

3D Perception based Imitation Learning under Limited Demonstration for Laparoscope Control in Robotic Surgery

Bin Li¹, Ruofeng Wei², Jiaqi Xu³, Bo Lu¹, Chi Hang Yee⁴, Chi Fai Ng⁴,
Pheng-Ann Heng³, *Senior Member, IEEE*, Qi Dou³, *Member, IEEE*, and Yun-Hui Liu¹, *Fellow, IEEE*

Abstract—Automatic laparoscope motion control is fundamentally important for surgeons to efficiently perform operations. However, its traditional control methods based on tool tracking without considering information hidden in surgical scenes are not intelligent enough, while the latest supervised imitation learning (IL)-based methods require expensive sensor data and suffer from distribution mismatch issues caused by limited demonstrations. In this paper, we propose a novel Imitation Learning framework for Laparoscope Control (ILLC) with reinforcement learning (RL), which can efficiently learn the control policy from limited surgical video clips. Specially, we first extract surgical laparoscope trajectories from unlabeled videos as the demonstrations and reconstruct the corresponding surgical scenes. To fully learn from limited motion trajectory demonstrations, we propose Shape Preserving Trajectory Augmentation (SPTA) to augment these data, and build a simulation environment that supports parallel RGB-D rendering to reinforce the RL policy for interacting with the environment efficiently. With adversarial training for IL, we obtain the laparoscope control policy based on the generated rollouts and surgical demonstrations. Extensive experiments are conducted in unseen reconstructed surgical scenes, and our method outperforms the previous IL methods, which proves the feasibility of our unified learning-based framework for laparoscope control.

I. INTRODUCTION

Robot-assisted minimally invasive surgery (RMIS) nowadays is widely used in a variety of laparoscopic procedures, which can reduce the patient's trauma and postoperative hospitalization [1]. A high-quality field of view (FoV) from the laparoscope is crucial for providing surgeons with suitable real-time visual feedback throughout the duration of a laparoscopic surgery. To help alleviate the burden of surgeons on periodically adjusting the laparoscope FoV, thus assisting them to improve surgical efficiency and performance, automatic laparoscope control has recently become popular for developing surgical automation techniques. However, this task has two main challenges. First, the laparoscope needs to




Fig. 1. Surgeons operate the laparoscope (ECM) by MTM to obtain the best FoV. Based on these limited demonstrations, we propose an ILLC to learn how to adjust FoV from raw recorded surgical videos.

be accurately controlled within a limited 3D space in body. The second is intelligent perception of 3D surgical scenes for deploying generalizable laparoscope control policies under complex surgical scenarios [2].

To address these challenges, previous efforts concentrate on the rule-based laparoscope control schemes with diverse perception forms, such as tool tracking [3], [4], [5], eye gaze [6], [7], etc. For example, Osa *et al.* [3] proposed a visual servoing method that controlled the laparoscope to keep the surgical tool at the center point of FoV. Li *et al.* [5] further extended the center point into a domain knowledge based area derived from procedural understanding of the surgical videos. However, these hand-crafted rules have a limited level of intelligence and can hardly be applied to complex surgical environments. Another solution is learning-based methods, which attempt to inject data-driven intelligence into the laparoscope control. Heuristics model [8] and Gaussian mixture models (GMM)-based method [9] are proposed to explore the connection between laparoscope motion and tool positions. Nevertheless, such methods require expensive extra sensory data as the state information (e.g., recording poses of tools and laparoscope) for behavior learning, which is typically unavailable in recorded surgical videos in clinical routine.

Recent advances in imitation learning (IL) based on visual states modelled by convolutional neural network (CNN) have shown promise to imitate expert behaviors from raw image observations. Some impressive success based on such image features has been demonstrated on general robot tasks, such as autonomous driving [10], [11], aerial filming [12], [13],

*This work is supported in part of the HK RGC under T42-409/18-R and 14202918, in part by Shenzhen Portion of Shenzhen-Hong Kong Science and Technology Innovation Cooperation Zone under HZQB-KCZYB-20200089, in part by the Multi-Scale Medical Robotics Centre, InnoHK, and in part by the VC Fund 4930745 of the CUHK T Stone Robotics Institute.

¹B. Li, B. Lu, and Y. H. Liu are with the T stone Robotics Institute, The Department of Mechanical and Automation Engineering, The Chinese University of Hong Kong.

²R. Wei is with the Department of Biomedical Engineering, City University of Hong Kong, Hong Kong.

³J. Xu, P.A. Heng, Q. Dou are with the Department of Computer Science and Engineering, The Chinese University of Hong Kong.

⁴C.H. Yee and C.F. Ng are with the SH Ho Urology Centre, Department of Surgery, The Chinese University of Hong Kong.

