Using Intermittent Synchronization to Compensate for Rhythmic Body Motion During Autonomous Surgical Cutting and Debridement

Vatsal Patel*[1], Sanjay Krishnan*[1], Aimee Goncalves[1], Carolyn Chen[1],
Walter Doug Boyd[2], Ken Goldberg[1]

***Abstract*— Anatomical structures are rarely static during a surgical procedure due to breathing, heartbeats, and peristaltic movements. Inspired by observing an expert surgeon, we propose an intermittent synchronization with the extrema of the rhythmic motion (i.e., the lowest velocity windows). We performed 2 experiments: (1) pattern cutting and (2) debridement. In (1), we found that the intermittent synchronization approach, while 1.8× slower than tracking motion, is significantly more robust to noise and control latency, and it reduces the max cutting error by 2.6× except when motion is along 3 or more orthogonal axes. In (2), a baseline approach with no synchronization succeeds in 62% of debridement attempts while intermittent synchronization achieves an 80% success rate.**

I. Introduction

Robotic Surgical Assistants (RSAs), such as Intuitive Surgical's da Vinci, can facilitate precise minimally invasive surgery [1]. RSAs are currently controlled by surgeons using pure tele-operation, often in a master-slave configuration. Operating such robots requires uninterrupted attention, control, and a fundamental understanding of the surgical system. Automation of surgical sub-tasks has the potential to reduce surgeon tedium and fatigue, operating time, and enable supervised tele-surgery over high-latency networks. Several recent papers have proposed techniques for introducing limited autonomy in surgery, although few have addressed this in a dynamic environment. [2–8].

Virtual simulators are widely used in surgical training to simulate anatomical motions such as breathing, heart beats, or peristaltic movements [9–14]. In this paper, we mount a surgical workspace on a 6 degrees of freedom miniaturized Stewart platform, allowing the workspace to physically rotate and translate ("SPRK: A Low-Cost Stewart Platform For Motion Study In Surgical Robotics" [15]).

Our results suggest that an intermittent synchronization policy, which targets robot motion around the extrema of the rhythmic motion, is much more robust than a full synchronization policy to kinematic uncertainty and control latency. This is because the minima and maxima of the rhythmic motions correspond to time points with the lowest velocities – thereby small errors in control have a lesser effect than at other times. However, this approach might result in a slower overall execution time. We used the Stewart platform in an initial pilot study, where an expert cardiac surgeon, W. Doug Boyd, performed a cutting task under rhythmic sinusoidal movement at 0.5 Hz. We hypothesized that he would attempt to *fully track* the movement of the platform, i.e., mentally model the motion and compensate for it in real time. However, he preferred an *intermittent synchronization* policy, where he synchronized his actions with the extrema of the rhythmic motion. This observation was counter-intuitive as intermittent synchronization is less efficient in terms of time than fully tracking the movement of the platform.

*Denotes equal contribution. Authors are affiliated with:
[1]The AUTOLAB at UC Berkeley; autolab.berkeley.edu
[2]UC Davis Health; ucdmc.ucdavis.edu

978-1-5386-2512-5/18/$31.00 ©2018 IEEE

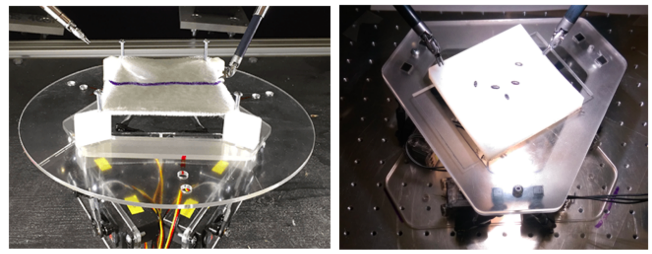

Fig. 1: Comparing full and intermittent synchronization for two tasks: (1) surgical cutting and (2) surgical debridement on a miniaturized Stewart platform performing rhythmic motions at up to 0.5 Hz with mounted gauze for cutting and black rice seeds on a silicone phantom for debridement.

