


# ORBIT-Surgical: An Open-Simulation Framework for Learning Surgical Augmented Dexterity

Qinxi Yu<sup>\*,1</sup>, Masoud Moghani<sup>\*,1</sup>, Karthik Dharmarajan<sup>2</sup>, Vincent Schorp<sup>2,3</sup>, William Chung-Ho Panitch<sup>2</sup>, Jingzhou Liu<sup>1,5</sup>, Kush Hari<sup>2</sup>, Huang Huang<sup>2</sup>, Mayank Mittal<sup>3,5</sup>, Ken Goldberg<sup>2</sup>, Animesh Garg<sup>1,4,5</sup>



**Fig. 1: ORBIT-Surgical Simulation Benchmark Tasks.** (1) Reach: dVRK Patient Side Manipulator (PSM) to reach a desired position (red sphere), (2) Reach with Obstacles: reach to a desired position (red sphere) with randomly placed obstacle in the scene (blue sphere object; object shape and size are customizable), (3) Suture Needle Lift: lift a suture needle to a desired position, (4) Peg Block Lift: lift a peg block to a desired position, (5) Pick and Place: pick and place a ring on a peg tower, (6) Dual-arm Reach: dual-arm reach to specific desired positions shown with red sphere, (7) Dual-arm Reach with Obstacles: dual-arm reach to specific desired positions (red sphere) with randomly placed obstacles in the scene, (8) Pick and Transfer: pick and transfer a peg block, (9) Needle Handover: handover and regrip a suture needle, (10) Threaded Needle Pass Ring: handover a threaded suture needle through a ring pole, (11) Gauze Cloth Pick: retrieve gauze and lift it to a desired location, (12) Shunt Insertion: insert a shunt (yellow tube) into larger blood vessel phantom (clear tube), (13) Multi-arm dVRK: needle handover task with camera input from the dVRK Endoscopic Camera Manipulator (ECM), (14) STAR Reach: STAR arm to reach a desired position. Best viewed in color. Please see task videos at [orbit-surgical.github.io](https://orbit-surgical.github.io)

**Abstract**—Physics-based simulations have accelerated progress in robot learning for driving, manipulation, and locomotion. Yet, a fast, accurate, and robust surgical simulation environment remains a challenge. In this paper, we present ORBIT-Surgical, a physics-based surgical robot simulation framework with photorealistic rendering in NVIDIA Omniverse. We provide 14 benchmark surgical tasks for the da Vinci Research Kit (dVRK) and Smart Tissue Autonomous Robot (STAR) which represent common subtasks in surgical training. ORBIT-Surgical leverages GPU parallelization to train reinforcement learning and imitation learning algorithms to facilitate study of robot learning to augment human surgical skills. ORBIT-Surgical also facilitates realistic synthetic data generation for active perception tasks. We demonstrate ORBIT-Surgical sim-to-real transfer of learned policies onto a physical dVRK robot.

Project website: [orbit-surgical.github.io](https://orbit-surgical.github.io)

## I. INTRODUCTION

Robotic surgical assistants (RSA) have been widely adopted since the early 2000s [1–3]. RSAs such as Intuitive Surgical’s da Vinci surgical systems enable minimally invasive robotic-assisted surgery, allowing surgeons to perform procedures using fewer, smaller incisions, reducing the risk of complications

and recovery time. With enhanced capabilities and prevalence of RSAs, there is an increasing interest in augmented surgical dexterity [4] for partial automation of specific, monotonous, and error-prone subtasks in surgical procedures. Augmented surgical dexterity has potential to decrease cognitive load, reduce the learning curve, and improve patient outcomes.

Progress in augmented surgical dexterity has evolved slower than applications for robotic manipulation [5–7], locomotion [8], and autonomous driving [9]. These advancements in robot learning often depend on the utilization of simulation environments for either synthetic data generation or interaction based learning [10–13]. While autonomous driving has progressed to Level 3 autonomy and is approaching Level 4, the majority of robotic surgical platforms still lack any form of autonomy [14].

A number of research efforts towards learning-based RSAs have led to simulating surgical subtasks [15–20]. Yet they lack accelerated physics to enable learning long-horizon manipulations. This results in a high barrier for researchers to contribute to the field, and impedes research for algorithmic development. Surgical environments are difficult to simulate due to the complicated interactions of rigid & deformable components and lack of physics simulation models. Additional challenges include motion restriction by the remote center

\*These authors contributed equally to this work

<sup>1</sup>University of Toronto, <sup>2</sup>University of California, Berkeley, <sup>3</sup>ETH Zurich, <sup>4</sup>Georgia Institute of Technology, <sup>5</sup>NVIDIA

Correspondence to: [moghani@cs.toronto.edu](mailto:moghani@cs.toronto.edu), [garg@cs.toronto.edu](mailto:garg@cs.toronto.edu)


of motion (RCM) in trocars for laparoscopic surgery, and simultaneous control of multiple surgical instruments [16]. Moreover, sensor simulation to achieve realistic endoscopic images also remains elusive.

