.rtnet.org), and OROCOS for real-time processing com-

Figure 6.2: Visualisation of the overall surgical theater developed in Unity.

ponents (www.orocos.org). The Windows PC (i7-3770 @ 3.40 GHz), which rendered the virtual scene to the Oculus Rift and deformed the soft tissue, utilised Unity 4.6 (Unity Technologies ApS), a game creation engine.

The following subsections provide more detail on the individual components and their implementation.

6.3.2 Virtual Surgical Theater

The virtual surgical theater, shown in Figure 6.2, was constructed using Unity 4.6, using 3D models created in Solidworks (Dassault Systemes S.A.) and Blender (Blender Foundation), and free models from 3D Warehouse (3dwarehouse.sketchup.com) and Archive3D (archive3d.net). The patient's body was created using MakeHuman (makehuman.org) and breathing motions were simulated by applying a sinusoidal scaling of the chest at a frequency of 0.2Hz, which is within resting adult respiratory rates [Lindh et al., 2009]. The peak motions in the patient's breathing are shown in Figure 6.3.

Within the surgical environment, a patient monitoring system, shown in Figure 6.4, was created, consisting of a heart and blood pressure monitor, a visualisation of the current US image, and an indicator of the desired points to be scanned. US scans were blurred based on the error in

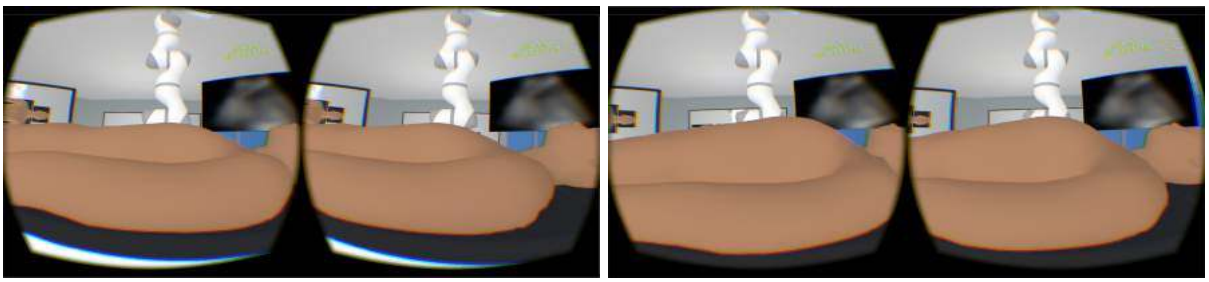

(a) Minimum point in the breathing motion. (b) Maximum point in the breathing motion.

Figure 6.3: Motion of the virtual patient's chest during breathing.

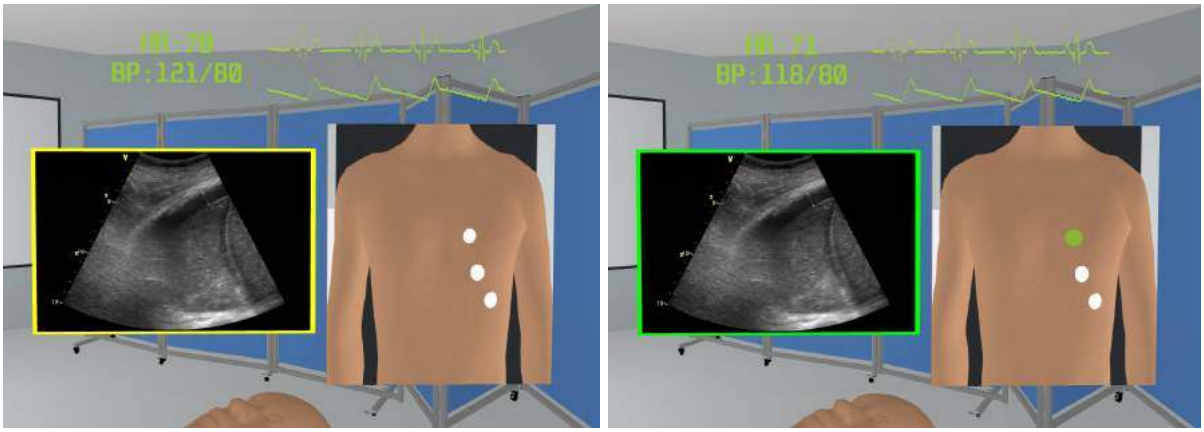

(a) (b)

Figure 6.4: The patient monitoring system used to inform the user about the status of the US.

the distance and orientation between the ultrasound probe and the target points to simulate the effects of misalignment of the US probe. As the user moved the tool closer to the correct position and orientation, the US image would increase in clarity.

The images were pre-blurred to reduce the computational load during run-time. From an original clear US image, increasingly blurred images were progressively generated by averaging the RGB values of the pixels within a set radius of each pixel. Within the simulation, the appropriate US image was chosen based on the following formula, which normalises and combines the position and orientation errors:

$$US_{index} = \begin{cases} 0 & : d < d_{min} \ \& \ \theta < \theta_{min} \\ n_{max} - 1 & : \frac{(d-d_{min})^2}{(d_{max}-d_{min})^2} + \frac{(\theta-\theta_{min})^2}{(\theta_{max}-\theta_{min})^2} > 1 \\ \lfloor (n_{max}-1)\left(\frac{(d-d_{min})^2}{(d_{max}-d_{min})^2} + \frac{(\theta-\theta_{min})^2}{(\theta_{max}-\theta_{min})^2}\right) \rfloor & : else \end{cases} \tag{6.1}$$

where d is the L^2 distance between the closest target point and the tool point, θ is the angle between the normal to the closest target point and the tool, n_{max} is the total number of images, d_{min} and θ_{min} are the error bounds for a successful scan and d_{max} and θ_{max} are the upper error bounds on the blurring. The parameters used in the experiments here are shown in Table 6.1.

Parameter	d_{min}	d_{max}	θ_{min}	θ_{max}	n_{max}
Value	$0.01m$	$0.05m$	$15°$	$30°$	21

Table 6.1: Parameter values used in simulating the blurring of the US due to misalignments in position and orientation between the US probe and the target points.

When the US probe was positioned such that the distance and orientation errors, d and θ, were below the set thresholds, d_{min} and θ_{min}, the border surrounding the US image turned yellow, as shown in Figure 6.4a. After $3.5s$, the border was turned green and a marker was placed on the indicator board, as seen in Figure 6.4b, to indicate to the user that a successful scan had been completed.

6.3.3 Head-Mounted Display

The virtual reality head-mounted display used in this setup was the Oculus Rift Development Kit 2 (Oculus VR, LLC.) shown in Figure 6.5. The Rift utilises an IR camera to track LEDs embedded in the headset in order to accurately update the position and orientation of the viewpoint cameras within the virtual scene. To further ensure that the headset visualisation matches the motion of the user's head and to reduce motion sickness caused by a mismatch between head motion and the camera poses, the Rift requires the visualisation to be updated at a fixed rate of 75Hz.

Figure 6.5: The Oculus Rift head-mounted display and tracking camera.

Figure 6.6 depicts the first person view of the surgical theater, seen through the Rift when using the system. Current solutions for tracking hand configurations, such as the Leap Motion Controller (Leap Motion Inc.), have difficulties in tracking the hand when grasping objects. Therefore, I utilised a simplified solution to improve immersion in which I assumed the user was always reaching for the tool. The user's virtual arm was attached to the Oculus Rift camera in approximately the shoulder position and the inverse kinematics of the arm joints were solved such that the hand grasped the tool. When the tool point was out of workspace of the arm, the arm was set to be reaching toward the tool point.

6.3.4 Integration of the Robot into the Virtual Scene

To allow the user to accurately maneuver the robot within the virtual scene, the physical and virtual robots must match in position and orientation. This is done by registering the Kuka and Oculus Rift and positioning the Oculus Rift camera frame relative to the virtual robot using the transformation matrix obtained during registration.

Points on the tracking camera, shown in Figure 6.7, were measured relative to the camera frame and touched with the robot end effector to measure the points in the robot frame. The rigid transformation matrix between the camera and robot frames was then computed using the SVD

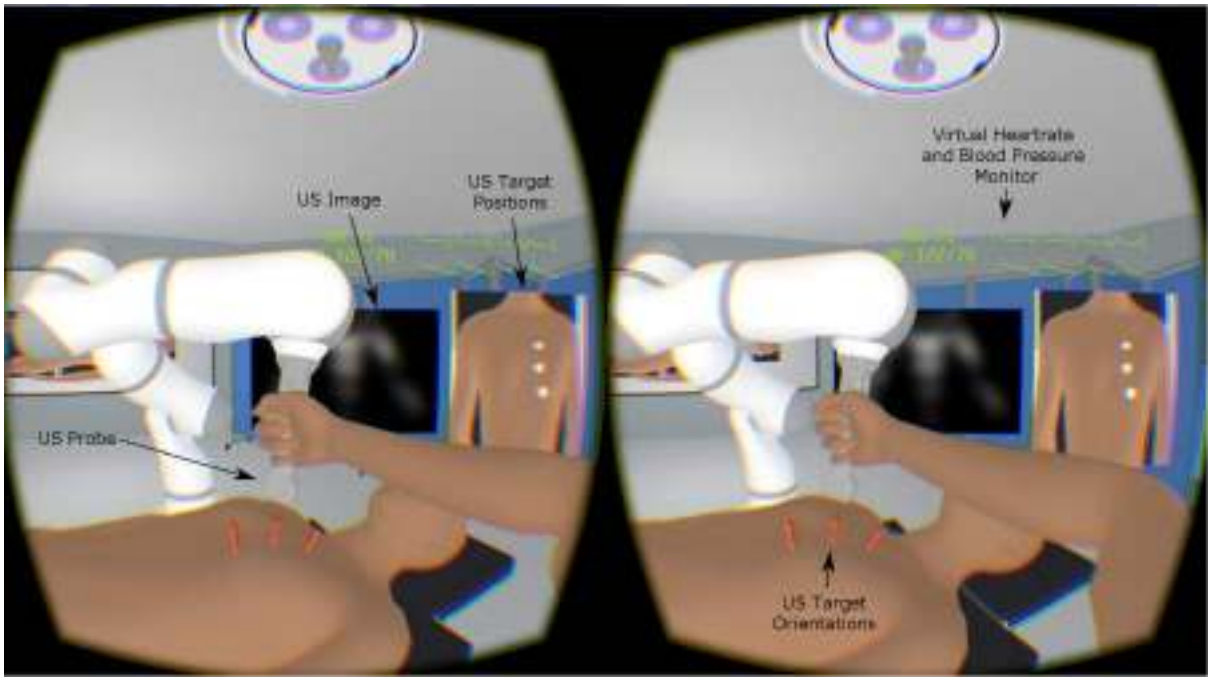

Figure 6.6: The Oculus Rift view of the surgical theater including the robot, patient, and monitoring system.

method [Bellekens et al., 2014]. The cross-correlation matrix, W, between the centers of the two point clouds was first computed,

$$W = \sum_{N} (r_i - r_m)(c_i - c_m)^T \tag{6.2}$$

where N is the total number of points, r_i and c_i are the i-th corresponding points in the robot and camera frames respectively, and $r_m = \frac{1}{N}\sum_{N} r_i$ and $c_m = \frac{1}{N}\sum_{N} c_i$ are the point cloud centers. The rotation matrix between the point sets was then computed as,

$$R = UMV^T \tag{6.3}$$

where U and V are computed from the SVD of $W = USV^T$ and $M = diag(1, 1, \det(UV^T))$ ensures that the resulting matrix contains only rotations and not reflections. The translation can be computed from this rotation matrix and the point clouds means.

$$t = r_m - Rc_m \tag{6.4}$$

Figure 6.7: Oculus Rift tracking camera with the points touched highlighted with spheres.

6.3.5 Ultrasound-Body Feedback

Within the Linux robot controller, a visualisation of which is shown in Figure 6.8, feedback forces and torques were implemented to simulate the interaction between the ultrasound and the patient's body. The true interaction of the tool with the patient involves forces and torques reciprocally applied to both the patient and the US probe, based on the particular areas of contact (i.e. ribs or stomach). Additionally, deformation of the patient's body would occur also and the amount of deformation would be based on the particular contacted area. Due to the computational complexity of such a soft tissue interaction, a simplified form was implemented, which ignored the deformation of the patient's body and assumed a uniform stiffness for the body. Instead, the US probe penetrated into the patient's body mesh and feedback forces were generated based on a linear stiffness parameter and the depth of this penetration.

To calculate the forces and torques of the tool-tissue interaction to apply on the robot to the user, first, the distance the tool was penetrating into the body mesh was found. This was done by computing the closest points on the triangular body mesh to 6 points on the US probe, seen in red in Figure 6.9, using conventional barycentric coordinate methodologies [Ericson, 2004]. The normal of the corresponding closest triangle was used to determine if the tool was inside or outside of the mesh. The point-to-mesh method was chosen over mesh-to-mesh algorithms to ensure a control rate of 1kHz. From the penetration algorithm, the feedback forces were computed for each point chosen on the ultrasound probe using the following linear force feedback law,

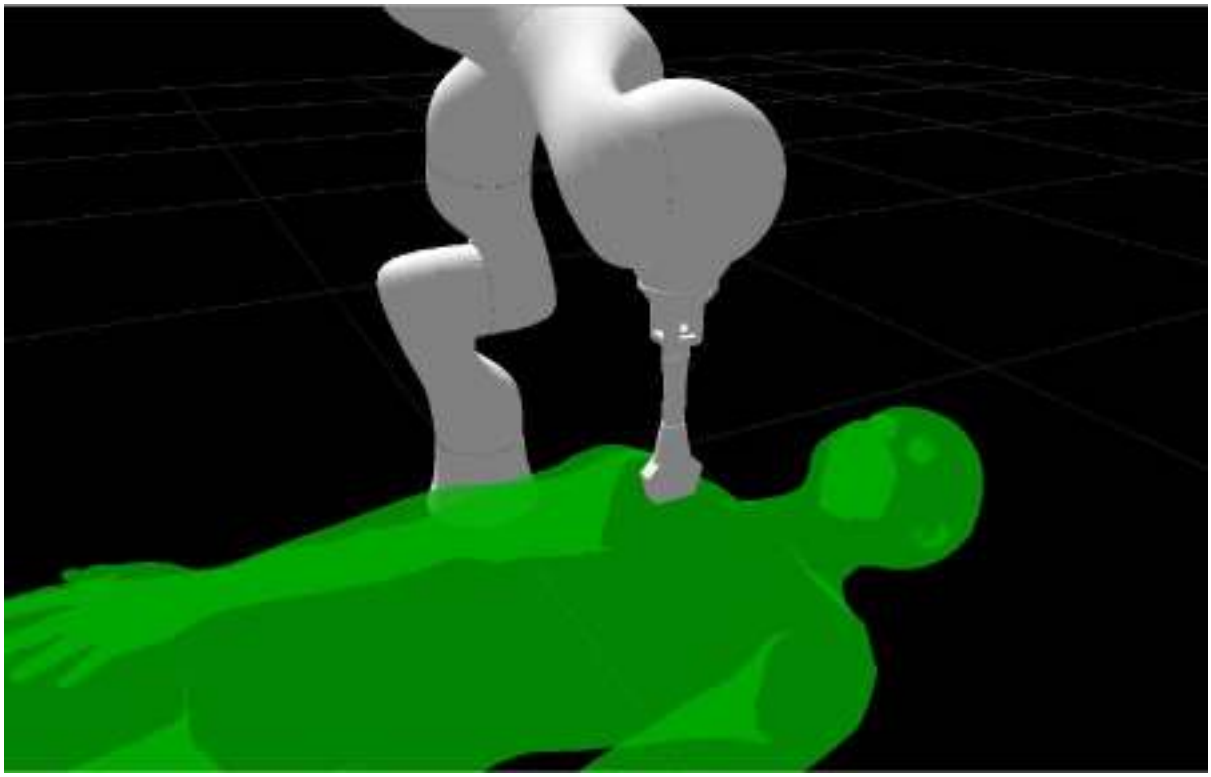

Figure 6.8: RVIZ visualisation of the Linux side robot and patient used for computing the body feedback forces and torque.

$$F_{p_i}^{bf} = \begin{cases} -k_p^{bf} d_p \hat{n} - k_d^{bf} J_{p_i} \dot{q} & : d_p > 0 \\ 0 & : d_p \leq 0 \end{cases} \tag{6.5}$$

where k_p^{bf} and k_d^{bf} are the stiffness and damping of the body feedback respectively, d_p is distance the point p_i has penetrated into the mesh, $\hat{n}$ is the normal vector from p_i to the closest mesh point, and $F_{p_i}^{bf}$ is the body feedback for a single point, p_i. The gains used in this experiment were empirically set to $k_p^{bf} = 100$ and $k_d^{bf} = 5$ to minimise penetration of the tool into the body mesh without any oscillatory instabilities occurring.

