Learning by Observation for Surgical Subtasks: Multilateral Cutting of 3D Viscoelastic and 2D Orthotropic Tissue Phantoms

Adithyavairavan Murali[1]*, Siddarth Sen[1]*, Ben Kehoe[2], Animesh Garg[3], Seth McFarland[2], Sachin Patil[1], W. Douglas Boyd[4], Susan Lim[5], Pieter Abbeel[1], Ken Goldberg[3]

***Abstract*— Automating repetitive surgical subtasks such as suturing, cutting and debridement can reduce surgeon fatigue and procedure times and facilitate supervised tele-surgery. Programming is difficult because human tissue is deformable and highly specular. Using the da Vinci Research Kit (DVRK) robotic surgical assistant, we explore a "Learning By Observation" (LBO) approach where we identify, segment, and parameterize sub-trajectories ("surgemes") and sensor conditions to build a finite state machine (FSM) for each subtask. The robot then executes the FSM repeatedly to tune parameters and if necessary update the FSM structure. We evaluate the approach on two surgical subtasks: debridement of 3D Viscoelastic Tissue Phantoms (3d-DVTP), in which small target fragments are removed from a 3D viscoelastic tissue phantom, and Pattern Cutting of 2D Orthotropic Tissue Phantoms (2d-PCOTP), a step in the standard Fundamentals of Laparoscopic Surgery training suite, in which a specified circular area must be cut from a sheet of orthotropic tissue phantom. We describe the approach and physical experiments, which yielded a success rate of 96% for 50 trials of the 3d-DVTP subtask and 70% for 20 trials of the 2d-PCOTP subtask.**

I. Introduction

Robotic surgical assistants (RSAs), such as Intuitive Surgical's da Vinci® system, have proven highly effective in facilitating precise minimally invasive surgery [10, 36]. Currently, these devices are primarily controlled by surgeons in a local tele-operation mode (master-slave with negligible time delays). Introducing autonomy of surgical subtasks has potential to assist surgeons, reduce fatigue, and facilitate supervised autonomy for remote tele-surgery.

Multilateral manipulation (with two or more arms) has potential to reduce the time required for surgical procedures, reducing the time patients are under anaesthesia and associated costs and contention for O.R. resources. Multilateral manipulation is also necessary for sub-tasks such as suturing; hand-off of tissue or tools between arms is common as each arm has limited dexterity and a workspace that may not cover the entire body cavity. Autonomous multilateral manipulation is of particular interest as surgical robot systems can be configured with 3, 4, or more arms.

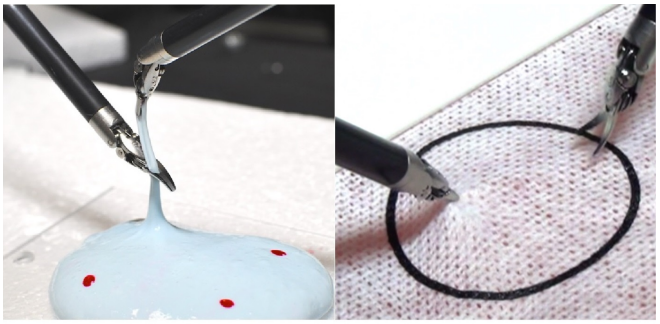

(a) (b)

Fig. 1: Autonomous multilateral surgical subtasks with the da Vinci Research Kit (DVRK). (a) Debridement of 3D Viscoelastic Tissue Phantoms (3d-DVTP) and (b) Pattern Cutting of 2D Orthotropic Tissue Phantoms (2d-PCOTP).

Automating manipulation and cutting presents challenges due to the difficulty of modeling the deformation behavior of highly nonlinear viscoelastic substances and the precision required for cutting.

We apply "Learning By Observation" (LBO) to two surgical subtasks relevant to surgical procedures: debridement of viscoelastic tissue phantoms (3d-DVTP) and pattern cutting of orthotropic deformable tissue phantoms (2d-PCOTP). Surgical debridement is a tedious subtask in which dead or damaged tissue is removed from the body to allow the remaining healthy tissue to heal [13]. 2d-PCOTP is one of five subtasks in the commonly-used Fundamentals of Laparoscopic Surgery training suite. Surgical residents are trained to perform precision pattern cutting, and it is used to evaluate the performance of surgeons [12, 30].

We use the da Vinci Research Kit (DVRK) [15, 16], a research platform built from mechanical components from the first-generation da Vinci robotic surgical system [4] and electronics and software from WPI and Johns Hopkins University.

II. Related Work

Robotic surgical systems have been used in a wide variety of surgical interventions [1, 5, 22, 33, 37]. In this work, we use the da Vinci Research Kit (DVRK) [15, 16], a research platform built from mechanical components from the first-generation of the da Vinci surgical system [4]. Padoy et al. [26] demonstrated execution of a human-robot collaborative suturing task on the DVRK platform. The DVRK platform has been used by a number of groups for tasks

*The two first authors contributed equally to this work.

[1]Department of Electrical Engineering and Computer Sciences; {adithya_murali, siddarthsen, animesh.garg, sachinpatil, pabbeel}@berkeley.edu

[2]Department of Mechanical Engineering; {benk@berkeley.edu, mcf.seth@gmail.com}

[3]IEOR and EECS; {animesh.garg, goldberg}@berkeley.edu

[1–3] University of California, Berkeley; Berkeley, CA 94720, USA

[4]Division of Cardiothoracic Surgery; University of California Davis Medical Center; Sacramento, CA 95817, USA; walter.boyd@ucdmc.ucdavis.edu

[5]Centre for Breast Screening and Surgery, Centre for Robotic Surgery, Singapore; susanlim@berkeley.edu

Fig. 2: Debridement of a 3D Viscoelastic Tissue Phantom (3d-DVTP) with a linear tumor target. This subtask consists of five surgemes: motion, penetration, grasping, retraction, and cutting. Multiple debridement operations are needed to remove a single target.

