Real-Time Active Constraint Generation and Enforcement for Surgical Tools using 3D Detection and Localisation Network

Spyridon Souipas [1], **Anh Nguyen** [2], **Stephen G. Laws** [1], **Brian L. Davies** [1] **and Ferdinando Rodriguez y Baena** [1]

[1] *Mechatronics in Medicine, Imperial College London, Mechanical Engineering, London, United Kingdom*
[2] *Department of Computer Science, University of Liverpool, Computer Science, London, United Kingdom*

Correspondence*:
Spyridon Souipas et al.
ss8413@imperial.ac.uk

ABSTRACT

Despite numerous improvements in orthopaedic surgical robots, there still exists a gap between bulky systems with haptic feedback capabilities, and miniaturised systems that only allow for boundary control. In addition, active constraint generation for haptic feedback is solely based on the use of optical trackers and is usually performed in the preoperative plan. This paper presents both a refined version of the previously presented Signature Robot, a three degress-of-freedom orthopaedic surgical collaborative system, but also a method for generating and enforcing an active constraint "on-the-fly", simply by using our previously presented monocular, RGB camera-based network, SimPS-Net. The robot was employed to explore constraint enforcement, allowing the operator to interact with a safe region, from which exit was discouraged, and a restricted region, where the robot was prevented from entering. With a network success rate of 54.7% $\pm$ 5.2% achieved in the constraint definition step, the same area was therefore used for both safe-zone and restricted-zone testing. In the former case, the robot was resisted upon exiting the defined region, and in the presence of haptic feedback, the mean distance from the boundary and mean duration of each exit was found to be 2.70mm $\pm$ 0.37mm and 0.76s $\pm$ 0.11s respectively. In the case of a restricted-zone, the operator was prevented from penetrating the defined region with a success rate of 100%. This paper ultimately demonstrates both the viability of the proposed robotic platform, along the promising results of a highly versatile constraint generation and enforcement pipeline.

Keywords: Active Constraints, Surgical Robotics, Surgical Tool Detection, 3D Pose Estimation, Surgical Tool Localisation

1 INTRODUCTION

With the invention of surgical robots, and their introduction to the operating theatre, there has also been a significant development across various forms of sensory feedback implementations. The majority of such feedback comes in the form of haptic algorithms, with some developments also being noted in visual feedback techniques. In the former case, the robotic system, be it a hands-on or a teleoperated, applies

guiding or restrictive forces to the operator when the robot end-effector is about to interact with a predefined boundary on the patient tissue, which is usually generated in the preoperative step. In the later case, the operator is simply able to better visualise the position of the robot with respect to either the human tissue or the aforementioned predefined boundary.

When examining haptic feedback, the term "active constraint" can be used to generally describe cases where some degree of assistance to the human during robotic manipulation is present in various forms. One such form involves the generation of a safety boundary and the application of forces to the operator upon trying to enter or exit the constraint geometry. For example, the user may wish to freely operate the robot within a region, and request the robot to apply restrictive forces upon trying to exit. Therefore, with the integration of these constraints along the workflow of robotic surgery, robots possess an inherently higher level of safety upon interaction with human tissue (1). In addition, active constraints generally describe a situation where the surgeon is not fully restricted, thus should they wish to consciously exit the safety region in operation, they may do so by overcoming the applied resistance.

Despite such developments of haptic feedback, however, the implementation requires robotic systems to incorporate motors that are strong enough to sustain the involved forces. This, in turn, may lead to surgical platforms that comprise large robot arms, thus making them cumbersome to manipulate. In addition, such designs suffer from a noteworthy footprint in the operating room due to the size of the components. Conversely, other robotic systems may opt for a miniaturised arrangement, which, while convenient to operate, cannot sustain any haptic feedback forces.

Irrespective of robot size and haptic capabilities, it is also important to consider the tracking systems that are usually integrated into the workflow of a robotic surgery. Namely, the majority of commercially available robots involve the use of an optical tracker, which serves the purpose of both patient registration, but can also be used for the purpose of active constraint definition in the preoperative plan (2). The safety region can be defined based on the scanned patient geometry and communicated to the robot via appropriate channels. Constraint definition can be done through devices that are distant from the patient, without having the need to interact with actual human tissue in order to define a constraint. Therefore, it is possible to generate regions of interest along the tissue of the patient, whilst ensuring that in operation, accurate interaction between robot and human tissue is achieved at all times. Nevertheless, such a process can be time consuming and is not dynamic, since at no point during surgery can the operator define regions of interest based solely on visual criteria in the operated tissue and apply some form of constraint accordingly.

Overall, the field of robotic surgery has enjoyed a steep increase in attention over the past years, both in terms of research, but also in terms of monetary investment. Surgical robots have been reported to provide increased industry benefits across the majority of availabel literature for joint surgery (3). Robots have become available for numerous types of surgeries, boasting different sizes, cutting capabilities, and overall increased accuracy (4) compared to conventional methods. In the sector of orthoapedics, current surgical robotic platforms deployed for cutting, as opposed to robots that support jig placement, can be arranged into three categories, depending on the cutting capabilities they provide.

The first method is simply autonomous function. As the name suggests, such systems function fully autonomously, without involving operators to a significant degree. An example of such a setup is the TSolution One (Think Surgical, USA) (5), which is an image-based milling apparatus that allows for more accurate implant placement in total knee arthroplasty (TKA). A cadaveric study on human heads demonstrated how this dynamic navigation system offers a reduction in operation duration and angular deflection compared to manual operation (6). Autonomous cutting, while useful, minimises the involvement

of trained surgeons in some cases, which in turn may raise ethical concerns, therefore not being adopted to a great extent.

The second form of cutting is the one that involves haptic feedback. As mentioned previously, in cases like this, an operator is required to manipulate the robot to perform cuts. The robot is allowed to move freely until interaction with a predefined safety border occurs, in which case restrictive or guiding forces are applied to the surgeon. One such example is the MAKO (Stryker, USA) (7), which is a six degrees-of-freedom system, used for unicompartmental knee arthroplasty (UKA), total knee arthroplasty (TKA), and total hip arthroplasty (THA). A surgeon manipulates surgical tools mounted on a robot arm, with the latter applying a restrictive force upon potential deviations from the constraint geometry. Studies have demonstrated that such a system can offer a significant degree of improvement across surgeries. For example, in a cadaveric study of TKA, the femoral anterior, distal, and posterior flexion values were found to be 0.4° 0.8° and 0.5° respectively, with the anterior and posterior flexion values demonstrating a significant reduction compared to the manual operation results of 4.7° and 2.3° respectively (8).

Lastly, robotic actions can be monitored through boundary control. In such cases, a surgeon is once again required to control a robot end-effector, with cutting being prevented or entirely terminated upon exiting a safety region. One such case is the NAVIO (Smith & Nephew, UK), used in TKA (9). This is a handheld tool that can either limit the speed of the surgical burr tooltip, or completely block the burr itself upon exiting a predefined boundary (10). Interestingly, considering how this robotic system has been shaped to be manipulated by one hand from the operator, the learning curve effect has been minimised (11), whilst cadaveric study results on TKA for this system demonstrated a mean femoral flexion of 2.0° (12), which is an improvement compared to standard methods. Similarly, the NAVIO achieved a femoral flexion of 1.23° in a cadaveric study for UKA (13). Furthermore, the robotic system CORI (Smith & Nephew, UK), which is the evolved version of NAVIO and is also dependent on an optical tracker (14), has demonstrated improved results compared to its predecessor (15). Another example of a boundary control capable robot is the newly developed VELYS system (DePuy Synthes, USA) (16) where, following the identification of ideal cut alignment of the proximal tibia during TKA, the robot prevents the operator from engaging the saw upon deviation from the established cutting path. The VELYS has reported some significantly accurate results in cadaveric studies, achieving satisfactory implant alignment levels throughout TKA with a femoral flexion of 2.2° (17) compared to conventional instrumentation.