The Kuka LWR utilises a single application point for end effector forces and torques and, therefore, the feedback forces must be transformed onto this application point.

$$F^{bf} = \sum_{p_{bf}} F_{p_i}^{bf} \tag{6.6}$$

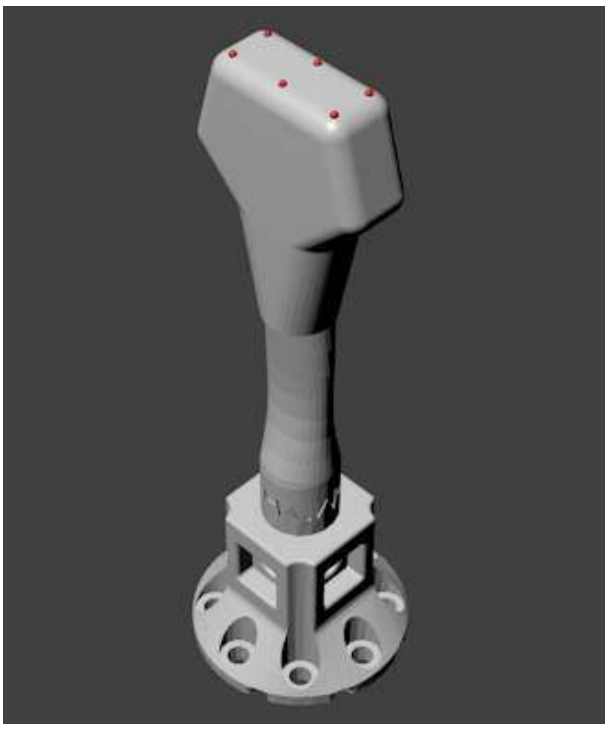

Figure 6.9: Rendering of the US probe tool with the points used in the closest point search for computing body feedback in red.

$$\tau^{bf} = \sum_{p_{bf}} (p_i - p_{cp}) \times F_{p_i}^{bf} \tag{6.7}$$

where F^{bf} and τ^{bf} are the body feedback forces and torques, respectively, applied to the control point, p_{cp}.

6.3.6 Ultrasound Targets

The ultrasound target points were estimated from clinical instructions detailing the use of thoracic ultrasound in pneumothorax [Reardon, 2015] and are shown in Figure 6.4. The normals to each target point were computed as the mesh normal of the closest point to the target on the patient's body.

The applied force required to obtain a good scan depends on the body mass index (BMI) of the patient. The average force has been shown to vary between $7.5N$ for normal BMI and $9.8N$ for high BMI [Dhyani et al., 2014]. The target points were therefore offset into the body along the

normal $1.5cm$ such that the resulting force from the body feedback when the US probe was at the target point was $9N$.

6.3.7 Active Constraints and Natural Motion

The dynamic active constraints strategy presented in Chapter 5 was used to prevent the US probe from penetrating too far into the body mesh and, consequently, harming the patient via force application which was too high. A set of constraint points were chosen such that the zero set of the generated implicit surface passed through the target points by offsetting from the target points along their respective normals and computing the threshold using the minimum protection distance formula from Chapter 5, $T = \psi(r_{min})$.

As discussed in Chapter 4, applying natural motion on the Kuka LWR 4+ requires the use of the Cartesian impedance control mode, due to limitations in applying null-space control torques in joint impedance mode whereas in Chapter 5, the joint impedance mode was used to implement dynamic active constraints. To apply these shared control strategies simultaneously, the dynamic active constraints strategy must be adapted to the Cartesian Impedance mode for the Kuka.

The methodology developed in Chapter 5 is a dynamics model-based control strategy which applies forces to control the acceleration of the end effector via joint torques. In Cartesian impedance mode, this is not possible, however, forces and torques can be set at the tool point as was done for body feedback. Therefore, I will convert the strategy to a direct force application at the end effector. In addition, I will expand the strategy to constrain multiple points at the end effector and to apply torques to constrain orientation. The same 6 points at the end effector used to compute the body feedback wrench (Figure 6.9) were used to compute the dynamic active constraints forces. For each of these points, the force feedback was computed in the implicit function task space and projected in the end effector frame,

$$F_{p_i}^{dac} = \frac{J_f^T(p_i)}{\|J_f(p_i)J(p_i)\|}\left(k_p^{dac}(f_{set} - f(p_i)) - k_d^{dac}J_{IF}(p_i)\dot{q}\right) \tag{6.8}$$

where $F_{p_i}^{dac}$ is the dynamic active constraints feedback for a single point p_i, $J_f(p_i)$ is the Jacobian

of constraint surface with respect to the spatial variables evaluated at p_i, $J(p_i)$ is Jacobian relating joint velocities to Cartesian velocities at p_i, k_p^{dac} and k_d^{dac} are the stiffness and damping of the dynamic active constraints respectively, f_{set} is the desired level set threshold, $f(p_i) = T - \sum_J \psi_{4,2}(\|p_i - p_{c,j}\|)$ is the current level set of p_i with respect to the set of J constraint points $p_{c,j}$, and $J_{IF}(p_i) = \frac{J_f(p_i)J(p_i)}{\|J_f(p_i)J(p_i)\|}$ is the Implicit Function Jacobian evaluated at p_i.

The desired constraint forces can then transformed onto the application point to find the total desired dynamic active constraints forces and torques.

$$F^{dac} = \sum_{p_{dac}} F_{p_i}^{dac} \tag{6.9}$$

$$\tau^{dac} = \sum_{p_{dac}} (p_i - p_{cp}) \times F_{p_i}^{dac} \tag{6.10}$$

where F^{dac} and τ^{dac} are the dynamic active constraint forces and torques, respectively, applied to the control point, p_{cp}. The parameters used in these experiment for the dynamic active constraints including the radius of effect R, the threshold T, and the desired level set f_{set} previously defined in Chapter 5.2 can be found in Table 6.2.

Parameter	k_p^{dac}	k_d^{dac}	R	T	f_{set}
Value	4	0.2	$0.05m$	0.3242	0

Table 6.2: Parameter values used in the dynamic active constraints controller.

In addition to limiting the forces applied by the user via dynamic active constraints, orientation constraints were used to ensure the US probe was correctly aligned with body normals to allow for a good scan. The desired orientation at each time step was computed using a spherical linear interpolation (SLERP) between the normal vectors of the target points at each instance in time. The tool orientation was commanded to this interpolation of the targets normals and not the body normal to simplify the control for the user and to prevent excessive rotations at the sides of the body. To handle the multiple targets and for situations where the tool was not near the target points, the interpolation was region based. The different regions, which are shown in Figure 6.10, were defined based on the y-coordinate of the tool and target points. When the end effector is in a region between two target points, the desired normal is an SLERP interpolation

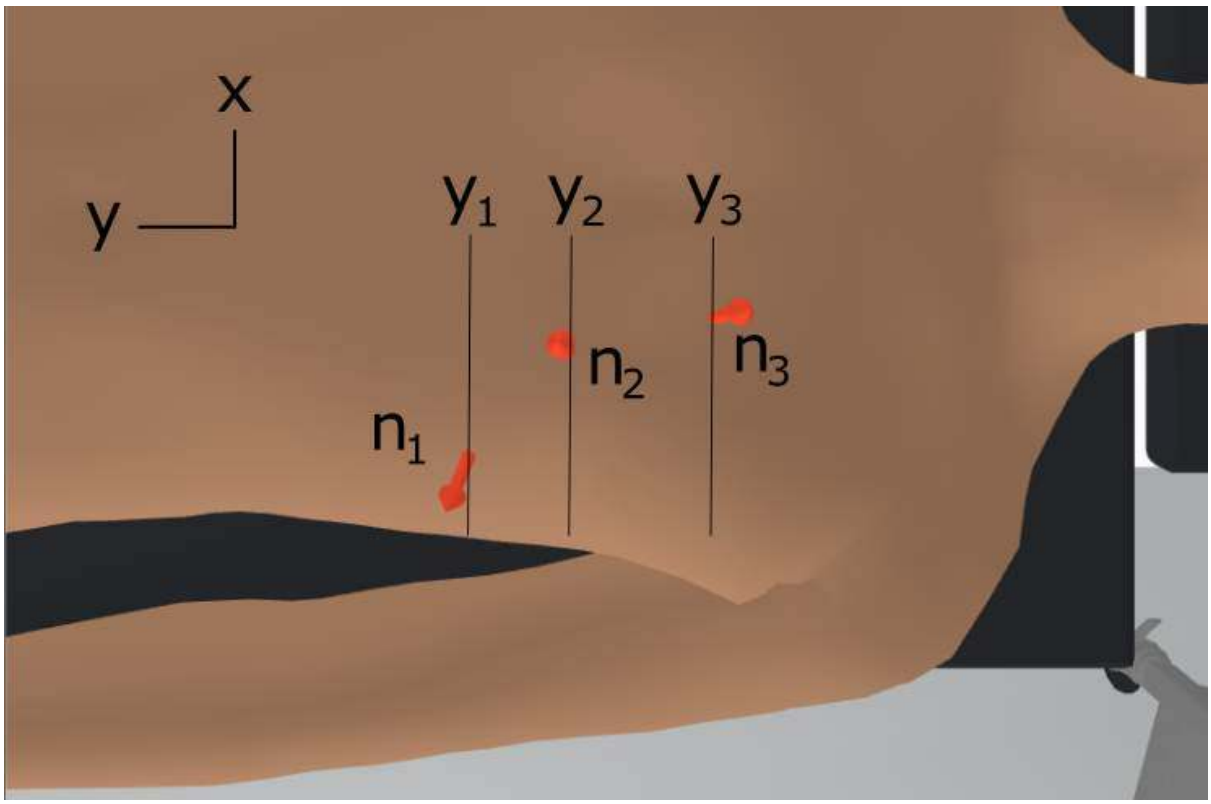

Figure 6.10: Labeled top down view of the patient's body and scan point orientation normals for computing the desired orientation of the US probe.

based on the y distance to each point. Outside these regions, the tool is commanded to the closest target normal. This strategy is summarised by the following equation, which computes the desired end effector normal.

$$n_{vec} = \begin{cases} n_1 & : y_{tool} \geq y_1 \\ \frac{\sin((1-t_{1,2})\phi_{1,2})}{\sin(\phi_{1,2})} n_1 + \frac{\sin(t_{1,2}\phi_{1,2})}{\sin(\phi_{1,2})} n_2 & : y_1 > y_{tool} \geq y_2 \\ \frac{\sin((1-t_{2,3})\phi_{2,3})}{\sin(\phi_{2,3})} n_2 + \frac{\sin(t_{2,3}\phi_{2,3})}{\sin(\phi_{2,3})} n_3 & : y_2 > y_{tool} \geq y_3 \\ n_3 & : y_3 > y_{tool} \end{cases} \tag{6.11}$$

where y_i is the y position of point i, y_{tool} is the y position of the tool point, n_i is the normal vector of point i, $\phi_{i,j} = \arccos(n_i \cdot n_j)$ is the angle between normal vectors n_i and n_j, and $t_{i,j} = \frac{y_i - y_{tool}}{y_i - y_j}$ is the normalised interpolation factor based on the distance the tool point is between points i and j in the y-coordinate. The orientation gains on the Kuka were empirically set to $k_p = 50$ and $k_d = 0.5$.