In this paper, we explore the differences between full and intermittent synchronization in the context of *autonomous* execution on two tasks: (1) surgical cutting, and (2) surgical debridement (Figure 1). (1) We consider cutting along a line while the platform translates 10 mm perpendicular to the line sinusoidally at 0.2 Hz. In our experiments, we constructed a simplified variant of the Fundamentals of Laparoscopic Surgery (FLS) cutting task, where we autonomously cut along a line and translated the platform perpendicular to the line at 0.2 Hz. The robot had to observe the movement using computer vision, estimate the frequency and phase, and execute a cut along the line. In our experiments, we found that the intermittent synchronization approach, while 1.8× slower, was significantly more robust. The maximum cutting error (maximum deviation from the desired trajectory) was reduced by 2.6×. (2) We consider surgical debridement where foreign bodies are removed from a tissue phantom that is translating 12.5 mm at 0.5 Hz. The robot had to observe the movement, estimate the frequency and phase, and grasp/remove 10 foreign bodies. A baseline approach achieved a 62% success rate for each removal, while intermittent synchronization was 80% successful.

II. Related Work

We use the da Vinci Research Kit (dVRK) [16, 17] as our RSA, which is a research platform based on Intuitive Surgical's da Vinci surgical system [18] and which has been frequently used in surgical robotics research [3, 4, 7, 19, 20].

A. Robot Assisted Surgery with Periodic Motion

Periodic motion has been studied in robotic surgery, including motion estimation of anatomical components [21–25] and robotic control for motion compensation [26–29]. Much of this work considers virtual surgical simulators, e.g., Duindam and Sastry [27], and proposes a full synchronization approach where the quasi-periodic motion of the anatomy is tracked. Other works, such as Moustris et al. [29], fully synchronize human input on real robot systems with stabilized virtual images, or passively compensate for motion using mounted devices, such as HeartLander [28]. Our work considers physical experiments with a physical simulator of periodic motion, which introduces additional challenges of state estimation from imperfect visual signals and control latency. We notice that in this setting an intermittent synchronization strategy is more reliable for autonomous execution – and is consistent with our observations of experts.

B. Experimental Automated Subtasks

For both cutting and debridement, we adapted state-of-the-art results for stationary settings to a new task setting with rythmic motion. Surgical robotic cutting has been studied in robotics [6, 7, 13, 30, 31] as well as in computer graphics and computational geometry [32, 33] with little to no consideration and compensation of periodic motion. We constructed a simplified variant of the FLS cutting task, where we autonomously cut along a line and translated the platform perpendicular to the line at 0.2 Hz. For this task, without movement, our baseline achieves state-of-the-art cutting accuracy with 100% success in 40 trials. This work seeks to extend prior work to a dynamic setting with periodic motion. Surgical debridement [34–36] is the process of removing dead or damaged tissue to allow the remaining part to heal. Automated surgical debridement was demonstrated in [7, 37–39] with a static target. As in cutting, our baseline debridement controller with a static target achieves an 85% success rate – competitive with the state-of-the-art – and we extend this controller to the dynamic setting.

III. Synchronization Modes

Next, we formalize the algorithmic problem and present three control strategies.

A. Problem Statement

We assume a periodic, 1-D, sinusoidal, quasi-static movement of the platform with negligible higher-order dynamical effects due to motion. The sampling frequency for tracking is assumed to be always greater than the Nyquist frequency f_N. For the full and intermittent synchronization modes, the tracking algorithm uses a colored circle marker on the platform in order to track its motion. Let $\mathcal{X} = [g_1, ..., g_k]$ denote a sequence of target positions with each $g_i \in SE(3)$ describing a point in a global inertial reference frame. Each target point is modulated by a discrete-time rhythmic motion $m[t] \in SE(3)$ that rotates and translates all of the points with respect to the global frame:

$$g_i^{(m)}[t] = m[t] \times g_i$$

$T = [T_x, T_y, T_z] \in SE(3)$ represents a vector of the axial periods of the motion, defined by $\min \|T\|_1$ where for each axis:

$$\exists T \in \mathbb{N} : \forall t \;\; m[t+T] = m[t],$$

and the amplitude is defined as the maximum translational motion due to $m[t]$:

$$\alpha = \max_t \sqrt{m[t].T_x^2 + m[t].T_y^2 + m[t].T_z^2}$$

We assume the robot has a sequence of positional targets and takes decisions of the form (u, τ) where $u \in SE(3)$ represents positional targets in the global coordinate frame and τ is the time at which the movement is completed. The robot takes k such decisions $\mathcal{U} = [(u_1, \tau_1), ..., (u_k, \tau_k)]$ corresponding to attempted movements to the target points. The error is defined by the cumulative error over all decisions:

$$\varepsilon = \sum_{i=1}^{k} \|g_i^{(m)}[\tau] - u_i\|$$

Given $\mathcal{X}$ and a position controller, we consider designing a control policy to generate $\mathcal{U}$ when $m[t]$ is unknown.