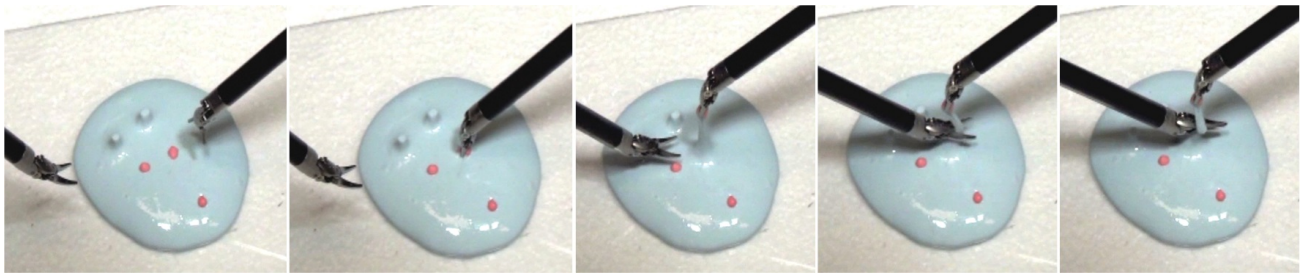

Fig. 3: Debridement of a 3D Viscoelastic Tissue Phantom (3d-DVTP) with spherical tumor targets. This subtask consists of five surgemes: motion, penetration, grasping, retraction, and cutting. The same finite state machine is used as in Figure 2.

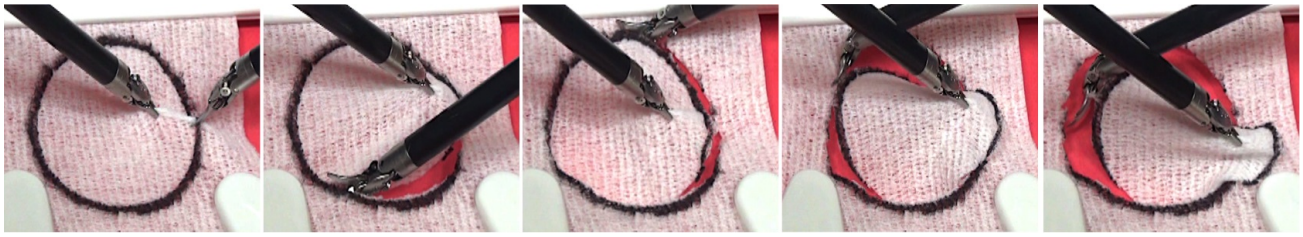

Fig. 4: Pattern Cutting of a 2D Orthotropic Tissue Phantom (2d-PCOTP). The finite state machine includes the following states: circle detection and estimation, warping, grasp OTP (push, grasp, retract), notch cutting (push, close, retract, twist), lower semicircle positioning, lower semicircle cutting, upper semicircle repositioning, upper semicircle cutting, active sensing for attachment detection (pull), and final cutting.

like tissue palpation using an ultrasound probe for tumor detection [6] and autonomous tool tracking in ultrasound images using an auxiliary da Vinci arm [21]. We evaluate the feasibility of autonomously performing surgical subtasks that involve cutting of deformable tissue phantoms using the DVRK. We consider debridement of viscoelastic tissue phantoms, which extends our previous work in autonomous debridement [17] with the Raven surgical robot [14], an open-architecture research platform similar to the DVRK.

Manipulation of deformable materials, particularly cutting, is an area of research interest in robotic surgery [23] and in computer graphics and computational geometry [9, 38]. However, high fidelity models of viscoelastic tissue deformations are computationally expensive due to the need for re-meshing and finite element simulations.

Prior work has explored the use of expert demonstrations to handle deformations in environment without explicit models and simulations. Reiley et al. [29] proposed a demonstration-based framework that used Gaussian Mixture Models (GMMs) for motion generation. Van den Berg et al. [34] proposed an iterative technique to learn a reference trajectory and execute it at higher than demonstration speeds for suture knot tying. This work was recently extended by Osa et al. [25] to deal with dynamic changes in the environment, but with an industrial manipulator. Mayer et al. [20] use principles of fluid dynamics and Schulman et al. [31] use non-rigid registration techniques to generalize human demonstrations to similar, yet previously unseen, initial conditions. These approaches are broadly classified under the category of Learning From Demonstrations (LfD) [2, 8], where demonstration trajectories are directly modified for generalizing to test situations. Learning By Observation (LBO), as defined by Dixon [11], includes Learning From Demonstrations (LfD) as one of the key components, as described in Section IV.

III. Subtasks for Case Study

A. Debridement of 3D Viscoelastic Tissue Phantoms (3d-DVTP)

Surgical debridement is a tedious surgical subtask in which dead or damaged tissue is removed from the body to allow the remaining healthy tissue to heal faster [3, 13]. As shown in Figures 2 and 3, we introduce an extension to the debridement task presented in our previous work [17].

We use a viscoelastic tissue phantom made from a mixture of Elmer's Glue, borax, and water. Embedded in the phantom are multiple targets of viscoelastic material of a tougher consistency. These targets represent damaged or tumorous tissue that must be removed from the surrounding phantom.

Autonomous surgical debridement of viscoelastic tissue requires perception to locate damaged tissue, grasp and motion planning to determine collision free trajectories for one or more arms and grippers to grasp them, and careful coordination of arms to retract and separate the damaged tissue from the VTP.

In this work, we consider targets that form convex regions, and in particular regions that, after any debridement operation, the targets left still form one or more convex regions. The maximum width of target material that can be removed in one debridement operation is d_w. We consider two types of convex regions: spherical regions of diameter $< d_w$, and linear regions of width $< d_w$.

