

Computer Integrated Surgery

Group 20: A reinforcement learning approach to robotic suturing

Team Members: Walee Attia (wattia1), Jocelyn Hsu (jhsu37), Jihoon Kim (jkim620)

Mentors: Dr. Anqi Liu, Dr. Adnan Munawar, Dr. Manish Sahu, Dr. Peter Kazanzides

February 23, 2022

Introduction

Reinforcement learning (RL) is a machine learning framework involved in creating artificial agents that fulfill various complex problems. Surgical robots have opened the door to surgical task automation, which has piqued the interest of RL research. However, no robust framework exists for RL tasks in surgical robotics environments. We propose an OpenAI Gym environment based on an autonomous robotic suturing simulator with benchmark algorithms to pave the way for future surgical automation.

Background

Reinforcement Learning has made significant progress in the robotics domain, enabled by open-source frameworks such as OpenAI Gym [1], which provides effortless implementation of complex algorithms in both simulation and real robots. RL has specifically seen success in robotic manipulation and grasping, with evidence that learned policies are transferable from simulation to real robots [2]. However, RL's success in robotics hinges on having lightweight and efficient simulation environments as it requires thousands to millions of simulated attempts to evaluate and explore policy options, which is crucial for real-world use due to the impracticality of running millions of attempts on a physical system [3].




Figure 1: Surgical Robotics Challenge Environment [4]

Thankfully, paradigms in the field of surgical robotics exist in the form of simulators. The 2021-2022 AccelNet Surgical Robotics Challenge [4] is a simulation platform featuring two seven degrees-of-freedom (DOF) needle drivers and a controllable based on the da Vinci Surgical System, a suturing phantom, and a needle with suture. The challenge is based on the Asynchronous Multi-Body Framework (AMBF) simulator [5], which is a real-time dynamics simulator for robots. The challenge is separated into three tasks: 1. Finding the pose of the suture needle with respect to the endoscope pose, 2. using the needle driver to grasp the needle and drive the needle through the suture phantom, and 3. use both needle drivers to suture the phantom.

Prior Work

As RL tasks are still emerging, there are limited papers in the intersection of surgical robotics and RL. However, there are two research endeavors that have attempted to bridge this gap: AMBF-RL [5] and dVRL [3].

AMBF-RL is a sandbox toolkit that can be used to design control algorithms for medical robotics in the AMBF simulator. Varier et al. were able to demonstrate the use of RL for debris removal on a simulated dVRK Patient Side Manipulator (PSM) using two state-of-the-art (SOTA) RL policy algorithms. However, AMBF-RL environment only supports kinematic tasks such as reaching and grasping, which are preliminary in nature compared to a dynamical suturing task.




Figure 2: dVRL Reach and Grasp Task [3]

Similarly, the dVRL project is the first open-source RL environment designed specifically for surgical robotics. One feature of dVRL is that their environment framework is functionally equivalent to OpenAI Gym’s API, which makes the dVRL environment more accessible to the public, given its open-sourced nature. Again, the dVRL environment can only support simple reach and grasp policies.

Goals and Relevance

Our goal is to develop an OpenAI Gym compatible interface for the Surgical Robotics Challenge (SRC) environment with efficient, accurate RL algorithms. An OpenAI Gym environment for the SRC will provide wide-spread use to the masses and will drive further innovation in the interdisciplinary fields of surgical robotics and RL. We will also provide baseline RL algorithms for the suturing tasks in the SRC for comparison and evaluation, with a potential NeurIPS 2023 paper submission on the Datasets and Benchmarks track.

Technical Approach

Reinforcement Learning




Figure 3: Agent Environment Loop [1]

Reinforcement learning is the process of training an agent to maximize the reward when completing a task. As an agent reaches closer to its end goal, it is rewarded positively, and if it missteps and deviates from its goal, it is punished with a negative reward. Due to the agent’s ability to manipulate

its environment, it will need to observe following each action to assess the subsequent action that will be taken. Our goal is to complete automated suturing with a reinforcement learning algorithm.