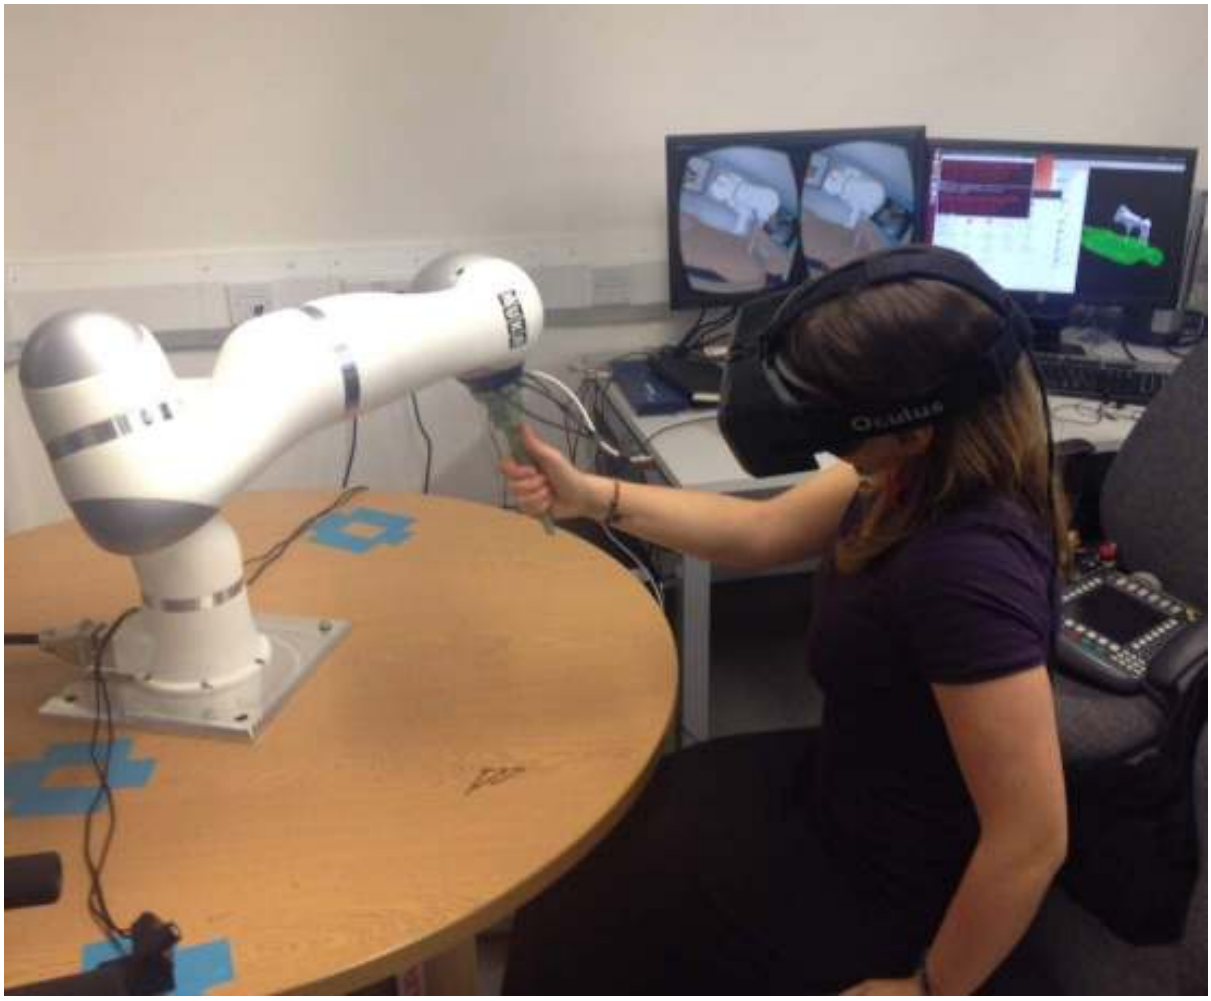

Figure 6.11: A subject performing US scans within the experimental environment.

Natural motion was applied in the null-space using the dynamic simulation strategy developed in Chapter 4. End effector mass and friction force were simultaneously optimised with a ratio 50%-50%. Similarly to Chapter 4, the parameters of (3.16) were set to $k_p = 50$, $k_d = 10$, and $\omega_{max} = 10 rad/s$, the potential velocity was set to $5 cm/s$, and the null-space parameters on the Kuka were set to $k_p^{null} = 100$ and $k_d^{null} = 0.1$.

6.4 Experimental Setup

A total of 9 non-surgeons participated the experiment, seven males and two females. Two of the subjects were left-handed and seven were right-handed. The ages of the participants ranged from 21-34. Each subject performed the US scans with the assistance modes 3 times each for

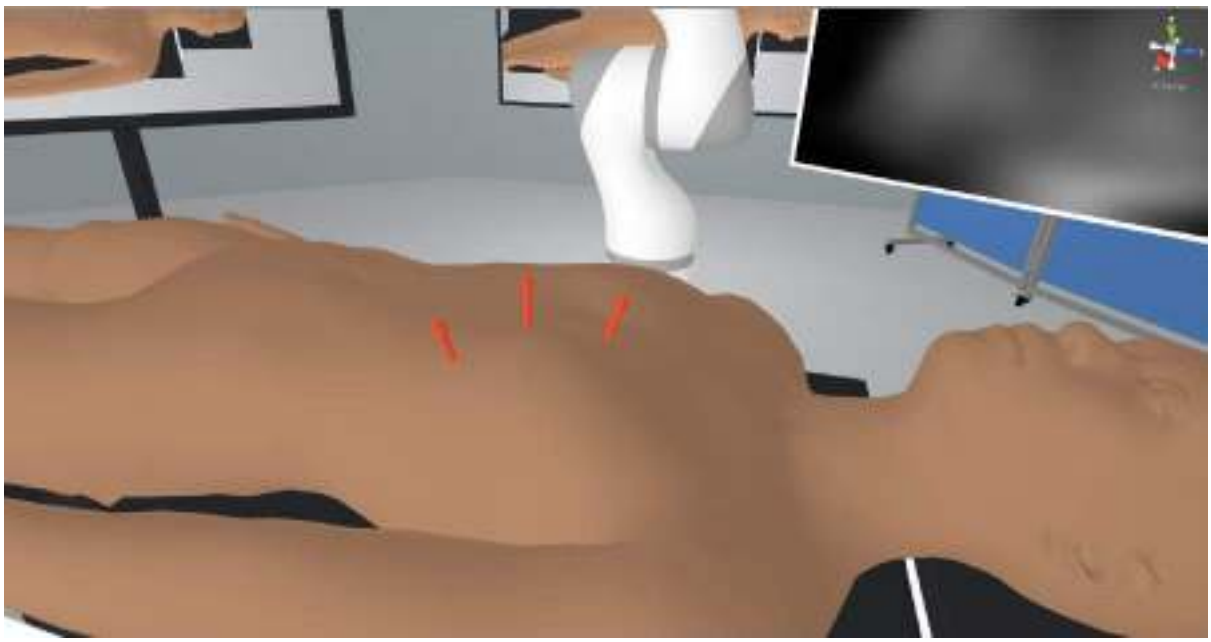

Figure 6.12: The visualisation of the body normals at the points to be scanned used in the AR and dynamics active constraint with AR assistance modes.

a total of 9 trials and 27 scanned points per subject. The order of the assistance modes was randomised to minimise possible learning effects. Prior to performing the trials, to reduce the possibility of training effects, users were given some time to familiarise themselves with the assistance methods and performing the ultrasound scan, but were not given details about the controllers and the aims of the experiment in general. The total time for each subject to perform the full experiment was approximately 20-30 minutes.

The experiments aimed to demonstrate the advantages of using robotic assistance, specifically dynamic active constraints, in performing US scans for detecting pleural effusion. 3 methods were compared: no assistance, augmented reality (AR), and dynamic active constraints with augmented reality (DACs with AR).

The no assistance mode aimed to mimic the use of a standard US probe. The robot was commanded with gravity compensation and the wrenches of the tool-tissue interaction were the only applied forces and torques to the user. Visually, the subjects only saw the US monitor and the desired scan points indicator.

AR aims to provide additional assistance to the user by embedding virtual information within the physical world. In these experiments, AR assistance provided a visualisation of the body

normals at the necessary scan points, as shown in Figure 6.12, to aid the user in correctly posing the US probe. This information is hypothesised to improve the accuracy in positioning and orientating the US probe as compared to no assistance. This mode aimed to mimic that of using a standard US probe with additional visual assistance. Similarly to the no assistance case, the robot was only commanded with gravity compensation and the forces and torques of the tool-tissue interaction.

While AR assistance aims to provide additional visual information to the surgeon, thereby making it easier to perform procedures, there is still a chance of hurting the patient or performing the procedure incorrectly due to the lack of enforcement of these guides. Robotic assistance aims to control some of these degrees-of-freedom while still leaving the surgeon in control of the overall procedure. The DACS with AR assistance mode utilises dynamic active constraints to limit the forces which the surgeon is applying to the body of the patient and controls the orientation to be aligned with the target points. In doing so, the likelihood for damage to the patient is reduced and a more accurate scan is guaranteed, leaving the surgeon to focus better on interpreting the US scan. In addition to the gravity compensation and the interaction forces and torques of the tool with the soft tissue, wrenches were applied at the end effector to control the orientation of the robot and to enforce the dynamic active constraint when the tool penetrated too deeply into the body mesh.

In each scenario, the robot was used with its end effector as the US probe, in order to apply the body feedback forces and torques. However, only in the dynamic active constraints scenario did the robot apply additional forces and torques to guide the user. Natural motion was applied in the null-space in all three modes to ease the user's movement of the tool.

6.5 Results

Four metrics were selected to assess the effectiveness of the assistance modes; average position error, average orientation error, average applied force, and trial length. Average position and orientation error are used as a measure of accuracy. Average applied force indicates the level of safety for the patient and trail length measures the speed at which the user can perform the

US scans. The results for these metrics across all subjects can be found in Table 6.3.

The average position error is the mean distance the end effector position was from the closest target point and its distribution across all subjects and assistance modes is shown in Figure 6.13a. Similarly, the average orientation error is the mean angle between the closest target normal and the tool and its distribution across all subjects and assistance modes is shown in Figure 6.13b. The average applied force is computed as the mean body feedback force and reflects the counter forces to the user's applied forces. Its distribution across all subjects and assistance modes is shown in Figure 6.13c. The average position error, orientation error, and applied force were only computed when the tool was penetrating the body mesh to avoid measuring movements between scan points. The trial length is the total time the user took to perform the scans starting from the first penetration of the body mesh with the tool and ending with the last scan. The distribution for this metric across all subjects and assistance modes can be found in Figure 6.13d.

Each of the metrics was found to be non-normally distributed using the Komogorov-Smirnov test. Therefore, the non-parametric Kruskal-Wallis one-way analysis of variance test was used and it was found that, for all four metrics, at least one of the distribution's medians was different from at least one other group at the 95% confidence level. To analyse specific pairs for dominance in the medians, pairwise comparisons of the assistance modes were made using the Dunn-Sidak test at the 95% confidence level. Additionally, the Brown-Forsythe test was used to make pairwise comparisons between the variances of the assistance modes at the 95% confidence level. The results of the pairwise median and variance comparisons are shown in Table 6.4.

6.6 Discussion

The experiments compared the performance of the subjects in performing US scans for diagnosing pleural effusion using three assistance modes; no assistance, augmented reality, and dynamic active constraints with augmented reality. The experimental results demonstrate that the DACS with AR significantly improve the accuracy, safety, and speed of the simulated ul-

Assistance Mode	Average Position Error (mm)	Average Orientation Error ($^\circ$)	Average Applied Force (N)	Trial Length (s)
None	9.62 ± 2.50	9.35 ± 3.00	12.22 ± 1.58	253.71 ± 248.94
AR	7.41 ± 1.41	6.14 ± 3.11	10.41 ± 1.57	94.87 ± 101.16
DACs and AR	7.24 ± 1.52	2.76 ± 1.58	8.70 ± 1.23	49.08 ± 24.34

Table 6.3: The mean and standard deviation of the metrics from the virtual reality experimental trials across all subjects for the three assistance modes. Note that the large standard deviations of the trial lengths are a result of the skewed nature of the non-normal distributions.

Assistance Mode	Average Position Error	Average Orientation Error	Average Applied Force	Trial Length
None vs AR	$\star$	$\star/\bullet$	$\star$	$\star/\bullet$
None vs DACs and AR	$\star$	$\star/\bullet$	$\star$	$\star/\bullet$
AR vs DACs and AR		$\star/\bullet$	$\star$	$\bullet$

Table 6.4: Significance results for the performance metrics across subjects for the three assistance modes. $\star$ represents a significant difference in the means for a given and assistance mode pairing at the 95% confidence interval. $\bullet$ represents a significant difference in the variances for a given and assistance mode pairing at the 95% confidence interval.

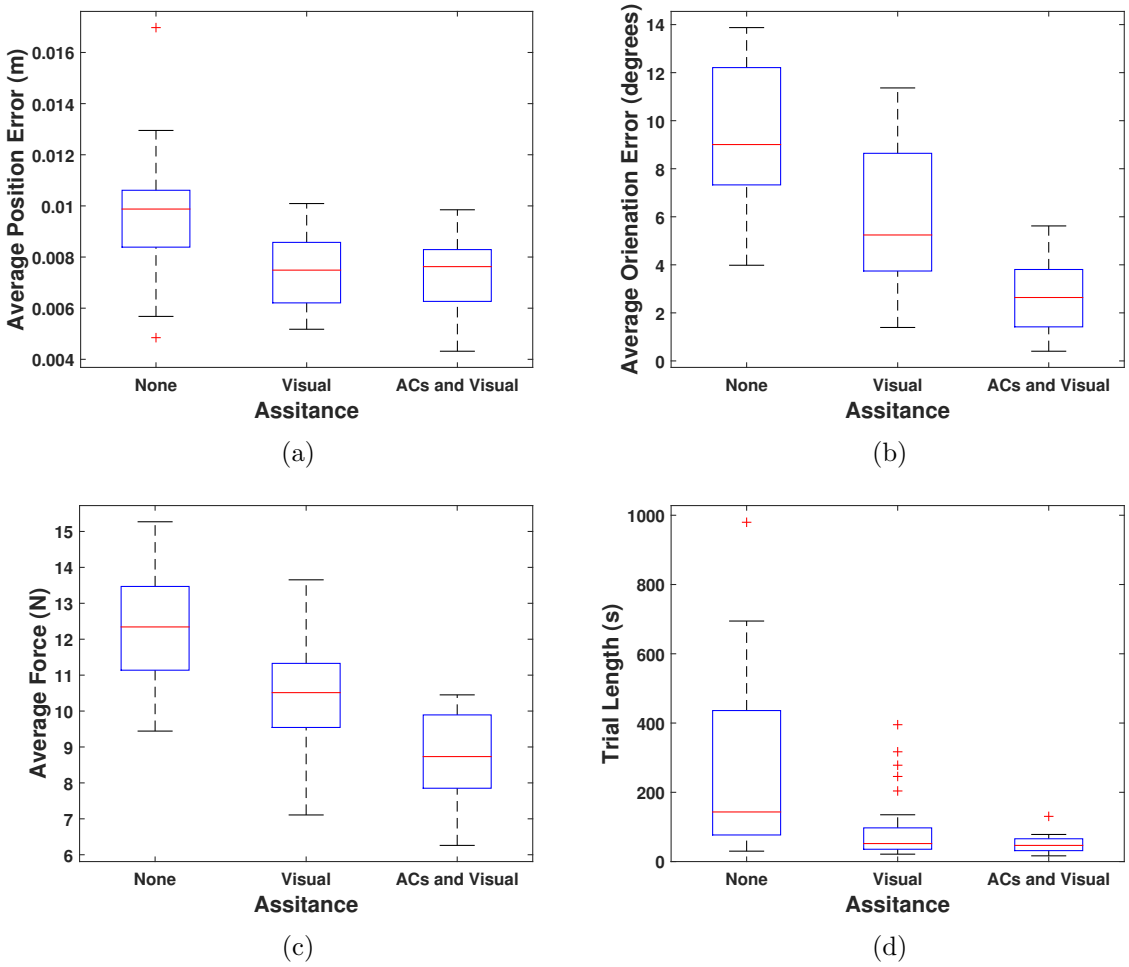

Figure 6.13: The median and IQR for the performance metrics across subjects for the three assistance modes. None signifies no assistance, Visual signifies augmented reality assistance, and ACs and Visual represents the dynamic active constraints with AR assistance.

trasound task.