Corresponding author: Bo Lu (bolu@cuhk.edu.hk)

virtual cinematography [14]. However, learning automatic laparoscope view adjustment skills from surgical videos with raw RGB-D images as the observation is largely unexplored. Moreover, the limited expert demonstration data extracted from the recorded videos poses challenges to learned models' generalizability to unseen scenarios. Although some online IL methods such as DAgger [15] can alleviate this problem, it still requires expensive labeling with the surgeon's efforts, which is inefficient. Very recently, another solution is to combine IL with reinforcement learning (RL) to achieve better generalizability through self-exploration, where the reward or policy can be recovered from expert demonstrations. Chi *et al.* show a GAIL-based approach for catheterization [16], and we also show that the bimanual peg transfer task can be learned with demonstrations [17]. Nonetheless, these vanilla IL methods may still result in policies that suffer from low performance and generalizability with image observations. Therefore, it is critical to design an efficient IL method using limited surgical expert data to learn the laparoscope control in realistic surgical environments.

In this paper, we propose a novel imitation learning framework for laparoscope control (ILLC), which to the best of our knowledge is the first to use IL with RL to effectively learn the laparoscope motion skills from limited real-world surgical videos, as shown in Fig. 1. Specifically, we first reconstruct 3D surgical scenes from surgical videos as the interactive RL environment and extract laparoscope motion trajectories as the expert demonstration data. To efficiently learn from the limited demonstration, we propose a shape-preserving trajectory augmentation (SPTA) to augment the expert data and use a prior policy to bootstrap the learning process. Afterward, the policy is reinforced using the adversarial learning approach through augmented demonstrations and continuous interactions in our simulated realistic surgical environment. Our main contributions are three folds:

- We propose a novel ILLC framework with RL to achieve intelligent autonomous laparoscope control by learning the joint information between the laparoscope motion and surgical scenes from raw surgical videos.
- We introduce a shape-preserving trajectory augmentation (SPTA) scheme and prior policy to overcome the limited demonstration problem in surgical field and bootstrap the training process to improve the performance of autonomous laparoscope control.
- We conduct validation experiments on controlling the laparoscope under unseen realistic surgical scenes to show the feasibility of our proposed framework and its superiority compared with previous IL baselines.

II. PROBLEM FORMULATION

We aim at building a learner that can imitate surgeons' behavior for laparoscope view adjustment based on the observed surgical scene. Formally, we model the laparoscope control problem as a Markov Decision Process (MDP), represented by $(\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{P}, p(s_0), \gamma)$, with state space \mathcal{S} , action space \mathcal{A} , reward function \mathcal{R} , transition probability \mathcal{P} , initial state distribution $p(s_0)$, and discount factor γ . The

objective of the learner is to find a policy $\pi_{\phi^*}(a|s)$ that can generate an action $a \in \mathcal{A}$ given the observed state $s \in \mathcal{S}$, and maximize the discounted sum of the rewards $r_{\theta}(s, a)$:

$$\max_{\pi_{\phi}} \mathbb{E}_{\pi_{\phi}} \left[\sum_{t=1}^T \gamma^t r_{\theta}(s_t, a_t) \right] \quad (1)$$

where the parameter θ is recovered from expert demonstration trajectories $\mathcal{D} = \{\tau_1, \dots, \tau_N\}$: $\max_{\theta} E_{\tau \sim \mathcal{D}} [\log p_{\theta}(\tau)]$, $p_{\theta}(\tau) \propto p(s_0) \prod_{t=0}^T p(s_{t+1}|s_t, a_t) e^{\gamma^t r_{\theta}(s_t, a_t)}$.

In this paper, we propose an adversarial imitation learning based approach [18] to learn the laparoscope control policy. Specifically, we use a state-of-art RL method to train the learning agents, which collect rollouts within the environment by observing surgical scenes and generating actions. Simultaneously, we train a discriminator to classify whether the trajectory is from expert demonstrations or agent rollouts. In this way, the discriminator can recover the reward function r_{θ} from two different collected data sources, maximizing the reward for RL agents that imitate surgeons' behaviors. Therefore, We obtain the IL policy by optimizing the RL agents and the discriminator alternatively until convergence, which is able to automate the laparoscope control and mimic the surgeon at the same time.

III. METHODS

This section describes our proposed ILLC framework for intelligent autonomous laparoscope control in detail, including 3D simulation environment construction, expert demonstration augmentation, and policy learning strategy, as shown in Fig. 2.

A. 3D perception-based surgical scene rendering

The first step is to build realistic surgical scenes based on the collected videos as the environment for RL agents to control laparoscope and interact within. Meanwhile, the expert trajectories are estimated from their corresponding videos as demonstrations. Specifically, the surgical environment is first reconstructed as the 3D object representation. Afterward, the scene image viewed by the laparoscope is rendered via a 3D projection with the knowledge of the camera pose.

