Learning a Library of Surgical Manipulation Skills for Robotic Surgery

Anonymous Author(s)
Affiliation
Address
email

Abstract: In this abstract, we present preliminary progress towards learning a library of surgical manipulation skills using the da Vinci Research Kit (dVRK) system. Such a library of skills may include picking-up various surgical tools, driving a needle into tissue, tying a suture knot, or cutting tissue, among many others. These skills may be composed in various ways to accomplish a long-horizon surgical task such as suturing, and each skill may be resused in other downstream applications in various novel settings. Our goal is to learn these skills using one of the most effective tools for robot learning, deep imitation learning. Initial results on two tasks are demonstrated: 1) picking-up a suture needle and 2) picking-up a suture needle and handing the needle over to the opposing arm. Promising initial results show that learning such diverse skills may be possible in simple settings. However, further innovation is necessary in adapting these systems to work in truly novel scenarios or even learn to recover from simple mistakes. Towards this end, a few points of improvement and promising future directions are discussed.

Keywords: CoRL, Robots, Learning

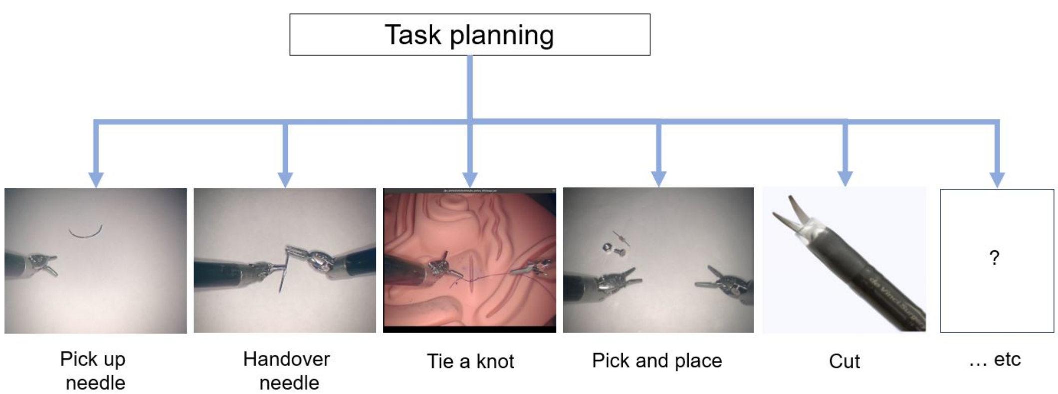

Figure 1: Learning a library of surgical skills using deep imitation learning

1 Introduction

In this abstract, we are concerned with the task of learning contact-rich surgical manipulation skills on the da Vinci Research Kit (dVRK). The goal is to learn a library of surgical skills, such as picking-up various surgical tools, tying a knot, grasping or cutting tissue, and many others, such that they may be composed together to solve a long-horizon surgical task. Towards this end, we demonstrate preliminary results on learning some of these tasks using deep imitation learning.

At a high-level, our approach uses transformer architecture which takes as input stereo images and generates a sequence of bimanual robot actions. At the current stage, we are concerned with solving the problem at hand as efficiently as possible using as little as 60 demonstrations. We demonstrate preliminary results on the two following tasks: 1) picking-up a suture needle, and 2) picking-up a suture needle and handing it over to the other gripper arm, as illustrated in Fig. 1.

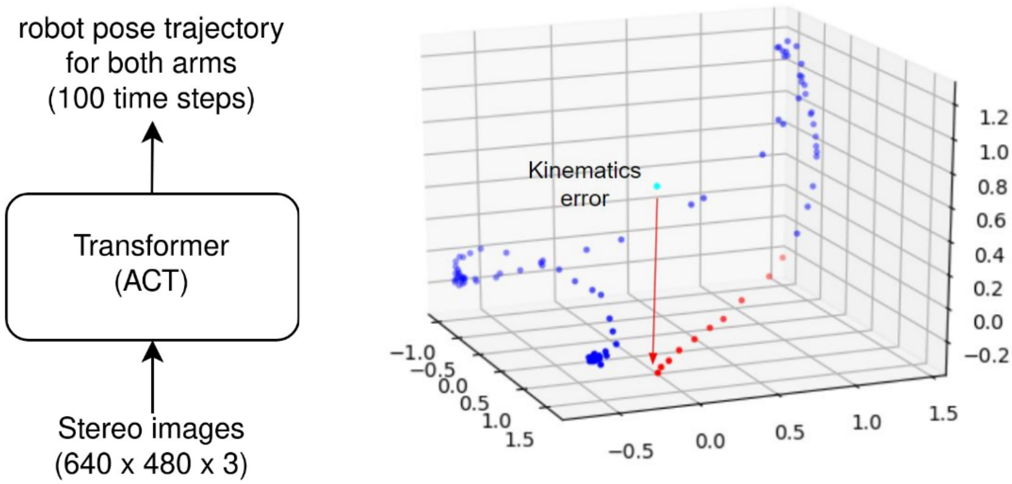

Figure 2: (Left) transformer network architecture used is shown. The inputs are stereo images and the outputs are robot pose trajectories for 100 timesteps (3s prediction horizon) for both gripper arms. (Right) illustration of kinematics errors due to the inherent design limtiation of the dVRK system. The cyan dot is the current robot position, and the other dots are the robot trajectory generated by the network. The red end of the trajectory indicates the starting point of the trajectory. To compensate for this gap, the entire predicted trajectory is simply translated to originate from the current robot position or the cyan dot.