B. Pattern Cutting of 2D Orthotropic Tissue Phantom (2d-PCOTP)

The second subtask, 2d-PCOTP, is shown in Figure 4. The Fundamentals of Laparoscopic Surgery (FLS) is the standard training regimen for medical students in laparoscopic surgery and consists of a suite of five subtasks of increasing complexity [30]. The second subtask in this suite is called "Pattern Cutting". This subtask features a 50 mm diameter, 2 mm thick circular pattern marked on a 4×4 inch square of surgical gauze suspended by clips. The subtask is complete once the circle has been cut from the surrounding gauze.

Metric: The FLS suite states that deviations under 2 mm from the line are not penalized. We define an inner circle, C_I, and an outer circle, C_O, at 2 mm inside and outside the pattern, respectively. We define the error E_I as the sum of the areas between C_I and the cut line falling inside C_I, and the error E_O as the area between C_O and the cut line falling outside C_O. With A the area of the annulus (Area of C_O-Area of C_I), we define the quality score as:

$$Q = 100\left[1 - \frac{E_I + E_O}{A}\right]$$

This quality corresponds to the symmetric difference of the cut and the circle (as a percentage).

IV. Learning By Observation

Learning By Observation (LBO) [11], involves observing human-operated demonstrations of a subtask and segmenting these demonstrations into subtrajectories and transition conditions. In contrast to the unsupervised learning used by Dixon, we consider supervised learning in which the user segments the trajectories.

Hager et al. propose a "Language of Surgery" with "surgemes" analogous to phonemes [19, 28, 35]. Manual segmentation of surgemes in demonstrations have been used for understanding and recognizing surgical skills and subtasks, and for evaluating surgeon skills [28, 35].

Our LBO approach is illustrated in Figure 5.

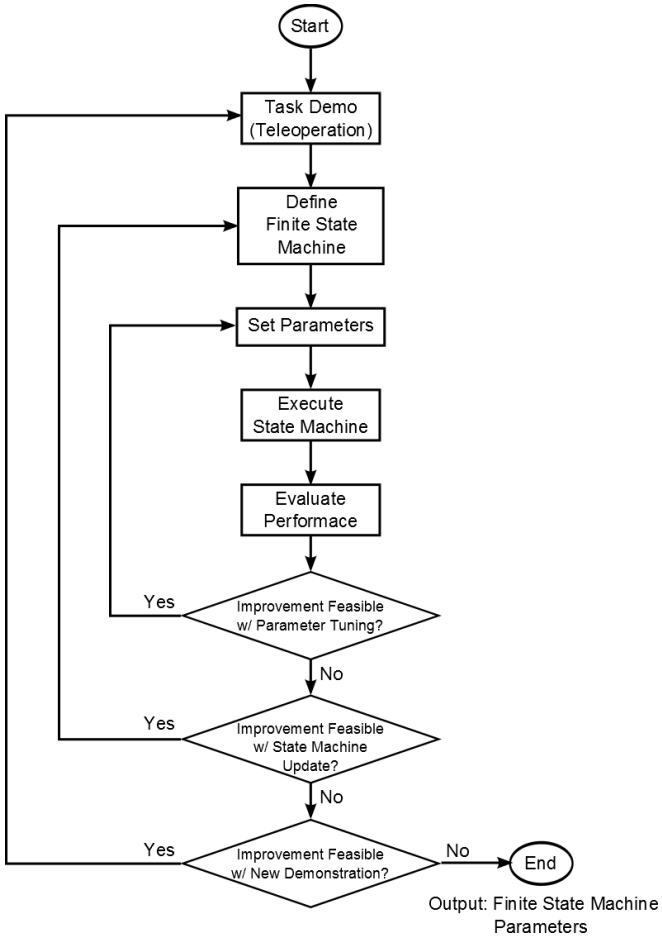

Fig. 5: The Learning By Observation process. First, a domain expert performs a demonstration using teleoperation. Observation of the demonstration is used to construct an FSM. Executions of the state machine are observed, and human observation is used to determine if the revision of parameters, state machine, or demonstration can feasibly improve performance.

V. System Architecture

A. da Vinci Research Kit (DVRK)

The da Vinci Research Kit (DVRK) is a development platform provided by Intuitive Surgical to advance research in teleoperated robotic surgical systems. It consists of proprietary hardware from the first-generation da Vinci "classic", and open-source electronics and software developed by WPI and Johns Hopkins University [16]. The robot hardware consists of two robotic laparoscopic arms, termed "Patient-Side Manipulators" (PSMs), and the Surgeon Console for teleoperating with a stereo viewer, two master controllers, termed "Master Tool Manipulators" (MTMs), and a set of foot pedals. The PSMs have interchangeable tools. We use two tools: the Large Needle Driver and the Curved Scissors. The Large Needle Driver is a grasper with 6 mm fingers. The Curved Scissors is a cutting instrument 10 mm in length. The PSM manipulates the attached instruments about a fixed point called the remote center of motion. The PSMs each have 6 degrees of freedom plus a grasp degree of freedom.

The electronics for the DVRK are provided by Fisher et al. at WPI and Kazanzides et al. at Johns Hopkins University

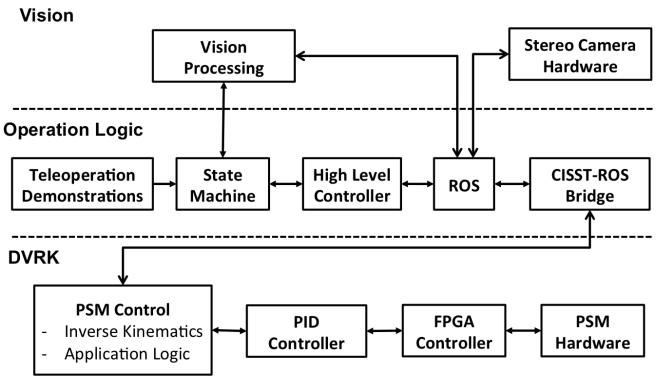

Fig. 6: Software Architecture. The software consists of three layers: vision, operation logic, and the DVRK system software. The open-source DVRK software is developed and supported by Johns Hopkins University [16]. The operation logic is based on code developed for our previous work on the Raven surgical robot [17], adapted for the da Vinci and the LBO framework. Our vision system, as described in Section V-C, consists of two Prosilica GigE GC1290C cameras with 6 mm focal length lenses and a 4.68 cm separation.

