

Early Failure Detection in Autonomous Surgical Soft-Tissue Manipulation via Uncertainty Quantification

Jordan Thompson, Ronald Koe, Anthony Le, Gabriella Goodman, Daniel S. Brown, and Alan Kuntz




Fig. 1: Overall pipeline and experimental setup. We sense a point cloud of the current geometry of the tissue using the Segment Anything Model (SAM) [23] to segment RGBD images of the scene. SAM takes as input the RGB components of the image along with a query point on the desired segmented object. The resulting segmentation mask is then converted into a point cloud using the depth component of the image along with the camera’s intrinsic matrix to give the three-dimensional geometry of the tissue. The current geometry, coupled with the goal geometry and manipulation point, are passed as input to a deep ensemble of DeformerNet [50] models to predict a robot action $\hat{A}$ (homogeneous transformation matrix) and an uncertainty estimate u . Using u , we determine if it is safe for the daVinci Research Kit (dVRK) robot to take action $\hat{A}$. If yes, the robot takes action $\hat{A}$, and the process repeats. Otherwise, the system requests a human to intervene and complete the task. We perform human interventions using the surgeon-side dVRK manipulators for teleoperation.

Abstract—Autonomous surgical robots are a promising solution to the increasing demand for surgery amid a shortage of surgeons. Recent work has proposed learning-based approaches for the autonomous manipulation of soft tissue. However, due to variability in tissue geometries and stiffnesses, these methods do not always perform optimally, especially in out-of-distribution settings. We propose, develop, and test the first application of uncertainty quantification to learned surgical soft-tissue manipulation policies as an early identification system for task failures. We analyze two different methods of uncertainty quantification, deep ensembles and Monte Carlo dropout, and find that deep ensembles provide a stronger signal of future task success or failure. We validate our approach using the physical daVinci Research Kit (dVRK) surgical robot to perform physical soft-tissue manipulation. We show that we are able to successfully detect out-of-distribution states leading to task failure and request human intervention when necessary while still enabling autonomous manipulation when possible. Our learned tissue manipulation policy with uncertainty-based early failure detection achieves a zero-shot sim2real performance improvement of 47.5% over the prior state of the art in learned soft-tissue manipulation. We also show that our method generalizes well to new types of tissue as well as to a bimanual soft-tissue manipulation task.

I. INTRODUCTION

Autonomous surgical robots have the potential to help solve the growing disparity between the population’s need for surgery and the number of available surgeons [55, 5]. However, surgical robot learning and automation is particularly challenging due to the nuanced and risk-sensitive nature of the tasks, the partially observable and deformable nature of the environment, and the scarcity of available data. In particular, because autonomous surgical system failures can be detrimental to patient health, it is crucial that these autonomous systems take into account uncertainty so that the system can safely cede control to the surgeon before causing any harm. Thus, we aim to develop an autonomous surgical soft-tissue manipulation system that is capable of reasoning over predictive uncertainty to detect task failure early in the robot’s operation (see Fig. 1).

In doing so, our goal is to mitigate the risk of system failures while still offloading the work of soft-tissue manipulation to the robot, whenever it is safe to do so.

Our main goal is to enable uncertainty-based task failure prediction in the domain of soft-tissue manipulation. To this end, we build on recent advances in robot learning for deformable object manipulation. In particular, we focus on the recently proposed, state-of-the-art DeformerNet framework [50] which uses large-scale self-supervised training in simulation to learn a closed-loop policy for manipulating deformable objects to desired goal geometries. While DeformerNet has been shown to perform well on in-distribution tissue, inter-human variability can lead to erroneous outputs, which can pose a safety risk. We propose to address this issue through quantifying the predictive uncertainty of DeformerNet.

One of the major obstacles when applying robot learning to surgical domains is the inability to do online learning in the real world. While we can collect demonstrations offline [40, 51, 12, 54] and can train policies in simulation [56, 36, 21, 46], offline imitation learning leads to compounding errors [43] and the sim2real gap is especially difficult to overcome for the types of deformable manipulation that are common in surgery [16]. Enabling a surgical robot to preemptively detect task failures addresses these challenges by enabling a human to correct the robot during real-world execution, thus preventing failures while also enabling the robot to potentially use human interventions to continually improve its policy and world model. While there has been prior work on interactive imitation learning and human interventions [22, 53, 48, 30, 18, 8], to the best of our knowledge, we are the first to study robot-requested interventions on the kinds of complex deformable object manipulation tasks common in surgical robotics.

Our main contributions are as follows: (1) We develop and study the first real-world surgical robotic system capable of uncertainty quantification and early task failure detection

during soft-tissue manipulation. (2) We extend and analyze two potential methods for quantifying predictive uncertainty during learned soft-tissue manipulation: deep ensembles and Monte Carlo dropout. We provide empirical evidence that ensembles provide more accurate uncertainty estimates and that using the slope of the ensemble variance rather than the raw variance more accurately discriminates between successful and unsuccessful tissue manipulation actions. (3) Across 40 real-world tissue manipulation trials, we show that our approach improves zero-shot sim2real performance of more than 47% when compared with DeformerNet by itself, the prior state-of-the-art approach for deformable tissue manipulation. Concretely, our method achieves a 95% system success rate compared with DeformerNet achieving 47.5% system success rate without our method. Further, our learned thresholds for failure detection successfully generalize to out-of-distribution tissue and to new tasks unseen during training and calibration.