B. Three Control Modes

Now, we describe three approaches to synthesizing $\mathcal{U}$.

1) No Synchronization (Open Loop): The baseline approach is an open-loop control strategy that ignores motion from the Stewart platform. This is a decision sequence $\mathcal{U}$ that is independent of time:

$$\mathcal{U} = [(g_1, \cdot), ..., (g_k, \cdot)]$$

We call this approach no synchronization as it ignores the time at which the the command will terminate.

2) Full Synchronization (Full Tracking): With this method, we consider translating or rotating the cutting arm to compensate for the motion. The natural first choice of a controller is a tracking controller – one that models the motion of the platform and tries to exactly compensate for it. We consider motion estimation, a two-step process where the system first learns a motion model and then uses that motion model to predict a translational offset.

First, we track a fixed point r on the platform for an observation period $T_{obs} \gg T$ and a tracking frequency faster than $f_N = \frac{1}{2T}$. This involves collecting tuples (r', r), where $r' = m[t] \times r$. We perform this tracking with a standard computer vision algorithm and the system tracks the motion of the Stewart platform at 15 fps for 1 minute. The algorithm segments a colored circle marker with standard OpenCV circle detection on the platform with an endoscopic stereo camera and registers the stereo frame to a pose in a fixed global coordinate frame.

We convert the positions and orientations into x, y, z and Euler angles y, p, y that represent the $SE(3)$ pose of the platform. Then, in each dimension (independently), we fit a 1-D sine curve:

$$C[t] = \alpha \cdot \sin(\omega(t + \phi))$$

At any point in time, we correct the robot's commanded motion in each dimension with these offsets:

$$u_i[\tau] = g_i - C[\tau].$$

This is a 1-D approximation to the problem, as modeling rotations and translations in the full $SE(3)$ space can be very challenging. The tracking controller is an efficient solution and has been widely applied in prior work [26–29]. However, it is quite hard to implement it in a physical setup. This requires precise estimation of ω, ϕ from data. Furthermore, it assumes that commands can be sent with predictable latency to the robot.

We find that in 10 experimental trials, the estimated frequency has a relative root-mean-square error (RMSE) of 3% in estimating the period due to kinematic variation in the platform and noise in the images. Furthermore, the phase has an RMSE of 0.22 seconds due to the frame rate of the camera and noise in the images. If we assume that these errors compound at the worst possible time in the trajectory, namely, when $\frac{dC}{dt}$ is at a maximum, then this noise alone could result in ≈ 1 mm of RMSE because the errors compound for t when $\underset{t}{\operatorname{argmax}}|\frac{dC}{dt}|$. Furthermore, we found that the total execution time varied by 0.576 seconds in latency on any particular cut. This makes it challenging to time the DVRK to a rhythmic motion. If this error occurred when $\frac{dC}{dt}$ is maximized, this would result in an additional > 3 mm of error.

3) Intermittent Synchronization: As before, this is a two-step process where first we learn a motion model by fitting a sinusoid to tracking data. Then, instead of correcting for the motion at any point of time, as was the procedure of the full synchronization controller, we synchronize the period of the robot with the motion of the platform. That is, we move the robot only so that it reaches the target position around the estimated minima or maxima of the sinusoid. As in control mode 2), we fit a sinusoid for each of the 6 degrees of freedom. Then, we identify the "dominant" degree of freedom, i.e., the one with the largest amplitude α. For this degree of freedom, we identify the first maximum or minimum and generate a sequence of times $[s_1, \ldots, s_k]$, based on the inferred frequency at which that optimum will occur in the future. We synchronize our positional commands with these times:

$$u_i[s_i] = g_i - C[s_i].$$

While this controller is sub-optimal in time, we hypothesize that it is more robust to estimation and control errors. For a sinusoid, the change in position is smallest around the optima $\frac{dC}{dt} \approx 0$. In this window, the uncertainty affects the motion the least. This basic idea applies more generally for any smooth continuous disturbance function. For any function, $\frac{dC}{dt} \approx 0$ around the optima. The second derivative around the optima ($\frac{d^2C}{d^2t}$) can be used to estimate the window width.

