

# Learning to Localize, Grasp, and Hand Over Unmodified Surgical Needles

Albert Wilcox\*, Justin Kerr\*, Brijen Thananjeyan, Jeff Ichnowski,  
Minho Hwang, Samuel Paradis, Danyal Fer, Ken Goldberg

\* equal contribution

**Abstract**—Robotic Surgical Assistants (RSAs) are commonly used to perform minimally invasive surgeries by expert surgeons. However, long procedures filled with tedious and repetitive tasks such as suturing can lead to surgeon fatigue, motivating the automation of suturing. As visual tracking of a thin reflective needle is extremely challenging, prior work has modified the needle with nonreflective contrasting paint. As a step towards automation of a suturing subtask without modifying the needle, we propose HOUSTON: Handoff of Unmodified, Surgical, Tool-Obstructed Needles, a problem and algorithm that uses a learned active sensing policy with a stereo camera to localize and align the needle into a visible and accessible pose for the other arm. To compensate for robot positioning and needle perception errors, the algorithm then executes a high-precision grasping motion that uses multiple cameras. In physical experiments using the da Vinci Research Kit (dVRK), HOUSTON successfully passes unmodified surgical needles with a success rate of 96.7% and is able to perform handover sequentially between the arms 32.4 times on average before failure. On needles unseen in training, HOUSTON achieves a success rate of 75 – 92.9%. To our knowledge, this work is the first to study handover of unmodified surgical needles. See <https://tinyurl.com/houston-surgery> for additional materials.

## I. INTRODUCTION

Robotic Surgical Assistants (RSAs) currently rely on human supervision for the entirety of surgical tasks, which can consist of many very repetitive subtasks such as suturing. Automation of surgical subtasks may reduce surgeon fatigue [37], with initial results in surgical cutting [23, 34], debridement [23, 30], suturing [7–10, 28, 31–33], hemostasis [26], and peg transfer [12, 13, 24, 27].

This paper considers automation of the bimanual regrasping subtask [7] of surgical suturing, which involves passing a surgical needle from one end effector to another. This handover motion is performed in between stitches during suturing, and is a critical step, as accurately positioning the needle in the end effector affects the stability of its trajectory when guided through tissue. Because varying cable tension in the cables driving the arms causes inaccuracies in motions, a high precision task such as passing a needle between the end effectors is challenging [12, 13, 21, 24, 25].

The task is also difficult because 3D pose information is critical for successfully manipulating needles, and surgical needles are challenging to perceive with RGB or active depth

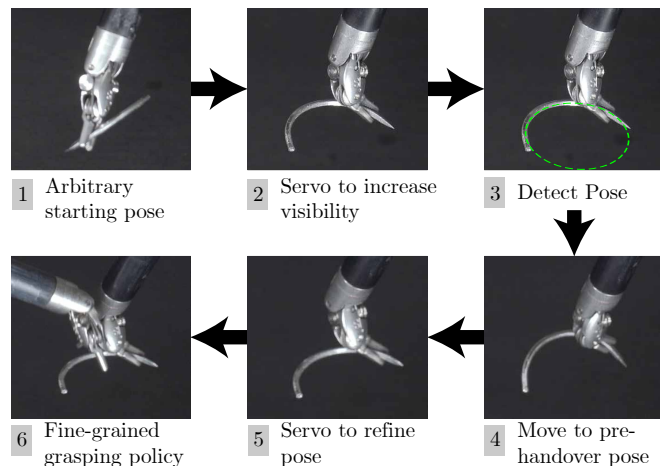


Fig. 1: **Algorithm overview:** The algorithm first servos the needle into a position that is easily visible to the stereo camera to produce a high confidence pose estimate. Using this estimate, it coarsely reorients the needle towards the grasping arm, and then refines this iteratively to correct for positioning errors. Finally, it executes a learned visual servoing grasping policy to complete the handover.

sensors due to their reflective surface, thin profile [17], and self-occlusions (Figure 2). Prior work has mitigated this by painting needles [7, 31] and using color segmentation, but this solution is not practical for clinical use.

To manipulate unmodified surgical needles, we combine recent advances in deep learning, active sensing, and visual servoing. We present HOUSTON: Handoff of Unmodified, Surgical, Tool-Obstructed Needles, a problem and algorithm for using stereo vision with coarse and fine-grained control policies (Figure 1) to sequentially localize, orient, and handover unmodified surgical needles. We present a localization method using stereo RGB with a deep segmentation network to output a point cloud of the needle in the workspace. This point cloud is used to define a coarse robot policy that uses visual servoing to reorient the needle for handover in a pose that is visible to the cameras and accessible by the other end effector. However, due to inaccuracies of robot positioning and the perception system, further corrections may be necessary. We train a fine robot policy from a small set of human demonstrations to perform these subtle but critical corrections from images for unmodified surgical needles.

This paper makes the following contributions:

- 1) A perception pipeline using stereo RGB to accurately estimate the pose of surgical steel needles in 3D space, enabling needle manipulation without active depth

sensors and painted needles.