OpenAI Gym Environment

To accomplish Surgical Robotics Challenge #2, we plan to divide the OpenAI Gym environment into three sub-environments: **Grasp**, **Insert**, and **Target**. Each environment will have the `step()`, `reset()`, `render()`, and `close()` methods implemented.

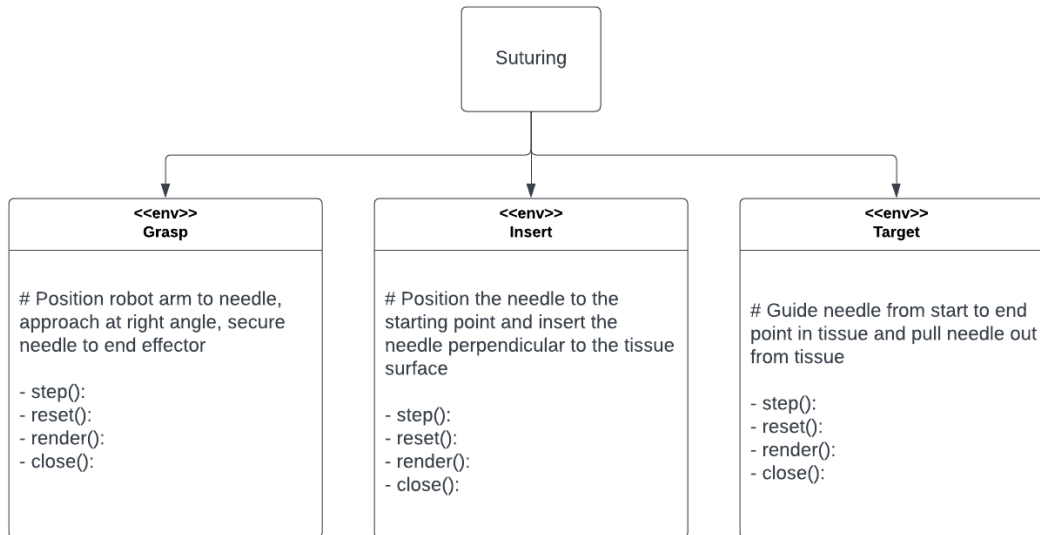


Figure 4: Planned Environment Architecture

Grasp

The **Grasp** environment will be responsible for handling the initial grabbing of the needle from the surface of the surgical table. Given the positions of the needle and the robotic arm, this environment will involve moving the robotic arm to the coordinates of where the needle is located, rotating the joint of the arm for an ideal grasping angle, securing the needle to the end effector of the robotic arm, and picking the needle up.

Insert

The **Insert** environment will move the needle already grasped in the robotic arm to the starting position provided. Once at the right position and angle, the robot will puncture the needle through the surface of the tissue.

Target

The **Target** environment guides the needle embedded in the tissue to the target end point. The needle will then be pulled through the end point to complete the single suture.

Suturing

As a maximum delivery, we hope to be able to combine the **Grasp**, **Insert**, and **Target** environments such that the robot performs **Grasp** once, at the beginning of the suturing procedure, and subsequently alternates between **Insert** and **Target** for a series of sutures.

Method development

For each of the sub-environments, the following methods will be developed:

- **step()**
 - Adjusts the robotic arm position based on the calculated action and returns the observation and reward resulting from the action. The robotic arm position can be backtracked to a previous state if a failure mode is encountered, such as dropping the needle after grasping.
 - Parameters:
 - * **self**: robotic arm position, needle position
 - * **action**: change in robotic arm position
 - Returns
 - * **observation**: robotic arm position, needle position
 - * **reward**: calculated reward score
 - * **terminated**: **True** if task is accomplished, **False** otherwise
 - * **truncated**: **True** if failing to accomplish task within certain number of iterations, **False** otherwise
 - * **info**: additional diagnostic information
- **reset()**
 - Sets all variables of the environment to its initial state.
 - Parameters:
 - * **self**: robotic arm position, needle position
 - * **seed** (optional): randomization seed for replication, if needed
 - Returns
 - * **observation**: initial robotic arm position, initial needle position
 - * **info**: additional diagnostic information
- **render()**
 - Visualize the state of the environment from the agent’s perspective.
 - Parameters:
 - * **self**: robotic arm position, needle position
 - Returns: **None**
- **close()**
 - Completes the simulation.
 - Parameters:
 - * **self**: robotic arm position, needle position
 - Returns:
 - * **Env.unwrapped**: raw state of the environment

Management Plan

Deliverables

The following contains an outline of our minimal, expected, and maximum deliverables.