II. RELATED WORK

A. Soft Tissue Manipulation

Much work has investigated automated tissue manipulation during surgery [4]. Several data-driven approaches have been proposed to learn soft tissue manipulation [41, 47, 38, 42, 50]. These learning-based approaches are made possible by taking advantage of recent advancements in high-fidelity deformable object simulation [9, 19, 28]. Other work has shown success in using model-independent deformation estimation techniques [1]. Our method primarily draws on the recently proposed DeformerNet framework [50]. DeformerNet takes a self-supervised learning approach to the problem of soft tissue manipulation. Given the current geometry of the object, the grasping point, and the desired geometry of the object, DeformerNet has been shown to have state-of-the-art performance for a variety of deformable object geometries and tasks [50]. However, as we show in our experiments, DeformerNet sometimes fails due to out-of-distribution inputs or bad grasp points. By quantifying the uncertainty in DeformerNet’s predictions, we enable informed early failure detection, taking an important step toward safe autonomy in soft-tissue surgical manipulation tasks.

B. Uncertainty Quantification

Uncertainty quantification for deep neural networks has been studied in recent years as a way to aid in the safe design and implementation of deep learning systems [15]. We study and compare the efficacy of ensembles and Monte Carlo dropout for uncertainty quantification for deformable robot manipulation tasks with a focus on soft-tissue manipulation. Ensembles are an approach by which multiple models are used in concert to make predictions and have been shown to improve overall model performance [27, 13]. Uncertainty is typically quantified using ensembles by measuring the variance in the ensemble’s outputs [37, 25].

Monte Carlo dropout is another technique utilized for uncertainty estimation in deep neural networks [33, 6]. By extending dropout to the inference phase and making multiple

inference passes through the network with different dropout instantiations, Monte Carlo dropout can be used to generate a distribution of outputs. Unlike ensembles, which require training multiple independent models, Monte Carlo dropout leverages the same trained model for inference, eliminating the need for repeated training [29].

By analyzing the distribution from an ensemble or Monte Carlo dropout, uncertainty metrics such as entropy or variance across predictions can be computed, providing insights into the model’s confidence and reliability [45, 31].

C. Human Interventions for Robot Policies

Offline learning from human demonstrations is a common approach for robot learning [40, 17, 2, 34, 3, 20, 35]; however, while learning from offline data has shown success in some domains [10, 49, 14, 7], executing policies learned from offline data is known to lead to suboptimal behavior due to compounding errors [43, 26], something that is unacceptable in surgical robotics. To remedy this problem, it is common practice to give control to a human supervisor who can provide an intervention. There are two strategies for deciding when to pass control from the autonomous agent to the human. In the first paradigm, the human decides when to intervene [22, 48, 30]. However, this imposes a large burden on the supervisor since it requires a human to continuously monitor the robot and relies on the human accurately predicting what the robot will do next. In the second paradigm, the robot actively requests human interventions [32, 53, 18, 8]. Similar to our approach, prior methods of this type typically use some form of novelty or uncertainty estimation when deciding to cede control to the human. However, prior work has focused mainly on simulated tasks with simulated human supervisors and simple control or manipulation tasks. By contrast, we study the efficacy of uncertainty quantification when deployed on a surgical dVRK robot performing complex, deformable tissue manipulation in the real world with access to real human interventions.

III. PROBLEM DEFINITION

Given a learned surgical robot policy, π_{robot} , for soft-tissue manipulation, we wish to develop an early identification system to determine whether the robot should execute its policy autonomously or request an intervention from a human supervisor. Following prior work on deformable tissue manipulation [50], we assume access to a partial-view point cloud of the current geometry of the soft tissue to be manipulated as well as a goal point cloud that defines how the soft-tissue should be manipulated.

The meta policy π_{meta} chooses whether to execute the learned robot policy or to hand over control to a human supervisor and can thus be represented as a binary classifier, $\pi_{\text{meta}} : \mathcal{O} \rightarrow \{0, 1\}$, where $\mathcal{O}$ is the observation space. A true positive (TP) is when the robot correctly requests an intervention and would have failed if it had operated autonomously, a false positive (FP) is when the robot wrongly asks for an intervention when it would have succeeded autonomously, a false negative (FN) is when the robot does not ask for an

intervention and fails at the task, and a true negative (TN) is when the robot does not request an intervention and succeeds at the task autonomously.