- 2) A visual servoing algorithm to perform coarse reorientation of a surgical needle for grasping.
- 3) A needle grasping policy that performs fine control of the needle learned from a small set of human demonstrations to compensate for robot positioning and needle sensing inaccuracies.
- 4) Combination of the pose estimator (1), the servoing algorithm (2) and needle controller (3) to perform bimanual surgical needle regrasping, where physical experiments on the da Vinci Research Kit (dVRK) [15] suggest a success rate of 96.7% on needles used in training, and 75 – 92.9% on needles unseen in training. On sequential handovers, HOUSTON successfully executes 32.4 handovers on average before failure.

## II. RELATED WORK

### A. Automation in Surgical Robotics

Automation of surgical subtasks is an active area of research with a rich history. Prior literature has studied automation of tasks related to surgical cutting [19, 23, 34], debridement [16, 23], hemostasis [26], peg transfer [12, 13, 24], and suturing [7, 9, 28, 31–33]. While automated suturing has been studied in prior work [28, 31], suturing without modifications such as painted fiducial markers is an open research problem. Recent work studies robust and general approaches to specific subproblems within suturing, including the precise manipulation of surgical needles during suturing from needle extraction [32] to bimanual regrasping [7], which is the focus of this work.

Needle manipulation is also studied by [32], where the approach studies the extraction of needles from tissue phantoms to compute robust grasps of the needle even in self-occluded configurations. Bimanual needle regrasping was studied in detail by Chiu *et al.* [7], with impressive results on simulation-trained policies that take needle end effector poses as input. We extend their problem definition to consider multiple handoffs of unmodified needles and end effectors, which requires perception of the needle and robot pose from images without color segmentation. Needle grasping has also been studied using visual servoing policies in D’Ettorre *et al.* [9], where the needle is painted with green markers to track its position during closed loop visual servoing. Varier *et al.* [35] study tabular RL policies for needle regrasping in a discretized space in a fixed setup with known needle pose and experiments without the needle on the dVRK. The experiments in [35] suggest that value iteration-trained policies can mimic expert trajectories used for inverse reinforcement learning. In contrast, we present an algorithm compatible with significantly varying initial needle and gripper poses using only image observations. We additionally present many physical experiments with a needle, evaluating the success rate and speed of the algorithm.

### B. Visual Servoing, and Active Perception

Visual servoing (VS) is a popular technique in robotics [11, 18], and has recently been applied to compensate for surgical

robot imprecision in the surgical peg transfer task [24]. While classical VS approaches typically make use of hand-tuned visual features and known system dynamics [5, 6], recent work proposes learning end-to-end visual servoing policies from examples [14, 20]. To reduce the need for tuned features and dynamics models and also reduce the number of training samples required to create a robust VS policy for bimanual needle regrasping, we present a hybrid approach that combines coarse motion planning with fine control, where a learned VS policy is only used in parts of the task where high precision is required. This framework, called intermittent visual servoing (IVS), was studied in detail by [24], where the system switches between a classical trajectory optimizer and imitation learning VS policy based on the precision required at the time. Inspired by this technique, we present an IVS approach to bimanual needle regrasping, that combines coarse perception and planning with fine VS control. Because this task requires reasoning about depth across several directions, we present a multi-view VS policy that learns to precisely hand over the needle based on several camera views.

Active perception is a popular technique with many variations to localize objects prior to manipulation by maximizing information gain about their poses [1, 2, 22, 29, 36]. This has been studied in the context of robot-assisted surgery, where the endoscope position is automatically adjusted via a policy learned from demonstrations to center the camera focus on inclusions during surgeon teleoperation. In this work, we actively servo the needle to highly visible poses to maximize the accuracy of its pose estimate. This is most similar to [1], where the authors propose an algorithm to actively select views of a grasping workspace to uncover enough information about unknown objects to plan grasps.

## III. PROBLEM FORMULATION

The HOUSTON problem extends and generalizes the previous problem definition from Chiu *et al.* [7] to include unmodified needles, occlusion, and multiple handoffs.

### A. Overview

In the HOUSTON problem, the surgical robot starts with a curved surgical needle with known curvature and radius grasped by one gripper and must accurately pass it to the other gripper and back. This is challenging due to the needle’s reflective surface, thin profile, and pathological configurations [7, 32] as depicted in Figure 2. Once the needle is successfully passed to the other end effector, it is passed back to the first end effector. This process is repeated  $N_{\max}$  times, or until the needle is dropped. We also consider a special case of this problem, the single-handover version, in which  $N_{\max} = 1$ .

### B. Notation

Let  $p_L(t)$  and  $p_R(t)$  denote the poses of the left and right grippers, respectively, at discrete timestep  $t$  with respect to a world coordinate frame. The needle has pose  $p_N(t)$  with respect to the world frame. Observations of the workspace are available via RGB images from a stereo camera,  $I_L(t)$




Fig. 2: **Visible and occluded configurations:** The left two frames depict needle orientations that are easily identifiable. However, the needle frequently reaches states that are occluded by the gripper or self-occluded, which makes estimating its state challenging.

and  $I_R(t)$ , or overhead monocular RGB images  $I_O(t)$  from an RGB camera. The left and right cameras in the stereo pair have world poses  $p_{st, \text{left}}$  and  $p_{st, \text{right}}$ , respectively. The overhead camera has pose in world frame  $p_O$ . Each trial starts with the needle in the left gripper and ends when the needle is dropped. Additionally, the trial terminates if no successful handoff occurs in  $\tau_{\text{max}}$  timesteps.