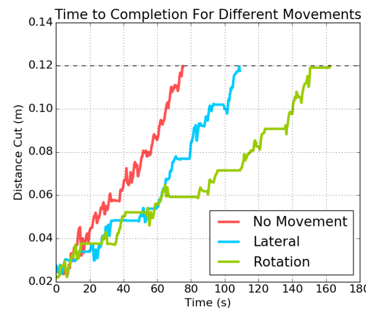

Fig. 2: Expert human surgeon cutting data from one trial period of experiment 1. Results show the distance cut over time along the gauze (12 cm) for the platform with (1a) no motion, (1b) lateral movement, (1c) rotational movement.

IV. Results: Expert Surgeon Demonstrations

Observations of a real surgeon motivated the design of the intermittent synchronization controller. In June 2016, co-author and expert cardiac surgeon Dr. W. Douglas Boyd performed an FLS task (pattern cutting) on the Stewart platform. We recorded trajectory data from these tasks at 90 measurements per second. This data consisted of the end-effector pose of each of the two DVRK arms and the corresponding video from the endoscope.

A. Experiment and Analysis

In experiment 1, Dr. Boyd cut along a 12 cm line under three types of motion: (1a) no movement, (1b) 0.51 cm, 0.5 Hz sinusoidal lateral movement, and (1c) 10^o, 0.5 Hz rotational movement. We hypothesized that the surgeon would apply some form of predictive control and try to compensate for the motion by anticipating where the platform would be. After running the experiment, we noticed that Dr. Boyd did not track the platform in his control strategy. Instead, he timed the period of the platform's movement, cutting at the lowest velocity times, i.e., minima and maxima. This behavior is evident in trajectories with both the rotational and translational motions but is most pronounced in the rotational motions (Figure 2). The plateaus seen in the data correspond to the extrema. There was no appreciable reduction in cutting accuracy when Dr. Boyd cut on the moving platform.

Figure 2 plots the distance cut along the line as a function of time. Dr. Boyd completed the baseline task (1a) without any movement in 67 seconds. With lateral movement (1b) the task was completed in 139 seconds. For the rotational movement (1c) the task was completed in 164 seconds. Waiting for these low-velocity intervals resulted in nearly a $3\times$ and more than a $2\times$ increase in task completion time with the rotational and translational movements, respectively.

V. Experiments: Cutting

We explore the differences between full and intermittent synchronization in the context of **autonomous** execution on two tasks: (1) surgical cutting, and (2) surgical debridement. First, we overview our results on (1).

TABLE I: 50 mm line gauze cutting experiments with 25 mm lateral movements at 0.2 Hz. We found that the intermittent synchronization approach, while 1.8× slower, was significantly more robust reducing the max cutting error by 2.6× and was successful in all four trials.

	Trial 1			Trial 2			Trial 3			Trial 4		
Controller	Finish	Err (mm)	Time	Finish	Err (mm)	Time	Finish	Err (mm)	Time	Finish	Err (mm)	Time
No Sync	No	N/A	104.51	No	N/A	100.78	No	N/A	92.42	No	N/A	91.54
Full Sync	Yes	5.72	**103.35**	No	N/A	102.21	No	N/A	97.56	Yes	2.53	**96.45**
Int. Sync	Yes	**2.70**	206.71	Yes	**1.52**	181.44	Yes	**1.76**	163.35	Yes	**1.12**	169.39

A. Single-Axis Motion

We illustrate 12 trials of cutting under 25 mm, 0.2 Hz sinusoidal motion in one degree of freedom (i.e., single axis). In the first set of experiments, we evaluated the three control modes for accuracy and reliability. For each controller, we ran 4 trials of cutting a 50 mm straight line (2 mm thick) in standard surgical gauze. We took a conservative approach for determining failure of the control policy. If, during the cutting process, the scissors disengaged from the gauze (either above or below), then the trial was marked as a failure. We measured the maximum error in cutting as the maximum displacement outside of the 2 mm line.