An analysis of the progress of commercially available surgical robotic technology can lead to the revelation of a pattern. In particular, there is a balance between the haptic capabilities of a robotic system against the space needed to integrate the platform in the operating theatre. On one side, there are systems like the NAVIO, which are miniaturised, and allow for freehand operation, without allowing for any haptic feedback. On the other side stand the bulkier systems, which allow for haptic feedback supported motion along fewer degrees-of-freedom, such as the MAKO. An interesting deviation from this pattern is the VELYS system, which while on the bulkier side, allows only for boundary control along three degrees-of-freedom. Despite the absence of haptic control, however, recent findings have demonstrated that the VELYS can achieve comparable accuracy (18) to MAKO (19) in terms of implant placement. Upon contrasting these two robot categories, a general disparity is evident between the haptic capabilities of smaller robots and the ease of deployment of bulkier setups.

In addition to deploying the robotic platform itself, further attention needs to be given to the deployment of optical trackers and other external sensors which require further room in the operating theatre. The optical trackers are necessary in the case of haptic feedback systems, as they can be employed to establish active constraint boundaries across the patient tissue.

In terms of haptic feedback, substantial research has been undertaken on the topic of active constraints, and a general summary is available (20). While extensive analysis of the available literature is beyond the scope of this paper, it is important to understand some concepts that are relevant to the implementation of constraints presented in this paper. Firstly, when exploring the concept of constraint definition, most research involves the employment of either optical trackers to generate a safety region (21), or it is assumed that the constraint has been defined a priori, or can be defined using some generated point-cloud (22). Another frequent implementation involves the generation and combination of primitive shapes, which combined can create a surface or volume that define the constraint (23). Therefore, the front of defining a constraint region has not received significant attention in the past.

When it comes to constraint enforcement, numerous techniques have been developed. A very relevant example is the definition of a constraint region, with the robot either being kept within the "safe zone", or the robot being prevented from entering the "restricted zone". The former can be considered as an attractive force being applied to the end-effector upon deviation from the safe zone boundary, guiding the operator back to the zone, whilst the latter, also referred to as a "forbidden region virtual fixture" (24), involves a repulsive force being applied to the end-effector when approaching the restricted region (25), thus limiting or entirely preventing motion within the aforementioned region. A final distinction to make concerns dynamic environments (26). The majority of constraints are static, thus assuming that the surrounding environmental changes do not affect the predefined boundaries in operation. A dynamic constraint, however, takes into account relevant tissue deviations, thus updating the constraint definition accordingly to patient tissue alterations.

The amalgamation of robotic platform shortcomings in terms of miniaturised systems not offering haptic feedback, along with the development of novel object detection and tracking techniques, has led to the proposed constraint definition pipeline, where surgeons can dynamically establish active constraint boundaries during operation. Specifically, instead of employing cumbersome optical trackers, a static virtual fixture can be defined "on-the-fly" by employing a monocular, RGB-based network capable of detecting and localising surgical tools in 3D space. In doing so, the operator can define safe regions or restricted regions simply by generating an area using a standard scalpel. This can ultimately guarantee the protection of chosen patient tissue regions from unintended damage due to substandard manipulation of a robotic end-effector. Additionally, this proposed workflow further benefits by the use of our previous work, SimPS-Net (27), which is capable of tracking standard surgical tools, which can be tracked by inexpensive RGB cameras, in order to swiftly identify regions of interest, thus eliminating the need of involving optical trackers in the dynamic constraint definition process, while also avoiding additional training for surgeons. This paper introduces a framework that enables a real-time construction of a safety boundary, as well as the implementation of the resulting active constraint in a novel robotic platform. This piece of research, therefore, comprises two fronts, namely the employment of our previously developed surgical object tracking network, SimPS-Net (27) for the purposes of constraint geometry generation, and the use of a refined version of the previously presented Signature Robot (28) for the purpose of exploring the generated constraints. The new version of the proposed, three degrees-of-freedom system has been assembled and refined with the purpose of addressing the chasm between bulky robots that offer haptic feedback, and miniaturised systems that solely offer boundary control.

2 METHODOLOGY

The aims of this research paper can be divided into three sections, namely, an analysis of the current state of the Signature Robot, along with relevant testing on plastic bone, followed by an examination of

an "on-the-fly" constraint area generation procedure, with the final purpose being the enforcement of the constraint on the robotic platform.

2.1 Signature Robot Assembly

The development of the Signature Robot came about as a successor to the Acrobot (29). In particular, the purpose of this novel platform is to address the aforementioned gap between bulky, haptic feedback robots and miniaturised, boundary control robots in orthopaedic surgery. With that in mind, the Signature Robot was design to move across three degrees-of-freedom, namely linear, yaw, and pitch, as shown in Figure 1. One of the main aims of this design is to allow for operators to simply manipulate the attached end-effector using precise finger motions or minor wrist motions. The range of motion is 150mm in linear translation and 90° in pitch and yaw, thus providing a sufficient range of motion for small joint surgery, such as UKA. Despite having the same degrees-of-freedom as the VELYS system, haptic feedback capabilities have also been implemented. In addition, the footprint consumed by the platform only, not including the power supply or other relevant components, is miniaturised, taking up only 200mm on each side.

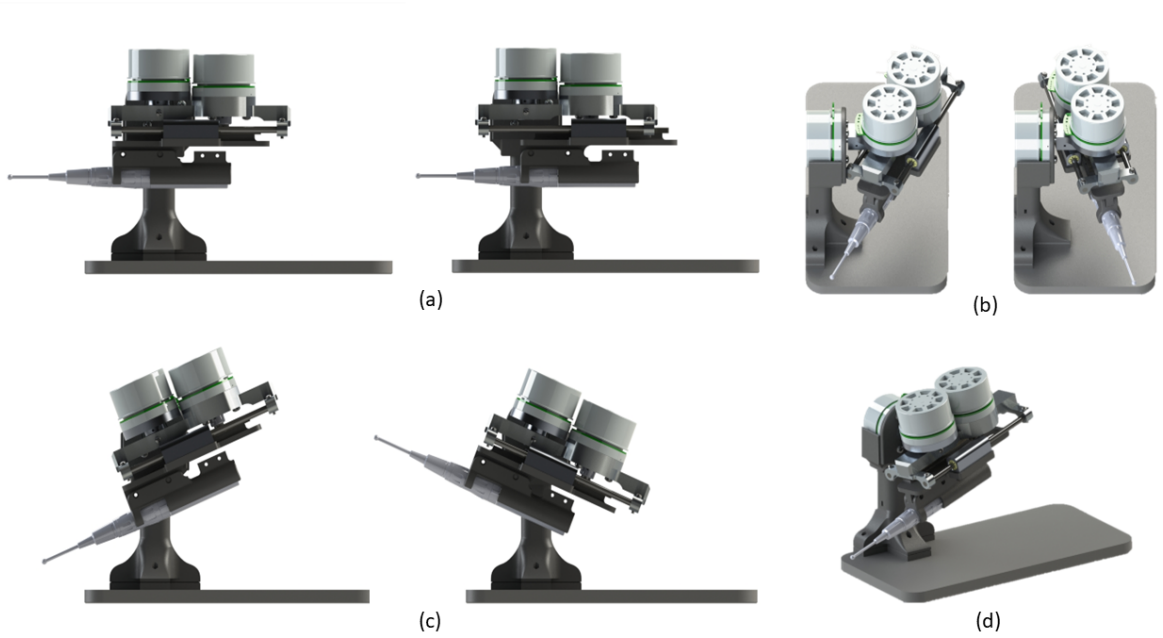

Figure 1. Demonstration of Signature Robot (d) Degrees-of-Freedom: Linear (a), Yaw (b), Pitch (c).

The original version of the Signature Robot (28) was developed to include numerous transmission components. The presence of those, and especially the incorporation of bevel gears, led to a significant degree of backlash. For that purpose, this presented design was constructed to minimise the number of such components, instead opting for a direct drive for each degree-of-freedom.