Accuracy in performing the US task is indicated by the average position and orientation error metrics, computed as the mean distance and angle respectively, between the end effector and the closest target point when the tool is penetrating the body mesh. Improving accuracy in posing the US probe ensures the desired regions are being properly scanned and is important in guiding needle insertion to remove air or fluid in the region surrounding the lungs. The no assistance condition performed significantly worse than both the AR and DACs with AR assistance modes, which lends evidence to the hypothesis that a reduction in information about the location of the target points results in low accuracy. The visual guides allowed the users to reduce the error perpendicular to the line passing through the target point and the closest point on the body mesh, leaving only the small error between the tool point and the target point along the normal.

The DACs with AR mode performed significantly better than the AR only mode when it comes to orientation error, however, there was no significant difference in position error between the two modes. The improvement in rotational error is consistent with the additional assistance provided by the DACs, which constrains the orientation of the tool. The DACs also prevent the US probe from being applied too forcefully to the patient by restricting the penetration depth of the tool. However, this failed to significantly affect the position error. The position error measurement here is the absolute distance between the target point and the tool point and, therefore, does not directly reflect the depth of the tool, which is the position restriction placed on the tool by the DACs.

Safety is assessed within the virtual US task by the average applied force, which is computed as the mean body feedback applied to the user in response to the tool penetrating the body mesh. By reducing the applied forces of the surgeon, musculoskeletal injuries can be reduced. The no assistance mode again performed significantly worse than the other modes in force application, with an average 40.5% increase in the worse case. Without any assistance in position, orientation, or force, the subjects could not easily estimate which of these parameters needs to be improved upon to achieve a clear scan. In other words, they could not discern whether the position of the device was to blame for the lack of clarity in the US image (indicating

the degree of error from the optimal pose), or any of the other parameters. When no guidance was provided on these parameters (no assistance mode), subjects developed search strategies, which included applying excessive force. In the assisted modes (AR and DAC, and DAC), subjects had guidance on both position and orientation and therefore, a narrower search space. This lead to more controlled searches and a lighter application of force.

There was a significant difference in the average force application across trials between AR modes where DACs were present and where they were not. The lack of DACs resulted in a 19.7% increase in average force application over AR with DACs and therefore, DACs helped reduce the force applied to the virtual patient, demonstrating increased safety in US scanning. As discussed above, from observing the trials, the AR assistance mode lessened the mean force over no assistance as it provided more information to the user, making it more clear what error needed to be corrected to achieve a good scan. However, there still was ambiguity in determining the correct pose and force and, therefore, users still applied higher forces than necessary when using the AR only assistance mode, whereas the assistance mode with DACs limited the application of force.

Speed in performing the US scans within the virtual environment is assessed by the trial length, which was computed as the time between the first penetration of the tool with the body mesh until the last point is successfully scanned. Increasing the speed at which the physician can perform US scans reduces the time to diagnosis or treatment. There was a significant increase in the trial length when performing the US scans with no assistance as compared to the AR and DACS with AR assistance modes. These results suggest that the subjects found it difficult to determine whether to translate, rotate, or apply force to achieve a good scan without additional assistance to reduce the dimensionally of the scanning, and were forced to find the appropriate parameters by trial and error.

There was a significant difference between the AR assistance mode and the DACs with AR mode in the variances of the trials lengths, however, there was not a significant difference in the medians. These results would suggest that providing visual feedback was the most significant factor in reducing the time it took to perform these scans, but, that additional constraints on the orientation and applied force did not significantly affect the speed.

The conclusions drawn from these experiments are limited by the realism of the soft tissue simulation, the testing on the control strategies on non-surgeons, and the selection of the controllers which were compared.

The feedback forces and torques from the tool interaction with the body were generated via a linear feedback law. Soft tissue simulation including factors such as varying stiffness and damping due to interactions with different bones and organs and deformation of the patient's body due to tool contact would provide a more realistic environment with which to evaluate the assistances modes. Soft tissue simulation software such as the Simulation Open Framework Architecture (SOFA) [Allard et al., 2007] could be incorporated to enhance the realism, however, are very computationally expensive to run, particularly to simulate a large portion of the body and effects such as cutting.

Surgeons demonstrate a higher level of dexterity and hand-eye coordination than the average population [Francis et al., 2001] and, therefore, the impact on accuracy, safety, and speed in these experiments performed by non-surgeons may be lessened when tested on clinicians. Due to the current clinical application of active constraints and the wide body of research demonstrating similar benefits, it would be expected that the constraints would still assist the surgeon, however, this must be explicitly demonstrated with these specific dynamic active constraints.

Due to time constraints, an assistance mode which applied dynamic active constraints on depth and orientation without augmented reality was not tested. It would be expected that the improved results on applied force and orientation error when using the DACs with AR would be preserved, however, the position error would increase due to the lack of visual information about the target points. Additionally, comparing the performance of the dynamic active constraints developed within this thesis with classical potential field constraint approaches in this task would evaluate the effectiveness of the proposed technique over a standard methodology.

6.7 Summary

This chapter has presented a virtual environment for testing hands-on robotic surgery control strategies. The physical robot and VR environment are registered such that virtual robot and

physical robot are aligned. This allows the user to interact with the robot and its control strategy and dynamics in the same way he or she would in a real scenario. As opposed to *in vivo* and phantom trials, the virtual environment allows for repeated testing of shared control strategies and surgical training through simulation of the soft tissue interaction with the tool. Additionally, such a system is more flexible than phantom trials, as multiple types of surgery can be more easily adapted and tested.

An ultrasound scenario for diagnosing pleural effusion and pneumothorax was developed in a virtual surgical environment to test the shared control strategies developed in the previous chapters. The technical details of the virtual environment and the experimental design were discussed. The experiment compared the performance in taking US scans with no assistance, augmented reality assistance, and robotic assistance with augmented reality. The results demonstrated that robotic assistance through dynamic active constraints and AR improved the user's ability to perform the US scans safely and accurately. The next chapter will draw conclusions about the thesis contributions with regards to the stated research aims, discuss the limitations of the work presented in this thesis, and identify future areas of research.

Chapter 7

Conclusions and Future Work

7.1 Conclusions

Hands-on robotic surgery is where the surgical tool is attached to the end effector of the robot and the surgeon maneuvers the instrument by directly applying forces and torques to the tool. The main advantage of performing surgery in this way is that the forces and torques of the interaction between the robot and the surgical environment can be felt through the tool, without the use of an additional force/torque sensor at the end effector. The surgeon can more precisely apply forces and torques to the patient using this natural feedback, as opposed to teleoperative systems, which have issues with stability and sterilisability in providing haptic feedback. Additionally, as compared to autonomous robotic systems for surgery, the surgeon remains more involved in the procedure, increasing its acceptance among physicians and regulatory bodies [Davies et al., 2004]. To further improve results, shared control strategies, in which the robot and user simultaneously apply control polices to accomplish a task, are used. This thesis has presented three contributions to shared control for hands-on robotic surgery; natural motion, dynamic active constraints through implicit surfaces generated from point clouds, and a virtual reality environment for testing shared control.

The first contribution made in this thesis is the concept of natural motion. Attaching the surgical instrument to the end effector of the robot imposes the dynamics of the robot at the tool point on the surgeon maneuvering the tool. These dynamics can include gear backlash,

joint friction, the mass and inertia of the robot links, and other disturbances. The additional dynamics can limit the ability of the surgeon to perform the surgery by reducing his or her dexterity and increasing fatigue. Natural motion aims to reduce these end effector dynamics and make the feeling of the surgeon's interaction with the robot closer to that of the surgical instrument unattached to the robot.

Current shared control strategies for hands-on robotic surgery only compensate for gravity. In the general robotics literature, cooperative control methodologies have been developed which aim to ease the user's ability to maneuver objects. These strategies utilise impedance controllers to effectively amplify the user's applied forces and torques, which enable the user to lift heavy objects. However, an increased risk of contact or coupling instability occurs when impedance controllers attempt to emulate dynamics that differ significantly from the intrinsic hardware dynamics [Colgate and Hogan, 1989]. Due to the complexity of the surgical environment, ensuring stability using these types of strategies would be difficult and, therefore, not applicable to clinical scenarios. In addition, these strategies increase the surgeon's applied forces and torques to the tissue, potentially leading to an unintentional injury to the patient. They also affect the ability of the surgeon to accurately feel the forces and torques of the tool-tissue interaction through the end effector, removing a substantial benefit of hands-on robotic surgery.

In Chapters 3 and 4, an optimisation strategy was demonstrated which reduces the mass, inertia, friction force, and friction torque at the end effector via control of the redundant degrees-of-freedom. By performing the optimisation in the null-space of the tool point, a more natural motion is created without affecting the surgeon's command of the surgical instrument. Simulation work demonstrated the effectiveness of the natural motion strategy in reducing these dynamic quantities in specific as well as general directions over the workspace of the robot. The experimental results of Chapter 4 prove a significant reduction in the work required to move the tool when using natural motion in a general point-to-point motion task, easing the user's ability to maneuver the surgical instrument. Additionally, natural motion was embedded in a realistic clinical environment in Chapter 7, demonstrating a seamless coexistence with other controllers, in this case, dynamic active constraints.

Within current clinically applied hands-on robotic systems, active constraints are used to protect

the patient and enable higher accuracy than what the surgeon can accomplish alone. Active constraints apply forces and torques to the surgeon to prevent the surgical tool from entering regions which have been deemed unsafe or unnecessary for the surgery. At present, however, clinical applications are limited to rigid bone surgeries, as the constraint representations in these scenarios are fixed and intra-operative tracking and registration of the patient's anatomy is simplified. Recently, research into dynamic active constraints, which can protect soft tissue deforming under such effects as tool interaction and heart beat, has begun in order to expand the range of surgeries active constraints can benefit.

The second contribution of this thesis has been developing a dynamic active constraints methodology which can directly generate, in real-time, a constraint surface from a streaming point cloud of the region to be protected and incorporates the dynamics of the manipulator into its enforcement. Current dynamic active constraint approaches restrict the allowed motion of the deforming region and/or require specific geometric representations of the patient's anatomy. To be effective across a wide range of surgeries, dynamic active constraints strategies need to be able to handle a variety of different modalities for intra-operative tracking and updating of the constraint region. The constraint representation must be able to take on arbitrary shape, in order to handle the unknown deformations which may occur due to pulsation, breathing, and tool-tissue interactions. Rebuilding the constraint surface from point clouds at each time step avoids issues in structured representations, such as meshes, in tracking and updating under complex deformations, such as indentation and cutting.

Recently, low cost 3D cameras and range finders have become ubiquitous and provide fast point cloud representations of the environment. Additionally, more standard sensing methodologies used in surgery, such US and CT scans, can be reduced to point clouds. Therefore, constructing the constraint from point clouds allows for the strategy's general application to many types of surgeries.

The methodology used to generate the dynamic constraints surface is similar to previous strategies for constraint based haptic rendering of dynamic environments from points clouds [Leeper et al., 2012] [Rydén and Chizeck, 2013]. The fundamental difference in the control strategy developed herein is that the dynamics of the robot have been incorporated into the enforcement

of the constraint. In haptic constraint enforcement, the master and slave device can be decoupled and the slave can be position controlled to remain outside the constraint region, thereby, ensuring the tool does not reach the soft tissue. Such a decoupling strategy is not possible in hands-on robotic surgery, as the surgeon directly applies forces and torques to the robot and impedance controlled devices are used to allow the surgeon to feel the tool-environment interaction through the robot. As a result, penetration will occur into the constraint region. Additionally, this physical contact between the robot and surgeon couples the dynamics of the two. Therefore, an impedance control strategy has been developed which incorporates the dynamics of the robot in order to provide an appropriate impedance to the surgeon's forces and torques when the tool has penetrated the constraint region.

Chapters 5 and 6 develop and demonstrate, experimentally, the efficacy of the implicit surface dynamic active constraints shared control strategy in protecting the patient and increasing accuracy in performing tasks on soft tissue.

The last contribution of this thesis is the development of a virtual reality environment for testing shared control strategies for hands-on robotic surgery. Preliminary testing of robotic control strategies is not performed *in vivo* due to safety concerns and expense. Soft tissue phantoms, which aim to mimic the motion of soft tissue, have been developed as an alternative method on which to test such strategies. However, phantoms are still limited in their ability to replicate soft tissue motion. It is difficult to devise a soft tissue phantom which can mimic the numerous mechanical properties of an organ in response to interactions such as indentation and cutting and exhibit deformation due to breathing and pulsation.

Similarly to soft tissue phantoms, a virtual environment offers a safe alternative to preliminary testing of shared control strategies, however, VR is more flexible and makes repeated trials easier. A virtual reality environment is limited by computational power rather than the ability to create a chemical recipe and it can be easily changed between different types of surgeries on various parts of the anatomy, once the soft simulation has been developed for a particular region of the body. In addition, the VR surgical environment can be easily reset to allow for repeated testing or training with hands-on shared control strategies, as opposed to soft tissue phantoms, which require a new phantom for each trial in scenarios with cutting or plastic

deformation.

In the virtual reality setup, the physical robot is registered and aligned with the virtual robot, allowing the complex dynamic interactions between the robot dynamics, the control strategy, and the user to be preserved. To emulate the tool's interaction with the soft tissue, forces and toques can be applied at the robot and soft tissue simulation can be used to simulate deformation. This VR surgical environment was used in Chapter 6 to evaluate the dynamic active constraints strategy of Chapter 5 and demonstrate the simultaneous application of natural motion and dynamic active constraints.

In addition to the potential benefits for hands-on robotic surgery, the contributions developed here have application to general robotics. Dynamic active constraints generated from point clouds could be applied to complex tasks on deforming objects to avoid particular regions of the item being worked on. Natural motion could be applied to general physical human-robot interaction to reduce the impact of the robot's end effector dynamics on the collaborating human. Similarly, the virtual reality environment designed for hands-on robotic surgery can be used for testing shared control strategies in human-robot interaction, providing an immersive environment which preserves the complex interaction between the robot, the control strategy, and the user.

7.2 Limitations and Future Work

There are several limitations to the research presented in this thesis which will be highlighted in this section. In addition, future research directions which could address these limitations are discussed.

The experimental validations of the shared control strategies presented within this thesis were performed by non-surgeons. Surgeons have been shown to have a higher level of hand-eye coordination and dexterity on average [Francis et al., 2001]. This above average performance may reduce the improvements in accuracy and safety demonstrated in Chapter 6 when utilising the dynamic active constraints control strategy, however, the benefits of active constraints shown in the general literature and the clinical application of these types of strategies would lead us

to conclude that these dynamic active constraints would still significantly improve performance and protect the patient.

Regarding the experimental validation of natural motion, the dynamics of a robotic manipulator have been shown to affect user movements both for surgeons and novices as compared to freehand movements and it was hypothesised that the surgeons formed an internal representation of these dynamics, which allowed them to maneuver more smoothly [Nisky et al., 2014]. However, as natural motion reduces the work required to move the tool in general, it is expected that this reduction would occur regardless of the user's compensation strategy and the results for surgeons would be similar to those of non-surgeons.

The dynamic active constraints methodology demonstrated herein focused solely on the construction of the constraint from point clouds and the control strategy to enforce the constraint. While these are important aspects in creating a viable strategy for protecting soft tissue, to implement this type of controller in a clinical scenario, several other areas of research relating to sensing the soft tissue deformation and tracking the constraint region must be pursued.