The no synchronization approach failed in all 4 trials (Table I). The full synchronization approach succeeded 2 out of the 4 times, but incurred a relatively high cutting error. The intermittent synchronization approach, while 1.8× slower than full synchronization, was successful in all four trials. It was also significantly more robust as it reduced the maximum cutting error by 2.6×.

B. Single-Axis: Increased Frequency

Next, we performed intermittent frequency cutting experiments over a range of frequencies in platform movement to test for control algorithm uncertainty. We defined the sinusoidal amplitude to be 5 mm and varied the frequency from 0 Hz to 0.3 Hz, which was the highest frequency we could cut at before failure, and measured the error. We found that for frequencies less than 0.3 Hz the error was relatively low (Figure 3, where each data point represents a single trial). As the frequencies got higher, it became harder to compensate for the motion, and the variability in control latency made it challenging to exactly time even the intermittent synchronizations. For example, the error jumps by more than 2× between 0.25 Hz and 0.3 Hz (most trials fail for 0.35 Hz motion).

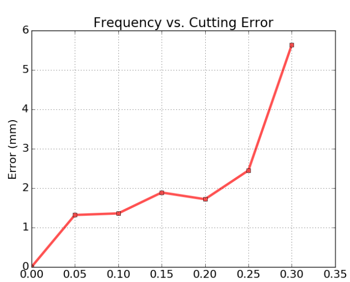

Fig. 3: For single-axis motion, the errors are relatively low until a frequency of 0.3 Hz.

C. Single-Axis: Alternate Materials

In the next set of experiments, we considered the same scenario but with alternative cutting materials to understand how material properties could affect the control technique. We considered a silicone tissue phantom to model skin and a nylon sheet to model tougher connective tissue.

Over the 10 trials run on the silicone phantoms, 7 trials had an error of 3 mm or less (Table II). It should be noted that one trial produced an error during cutting so the trial terminated without completion.

We also repeated the same experiment on 10 nylon sheets (Table II). The texture of nylon consists of parallel grains; the cutting trajectories were drawn to go with the grain. All of the trials were successful and 8 out of 10 trials had an error of 3 mm or less. We attribute the larger cutting deviation errors for both sets of trials to frequency estimation errors.

Overall, we found that the material properties (gauze vs. silicone vs. nylon) did not affect the error as much as the estimation errors. In future work, we hope to explore more sophisticated techniques to register the workspace and estimate its movement.

TABLE II: Results from 10 silicone tissue phantom and 10 nylon sheet cutting tests. Out of the 20 cutting trials for the two materials, only one trial terminated without completion.

Silicone Tissue Phantom			**Nylon Sheet**		
Finish	Trial Err (mm)	Estim Freq (Hz)	Finish	Trial Err (mm)	Estim Freq (Hz)
Yes	3	0.080	Yes	2	0.078
Yes	6	0.075	Yes	1	0.075
Yes	1	0.080	Yes	3	0.075
No	**N/A**	**0.075**	Yes	5	0.080
Yes	2	0.079	Yes	0	0.078
Yes	2	0.079	Yes	4	0.079
Yes	1	0.079	Yes	0	0.077
Yes	4	0.077	Yes	1	0.079
Yes	2	0.050	Yes	2	0.079
Yes	2	0.076	Yes	3	0.078

D. Multi-Axis Motions

We next characterize the performance of intermittent synchronization on multi-axis periodic motions. This means that the periodic motion occurs in six degrees of freedom at the same frequency and phase but with different amplitudes. We compare the following motion modes: No Motion, Motion along X only (which translates orthogonally to the cutting line), 3D Translation, and 6D Translation and Rotation. The motions along X only were sinusoidal with a 2.5 cm amplitude and a frequency of 0.2 Hz. For the 3D and 6D translations, we generated amplitudes in each of the translation dimensions where the total *norm* of the amplitudes was equal to the amplitude we used for motion along X only. Additionally,

TABLE III: Results from 5 trials for each motion mode showing the average deviation from cutting line.

Motion Mode	Error (mm)
No Motion	0.51 +/- 0.84
X Only	1.66 +/- 1.04
3D Translation	3.91 +/- 3.01
6D Trans. and Rot.	4.41 +/- 2.97

for 6D translation and rotation, we generated amplitudes in each of the rotational dimensions where the total *norm* of the amplitudes was 15°. The average deviations of a cut from the 2 mm marked line over 5 trials for each with the 1 standard deviation error are listed in Table III. The results show that performance degrades with periodic motions that rotate and translate in all of SE(3), and that the intermittent synchronization is best suited for rhythmic translations in a single dominant direction.