Because we are focused on surgical robotic applications, we are most concerned with failures. Thus, we want to minimize failures associated with false negatives. However, this is not the only source of potential failures—we also want to minimize false positives. False positives may also lead to failure because they unnecessarily distract a human expert and waste time that could be spent addressing other concerns (e.g., assisting a different patient or robot) and preventing other failures. Thus, given a cost of failure c_f , we want to find the following meta policy:

$$\pi_{\text{meta}}^* = \arg \min_{\pi_{\text{meta}}} \mathbb{E}[c_f \cdot FN + c_f \cdot P(\text{failure}|\text{FP}) \cdot \text{FP}], \quad (1)$$

where $P(\text{failure}|\text{FP})$ is the probability of a failure occurring somewhere else due to the human supervisor being unnecessarily distracted or burdened. Since $c_f > 0$, we can just minimize $FN + P(\text{failure}|\text{FP}) \cdot \text{FP}$. We do this by learning the optimal classification threshold for π_{meta} using a calibration set [11, 24, 39] to appropriately tradeoff between FN and FP. Note that the appropriate tradeoff will depend on the particular context and will be problem dependent.

IV. METHODOLOGY

A. DeformerNet

For this work, we modify and endow the DeformerNet framework [50] with uncertainty awareness and the ability to preemptively detect task failures conditioned on that uncertainty. DeformerNet is a deep learning based shape-servoing model for manipulating deformable objects. The model takes as input the current geometry P_c and goal geometry P_g of the deformable object as partial-view point clouds along with the manipulation point $\mathbf{m}$ on the object. The model then predicts an action $\hat{A}$ as a homogeneous transformation matrix for the robot’s end-effector, $\hat{A} = \begin{bmatrix} \hat{R} & \hat{\mathbf{p}} \\ \mathbf{0} & 1 \end{bmatrix}$, where $\hat{R}$ and $\hat{\mathbf{p}}$ are the predicted change in end-effector orientation and position respectively. DeformerNet is a closed loop control method that reasons over the deformable object’s geometry and the goal shape’s geometry leveraging encoders based on PointConv [52]. DeformerNet is trained in a self-supervised fashion using large-scale simulated data. In simulation, given a manipulation point $\mathbf{m}$ on a deformable object with current geometry P_c , a random homogeneous transformation matrix A is applied to the robot’s end effector. The geometry of the deformable object is then sensed and is used as a self-supervised goal, P_g . DeformerNet is trained to predict the homogeneous transformation that led to P_g over a data set consisting of instances of the following form: $(P_c, P_g, \mathbf{m}, A)$.

B. Uncertainty quantification

We aim to augment the DeformerNet model with uncertainty quantification to enable early task failure detection. To this end, we propose the use of a deep ensemble consisting of

multiple DeformerNet models. Each model predicts a change in position and orientation of the robot’s end effector that will deform the object being manipulated toward the goal shape. At inference time, the positional component $\hat{\mathbf{p}}$ of the deep ensemble’s prediction is computed by averaging the positional component of the individual models,

$$\hat{\mathbf{p}} = \frac{1}{N} \sum_{i=1}^N \hat{\mathbf{p}}_i, \quad (2)$$

where N is the number of models in the ensemble, and $\hat{\mathbf{p}}_i$ is the positional output of model i .

At inference time, the orientation component of the deep ensemble’s output is found by first computing the arithmetic mean S of the set of rotation matrices produced by the individual models of the ensemble, $S = \frac{1}{N} \sum_{i=1}^N \hat{R}_i$, where N is the number of models, and $\hat{R}_i$ is the rotation matrix output by model i . Matrix S , however, will almost certainly not satisfy the constraints required to form a valid rotation matrix. To obtain a valid rotation matrix, we take the singular value decomposition, $S = UDV'$ and multiply U and V' to obtain a rotation matrix that minimizes the Euclidean norm to S and is therefore the rotation matrix that minimizes the average geodesic distance to each component rotation matrix [44]:

$$\hat{R} = UV'. \quad (3)$$

We compute the ensemble’s predictive variance for both position and rotation individually. Positional variance is the average squared distance to the deep ensemble’s positional prediction ($\hat{\mathbf{p}}$ in Equation (2)) for each component network:

$$\text{Var}(\hat{\mathbf{p}}) = \frac{1}{N} \sum_{i=1}^N \|\hat{\mathbf{p}}_i - \hat{\mathbf{p}}\|^2. \quad (4)$$

For rotational variance, we measure the average geodesic distance to the ensemble’s predicted rotation matrix ($\hat{R}$ in Equation (3)):

$$\text{Var}(\hat{R}) = \frac{1}{N} \sum_{i=1}^N \arccos\left(\frac{\text{Tr}(\hat{R}_i \hat{R}) - 1}{2}\right). \quad (5)$$

We define the predictive uncertainty of DeformerNet at time-step t as a vector containing the change in both positional and rotational variances from equations (4) and (5).

$$\mathbf{u}(t) = [\text{Var}(\hat{\mathbf{p}}_t) - \text{Var}(\hat{\mathbf{p}}_{t-1}), \text{Var}(\hat{R}_t) - \text{Var}(\hat{R}_{t-1})] \quad (6)$$

While standard approaches to uncertainty quantification use raw variance values as their uncertainty metric [18, 15, 25], we find in our experiments that these raw variance values are not good indicators for downstream task performance. However, when we measure the slopes of these variance values over subsequent time steps, we find a significantly stronger signal of downstream task performance.