With the assembly of this system, some form of haptic feedback was also integrated. Specifically, the system has been developed to allow the operator to generate some primitive shape at a location of their choosing. The end-effector can be manipulated in an unrestricted manner within the static safe region, with a viscoelastic force being applied to the tooltip upon exiting this region, as demonstrated in Equation 1:

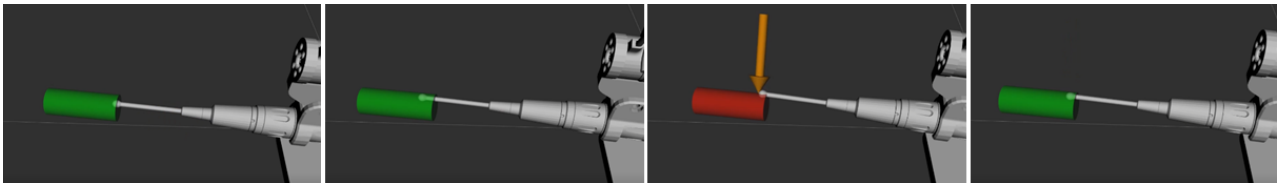

Figure 2. Demonstration of Active Constraint on Signature Robot.

$$\mathbf{F} = \mathbf{k} \cdot \mathrm{d} \cdot \hat{\mathbf{p}} + \mathbf{c} \cdot \dot{\mathbf{x}} \tag{1}$$

where $\mathbf{F}$ is the force vector on the robot tooltip, $\mathbf{k}$ is the elastic constant vector, d is the robot distance from the boundary, $\hat{\mathbf{p}}$ is the unit penetration vector, $\mathbf{c}$ is the damping vector and $\dot{\mathbf{x}}$ is the robot velocity vector. The elastic and damping vectors are defined as $\mathbf{k} = [\mathrm{k_x}, \mathrm{k_y}, \mathrm{k_z}]$ and $\mathbf{c} = [\mathrm{c_x}, \mathrm{c_y}, \mathrm{c_z}]$ respectively. Both of these were set empirically, and highly depend on both the application and operator preference. Higher values of $\mathbf{k}$ will lead to a stiffer constraint, thus minimising distance from boundary upon exit, however potentially limiting surgeon control in operation. Higher damping values, on the other hand, restrict high speed motions. An inherent shortcoming of the viscoelastic algorithm is that it is an energy storing constraint, meaning that the system stores potential energy when moving outside the boundary and, in the accidental case of release, the robot tooltip can move in unpredictable or even dangerous motion. The damping component can address this, to a degree, by limiting high speeds. Empirically, the values for these vectors were set to $\mathbf{k} = [2.5, 2.3, 2.0]\mathrm{kN} \cdot \mathrm{m}^{-1}$ and $\mathbf{c} = [0.3, 0.3, 0.2]\mathrm{kN} \cdot \mathrm{s} \cdot \mathrm{m}^{-1}$. Still examining Equation 1, $\hat{\mathbf{p}}$ is the unit vector between the current robot position and the constraint centre or its central axis. For example, in the case of a cylindrical constraint, as the one shown in Figure 2, the force will always point towards the cylinder axis. As shown in the figure, where the digital twin of the designed robot mimics the motions of the real system, while the robot is within the safe region, deemed as green, no forces act on the system. However, upon deviation from the boundary, some visual feedback is provided to the operator in the form of a colour change to red, and a force is applied to the tooltip to steer the surgeon back to the safe region. Once that happens, the motor efforts are once again set to zero, until another exit from the boundary occurs.

With both the design and the haptic constraint being established, a comparison between freehand operation and controlled operation was further explored. Specifically, 4 tests were undertaken, each with a different volunteer. Each test involved the cutting of a cylindrical section of 15mm diameter on plastic bone using the Signature Robot under four different modes of operation, namely freehand, visual feedback, haptic feedback, and complete feedback modes. In freehand operation, the volunteers were requested to define the starting point of a cylinder on the plastic bone and then proceed with the cut, with no form of feedback being available. In visual feedback mode, the same was done, however, operators had access to the digital twin, thus receiving feedback identical to the one of Figure 2. In the third mode, haptic feedback solely was present, with the final mode of operation combining visual and haptic feedback access. The robot was mounted with a deSoutter (UK) surgical 5mm burr, and aligned within range of the plastic bone, as demonstrated in Figure 3. The users were asked to cut for up to 120 seconds.

Robot position, cylinder starting point, and cylinder axis were all recorded and, upon completing each test, relevant findings were calculated and ultimately amalgamated. The metrics can be divided into two parts, namely boundary exit related metrics and volume removal metrics. The former simply indicates parameters associated with each exit, like duration and distance from the boundary, while the latter demonstrates the amount of volume burred within and outside of the safe region in each mode of operation. To measure the

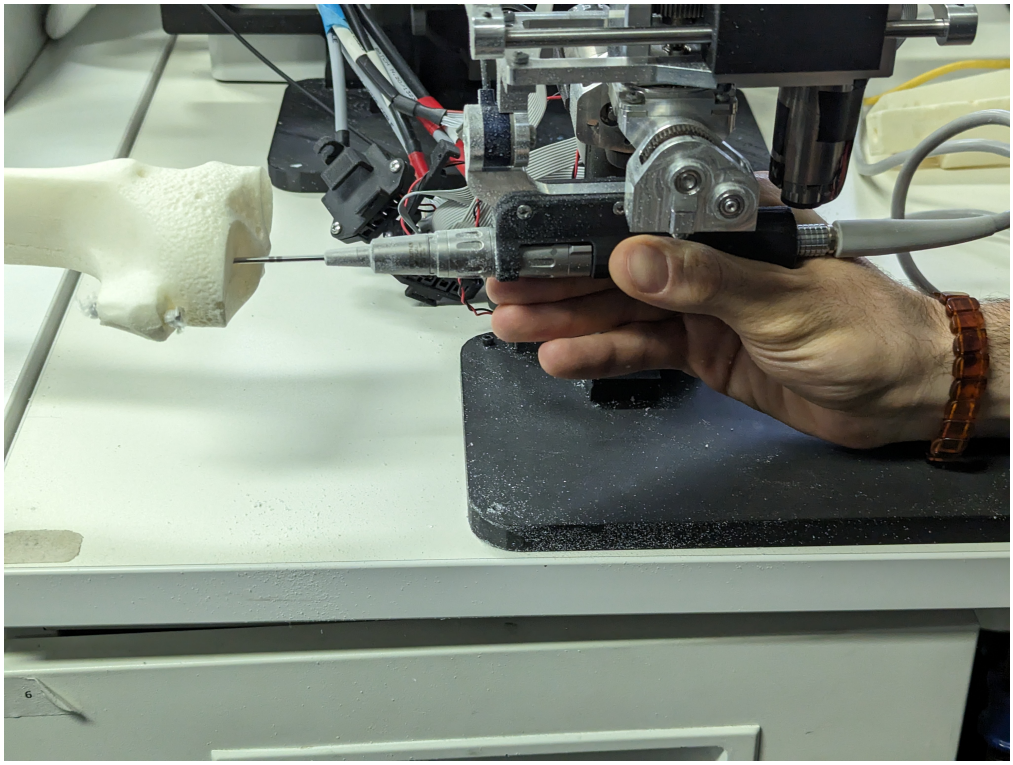

Figure 3. Experimental Setup for Plastic Knee Cutting.

volume burred, for each test, a gridspace was generated based on the maximum and minimum coordinates that the robot reached. The grid points were set to have a fixed distance of 0.5mm. When iterating across the recorded trajectory of a test, the nearest grid point was located and defined as cut. In addition, all grid points in a distance equal to the radius of the bur were set to be in a similar state. Once all relevant grid points were identified, the volume within the safe region and the total volume burred were measured. For one of the participants operating the robot in freehand mode, Figure 4 shows the total volume burred and the volume burred within the active constraint. The red points indicate grid marks that the robot tooltip removed from the plastic bone during cutting. The volume was simply calculated by creating cubes around each grid point, calculating the individual volumes, and ultimately summing them up.