The experiments utilised simple point markers to track the constraint surface. Data generated from 3D cameras and range finders can be very dense, slowing down construction of the implicit constraint surface. The constraint region must also be identified from the pre-operative plan, further slowing down its construction. In such cases, the data must be partitioned and possibly sub-sampled to ensure real-time applicability. Similar work has demonstrated real-time application of similar techniques [Leeper et al., 2012] and, therefore, I expect the problem to be solvable. To use imaging modalities which do not directly produce point clouds, strategies for choosing the appropriate point clouds must be developed. An algorithm which can, in real time, identify the critical area to protect and choose points in an appropriate density must be developed.

In addition, robust real-time soft tissue deformation registration solutions must be developed for current sensing modalities, such as US and CT, to enable the usage of these dynamic active constraints with these devices. At present, research is ongoing in these areas to improve accuracy in surgical procedures.

The methodology with which to define the constraint region has not been discussed. Current

clinical systems for rigid bone surgeries use preplanned geometries defined by the surgeon. It is likely a similar strategy could be employed to define the constraint region for dynamic active constraints from pre-operative imaging. Further research could enable additional constraints to be generated to automatically protect sensitive structures, such as blood vessels, using machine learning techniques on intra-operative imaging.

Issues relating to noise in the data produced from the real-time sensing and its effect on the accuracy of the dynamic active constraints surface and stability of the control algorithm under such conditions has not been examined. The impact of this noise would be dependent on relative magnitude of the noise and the accuracy required for the procedure at hand.

Finally, issues of occlusion have not been addressed within this thesis. Occlusion occurs when a tool or other item blocks the view of the region to be protected from the sensing device. Appropriate filtering techniques must be developed to detect occlusion, and prediction strategies may be used to estimate the constraint surface during these cases, in order to ensure the patient is protected.

The natural motion experiments presented here demonstrated a reduction in work when performing a point-to-point motion task. However, these results, while showing a reduction in work for general motions, do not directly reflect a reduction in work when performing a clinical procedure. Surgical tasks operate on a smaller scale and over a longer time period than is possible to perform on the Kuka LWR 4+. The true efficacy of this strategy in reducing surgeon fatigue and restoring dexterity remains to be shown in a clinical task.

Additionally, due the large torques required to rotate the Kuka end effector by a significant amount, the experimental task in Chapter 4 focused on translational motion only. The efficacy of the natural motion strategy in improving rotational quantities was demonstrated in simulation, however, the ability of the strategy to improve rotational quantities on a physical system must be shown. Furthermore, the ability to optimise rotational and translation quantities simultaneously was not investigated.

The use of natural motion within a surgical theater would require additional safety features incorporated into the strategy including joint limit and obstacle avoidance. The formalisation of natural motion as a null-space based gradient descent allows for the addition of other weighted

optimisation criteria. However, these additional features remain to be tested.

Lastly, all natural motion simulations and experiments were performed on the Kuka LWR 4+. Significant benefits were shown on this particular robot, however, the improvements the null-space strategy can provide are highly dependent on the kinematics and dynamics of the specific robot being optimised. The application of natural motion to various kind of redundant robots should be pursued to determine the general effectiveness of this strategy.

The final experiments of Chapter 6 were demonstrated in a virtual reality environment in which the feedback from the tool's interaction with the patient's body was simulated using a simple linear feedback law. This feedback provided a good preliminary validation of the proposed dynamic active constraints and their simultaneous application with natural motion optimisation. However, a more accurate simulation of this interaction which can, for example, include deformation of the patient's body due tool contact and feedback forces and torques which vary based on the organs and bone structures contacted by the tool, would provide a more realistic evaluation of the proposed strategies. Furthermore, the true efficacy of these methods could only be demonstrated in a clinical task performed on a real patient.

Bibliography

[Abbott et al., 2007] Abbott, J. J., Marayong, P., Okamura, A. M., Panadda, M., and Allison, O. (2007). Haptic Virtual Fixtures for Robot-Assisted Manipulation. In *Robotics Research*, volume 28, pages 49–64. Springer Berlin Heidelberg.

[ACTIVE, 2012] ACTIVE (2012). Active Constraint Technologies for Ill-Defined or Volatile Environments. `http://www.active-fp7.eu/`.

[Adamson and Alexa, 2003] Adamson, A. and Alexa, M. (2003). Approximating and Intersecting Surfaces from Points. In *Eurographics Symposium on Geometry Processing*.

[Aguirre-Ollinger et al., 2012] Aguirre-Ollinger, G., Colgate, J. E., Member, S., Peshkin, M. A., and Goswami, A. (2012). Inertia Compensation Control of a One-Degree-of- Freedom Exoskeleton for Lower-Limb Assistance : Initial Experiments. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 20(1):68–77.

[Alaraj et al., 2011] Alaraj, A., Lemole, M. G., Finkle, J. H., Yudkowsky, R., Wallace, A., Luciano, C., Banerjee, P. P., Rizzi, S. H., and Charbel, F. T. (2011). Virtual reality training in neurosurgery: Review of current status and future applications. *Surgical neurology international*, 2(52).

[Albu-Schäffer et al., 2007] Albu-Schäffer, a., Haddadin, S., Ott, C., Stemmer, A., Wimböck, T., and Hirzinger, G. (2007). The DLR lightweight robot: design and control concepts for robots in human environments. *Industrial Robot: An International Journal*, 34(5):376–385.

[Alexander, 1997] Alexander, J. (1997). Pain after laparoscopy. *British Journal of Anaesthesia*, 79:369–378.

[Allard et al., 2007] Allard, J., Cotin, S., Faure, F., Bensoussan, P.-J., Poyer, F., Duriez, C., Delingette, H., and Grisoni, L. (2007). SOFA - an open source framework for medical simulation. *MMVR 15 - Medicine Meets Virtual Reality*, 125:13–18.

[Ammar et al., 2010] Ammar, L. I., Kaddouh, B. Y., Mohanna, M. K., and Elhajj, I. H. (2010). SAS: SMA aiding sleeve. In *2010 IEEE International Conference on Robotics and Biomimetics, ROBIO 2010*, pages 1596–1599.

[Beasley, 2012] Beasley, R. A. (2012). Medical Robots: Current Systems and Research Directions. *Journal of Robotics*, 2012(1):1–14.

[Bebek and Çavusoglu, 2007] Bebek, O. and Çavusoglu, M. (2007). Intelligent control algorithms for robotic-assisted beating heart surgery. *IEEE Transactions on Robotics*, 23(3):468–480.

[Beckh et al., 2002] Beckh, S., Bolcsket, P., and Lessnau, K. (2002). Real-time chest ultrasonography: A comprehensive review for the pulmonologist. *Chest*, 122:1759–73.

[Bellekens et al., 2014] Bellekens, B., Spruyt, V., Berkvens, R., and Weyn, M. (2014). A Survey of Rigid 3D Pointcloud Registration Algorithms. In *The Fourth International Conference on Ambient Computing, Applications, Services, and Technologies*, pages 8–13.

[Benabid et al., 1987] Benabid, A., Cinquin, P., Lavalle, S., Le Bas, J., Demongeot, J., and de Rougemont, J. (1987). Computer-driven robot for stereotactic surgery connected to ct scan and magnetic resonance imaging. technological design and preliminary results. *Appl. Neurophysiol.*, 50:153–154.

[Bergeles and Yang, 2014] Bergeles, C. and Yang, G.-z. (2014). From Passive Tool Holders to Microsurgeons : Safer, Smaller, Smarter Surgical Robots. *IEEE Transactions on Bio-Medical Engineering*, 61(5):1565–1576.

[Bettini et al., 2004] Bettini, A., Marayong, P., Lang, S., Okamura, A. M., and Hager, G. D. (2004). Vision-assisted control for manipulation using virtual fixtures. *IEEE Transactions on Robotics*, 20(6):953–966.

[Blinn, 1982] Blinn, J. F. (1982). A generalization of algebraic surface drawing. *ACM SIGGRAPH Computer Graphics*, 16(3):273.

[Bona and Indri, 2005] Bona, B. and Indri, M. (2005). Friction compensation in robotics: An overview. In *Proceedings of the 44th IEEE Conference on Decision and Control, and the European Control Conference*, pages 4360–4367.

[Bowling and Khatib, 1998] Bowling, A. and Khatib, O. (1998). The motion isotropy hypersurface: a characterization of acceleration capability. In *1998 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 965–971.

[Bowyer, 2014] Bowyer, S. (2014). *Active Constraints for Robotic Surgery in Deforming Tissue*. PhD thesis, Imperial College London.

[Bowyer et al., 2014] Bowyer, S. A., Davies, B. L., and Rodriguez y Baena, F. (2014). Active Constraints / Virtual Fixtures : A Survey. *IEEE Transactions on Robotics*, 30(1):138–157.

[Bowyer and Rodriguez y Baena, 2014] Bowyer, S. A. and Rodriguez y Baena, F. (2014). Dynamic frictional constraints in translation and rotation. In *2014 IEEE International Conference on Robotics and Automation*, pages 2685–2692.

[Brisson et al., 2004] Brisson, G., Kanade, T., Digioia, A., and Jaramaz, B. (2004). Precision Freehand Sculpting of Bone. In *Medical Image Computing and Computer Assisted Intervation*, pages 105–112.

[Buerger and Hogan, 2007] Buerger, S. P. and Hogan, N. (2007). Complementary Stability and Loop Shaping for Improved Human Robot Interaction. *IEEE Transactions on Robotics*, 23(2):232–244.

[Cao et al., 2008] Cao, A., Thompson, R. C., Dumpuri, P., Dawant, B. M., Galloway, R. L., Ding, S., and Miga, M. I. (2008). Laser range scanning for image-guided neurosurgery: investigation of image-to-physical space registrations. *Medical physics*, 35(4):1593–1605.

[Castillo-Cruces and Wahrburg, 2010] Castillo-Cruces, R. a. and Wahrburg, J. (2010). Virtual fixtures with autonomous error compensation for humanrobot cooperative tasks. *Robotica*, 28(2):267.

[Cattilino et al., 2014] Cattilino, M., Secoli, R., Galvan, S., Forte, A. E., Dini, D., and Baena, F. R. (2014). Development of a Dynamic Soft Tissue Phantom for Cooperative Control Testing in Robotic Surgery. In *Hamltn Symposium*.

[Chen et al., 2010] Chen, S. J. S., Hellier, P., Gauvrit, J. Y., Marchal, M., Morandi, X., and Collins, D. L. (2010). An anthropomorphic polyvinyl alcohol triple-modality brain phantom based on Colin27. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 6362 LNCS(PART 2):92–100.

[Chiacchio, 2000] Chiacchio, P. (2000). A new dynamic manipulability ellipsoid for redundant manipulators A new dynamic manipulability ellipsoid for redundant manipulators. *Robotica*, 18(04):381–387.

[Chiu, 1988] Chiu, S. (1988). Task compatibility of manipulator postures. *The International Journal of Robotics Research*, 7(5):13–21.

[Choi et al., 2012] Choi, H.-b., Ryu, J., and Ieee, M. (2012). Convex Hull-based Power Manipulability Analysis of Robot Manipulators. In *2012 IEEE International Conference on Robotics and Automation*, pages 2972–2977.

[Choi and Kim, 1999] Choi, S. I. and Kim, B. K. (1999). Obstacle Avoidance Control for Redundant Manipulators Using Collidability Measure. In *1999 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1816–1821.

[Citak et al., 2013] Citak, M., Suero, E. M., Citak, M., Dunbar, N. J., Branch, S. H., Conditt, M. a., Banks, S. a., and Pearle, A. D. (2013). Unicompartmental knee arthroplasty: is robotic technology more accurate than conventional technique? *The Knee*, 20(4):268–71.

[Cobb et al., 2006] Cobb, J., Henckel, J., Gomes, P., Harris, S., Jakopec, M., Rodriguez, F., Barrett, A., and Davies, B. (2006). Hands-on robotic unicompartmental knee replacement: a prospective, randomised controlled study of the acrobot system. *The Journal of Bone and Joint Surgery - British Volume*, 88(2):188–97.

[Colgate and Hogan, 1989] Colgate, E. and Hogan, N. (1989). An Analysis of Contact Instability in Terms of Passive Physical Equivalents. In *1989 IEEE International Conference on Robotics and Automation*, pages 404–409.

[Constantinescu et al., 2005] Constantinescu, D., Salcudean, S. E., and Croft, E. a. (2005). Haptic rendering of rigid contacts using impulsive and penalty forces. *IEEE Transactions on Robotics*, 21(3):309–323.

[Cossu et al., 2005] Cossu, M., Cardinale, F., Castana, L., Citterio, A., Francione, S., Tassi, L., Benabid, A. L., and Lo Russo, G. (2005). Stereoelectroencephalography in the presurgical evaluation of focal epilepsy: a retrospective analysis of 215 procedures. *Neurosurgery*, 57(4):706–718.

[Dahl, 1968] Dahl, P. R. (1968). Solid Friction Model. Technical report.

[Dargar et al., 2015] Dargar, S., Kennedy, R., Lai, W., Arikatla, V., and De, S. (2015). Towards immersive virtual reality (ivr): a route to surgical expertise. *Journal of Computational Surgery*, 2(1):2.

[Davies, 1996] Davies, B. (1996). *A discussion of safety issues for medical robots*, pages 287–296. MIT Press.

[Davies et al., 1999] Davies, B., Harris, S., Jakopec, M., and Cobb, J. (1999). Novel hands-on robot for knee replacement. In *European Robotic Systems Workshop on Medical Robotics*, pages 98–102.

[Davies et al., 2004] Davies, B. L., Harris, S. J., Rodriguez, F., Gomes, P., and Jakopec, M. (2004). Hands-On Robotic Surgery : Is This the Future ? In *Medical Imaging and Augmented Reality*, pages 27–37. Springer-Verlag Berlin Heidelberg.

[Davies et al., 1991] Davies, B. L., Hibberd, R. D., Ng, W. S., Timoney, A. G., and Wickham, J. E. A. (1991). The Development of a Surgeon Robot for Prostatectomies. *Proceedings of the Institution of Mechanical Engineers , Part H : Journal of Engineering in Medicine*, 205(1):35–38.

[Dhyani et al., 2014] Dhyani, M., Gilbertson, M., Anvari, A., Anthony, B., and Samir, A. (2014). Precise Quantification of Sonographic Forces: A First Step Toward Reducing Ergonomic Injury. *AIUM Annual Converntion, J Ultrasound Med*, 33(suppl):.

[Ding et al., 2011] Ding, S., Miga, M. I., Pheiffer, T. S., Simpson, A. L., Thompson, R. C., and Dawant, B. M. (2011). Tracking of vessels in intra-operative microscope video sequences for cortical displacement estimation. *IEEE Transactions on Biomedical Engineering*, 58(7):1985–1993.

[Doelken and Mayo, 2008] Doelken, P. and Mayo, P. (2008). Thoracic ultrasound: Indications, advantages, and technique. `http://www.uptodate.com/contents/thoracic-ultrasound-indications-advantages-and-technique`. Accessed: 2015-05-22.