as a replacement for the proprietary Intuitive Surgical electronics, which are not provided as a part of the DVRK. The electronics entail a FPGA for centralized computation and distributed I/O architecture while an IEEE-1394 (FireWire) bus is used for a low-latency interface with the mechanical hardware. These electronics, developed under an NSF grant, are open-source and not specific to the da Vinci system. We use a hardware mount designed by Okamura's group.

B. System Software

Software to control the da Vinci hardware is provided for the DVRK by JHU with their cisst/SAW libraries. This component-based framework provides DVRK-specific components to communicate with the electronics, along with generic components to enable PID control, teleoperation, recording, GUI development, and integration with ROS. The changes to the DVRK made as part of this research are openly shared with the DVRK community.

ROS [27] allowed us to easily port our software for the Raven II open-source surgical robotics platform, used in our previous work [17], to provide high-level state machines and motion. The da Vinci and Raven kinematic chains differ only in parameter values, which meant only the robot model needed to be changed and the ROS communication adapted for the DVRK's different input API. The high level architecture for our system is shown in Figure 6. We use the inverse kinematics and PID controllers from the DVRK system software. This allows us to control the robot using pose commands, working directly in Cartesian space.

C. Vision System

Due to tissue specularity, perception using RGBD sensing is not feasible. We use two fixed stereo camera pairs, each composed of two Prosilica GigE GC1290C cameras with 6 mm focal length lenses and a 4.68 cm separation. We use HSV (Hue, Saturation, Value) separation, where different

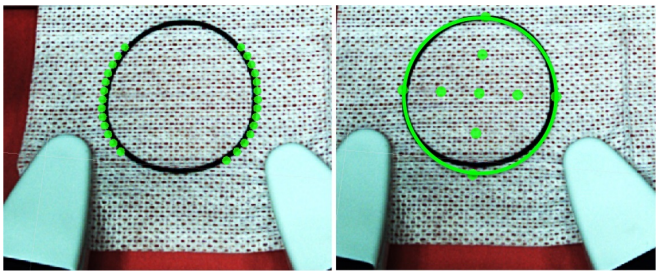

Fig. 7: Estimation of elliptical pattern on the Orthotropic Tissue Phantom in the 2d-PCOTP subtask. Due to deformation, the pattern is not a perfect ellipse. The perception process starts with a raw image and uses HSV thresholding to find a contour of the pattern outline. Then, stereo matching is performed to estimate the 3D location of points along the left and right quadrants of the circle (a). Finally, using the set of 3D points on the pattern boundary, an ellipse is fit to these 3D points using least-squares (b).

materials occupy disjoint ranges in the color space. We use a click interface to manually select pixels of each material in view of the camera. For each of the materials, we find a range of HSV values that contains all of its pixels, while excluding the HSV values of pixels from other materials, which for the subtasks in this work was always possible due to a small number of different materials, the color of which we had control over. We perform this process after any change to the material properties or lighting conditions. We use the open-source OpenCV library [7].

VI. Algorithmic Approach

A. Debridement of 3D Viscoelastic Tissue Phantom (3d-DVTP)

In the 3d-DVTP subtask, we first find the 3D centroids of the targets. To do this, we find the 2D centroids of each target in each of the left and right stereo images. For each image, we use HSV separation to find the pixels corresponding to the VTP. We use this as a bounding box to find the pixels of targets.

We use OpenCV to find the contours among the target pixels, rejecting contours that contain less than $p_1 = 100$ pixels. These contours often have holes, which impairs the centroid estimation of the contours. To remove holes in contours, we first apply dilation: for each pixel in the image, a square window with $p_2 = 3$ pixels per side is taken, and the value of the center pixel is taken to be the maximum value of all the pixels in the window. This "dilates" the targets in the image.

For each contour in each of the two stereo images, we find the centroid. We take the left-most centroid in the left image and find the corresponding centroid in the right image based on a sliding window. The resulting disparity gives the 3D centroid. We repeat this procedure until all matching centroids have been found. If no matching centroids are found, we consider the subtask to be complete, unless no targets have been debrided, in which case we consider the subtask to have failed.

Through teleoperation, we found that the surgemes necessary for debridement of the VTP were well-represented by

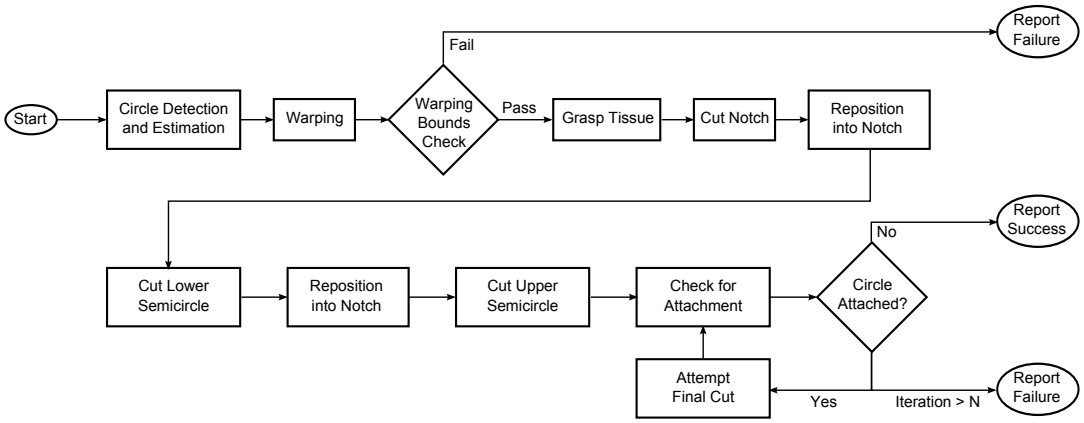

Fig. 8: Finite State Machine for 2d-PCOTP. This subtask consists of ten states, with the surgemes for states noted in parentheses: circle detection and estimation, warping, grasp tissue (push, grasp, retract), notch cutting (push, close, retract, twist), lower semicircle positioning, lower semicircle cutting, upper semicircle repositioning, upper semicircle cutting, active sensing for attachment detection (pull), and final cutting. The iteration limit N is the parameter of the finite state machine $p_{9,3} = 2$.