In the scene reconstruction, as the laparoscope and surgical tools cannot be manipulated simultaneously in the da Vinci surgical system [19], we assume the surgical environment remains static when the laparoscope moves, which allows us to reconstruct the 3D scene based on the rigid body assumption. For each frame, we first estimate the dense depth information based on the hierarchical deep stereo matching (HSM) [20] and calculate the camera pose using the iterative closest point (ICP), with the stereo calibration parameters [21]. Finally, the 3D reconstructed scene for each video clip is obtained by jointly merging the object from all frames. You can refer to our previous work [22] for more details. The laparoscope motion clips are extracted from our HKPWH dataset [23] containing real-world stereo prostatectomy videos performed by experienced surgeons.

Based on the reconstructed scene, we build our ECM Simulation environment *ECMSim* for laparoscope control,

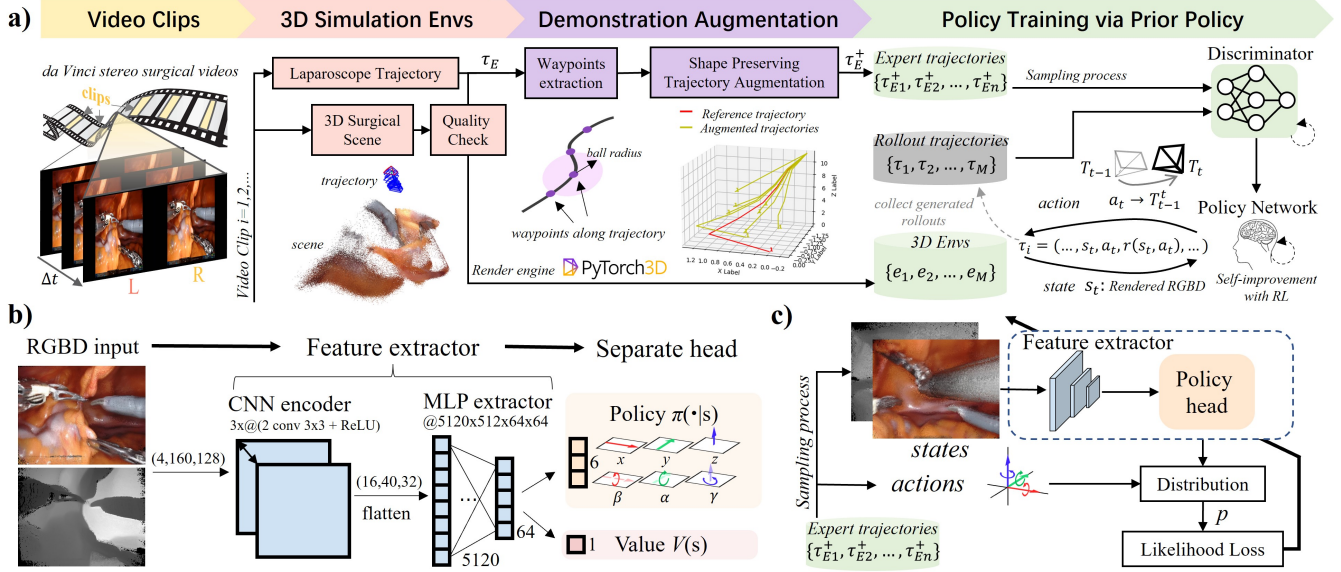


Fig. 2. **The proposed ILLC framework for laparoscope control.** (a) Our ILLC framework comprises 3D scene reconstruction, expert data augmentation, and policy training. (b) The RL policy network consists of a shared CNN encoder, MLP feature extractor, and two separate policy and value heads. (c) Prior policy generated from expert demonstrations with likelihood loss.

which can render realistic RGB-D images when the laparoscope moves within the environment at any position, with *PyTorch3D* [24] as our rendering engine. The dynamic physical interaction between the laparoscope and the tissue is not considered in this work. To maximize the rendering realism, we compared two commonly-used 3D object representations, i.e., origin point cloud and its corresponding polygonal mesh, as illustrated in Fig. 3.

Although the polygonal mesh is theoretically superior to the point cloud, it needs to be generated from the high-quality point cloud. Due to the short duration of the video clips and subtle estimation errors of the depth maps, the quality of the reconstructed point cloud is enough for rendering but suboptimal for high-quality mesh generation. As in Fig. 3, the rendered images from the point cloud is almost the same with the original video images, which shows the accurate estimated trajectory and reconstructed environment scene. Moreover, we can observe that the visualization performance from point cloud is superior than the rendered images with smearing effects from the generated mesh. Therefore, we adopt the reconstructed point cloud as the 3D object representation of the scene for rendering.

B. Demonstration augmentation with SPTA

The expert state-action trajectories τ_E extracted during 3D scene reconstruction is of limited number, which may hinder the generalization of the learned model. To solve this problem, we propose a shape-preserving trajectory augmentation (SPTA) method to augment the collected expert data and build an enriched expert demonstration database τ_E^+ .