We present **ORBIT-Surgical**, an open-source surgical robot simulator, built on top of NVIDIA Isaac Sim [13]. It provides stable GPU-accelerated physics with ray-traced rendering, rich suite of rigid and deformable object models, as well as a standardized, intuitive, and modular programming interface. ORBIT-Surgical supports a variety of robot learning workflows: reinforcement learning (RL) and imitation learning (IL), teleoperation, and synthetic data generation. ORBIT-Surgical currently supports two surgical robotic platforms; da Vinci Research Kit (dVRK) [24] and Smart Tissue Autonomous Robot (STAR) [25]. ORBIT-Surgical includes simulated surgical subtasks intended for basic surgical maneuver training, along with examples of rigid and deformable surgical materials. The environments are modular to enable the robot learning community to quickly evaluate new algorithms on detailed surgical scenes. Summary of contributions:

- 1) **GPU-based surgical simulation:** An open-source surgical robotics simulator with GPU-accelerated physics, contact-rich physical interaction, and high-fidelity rendering, extending ORBIT [13] framework for robotic surgery.
- 2) **Surgical task benchmarks:** The simulation supports realistic interactions with both rigid and deformable material, a variety of standard implementations of RL and IL algorithms, and various input peripherals including keyboard, spacemouse, gamepad, VR controller, and dVRK Master Tool Manipulator (MTM) for human-expert demonstration collection.
- 3) **Simulation experiments:** Experiments using RL and IL methods demonstrating proof-of-concept policies for the provided surgical tasks. We also evaluate sub-policy learning for long-horizon surgical robotic tasks.
- 4) **Synthetic data generation:** We evaluate ORBIT-Surgical for generative synthetic images to train a needle segmentation framework with a combination of real and synthetic data to improve perception performance by over  $2\times$ .
- 5) **Real robot experiments:** Experiments demonstrating sim-to-real transfer of state-based and RL policies from ORBIT-Surgical to a physical dVRK platform, showing transferability for real-world policy deployment.

## II. RELATED WORK

**Simulators in Robot Learning** Several robot learning efforts have introduced domain specific frameworks, often catering to specialized needs. Many prior frameworks [26, 27] using MuJoCo [10] or Bullet [11] focus mainly on rigid object manipulation tasks. On the other hand, frameworks for deformable bodies [28, 29] mainly employ Bullet [11] or FleX [21], which use particle-based dynamics for soft bodies and cloth simulation. Most of these physics engines are CPU-based, relying on CPU clusters for parallelization. However, limited tooling exists for unified frameworks to enable content development specific to domains in surgery. ORBIT-Surgical builds on recently released framework ORBIT,



**Fig. 2:** (a) Physical da Vinci Research Kit (dVRK) platform, (b) ORBIT-Surgical simulation.

which provides a modular and unified simulation interface for robot learning. ORBIT-Surgical relies on GPU-accelerated physics engines, signed-distance field (SDF) collision checking, and stable solvers based on FEM for deformable body simulation [30, 31].

**Learning in Surgical Robotics** Robot learning for surgical tasks has been attempted for a number of subtasks with varying levels of autonomy [14, 32, 33] such as shape cutting [34, 35], suturing [36, 37], debridement [34], as well as dissection and tissue retraction [32]. However, learning at scale has been limited due to lack of easy to use, high quality simulation frameworks for surgical robotics. Prior state-of-the-art surgical robotics simulators have suffered from issues such as limited rendering capabilities [16], few available surgical environments [18, 19], and lack of native support for deformable objects [17]. But most of all, the lack of GPU-accelerated physics prevents learning [38]. Further, there are numerous proprietary systems such as SimNow [39] and Mimic & Symbionix [40]. These offer virtual environments primarily for teleoperated skill and procedural practice, while lacking programmatic access for learning based methods, and high-speed throughput. As a result, a surgical robotic framework that integrates precise physics, delivers high speed realistic rendering, and accommodates various robot learning tools continues to be an ongoing challenge. TABLE I provides a comparative overview of open simulation frameworks for surgical robotics. In addition, we compare the throughputs of different simulators on environments with similar complexity; namely, robot reach task, rigid object manipulation, and soft object manipulation. We provide the result in TABLE II.

## III. ORBIT-SURGICAL

### A. Framework Overview

ORBIT-Surgical uses NVIDIA Isaac Sim and builds on the ORBIT framework enabling fast and accurate physics interactions, scene rendering, as well as access to robot learning libraries geared toward surgical subtask learning. ORBIT-Surgical supports parallelized GPU simulation [12], contact-rich interactions (i.e. SDF-based collision meshes [41] and convex decomposition [42]), and APIs for object definition (i.e. rigid body, FEM-based deformable body, and well as particle based simulation).