Fig. 2: Distributions of raw variance values (computed via Equations 4 and 5) and initial slopes (computed via Equation 6 at $t = 1$) for both position and rotation across 30 trials at inference time on the dVRK surgical robot. We find that the distributions of raw variance values for position and rotation across successful and failure trials only have Kullback-Leibler (KL) divergences of 0.357 and 0.507 respectively. However, the initial slopes of the variance values have KL-divergence values of 3.107 and 4.078 respectively. We set a threshold value on the slope of the initial positional and rotational variance values at -0.310 and -0.487 respectively as indicated by the red lines in the graphs.

C. Monte Carlo Dropout

We also investigate uncertainty quantification using Monte Carlo dropout. Using stochastic dropout at inference time allows us to sample from a distribution of potential changes in end-effector pose. We then use equations (2) – (6) to get the model predictions and variances.

V. PHYSICAL EXPERIMENTAL SETUP

We train a deep ensemble DeformerNet consisting of five models. We use the source code and training and testing datasets provided by Thach et al. [50], consisting of 11,566 training and 1,285 test examples of manipulations of a deformable box object. Following prior work [27, 25, 13], each ensemble component model was trained using the same training dataset with different random weight initializations. We train a Monte Carlo dropout version of the DeformerNet model with dropout added to the last layer at varying levels of dropout percentage: 25%, 50%, and 75%. We collect 100 samples from the model to construct the predictive distributions.

The deep ensemble and the Monte Carlo dropout models are implemented using a zero-shot sim2real framework. Trained entirely on a simulated box-shaped deformable object, the system is tasked with manipulating both ex vivo chicken and bovine tissue of varying geometries. To generate goal shapes for evaluation purposes only, using ex vivo tissue we teleoperated the robot to manually manipulate the tissue to a desired goal geometry. We then reset the system and task the models with manipulating the tissue to the same desired geometry (with no knowledge of how the shape was generated). We use an Intel Realsense D405 camera for tracking point cloud representations of the tissue geometry both in goal generation for evaluation and during method execution. We track the model variance as it manipulates the tissue toward the goal shape and measure task success as whether the model converges to the desired geometry. We define a method termination criteria as when $\|\hat{p}\| < 0.001$. At this point, the method has effectively stopped moving the robot’s end effector. We then compare the resulting tissue geometry with the desired geometry to test for task success. Figure 1 shows an image of the experimental setup. We use the patient-side dVRK manipulator(s) to manipulate the tissue both autonomously and by teleoperation.

VI. PHYSICAL SOFT-TISSUE MANIPULATION RESULTS

A. DeformerNet Performance Validation

a) *In-Distribution Performance:* In prior work, DeformerNet has been shown to be capable of achieving 100% success rates on in-distribution cases [50]. To validate that our ensemble is capable of achieving similar performance, we performed 20 in-distribution trials on ex vivo chicken muscle tissue with optimal grasping points and goal geometries that only required local control. Our ensemble successfully manipulated the tissue in all 20 in-distribution trials. We also performed this same experiment across 10 in-distribution cases using the Monte Carlo dropout model. We find that the Monte Carlo dropout version of DeformerNet fails to complete the task in all cases. Thus, applying Monte Carlo dropout at inference time severely hinders the performance of DeformerNet. As a result, the following experiments are performed solely using the deep ensemble.

b) *Out-of-Distribution Performance:* While DeformerNet has 100% task success on in-distribution cases, out-of-distribution cases frequently result in task failure. We performed 15 trials spanning 3 different out-of-distribution cases: sub-optimal grasping points, non-local control, and out-of-distribution tissue geometries. Sub-optimal grasping points can occur when the robot grasps the tissue at a location that does not optimally facilitate the manipulation of the tissue to the goal. Non-local control cases occur when the requested goal geometry requires more than one type of manipulation in sequence, i.e. folding the tissue which requires first lifting an edge then subsequently placing it back down elsewhere. Out-of-distribution tissue geometries occur when the geometry of the tissue being manipulated does not align with the tissue geometries seen during model training. Our ensemble only succeeded in 1 out-of-distribution geometry case while failing in the remaining 14 cases. Due to the high variability and lack of available data on surgical soft-tissue manipulation, these out-of-distribution cases will inevitably occur in practice. Thus, there is a great need to detect out-of-distribution cases prior to the manipulation of the tissue.

B. Uncertainty Quantification

We collected 30 trials containing both in-distribution and out-of-distribution cases of DeformerNet on ex vivo chicken

	Intervention Requests											
	Chicken Tissue						Bovine Tissue					
	Positional		Rotational		Both		Positional		Rotational		Both	
	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No
Intervention Needed	10	1	9	2	10	1	9	1	10	0	10	0
Intervention Not Needed	3	6	5	4	6	3	1	9	3	7	3	7

TABLE I: Results from using variance slope thresholds for requests for human intervention on ex vivo chicken muscle tissue and ex vivo bovine tissue across 40 trials. Placing a threshold on the change in positional variance gives an accuracy of 85% across ex vivo chicken and bovine tissue with a false positive rate of 21.1% and a false negative rate of 9.5%. Using rotational variance, we achieve an accuracy of 75% with a false positive rate of 42.1% and a false negative rate of 9.5%. Simultaneously thresholding position and rotation yields an accuracy of 75% with a false positive rate of 47.4% and a false negative rate of 4.8%.