[Dogangil et al., 2010] Dogangil, G., Davies, B. L., Baena, F. R., and Rodriguez y Baena, F. (2010). A review of medical robotics for minimally invasive soft tissue surgery. *Proceedings of the Institution of Mechanical Engineers. Part H, Journal of Engineering in Medicine*, 224(5):653–679.

[Dormand and Prince, 1980] Dormand, J. and Prince, P. (1980). A family of embedded Runge-Kutta formulae. *Journal of Computational and Applied Mathematics*, 6(1):19–26.

[Doty et al., 1995] Doty, K. L., Melchiom, C., Schwartz, E. M., and Bonivento, C. (1995). Robot Manipulability. *IEEE Transactions on Robotics and Automation*, 11(3):462–468.

[D'Souza et al., 2001] D'Souza, W. D., Madsen, E. L., Unal, O., Vigen, K. K., Frank, G. R., and Thomadsen, B. R. (2001). Tissue mimicking materials for a multi-imaging modality prostate phantom. *Medical physics*, 28(4):688–700.

[Dupont et al., 2002] Dupont, P., Hayward, V., Armstrong, B., and Altpeter, F. (2002). Single State Elastoplastic Friction Models. *IEEE Transactions on Automatic Control*, 47(5):787–792.

[Ericson, 2004] Ericson, C. (2004). *Real-Time Collision Detection*. Number v. 1 in Morgan Kaufmann series in interactive 3D technology. Taylor & Francis.

[Faria et al., 2014] Faria, C., Sadowsky, O., Bicho, E., Ferrigno, G., Joskowicz, L., Shoham, M., Vivanti, R., and De Momi, E. (2014). Validation of a stereo camera system to quantify brain deformation due to breathing and pulsatility. *Medical Physics*, 41(113502):1–11.

[Farina and Merletti, 2000] Farina, D. and Merletti, R. (2000). Comparison of algorithms for estimation of EMG variables during voluntary isometric contractions. *Journal of Electromyography and Kinesiology*, 10(5):337–349.

[Faulkner et al., 2015] Faulkner, A. R., Bourgeois, A. C., Bradley, Y. C., and Pasciak, A. S. (2015). A robust and inexpensive phantom for fluoroscopically guided lumbar puncture training. *Simul. Healthc.*, 10:54–58.

[Feller-Kopman, 2007] Feller-Kopman, D. (2007). Therapeutic thoracentesis: the role of ultrasound and pleural manometry. *Current opinion in pulmonary medicine*, 13(4):312–318.

[Ferrant et al., 2002] Ferrant, M., Nabavi, A., Macq, B., Black, P. M., Jolesz, F. A., Kikinis, R., and Warfield, S. K. (2002). S erial registration of intraoperative MR images of the brain. *Medical Image Analysis*, 6:337–359.

[Ficuciello et al., 2014] Ficuciello, F., Romano, A., Villani, L., and Siciliano, B. (2014). Cartesian Impedance Control of Redundant Manipulators for Human-Robot Co-manipulation. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2120–2125. Ieee.

[Forte et al., 2014] Forte, A. E., Galvan, S., Manieri, F., and Dini, D. (2014). A novel composite phantom for brain tissue. In *Annual International IEEE EMBS Conference*.

[Francis et al., 2001] Francis, N. K., Hanna, G. B., Cresswell, A. B., Carter, F. J., and Cuschieri, A. (2001). The performance of master surgeons on standard aptitude testing. *American Journal of Surgery*, 182(1):30–33.

[Gautier et al., 2013] Gautier, M., Jubien, A., Janot, A., and Robet, P.-p. (2013). Dynamic Identification of flexible joint manipulators with an efficient closed loop output error method based on motor torque output data. In *2013 IEEE Conference on Robotics and Automation*, pages 2934–2940.

[Gibo et al., 2009] Gibo, T. L., Verner, L. N., Yuh, D. D., and Okamura, A. M. (2009). Design considerations and human-machine performance of moving virtual fixtures. In *2009 IEEE International Conference on Robotics and Automation*, pages 671–676. Ieee.

[Ginhoux et al., 2005] Ginhoux, R., Gangloff, J., Mathelin, M. D., Member, S., Soler, L., Sanchez, M. M. A., and Marescaux, J. (2005). Active Filtering of Physiological Motion in Robotized Surgery Using Predictive Control. *IEEE Transactions on Robotics*, 21(1):67–79.

[Gopura et al., 2011] Gopura, R. a. R. C., Kiguchi, K., and Bandara, D. S. V. (2011). A brief review on upper extremity robotic exoskeleton systems. *2011 6th International Conference on Industrial and Information Systems, ICIIS 2011 - Conference Proceedings*, 8502:346–351.

[Gopura et al., 2009] Gopura, R. a. R. C., Kiguchi, K., and Li, Y. (2009). SUEFUL-7: A 7DOF upper-limb exoskeleton robot with muscle-model-oriented EMG-based control. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1126–1131.

[Graettinger and Krogh, 1988] Graettinger, T. J. and Krogh, B. H. (1988). The Acceleration Radius : A Global Performance Measure for Robotic Manipulators. *IEEE Journal of Robotics and Automation*, 4(I):60–69.

[Gryminski et al., 1976] Gryminski, J., Krakowka, P. D., and Lypacewicz, G. (1976). The Diagnosis of Pleural Effusion by Ultrasonic and Radiologic Techniques. *Chest*, 70(1):33–37.

[Guthart and Salisbury, 2000] Guthart, G. S. and Salisbury, J. K. (2000). The Intuitive Telesurgery System : Overview and Application. In *2000 IEEE International Conference on Robotics and Automation*, pages 618–621.

[Hagag et al., 2011] Hagag, B., Abovitz, R., Kang, H., Schmitz, B., and Conditt, M. (2011). Rio: Robotic-arm interactive orthopedic system makoplasty: User interactive haptic orthopedic robotics. In Rosen, J., Hannaford, B., and Satava, R. M., editors, *Surgical Robotics*, pages 219–246. Springer US.

[Hager et al., 2008] Hager, B. Y. G. D., Okamura, A. M., Kazanzides, P., Whitcomb, L. L., Fichtinger, G., and Taylor, R. H. (2008). Surgical and Interventional Robotics: Part III. *IEEE Robotics and Automation Magazine*, 15(4):84–93.

[Hagn et al., 2010] Hagn, U., Konietschke, R., Tobergte, A., Nickl, M., Jörg, S., Kübler, B., Passig, G., Gröger, M., Fröhlich, F., Seibold, U., Le-Tien, L., Albu-Schäffer, A., Nothhelfer, A., Hacker, F., Grebenstein, M., and Hirzinger, G. (2010). DLR MiroSurge: a versatile

system for research in endoscopic telesurgery. *International Journal of Computer Assisted Radiology and Surgery*, 5(2):183–93.

[Hagn et al., 2008] Hagn, U., Nickl, M., Jörg, S., Passig, G., Bahls, T., Nothhelfer, a., Hacker, F., Le-Tien, L., Albu-Schäffer, a., Konietschke, R., Grebenstein, M., Warpup, R., Haslinger, R., Frommberger, M., and Hirzinger, G. (2008). The DLR MIRO: a versatile lightweight robot for surgical applications. *Industrial Robot: An International Journal*, 35(4):324–336.

[Hannaford et al., 2013] Hannaford, B., Rosen, J., Friedman, D. W., King, H., Roan, P., Cheng, L., Glozman, D., Ma, J., Kosari, S. N., and White, L. (2013). Raven-II: An open platform for surgical robotics research. *IEEE Transactions on Biomedical Engineering*, 60(4):954–959.

[Harris et al., 1997] Harris, S. J., Arambula-Cosio, F., Mei, Q., Hibberd, R. D., Davies, B. L., Wickham, J. E. a., Nathan, M. S., and Kundu, B. (1997). The Probot–an active robot for prostate resection. *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine*, 211(4):317–325.

[Havelock et al., 2010] Havelock, T., Teoh, R., Laws, D., and Gleeson, F. (2010). Pleural procedures and thoracic ultrasound: British Thoracic Society Pleural Disease Guideline 2010. *Thorax*, 65 Suppl 2:ii61–ii76.

[Hayward et al., 2009] Hayward, V., Armstrong, B. S. R., Member, S., Altpeter, F., Dupont, P. E., and Member, S. (2009). Discrete-Time Elasto-Plastic Friction Estimation. *IEEE Transactions on Control Systems Technology*, 17(3):688–696.

[He et al., 2012] He, X., Roppenecker, D., Gierlach, D., Balicki, M., Olds, K., Gehlbach, P., Honda, J., Taylor, R., and Iordachita, I. (2012). Toward Clinically Applicable Steady-Hand Eye Robot for Vitreoretinal Surgery. In *ASME 2012 International Mechanical Engineering Congress and Exposition*, pages 1–9.

[Health Quality Ontario, 2010] Health Quality Ontario (2010). Robotic-assisted minimally invasive surgery for gynecologic and urologic oncology: an evidence-based analysis. *Ontario Health Technology Assessment Series*, 10(27):1–118.

[Hennekens et al., 2008] Hennekens, D., Constantinescu, D., and Steinbuch, M. (2008). Continuous Impulsive Force Controller for Forbidden-Region Virtual Fixtures. In *2008 IEEE*

International Conference on Robotics and Automation, pages 2890–2895.

[Ho et al., 1995] Ho, C., Hibberd, R. D., and Davies, B. L. (1995). Robot Assisted Knee Surgery. *IEEE Engineering in Medicine and Biology*, 14(3):292–300.

[Hogan, 1984] Hogan, N. (1984). Impedance Control: An Approach to Manipulation. In *American Control Conference*, pages 304–313.

[Holden et al., 2014] Holden, M. S., Ungi, T., Sargent, D., McGraw, R. C., Chen, E. C. S., Ganapathy, S., Peters, T. M., and Fichtinger, G. (2014). Feasibility of real-time workflow segmentation for tracked needle interventions. *IEEE Transactions on Biomedical Engineering*, 61(6):1720–1728.

[Hundtofte et al., 2002] Hundtofte, C., Hager, G., and Okamura, A. (2002). Building a task language for segmentation and recognition of user input to cooperative manipulation systems. In *Proceedings 10th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. HAPTICS 2002.*

[Ikeura and Inooka, 1995] Ikeura, R. and Inooka, H. (1995). Variable impedance control of a robot for cooperation with a human. In *1995 IEEE International Conference on Robotics and Automation*, pages 3097–3102.

[Ikeura et al., 1994] Ikeura, R., Monden, H., and Inooka, H. (1994). Cooperative Motion Control of a Robot and a Human. In *1994 IEEE International Workshop on Robot and Human Communication*, pages 112–117.

[Ikeura et al., 2002] Ikeura, R., Moriguchi, T., and Mizutani, K. (2002). Optimal Variable Impedance Control for a Robot and Its Application To Lifting an Object with a Human. In *2002 IEEE International Workshop on Robot and Human Interactive Communication*, pages 500–505.

[Intuitive Surgical Ltd., 2015] Intuitive Surgical Ltd. (2015). Intuitive Surgical Ltd. `http://www.intuitivesurgical.com/`.

[Jakopec et al., 2003] Jakopec, M., Harris, S. J., Baena, F. R. Y., Gomes, P., and Davies, B. L. (2003). The Acrobot system for total knee replacement. *Industrial Robot: An International*

Journal, 30(1):61–66.

[Ji et al., 2014] Ji, S., Fan, X., Roberts, D. W., Hartov, A., and Paulsen, K. D. (2014). Cortical surface shift estimation using stereovision and optical flow motion tracking via projection image registration. *Medical Image Analysis*, 18(7):1169–1183.

[Joice et al., 1998] Joice, P., Hanna, G. B., and Cuschieri, A. (1998). Errors enacted during endoscopic surgery–a human reliability analysis. *Applied Ergonomics*, 29(6):409–414.

[Jong Yoon et al., 2013] Jong Yoon, W., Velasquez, C. a., White, L. W., Hannaford, B., Sang Kim, Y., Lendvay, T. S., and Yoon, W. J. (2013). Preliminary Articulable Probe Designs With RAVEN and Challenges: Image-Guided Robotic Surgery Multitool System. *Journal of Medical Devices*, 8(1):1–6.

[Jordan et al., 2001] Jordan, J. a., Gallagher, a. G., McGuigan, J., and McClure, N. (2001). Virtual reality training leads to faster adaptation to the novel psychomotor restrictions encountered by laparoscopic surgeons. *Surgical Endoscopy*, 15(10):1080–1084.

[Kazanzides et al., 1992] Kazanzides, P., Zuhars, J., Mittelstadt, B., and Taylor, R. H. (1992). Force Sensing and Control for a Surgical Robot. In *1992 IEEE International Conference on Robotics and Automation*, pages 612–617.

[Kazerooni et al., 2005] Kazerooni, H., Racine, J.-l., Huang, L., and Steger, R. (2005). On the Control of the Berkeley Lower Extremity Exoskeleton (BLEEX). In *2005 IEEE International Conference on Robotics and Automation*, pages 4353–4360.

[Kazerooni and Snyder, 1995] Kazerooni, H. and Snyder, T. J. (1995). Case Study on Haptic Devices - Human-Induced Instability in Powered Hand Controllers. *Journal of Guidance, Control, and Dynamics*, 18(1):108–113.

[Kerkhof et al., 2009] Kerkhof, E. M., van der Put, R. W., Raaymakers, B. W., van der Heide, U. a., Jürgenliemk-Schulz, I. M., and Lagendijk, J. J. W. (2009). Intrafraction motion in patients with cervical cancer: The benefit of soft tissue registration using MRI. *Radiotherapy and Oncology*, 93(1):115–121.

[Kermani et al., 2007] Kermani, M. R., Patel, R. V., and Moallem, M. (2007). Friction Identification and Compensation in Robotic Manipulators. *IEEE Transactions on Instrumentation and Measurement*, 56(6):2346–2353.

[Khatib, 1987] Khatib, O. (1987). A Unified Approach for Motion and Force Control of Robot Manipulators. *IEEE Journal of Robotics and Automation*.

[Khatib, 1995] Khatib, O. (1995). Inertial properties in robotic manipulation: An object-level framework. *The International Journal of Robotics Research*, 14(1):19–36.

[Khatib and Burdick, 1987] Khatib, O. and Burdick, J. (1987). Optimization of dynamics in manipulator design: The operational space formulation. *International Journal of Robotics and Automation*, 2(2):90–98.

[Kikuuwe et al., 2008] Kikuuwe, R., Takesue, N., and Fujimoto, H. (2008). A control framework to generate nonenergy-storing virtual fixtures: Use of simulated plasticity. *IEEE Transactions on Robotics*, 24(4):781–793.

[Kikuuwe et al., 2006] Kikuuwe, R., Takesue, N., Sano, A., Mochiyama, H., and Fujimoto, H. (2006). Admittance and impedance representations of friction based on implicit Euler integration. *IEEE Transactions on Robotics*, 22(6):1176–1188.