At timestep  $t$ , the algorithm is provided observation  $y(t)$  which contains images from the sensors:  $y(t) = (I_L(t), I_R(t), I_O(t))$ . The algorithm outputs a target pose and jaw state for each end effector  $u(t) = ((p_L(t+1), \varphi_L(t+1)), (p_R(t+1), \varphi_R(t+1)))$ , where  $\varphi_L(t+1) \in \{0, 1\}$  indicates the whether the left jaw is closed at timestep  $t+1$ . *C. Assumptions*

In order to deterministically evaluate HOUSTON policies in a wide variety of needle configurations, we discretize the needle-in-gripper pose possibilities by choosing a number of categories across three degrees of freedom:

- 1) The needle’s curve can face either towards or away from the camera, providing 2 possibilities
- 2) The gripper may hold the needle either at the tip, or  $30^\circ$  inwards following the curvature of the needle. This degree of freedom has 2 possibilities.
- 3) The rotation of the needle about the tangent line  $\omega$  to the point the gripper intersects ranges from  $0^\circ$  to  $180^\circ$ , and is discretized into 7 bins as in Figure 3.

This gives a total of 28 possible needle configurations and we perform a grid search over these possibilities. We chose these configurations to be representative of those seen in suturing tasks post-needle extraction.

While the robot encoders provide an estimate of the gripper poses  $p_L(0)$  and  $p_R(0)$ , the precise needle pose is unknown due to cabling effects of the arms. We assume access to a stereo RGB pair in the workspace, an overhead RGB camera, and the transforms between the coordinate frames of these cameras and the robot arms.

#### D. Evaluation Metrics

We evaluate HOUSTON by recording: *i*) the number of successful handoffs in a multi-handoff trial, *ii*) the success rate per arm of single handoffs beginning from each configuration in III-C, and *iii*) the average time for each handoff.

### IV. HOUSTON ALGORITHM

HOUSTON uses active stereo visual servoing with both a coarse-motion and fine-motion learned policy for the bimanual regrasping task.




Fig. 3: **Needle Reset Degrees of Freedom:** We vary the starting configurations of the needle relative to the gripper in three degrees of freedom as described in III-C. **Left:** We rotate the needle into 7 discretized states about the axis  $\omega$  tangent to the needle. **Middle:** An example holding the middle  $30^\circ$  inward from the tip. **Right:** An example where the needle’s arc faces towards the camera.

#### A. Phase 1: Active Needle Presentation

In the first phase, the algorithm repositions the needle to a pose where the other arm can easily grasp it without collisions and the cameras can clearly view it. Throughout execution of the coarse policy, we parameterize the needle as a circle of known radius, and measure its state in world frame as a center point  $c_p$ , normal vector  $c_n$ , and a needle tip point  $c_t$  as shown in Figure 6. Active Needle Presentation consists of two stages: *needle acquisition* and *handover positioning*. The needle acquisition stage moves the needle to maximize visibility, and the positioning stage uses visual servoing to move the needle into a graspable state.

1) *Needle State Estimation:* The state estimator passes stereo images into a fully convolutional neural network that is trained to output segmentation masks for the needle in each image. See the project website for architecture details. It computes a distance transform of the segmentation mask to label each pixel with its distance to the nearest unactivated pixel. Next, it finds peaks in the distance transform along horizontal lines in each image, which correspond to points near the center of activated patches. It then triangulates all pairs of peaks along each horizontal line in the images to obtain a point-cloud as in Figure 4. Because this may triangulate outlier points from the gripper or incorrectly match points on different parts of the needle, RANSAC is applied to filter out incorrect point correspondences and returns the final predicted needle state. At each iteration, it samples a set of 3 points, to which a plane is fit. Each subset of 3 points generates 2 candidate circles in the plane, corresponding to the two circles which pass through one of the pairs. RANSAC uses an inlier radius of 1mm and runs for 300 iterations.

The network is first trained on a dataset of 2000 simulated stereo images of randomly-placed, textured and scaled floating needles and random objects [4] above a surface plane generated in Blender 2.92. Lighting intensity, size, and position, and stereo camera position are also randomized. The segmentation network is fine-tuned on a dataset of 200 manually-labeled images of the surgical needle in the end effector. Training a network on a PC with an NVIDIA V100 GPU takes 3 hours, and fine tuning takes 10 minutes.

2) *Visual servoing:* HOUSTON uses Algorithm 1 to compute updates for visual servoing in both the needle presentation and the handover positioning phases. Algorithm 1 is a fixed point iteration method that uses a state estimator




Fig. 4: **Needle state estimation:** Execution of the stereo RGB pipeline on a highly visible needle. The network takes in raw stereo images as input, producing segmasks of the needle. The triangulated segmasks produce a point-cloud to which a circle is fit with RANSAC (3rd panel). Inliers are shown in blue, the best-fit circle in green, and outliers in red. The resulting observation reprojected into the left image is shown in the final panel.

to iteratively visually servo to a target state. Similar to first and second order optimization algorithms, it computes a global update based on its current state and iterates until the computed update is zero. In each iteration, it queries the current 3D state estimate of the needle and then computes an update step in the direction of the target state. To compensate for estimation errors due to challenging needle poses, this process is repeated at each iteration until the algorithm converges to a local optimum within a pose error tolerance.