Fig. 3: False positive and negative rates on 20 ex vivo chicken tissue trials for various positional and rotational slope thresholds. As the thresholds are raised, the false positive rates decrease while the false negative rates increase.

muscle tissue and measured the predictive variances across the trajectories using Equations 4 and 5. Figure 2 shows the distributions of these raw variance values for successful and unsuccessful trials. We find that raw variance values are not good indicators of task success, as there is no clear distinction between the successful and unsuccessful trials.

Using our uncertainty definition from Equation 6, however, we find that the slopes in the variance values at $t = 1$ provide a strong predictive signal of task success. We use Kullback-Leibler (KL) divergence as a quantitative metric of distribution difference. In the case of raw variance values, we observe KL divergences of 0.357 and 0.507 for position and rotation respectively. However, the slopes in variance at $t = 1$ give KL divergences of 3.107 and 4.078 for position and rotation respectively. These results demonstrate that in successful cases, the positional and rotational variances tend to decrease moreso than in failure cases at the initial time step.

a) Calibration of the Meta Policy: To detect task failures, we set a threshold on the slope of the positional and rotational variances at time-step 1 in the trajectories ($u(1)$ from Eq. 6) by calibrating to desired false positive and negative rates on the above 30 trials (false positives are detecting task failure in successful cases, false negatives are failing to detect task failure when necessary). False positive and negative rates for different threshold values on 20 test trials can be seen in Fig. 3.

b) Evaluation on Chicken Tissue: Placing thresholds at -0.310 and -0.487 for the positional and rotational components

of $u(1)$ respectively, we ran 20 trials on ex vivo chicken muscle tissue. The system detects task failure when the specified threshold is violated. Once the system detects task failure, the human teleoperates the remainder of the task. The results of these trials are summarized in Table I. When only using the positional threshold, we successfully classified the success of the task in 16 out of the 20 trials. In only 1 out of the 20 trials we did not detect a task failure. When only using the rotational threshold, we successfully classified task success in 13 out of 20 trials while not detecting task failure in 2 of the trials. When using a threshold on both position and rotation simultaneously, we successfully classify task success in 13 trials and do not detect task failure in 1 trial.

c) Generalization to Bovine Tissue: To test the generalization of our method to new tissue types, we ran 20 trials on ex vivo bovine tissue with no re-calibration of the chosen thresholds. The results for the bovine tissue are summarized in Table I. We successfully classify task completion in 18 and 17 of the trials for positional and rotational intervention respectively while only failing to detect a task failure in 1 of the positional intervention trials. When using both thresholds, we successfully classify task completion in 17 trials and never fail to detect a task failure.

Across all 40 physical trials, the baseline DeformerNet failed to reach the desired geometry in 52.5% of trials. With the introduction of uncertainty thresholds, we successfully reduced the system failure rate to 5% while using either positional or rotational thresholds constituting a 47.5% improvement over base DeformerNet while still allowing autonomous manipulation in approximately 78.9% of the trials for position, 57.9% of the trials for rotation, and 52.6% of the trials for both. Figure 4 shows examples of both successful and failed manipulations with and without human intervention. We see that we are successfully able to enable autonomous tissue manipulation when appropriate while preventing system failures by requesting human intervention when necessary.

d) Task Generalization: We also tested our method’s ability to generalize to new tasks by applying our thresholds to a bimanual shape-servoing task. In this task, the left patient-side manipulator of the dVRK is always teleoperated by the human while the right patient-side manipulator is controlled via our DeformerNet ensemble. The human operator is given control of the automated arm when an uncertainty threshold is




Fig. 4: Example soft-tissue manipulations with and without requests for human intervention. The images show the view of the RealSense camera at subsequent steps of the trajectory. The point clouds show the corresponding current (black) and goal (red) geometries of the tissue while the grasp point on the tissue is shown as a blue dot. (Top) A successful manipulation where human intervention was not requested. (Middle) A successful manipulation where human intervention was requested. (Bottom) A failed manipulation where the trajectory from the middle panel is allowed to complete autonomously instead of having a human intervene.

violated. Similarly to the single-arm task, example bimanual soft-tissue manipulations can be seen in Fig. 5. Across 10 trials with positional interventions, we achieved a 100% system success rate with a false positive rate of 66.6% and a false negative rate of 0% (7 requested interventions while only 1 intervention was necessary). Using rotational interventions, we again achieved a 100% system success rate with 0% false positive and negative rates (1 requested intervention in a case where it was necessary). These results show that while our rotational threshold was transferable to this new task with no necessary recalibration, our positional threshold is overly conservative in this new task and would need to be recalibrated to achieve similar performance to the single-arm task.