Raw expert trajectory preprocessing: For raw expert trajectories extracted during the 3D scene reconstruction, we observe that there exist local jittering fluctuations and inconsistent speeds. Thus, we use the first-order Savitzky-Golay filter to smooth the raw expert trajectories, both




Fig. 3. **Reconstructed surgical scene.** Left: Screenshots of the reconstructed surgical scenes in the point cloud modality (up) and the mesh modality (bottom). Right: The sequences are origin video image (up), rendered image at estimated camera positions in point cloud scene (middle) and mesh scene (bottom).

in position (x, y, z) and orientation (α, β, γ) . To acquire consistent speed during motion, the smoothed trajectory of m points $P_{i \in \{1, \dots, m\}}$ is then generated into n equally-distributed waypoints $W_{j \in \{1, \dots, n\}}$ at a fixed separation distance d_{fixed} , based on the linear interpolation as follows:

$$W_j^{(\cdot)} = W_{j-1}^{(\cdot)} + \frac{d_{fixed}}{\|P_{i^*}^{(\cdot)} - W_{j-1}^{(\cdot)}\|_2} (P_{i^*}^{(\cdot)} - W_{j-1}^{(\cdot)}), \quad (2)$$

$$i^* = \arg \min_{i \in \{1, \dots, n\}} (\|P_i^{(\cdot)} - W_{j-1}^{(\cdot)}\|_2 > d_{fixed}).$$

where the initial waypoint $W_0^{(\cdot)} = P_0^{(\cdot)}$, and (\cdot) can be any of $\{x, y, z, \alpha, \beta, \gamma\}$. Then we calculate the j^{th} waypoint's camera pose relative to the initial camera pose T_0 , represented by a transformation matrix T_0^j with rotation matrix $R(\alpha, \beta, \gamma)$ and translation vector $t(x, y, z)$. Its observed state s_j as an RGB-D image is rendered by our *ECMSim*. The corresponding action a_j is calculated as: $a_j \rightarrow T_j^{j+1} = inv(T_0^j) \cdot T_0^{j+1}$, and the action a is defined as the transformation $(x, y, z, \alpha, \beta, \gamma)$ relative to the current camera pose.

Shape Preserving Trajectory Augmentation (SPTA):

To augment surgeons' demonstrations and mimic the expert behavior, we consider two criteria:

Criterion 1: The augmented laparoscope motion trajectory should maintain a similar movement pattern to the expert trajectory for each surgical scene.

Criterion 2: The laparoscope needs to stop near the endpoint in the expert trajectory to obtain a desirable view similar to the expert for any sampled start positions.

To enable this, we first generate a 3D cubic workspace WS from the trajectory $\tau = W_{j \in \{1, \dots, end\}}$: $WS = Cube\{W_{1:k}\}$, where k is randomly chosen and slightly larger than 1, so the generated WS is near the starting point W_1 . Then the starting point S is uniformly sampled within WS for diversity. Afterward, we generate an augmented trajectory from S to W_{end} of a similar movement pattern with the original expert demonstration.

To achieve a balance between the local and global similarity, we acquire the waypoint point W_{j^*} closest to sampled start point S and use the trajectory segment from this point to the trajectory endpoint $W_{j^*:end}$ as the reference for augmentation. Therefore, the influence of unused motion pattern in the initial part $W_{1:j^*}$ far from the start point can be avoided, which is similar to the idea of reward-to-go [25]. Next, we treat it as a non-linear fitting problem that use the exponential decay to gradually reduce the initial distance e_0 between S and W_{j^*} as follows:

$$dist = \|S - W_{j^*}\|_2, j^* = \arg \min_{j \in \{1, \dots, end\}} (\|S - W_j\|_2) \quad (3)$$

$$\begin{aligned} W'_{j'} &= W_{j^*+j'} + (1 - \epsilon) \cdot \Delta e_{j'}, \\ \Delta e_{j'} &= k_1 \cdot e^{\gamma \cdot j'} + k_2, j' = 0, 1, \dots, end - j^*. \end{aligned} \quad (4)$$

where the coefficients k_1, k_2 are solved with two boundary conditions: $e_0 = dist$ and $e_{end} = \epsilon$, ϵ is a random small tolerance; γ is the hyper-parameter which tunes the approaching speed, and $W'_{j'}$ is the augmented waypoint.

C. Policy learning via prior policy

Inspired by [18], we use adversarial learning approach for policy learning with RGB-D image input in the ILLC framework. The RL laparoscope agent policy $\pi_\phi(\cdot)$ parameterized by ϕ is modeled as the generator and is trained against the discriminator adversarially. The discriminator D_θ is trained to classify surgeon demonstration trajectories $(s_t, a_t) \sim \tau_E^+$ and agent rollouts $(s'_t, a'_t) \sim \tau$, with the cross-entropy loss:

$$\begin{aligned} L_\theta(D) &= -\mathbb{E}_{\tau_E^+} [\log(D_\theta(s_t, a_t))] \\ &\quad - \mathbb{E}_\tau [\log(1 - D_\theta(s'_t, a'_t))] \end{aligned} \quad (5)$$