Upon completing these measurements, a qualitative test was also performed in the form of completing a NASA TLX assessment (30) to further understand the perceived workload, with the purpose of further improving the robot in future iterations. With the proposed platform assembled and tested, the next step is to deploy it alongside an active constraint definition vision-based system in order to explore the capabilities of the constraint pipeline presented in this paper.

2.2 Safety Region Generation

With the robotic platform established, the initial step in generating a constraint in real-time involves the region definition. This constraint region can be subsequently used for different types of virtual fixtures, as will be discussed in the next subsection. In order to allow the operator to define the required volume, we employed the previously mentioned SimPS-Net in order to achieve real-time surgical tool localisation in 3D space. The chosen network is an expansion of the Mask-RCNN (31), a widely used network for the purpose of pixelwise object detection. The purpose of SimPS-Net was to allow the operator to define a region on a phantom knee using a standard, surgical scalpel. The process of recording this area involved

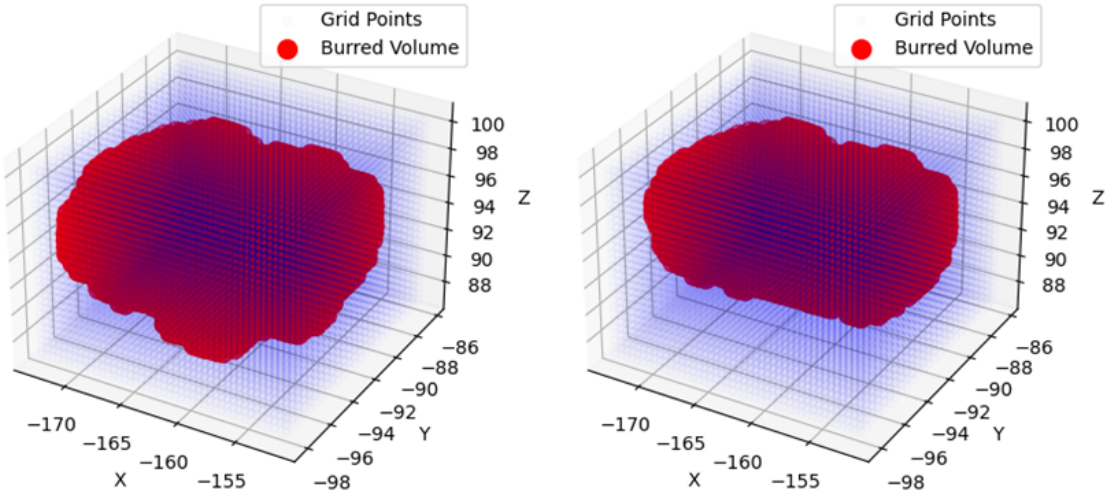

Figure 4. Volume Grid, with Total Volume (left) and Volume Within Constraint (right).

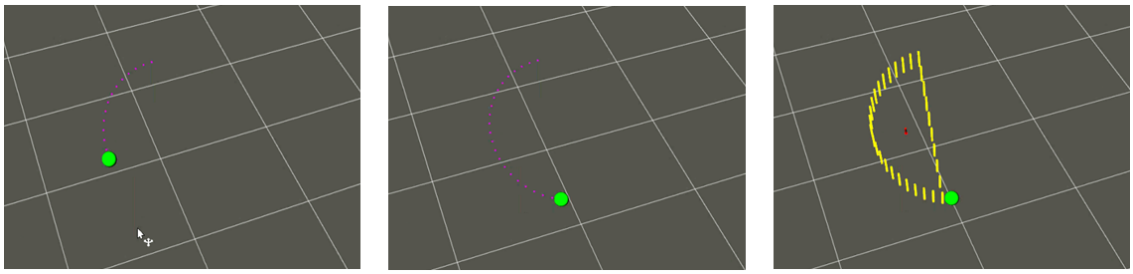

Figure 5. Example of safety volume generation with tool position shown in green, tracked points in purple (left, middle), and final volume in yellow (right).

two different sensors. A RealSense D415 camera set in RGB mode was used for the purpose of surgical tool localisation, while a stk300 optical tracker from Atracsys (Switzerland) was employed to record the ground truth positions of the scalpel. Two different thresholds were defined, a positional and an angular one, each of value 8mm and 5 degrees respectively. The 3D pose obtained via network was compared to the ground truth values, and any values exceeding the predefined thresholds were rejected. The network success rate in real-time was therefore defined as the ratio of accepted network measurements over total network measurements. The operator was required to generate a collection of non-coplanar points in 3D space using a scalpel. Figure 5 demonstrates the process of collecting an area. Each accepted position of the robot was collected at a fixed rate, shown in purple. Once the operator decided to terminate the recording, the final point was connected to the initial point of the recorded trajectory. Subsequently, principal component analysis (32) was applied to the collection of recorded points, thus identifying the theoretically optimal plane, as well as the normal to the plane, $\hat{\mathbf{n}}$. The normal vector was placed at the centroid of all collected points. The recorded points were subsequently extended in space parallel to this vector by some fixed depth, thus generating a volume in space, which was stored in camera 3D coordinates. For each of the nine repetitions that the active constraint experiment was undertaken, a new area was generated. The network achieved a frame rate of 4.5 frames per second using a NVIDIA GeForce RTX 3070 and Tensorflow 2.0.

2.3 Robot Constraint Enforcement

The final step of the proposed pipeline involves the enforcement of the active constraint on the Signature Robot. With the constraint having been defined in the previous step of the process, two different constraint enforcements were chosen for testing purposes. Specifically, the first enforcement type was a "safe region" type which, as previously mentioned, involves the application of a force on the tooltip only when the operator attempts to move the robot beyond the safety boundary. The second enforcement type was a "restricted zone", where the same constraint region was employed as a region which the robot was not allowed to penetrate. Both these experiments shared a common constraint definition at each repetition.

The robot was placed in a fixed position with respect to the phantom knee in a cadaveric lab, as shown in Figure 7. The optical tracker was used in order to register the fixed robot base, thus allowing for the transformation of tracker or camera coordinates to robot coordinates. Upon doing so, the collected constraint region was hence translated to robot coordinates and communicated to the Signature Robot embedded computer via Ethernet.

2.3.1 Safe Zone Testing

With the area established, the robot was placed in position, starting within the safe zone, and the operator was asked to move within the region. The test was repeated twice, once in freehand mode, where no constraint was applied, and once in controlled mode, where a constraint was engaged. In particular, upon deviation from the constraint boundary when operating in controlled mode, a force was applied to the tooltip with the purpose of returning the user back into the safe zone. The force was calculated in the same manner as the plastic bone robot testing experiments of Section 2.1, with Equation 1 being used to calculate the force. The only difference was that, since the generated safe zone was not continuous, the nearest point of the area was selected, and the distance, d, was defined as the distance of the robot tooltip from that identified point. This is demonstrated in Figure 6, where the position of the robot turns from purple to red upon boundary exit. The blue point is the nearest point of the volume to the robot, and the force is thus calculated using the distance between the two points. It should be pointed out that a threshold of 100ms was imposed on the experiments, below which an exit was not taken into consideration.

2.3.2 Restriced Zone Testing

Upon completing the measurements of safe zone testing, restricted-zone testing was undertaken. The robot tooltip was positioned outside the constraint region, and the restricted zone constraint was explored.