**Simulation Environment Design:** The framework architec-

**TABLE I: Comparison Between Free Open-source Surgical Robotic Simulation Frameworks Available for Robot Learning.**

Name	Physics Engine	Renderer	Collider Approximation	GPU Acceleration	Supported Dynamics				Active Development
					Rigid	Soft	Cloth	Fluid	
LapGym [16]	SOFA	OpenGL	Surface meshes	✗*	✓*	✓	✓	✗	✓
SurRol [17]	Bullet [11]	OpenGL	Convex decomposition	✗	✓	✗	✗	✗	✓
AMBF-RL [18]	Bullet [11]	OpenGL	Surface meshes	✗	✓	✓	✗	✗	✗
UnityFlexML [19]	Flex [21]	Unity3D	Not specified	✗	✗	✓	✗	✗	✗
dVRL [20]	V-REP [22]	OpenGL	✗	✗	✓	✗	✗	✗	✗
Surgical Gym [23]	PhysX 5.1	Photorealistic	SDF	✓	✓	✗	✗	✗	✓
		Ray-tracing (OV-RTX)	Convex decomposition Convex hull						
ORBIT-Surgical	PhysX 5.3	Photorealistic Ray-tracing (OV-RTX)	SDF Convex decomposition Convex hull	✓	✓	✓	✓	✓	✓

\* The check (✓) and cross (✗) denote presence or absence of the feature.

**TABLE II: Frame Rate Comparison Between Surgical Robotic Simulation Frameworks.** The frame rate in frames per second (FPS) for different simulators are evaluated on three environments: robot reach task (Robot), rigid object manipulation (Rigid), and soft object manipulation (Soft). We evaluate all the environments based on the default environment configurations, and the throughput is calculated based on 100 simulation steps. For consistency in hardware comparison, only simulators compatible with the Ubuntu 20.04 LTS are assessed.

Framework	Robot	Rigid	Soft
SurRol [17]	129	92	N.A.
LapGym [16]	621	617	77
ORBIT-Surgical	112444	87805	2048
Speedup	181×	142×	27×

ture consists of a world and an agent, similar to how an operating system runs on a real-world robot to allow interaction with its environment. The agent receives observations from the world (i.e. environment), and computes actions to be applied on the surgical robot to interact with the world accordingly. Here, the world is defined as an entity consisting of the robot, sensory elements, and surrounding objects. The world description is modular in its core and can be designed in standardized scripts or through the graphical interface of NVIDIA Isaac Sim. This modularity and flexibility will allow the community to easily alter and introduce new elements into the existing world environments to meet the needs and features of their surgical applications. The surgical robot is the main component of the simulated surgical scenes. It embodies realistic joint articulations, actuator models, controllers, and collision properties to enable rich interactions with the environment and facilitate the sim-to-real transfer of control policies onto real-world robots.

**Controllers:** We inherit GPU-based implementations of differential inverse kinematic and joint-level controllers from the ORBIT framework. The inverse kinematic models determine the joint positions needed to achieve the desired gripper pose in the Cartesian space. The inverse kinematic methods used to compute the inverse of the Jacobian include Moore-Penrose pseudo-inverse, adaptive SVD, Jacobian transpose, as well as damped least square (Levenberg-Marquardt).

**Action Space:** We support both the Cartesian-space control

and the joint-space control as our action space, enhancing the adaptability and transferability of the trained policies to a real-world dVRK platform. In Cartesian-space, the orientation of the gripper can be represented either as Euler angles or quaternions. Therefore, the quaternion-based action space involves  $(d_x, d_y, d_z, q_w, q_x, q_y, q_z, d_g)$ , where  $d_x, d_y, d_z$  determine the position,  $q_w, q_x, q_y, q_z$  determine the orientation of the gripper in the Cartesian space, and  $d_g$  determines if the gripper jaws are open/closed, respectively. The dVRK Patient Side Manipulator (PSM) has seven degrees-of-freedom (DOF). The first six DOFs correspond to the arm movement and the last DOF determines the gripper state. The arm DOFs consist of revolute (R) and prismatic (P) joints articulated in RRPRRR sequence [17]. In joint-space mode, the action space is defined as  $(q_1, q_2, q_3, q_4, q_5, q_6, d_g)$  where  $q_1, q_2, q_3, q_4, q_5, q_6$  determine the joint positions of the arm, and  $d_g$  determines the state of the gripper jaws, respectively.