VI. Experiments: Debridement Task

Next, we consider surgical debridement where foreign bodies are removed from a tissue phantom that is moving with a 12.5 mm amplitude at 0.5 Hz along a single axis. The robot had to observe the movement, estimate the frequency and phase, and grasp/remove 10 inclusions. The inclusions were black rice seeds that were 5 mm along their longest axis and 2 mm along the two other axes. We used surgical grippers with a 7 mm gripper throw. A computer vision system observes the foreign bodies on a silicone phantom and segments the seeds using a standard contour detector. Then, each seed is registered to the global coordinate frame. Chessboard experiments suggest an inherent error of 2.25 mm in the registration system alone. The robot then controls to the center of mass of the seed and then positions the gripper orthogonal to the long axis. Seeds were placed in random positions and orientations on the surface.

Unsuccessful grasps are unavoidable due to perceptual mistakes and registration/kinematic uncertainty. The task allows retrials and we give the robot a maximum of 20 attempted grasps to clear 10 inclusions. We compared using intermittent synchronization to a baseline approach of no synchronization on 10 trials over 10 inclusions (Table IV). The baseline approach achieves a 62% success rate for each removal, while intermittent synchronization achieves 80% success. The intermittent synchronization runs at an average speed of 7.2 grasps a minute while the baseline runs at 10.1 grasps a minute.

TABLE IV: Results from 10 debridement trials with no synchronization (baseline) and with intermittent synchronization each.

No Synchronization			**Intermittent Synchronization**		
Finish	Attempts	Successes	Finish	Attempts	Successes
No	20	9	Yes	16	10
Yes	16	10	Yes	13	10
Yes	10	10	Yes	10	10
Yes	15	10	Yes	10	10
No	20	8	Yes	13	10
Yes	19	10	Yes	10	10
Yes	14	10	Yes	11	10
Yes	13	10	Yes	10	10
Yes	13	10	Yes	10	10
Yes	12	10	No	20	9

VII. Conclusions

In this paper, we explored 3 control modes for automated surgical subtasks in a dynamic environment using a Stewart platform. We analyzed the results of an expert cardiac surgeon cutting gauze on the Stewart platform and extrapolated an intermittent synchronization control strategy, which favored cutting along a trajectory at windows of low velocity. In our experiments, we found that the intermittent synchronization approach, while slower, was significantly more robust.

Acknowledgment

This research was performed at the AUTOLAB at UC Berkeley in affiliation with the Real-Time Intelligent Secure Execution (RISE) Lab, Berkeley AI Research (BAIR) Lab, and the CITRIS "People and Robots" (CPAR) Initiative: http://robotics.citris-uc.org in affiliation with UC Berkeley's Center for Automation and Learning for Medical Robotics (Cal-MR). This research is supported in part by DHS Award HSHQDC-16-3-00083, NSF CISE Expeditions Award CCF-1139158, by the Scalable Collaborative Human-Robot Learning (SCHooL) Project NSF National Robotics Initiative Award 1734633, and donations from Alibaba, Amazon, Ant Financial, CapitalOne, Ericsson, GE, Google, Huawei, Intel, IBM, Microsoft, Scotiabank, VMware, Siemens, Cisco, Autodesk, Toyota Research, Samsung, Knapp, and Loccioni Inc. We also acknowledge a major equipment grant from Intuitive Surgical and by generous donations from Andy Chou and Susan and Deepak Lim.

References

[1] 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 (cit. on p. 1).

[2] T. Osa, N. Sugita, and M. Mamoru, "Online Trajectory Planning in Dynamic Environments for Surgical Task Automation", in *Robotics: Science and Systems (RSS)*, 2014 (cit. on p. 1).

[3] 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 *IEEE International Conference on Robotics and Automation (ICRA)*, 2010 (cit. on pp. 1, 2).

[4] S. Sen, A. Garg, D. V. Gealy, S. McKinley, Y. Jen, and K. Goldberg, "Automating Multiple-Throw Multilateral Surgical Suturing with a Mechanical Needle Guide and Sequential Convex Optimization", in *IEEE International Conference on Robotics and Automation (ICRA)*, 2016 (cit. on pp. 1, 2).