VII. CONCLUSION

Robot learning for the autonomous completion of surgical tasks is a particularly challenging domain due to data




Fig. 5: Example bimanual soft-tissue manipulations. The images show the view of the RealSense camera at subsequent steps of the trajectory. The point clouds show the corresponding current (black) and goal (red) geometries of the tissue while the grasp point on the tissue is shown as a blue dot. The left patient-side manipulator of the dVRK is always teleoperated by a human while the right patient-side manipulator is autonomously controlled by DeformerNet; however, the right patient-side manipulator can request a human to teleoperate when needed. (Top) A successful manipulation where human intervention was not requested. (Middle) A successful manipulation where human intervention was requested. (Bottom) A failed manipulation where the trajectory from the middle panel was instead allowed to complete the task autonomously despite requesting human intervention.

scarcity, out-of-distribution observations, and safety concerns. To address this challenge, we introduce a novel framework for the safe implementation of an autonomous soft-tissue manipulation system for uncertainty quantification and early failure detection. We apply our method on the dVRK surgical robot system for a soft-tissue manipulation task using both in- and out-of-distribution ex-vivo tissue types and demonstrate with high accuracy the ability to request interventions in cases that would have otherwise failed. In the future we plan to study bi-directional human-robot hand-offs during soft tissue manipulation as well as study when humans feel confident ceding control back to the robot and how to visualize model uncertainty to enhance interpretability and trust.

REFERENCES

- [1] Farshid Alambeigi, Zerui Wang, Yun-hui Liu, Russell H Taylor, and Mehran Armand. Toward semi-autonomous cryoablation of kidney tumors via model-independent deformable tissue manipulation technique. *Annals of biomedical engineering*, 46:1650–1662, 2018.
- [2] Brenna D Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot learning from demonstration. *Robotics and autonomous systems*, 57(5):469–483, 2009.
- [3] Saurabh Arora and Prashant Doshi. A survey of inverse reinforcement learning: Challenges, methods and progress. *arXiv preprint arXiv:1806.06877*, 2018.
- [4] Aleks Attanasio, Bruno Scaglioni, Matteo Leonetti, Alejandro F Frangi, William Cross, Chandra Shekhar Biyani, and Pietro Valdastrì. Autonomous tissue retraction in robotic assisted minimally invasive surgery—a feasibility study. *IEEE Robotics and Automation Letters*, 5(4):6528–6535, 2020.
- [5] Aleks Attanasio, Bruno Scaglioni, Elena De Momi, Paolo Fiorini, and Pietro Valdastrì. Autonomy in surgical robotics. *Annual Review of Control, Robotics, and Autonomous Systems*, 4:651–679, 2021.
- [6] Robin Camarasa, Daniel Bos, Jeroen Hendrikse, Paul Nederkoorn, Eline Kooi, Aad Van Der Lugt, and Marleen De Bruijne. Quantitative comparison of monte-carlo dropout uncertainty measures for multi-class segmentation. In *Uncertainty for Safe Utilization of Machine Learning in Medical Imaging, and Graphs in Biomedical Image Analysis: Second International Workshop, UNSURE 2020, and Third International Workshop, GRAIL 2020, Held in Conjunction with MICCAI 2020, Lima, Peru, October 8, 2020, Proceedings 2*, pages 32–41. Springer, 2020.
- [7] Felipe Codevilla, Eder Santana, Lopez Antonio M., and Adrien Gaidon. Exploring the limitations of behavior cloning for autonomous driving. *International Conference on Computer Vision*, 2019.
- [8] Gaurav Datta, Ryan Hoque, Anrui Gu, Eugen Solowjow, and Ken Goldberg. Iifl: Implicit interactive fleet learning from heterogeneous human supervisors. In *Conference on Robot Learning*, pages 2340–2356. PMLR, 2023.
- [9] Kenny Erleben. Non-smooth newton methods for deformable multi-body dynamics. 2019.
- [10] Chelsea Finn, Tianhe Yu, Tianhao Zhang, Pieter Abbeel, and Sergey Levine. One-shot visual imitation learning via meta-learning. *Conf. on Robot Learning (CoRL)*, 2017.
- [11] Peter A Flach. Roc analysis. In *Encyclopedia of machine learning and data mining*, pages 1–8. Springer, 2016.
- [12] Pete Florence, Corey Lynch, Andy Zeng, Oscar A Ramirez, Ayzaan Wahid, Laura Downs, Adrian Wong, Johnny Lee, Igor Mordatch, and Jonathan Tompson. Implicit behavioral cloning. In *Conference on Robot Learning*, pages 158–168. PMLR, 2022.
- [13] Stanislav Fort, Huiyi Hu, and Balaji Lakshminarayanan. Deep ensembles: A loss landscape perspective. *arXiv preprint arXiv:1912.02757*, 2019.
- [14] Aditya Ganapathi, Priya Sundaresan, Brijen Thananjeyan, Ashwin Balakrishna, Daniel Seita, Jennifer Grannen, Minho Hwang, Ryan Hoque, Joseph E Gonzalez, Nawid Jamali, Katsu Yamane, Soshi Iba, and Ken Goldberg. Learning Dense Visual Correspondences in Simulation to Smooth and Fold Real Fabrics. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2021.
- [15] Jakob Gawlikowski, Cedric Rovee Njietcheu Tassi, Mohsin Ali, Jongseok Lee, Matthias Humt, Jianxiang Feng, Anna Kruspe, Rudolph Triebel, Peter Jung, Ribana Roscher, et al. A survey of uncertainty in deep neural networks. *Artificial Intelligence Review*, 56(Suppl 1): 1513–1589, 2023.
- [16] Mustafa Haiderbhai, Radian Gondokaryono, Thomas Looi, James M Drake, and Lueder A Kahrs. Robust sim2real transfer with the da vinci research kit: A study on camera, lighting, and physics domain randomization. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3429–3435. IEEE, 2022.
- [17] Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. In *Advances in Neural Information Processing Systems*, 2016.
- [18] Ryan Hoque, Ashwin Balakrishna, Ellen Novoseller, Albert Wilcox, Daniel S Brown, and Ken Goldberg. Thriftydagger: Budget-aware novelty and risk gating for interactive imitation learning. *arXiv preprint arXiv:2109.08273*, 2021.
- [19] Isabella Huang, Yashraj Narang, Clemens Eppner, Balakumar Sundaralingam, Miles Macklin, Tucker Hermans, and Dieter Fox. Defgraspsim: Simulation-based grasping of 3d deformable objects. *arXiv preprint arXiv:2107.05778*, 2021.
- [20] Auke Jan Ijspeert, Jun Nakanishi, Heiko Hoffmann, Peter Pastor, and Stefan Schaal. Dynamical movement primitives: learning attractor models for motor behaviors. *Neural computation*, 25(2), 2013.
- [21] Manuel Kaspar, Juan D Muñoz Osorio, and Jürgen Bock. Sim2real transfer for reinforcement learning without dynamics randomization. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4383–4388. IEEE, 2020.
- [22] Michael Kelly, Chelsea Sidrane, Katherine Driggs-Campbell, and Mykel J. Kochenderfer. HG-Dagger: Interactive imitation learning with human experts. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2019.
- [23] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023.
- [24] Oluwasanmi O Koyejo, Nagarajan Natarajan, Pradeep K