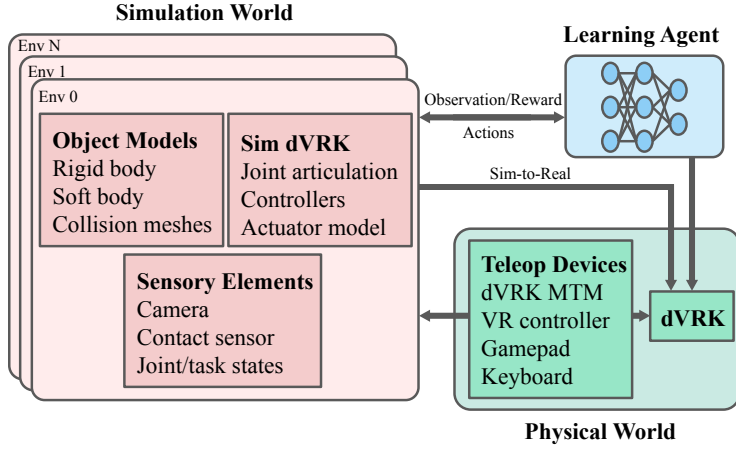
**Observation Space:** We support working with diverse observations methods including the ground-truth states from the simulation environment (e.g. 6D pose of the gripper, object position in the Cartesian space) as well as various sensor simulations (e.g. cameras, height scanners, range sensors, contact sensors). The observation modalities are defined according to the specific needs of each task environment. In this work, we mainly focus on state-based observations for surgical skill learning to provide benchmark results.

**Object Assets:** The dVRK robot is based on meshes from AMBF [38] and SurRol [17] projects. The STAR arm is from Surgical Gym [23]. Surgical assets are partially created in Blender and partially from SurRol [17]. The meshes are converted to USD file format in NVIDIA Isaac Sim. USD files include material definitions, geometric representations, and collision meshes among other asset attributes.

#### B. ORBIT-Surgical: Task Benchmark

ORBIT-Surgical provides surgical training task benchmarks (see Fig. 1) inspired by Fundamentals of Laparoscopic Surgery (FLS) training curriculum [43] to reflect the technical skills for basic surgical maneuvers. Each task is provided in a standalone environment for benchmarking. These environments can be used to generate hand-scripted trajectories, collect human-expert samples using multiple input peripherals, as





**Fig. 3: ORBIT-Surgical Framework Architecture.** The ORBIT framework comprises of a simulation world, intelligent agents, and transfer of intelligence trained in simulation to the real world. The ORBIT-Surgical simulation world includes realistic clones of surgical robots with joint articulation and low-level controllers. The object models include a rich suite of rigid and soft objects with precise collision properties. Sensory element in the simulation imitates the real world sensors to provide training data. Peripheral I/O devices including the dVRK MTM are connected to the simulation environment to enable teleoperation of the simulated surgical robots. This setup facilitates the real-time capture of inputs from human experts, leveraging real-world demonstrations for policy learning in simulation.

well as training RL and IL policies to facilitate advancements in surgical autonomy. Up to 8000 environments can be simulated in parallel on a single NVIDIA RTX 3090 GPU for efficient data collection and fast surgical robot learning. Environments are modular and can be modified or integrated with other robots or tasks. An illustration of each task along with a short description is provided in Fig. 1.

### C. Robot Learning Workflows

**Reinforcement Learning:** We support various reinforcement learning (RL) frameworks; rl-games [44], RSL-rl [8], stable-baselines-3 (SB3) [45]. In this work, we mainly utilise the RSL-rl library, as it is optimized for vectorized simulation and GPU training. We use the Proximal Policy Optimization (PPO) [46] algorithm to establish a learning baseline.

**Teleoperation and Behavior Cloning:** Expert recorded samples can help to investigate learning for challenging and long-horizon surgical tasks that are hard to solve with the current RL algorithms. We support the RoboMimic library [47] which encompasses IL algorithms including behavior cloning (BC) and offline reinforcement learning (BCQ). ORBIT-Surgical offers support for various I/O peripheral devices to teleoperate the dVRK gripper in simulation in real-time and record human expert data to train IL algorithms. Currently, our peripheral devices include keyboard, spacemouse, gamepad, VR controller, as well as the dVRK Master Tool Manipulator (MTM) to provide input commands and control the dVRK gripper in Cartesian space.

## IV. EXPERIMENTAL RESULTS

In this section, we investigate task learning for the proposed surgical environments entirely in simulation and evaluate how well the trained policies transfer to a real world dVRK platform. To this end, we are answering the following questions:

- 1) How accurately do ORBIT-Surgical’s environments simulate physical interactions between the dVRK surgical robot and deformable objects?
- 2) How well do RL algorithms perform for learning simple surgical subtasks?
- 3) What is the benefit of the data collection from human-expert or code generated trajectories for learning from

demonstration for long-horizon tasks?

- 4) How well do trained policies in ORBIT-Surgical transfer from simulation environments to a physical dVRK surgical robot in the lab?


*Experimental Setup:* The visualizations of experimentation environments are depicted in Fig. 1. Each environment consists of either single or dual arms in a fixed location. The initial and desired states of tasks are randomized across the parallelized environments. The setup for the real world experiments consists of two arms mounted on a table. The surgical robot arms can be controlled either programmatically or through the MTM surgeon control platform.