[5] 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 International Conference on Intelligent Robots and Systems (IROS)*, 2013, pp. 4111–4117 (cit. on p. 1).

[6] S. Krishnan, A. Garg, R. Liaw, B. Thananjeyan, L. Miller, F. T. Pokorny, and K. Goldberg, "SWIRL: A sequential windowed inverse reinforcement learning algorithm for robot tasks with delayed rewards.", in *Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2016 (cit. on pp. 1, 2).

[7] A. Murali, S. Sen, B. Kehoe, A. Garg, S. McFarland, S. Patil, W. D. Boyd, S. Lim, P. Abbeel, and K. Goldberg, "Learning by Observation for Surgical Subtasks: Multilateral Cutting of 3d Viscoelastic and 2d Orthotropic Tissue Phantoms", in *IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 1202–1209 (cit. on pp. 1, 2).

[8] A. Garg, S. Sen, R. Kapadia, Y. Jen, S. McKinley, L. Miller, and K. Goldberg, "Tumor localization using automated palpation with gaussian process adaptive sampling", in *IEEE Conference on Automation Science and Engineering (CASE)*, 2016 (cit. on p. 1).

[9] D. Bielser, P. Glardon, M. Teschner, and M. Gross, "A state machine for real-time cutting of tetrahedral meshes", *Graphical Models*, vol. 66, no. 6, pp. 398–417, 2004 (cit. on p. 1).

[10] D. Steinemann, M. Harders, M. Gross, and G. Szekely, "Hybrid cutting of deformable solids", in *Virtual Reality Conference, 2006*, IEEE, 2006, pp. 35–42 (cit. on p. 1).

[11] C. Mendoza and C. Laugier, "Simulating soft tissue cutting using finite element models", in *IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, vol. 1, 2003, pp. 1109–1114 (cit. on p. 1).

[12] E. Sifakis, K. G. Der, and R. Fedkiw, "Arbitrary cutting of deformable tetrahedralized objects", in *Proceedings of the 2007 ACM SIGGRAPH/Eurographics symposium on Computer animation*, Eurographics Association, 2007, pp. 73–80 (cit. on p. 1).

[13] H.-W. Nienhuys and A. F. van der Stappen, "A Surgery Simulation supporting Cuts and Finite Element Deformation", in *International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2001 (cit. on pp. 1, 2).

[14] A. Moglia, V. Ferrari, L. Morelli, M. Ferrari, F. Mosca, and A. Cuschieri, "A systematic review of virtual reality simulators for robot-assisted surgery", *European urology*, vol. 69, no. 6, pp. 1065–1080, 2016 (cit. on p. 1).

[15] V. Patel*, S. Krishnan*, A. Goncalves, and K. Goldberg, "SPRK: A low-cost stewart platform for motion study in surgical robotics", in *International Symposium on Medical Robotics (ISMR)*, In Press, 2018 (cit. on p. 1).

[16] Intuitive Surgical, *da Vinci Research Kit*, 2014 (cit. on p. 2).

[17] 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 *IEEE International Conference on Robotics and Automation (ICRA)*, 2014 (cit. on p. 2).

[18] G. Ballantyne and F. Moll, "The da Vinci Telerobotic Surgical System: The Virtual Operative Field and Telepresence Surgery", *Surgical Clinics of North America*, 2003 (cit. on p. 2).

[19] S. Krishnan, R. Fox, and K. Goldberg, "DDCO: Discovery of Deep Continuous Options for Robot Learning from Demonstrations", in *Conference on Robot Learning (CoRL)*, 2017 (cit. on p. 2).

[20] B. Thananjeyan, A. Garg, S. Krishnan, C. Chen, L. Miller, and K. Goldberg, "Multilateral surgical pattern cutting in 2d orthotropic gauze with deep reinforcement learning policies for tensioning", in *IEEE International Conference on Robotics and Automation (ICRA)*, 2017 (cit. on p. 2).

[21] T. Ortmaier, M. Groger, D. H. Boehm, V. Falk, and G. Hirzinger, "Motion estimation in beating heart surgery", *IEEE Transactions on Biomedical Engineering*, vol. 52, no. 10, pp. 1729–1740, 2005 (cit. on p. 2).