Minimal: Open-source OpenAI Gym environment to train RL algorithms (expected by March 17, 2023)

- Transfer of SRC to OpenAI Gym environment (sandbox env)
- Core functionality for OpenAI Gym: `make()`, `reset()`, `step()`, `render()`
- Environments built to accomplish SRC Challenge #2: grasp needle and drive through tissue
- Documentation of our work through entire development pipeline
- Compatibility with dVRK in the LCSR lab

Expected: Benchmarked performances of the SOTA RL algorithms (expected by April 28, 2023)

- Literature review of state-of-the-art (SOTA) RL Algorithms
- RL algorithm training for automated suturing task (SRC #2)
- Performance evaluation (accuracy and efficiency) of RL algorithms & 2021-2022 SRC winners

Maximum: Submission as NeurIPS (Datasets and Benchmark track), additional SRC tasks (expected by TBD)

- Paper for NeurIPS conference (database and benchmarking track), deadline TBA [6]
- Additional SRC challenges [4]:
 - Challenge 1: Find the needle
 - Challenge 3: Suture phantom

Technical Dependencies

Below are the technical dependencies needed to complete our project. All dependencies have already been discussed with our mentors and will be provided to us by our date needed.

Dependency	Need	Source	Date Needed	Status	Contingency Plan
Swipe access to Robotarium, LCSR	Environment Development	Dr. Adnan Munawar	3/1/2023	In-progress	N/A
Access to dVRK systems at the Robotarium, LCSR	Test simulation environment	Dr. Adnan Munawar	3/15/2023	In-progress	Linux Virtual Machine w/ AMBF + ROS
Rockfish GPU Access	Benchmarking	Dr. Anqi Liu	4/1/2023	In-progress	Google cloud
SRC Winning Algorithms	Benchmarking	Dr. Adnan Munawar	4/1/2023	In-progress	N/A

Key Dates & Responsibilities

We plan to accomplish our deliverables according to the following timeline.

Milestone	Output	Start Date	Soft Deadline	Hard Deadline
Transfer SRC Environment to OpenAI Gym	GitHub Repo w/ MVP of environment	2/23/23	3/1/2023	3/6/2023
Reward functions: Challenge #2 - Reward functions for needle grasp, insert, target - Test reward functions exhaustively	Tested code	3/6/2023	3/13/2023	3/17/2023
Benchmarking: Challenge #2 - Trained RL algorithm for challenge 2, - Evaluate against SRC 2022 winners	Model Results Evaluation Sheet	3/27/2023	4/21/2023	4/28/2023
Env Dev + Benchmarking: Challenge #1 - CV algo for needle finding, - Evaluate against SRC 2022 winners	Tested code Model Results Evaluation Sheet	5/1/2023	5/10/2023	5/12/2023
Env Dev + Benchmarking: Challenge #3 - Suturing reward functions, - Benchmark RL algo for suturing	Tested code Model Results Evaluation Sheet	5/15/2023	5/26/2023	5/31/2023
NeurIPS - Paper submission - Datasets and Benchmark Track	Tested code Model Results Evaluation Sheet	5/29/2023	TBD	TBD