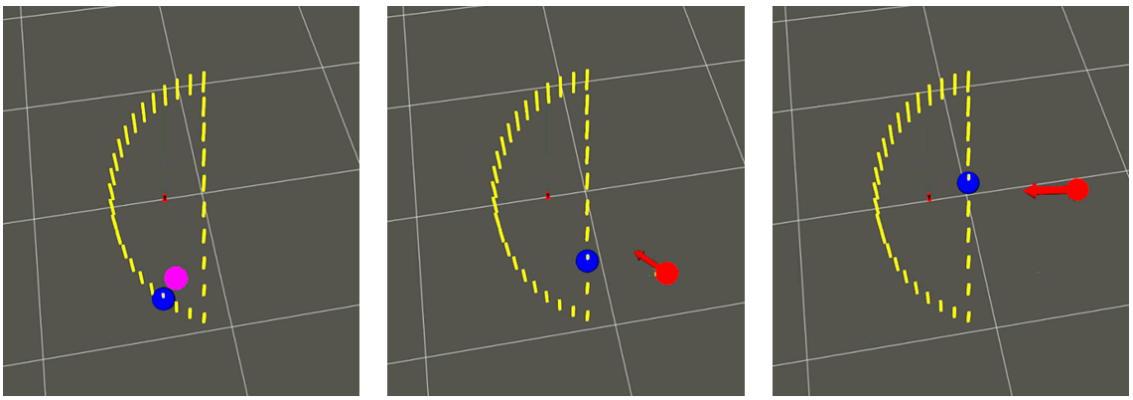

Figure 6. Example of force application on tooltip. The tooltip is shown in pink when within the safety region (left) and red when outside (middle, right). Shown in blue is the nearest point on the safety volume.

In this case, the operator was allowed to move freely outside of the constraint region, with a repulsive force applied on the tooltip when an attempt to enter the region was made. Specifically, the force on the tooltip is calculated using an inverse power law:

$$\mathbf{F} = \frac{\psi}{\mathrm{d}^2} \cdot \hat{\mathbf{q}} \tag{2}$$

Where ψ is a constant used for tuning the force magnitude, ensuring that the forces experienced by the operator were sensible, d is the distance of the current robot position from the nearest point along the restricted region boundary, and $\hat{\mathbf{q}}$ is the vector between the robot position and the normal vector, $\hat{\mathbf{n}}$, but points away from the restricted boundary. The value of ψ was empirically set to 1.5 $kN \cdot m$. In order to avoid reaching unsustainable forces, a threshold of 0.1mm was set. In case the end-effector reached a distance to the boundary lower than this, the distance was still locked to equal this value, thus generating a force plateau at small distances. For all nine repetitions, the operator was requested to attempt to penetrate the restricted zone, with force and distance from the boundary being recorded. Similarly to the safe zone testing, no visual feedback was present.

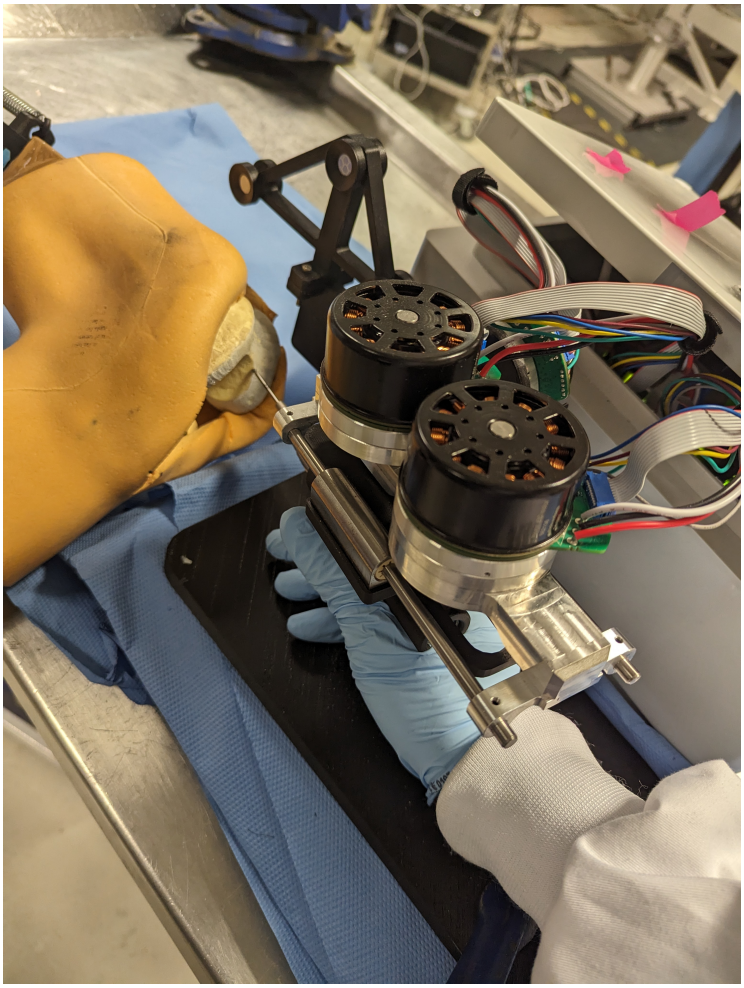

Figure 7. Experimental setup of robot and phantom knee.

3 RESULTS

We employ a total of 4 volunteers to perform the plastic bone cutting experiments undertaken for testing of the Signature Robot platform. In addition, nine repetitions were completed for the process of area generation and "on-the-fly" active constraint testing. The following subsections list the results for these two different sets of tests.

3.1 Signature Robot Testing

Four different modes of operation were explored by 4 volunteers in order to understand the capabilities of the current version of Signature Robot. Volunteers were asked to define the start of a cylinder on a plastic bone, and then attempt to cut out the defined cylindrical section. Across the participants, some important parameters, namely the mean exit duration, the maximum exit duration, and the mean and maximum deviations from the constraint boundary were collected. In addition, the total volume burred in operation, along with the volume contained within the constraint region were calculated, and their ratio was also measured as a final result. Table 1 lists the average results of for all modes of operation, along with the relevant uncertainties.

Table 1. Table of Results of Plastic Bone Cutting Experiments using the Signature Robot

Test	Mean Exit Duration (s)	Maximum Exit Duration (s)	Mean Distance of Exit (mm)	Maximum Distance of Exit (mm)	Inner Volume/Total Volume (%)
Test 1 - Freehand	2.76 ± 1.26	4.70 ± 1.46	1.97 ± 0.30	3.23 ± 0.47	63.5 ± 23.8
Test 2 - Visual Feedback	1.53 ± 0.97	2.63 ± 1.13	1.06 ± 0.27	1.57 ± 0.42	91.4 ± 0.28
Test 3 - Haptic Feedback	2.39 ± 0.97	3.53 ± 1.12	1.13 ± 0.64	2.23 ± 1.57	87.6 ± 6.11
Test 4 - Combined Feedback	0.70 ± 0.32	1.23 ± 0.36	0.35 ± 0.03	0.57 ± 0.07	95.7 ± 1.64

Furthermore, each participant filled out a NASA TLX form at the end of the tests, with the results being amalgamated and averaged out on Table 2. This standard form addresses several parameters that, while subjective based on the experience of each participant, may provide some insightful conclusions in terms of ease of operation. Such feedback can, in turn, be incorporated into future design improvements or additional features. The form was completed by each participant upon completing the plastic bone burring process in all four modes of operation.

Table 2. Table of Average Collected NASA TLX Responses

Mental Demand	Physical Demand	Temporal Demand	Performance	Effort	Frustration
4.7 ± 1.1	12.8 ± 1.8	2.7 ± 0.7	16.3 ± 1.6	6.0 ± 3.1	4.7 ± 2.6

3.2 Active Constraint Generation Testing

Upon conducting nine reiterations, the average network success rate was found to be $\mathbf{54.7\% \pm 5.2\%}$. Following the area generation outputs, the results of boundary interactions in safe-zone testing and restricted-zone testing were obtained. Firstly, when examining the safe-zone experiments, metrics for both freehand operation and haptic feedback operation were collected, with the combined results presented in

Table 3. The parameters of interest involve the mean and maximum duration of each exit, as well as the mean and maximum distance from the constraint region. Number of exits were also calculated.