[22] T. J. Franke, O. Bebek, and M. C. Cavusoglu, "Improved prediction of heart motion using an adaptive filter for robot assisted beating heart surgery", in *Intelligent Robots and Systems, 2007. IROS 2007. IEEE/RSJ International Conference on*, IEEE, 2007, pp. 509–515 (cit. on p. 2).

[23] T. J. Ortmaier, "Motion compensation in minimally invasive robotic surgery", PhD thesis, Universität München, 2003 (cit. on p. 2).

[24] R. Richa, A. P. Bó, and P. Poignet, "Beating heart motion prediction for robust visual tracking", in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, IEEE, 2010, pp. 4579–4584 (cit. on p. 2).

[25] C. Sloth, R. Wisniewski, J. Larsen, J. Leth, and J. Poulsen, "Model of beating-heart and surgical robot", Technical report, Aalborg University, Aalborg, Tech. Rep., 2012 (cit. on p. 2).

[26] M. C. Cavusoglu, J. Rotella, W. S. Newman, S. Choi, J. Ustin, and S. S. Sastry, "Control algorithms for active relative motion cancelling for robotic assisted off-pump coronary artery bypass graft surgery", in *Advanced Robotics, 2005. ICAR'05. Proceedings., 12th International Conference on*, IEEE, 2005, pp. 431–436 (cit. on pp. 2, 3).

[27] V. Duindam and S. Sastry, "Geometric motion estimation and control for robotic-assisted beating-heart surgery", in *Intelligent Robots and Systems, 2007. IROS 2007. IEEE/RSJ International Conference on*, IEEE, 2007, pp. 871–876 (cit. on pp. 2, 3).

[28] C. N. Riviere, J. Gangloff, and M. De Mathelin, "Robotic compensation of biological motion to enhance surgical accuracy", *Proceedings of the IEEE*, vol. 94, no. 9, pp. 1705–1716, 2006 (cit. on pp. 2, 3).

[29] G. Moustris, A. Mantelos, and C. Tzafestas, "Active motion compensation in robotic cardiac surgery", in *Control Conference (ECC), 2013 European*, IEEE, 2013, pp. 2311–2316 (cit. on pp. 2, 3).

[30] B. Thananjeyan, A. Garg, S. Krishnan, C. Chen, L. Miller, and K. Goldberg, "Multilateral surgical pattern cutting in 2d orthotropic gauze with deep reinforcement learning policies for tensioning", in *IEEE International Conference on Robotics and Automation (ICRA)*, 2017 (cit. on p. 2).

[31] M. Mahvash, L. Voo, D. Kim, K. Jeung, J. Wainer, and A. M. Okamura, "Modeling the forces of cutting with scissors", *IEEE transactions on bio-medical engineering*, vol. 55.3, pp. 848–856, 2008 (cit. on p. 2).

[32] H. Zhang, S. Payandeh, and J. Dill, "On cutting and dissection of virtual deformable objects", in *IEEE International Conference on Robotics and Automation (ICRA)*, 2004 (cit. on p. 2).

[33] 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", *ACM Transactions on Graphics*, vol. 28, no. 3, 2009 (cit. on p. 2).

[34] 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 (cit. on p. 2).

[35] 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 (cit. on p. 2).

[36] K. A. Nichols, A. Murali, S. Sen, K. Goldberg, and A. M. Okamura, "Models of Human-Centered Automation in a Debridement Task", in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015 (cit. on p. 2).

[37] 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 *IEEE International Conference on Robotics and Automation (ICRA)*, 2014 (cit. on p. 2).

[38] J. Mahler, S. Krishnan, M. Laskey, S. Sen, A. Murali, B. Kehoe, S. Patil, J. Wang, M. Franklin, P. Abbeel, and K. Goldberg, "Learning Accurate Kinematic Control of Cable-Driven Surgical Robots Using Data Cleaning and Gaussian Process Regression.", in *IEEE Conference on Automation Science and Engineering (CASE)*, 2014 (cit. on p. 2).

[39] D. Seita, S. Krishnan, R. Fox, S. McKinley, J. F. Canny, and K. Goldberg, "Fast and Reliable Autonomous Surgical Debridement with Cable-Driven Robots Using a Two-Phase Calibration Procedure", in *IEEE International Conference on Robotics and Automation (ICRA)*, In Press, 2018 (cit. on p. 2).