linear paths. The debridement task consists of five surgemes: motion, penetration, grasping, retraction, and cutting, with parameter vector $p = (p_1, \ldots, p_9)$, with p_1 and p_2 the vision parameters given above.

The motion surgeme consists of moving the gripper to a point $p_3 = 1.2$ cm directly above the target. The penetration surgeme moves vertically down until the tips of the gripper fingers are approximately $p_4 = 4$ mm into the tissue. The grasping surgeme is the closure of the gripper on the target without movement. The retraction surgeme pulls the target and surrounding material $p_5 = 3$ cm vertically. In the cutting surgeme, the cutting tool moves until the tips of the cutting tool are $p_6 = 4$ mm past the center of the retracted material, $p_7 = 2.8$ cm below the gripper and then the scissors are closed. In the final discarding surgeme, the removed material is dropped into a receptacle.

For the retraction surgeme, the gripper moves at a speed of $p_8 = 0.5$ cm/s to prevent the VTP material from snapping. For the other surgemes, the tools move at $p_9 = 3$ cm/s.

B. Pattern Cutting of 2D Orthotropic Tissue Phantom (2d-PCOTP)

Using LBO, we construct the finite state machine shown in Figure 8. Each state may execute multiple surgemes, and may have one or more parameters. We identified ten states, two sensor conditions, and thirteen surgemes.

1) Circle Detection and Estimation: The first state is Circle Detection and Estimation. In this state, we find the 3D position of the circle pattern.

We find the outer contour of the circle using thresholding and a dilation step similar to the one in Section VI-A, but with a $p_{1,1} = 5$ pixel window for a greater dilation effect due to the larger size of the contour. We then fit an ellipse to the circle contour. To find 3D points that we can use for ellipse fitting from the contour of the image, we find correspondences between the contours in the left and right stereo images. We begin by finding the leftmost pixel in the left contour, which we call the *anchor pixel*. We take the intersection of the left and right contours with $p_{1,2} = 15$ horizontal lines across each image. One of these lines in through the anchor pixel. All lines are spaced $p_{1,3} = 15$ pixels apart, with seven above and seven below the anchor pixel. Each line produces a pair of intersections with the contour in each image, which are matched with the corresponding pair in the other image, giving two 3D points (see Figure 7a). Because the stereo cameras are separated horizontally, it is difficult to get stereo correspondence for points along the top and bottom areas of the circle. We assume that all the points lie in the horizontal plane, so the points are projected onto that plane. We then use least squares to fit an ellipse to the points in the plane (see Figure 7b).

2) Warping: LBO contains a trajectory translation step, which we include as part of our finite state machine. In this step, we find a rigid translation between the circle pattern detected for the demonstration and the current circle pattern as detected in the previous state. This allows us to translate our surgeme trajectories to the new environment.

We manually set bounds on the translation to ensure safety of the resulting trajectories. Due to the close proximity of the OTP fixtures, translation in the y direction was limited to $p_{2,1} = 5$ mm. If the transformation produced by the previous state was beyond these bounds, the system reports failure.

3) Grasp Tissue: The third state, Grasp Tissue, has the purpose of pulling the OTP taut to allow the cutting tool to make an incision. We used demonstration surgemes for this state. We observed three surgemes in the demonstration. First, the open gripper pushes the material down until it contacts the surface below. Then, the gripper closes, folding and grasping a section of the OTP. Finally, the gripper lifts vertically to pull the OTP taut. This forms a ridge, which the cutting tool takes advantage of to cut a notch in the material, as can be seen in Figure 4a.

4) Cut Notch: The fourth state, Cut Notch, is the most complicated of the demonstration trajectories. The FLS rules allow a surgeon to either cut a notch to begin the subtask or to cut in from the edge.

We observed a cycle of surgemes in the demonstration. First, with the cutting tool open, the tool pushes down on the OTP. Second, the cutting tool closes. Third, the cutting tool retracts. This process is repeated three times. However, on the second execution, after retraction, a complicated twisting maneuver was observed in the demonstrations.

5) Reposition into Notch for Lower Semicircle: The fifth state, Reposition into Notch for Lower Semicircle, uses a single demonstration surgeme. We observed that this trajectory approached the notch along the line that it would start cutting.

6) Cut Lower Semicircle: The sixth state, Cut Lower Semicircle, uses a single demonstration surgeme. The surgeme cuts along the lower arc of the pattern approximately halfway around.

7) Reposition into Notch for Upper Semicircle: The seventh state, Reposition into Notch for Upper Semicircle, uses a single demonstration surgeme. Similar to the earlier repositioning, the demonstration trajectory slowly approached the point to be cut along the line of cutting.

8) Cut Upper Semicircle: The eighth state, Cut Upper Semicircle, uses a single demonstration surgeme. The surgeme cuts along the lower arc of the pattern approximately halfway around.