During **needle acquisition**, the arm moves to a home pose, then rotates around the world  $z$  and  $x$  axes, stopping when the state estimator observes at threshold number of inlier points  $n_{\text{thresh}} = 20$  from RANSAC circle fitting. During trials, at most 2 consecutive rotations sufficed to resolve the state to this degree. After this initial acquisition, we apply Algorithm 1 to align  $c_n$  towards the left stereo camera position  $p_{\text{st,left}}^{\text{pos}}$ , with  $d$  defined as a rotation about the axis  $c_n \times (p_{\text{st,left}}^{\text{pos}} - c_p)$ . Once clearly presented to the camera, we measure  $c_t$  by choosing the inlier point from the circle fitting step which is furthest from the gripper in 3D space.

Subsequently, during **handover positioning**, we compute inverse kinematics to move  $c_p$  towards the center of the workspace with  $c_t$  pointing towards the other gripper and  $c_n$  orthogonal to the table plane. This flat configuration is critical for the grasping step, since the dVRK arm is primarily designed for top-down grasps near the center of its workspace. Because the arm only has 5 rotational degrees of freedom, we use a numerical IK solver from [3] and attempt to find a configuration minimizing rotational error within a  $6 \times 6 \times 8$  cm tolerance region on end effector translation. After moving to this pose, we repeat Algorithm 1, with  $d$  defined as a rotation aligning  $c_t$  towards the other gripper and  $c_n$  orthogonal to the table. We calculate IK to a configuration with needle curvature towards the camera and one with curvature away from the camera, then pick the configuration which minimizes rotational error to the goal.

### B. Phase 2: Executing a Grasping Policy

After the active presentation phase described in Section IV-A and the pose of the needle is relatively accurately known and accessible to the grasping arm, we execute a grasping policy  $\pi_{\text{grasp}}$  to grasp the needle. However, the needle pose estimate after the first phase may not be perfect, so we must visually servo to compensate for these errors when grasping. Even if the needle pose was perfectly known, reliable grasping of a small needle is still challenging due to the positioning

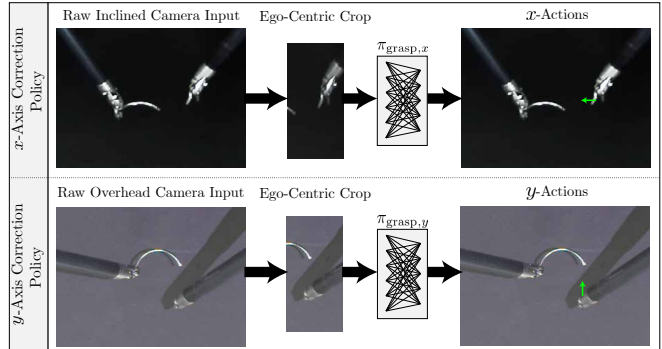


Fig. 5: **Fine-grained grasping policy:** We split corrective actions along the  $x$ - and  $y$ -axes and learn two corresponding policies  $\pi_{\text{grasp},x}$  and  $\pi_{\text{grasp},y}$ . Each policy begins by performing an ego-centric crop by projecting the gripper’s kinematically calculated approximate location into the input image and cropping around it. Then, it feeds the cropped image through a neural network to predict a correction direction.

errors of the robot, which are a result of its cable-driven arms [12, 13, 21, 24, 25, 30]. The policy splits corrective actions between the  $x$ - and  $y$ -axes with two sub-policies,  $\pi_{\text{grasp},x}$  and  $\pi_{\text{grasp},y}$ . Each policy uses RGB inputs ego-centrally cropped around the grasping arm, with the  $x$ - and  $y$ -axis policies using  $140 \times 200$  pixel crops from the inclined camera and  $70 \times 200$  pixel crops from the overhead camera respectively. The cropping forces the policy to condition based on the relative position of the gripper and needle without the ability to overfit to texture cues from other parts of the scene. The fine-grained grasping policy and image crops are displayed in Figure 5. We ablate different design choices for the grasping correction policy and also present open-loop grasping results in Section V.

The grasping subpolicy  $\pi_{\text{grasp},y}$  is a neural network classifier that outputs whether the grasping arm should move in the  $+y$  (down in the crop) or  $-y$  (up in the crop) direction.  $\pi_{\text{grasp},x}$  is trained similarly to output whether the grasping arm should move in the  $+x$  or  $-x$  directions. The policies are trained by collecting offline human demonstrations through two methods: 1) we sample poses for arms in the workspace such that needle orientation is perturbed by  $10^\circ$  about each axis, then move the robot to a good grasping position via a keyboard teleoperation interface. 2) we execute the pre-handover positioning routine and position the robot in the desired grasp location by hand, after which the robot autonomously iterates through offsets in the  $\pm x$  and  $\pm y$  directions, labeling actions according to the offset from goal position. We experimentally find that separating the policy across two axes significantly improves grasp accuracy (Section V). A separate grasping policy is trained for each arm on 100 demonstrations each. Each demonstration takes 5-10 actions and each dataset takes about an hour to collect. The policies  $\pi_{\text{grasp},x}$  and  $\pi_{\text{grasp},y}$  are each represented by voting ensembles of 5 classifiers, each of which have three convolutional layers and two fully connected layers. Details about model architectures are located in the project website.