### A. Reinforcement Learning

To demonstrate the high learning speed facilitated by ORBIT-Surgical, RL policies were trained for surgical subtasks in ORBIT-Surgical and were compared to those of LapGym [16], one of the most recently released surgical simulators. Two tasks, i.e. Reach (Task1) and Suture Needle Lift (Task3) were selected for this comparison to study both simulators with and without object manipulation tasks. Both simulators have similar Reach tasks, however, the Pick and Place task from LapGym was selected to be of comparable complexity (i.e. requiring object interaction and data processing) with the ORBIT-Surgical Suture Needle Lift task to study Rigid Object Manipulation. For the Reach task, both simulators learned the simple task quite fast since it does not involve interactions with objects. However, in Rigid Object Manipulation tasks, ORBIT-Surgical started to learn to complete the task in around 100 minutes, whereas LapGym was not able to learn the task at all during the two-hour experiment time. Training rigid object manipulation tasks in ORBIT-Surgical averaged 71700 FPS with RSL-RL while only getting 617 FPS throughput in LapGym with SB3 library on an NVIDIA 3090 RTX GPU. During training, ORBIT-Surgical was able to parallelize 4096 environments. The success rate comparison is illustrated in Fig. 4.

### B. Behavior Cloning from Demonstration


Surgical manipulation tasks involve a high level of precision. Current RL algorithms, at times, may not be capable of



**Fig. 4: Reinforcement Learning.** Success rate versus wall-clock time for RL policy training using the PPO algorithm on Reach task (above), and Rigid Object Manipulation task (below) evaluated for both ORBIT-Surgical and LapGym across 10 trials for each saved checkpoint in simulation.

learning these tasks with optimal efficiency. In such instances, leveraging expert demonstrations for IL policies can be an effective strategy. ORBIT-Surgical supports multiple teleoperation interfaces for data collection, including keyboard, VR controller, and MTM, where the collected data from an expert can be used to train BC models. Here, we evaluate the efficacy of BC policies for the single arm Suture Needle Lift task, where the initial position and goal states of the needle are either fixed or randomized. We collected 100, 500, and 1000 trajectories and subsequently trained BC policies. Due to the large number of demonstrations required for comparison, we collected demonstrations using hand-scripted state trajectories with variations in initial states and goal states. Fig. 5 compares the success rate of policy execution for the Suture Needle Lift task in the simulation environment.

Long-horizon tasks are computationally expensive and can be hard to solve by plain BC algorithms. Instead, multi-stage BC policy can be learned efficiently, where long-horizon tasks are divided into subtasks, and sub-policies are trained to perform each subtask. All the subtasks are then chained for end-to-end policy execution. As an example, we show learning from demonstration for the Needle Handover task. We collected 1000 long-horizon trajectories of transferring a needle from one gripper to the other, and trained a single BC policy for the complete task which scored a 0 % success rate. We then divided the task into three stages, i.e. Lift, Handover, and Reach and trained BC policies to perform each subtask, respectively. All three subtasks were chained for end-to-end policy execution. For data collection, we recorded 100 trajectories for both Lift and Handover, and 300 trajectories for Reach. Policy execution for the Needle Handover task across 10 trials achieved a 40% success rate in simulation. Please visit the website for a successful needle handover policy execution in simulation.



**Fig. 5: Generalization of BC Policies in Simulation.** We collected demonstrations using hand-scripted state trajectories with variations in initial states and goal states for Suture Needle Lift (Fig. 1: Task3) in three different settings where only initial states are randomized (I), only goal states are randomized (G), and both initial and goal states are randomized (IG). Here, we compare the success rate of each trained policy across 10 trials in simulation. We note improved behavior for the policies trained on demonstrations with randomized initial and goal states.

**TABLE III: Real World Needle Segmentation Evaluation.** We note that dataset augmentation with simulation to a set of real images (50) improves segmentation IoU by over  $2\times$ .

Training Regime	Sim Dataset Size	Real Dataset Size	IoU
Sim Only	2900	0	0.06
Real Only	0	100	0.15
Sim+Real (20:1)	1000	50	<b>0.42</b>
Sim+Real (29:1)	2900	100	<b>0.40</b>


### C. Synthetic Data Generation: Needle Segmentation

The fusion of photorealistic synthetic and real-world images enhances surgical tool segmentation models. 2900  $640 \times 480$  RGB images and corresponding segmentation masks in the ORBIT-Surgical simulator were paired with 100 RGB image-mask pairs of a needle in the real world, as shown in Fig. 6. Using this data in varying proportions, several models with the same hyperparameters were trained and compared in TABLE III. Intersection over union (IoU) metric was considered to evaluate the performance of segmentation models on a separate evaluation set of 100 real-world RGB images by comparing the predicted segmentation to ground truth masks. The models trained exclusively on the simulation data generated segmentation masks that included other objects in the image such as the gripper, resulting in a low IoU. Combining simulation and real data resulted in predictions that do not select the gripper in the segmentation mask, while covering most of the needle area, leading to higher IoU scores than training exclusively on either simulation or real data.

### D. Sim-to-Real: Deployment on the dVRK

We developed a state machine for two tasks involving deformable objects, Threaded Needle Pass Ring and Insert Tube, and deployed the trajectories recorded in the simulation environment on a real-world dVRK platform.