9) Check for Attachment: The ninth state, Check for Attachment, occurs at the end of the Cut Upper Semicircle state. It combines two sensor measurements with a single surgeme. The purpose of the state is to determine if the circle pattern has been successfully separated from the surrounding OTP. To determine this, we use active sensing to attempt to deform the OTP by moving the circle. First, we image a $p_{9,1} = 260 \times 200$ pixel window at a known offset from the center of the circle. This window falls on the left edge of the OTP. Then, the gripper moves 8 mm to the right (that is, away from the left edge). Finally, we re-image the window. Then, we use the `matchTemplate` function of OpenCV to compute a difference metric for the two images [24]. Repeated observations guide us to set the threshold at $p_{9,2} = 0.9$ for determining that the edge has not moved.

If the edge has not moved, the state machine terminates, reporting success. If the edge has moved, we judge that the circle pattern has not been successfully separated, and so we attempt the last state, Final Cutting. This forms a loop, re-checking the attachment after the Final Cutting state. However, this loop is executed a maximum of two times. If the circle is judged to attached after $p_{9,3} = 2$ Final Cutting attempts, the state machine terminates, declaring failure.

10) Final Cutting: The tenth state is Final Cutting, which consists of a single surgeme comprising a multi-arm maneuver. The cutting tool moves $p_{10,1} = 2$ cm forward (that is, continuing along the arc it started with Cut Upper Semicircle). To avoid colliding with the gripper arm, the gripper arm moves $p_{10,2} = 1.5$ cm in the same direction, and $p_{10,3} = 1$ cm towards the cutting tool.

Trial	Length (mm)	Outcome	Retractions	Cut Failures	Time (s)	
					Total	Mean
1	21	Success	3	0	70	20.3
2	22	Success	3	0	70	20.3
3	27	Success	3	0	73	21.3
4	27	Success	4	1	94	20.5
5	24	Success	3	0	73	21.3
					76	20.8

TABLE I: Results for 3d-DVTP with linear tumor targets. We performed five trials, all of which succeeded in fully debriding the targets. Four trials required three retractions to complete, while one required four retractions and also experienced a cut failure. The average total time was 76 seconds, with a standard deviation of 10.2. The mean time of debridement per target was 20.8 seconds.

Trial	Targets	Failures		Time (s)	
		Detection	Cut	Total	Mean
1	5	0	0	128	23.2
2	5	0	0	127	23.0
3	5	0	0	125	22.6
4	5	0	0	128	23.2
5	5	0	0	128	23.2
6	5	0	0	127	23.0
7	5	1	1	103*	23.5
8	5	0	0	125	22.6
9	5	0	0	125	22.6
10	5	0	0	124	22.4
	50	1	1	—	22.3

TABLE II: Results for 3d-DVTP with spherical tumor targets 2 mm in diameter. Ten trials were performed, with five targets in each trial. In nine of the ten trials, all five targets were successfully debrided. In the remaining trial, the fourth target experienced a cut failure, where target was not entirely severed from the VTP. Subsequent to this, the fifth target failed to be detected. This detection failure caused the total time, marked above by *, to be lower than other trials. The total success rate was 48 out of 50 targets, or 96%. The average time per target was 25.3 seconds. The adjusted mean was 22.3 seconds.

VII. Experimental Evaluation

A. Debridement of 3D Viscoelastic Tissue Phantom (3d-DVTP)

Using the FSM from Section VI-A, we performed the 3d-DVTP subtask with two kinds of tumor targets: linear and spherical.

1) Linear Tumor Targets: We used targets of lengths between 21 and 27 mm. We ran 5 trials with the targets in four different orientations. Four trials required three retractions to complete, while one required four retractions. The average total time was 76 seconds, with a standard deviation of 10.2. The mean time of debridement per target was 20.8 seconds.

2) Spherical Tumor Targets: We performed a randomization procedure to place spherical targets in the VTP. The debris was placed on the VTP using randomly-generated coordinates. Using the visual segmenting described in Section V-C, we uniformly sampled a rectangle containing the VTP contour and kept only samples falling inside the contour. We also rejected samples falling within 40 pixels of existing samples, so that the targets were not too close together. The coordinates were overlaid on the picture, and this was used to place the targets in the VTP.

We performed 10 trials with 5 spherical targets each. The

Trial	Success	Score	Failed State	Transl. (mm)		Total Time
				x	y	
Demonstration		99.86		0.0	0.0	263
1	Success	99.81	—	26.4	-1.0	284
2	Failure	—	Notch	2.0	-0.5	130*
3	Failure	—	Notch	1.2	-3.0	120*
4	Success	94.52	—	4.5	-2.1	289
5	Failure	—	L.S.	2.0	-1.4	115*
6	Success	97.32	—	-1.2	-2.2	283
7	Success	99.12	—	4.0	-0.9	282
8	Failure	—	Notch	3.6	-0.9	131*
9	Failure	—-	U.S.	8.1	0.2	248*
10	Success	98.89	—	5.6	-0.4	279
11	Failure	—-	Notch	8.5	-1.8	129*
12	Success	99.87	—	5.6	-0.8	279
13	Success	100.00	—	6.6	0.4	284
14	Success	99.96	—	2.3	-1.6	285
15	Success	99.86	—	3.0	0.3	283
16	Success	98.96	—	9.3	-0.4	284
17	Success	98.39	—	8.5	-0.7	285
18	Success	98.94	—	10.5	-0.7	284
19	Success	98.85	—	9.3	0.5	284
20	Success	99.98	—	6.8	0.8	284
Mean	70%	98.89		6.5	1.0	284
Std. Dev.		1.47		5.6	0.8	2.5

TABLE III: Results for 2d-PCOTP. Twenty trials were performed, with a 70% success rate. The mean completion time for the successful trials (excluded times marked with an asterisk) was 284 seconds, less than the required limit of 300 seconds. The mean quality of successful trials was 99.89. For the Failed State column, "L.S." and "U.S" stand for the Lower and Upper Semicircle Cutting states, respectively. The average translation of the circle from its position in the demonstration was 6.5 mm in the x direction and 1.0 mm in the y direction.

results are shown in Table II. We measured the total runtime of the debridement, from which the mean per-target time was calculated.