[Klein and Huang, 1983] Klein, C. A. and Huang, C.-H. (1983). Review of Pseudoinverse Control for Use with Kinematically Redundant Manipulators. *IEEE Transactions on Systems, Man, and Cybernetics*, 13(3):245–250.

[Kobayashi et al., 2013] Kobayashi, Y., Minami, M., Yanou, A., Maeba, T., and Manipulability, A. D. (2013). Dynamic Reconfiguration Manipulability Analyses of Humanoid Bipedal Walking. In *2013 IEEE Conference on Robotics and Automation*, pages 4764–4769.

[Konietschke et al., 2009] Konietschke, R., Hagn, U., Nickl, M., Stefan, J., Tobergte, A., Passig, G., Seibold, U., Le-tien, L., Bernhard, K., Gr, M., Fr, F., Rink, C., Albu-sch, A., Grebenstein, M., Ortmaier, T., and Hirzinger, G. (2009). The DLR MiroSurge A Robotic System for Surgery. In *2009 IEEE International Conference on Robotics and Automation*, pages 1589–1590.

[Konietschke et al., 2006] Konietschke, R., Ortmaier, T., Ott, C., Hagn, U., Le-Tien, L., and Hirzinger, G. (2006). Concepts of Human-Robot Cooperation for a New Medical Robot. In *Second International Workshop on Human Centered Robotic Systems*, pages 1–6.

[Kosuge and Kazamura, 1997] Kosuge, K. and Kazamura, N. (1997). Control of a Robot Handling an Object in Cooperation with a Human. In *1997 IEEE International Workshop on Robot and Human Communication*, pages 142–147.

[Koulalis et al., 2011] Koulalis, D., O'Loughlin, P. F., Plaskos, C., Kendoff, D., Cross, M. B., and Pearle, A. D. (2011). Sequential versus automated cutting guides in computer-assisted total knee arthroplasty. *The Knee*, 18(6):436–42.

[Kumar et al., 2000] Kumar, R., Berkelman, P., Gupta, P., Barnes, A., Jensen, P. S., Whitcomb, L. L., and Taylor, R. H. (2000). Preliminary Experiments in Cooperative Human / Robot Force Control for Robot Assisted Microsurgical Manipulation. In *2000 IEEE International Conference on Robotics and Automation*, pages 610–617.

[Kumar et al., 1999] Kumar, R., Jensen, P., and Taylor, R. H. (1999). Experiments with a Steady Hand Robot in Constrained Compliant Motion and Path Following. In *1999 IEEE International Workshop on Robot and Human Interaction*, pages 92–97.

[Kwoh et al., 1988] Kwoh, Y. S., Hou, J., Jonckheere, E. A., and Hayati, S. (1988). A Robot with Improved Absolute Positioning Accuracy for CT Guided Stereotactic Brain Surgery. *IEEE Transactions on Bio-Medical Engineering*, 35(2):153–160.

[Lamouche et al., 2012] Lamouche, G., Kennedy, B. F., Kennedy, K. M., Bisaillon, C.-E., Curatolo, A., Campbell, G., Pazos, V., and Sampson, D. D. (2012). Review of tissue simulating phantoms with controllable optical, mechanical and structural properties for use in optical coherence tomography. *Biomedical Optics Express*, 3(6):1381.

[Lang et al., 2011] Lang, J. E., Mannava, S., Floyd, a. J., Goddard, M. S., Smith, B. P., Mofidi, A., M. Seyler, T., and Jinnah, R. H. (2011). Robotic systems in orthopaedic surgery. *Journal of Bone and Joint Surgery - British Volume*, 93-B(1):1296–1299.

[Lawson et al., 2007] Lawson, J. D., Schreibmann, E., Jani, A. B., and Fox, T. (2007). Quantitative evaluation of a cone-beam computed tomography-planning computed tomography

deformable image registration method for adaptive radiation therapy. *Journal of applied clinical medical physics / American College of Medical Physics*, 8(4):2432.

[Lee et al., 2012] Lee, H., Kim, W., Han, J., and Han, C. (2012). The technical trend of the exoskeleton robot system for human power assistance. *International Journal of Precision Engineering and Manufacturing*, 13(8):1491–1497.

[Leeper et al., 2012] Leeper, A., Chan, S., and Salisbury, K. (2012). Point clouds can be represented as implicit surfaces for constraint-based haptic rendering. In *2012 IEEE International Conference on Robotics and Automation*, pages 5000–5005.

[Li et al., 2007] Li, M., Ishii, M., and Taylor, R. H. (2007). Spatial Motion Constraints Using Fixtures Generated by Anatomy. *IEEE Transactions on Robotics*, 23(1):4–19.

[Li and Okamura, 2003] Li, M. and Okamura, A. M. (2003). Recognition of operator motions for real-time assistance using virtual fixtures. In *11th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 125–131. IEEE Comput. Soc.

[Li et al., 2002] Li, Q. H., Zamorano, L., Pandya, A., Perez, R., Gong, J., and Diaz, F. (2002). The application accuracy of the NeuroMate robot - A quantitative comparison with frameless and frame-based surgical localization systems. *Computer Aided Surgery*, 7(2):90–98.

[Liegeois, 1977] Liegeois, A. (1977). Automatic supervisory control of the configuration and behavior of multibody mechanisms. *IEEE Transactions on Systems, Man, and Cybernetics*, 7(12):868–871.

[Lindh et al., 2009] Lindh, W., Pooler, M., Tamparo, C., and Dahl, B. (2009). *Delmar's Comprehensive Medical Assisting: Administrative and Clinical Competencies*, page 573. Cengage Learning.

[Liu and Curet, 2015] Liu, M. and Curet, M. (2015). A review of training research and virtual reality simulators for the da vinci surgical system. *Teaching and Learning in Medicine*, 27(1):12–26.

[Lonner and Kerr, 2012] Lonner, J. H. and Kerr, G. J. (2012). Robotically Assisted Unicompartmental Knee Arthroplasty. *Operative Techniques in Orthopaedics*, 22(4):182–188.

[Lum et al., 2009] Lum, M. J. H., Friedman, D. C. W., Sankaranarayanan, G., King, H., Fodero, K., Leuschke, R., Hannaford, B., Rosen, J., and Sinanan, M. N. (2009). The RAVEN: Design and Validation of a Telesurgery System. *The International Journal of Robotics Research*, 28(9):1183–1197.

[Ma and Angeles, 1990] Ma, O. and Angeles, J. (1990). The concept of dynamic isotropy and its applications to inverse kinematics and trajectory planning. In *1990 IEEE International Conference on Robotics and Automation*, pages 481–486.

[Maciejewski and Klein, 1985] Maciejewski, A. A. and Klein, C. A. (1985). Obstacle Avoidance for Kinematically Redundant Manipulators in Dynamically Varying Environments. *The International Journal of Robotics Research*, 4(3):109–116.

[Maclachlan et al., 2012] Maclachlan, R. A., Becker, B. C., Member, S., Tabar, J. C., Podnar, G. W., Lobes, L. A., Riviere, C. N., and Member, S. (2012). Micron: An Actively Stabilized Handheld Tool for Microsurgery. *IEEE Transactions on Robotics*, 28(1):195–212.

[Mahalil et al., 2014] Mahalil, I., Rusli, M., Yusof, A., Mohd Yusoff, M., and Razieff Zainudin, A. (2014). Study of immersion effectiveness in vr-based stress therapy. In *Information Technology and Multimedia (ICIMU), 2014 International Conference on*, pages 380–384.

[Mansouri and Ouali, 2011] Mansouri, I. and Ouali, M. (2011). The power manipulability A new homogeneous performance index of robot manipulators. *Robotics and Computer-Integrated Manufacturing*, 27(2):434–449.

[Marchal-Crespo and Reinkensmeyer, 2009] Marchal-Crespo, L. and Reinkensmeyer, D. J. (2009). Review of control strategies for robotic movement training after neurologic injury. *Journal of Neuroengineering and Rehabilitation*, 6(20):1–15.

[Mercaldi and Lanes, 2013] Mercaldi, C. J. and Lanes, S. F. (2013). Ultrasound guidance decreases complications and improves the cost of care among patients undergoing thoracentesis and paracentesis. *Chest*, 143(2):532–538.

[Minami and Takahara, 2003] Minami, M. and Takahara, M. (2003). Avoidance manipulability for redundant manipulators. In *2003 IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, pages 314–319.

[Mirk et al., 1999] Mirk, P., Magnavita, N., Masini, L., Bazzocchi, M., and Fileni, A. (1999). Frequency of musculoskeletal symptoms in diagnostic medical sonographers. results of a pilot survey. *Radiol. Med.*, 98:236–241.

[Mitchell et al., 2007] Mitchell, B., Koo, J., Iordachita, I., Kazanzides, P., Kapoor, A., Handa, J., Hager, G., and Taylor, R. (2007). Development and Application of a New Steady-Hand Manipulator for Retinal Surgery. In *2007 IEEE International Conference on Robotics and Automation*, pages 623–629.

[Morse et al., 2005] Morse, B. S., Yoo, T. S., Rheingans, P., Chen, D. T., and Subramanian, K. R. (2005). Interpolating Implicit Surfaces From Scattered Surface Data Using Compactly Supported Radial Basis Functions. In *ACM SIGGRAPH*.

[Nagendran et al., 2013] Nagendran, M., Gurusamy, K. S., Aggarwal, R., Loizidou, M., and Davidson, B. R. (2013). Virtual reality training for surgical trainees in laparoscopic surgery. *Cochrane Database of Systematic Reviews*, 8(8).

[Nakamura et al., 2001] Nakamura, Y., Kishi, K., and Kawakami, H. (2001). Heartbeat synchronization for robotic cardiac surgery. In *2001 IEEE International Conference on Robotics and Automation*, pages 2014–2019.

[Navkar et al., 2012] Navkar, N. V., Deng, Z., Shah, D. J., Bekris, K. E., and Tsekos, N. V. (2012). Visual and force-feedback guidance for robot-assisted interventions in the beating heart with real-time MRI. In *2012 IEEE International Conference on Robotics and Automation*, pages 689–694. Ieee.

[Nawabi et al., 2013] Nawabi, D. H., Conditt, M. a., Ranawat, a. S., Dunbar, N. J., Jones, J., Banks, S., and Padgett, D. E. (2013). Haptically guided robotic technology in total hip arthroplasty: a cadaveric investigation. *Proceedings of the Institution of Mechanical Engineers. Part H - Journal of Engineering in Medicine*, 227(3):302–309.

[Nisky et al., 2014] Nisky, I., Okamura, A. M., and Hsieh, M. H. (2014). Effects of robotic manipulators on movements of novices and surgeons. *Surgical Endoscopy and Other Interventional Techniques*, 28(1):2145–2158.

[Okamura, 2009] Okamura, A. M. (2009). Haptic feedback in robot-assisted minimally invasive surgery. *Curr. Opin. Urol.*, 19:102–107.

[Okamura et al., 2010] Okamura, a. M., Verner, L. N., Reiley, C. E., and Mahvash, M. (2010). Haptics for robot-assisted minimally invasive surgery. *Springer Tracts in Advanced Robotics*, 66:361–372.

[Olsson et al., 1998] Olsson, H., Åström, K. J., Canudas De Wit, C., Gäfvert, M., and Lischinsky, P. (1998). Friction Models and Friction Compensation. *European Journal of Control*, 4(1):176–195.

[Ortmaier et al., 2005] Ortmaier, T., Gröger, M., Boehm, D. H., Falk, V., and Hirzinger, G. (2005). Motion estimation in beating heart surgery. *IEEE Transactions on Bio-Medical Engineering*, 52(10):1729–40.

[Ortmaier et al., 2006] Ortmaier, T., Weiss, H., Hagn, U., Grebenstein, M., Nickl, M., Albu-Schaffer, A., Ott, C., Jorg, S., Konietschke, R., and Hirzinger, G. (2006). A hands-on-robot for accurate placement of pedicle screws. In *2006 IEEE International Conference on Robotics and Automation*, pages 4179–4186.

[Park et al., 2008] Park, D., Hoffmann, H., and Schaal, S. (2008). Combining dynamic movement primitives and potential fields for online obstacle avoidance. *... Motion of Animals and ...*, pages 2–3.

[Patel et al., 2012] Patel, P. A., Ernst, F. R., and Gunnarsson, C. L. (2012). Ultrasonography guidance reduces complications and costs associated with thoracentesis procedures. *Journal of Clinical Ultrasound*, 40(3):135–141.

[Paul et al., 1992] Paul, H. A., Bargar, W. L., Mittlestadt, B., Musits, B., Taylor, R. H., Kazanzides, P., Zuhars, J., Williamson, B., and Hanson, W. (1992). Development of a Surgical Robot for Cementless Total Hip Athropasty. *Clinical Orthopaedics and Related Research*, 285(1):57–66.

[Pausch et al., 1997] Pausch, R., Proffitt, D., and Williams, G. (1997). Quantifying immersion in virtual reality. In *Proceedings of the 24th Annual Conference on Computer Graph-*

ics and Interactive Techniques, SIGGRAPH '97, pages 13–18, New York, NY, USA. ACM Press/Addison-Wesley Publishing Co.

[Petersen and Rodriguez Baena, 2013] Petersen, J. G. and Rodriguez Baena, F. (2013). A dynamic active constraints approach for hands-on robotic surgery. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1966–1971.

[Pezzementi et al., 2007] Pezzementi, Z., Okamura, A. M., and Hager, G. D. (2007). Dynamic guidance with pseudoadmittance virtual fixtures. In *2007 IEEE International Conference on Robotics and Automation*, pages 1761–1767.

[Philippsen et al., 2011] Philippsen, R., Sentis, L., and Khatib, O. (2011). An Open Source Extensible Software Package to Create Whole-Body Compliant Skills in Personal Mobile Manipulators. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1036–1041.

[Plaskos et al., 2005] Plaskos, C., Cinquin, P., Lavallée, S., and Hodgson, a. J. (2005). Praxiteles: a miniature bone-mounted robot for minimal access total knee arthroplasty. *The International Journal of Medical Robotics and Computer Assisted Surgery*, 1(4):67–79.

[Pransky, 1997] Pransky, J. (1997). ROBODOC surgical robot success story. *Industrial Robot*, 24(3):231–233.

[Prasad et al., 2004] Prasad, S. M. S. M., Maniar, H. S., Chu, C., Schuessler, R. B., and Damiano, R. J. (2004). Surgical robotics: impact of motion scaling on task performance. *Journal of the American College of Surgeons*, 199(6):863–8.

[Puangmali et al., 2008] Puangmali, P., Althoefer, K., Seneviratne, L. D., Murphy, D., and Dasgupta, P. (2008). State-of-the-art in force and tactile sensing for minimally invasive surgery. *IEEE Sensors Journal*, 8(4):371–380.

[Rahman et al., 1999] Rahman, M., Ikeura, R., and Mizutani, K. (1999). Investigating the impedance characteristic of human arm for development of robots to co-operate with human operators. In *1999 IEEE International Conference on Systems, Man, and Cybernetics*, pages 676–681.