With the following form for discriminator: $D_\theta(s, a) = \frac{\pi_\phi(a|s)}{\exp\{r_\theta(s, a)\} + \pi_\phi(a|s)} = f(\log \pi_\phi(a|s) - r_\theta(s, a))$, where $f(\cdot)$ is the softmax function, $r_\theta(\cdot)$ is the reward function to be recovered. Meanwhile, the policy π_ϕ is trained to minimize the entropy-regularized discriminative reward: $\hat{r}(s, a) = \log \pi_\phi(a|s) - r_\theta(s, a)$. The entropy-regularized RL objective is obtained by summing $\hat{r}(s_t, a_t)$ over entire trajectories:

$$\mathbb{E}_{\pi_\phi} [\sum \hat{r}(s_t, a_t)] = \mathbb{E}_{\pi_\phi} [\sum \log \pi_\phi(a_t|s_t) - r_\theta(s_t, a_t)] \quad (6)$$

And the global optima of the discriminator objective is achieved when $\pi_\theta = \pi_E$, as the learned policy π_θ converges to the expert policy π_E .

Considering that the current view alone can determine the next action in the laparoscope control, we assume that the ground truth reward is only related to the current state (RGB-D image), which meets the basis of the algorithm [18]. The stochastic Gaussian policy is used for continuous laparoscope actions, $\pi_\phi(a|s) = \mu_\phi(s) + \sigma_\phi(s) \odot \epsilon$, where mean and variance are generated by deterministic functions μ_ϕ, σ_ϕ , $\epsilon \sim \mathcal{N}(0, 1)$, and \odot denotes the element-wise product.

During training, all training scenes are loaded into multiple threads simultaneously to accelerate the data collection, and each scene acts as an independent scene-agent environment. We collect the agents' rollout trajectories from all training environments and combine them with expert demonstrations to optimize the discriminator D_θ . The accordingly updated rewards with r_θ in the previous rollouts are then used to update the agent policy π_ϕ using Eq. 6.

Bootstrapping with prior policy: Generally, the adversarial training will gradually optimize the agent policy and the discriminator. However, as laparoscope control requires high precision, if the policy starts from random, a large amount of valueless rollout in the early stage will make policy and discriminator get stuck with low performance. To bootstrap the training process and sidestep the exploration challenge, we utilize the augmented expert demonstration π_E^+ to pre-train the policy and get a stable prior policy $\pi_{\phi_{init}}$, inspired by [26], followed with the adversarial training to further fine-tune the parameters. Formally, with the expert state-action tuples $\{(s_i, a_i)\} \in \tau_E^+$, our objective is to find the parameters ϕ_{init} , which best fit the expert state-action pairs:

$$\phi_{init} = \arg \max_{\phi} \prod \pi_\phi(a_i|s_i) \quad (7)$$

where $\pi_\phi(a_i|s_i)$ represent the probability of each expert action a_i in the trained policy's Gaussian distribution $\pi_\phi(\cdot|s_i)$. Therefore we can use Adam to solve ϕ_{init} based on the gradient to seek the maximum-likelihood.

D. Implementation

The RL policy π_ϕ is realized by a deep neural network, which consists of a shared component of 6 CNN layers, MLP layers, and a separate policy head and value head of output size 6 and 1, respectively, as shown in Fig. 2 (b). Considering the balance of training speed and accuracy, the input size are scaled to 160×128 and the proximal policy optimization (PPO) [27] is used. Actions are scaled to $[-1, 1]$ by the tanh layer in the PPO policy network and then rescaled to their appropriate range $[-1.5, 1.5] mm$ and $[-3^\circ, 3^\circ]$ to move the laparoscope. The learning rate of PPO agent is $1e-5$.

For discriminator D_θ , the reward function r_θ is released by a basic reward network $r_b(\cdot)$ and shaping term $h(\cdot)$ with the same network structure but separate weights, consisting of 2 CNN layers followed by $[32, 32, 1]$ MLP layers. The learning rate for the discriminator is $3e-4$. The batch size is 64, and the rollout buffer size is 4096.

IV. EXPERIMENT RESULTS

A. Experiment setup

Since the 3D reconstruction may fail in severe conditions, such as lens contamination, bleeding, we manually check and collect 95 high-quality reconstructed surgical scenes as our dataset. Among them, 80 scenes are randomly chosen for training and the rest 15 scenes for the test, as shown in Fig. 4. The statistics of video clips and reconstruction details (scene point cloud, trajectory) are shown in Table I. On average, the laparoscope moves 8.23 *mm* distance within 0.49 *s*, which shows the challenge of this task because it requires precise laparoscope control in a short time.

We built our proposed simulation environment based on the gym interface [28] and use *PyTorch3D* as the render engine. To improve the generality, the initial laparoscope position is randomly sampled in the origin trajectory workspace, also with the orientation variation. An episode is terminated when the deviations of both the position and orientation from the expert endpoint are below particular thresholds, or when the episode steps reach a maximum length. In our experiments, we set thresholds of position and orientation as 2 *mm* and 5°, and the maximum length to be 16.