In nine of the ten trials, all five targets were successfully debrided. In the remaining trial, the fourth target experienced a cut failure. Subsequent to this, the fifth target failed to be detected. The total success rate was 48 out of 50 targets, or 96%. The mean time of debridement per target was 22.3 seconds.

B. Pattern Cutting of Orthotropic Tissue Phantom (2d-PCOTP)

For 2d-PCOTP, we used the equipment from the FLS kit using one layer of gauze.

Using the FSM and parameter vector, the system performed 20 trials. The translation in the circle from its position during the demonstration, as detected by our ellipse-fitting algorithm, was up to 26.4 mm in the x direction (left to right), and up to 3 mm in the y direction, with an average of 6.5 mm and 1.0 mm, respectively. This was very constrained by the clips holding the material in place, which the tool approaches closely in the demonstration trajectory.

The results are shown in Table III. Of the six trials that failed, four of these failed in the fourth state, notch cutting. One trial failed for each of the upper and lower semicircle cutting states due to the deformation of the material not matching the deformation experienced during the demonstration. The average execution time for the successful trials was 284 seconds, with a standard deviation of 2.5 seconds. The variation in time was due to differences in the execution of the final state. All of the trials completed in under 300 seconds. In FLS, "expert proficiency" is granted when the task is completed in 162 seconds with no errors [12].

We found the average quality to be 99.89, with a standard deviation of 1.47. This is slightly higher than the quality of the demonstration, which was 99.86. The minimum quality of the autonomous system for a successful trial was 94.52.

VIII. Conclusion and Future Work

Initial experiments suggest that thse subtasks can be reliably automated using LBO but performance times are at least twice as slow as expert human teleoperation. The next steps are to continue parameter tuning to further reduce execution time, extend the vision system to allow denser arrangements of tumor targets, develop analytic models of deformable tissue, and apply the LBO framework to other FLS subtasks, axillary dissection [18], and cardiothoracic vein harvesting [32].

Acknowledgments

We thank our collaborators, in particular Allison Okamura, Greg Hager, Blake Hannaford, and Jacob Rosen. We thank Intuitive Surgical, and in particular Simon DiMao, for making the DVRK possible. We also thank the DVRK community, including Howie Choset, Anton Deguet, James Drake, Greg Fisher, Peter Kazanzides, Tim Salcudean, Nabil Simaan, and Russ Taylor. We also thank Aliakbar Toghyan, Barbara Gao, Raghid Mardini, and Sylvia Herbert for their assistance on this project. This work is supported in part by a seed grant from the UC Berkeley Center for Information Technology in the Interest of Science (CITRIS), by the U.S. National Science Foundation under Award IIS-1227536: Multilateral Manipulation by Human-Robot Collaborative Systems, by AFOSR-YIP Award #FA9550-12-1-0345, and by Darpa Young Faculty Award #D13AP00046.

References

[1] R. Alterovitz and K. Goldberg, *Motion Planning in Medicine: Optimization and Simulation Algorithms for Image-guided Procedures*. Springer, 2008.

[2] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robotics and Autonomous Systems*, vol. 57, no. 5, pp. 469–483, 2009.

[3] C. E. Attinger, E. Bulan, and P. A. Blume, "Surgical debridement: The key to successful wound healing and reconstruction," *Clinics in podiatric medicine and surgery*, vol. 17, no. 4, p. 599, 2000.

[4] G. Ballantyne and F. Moll, "The da Vinci Telerobotic Surgical System: The Virtual Operative Field and Telepresence Surgery," *Surgical Clinics of North America*, vol. 83, no. 6, pp. 1293–1304, 2003.

[5] R. A. Beasley, "Medical Robots: Current Systems and Research Directions," *Journal of Robotics*, vol. 2012, 2012.

[6] S. Billings, N. Deshmukh, H. Kang, R. Taylor, and E. Boctor, "System for Robot-Assisted Real-Time Laparoscopic Ultrasound Elastography," in *SPIE Medical Imaging*, 2012.

[7] G. Bradski and A. Kaehler, *Learning OpenCV: Computer vision with the OpenCV library*. O'Reilly, 2008.

[8] S. Calinon, "Robot Programming by Demonstration," in *Springer handbook of robotics*, 2008, pp. 1371–1394.

[9] N. Chentanez, R. Alterovitz, D. Ritchie, L. Cho, K. Hauser, K. Goldberg, J. R. Shewchuk, and J. F. O'Brien, *Interactive Simulation of Surgical Needle Insertion and Steering*, 2009, vol. 28, no. 3.

[10] S. A. Darzi and Y. Munz, "The impact of minimally invasive surgical techniques," in *Annu Rev Med.*, vol. 55, 2004, pp. 223–237.

[11] K. Dixon, "Inferring User Intent for Learning by Observation," Ph.D. dissertation, Carnegie Mellon University, 2004.

[12] G. Dulan, R. V. Rege, D. C. Hogg, K. M. Gilberg-Fisher, N. A. Arain, S. T. Tesfay, and D. J. Scott, "Developing a comprehensive, proficiency-based training program for robotic surgery," *Surgery*, vol. 152, no. 3, pp. 477–488, 2012.

[13] M. Granick, J. Boykin, R. Gamelli, G. Schultz, and M. Tenenhaus, "Toward a common language: Surgical wound bed preparation and debridement," *Wound repair and regeneration*, vol. 14, no. s1, pp. 1–10, 2006.

[14] B. Hannaford, J. Rosen, D. C. Friedman, H. King, P. Roan, L. Cheng, D. Glozman, J. Ma, S. Kosari, and L. White, "Raven-II: An open platform for surgical robotics research," *IEEE Transactions on Biomedical Engineering*, vol. 60, pp. 954–959, Apr. 2013.