Table 3. Table of Results of Freehand and Controlled Modes of Operation

Mode	Mean Exit Duration (s)	Maximum Exit Duration (s)	Number of Exits	Mean Distance of Exit (mm)	Mean Maximum Exit (mm)
Freehand	1.32 ± 0.56	3.69 ± 1.97	13 ± 6	5.83 ± 1.65	20.74 ± 6.78
Controlled	0.76 ± 0.11	2.03 ± 0.51	19 ± 5	2.70 ± 0.37	9.58 ± 2.48

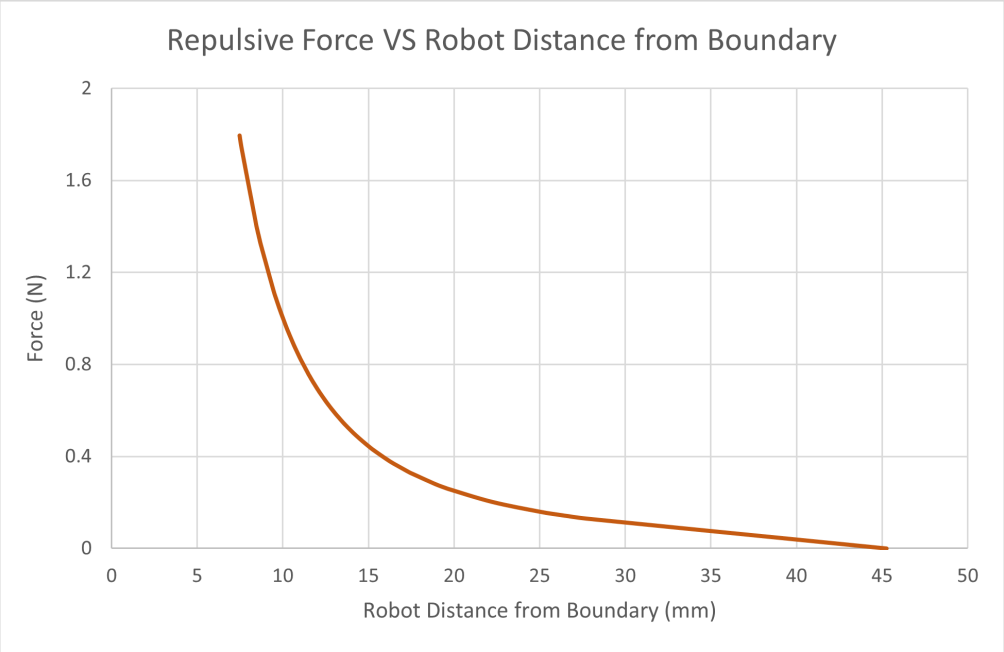

Figure 8. Force on Tooltip in Restricted Zone Testing.

Subsequently, the restricted-zone tests mostly focused on proving that the robot can be completely prevented from entering the defined region. Indeed, throughout all iterations of the test, not a single penetration event was noted. Figure 8 demonstrates the force increase the closer the robot gets to the region.

4 DISCUSSION

The main purposes of this research were to both demonstrate how the Signature Robot was further improved from previous iterations, as well as establish an easily generated constraint definition method, and ultimately enforce it on the aforementioned robotic platform. The constraint generation technique presented in the previous sections offers a noteworthy level of versatility, as it does not involve the use of any optical trackers, whilst also allowing operators to seamlessly produce a constraint region in the course of surgery, by employing standard surgical tools. This latter point suggests that surgeons need not familiarise themselves with unknown tools in order to achieve the process of constraint definition. Ultimately, the defined region can be easily communicated to the robotic platform and explored in operation.

4.1 Signature Robot

The Signature Robot was developed for orthopaedic surgery, with the purpose of maintaining a small footprint in the operating room, along with ease of deployment, while also providing haptic feedback to the operator. With the presented version employing a direct drive for each of the three degrees-of-freedom, the system is indeed capable of providing a sufficient force, which remained at about 6N when testing on plastic bone, but can be further increased according to the needs of the operator. The static constraint tested in the form of a cylinder is a promising starting point, and the highly tunable force algorithm presented in Equation 1 suggests that with further qualitative testing, an optimal value of elastic and damping coefficients can be identified.

By observing Table 1, it is possible to draw some general conclusions. Specifically, it is no surprise that the metrics of freehand operation are principally higher, both in terms of exit duration but also in terms of distance from the boundary. Without any form of feedback, the operators were dependant solely on spatial awareness of the location of the cylinder, and any deviation from the initial shape could not be limited. This is further underlined in the metric of inner volume over total volume. This metric indicates the ratio of the correctly removed plastic bone over the total volume burred, and therefore the lower the percentage, the higher the amount of volume removed that should not have been burred out during operation. The volume ratio is highly variable in freehand operation since some volunteers misjudged the orientation of the generated cylinder and proceeded to perform cuts that were notably outside the constraint.

Some interesting findings, however, are observed when comparing visual feedback and haptic feedback. Specifically, the former seems to outperform the latter, even if marginally in some cases. This suggests that in this application, volunteers were more responsive to visual information in the form of colour changes, thus reacting faster upon exiting the boundary, as indicated by the lower mean and maximum exit duration. Despite that, the mean distance of each exit remains comparable under the two modes of operation. One explanation for this could be that, under visual feedback, the volunteers focused more on avoiding boundary exits using their own skill, whereas in the case of haptic feedback, the level of dependence on the robot led to some degree of overconfidence, which thus resulted in higher errors. One way to address this is to further increase the elastic constants, hence leading to a stiffer force being applied sooner upon exiting the constraint area. However, a balance is needed between constraint stiffness and freedom of operation, since surgeons need to remain in control, instead of being restricted by the robotic platform. Finally, the results upon combining visual and haptic feedback are indeed lower overall. This case is the most realistic, since both types of feedback are usually available in robotic surgery. Under this condition, the mean exit duration remained below 1s, and the exit distances are also minimised to sustainably be below a millimeter. This, in turn, leads to the volume being burred being almost completely within the constraint region. Overall, these findings demonstrate that even though visual feedback is considerably useful by itself, it is in no way sufficient. Merging the two types of feedback is what provides the optimal result when deploying this system for burring primitive shapes in geometry. Since the Signature Robot can sustain higher forces than the ones involved in plastic bone cutting, it can be hypothesised that a similar pattern of results can be noted when tackling cadaveric bone specimens.

In addition to the numerical results discussed above, the volunteers were also asked to complete a TLX form, with the amalgamated results shown in Table 2. With a brief examination, one may conclude that the main shortcoming of this platform is the physical demand of operation. Indeed, the direct drive design suffers from an increased weight, which is directly applied to the hand of the operator. This finding suggests that the robot would greatly benefit from the implementation of a gravity compensation algorithm, thus alleviating the load on the operator during surgery. In contrast, most of the volunteers stated to be

sufficiently satisfied with the performance when manipulating the robot, whilst doing so with low effort, meaning that despite being untrained with the system, volunteers achieved the desired results. Lastly, the cases of feedback testing, be it haptic, visual, or both, led to a significantly low value of mental and temporal demand, which is highly encouraging.

4.2 Active Constraint Definition

With regard to the constraint definition and enforcement experiments, the relevant results also allow for some generalisations. Firstly, as mentioned in the previous section, a real-time success rate of about 55% was achieved for the employed network. Even though the positional and angular thresholds, set to 8mm and 5 degrees respectively, were potentially too generous for orthoapedic surgery applications, the results are indeed promising, since one needs to take into account some shortcomings of this examined implementation. First and foremost, SimPS-Net has been developed and trained under a specific dataset (33). The dataset comprises images from the same cadaveric lab as the one utilised in this presented research, however, the specimen used for the dataset was a cadaveric knee. In contrast, the experiment presented here was performed on a phantom knee specimen. The differences are multiple, and this can be the main source of error in network inferences. Furthermore, the network was trained on four tools, a scalpel and an electric burr being among them. The latter was mounted on the robot during tool tracking, so even though the purpose of the network was to solely detect and localise the scalpel, the presence of the burr in the frame could affect the localisation results of the scalpel.