	Member	February				March				April				May				June	
		1	2	3	4	1	2	3	4	1	2	3	4	1	2	3	4	1	2
Preliminary Research																			
Set up OpenAI Gym, ROS, AMBF	All	█	█	█															
Lit review of SRC, AMBF	All	█	█	█	█														
RL online course	All	█	█	█	█														
Env Dev: Challenge 2																			
Code Documentation	All					█	█	█	█	█									
Transfer SRC env to OpenAI Gym	All				█														
Develop grasp reward function	Jihoon					█	█	█	█										
Develop insert reward function	Walee					█	█	█	█										
Develop target reward function	Jocelyn					█	█	█	█										
Test reward functionality & train RL algorithm	All					█	█	█	█	█									
Benchmarking: Challenge 2																			
Benchmarking Documentation	All								█	█	█	█							
Lit review on SOTA RL models	All								█	█	█	█							
Evaluate RL algorithm	All									█	█	█	█						
Env Dev + Benchmark: Challenge 1																			
Code Documentation	Jihoon													█	█				
Develop needle finding CV algorithm	Jihoon													█	█				
Test functionality	Jihoon													█	█				
Evaluate algorithm w/ SRC 2022	Jihoon													█	█				
Env Dev + Benchmark: Challenge 3																			
Code Documentation	Jocelyn + Walee													█	█	█	█		
Develop suture reward function	Jocelyn + Walee													█	█	█	█		
Test reward functionality	Jocelyn + Walee													█	█	█	█		
Benchmark RL algorithm on Challenge 3	Jocelyn + Walee													█	█	█	█		
Compare RL w/ SRC 2022	Jocelyn + Walee														█	█	█		
NeurIPS																			
Write paper	All																█	█	
Code cleanup for NeurIPS submission	All																█	█	

Figure 5: Timeline and Responsibilities

Meeting Plan

- Undergraduate meetings: Weekly - Mondays and Fridays @ 11AM
- Team meetings with mentors: Weekly - Wednesdays @ 9AM (w/ Dr. Liu, Dr. Munawar, Dr. Sahu)
- Communication: Microsoft Teams and JHU email (mentors + students), Discord (students)

Reading List

- Richter, F., Orosco, R. K., & Yip, M. C. (2019). Open-sourced reinforcement learning environments for surgical robotics. arXiv preprint arXiv:1903.02090. [3]
- Introduction to reinforcement learning with David Silver. DeepMind. (n.d.). Retrieved February 19, 2023, from <https://www.deepmind.com/learning-resources/introduction-to-reinforcement-learning-with-david-silver> [7]
- V. M. Varier, D. K. Rajamani, F. Tavakkolmoghaddam, A. Munawar and G. S. Fischer, "AMBF-RL: A real-time simulation based Reinforcement Learning toolkit for Medical Robotics," 2022 International Symposium on Medical Robotics (ISMR), GA, USA, 2022, pp. 1-8, doi: 10.1109/ISMR48347.2022.9807609. [5]
- Medical Open Network for Artificial Intelligence. MONAI. (n.d.). Retrieved February 19, 2023, from <https://monai.io/index.html> [8]
- Gymnasium documentation. Basic Usage. (n.d.). Retrieved February 19, 2023, from https://gymnasium.farama.org/content/basic_usage/ [1]

References

- [1] OpenAI, "Openai gym documentation: Basic usage." [Online]. Available: https://www.gymnasium.dev/content/basic_usage/
- [2] X. B. Peng, M. Andrychowicz, W. Zaremba, and P. Abbeel, "Sim-to-real transfer of robotic control with dynamics randomization," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, May 2018. [Online]. Available: <https://doi.org/10.1109/icra.2018.8460528>
- [3] F. Richter, R. K. Orosco, and M. C. Yip, "Open-sourced reinforcement learning environments for surgical robotics," *arXiv preprint arXiv:1903.02090*, 2019.
- [4] C. R. T. (CRTK), "2021-2022 accelnet surgical robotics challenge." [Online]. Available: <https://collaborative-robotics.github.io/surgical-robotics-challenge/challenge-2021.html>
- [5] V. M. Varier, D. K. Rajamani, F. Tavakkolmoghaddam, A. Munawar, and G. S. Fischer, "Ambf-rl: A real-time simulation based reinforcement learning toolkit for medical robotics," in *2022 International Symposium on Medical Robotics (ISMR)*. IEEE, 2022, pp. 1–8.
- [6] NeurIPS, "Neurips — 2023." [Online]. Available: <https://nips.cc/Conferences/2023>
- [7] D. Silver, "Introduction to reinforcement learning with david silver." [Online]. Available: <https://www.deepmind.com/learning-resources/introduction-to-reinforcement-learning-with-david-silver>
- [8] P. MONAI, "Medical open network for artificial intelligence." [Online]. Available: <https://monai.io/index.html>