[15] Intuitive Surgical, "da Vinci Research Kit," 2014. [Online]. Available: http://research.intusurg.com/dvRK

[16] P. Kazanzides, Z. Chen, A. Deguet, G. Fischer, R. Taylor, and S. DiMaio, "An Open-Source Research Kit for the da Vinci Surgical System," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2014.

[17] B. Kehoe, G. Kahn, J. Mahler, J. Kim, A. Lee, A. Lee, K. Nakagawa, S. Patil, W. Boyd, P. Abbeel, and K. Goldberg, "Autonomous Multilateral Debridement with the Raven Surgical Robot," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2014.

[18] S. M. Lim, C. K. Kum, and F. L. Lam, "Nerve-sparing Axillary Dissection using the da Vinci Surgical System," *World Journal of Surgery*, vol. 29, no. 10, pp. 1352–1355, 2005.

[19] H. C. Lin, I. Shafran, D. Yuh, and G. D. Hager, "Towards Automatic Skill Evaluation: Detection and Segmentation of Robot-assisted Surgical Motions," *Computer Aided Surgery*, vol. 11, no. 5, pp. 220–230, 2006.

[20] H. Mayer, I. Nagy, D. Burschka, A. Knoll, E. Braun, R. Lange, and R. Bauernschmitt, "Automation of Manual Tasks for Minimally Invasive Surgery," in *Int. Conf. on Autonomic and Autonomous Systems*, 2008, pp. 260–265.

[21] O. Mohareri, P. Black, and S. E. Salcudean, "da Vinci Auxiliary Arm as a Robotic Surgical Assistant for Semi-Autonomous Ultrasound Guidance during Robot-Assisted Laparoscopic Surgery," in *Hamlyn Symposium on Medical Robotics (HSMR)*, 2014.

[22] G. Moustris, S. Hiridis, K. Deliparaschos, and K. Konstantinidis, "Evolution of Autonomous and Semi-Autonomous Robotic Surgical Systems: A Review of the Literature," *Int. Journal of Medical Robotics and Computer Assisted Surgery*, vol. 7, no. 4, pp. 375–392, 2011.

[23] H.-W. Nienhuys and A. F. van der Stappen, "A surgery simulation supporting cuts and finite element deformation," in *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2001*. Springer, 2001, pp. 145–152.

[24] OpenCV, "Template matching." [Online]. Available: http://docs.opencv.org/doc/tutorials/imgproc/histograms/template_matching/template_matching.html

[25] T. Osa, N. Sugita, and M. Mamoru, "Online Trajectory Planning in Dynamic Environments for Surgical Task Automation," in *Robotics: Science and Systems (RSS)*, 2014.

[26] N. Padoy and G. Hager, "Human-Machine Collaborative Surgery using Learned Models," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2011, pp. 5285–5292.

[27] M. Quigley, B. Gerkey, K. Conley, J. Faust, T. Foote, J. Leibs, E. Berger, R. Wheeler, and A. Ng, "ROS: an open-source Robot Operating System," in *ICRA Workshop on Open Source Software*, 2009.

[28] C. E. Reiley and G. D. Hager, "Task versus Subtask Surgical Skill Evaluation of Robotic Minimally Invasive Surgery," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2009, pp. 435–442.

[29] C. E. Reiley, E. Plaku, and G. D. Hager, "Motion generation of robotic surgical tasks: Learning from expert demonstrations," in *Engineering in Medicine and Biology Society (EMBC), 2010 Annual International Conference of the IEEE*. IEEE, 2010, pp. 967–970.

[30] E. Ritter and D. Scott, "Design of a proficiency-based skills training curriculum for the fundamentals of laparoscopic surgery," vol. 14, no. 2, pp. 107–112, 2007.

[31] J. Schulman, A. Gupta, S. Venkatesan, M. Tayson-Frederick, and P. Abbeel, "A Case Study of Trajectory Transfer through Non-Rigid Registration for a Simplified Suturing Scenario," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2013, pp. 4111–4117.

[32] K. D. Stahl, W. D. Boyd, T. A. Vassiliades, and H. L. Karamanoukian, "Hybrid Robotic Coronary Artery Surgery and Angioplasty in Multivessel Coronary Artery Disease," *The Annals of thoracic surgery*, vol. 74, no. 4, pp. 1358–1362, 2002.

[33] R. Taylor, A. Menciassi, G. Fichtinger, and P. Dario, "Medical Robotics and Computer-Integrated Surgery," *Springer Handbook of Robotics*, pp. 1199–1222, 2008.

[34] J. Van Den Berg, S. Miller, D. Duckworth, H. Hu, A. Wan, X. Fu, K. Goldberg, and P. Abbeel, "Superhuman Performance of Surgical Tasks by Robots using Iterative Learning from Human-Guided Demonstrations," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2010, pp. 2074–2081.

[35] B. Varadarajan, C. Reiley, H. Lin, S. Khudanpur, and G. Hager, "Data-derived Models for Segmentation with Application to Surgical Assessment and Training," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2009, pp. 426–434.

[36] R. Veldkamp, E. Kuhry, W. Hop, J. Jeekel, G. Kazemier, H. J. Bonjer, E. Haglind, L. Pahlman, M. A. Cuesta, S. Msika, *et al.*, "Laparoscopic surgery versus open surgery for colon cancer: short-term outcomes of a randomised trial," *Lancet Oncol*, vol. 6, no. 7, pp. 477–484, 2005.

[37] A. Wolf and M. Shoham, "Medical Automation and Robotics," in *Springer Handbook of Automation*, 2009, pp. 1397–1407.

[38] H. Zhang, S. Payandeh, and J. Dill, "On Cutting and Dissection of Virtual Deformable Objects," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, vol. 4, 2004, pp. 3908–3913.