- Ravikumar, and Inderjit S Dhillon. Consistent binary classification with generalized performance metrics. *Advances in neural information processing systems*, 27, 2014.
- [25] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.
- [26] Michael Laskey, Jonathan Lee, Roy Fox, Anca Dragan, and Ken Goldberg. DART: Noise injection for robust imitation learning. In *Conf. on Robot Learning (CoRL)*, 2017.
- [27] Stefan Lee, Senthil Purushwalkam, Michael Cogswell, David Crandall, and Dhruv Batra. Why m heads are better than one: Training a diverse ensemble of deep networks. *arXiv preprint arXiv:1511.06314*, 2015.
- [28] Jacky Liang, Viktor Makoviychuk, Ankur Handa, Nuttapong Chentanez, Miles Macklin, and Dieter Fox. Gpu-accelerated robotic simulation for distributed reinforcement learning. In *Conference on Robot Learning*, pages 270–282. PMLR, 2018.
- [29] Yuki Mae, Wataru Kumagai, and Takafumi Kanamori. Uncertainty propagation for dropout-based bayesian neural networks. *Neural Networks*, 144:394–406, 2021.
- [30] Ajay Mandlekar, Danfei Xu, Roberto Martín-Martín, Yuke Zhu, Li Fei-Fei, and Silvio Savarese. Human-in-the-loop imitation learning using remote teleoperation, 2020.
- [31] José Mena, Oriol Pujol, and Jordi Vitria. A survey on uncertainty estimation in deep learning classification systems from a bayesian perspective. *ACM Computing Surveys (CSUR)*, 54(9):1–35, 2021.
- [32] Kunal Menda, Katherine Driggs-Campbell, and Mykel J. Kochenderfer. EnsembleDagger: A Bayesian Approach to Safe Imitation Learning. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2019.
- [33] Daily Milanés-Hermosilla, Rafael Trujillo Codorníu, René López-Baracaldo, Roberto Sagaró-Zamora, Denis Delisle-Rodriguez, John Jairo Villarejo-Mayor, and José Ricardo Núñez-Álvarez. Monte carlo dropout for uncertainty estimation and motor imagery classification. *Sensors*, 21(21):7241, 2021.
- [34] Takayuki Osa, Joni Pajarinen, Gerhard Neumann, J Andrew Bagnell, Pieter Abbeel, and Jan Peters. An algorithmic perspective on imitation learning. *arXiv preprint arXiv:1811.06711*, 2018.
- [35] Alexandros Paraschos, Christian Daniel, Jan R Peters, and Gerhard Neumann. Probabilistic movement primitives. In *Advances in neural information processing systems*, pages 2616–2624, 2013.
- [36] Alexander Pashevich, Robin Strudel, Igor Kalevatykh, Ivan Laptev, and Cordelia Schmid. Learning to augment synthetic images for sim2real policy transfer. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2651–2657. IEEE, 2019.
- [37] Tim Pearce, Alexandra Brintrup, Mohamed Zaki, and Andy Neely. High-quality prediction intervals for deep learning: A distribution-free, ensembled approach. In *International conference on machine learning*, pages 4075–4084. PMLR, 2018.
- [38] Sahba Aghajani Pedram, Peter Walker Ferguson, Changyeob Shin, Ankur Mehta, Erik P Dutson, Farshid Alambeigi, and Jacob Rosen. Toward synergic learning for autonomous manipulation of deformable tissues via surgical robots: An approximate q-learning approach. In *2020 8th IEEE RAS/EMBS International Conference for Biomedical Robotics and Biomechatronics (BioRob)*, pages 878–884. IEEE, 2020.
- [39] Geoff Pleiss, Manish Raghavan, Felix Wu, Jon Kleinberg, and Kilian Q Weinberger. On fairness and calibration. *Advances in neural information processing systems*, 30, 2017.
- [40] Dean A Pomerleau. Efficient training of artificial neural networks for autonomous navigation. *Neural Computation*, 3(1), 1991.
- [41] Ameya Pore, Eleonora Tagliabue, Marco Piccinelli, Diego Dall’Alba, Alicia Casals, and Paolo Fiorini. Learning from demonstrations for autonomous soft-tissue retraction. In *2021 International Symposium on Medical Robotics (ISMR)*, pages 1–7. IEEE, 2021.
- [42] Manuel Retana, Kunal Nalamwar, Drew T Conyers, S Farokh Atashzar, and Farshid Alambeigi. Autonomous data-driven manipulation of an unknown deformable tissue within constrained environments: A pilot study. In *2022 International Symposium on Medical Robotics (ISMR)*, pages 1–7. IEEE, 2022.
- [43] Stephane Ross, Geoffrey J Gordon, and J Andrew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2011.
- [44] Soheil Sarabandi, Arya Shabani, Josep M Porta, and Federico Thomas. On closed-form formulas for the 3-d nearest rotation matrix problem. *IEEE Transactions on Robotics*, 36(4):1333–1339, 2020.
- [45] Murat Sensoy, Lance Kaplan, and Melih Kandemir. Evidential deep learning to quantify classification uncertainty. *Advances in neural information processing systems*, 31, 2018.
- [46] Satvik Sharma, Ellen Novoseller, Vainavi Viswanath, Zaynah Javed, Rishi Parikh, Ryan Hoque, Ashwin Balakrishna, Daniel S Brown, and Ken Goldberg. Learning switching criteria for sim2real transfer of robotic fabric manipulation policies. In *2022 IEEE 18th International Conference on Automation Science and Engineering (CASE)*, pages 1116–1123. IEEE, 2022.
- [47] Changyeob Shin, Peter Walker Ferguson, Sahba Aghajani Pedram, Ji Ma, Erik P Dutson, and Jacob Rosen. Autonomous tissue manipulation via surgical robot using learning based model predictive control. In *2019 International conference on robotics and automation (ICRA)*, pages 3875–3881. IEEE, 2019.