2 Method

As shown in Fig. 2, we used a transformer network adopted from [1]. We chose this implementation as it demonstrated strong results in solving contact-rich manipulation tasks in various challenging settings, though on a different robot hardware. A few changes were made to tailor the architecture for the dVRK system. The input dimension was changed such that it accepted stereo images. The action space was changed to generate gripper position, orientation (unit quaternion), and jaw angle for both arms; the original implementation used joint values of the robot arms instead.

One additional challenge we faced was dealing with the kinematics error of the dVRK system. Since the dVRK system is a cable-driven robot and the arms themselves are not perfectly rigid, the robot kinematics readings may be inconsistent. This led to situations where the robot trajectory generated by the network did not originate from the current gripper position but instead originated at an offset location, as shown in Fig. 2. To compensate for this error, the current robot kinematics position was provided as an additional input to the network; however, this did not make much difference. Instead, we simply translated the predicted trajectory to originate from the current end-effector position, and this was sufficient to make the system work.

The experimental setup consists of the dVRK system with its stereo imaging system with an integrated light source, and the grippers are large needle driver and DeBakey forceps, as shown in Fig 3. The workspace consisted of a small box with a white piece of paper covering the top, to provide a simple homogeneous background. To constraint the workspace, the gripper arms and the needle position was constrained to originate within their respective boxes as shown in Fig. 3. Note that for the task of picking-up a suture needle, which only requires a single arm, the output space of the network was resized to accomodate a single arm.

For both tasks, 60 demonstration trajectories were collected by a human. All demonstration data consisted of successful trials and no failures. 80% of the data was used for training and the rest for testing. The rest of the training details follow the original paper.

Figure 3: (Left) close up view of the dVRK arms, camera, and the integrated light source, (middle) experimental setup, with a box as a platform and piece of paper as a background, (right) workspace view from the camera perspective; blue squares indicate the general starting locations of the grippers and the suture needle

3 Results and Discussions

We present preliminary results for the aforementioned tasks. For the task of picking-up a suture needle using a single arm, 3 out of 6 trials were successful. For the task of picking-up a suture needle and handing it over to the other gripper arm, 1 trial out of 4 trials was successful. An example of a successful trials for the latter task is shown in Fig 4.

While the results are encouraging, note that this performance is not fully reflective of what is expected as there is much room for improvement on the robot implementation side. First, due to implementation constraints, the network was inferenced 2-3 times throughout the entire task. Ideally, the network should be running in a receding-horizon manner such that it may perform closed-loop control and recover from any potential errors observed during execution. Also, the camera configuration during training and testing were not identical; the dVRK system is a shared system across various labs and given the system design, there was no reliable way to replicate the setup to be identical to the training setting. Although the difference may be slight, such change in camera viewpoint may have also contributed to the diminished success rate for the system.

There is significant room for future work. First and foremost is increasing the inference rate so that the robot may perform closed-loop control in real-time. To compensate for shifted views due to change in camera configuration, appropriate data augmnetation should be employed. Another important problem is improving the cognitive capabilities of the system. For instance, during failure trials, it was often observed that even though the needle was not successfully picked-up, the robot continued blindly with the remaining motion, rather than reattempting to pick-up the needle. A simple strategy to improve may be to include examples where the robot learns to recover from such mistakes in the training dataset. However, if another novel failure is encountered, the robot may continue to make similar errors thus this solution is not scalable. As the workshop abstract suggests, equipping the robots with cognitive capabilities, especially at the levels of conscious and introspective, is essential to improving their abilities in novel erroneous situations. One interesting way to deal with this problem in a scalable manner is to utilize world models for augmenting the training data, which is another direction we are considering to explore.