[Rahman and Ikeura, 2010] Rahman, S. and Ikeura, R. (2010). Controlling a power assist robot for lifting objects considering human's unimanual, bimanual and cooperative weight perception. In *2010 IEEE International Conference on Robotics and Automation*, pages 2356–2362.

[Reardon, 2015] Reardon, R. (2015). Thoracic ultrasound for pneumothorax. `http://hqmeded.com/thoracic-ultrasound-for-pneumothorax-2/`. Accessed: 2015-05-22.

[Reinertsen and Collins, 2006] Reinertsen, I. and Collins, D. L. (2006). A realistic phantom for brain-shift simulations. *Medical physics*, 33(9):3234–3240.

[Ren et al., 2008] Ren, J., Patel, R. V., McIsaac, K. a., Guiraudon, G., and Peters, T. M. (2008). Dynamic 3-D virtual fixtures for minimally invasive beating heart procedures. *IEEE Transactions on Medical Imaging*, 27(8):1061–70.

[Riviere and Thakor, 1996] Riviere, C. and Thakor, N. (1996). Modeling and canceling tremor in human-machine interfaces. *IEEE Engineering in Medicine and Biology Magazine*, 15(3):29–36.

[Riviere et al., 1998] Riviere, C. N., Rader, R. S., and Thakor, N. V. (1998). Adaptive canceling of physiological tremor for improved precision in microsurgery. *IEEE Transactions on Bio-Medical Engineering*, 45(7):839–46.

[Rosen et al., 2006] Rosen, J., Brown, J. D., Chang, L., Sinanan, M. N., and Hannaford, B. (2006). Generalized approach for modeling minimally invasive surgery as a stochastic process using a discrete Markov model. *IEEE Transactions on Biomedical Engineering*, 53(3):399–413.

[Rosenberg, 1992] Rosenberg, L. B. (1992). The Use of Virtual Fixtures as Perceptual Overlays to Enhance Operator Performance in Remote Environments. Technical report, Stanford University.

[Rosenstein and Grupen, 2002] Rosenstein, M. and Grupen, R. (2002). Velocity-dependent dynamic manipulability. In *2002 IEEE International Conference on Robotics and Automation*, pages 2424–2429.

[Rotella, 2005] Rotella, J. (2005). *Predictive Tracking of Quasi Periodic Signals for Active Relative Motion Cancellation in Robotic Assisted Coronary Artery Bypass Graft Surgery.* PhD thesis, Case Western Reserve University.

[Ruurda et al., 2004] Ruurda, J. P., Broeders, I. A. M. J., Pulles, B., Kappelhof, F. M., Van Der Werken, C., and Werken, C. V. D. (2004). Manual robot assisted endoscopic suturing: Time-action analysis in an experimental model. *Surgical Endoscopy and Other Interventional Techniques*, 18(8):1249–1252.

[Rydén and Chizeck, 2012] Rydén, F. and Chizeck, H. (2012). Forbidden-region virtual fixtures from streaming point clouds: Remotely touching and protecting a beating heart. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3308–3313.

[Rydén and Chizeck, 2013] Rydén, F. and Chizeck, H. (2013). A Method for Constraint-Based Six Degree-of-Freedom Haptic Interaction with Streaming Point Clouds. In *2013 IEEE Conference on Robotics and Automation*, pages 2345–2351.

[Rydén et al., 2011] Rydén, F., Kosari, S., and Chizeck, H. (2011). Proxy method for fast haptic rendering from time varying point clouds. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2614–2619.

[San Martin et al., 2007] San Martin, J., Trivino, G., and Bayona, S. (2007). Mechanical Design of a Minimally Invasive Surgery Trainer Using the Manipulability as Measure of Optimization. In *International Conference on Mechatronics*, pages 1–5.

[Savchenko et al., 1995] Savchenko, V. V., Pasko, A. a., Okunev, O. G., and Kunii, T. L. (1995). Function Representation of Solids Reconstructed from Scattered Surface Points and Contours. *Computer Graphics Forum*, 14(4):181–188.

[Schneider and Troccaz, 2001] Schneider, O. and Troccaz, J. (2001). A six-degree-of-freedom passive arm with dynamic constraints (PADyC) for cardiac surgery application: preliminary experiments. *Computer aided surgery*, 6(6):340–51.

[Schreiber et al., 2010] Schreiber, G., Stemmer, A., and Bischoff, R. (2010). The Fast Research Interface for the KUKA Lightweight Robot. In *2010 IEEE ICRA Workshop on Innovative Robot Control Architectures.*

[Sentis, 2007] Sentis, L. (2007). *Synthesis and Control of Whole-Body Behaviors in Humanoid Systems*. PhD thesis, Stanford University.

[Seymour et al., 2002] Seymour, N. E., Gallagher, A. G., Roman, S. a., O'Brien, M. K., Bansal, V. K., Andersen, D. K., and Satava, R. M. (2002). Virtual reality training improves operating room performance: results of a randomized, double-blinded study. *Annals of surgery*, 236(4):458–463; discussion 463–464.

[Shoham et al., 2003] Shoham, M., Burman, M., Zehavi, E., Joskowicz, L., Member, S., and Batkilin, E. (2003). Bone-Mounted Miniature Robot for Surgical Procedures : Concept and Clinical Applications. *IEEE Transactions on Robotics and Automation*, 19(5):893–901.

[Shoham et al., 2007] Shoham, M., Lieberman, I. H., Benzel, E. C., Togawa, D., Zehavi, E., Zilberstein, B., Roffman, M., Bruskin, a., Fridlander, a., Joskowicz, L., Brink-Danan, S., and Knoller, N. (2007). Robotic assisted spinal surgery-from concept to clinical practice. *Computer Aided Surgery*, 12(2):105–115.

[Siebert et al., 2002] Siebert, W., Mai, S., Kober, R., and Heeckt, P. F. (2002). Technique and first clinical results of robot-assisted total knee replacement. *The Knee*, 9(3):173–180.

[Slotty et al., 2012] Slotty, P. J., Kamp, M. a., Wille, C., Kinfe, T. M., Steiger, H. J., and Vesper, J. (2012). The impact of brain shift in deep brain stimulation surgery: Observation and obviation. *Acta Neurochirurgica*, 154(11):2063–2068.

[Smith et al., 2014] Smith, J. R., Riches, P. E., and Rowe, P. J. (2014). Accuracy of a free-hand sculpting tool for unicondylar knee replacement. *The International Journal of Medical Robotics and Computer Assisted Surgery*, 10(2):162–169.

[Sousa and Cortesão, 2014] Sousa, C. D. and Cortesão, R. (2014). Physical feasibility of robot base inertial parameter identification: A linear matrix inequality approach. *The International Journal of Robotics Research*, 33(6):931–944.

[Stasse et al., 2008] Stasse, O., Escande, A., Mansard, N., Miossec, S., Evrard, P., and Kheddar, A. (2008). Real-Time (Self)-Collision Avoidance Task on a HRP-2 Humanoid Robot. In *2008 IEEE International Conference on Robotics and Automation*, pages 3200–3205.

[Sugiura et al., 2006] Sugiura, H., Gienger, M., Janssen, H., and Goerick, C. (2006). Real-Time Self Collision Avoidance for Humanoids by means of Nullspace Criteria and Task Intervals. In *2006 IEEE-RAS International Conference on Humanoid Robots*, pages 575–580.

[Sun et al., 2007] Sun, L. W., Van Meer, F., Schmid, J., Bailly, Y., Thakre, A. A., and Yeung, C. K. (2007). Advanced da Vinci surgical system simulator for surgeon training and operation planning. *The international journal of medical robotics and computer assisted surgery*, 3(1):245–251.

[Sutherland et al., 2003] Sutherland, G. R., McBeth, P. B., and Louw, D. F. (2003). Neuroarm: an MR compatible robot for microsurgery. *Computer Assisted Radiology and Surgery*, 1256(1):504–508.

[Sutherland et al., 2013] Sutherland, G. R., Wolfsberger, S., Lama, S., and Zarei-nia, K. (2013). The evolution of neuroArm. *Neurosurgery*, 72(1):27–32.

[Swapp et al., 2006] Swapp, D., Pawar, V., and Loscos, C. (2006). Interaction with co-located haptic feedback in virtual reality. *Virtual Reality*, 10(1):24–30.

[Takubo et al., 2002] Takubo, T., Arai, H., Hayashibara, Y., and Tanie, K. (2002). Human-Robot Cooperative Manipulation Using a Virtual Nonholonomic Constraint. *The International Journal of Robotics Research*, 21(5-6):541–553.

[Tanaka et al., 2005] Tanaka, H., Minami, M., and Mae, Y. (2005). Trajectory Tracking of Redundant Manipulators Based on Avoidance Manipulability Shape Index. In *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4083–4088.

[Taylor et al., 1999] Taylor, R., Jensen, P., Whitcomb, L., Barnes, A., Kumar, R., Stoianovici, D., Gupta, P., Wang, Z., Dejuan, E., and Kavoussi, L. (1999). A Steady-Hand Robotic System for Microsurgical Augmentation. *The International Journal of Robotics Research*, 18(12):1201–1210.

[Tenzer et al., 2012] Tenzer, Y., Davies, B. L., and Baena, F. R. Y. (2012). Four-state rotary joint control: Results with a novel programmable brake. *IEEE/ASME Transactions on Mechatronics*, 17(5):915–923.

[The Royal College of Radiologists, 2012] The Royal College of Radiologists (2012). Ultrasound training recommendations for medical and surgical specialties, Second edition. Technical report, The Royal College of Radiologists.

[Trejos et al., 1999] Trejos, A. L., Salcudean, S. E., Sassani, F., and Lichtenstein, S. (1999). On the Feasibility of a Moving Support for Surgery on the Beating Heart. In *Medical Image Computing and Computer Assisted Intervation*, pages 1088–1098.

[Troccaz and Delnondedieu, 1996] Troccaz, J. and Delnondedieu, Y. (1996). Semi-active guiding systems in surgery. A two-dof prototype of the passive arm with dynamic constraints (PADyC). *Mechatronics*, 6(4):399–421.

[Tsumugiwa et al., 2003] Tsumugiwa, T., Sakamoto, A., Yokogawa, R., and Hara, K. (2003). Switching Control of Position/Torque Control for Human-Robot Cooperative Task. In *2003 IEEE International Conference on Robotics and Automation*, pages 1933–1939.

[Tsumugiwa et al., 2002] Tsumugiwa, T., Yokogawa, R., and Hara, K. (2002). Variable impedance control based on estimation of human arm stiffness for human-robot cooperative calligraphic task. In *2002 IEEE International Conference on Robotics and Automation*, pages 644–650.

[Turro et al., 2001] Turro, N., Coste-maniere, E., and Antipolis, S. (2001). Haptically Augmented Teleoperation. In *2001 IEEE International Conference on Robotics and Automation*, pages 386–392.

[Varma and Eldridge, 2006] Varma, T. R. K. and Eldridge, P. (2006). Use of the NeuroMate stereotactic robot in a frameless mode for functional neurosurgery . *The International Journal of Medical Robotics and Computer Assisted Surgery*, 2(2):107–113.

[Wagner et al., 2002] Wagner, C., Stylopoulos, N., and Howe, R. (2002). The role of force feedback in surgery: analysis of blunt dissection. In *10th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 8–14. IEEE Comput. Soc.

[Wein et al., 2008] Wein, W., Brunke, S., Khamene, A., Callstrom, M. R., and Navab, N. (2008). Automatic CT-ultrasound registration for diagnostic imaging and image-guided intervention. *Medical Image Analysis*, 12(5):577–585.

[Wein et al., 2013] Wein, W., Ladikos, A., Fuerst, B., Shah, A., Sharma, K., and Navab, N. (2013). Global registration of ultrasound to MRI using the LC2 metric for enabling neurosurgical guidance. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 8149 LNCS(PART 1):34–41.

[Wendland, 1995] Wendland, H. (1995). Piecewise polynomial, positive definite and compactly supported radial functions of minimal degree. *Advances in Computational Mathematics*, 4(1):389–396.

[Weon et al., 2015] Weon, C., Hyun Nam, W., Lee, D., Lee, J. Y., and Ra, J. B. (2015). Position tracking of moving liver lesion based on real-time registration between 2D ultrasound and 3D preoperative images. *Medical Physics*, 42:335–347.

[Westebring-van der Putten et al., 2008] Westebring-van der Putten, E. P., Goossens, R. H. M., Jakimowicz, J. J., and Dankelman, J. (2008). Haptics in minimally invasive surgery–a review. *Minimally invasive therapy & allied technologies : MITAT : official journal of the Society for Minimally Invasive Therapy*, 17(1):3–16.

[Wit et al., 1995] Wit, C. C. D., Olsson, H., Åström, K. J., and Lischinsky, P. (1995). A New Model for Control of Systems with Friction. *IEEE Transactions on Automatic Control*, 40(3):419–425.

[Wottawa et al., 2013] Wottawa, C. R., Cohen, J. R., Fan, R. E., Bisley, J. W., Culjat, M. O., Grundfest, W. S., and Dutson, E. P. (2013). The role of tactile feedback in grip force during laparoscopic training tasks. *Surgical Endoscopy and Other Interventional Techniques*, 27(1):1111–1118.

[Yoshikawa, 1985a] Yoshikawa, T. (1985a). Dynamic manipulability of robot manipulators. In *1985 IEEE International Conference on Robotics and Automation*, pages 1033–1038.

[Yoshikawa, 1985b] Yoshikawa, T. (1985b). Manipulability and redundancy control of robotic mechanisms. In *1985 IEEE International Conference on Robotics and Automation*, pages 1004–1009.

[Yu et al., 1993] Yu, C. J., Yang, P. C., Wu, H. D., Chang, D. B., Kuo, S. H., and Luh, K. T. (1993). Ultrasound study in unilateral hemithorax opacification. Image comparison with

computed tomography. *The American review of respiratory disease*, 147:430–434.

[Yu and Liang, 2012] Yu, Y. and Liang, W. (2012). Design Optimization for Lower Limb Assistive Mechanism Based on Manipulability Inclusive Principle. In *2012 IEEE International Conference on Robotics and Biomimetics*, pages 174–180.

[Yuen et al., 2009] Yuen, S. G., Kettler, D. T., Novotny, P. M., Plowes, R. D., and Howe, R. D. (2009). Robotic Motion Compensation for Beating Heart Intracardiac Surgery. *The International Journal of Robotics Research*, 28(10):1355–1372.

[Zhu et al., 2014] Zhu, E., Hadadgar, A., Masiello, I., and Zary, N. (2014). Augmented reality in healthcare education: an integrative review. *PeerJ*, 2:e469.

[Ziherl and Munih, 2009] Ziherl, J. and Munih, M. (2009). Pick to place trajectories in human arm training environment. *Technology and Health Care*, 17(4):440–443.

[Zoss et al., 2005] Zoss, A., Kazerooni, H., and Chu, A. (2005). On the Mechanical Design of the Berkeley Lower Extremity Exoskeleton (BLEEX). In *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3132–3139.