TABLE I

STATISTICS OF LAPAROSCOPY VIDEO CLIPS AND 3D RECONSTRUCTION

Split	num	duration / <i>sec</i>	point number / <i>1e5</i>	distance / <i>mm</i>	steps
Train	80	0.50±0.30	1.92±0.53	8.35±4.24	7.79±5.80
Test	15	0.42±0.11	1.85±0.28	7.57±4.22	6.33±3.17
Total	95	0.49±0.28	1.91±0.50	8.23±4.25	7.56±5.50

B. Evaluation metrics

Success rate (SR) is used as the main metric to assess the performance of different policies. Besides, we observe the motion pattern is not efficient in some successful episodes, e.g., the zig-zag movement. Hence, we propose another metric action efficiency (a_{eff}) to measure the motion efficiency in successful episodes. These metrics are calculated as:

$$SR = \frac{\sum_i^N sr_i}{N}, \quad sr = \mathbb{I}(\Delta p^T < \epsilon_p \cap \Delta r^T < \epsilon_r), \quad (8)$$

$$\Delta p = \|P_{xyz} - P_{xyz}^*\|_2, \quad \Delta r = \cos^{-1}\left(\frac{tr(R_{\alpha\beta\gamma} \cdot R_{\alpha\beta\gamma}^*)}{2}\right)$$

$$a_{eff} = \frac{\sum_i^N \left(sr_i \cdot \|P_{xyz}^1 - P_{xyz}^{T_i}\|_2 / T_i \right)}{\sum_i^N sr_i} \quad (9)$$

where N is the episode number, $\mathbb{I}(\cdot)$ is the indicator function, $P^{(*)}$ and $R^{(*)}$ are the endpoint position and orientation in rollouts ($*$ represent in demonstrations), ϵ_p and ϵ_r denote the corresponding thresholds, $tr(\cdot)$ is the trace of a matrix, T is the total steps within one episode.

C. Results on validation

1) *Comparison with IL baselines*: We run 50 episodes for policies with random initialization under each unseen test scene, and report the mean and the standard deviations for 15 test scenes. The experimental results are listed in




Fig. 4. Screenshots of some typical training and test scenes.

TABLE II

EVALUATION RESULT OF DIFFERENT POLICIES ON TEST SCENES

Metrics	steps (↓)	a_{eff} / <i>mm</i> (↑)	SR / % (↑)
Expert policy	5.63±3.45	1.44±0.25	100
Random policy	16.00±0.00	0	0
BC policy [29]	10.98±6.06	0.90±0.78	39.73±43.25
GAIL policy [30]	9.70±5.77	0.86±0.53	43.86±36.22
ILLC (ours)	8.57±5.37	1.04±0.34	66.47±25.62

Table II. We choose the random policy with random actions to serve as a lower bound. Two learning-based methods, i.e., behavior cloning (BC) [29] and GAIL [30] are selected and trained with our augmented expert data for a fair comparison. From the result, BC policy only achieves a success rate of 39.73% with a large variance of 43.25%, since it only remembers the actions in the training scenes and cannot generalize well to the test scenes. GAIL policy has achieved a higher success rate of 43.86%, however, it still fails in many cases as it is difficult to deal with unknown dynamics in various unseen scenes. By contrast, our ILLC policy achieves the highest success rate of 66.47% with the smallest variance and succeeds under a majority of unseen scenes, which demonstrates that the ILLC policy can learn the expert motion pattern with generality. By illustrating the typical laparoscope trajectories generated by the random, expert, and our policy in Fig. 5(a), we can observe that our method can output similar laparoscope trajectories compared to the expert ones, and can also terminate with the same target view. Moreover, even in some positions that deviated from the expert trajectories, our policy can still correct its direction and move to the corresponding ideal endpoint. The screenshots in Fig. 5(b) also indicate that the learned policy tend to achieve an appropriate FoV that is consistent with expert, e.g., making tools located in the middle region, and making the depth appropriate for the whole scene, etc.

2) *Ablation study*: We design experiments to analyze the impact of three factors: 1) SPTA, 2) prior policy, 3) observed state. The policy performance with SPTA under different augmentation rates R (SPTA@ R x) is compared in Table III and also visualized in Fig. 6. Results show that SPTA can largely improve the success rate, e.g., from 48.27% without SPTA to 66.47% with SPTA@32x, of a smaller variance and improved action efficiency. The reason is that SPTA can generate more expert-like data, which enables the IL agent to




Fig. 5. (a) Examples of laparoscope motion trajectories with different policies under test scenes: random, expert, trained policy are colored by red, green, blue respectively. (b) Screenshots of the #steps in the two test scenes running with expert (first row), ILLC (second row), random policies (last row).