For the Threaded Needle Pass Ring task, we de-




**Fig. 6: Synthetic Data Generation.** Simulated images used as training data for needle segmentation (top). Real images of the workspace consisting of the needle are used for evaluation (bottom left). A model trained with both sim and real images is evaluated with a real dataset, resulting in masks shown on the bottom right.


signed an environment containing two dVRK arms, a needle attached to a thread as well as a loop on a pole. We programmed the state machine to grasp the needle using the right dVRK arm, rotate it through the ring and regrasp the needle with the left dVRK arm as seen in Fig. 7. The history of joint and gripper jaw states for both arms was recorded in ORBIT-Surgical and transferred to the real robot. The joint and gripper state trajectories were adapted to match the real robot position offset and jaw command range. We used a surgical suture with a GS-21 needle for the real setup and replayed the trajectory on both arms. A qualitative comparison is shown in Fig. 7.

The Insert Tube task consists of two horizontal flexible tubes of equal length but different diameters and a single dVRK arm. This experiment shows the capacity to simulate shunt insertion maneuvers. The smaller diameter, more flexible tube was grasped by the surgical robot arm and inserted into the stiffer, larger diameter tube as shown in Fig. 1: Task 12. Similar to the Threaded Needle Pass Ring experiment, the joints trajectory and gripper states were recorded in ORBIT-Surgical and replayed on the real robot. For the real setup, we used latex surgical tubing for the shunt and a clear vinyl tube for the artery. The artery is held in place using a fixed clamp. Please refer to the video for real and simulated robot experiments.

We deployed a PPO policy trained with RSL-rl on the Suture Needle Lift task in Cartesian space on a real dVRK setup. The needle was placed below one of the grippers and on top of a green piece of foam for safety reasons shown in Fig. 8. All of the real-world observations for this task were derived from the real dVRK’s joint encoders and other known values except for the needle position. For receiving the needle’s center position, we manually performed initial measurements to determine its location, and hard coded this value in the observation space. When the gripper grasped the needle, we used the last known relative pose between the gripper and the needle along with the end effector position to derive the current needle position. We performed 20 trials, and observe that the transferred policy achieved a 50% success



**Fig. 7: Sim-to-Real.** Threaded Needle Pass Ring task in ORBIT-Surgical (top) and on a real physical dVRK robot (bottom). (a) Right gripper grasps and lifts a threaded needle, (b) right gripper hands over the needle to the left gripper through the ring, (c) handover is complete and the left gripper moves away to a desired final location. Note how suture thread deformation in sim closely matches that in real.



**Fig. 8: Zero-shot Suture Needle Lift Policy Transfer.** (1) Starting position of the dVRK arm and suture needle, (2) arm moving towards the needle, (3) grasping the needle, (4) successfully lifting the needle. The policy was fully trained in ORBIT-Surgical and deployed on the physical robot.

rate. One execution run of the policy is shown in Fig. 8.

## V. CONCLUSIONS AND FUTURE WORK

This paper introduces ORBIT-Surgical, an open physics-based simulator which provides contact-rich surgical simulation environments for robot learning. GPU parallelization enables ORBIT-Surgical to train reinforcement learning (RL) policies for surgical robotics tasks on the order of hours. Through multiple teleoperation interfaces, ORBIT-Surgical leverages real-world demonstrations for policy learning in simulation. ORBIT-Surgical enables high-fidelity synthetic data generation by harnessing GPU-accelerated physics with ray-traced rendering. Together with high-fidelity rendering and high-speed simulation, we hope to accelerate progress in learning-based assistive autonomy in robotic surgery.

While ORBIT-Surgical provides high throughput in state-based environments, vectorized image rendering is still under development. This will enable us to incorporate visual inputs into training RL policies. ORBIT-Surgical also does not model the cable stretching effects found on real dVRK surgical robots, as demonstrated by the success rate of 50% when deploying the RL policy for the Suture Needle Lift task. In future work, we will incorporate cutting of soft material [48] and provide algorithms to solve longer-horizon end-to-end tasks such as suturing and knot tying.