4.3 Active Constraint Enforcement

With the area having been generated, the robot was tested under the safe-zone and the restricted-zone constraints. In the former, the end-effector was freely manipulated within the constraint region, and attractive forces were only applied upon exiting the boundary, thus guiding the operator back within the safety region. The restricted-zone testing was enforced in a different test, with the purpose of preventing a robot from entering the generated region.

4.3.1 Safe-Zone Testing

Firstly examining the safe-zone results on Table 3, the user was allowed to explore the generated area, initially in freehand mode, where no forces were involved, and subsequently in controlled mode, where haptic feedback was enforced. Similarly to the aforementioned plastic bone cutting experiments performed using the same robot, in the case of freehand operation, the user could only depend on spatial awareness, as visual feedback was absent in both cases. By examining the results, it can be stated that the presence of tactile feedback overall improves the performance of an untrained operator, since the exit duration and the distance from the boundary both remain consistently lower. Indeed, in the absence of haptic feedback, the average duration of individual exits increased by 74%, with the maximum duration rising by 82% as well. In contrast, controlled operation ensured the mean duration of each exit remained below one second, thus demonstrating the presence of haptic feedback greatly reduces the exit duration.

Similar findings are noted in terms of exit distance. Precisely, without any tactile information from the robotic platform, the average distance of the robot from the constraint boundary increased by 116%, with a similar rise in the maximum distance by 117%. A mean exit of 2.70mm under controlled testing conditions is a satisfactory result, however, it should be further improved. One way of improving this result is, once again, further tuning the force coefficients involved in Equation 1. In doing so, a stiffer response can lead to even lower exit distances. Another possible improvement is the introduction of visual feedback to the operator, which has been shown to improve controlled operation results in the plastic bone experiments.

A final point to make regarding safe-zone testing involves the number of exits. Specifically, the number of exits is shown to be higher in controlled mode. However, the fact that exits last less and the distances are lower suggests that the operators once again develop a level of dependency on the robot which quickly replaces their reliance on spatial awareness exclusively. As shown in the plastic bone cutting results, a high degree of such reliance, in the absence of visual feedback, can cause a notable degree of error on the exit distances. Nevertheless, the fact that more exits were noted under controlled operation does not indicate sub-optimal performance, it simply demonstrates that a confident operator attempted to more freely explore the constraint region and, upon exiting, the restorative forces functioned as intended, by minimising the distance from the boundary.

4.3.2 Restricted-Zone Testing

Examining the restricted-zone testing, the main result is the fact that across all nine repetitions, no successful penetration of the constraint region was achieved, as intended. As mentioned, a force plateau was enforced, thus creating a maximum cap on the possible force. However, the operator never reached the threshold of 0.1mm, since the force exponentially increased at a distance of about 5mm from the boundary. This was achieved by correctly tuning the value of ψ in Equation 2. As expected with a power law, the force increased at a high rate the closer the robot tooltip was positioned to the boundary. Therefore, this experiment demonstrates that a restricted-zone constraint can be imposed on the Signature Robot and successfully prevent the end-effector from penetrating the defined region. One point to address is that in cases where the operator may wish to enter the defined zone for some unexpected condition in operation, the proposed system should be altered in order to allow for bypass of the constraint, thus allowing the surgeon to remain in control during operation. Irrespective of these additions, however, it is promising to note that no penetration was achieved across the boundary, and thus the purpose of this restriction was fully achieved.

4.4 Conclusion & Future Work

Throughout this paper, both the assembly of an orthopaedic surgical robot and the construction and enforcement of an active constraint "on-the-fly" were presented and analysed, with the constraint ultimately being enforced on the Signature Robot.

In terms of robot development, it was shown that the haptic capabilities of the system can definitely improve the performance in operation, even when it comes to untrained volunteers. Furthermore, the results when operating on plastic bone can be enhanced through the fusion of haptic and visual feedback, thus resulting in submillimeter deviations from a constraint boundary. Ultimately, this novel platform is a promising starting point that addresses the absence of miniaturised orthopaedic robots with haptic feedback capabilities.

Despite these results, the presented robotic platform would definitely benefit from the implementation of gravity compensation, which will significantly lower the physical demand of operating the robot, as the latter was reported to be noteworthy by the volunteers.

Regarding the constraint generation and enforcement, it was demonstrated that a real-time, monocular, RGB camera-based network capable of surgical tool localisation, SimPS-Net is indeed sufficient for the purpose of constraint definition. A constraint area was produced using this technique, with a satisfactory success rate of localisation results, thus allowing the operator to generate regions of interest across patient tissue in operation, without the need to access the preoperative data or to employ cumbersome sensors, such as optical trackers. Furthermore, the absence of optical trackers offers the benefit that the operators need not

train on tools they may be unfamiliar with, nor do the standard surgical tools need to be redesigned in order to accommodate trackable bodies. The constraint enforcement was initially explored in the form of a safe-zone, and it was demonstrated that the operator manipulated the end-effector within the region significantly more successfully in the presence of haptic feedback. Subsequently, a restricted-zone constraint was tested, with no penetration occurring across the defined boundary, thus successfully preventing the robot from entering the region. Both these experiments demonstrate that a constraint can be successfully imposed on a robotic platform after being generated "on-the-fly" using the presented technique. Therefore, patient tissue safety can be further enhanced by integrating this technique along the robotic surgery workflow.

Despite the promising results in constraint generation and enforcement, however, there are several improvements that should be addressed in the future. Firstly, regarding constraint definition, either SimPS-Net will need to be further trained to address different environmental conditions, such as different lighting, various skin colours, or even different body parts. This will increase the versatility of the network, and also improve the success rate noted. In addition, network simplifications can be undertaken in order to improve the achieved framerate in real-time deployment. In terms of constraint enforcement, one future focus will be the identification of the optimal proportionality constant, ψ. This could be done in a manner that each operator may define their own ideal value, instead of trying to identify one general optimal value. This means that a calibration step would be required, which would allow the user to identify this value. Another improvement is the introduction of visual feedback along the pipeline, which would reduce both the duration and the distance from the boundary of each exit.

A final point to make is the need to implement a more robust algorithm for constraint enforcement, both in terms of plastic bone testing and also safe-zone testing. As mentioned, the viscoelastic model stores energy, and can lead to patient tissue injury upon releasing the robot end-effector unintentionally. More robust algorithms, such as the frictional model (34), have been proposed in the literature, and therefore integrating them in the haptic model of Signature Robot would make the entire system safer and easier to manipulate.

REFERENCES

1 .Antonia F. Chen, Gregory S. Kazarian, Galen W. Jessop, and Asim Makhdom. Current concepts review: Robotic technology in orthopaedic surgery. *Journal of Bone and Joint Surgery - American Volume*, 100(22):1984–1992, 2018.

2 .Angela Sorriento, Maria Bianca Porfido, Stefano Mazzoleni, Giuseppe Calvosa, Miria Tenucci, Gastone Ciuti, and Paolo Dario. Optical and Electromagnetic Tracking Systems for Biomedical Applications: A Critical Review on Potentialities and Limitations, 2020.

3 .Johanna Elliott, Jobe Shatrov, Brett Fritsch, and David Parker. Robotic-assisted knee arthroplasty: an evolution in progress. A concise review of the available systems and the data supporting them, 12 2021.

4 .S Walgrave, S Oussedik, and Bone Jt. Comparative assessment of current robotic-assisted systems in primary total knee arthroplasty. *Bone Jt Open*, 4(1):13–18, 2022.

5 .Ming Han Lincoln Liow, Pak Lin Chin, Hee Nee Pang, Darren Keng Jin Tay, and Seng Jin Yeo. THINK surgical TSolution-One®(Robodoc) total knee arthroplasty, 2017.