TABLE III
ABLATION STUDY ON POLICY TRAINING

Metrics	steps (↓)	a_{eff} / mm (↑)	SR / % (↑)
ILLC (SPTA@32x)	8.57±5.37	1.04±0.34	66.47±25.62
w/o SPTA	10.26±5.32	1.00±0.63	48.27±38.38
w/ SPTA@8x	9.78±4.70	1.23±0.44	54.00±31.73
w/o prior policy	10.57±5.37	0.89±0.32	46.21±39.20
w/o depth	9.83±5.12	0.98±0.56	54.29±39.34




Fig. 6. Policy performance with SPTA under different augmentation rates.

fully explore the environment and achieve a better generality. Table III also compares the policy trained without the prior policy, we can observe that the prior policy can improve the success rate by 20.26% since it can provide more meaningful rollouts in the early stage and bootstraps the overall training process, hence proving its importance in our policy. Finally, we discuss the impact of depth in the state representation on the learned policy. Results show that including the depth in the input can improve the success rate by 12.18%, which demonstrates its benefits to the FoV control.

3) *Visualization analysis*: To explore the hidden knowledge learned by our proposed ILLC policy when performing the laparoscope control task and intuitively illustrate such information, we use t-SNE [31] to visualize the embedding features in the last hidden layer of the ILLC policy network. The input observed states to the network are collected in test scenes for evaluation with random, expert, and ILLC policy, respectively. As shown in Fig. 7(a), the distributions of ILLC policy and expert policy are thoroughly mixed and far away from the random policy ones, which indicates that to same extent the ILLC policy controls the laparoscope in a similar

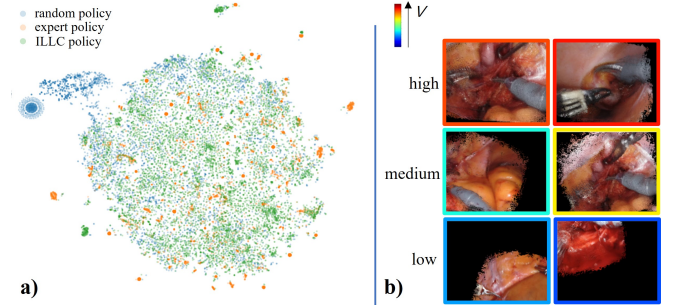


Fig. 7. (a) 2D t-SNE embeddings of policy network features for data collected by different policies: random, expert, and our ILLC policy. (b) Visualization of predicted values by trained policy; the color surrounding the image represents the state values from lowest (blue) to highest (red).

way to the experts. We also visualize the estimated values from the policy network with observed images in Fig. 7(b), where high values indicate the observed states are promising outcomes. As expected, the ILLC policy maps the proper laparoscopic views to high estimated values across different test scenes. For example, in Fig. 7(b), the top row with appropriate views has higher state values, and the bottom row far from the ideal viewpoints has low values.

V. CONCLUSIONS

In this paper, we propose a novel ILLC framework to achieve intelligent autonomous laparoscope control. Unlike conventional tracking-based control approaches, the learned policy generates actions considering the RGB-D surgical scene context. Meanwhile, the policy can capture the joint connection between the laparoscope motion and observed scenes from raw surgical videos. The learned policy shows superior laparoscope control performance and generality to unseen test scenes compared with previous IL baselines.

In the future, we will enrich the surgical scenes from surgical videos and improve the 3D reconstruction quality to make the interactive environment more realistic. In addition, we will optimize the policy network, such as adding temporal information to further improve performance.