# REFERENCES

- [1] J.-F. Onnasch, F. Schneider, V. Falk, M. Mierzwa, J. Bucerius, and F. W. Mohr, "Five years of less invasive mitral valve surgery: from experimental to routine approach," in *The heart surgery forum*, vol. 5, no. 2, 2002, pp. 132–135.
- [2] A. Tewari, J. Peabody, R. Sarle, G. Balakrishnan, A. Hemal, A. Shrivastava, and M. Menon, "Technique of da Vinci robot-assisted anatomic radical prostatectomy," *Urology*, vol. 60, no. 4, pp. 569–572, 2002.
- [3] J. Bodner, H. Wykypiel, G. Wetscher, and T. Schmid, "First experiences with the da Vinci™ operating robot in thoracic surgery," *European Journal of Cardio-thoracic surgery*, vol. 25, no. 5, pp. 844–851, 2004.
- [4] K. Goldberg. (2023) Augmented Dexterity: How Robots Can Enhance Surgeon Dexterity. [Online]. Available: <https://bit.ly/Augmented-Dexterity-S24>
- [5] A. Allshire, M. Mittal, V. Lodaya, V. Makoviychuk, D. Makoviichuk, F. Widmaier, M. Wüthrich, S. Bauer, A. Handa, and A. Garg, "Transferring dexterous manipulation from GPU simulation to a remote real-world trifinger," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 11 802–11 809.
- [6] A. Handa, A. Allshire, V. Makoviychuk, A. Petrenko, R. Singh, J. Liu, D. Makoviichuk, K. Van Wyk, A. Zhurkevich, B. Sundaralingam *et al.*, "Dextreme: Transfer of agile in-hand manipulation from simulation to reality," *arXiv preprint arXiv:2210.13702*, 2022.
- [7] J. Mahler, M. Matl, V. Satish, M. Danielczuk, B. DeRose, S. McKinley, and K. Goldberg, "Learning ambidextrous robot grasping policies," *Science Robotics*, vol. 4, no. 26, p. eaau4984, 2019.
- [8] N. Rudin, D. Hoeller, P. Reist, and M. Hutter, "Learning to walk in minutes using massively parallel deep reinforcement learning," in *Conference on Robot Learning*. PMLR, 2022, pp. 91–100.
- [9] P. Kaur, S. Taghavi, Z. Tian, and W. Shi, "A survey on simulators for testing self-driving cars," in *2021 Fourth International Conference on Connected and Autonomous Driving (MetroCAD)*. IEEE, 2021, pp. 62–70.
- [10] E. Todorov, T. Erez, and Y. Tassa, "Mujoco: A physics engine for model-based control," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2012, pp. 5026–5033.
- [11] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning," 2016–2021.
- [12] V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa *et al.*, "Isaac Gym: High performance GPU-based physics simulation for robot learning," *arXiv preprint arXiv:2108.10470*, 2021.
- [13] M. Mittal, C. Yu, Q. Yu, J. Liu, N. Rudin, D. Hoeller, J. L. Yuan, R. Singh, Y. Guo, H. Mazhar *et al.*, "ORBIT: A unified simulation framework for interactive robot learning environments," *IEEE Robotics and Automation Letters*, 2023.
- [14] A. Attanasio, B. Scaglioni, E. De Momi, P. Fiorini, and P. Valdastrì, "Autonomy in surgical robotics," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 4, pp. 651–679, 2021.
- [15] A. T. Bourdillon, A. Garg, H. Wang, Y. J. Woo, M. Pavone, and J. Boyd, "Integration of reinforcement learning in a virtual robotic surgical simulation," *Surgical Innovation*, vol. 30, no. 1, pp. 94–102, 2023.
- [16] P. M. Scheikl, B. Gyenes, R. Younis, C. Haas, G. Neumann, M. Wagner, and F. Mathis-Ullrich, "LapGym—an open source framework for reinforcement learning in robot-assisted laparoscopic surgery," *arXiv preprint arXiv:2302.09606*, 2023.
- [17] J. Xu, B. Li, B. Lu, Y.-H. Liu, Q. Dou, and P.-A. Heng, "SurRol: An open-source reinforcement learning centered and dvrk compatible platform for surgical robot learning," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021.
- [18] V. M. Varier, D. K. Rajamani, F. Tavakkolmoghadam, 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.
- [19] E. Tagliabue, A. Pore, D. Dall'Alba, E. Magnabosco, M. Piccinelli, and P. Fiorini, "Soft tissue simulation environment to learn manipulation tasks in autonomous robotic surgery," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 3261–3266.
- [20] F. Richter, R. K. Orosco, and M. C. Yip, "Open-sourced reinforcement learning environments for surgical robotics," *arXiv preprint arXiv:1903.02090*, 2019.
- [21] NVIDIA, "Flex: A particle-based simulation library," May 2017. [Online]. Available: <https://github.com/NVIDIAGameWorks/Flex>
- [22] G. A. Fontanelli, M. Selvaggio, M. Ferro, F. Ficuciello, M. Vendittelli, and B. Siciliano, "A V-REP simulator for the da Vinci Research Kit robotic platform," in *2018 7th IEEE International Conference on Biomedical Robotics and Biomechatronics (Biorob)*. IEEE, 2018, pp. 1056–1061.
- [23] S. Schmidgall, A. Krieger, and J. Eshraghian, "Surgical Gym: A high-performance GPU-based platform for reinforcement learning with surgical robots," *arXiv preprint arXiv:2310.04676*, 2023.
- [24] P. Kazanzides, Z. Chen, A. Deguet, G. S. Fischer, R. H. Taylor, and S. P. DiMaio, "An open-source research kit for the da Vinci® Surgical System," in *2014 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2014, pp. 6434–6439.
- [25] A. Shademan, R. S. Decker, J. D. Opfermann, S. Leonard, A. Krieger, and P. C. Kim, "Supervised autonomous robotic soft tissue surgery," *Science translational medicine*, vol. 8, no. 337, pp. 337ra64–337ra64, 2016.