During policy execution, we iteratively sample actions first from  $\pi_{\text{grasp},x}$ , then  $\pi_{\text{grasp},y}$ , waiting for each to converge






Fig. 6: **Acquisition Stage Rollout:** Rollout of the phase described in IV-A. Images are taken from the left camera and cropped to the gripper, with the circle observation projected in green. The pose goal is one in which the needle faces towards the camera. Note how uncertainty in the second image is resolved in later images as the needle reaches a more observable configuration.

before continuing. We multiply the action magnitude by a scalar  $\beta_{decay} = 0.5$  every time the network outputs the opposite action of the previous timestep. Servoing terminates when action magnitude decays to under 0.2mm. This enables implicit convergence to the goal without explicitly training the policy to stop. After  $x$  and  $y$  convergence, we execute a simple downward motion of 1cm to grasp the needle.

---

#### Algorithm 1 Presentation Visual Servoing Policy

---

**Require:** State estimator  $\Phi$ , target needle state  $x_N^{\text{targ}}$ , arm grasping needle  $a \in \{L, R\}$ , current time  $t$ , number of iterations  $N_{\text{max}}$ , tolerance  $\epsilon$ , distance metric  $d$ .

- 1: **for**  $i = 0; i < N_{\text{max}}; i = i + 1$  **do**
- 2:   Predict needle state  $\hat{x}_N(t + i) = \Phi(o(t + i))$
- 3:   Compute gripper update  $p_\delta = d(\hat{x}_N(t + i), x_N^{\text{targ}})$
- 4:   Update arm  $a$  pose:  $p_a(t + i + 1) = p_a(t + i) + p_\delta$
- 5:   **if**  $|p_\delta| < \epsilon$  **then**
- 6:     **break**
- 7:   **end if**
- 8: **end for**

---

## V. PHYSICAL EXPERIMENTS

The experiments aim to answer the following question: how efficient and reliable is HOUSTON compared to baseline approaches? We also perform several ablation studies of method components in this section.

### A. Baselines

To evaluate the method for the task of active needle presentation, we compare to the following baselines:

- *Depth-based presentation:* Instead of using the stereo RGB network to detect needle pose, we use a depth image-based detection algorithm to detect the needle pose and servo it to the flat grasping pose. This method takes the depth image from the built in depth calculation from the stereo camera,  $I_D(t)$ , as input, masks the gripper out of the depth image using the dVRK’s forward kinematics, then performs a volume crop around the end effector and fits a circle of known radius to the points using RANSAC to extract the state.
- *No-Sim-Data:* This is an ablation of the RGB stereo segmentation network that is only trained on the small dataset of real data.



Fig. 7: **Needles used:** The needles are shown with a coin for scale. All models were trained using needle 1, while needles 2, 3 and 4 were also used for testing. Needles 1 and 3 have a radius of 1.25cm, and needles 2 and 4 have radii of 1.75cm and 0.75cm

To evaluate the design choices used in the fine-grained grasping policy, we compare to the following baselines:

- *Open Loop:* Executes an open loop grasping motion to grasp the needle based only on needle geometry and inverse kinematics.
- *Shared  $(x, y)$  Grasp Policy:* Trains a single policy to output  $x$  and  $y$  displacements, and takes both  $I_O(t)$  and  $I_L(t)$  as input.

To evaluate whether the system can transfer to needles *unseen in training*, we evaluate HOUSTON on three additional needles as in Figure 7.

### B. Experimental Setup

We perform experiments using the daVinci Research Kit (dVRK) surgical robot [15], a cable-driven surgical robot with two needle drivers, a gripper which can open 1cm.

For perception, the setup includes a Zed Mini stereo camera angled slightly downwards to face the arms, and an overhead Zivid One Plus M camera facing directly down. Stereo images are captured at 2K resolution, and overhead images are captured at 1080p. Locations of the arms relative to each of the cameras is statically calibrated.

1) *Single handover:* For single handover experiments, we manually vary the orientation of the gripper before each trial to the orientations described in III-C. A handoff is considered successful if the needle switches from one gripper to the other, and at the end is fully supported by the other gripper.

2) *Multiple handovers:* For multiple handover experiments, we start the needle in the left gripper in a visible configuration to the camera, so that all errors are a result of handoffs rather than initialization. We evaluate two configurations: one where the needle arc ends facing the stereo camera in the grasping configuration (**Towards**), and one where it faces the opposite direction (**Away**). This configuration is typically maintained throughout each multi-handover trial because of the consistency of the needle presentation step.

### C. Single Handover Results

We evaluate HOUSTON and baselines on the single handover task in Table I, and perform multiple systematic passes over the 28 starting configurations described in Section III-C. We find that HOUSTON is able to more reliably perform the task than comparisons, which either experience many presentation errors or many grasp positioning errors.