REFERENCES

- [1] V. Vitiello, K.-W. Kwok, and G.-Z. Yang, "Introduction to robot-assisted minimally invasive surgery (mis)," in *Medical Robotics*. Elsevier, 2012, pp. 1–P1.
- [2] A. Pandya, L. A. Reisner, B. King, N. Lucas, A. Composto, M. Klein, and R. D. Ellis, "A review of camera viewpoint automation in robotic and laparoscopic surgery," *Robotics*, vol. 3, no. 3, pp. 310–329, 2014.
- [3] T. Osa, C. Staub, and A. Knoll, "Framework of automatic robot surgery system using visual servoing," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2010, pp. 1837–1842.
- [4] B. Yang, W. Chen, Z. Wang, Y. Lu, J. Mao, H. Wang, and Y.-H. Liu, "Adaptive fov control of laparoscopes with programmable composed constraints," *IEEE Transactions on Medical Robotics and Bionics*, vol. 1, no. 4, pp. 206–217, 2019.
- [5] B. Li, B. Lu, Y. Lu, Q. Dou, and Y.-H. Liu, "Data-driven holistic framework for automated laparoscope optimal view control with learning-based depth perception," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 12 366–12 372.
- [6] K. Fujii, G. Gras, A. Salerno, and G.-Z. Yang, "Gaze gesture based human robot interaction for laparoscopic surgery," *Medical image analysis*, vol. 44, pp. 196–214, 2018.
- [7] G. Gras, K. Leibrandt, P. Wisanuvej, P. Giataganas, C. A. Seneci, M. Ye, J. Shang, and G.-Z. Yang, "Implicit gaze-assisted adaptive motion scaling for highly articulated instrument manipulation," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 4233–4239.
- [8] A. Bihlmaier and H. Woern, "Automated endoscopic camera guidance: A knowledge-based system towards robot assisted surgery," in *ISR/Robotik 2014; 41st International Symposium on Robotics*. VDE, 2014, pp. 1–6.
- [9] I. Rivas-Blanco, C. J. Perez-del Pulgar, C. López-Casado, E. Bauzano, and V. F. Muñoz, "Transferring know-how for an autonomous camera robotic assistant," *Electronics*, vol. 8, no. 2, p. 224, 2019.
- [10] A. O. Ly and M. Akhloufi, "Learning to drive by imitation: An overview of deep behavior cloning methods," *IEEE Transactions on Intelligent Vehicles*, vol. 6, no. 2, pp. 195–209, 2020.
- [11] H. Xu, Y. Gao, F. Yu, and T. Darrell, "End-to-end learning of driving models from large-scale video datasets," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2174–2182.
- [12] C. Huang, Z. Yang, Y. Kong, P. Chen, X. Yang, and K.-T. T. Cheng, "Learning to capture a film-look video with a camera drone," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 1871–1877.
- [13] C. Huang, Y. Dang, P. Chen, X. Yang *et al.*, "One-shot imitation filming of human motion videos," *arXiv preprint arXiv:1912.10609*, 2019.
- [14] H. Jiang, B. Wang, X. Wang, M. Christie, and B. Chen, "Example-driven virtual cinematography by learning camera behaviors," *ACM Transactions on Graphics (TOG)*, vol. 39, no. 4, pp. 45–1, 2020.
- [15] S. Ross, G. Gordon, and D. Bagnell, "A reduction of imitation learning and structured prediction to no-regret online learning," in *Proceedings of the fourteenth international conference on artificial intelligence and statistics*. JMLR Workshop and Conference Proceedings, 2011, pp. 627–635.
- [16] W. Chi, G. Dagnino, T. M. Kwok, A. Nguyen, D. Kundrat, M. E. Abdelaziz, C. Riga, C. Bicknell, and G.-Z. Yang, "Collaborative robot-assisted endovascular catheterization with generative adversarial imitation learning," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 2414–2420.
- [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," *arXiv preprint arXiv:2108.13035*, 2021.
- [18] J. Fu, K. Luo, and S. Levine, "Learning robust rewards with adversarial inverse reinforcement learning," in *International Conference on Learning Representations*, 2018.
- [19] S. DiMaio, M. Hanuschik, and U. Kreaden, "The da vinci surgical system," in *Surgical Robotics*. Springer, 2011, pp. 199–217.
- [20] G. Yang, J. Manela, M. Happold, and D. Ramanan, "Hierarchical deep stereo matching on high-resolution images," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 5515–5524.
- [21] Z. Zhang, "A flexible new technique for camera calibration," *IEEE Transactions on pattern analysis and machine intelligence*, vol. 22, no. 11, pp. 1330–1334, 2000.
- [22] R. Wei, B. Li, H. Mo, B. Lu, Y. Long, B. Yang, Q. Dou, Y. Liu, and D. Sun, "Stereo dense scene reconstruction and accurate laparoscope localization for learning-based navigation in robot-assisted surgery," *arXiv preprint arXiv:2110.03912*, 2021.
- [23] Y. Long, Z. Li, C. H. Yee, C. F. Ng, R. H. Taylor, M. Unberath, and Q. Dou, "E-dssr: Efficient dynamic surgical scene reconstruction with transformer-based stereoscopic depth perception," *arXiv preprint arXiv:2107.00229*, 2021.
- [24] N. Ravi, J. Reizenstein, D. Novotny, T. Gordon, W.-Y. Lo, J. Johnson, and G. Gkioxari, "Accelerating 3d deep learning with pytorch3d," *arXiv:2007.08501*, 2020.
- [25] A. Tamar, D. Di Castro, and S. Mannor, "Learning the variance of the reward-to-go," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 361–396, 2016.
- [26] K. He, R. Girshick, and P. Dollár, "Rethinking imagenet pre-training," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 4918–4927.
- [27] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [28] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "Openai gym," 2016.
- [29] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang *et al.*, "End to end learning for self-driving cars," *arXiv preprint arXiv:1604.07316*, 2016.
- [30] J. Ho and S. Ermon, "Generative adversarial imitation learning," *Advances in neural information processing systems*, vol. 29, pp. 4565–4573, 2016.
- [31] L. Van der Maaten and G. Hinton, "Visualizing data using t-sne," *Journal of machine learning research*, vol. 9, no. 11, 2008.