- [26] Y. Zhu, J. Wong, A. Mandlekar, and R. Martín-Martín, “robosuite: A modular simulation framework and benchmark for robot learning,” in *arXiv preprint arXiv:2009.12293*, 2020.
- [27] S. James, Z. Ma, D. Rovick Arrojo, and A. J. Davison, “RLBench: The robot learning benchmark & learning environment,” *IEEE Robotics and Automation Letters*, 2020.
- [28] R. Antonova, P. Shi, H. Yin, Z. Weng, and D. K. Jensfelt, “Dynamic environments with deformable objects,” in *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- [29] X. Lin, Y. Wang, J. Olkin, and D. Held, “SoftGym: Benchmarking deep reinforcement learning for deformable object manipulation,” in *Conference on Robot Learning (CoRL)*. PMLR, 2020.
- [30] M. Macklin and M. Muller, “A constraint-based formulation of stable neo-hookean materials,” in *ACM SIGGRAPH Conference on Motion, Interaction and Games*, New York, NY, USA, 2021.
- [31] M. Macklin, K. Storey, M. Lu, P. Terdiman, N. Chentanez, S. Jeschke, and M. Müller, “Small steps in physics simulation,” in *ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, 2019.
- [32] T. D. Nagy and T. Haidegger, “Autonomous surgical robotics at task and subtask levels,” in *Advanced Robotics and Intelligent Automation in Manufacturing*. IGI global, 2020, pp. 296–319.
- [33] F. Ficuciello, G. Tamburrini, A. Arezzo, L. Villani, and B. Siciliano, “Autonomy in surgical robots and its meaningful human control,” *Paladyn, Journal of Behavioral Robotics*, vol. 10, no. 1, pp. 30–43, 2019.
- [34] 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 *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 1202–1209.
- [35] 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 *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 2371–2378.
- [36] S. Sen, A. Garg, D. V. Gealy, S. McKinley, Y. Jen, and K. Goldberg, “Automating multi-throw multilateral surgical suturing with a mechanical needle guide and sequential convex optimization,” in *2016 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2016, pp. 4178–4185.
- [37] S. Krishnan, A. Garg, S. Patil, C. Lea, G. Hager, P. Abbeel, and K. Goldberg, “Transition state clustering: Unsupervised surgical trajectory segmentation for robot learning,” *International Journal of Robotics Research (IJRR)*, 2017.
- [38] A. Munawar, Y. Wang, R. Gondokaryono, and G. S. Fischer, “A real-time dynamic simulator and an associated front-end representation format for simulating complex robots and environments,” in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 1875–1882.
- [39] Intuitive Surgical. SimNow Simulation Software. [Online]. Available: <https://www.intuitive.com/en-us/products-and-services/da-vinci/learning/simnow>
- [40] Surgical Science. Mimic and Simbionix. [Online]. Available: <https://surgiscience.com/custom-made-simulation/robotic-surgery-simulation/>
- [41] Y. Narang, K. Storey, I. Akinola, M. Macklin, P. Reist, L. Wawrzyniak, Y. Guo, A. Moravanszky, G. State, M. Lu *et al.*, “Factory: Fast contact for robotic assembly,” *arXiv preprint arXiv:2205.03532*, 2022.
- [42] K. Mamou, E. Lengyel, and A. Peters, “Volumetric hierarchical approximate convex decomposition,” in *Game Engine Gems 3*. CRC Press, 2016.
- [43] Fundamentals of Laparoscopic Surgery. [Online]. Available: <https://www.flsprogram.org/>
- [44] D. Makoviichuk and V. Makoviychuk, “rl-games: A High-performance Framework for Reinforcement Learning,” May 2021. [Online]. Available: [https://github.com/Denys88/rl\\_games](https://github.com/Denys88/rl_games)
- [45] A. Raffin, A. Hill, A. Gleave, A. Kanervisto, M. Ernestus, and N. Dormann, “Stable-baselines3: Reliable reinforcement learning implementations,” *The Journal of Machine Learning Research*, 2021.
- [46] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” *arXiv preprint arXiv:1707.06347*, 2017.
- [47] A. Mandlekar, D. Xu, J. Wong, S. Nasiriany, C. Wang, R. Kulkarni, L. Fei-Fei, S. Savarese, Y. Zhu, and R. Martín-Martín, “What matters in learning from offline human demonstrations for robot manipulation,” *arXiv preprint arXiv:2108.03298*, 2021.
- [48] E. Heiden, M. Macklin, Y. Narang, D. Fox, A. Garg, and F. Ramos, “Disect: A differentiable simulator for parameter inference and control in robotic cutting,” 2022.