TABLE I: **Single Handover Physical Experiments:** We report success rate, 95% confidence intervals and durations taken over a grid search of the 28 start configurations described in Section III-C for the full surgical needle bimanual regrasping task. We report the frequency of three failure modes: (**P**) error in the presentation procedure, (**X**) error along the  $x$ -axis and (**Y**) error along the  $y$ -axis. HOUSTON significantly outperforms baselines, which either have many presentation failures or grasp positioning failures. The No-Sim-Data ablation also has a high success rate, but we find that the segmentation masks are qualitatively less accurate and have more false positives in the workspace. We present results with HOUSTON using three needles (Figure 7) that were unseen in training samples. HOUSTON performs best on the larger two needles (2 and 3), and performs worse on the smaller needle 4, where occlusions with the gripper are more severe.

	SUCCESS RATE		95% CONF. INT.		COMPLETION TIME (S)	FAILURES		
	Successes / Total	% Success	Low	High		P	X	Y
OPEN LOOP	28/56	50.0	36.3	63.7	<b>14.85 ± 2.43</b>	0	0	28
SHARED ( $x, y$ ) GRASP POLICY	2/28	7.1	0.9	23.5	19.41 ± 4.23	0	0	26
NO SIM DATA	53/56	94.6	85.1	98.9	24.58 ± 3.39	3	0	0
DEPTH-BASED PRESENTATION	6/28	21.4	8.3	41.0	32.61 ± 7.42	22	0	0
HOUSTON (LEFT TO RIGHT)	<b>108/112</b>	<b>96.3</b>	91.1	99.0	21.45 ± 3.08	4	0	0
HOUSTON (RIGHT TO LEFT)	<b>108/112</b>	<b>96.3</b>	91.1	99.0	23.93 ± 3.20	1	0	3
NEEDLES UNSEEN IN TRAINING								
HOUSTON (RIGHT TO LEFT, NEEDLE 2)	26/28	92.9	76.5	99.1	23.79 ± 2.55	1	0	1
HOUSTON (RIGHT TO LEFT, NEEDLE 3)	25/28	89.3	71.8	97.7	23.15 ± 3.46	3	0	0
HOUSTON (RIGHT TO LEFT, NEEDLE 4)	21/28	75.0	55.1	89.3	23.44 ± 4.13	5	0	2

TABLE II: **Multiple Handover Physical Experiments:** To evaluate the consistency of the HOUSTON, we evaluate it on the full multiple handover HOUSTON task with maximum handoffs  $N_{\max} = 50$ . We run HOUSTON with two different needle orientations during the trial, and report the number of successful handovers (**Num**), time per handover (**Time**), and failure mode (**F**). HOUSTON averages 26.20 and 38.60 successful passes in each configuration, and has three runs with no failures.

	TRIAL 1			TRIAL 2			TRIAL 3			TRIAL 4			TRIAL 5			AVG.	
	NUM	TIME	F	NUM	TIME	F	NUM	TIME	F	NUM	TIME	F	NUM	TIME	F	NUM	T/H
Away config.	12	25.40	Y	45	25.34	P	23	25.42	P	<b>35</b>	24.71	P	16	25.28	X	26.20	25.13
Towards config.	<b>50</b>	<b>26.57</b>	-	<b>50</b>	26.76	-	<b>50</b>	27.00	-	20	25.08	Y	23	26.48	P	38.60	26.56

#### D. Multiple Handover Results

We evaluate HOUSTON on the multiple handover task with  $N_{\max} = 50$  with two different starting configurations (Table II). We observe that in the first configuration, the algorithm completes 26.20 successful handovers on average and 38.60 in the second. In three trials, no errors occur, and we manually terminate them after 50 successful handovers.

#### E. Failure Analysis

HOUSTON encounters three failure modes:

- **P:** Presentation error: the robot fails to present the needle in an orientation that is in the plane of the table with the needle tip pointing toward the grasping arm. This may lead to grasping angles that are unreachable or out of the training distribution for the grasping arm.
- **X:** Grasping positioning error (X): the  $x$ -axis grasping policy fails to line up with the needle prior to executing the  $y$ -axis grasping policy.
- **Y:** Grasping positioning error (Y): the  $y$ -axis grasping policy fails line up with the needle prior to grasping.

We categorize all of the failure modes encountered in Table I. We find that the open loop grasping policies are not able to consistently position well for grasping. HOUSTON has failures that are evenly distributed across the failure modes. Grasp policy servoing errors stem mainly from needle configurations that are far outside the distribution seen in training. Presentation phase failures stem primarily from mis-detection of the needle true tip, either because of incomplete

segmentation masks or from drift in robot kinematics causing the most distal needle point to not be the tip. This causes the servoing policy to rotate the needle away from the camera, after which sometimes it loses visibility and fails to bring the needle to the pre-handover pose. Multi-handoff failures most frequently arise because of subtle imperfections in grasp execution where the needle rotates to a difficult angle. During the subsequent handover the needle can become obstructed by the holding gripper, inhibiting the grasping policy.