6 .O. Dianat, A. Nosrat, B. Mostoufi, J. B. Price, S. Gupta, and F. C. Martinho. Accuracy and efficiency of guided root-end resection using a dynamic navigation system: a human cadaver study. *International Endodontic Journal*, 54(5):793–801, 5 2021.

7 .Martin Roche. The MAKO robotic-arm knee arthroplasty system. *Archives of Orthopaedic and Trauma Surgery*, 141(12):2043–2047, 12 2021.

8 .Emily L. Hampp, Morad Chughtai, Laura Y. Scholl, Nipun Sodhi, Manoshi Bhowmik-Stoker, David J. Jacofsky, and Michael A. Mont. Robotic-Arm Assisted Total Knee Arthroplasty Demonstrated Greater Accuracy and Precision to Plan Compared with Manual Techniques. *Journal of Knee Surgery*, 32(3):239–250, 2019.

9 .Mark Wu, Lefko Charalambous, Colin Penrose, Elshaday Belay, and Thorsten M. Seyler. Imageless Robotic Knee Arthroplasty. *Operative Techniques in Orthopaedics*, 31(4), 12 2021.

10 .Fabio Mancino, Christopher W. Jones, Francesco Benazzo, Alessandro Singlitico, Alessandro Giuliani, and Ivan De Martino. Where are We Now and What are We Hoping to Achieve with Robotic Total Knee Arthroplasty? A Critical Analysis of the Current Knowledge and Future Perspectives, 2022.

11 .Kade Collins, Paul A. Agius, Andrew Fraval, and Josh Petterwood. Initial Experience with the NAVIO Robotic-Assisted Total Knee Replacement-Coronal Alignment Accuracy and the Learning Curve. *Journal of Knee Surgery*, 35(12):1295–1300, 10 2022.

12 .Michael Casper, Riddhit Mitra, Rahul Khare, Branislav Jaramaz, Brian Hamlin, Brian McGinley, David Mayman, Jeff Headrick, Kenneth Urish, Mark Gittins, Stephen Incavo, and Vivek Neginhal. Accuracy assessment of a novel image-free handheld robot for Total Knee Arthroplasty in a cadaveric study. *Computer Assisted Surgery*, 23(1):14–20, 1 2018.

13 .Rahul Khare, Branislav Jaramaz, Brian Hamlin, and Kenneth L. Urish. Implant orientation accuracy of a hand-held robotic partial knee replacement system over conventional technique in a cadaveric test. *Computer Assisted Surgery*, 23(1):8–13, 1 2018.

14 .Peter Bollars, Prashant Meshram, Saeed Al Thani, Martijn G.M. Schotanus, and Ali Albelooshi. Achieving functional alignment in total knee arthroplasty: early experience using a second-generation imageless semi-autonomous handheld robotic sculpting system. *International Orthopaedics*, 47(2):585–593, 2 2023.

15 .Olga Adamska, Krzysztof Modzelewski, Jakub Szymczak, Jakub Świderek, Bartosz Maciąg, Paweł Czuchaj, Małgorzata Poniatowska, and Artur Wnuk. Robotic-Assisted Total Knee Arthroplasty Utilizing NAVIO, CORI Imageless Systems and Manual TKA Accurately Restore Femoral Rotational Alignment and Yield Satisfactory Clinical Outcomes: A Randomized Controlled Trial. *Medicina (Lithuania)*, 59(2), 2 2023.

16 .Gary Doan. *An Accuracy and Precision Analysis of the VELYS™ Robotic Assisted Solution for Total Knee Arthroplasty Advisor*. PhD thesis, University of Denver, Denver, 2021.

17 .Gary W. Doan, R. Patrick Courtis, Joseph G. Wyss, Eric W. Green, and Chadd W. Clary. Image-Free Robotic-Assisted Total Knee Arthroplasty Improves Implant Alignment Accuracy: A Cadaveric Study. *Journal of Arthroplasty*, 37(4):795–801, 4 2022.

18 .Vivek Singh, Greg M. Teo, and William J. Long. Versatility and accuracy of a novel image-free robotic-assisted system for total knee arthroplasty, 12 2021.

19 .James D. Sires, Johnathan D. Craik, and Christopher J. Wilson. Accuracy of Bone Resection in MAKO Total Knee Robotic-Assisted Surgery. *Journal of Knee Surgery*, 34(7):745–748, 6 2021.

20 .Stuart A. Bowyer, Brian L. Davies, and Ferdinando Rodriguez Y Baena. Active constraints/virtual fixtures: A survey. *IEEE Transactions on Robotics*, 30(1):138–157, 2014.

21 .Joshua Petersen and Ferdinando Rodriguez y Baena. A Dynamic Active Constraints Approach for Hands-On Robotic Surgery. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 11 2013.

22 .Andrew Sharp and Mitch Pryor. Virtual Fixture Generation for Task Planning with Complex Geometries. *Journal of Computing and Information Science in Engineering*, 21(6), 12 2021.

23 .Alessandro Bettini, Panadda Marayong, Samuel Lang, Allison M. Okamura, and Gregory D. Hager. Vision-assisted control for manipulation using virtual fixtures. *IEEE Transactions on Robotics*, 20(6):953–966, 12 2004.

24 .Jake I Abbott and Allison M Okamura. Virtual Fixture Architectures for Telemanipulation. *Proceedings of the 2003 IEEE International Conference on Robotics &Automation*, 2003.

25 .Murilo Marques Marinho, Bruno Vilhena Adorno, Kanako Harada, and Mamoru Mitsuishi. Dynamic Active Constraints for Surgical Robots Using Vector-Field Inequalities. *IEEE Transactions on Robotics*, 35(5):1166–1185, 10 2019.

26 .Jing Ren, Rajni V. Patel, Kenneth A. McIsaac, Gerard Guiraudon, and Terry M. Peters. Dynamic 3-D virtual fixtures for minimally invasive beating heart procedures. *IEEE Transactions on Medical Imaging*, 27(8):1061–1070, 8 2008.

27 .Spyridon Souipas, Anh Nguyen, Stephen G. Laws, Brian L. Davies, and Ferdinando Rodriguez Y Baena. SimPS-Net: Simultaneous Pose and Segmentation Network of Surgical Tools. *IEEE Transactions on Medical Robotics and Bionics*, 5(3):614–622, 8 2023.

28 .Spyridon Souipas, Stephen G Laws, Ferdinando Rodriguez Y Baena, and Brian L Davies. Towards Miniaturised Collaborative Haptic Robots for Computer Aided Knee Surgery: Signature Robot. *Health Sciences EPiC Series in Health Sciences*, 5, 2022.

29 .Matjaz Jakopec, Ferdinando Rodriguez y Baena, Simon J. Harris, Paula Gomes, Justin Cobb, and Brian L. Davies. The Hands-On Orthopaedic Robot "Acrobot": Early Clinical Trials of Total Knee Replacement Surgery. *IEEE Transactions on Robotics and Automation*, 19(5):902–911, 10 2003.

30 .Sandra G Hart, Moffett Field, California Lowell, and E Staveland. Development of NASA-TLX (Task Load Index)" Results of Empirical and Theoretical Research. Technical report, NASA, 1998.

31 .Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(2):386–397, 2017.

32 .Takio Kurita. Principal Component Analysis (PCA). In *Computer Vision*, pages 1–4. Springer International Publishing, Cham, 2020.

33 .Spyridon Souipas, Anh Nguyen, Stephen Laws, Brian Davies, and Ferdinando Rodriguez Y Baena. 3dStool - A 3D Surgical Tool dataset for detection and pose estimation. https://github.com/SpyrosSou/3dStool, 2 2023.

34 .Stuart Bowyer and Ferdinando Rodriguez Y Baena. Dynamic Frictional Constraints in Translation and Rotation. *Robotics and Automation (ICRA) 2014 IEEE International Conference*, page 6822, 5 2014.