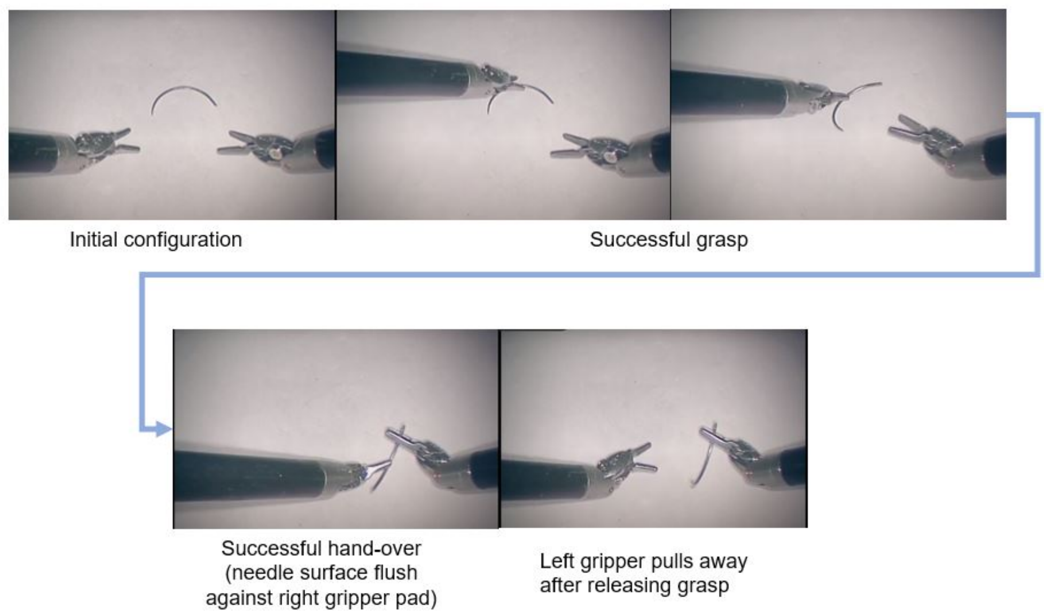

Figure 4: A successful roll-out of the bimanual task is shown. Initially, the left gripper must pick-up the suture needle, bring the needle between the right gripper jaws, close the right gripper jaw, release the left gripper jaw, and then the left gripper must pull away. During the hand-over step, it is important to ensure that the needle surface is flush against the right gripper jaw surface, otherwise a stable grasp may not be established. Also, when the left gripper pulls away, it must release its jaws, otherwise the two arms may fight for the needle which is undesirable as it may overload the robot.

4 Related Work

A classic approach to tackling such bimanual manipulation problem as posed here would require heavy hand-crafted work, such as building a perception system to extract the state of the objects in the scene (and also deciding what the state space should even be), developing an accurate dynamics model to model interactions between various objects, and additionally performing motion planning using the developed models [2], [3]. However, these approaches may not scale easily across various tasks and would require much more significant effort. An alternative approach, as proposed here, is to employ learning from human demonstrations [4] [1]. There also exists a body of work addressing fine manipulation tasks such as handling clothes or untying a knot, to list a few [5], [6]. In the surgical community, there has been some effort towards autonomy such as [7], which demonstrated an impressive feat of performing suturing tasks on a live pig. However, the suturing task was simplified to a reach task by employing an automatic suture device. In general, contact-rich manipulation in the surgical community, and also explicitly in the dVRK community, is scarce.

5 Conclusion

In this work, we demonstrated preliminary steps towards developing a library of surgical skills for the dVRK surgical system. We demonstrated reasonable success rates on a relatively simple task of picking-up a suture needle and a more challenging task of picking-up the suture needle and handing it over to the other gripper arm. In addition several points of improvement as suggested in the discussion section, we will continue to push towards evaluating the network on various other skills such as tying a knot, cutting tissue, and other challenging manipulation tasks, towards the final goal of learning a library of surgical skills.

References

[1] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn. Learning fine-grained bimanual manipulation with low-cost hardware, 2023.

[2] C. Smith, Y. Karayiannidis, L. Nalpantidis, X. Gratal, P. Qi, D. V. Dimarogonas, and D. Kragic. Dual arm manipulation—a survey. *Robotics and Autonomous Systems*, 60(10):1340–1353, 2012. ISSN 0921-8890. doi:https://doi.org/10.1016/j.robot.2012.07.005. URL https://www.sciencedirect.com/science/article/pii/S092188901200108X.

[3] S. S. M. Salehian, N. Figueroa, and A. Billard. A unified framework for coordinated multi-arm motion planning. *The International Journal of Robotics Research*, 37(10):1205–1232, 2018. doi:10.1177/0278364918765952. URL https://doi.org/10.1177/0278364918765952.

[4] S. S. M. Salehian, N. Figueroa, and A. Billard. A unified framework for coordinated multi-arm motion planning. *The International Journal of Robotics Research*, 37(10):1205–1232, 2018. doi:10.1177/0278364918765952. URL https://doi.org/10.1177/0278364918765952.

[5] A. Ganapathi, P. Sundaresan, B. Thananjeyan, A. Balakrishna, D. Seita, J. Grannen, M. Hwang, R. Hoque, J. E. Gonzalez, N. Jamali, K. Yamane, S. Iba, and K. Goldberg. Learning dense visual correspondences in simulation to smooth and fold real fabrics, 2020.

[6] J. Grannen, P. Sundaresan, B. Thananjeyan, J. Ichnowski, A. Balakrishna, M. Hwang, V. Viswanath, M. Laskey, J. E. Gonzalez, and K. Goldberg. Untangling dense knots by learning task-relevant keypoints. 2020.

[7] H. Saeidi, J. D. Opfermann, M. Kam, S. Wei, S. Leonard, M. H. Hsieh, J. U. Kang, and A. Krieger. Autonomous robotic laparoscopic surgery for intestinal anastomosis. *Science Robotics*, 7(62):eabj2908, 2022. doi:10.1126/scirobotics.abj2908. URL https://www.science.org/doi/abs/10.1126/scirobotics.abj2908.