## VI. DISCUSSION

In this work we present HOUSTON, a problem and an algorithm for reliably completing the bimanual regrasping task on unpainted surgical needles. To our knowledge, this work is the first to study the unmodified variant of the regrasping task. The main limitations of this approach are its reliance on human demonstrations to learn the grasping policy, and sensitivity to needle and environment appearance. We hypothesize that the former could be mitigated via self-supervised demonstration collection, or by exploring unsupervised methods for fine-tuning behavior cloned policies. Future work will address the latter issue by exploring more powerful network architectures leveraging stereo disparity such as [17], and designing more autonomous data collection techniques which can label real needle data without human input. In future work, we will also study how to reorient needles between handovers for precise control of needle-in-hand pose and attempt to make needle tracking more robust to occlusions from tissue phantoms.

## REFERENCES

- [1] E. Arruda, J. Wyatt, and M. Kopicki, "Active vision for dexterous grasping of novel objects," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2016, pp. 2881–2888.
- [2] R. Bajcsy, "Active perception," *Proceedings of the IEEE*, vol. 76, no. 8, pp. 966–1005, 1988.
- [3] P. Beeson and B. Ames, "Trac-ik: An open-source library for improved solving of generic inverse kinematics," in *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*, 2015, pp. 928–935.
- [4] B. Calli, A. Singh, J. Bruce, A. Walsman, K. Konolige, S. Srinivasa, P. Abbeel, and A. M. Dollar, "Yale-cmu-berkeley dataset for robotic manipulation research," *The International Journal of Robotics Research*, vol. 36, no. 3, pp. 261–268, 2017.
- [5] G. Caron, E. Marchand, and E. M. Mouaddib, "Photometric visual servoing for omnidirectional cameras," *Autonomous Robots*, vol. 35, no. 2, pp. 177–193, 2013.
- [6] F. Chaumette and S. Hutchinson, "Visual servo control. i. basic approaches," *IEEE Robotics & Automation Magazine*, vol. 13, no. 4, pp. 82–90, 2006.
- [7] Z.-Y. Chiu, F. Richter, E. K. Funk, R. K. Orosco, and M. C. Yip, "Bimanual regrasping for suture needles using reinforcement learning for rapid motion planning," *arXiv preprint arXiv:2011.04813*, 2020.
- [8] D.-L. Chow and W. Newman, in *Improved Knot-Tying Methods for Autonomous Robot Surgery*, IEEE Conference on Automation Science and Engineering (CASE), 2013.
- [9] C. D'Ettore, G. Dwyer, X. Du, F. Chadebecq, F. Vasconcelos, E. De Momi, and D. Stoyanov, in *Automated Pick-up of Suturing Needles for Robotic Surgical Assistance*, IEEE International Conference on Robotics and Automation (ICRA), 2018.
- [10] *Definition of superhuman by merriam webster*, <https://www.merriam-webster.com/dictionary/superhuman>, Accessed: 2021-05-24.
- [11] S. Hutchinson, G. D. Hager, and P. I. Corke, "A tutorial on visual servo control," *IEEE transactions on robotics and automation*, vol. 12, no. 5, pp. 651–670, 1996.
- [12] M. Hwang, D. Seita, B. Thananjeyan, J. Ichnowski, S. Paradis, D. Fer, T. Low, and K. Goldberg, in *Applying Depth-Sensing to Automated Surgical Manipulation with a da Vinci Robot*, International Symposium on Medical Robotics (ISMR), 2020.
- [13] M. Hwang, B. Thananjeyan, S. Paradis, D. Seita, J. Ichnowski, D. Fer, T. Low, and K. Goldberg, in *Efficiently Calibrating Cable-Driven Surgical Robots With RGBD Fiducial Sensing and Recurrent Neural Networks*, IEEE Robotics and Automation Letters (RA-L), 2020.
- [14] D. Kalashnikov, A. Irpan, P. Pastor, J. Ibarz, A. Herzog, E. Jang, D. Quillen, E. Holly, M. Kalakrishnan, V. Vanhoucke, and S. Levine, in *QT-Opt: Scalable Deep Reinforcement Learning for Vision-Based Robotic Manipulation*, Conference on Robot Learning (CoRL), 2018.
- [15] P. Kazanzides, Z. Chen, A. Deguet, G. Fischer, R. Taylor, and S. DiMaio, in *An Open-Source Research Kit for the da Vinci Surgical System*, IEEE International Conference on Robotics and Automation (ICRA), 2014.
- [16] B. Kehoe, G. Kahn, J. Mahler, J. Kim, A. Lee, A. Lee, K. Nakagawa, S. Patil, W. D. Boyd, P. Abbeel, et al., "Autonomous multilateral debridement with the raven surgical robot," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2014, pp. 1432–1439.
- [17] T. Kollar, M. Laskey, K. Stone, B. Thananjeyan, and M. Tjersland, "Simnet: Enabling robust unknown object manipulation from pure synthetic data via stereo," *arXiv preprint arXiv:2106.16118*, 2021.
- [18] D. Kragic, H. I. Christensen, et al., "Survey on visual servoing for manipulation," *Computational Vision and Active Perception Laboratory, Fiskartorpsv*, vol. 15, p. 2002, 2002.