- [48] Jonathan Spencer, Sanjiban Choudhury, Matthew Barnes, Matthew Schmittle, Mung Chiang, Peter Ramadge, and Siddhartha Srinivasa. Learning from interventions: Human-robot interaction as both explicit and implicit feedback. In *Proc. Robotics: Science and Systems (RSS)*, 2020.
- [49] Priya Sundaresan, Priya Grannen, Brijen Thananjeyan, Ashwin Balakrishna, Michael Laskey, Kevin Stone, Joseph E. Gonzalez, and Ken Goldberg. Learning interpretable and transferable rope manipulation policies using depth sensing and dense object descriptors. In *International Conference on Robotics and Automation (ICRA)*. IEEE, 2020.
- [50] Bao Thach, Brian Y Cho, Tucker Hermans, and Alan Kuntz. Deformernet: Learning bimanual manipulation of 3d deformable objects. *arXiv preprint arXiv:2305.04449*, 2023.
- [51] Faraz Torabi, Garrett Warnell, and Peter Stone. Behavioral cloning from observation. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, 7 2018.
- [52] Wenxuan Wu, Zhongang Qi, and Li Fuxin. Pointconv: Deep convolutional networks on 3d point clouds. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, pages 9621–9630, 2019.
- [53] Jiakai Zhang and Kyunghyun Cho. Query-efficient imitation learning for end-to-end autonomous driving. In *Association for the Advancement of Artificial Intelligence (AAAI)*, 2017.
- [54] Wenjia Zhang, Haoran Xu, Haoyi Niu, Peng Cheng, Ming Li, Heming Zhang, Guyue Zhou, and Xianyuan Zhan. Discriminator-guided model-based offline imitation learning. In *Conference on Robot Learning*, pages 1266–1276. PMLR, 2023.
- [55] Xiaoming Zhang, Daniel Lin, Hugh Pforsich, and Vernon W Lin. Physician workforce in the united states of america: forecasting nationwide shortages. *Human resources for health*, 18(1):1–9, 2020.
- [56] Wenshuai Zhao, Jorge Peña Queralta, and Tomi Westerlund. Sim-to-real transfer in deep reinforcement learning for robotics: a survey. In *2020 IEEE symposium series on computational intelligence (SSCI)*, pages 737–744. IEEE, 2020.