- [19] 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," *The International Journal of Robotics Research*, vol. 38, no. 2-3, pp. 126–145, 2019.
- [20] S. Levine, P. Pastor, A. Krizhevsky, J. Ibarz, and D. Quillen, "Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection," *The International Journal of Robotics Research*, vol. 37, no. 4-5, pp. 421–436, 2018.
- [21] J. Mahler, S. Krishnan, M. Laskey, S. Sen, A. Murali, B. Kehoe, S. Patil, J. Wang, M. Franklin, P. Abbeel, and K. Goldberg, in *Learning Accurate Kinematic Control of Cable-Driven Surgical Robots Using Data Cleaning and Gaussian Process Regression.*, IEEE Conference on Automation Science and Engineering (CASE), 2014.
- [22] L. Mihaylova, T. Lefebvre, H. Bruyninckx, K. Gadeyne, and J. De Schutter, "A comparison of decision making criteria and optimization methods for active robotic sensing," in *International Conference on Numerical Methods and Applications*, Springer, 2002, pp. 316–324.
- [23] A. Murali, S. Sen, B. Kehoe, A. Garg, S. McFarland, S. Patil, W. D. Boyd, S. Lim, P. Abbeel, and K. Goldberg, in *Learning by Observation for Surgical Subtasks: Multilateral Cutting of 3D Viscoelastic and 2D Orthotropic Tissue Phantoms*, IEEE International Conference on Robotics and Automation (ICRA), 2015.
- [24] S. Paradis, M. Hwang, B. Thananjeyan, J. Ichnowski, D. Seita, D. Fer, T. Low, J. E. Gonzalez, and K. Goldberg, in *Intermittent Visual Servoing: Efficiently Learning Policies Robust to Instrument Changes for High-precision Surgical Manipulation*, IEEE International Conference on Robotics and Automation (ICRA), 2021.
- [25] H. Peng, X. Yang, Y.-H. Su, and B. Hannaford, in *Real-time Data Driven Precision Estimator for RAVEN-II Surgical Robot End Effector Position*, IEEE International Conference on Robotics and Automation (ICRA), 2020.
- [26] F. Richter, S. Shen, F. Liu, J. Huang, E. K. Funk, R. K. Orosco, and M. C. Yip, in *Autonomous Robotic Suction to Clear the Surgical Field for Hemostasis using Image-based Blood Flow Detection*, arXiv preprint arXiv:2010.08441, 2020.
- [27] J. Rosen and J. Ma, in *Autonomous Operation in Surgical Robotics*, vol. 137, Mechanical Engineering, 2015.
- [28] H. Saeidi, H. N. D. Le, J. D. Opfermann, S. Leonard, A. Kim, M. H. Hsieh, J. U. Kang, and A. Krieger, in *Autonomous Laparoscopic Suturing with a Novel Actuated Suturing Tool and 3D Endoscope*, IEEE International Conference on Robotics and Automation (ICRA), 2019.
- [29] P. Salaris, R. Spica, P. R. Giordano, and P. Rives, "Online optimal active sensing control," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2017, pp. 672–678.
- [30] D. Seita, S. Krishnan, R. Fox, S. McKinley, J. Canny, and K. Goldberg, in *Fast and Reliable Autonomous Surgical Debridement with Cable-Driven Robots Using a Two-Phase Calibration Procedure*, IEEE International Conference on Robotics and Automation (ICRA), 2018.
- [31] S. Sen, A. Garg, D. V. Gealy, S. McKinley, Y. Jen, and K. Goldberg, in *Automating Multiple-Throw Multilateral Surgical Suturing with a Mechanical Needle Guide and Sequential Convex Optimization*, IEEE International Conference on Robotics and Automation (ICRA), 2016.
- [32] P. Sundaresan, B. Thananjeyan, J. Chiu, D. Fer, and K. Goldberg, in *Automated Extraction of Surgical Needles from Tissue Phantoms*, IEEE Conference on Automation Science and Engineering (CASE), 2019.
- [33] B. Thananjeyan, A. Balakrishna, U. Rosolia, F. Li, R. McAllister, J. E. Gonzalez, S. Levine, F. Borrelli, and K. Goldberg, in *Safety Augmented Value Estimation from Demonstrations (SAVED): Safe Deep Model-Based RL for Sparse Cost Robotic Tasks*, IEEE Robotics and Automation Letters (RA-L), 2020.
- [34] B. Thananjeyan, A. Garg, S. Krishnan, C. Chen, L. Miller, and K. Goldberg, in *Multilateral Surgical Pattern Cutting in 2D Orthotropic Gauze with Deep Reinforcement Learning Policies for Tensioning*, IEEE International Conference on Robotics and Automation (ICRA), 2017.
- [35] V. M. Varier, D. K. Rajamani, N. Goldfarb, F. Tavakkolmoghaddam, A. Munawar, and G. S. Fischer, "Collaborative suturing: A reinforcement learning approach to automate hand-off task in suturing for surgical robots," in *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, IEEE, 2020, pp. 1380–1386.
- [36] S. D. Whitehead and D. H. Ballard, "Active perception and reinforcement learning," in *Machine Learning Proceedings 1990*, Elsevier, 1990, pp. 179–188.
- [37] M. Yip and N. Das, in *Robot Autonomy for Surgery*, The Encyclopedia of Medical Robotics, 2017.