

Colonoscopy Navigation using End-to-End Deep Visuomotor Control: A User Study

Ameya Pore^{1,2,*}, Martina Finocchiario^{2,3}, Diego Dall’Alba¹, Albert Hernansanz², Gastone Ciuti³,
Alberto Arezzo⁴, Arianna Menciassi³, Alicia Casals², Paolo Fiorini¹

Abstract—Flexible Endoscopes (FEs) for colonoscopy present several limitations due to their inherent complexity, resulting in patient discomfort and lack of intuitiveness for clinicians. Robotic FEs with autonomous control represent a viable solution to reduce the workload of endoscopists and the training time while improving the procedure outcome. Prior works on autonomous FE control use heuristic policies that limit their generalisation to the unstructured and highly deformable colon environment and require frequent human intervention. This work proposes an image-based FE control using Deep Reinforcement Learning, called Deep Visuomotor Control (DVC), to exhibit adaptive behaviour in convoluted sections of the colon. DVC learns a mapping between the images and the FE control signal. A first user study of 20 expert gastrointestinal endoscopists was carried out to compare their navigation performance with DVC using a realistic virtual simulator. The results indicate that DVC shows equivalent performance on several assessment parameters, being more safer. Moreover, a second user study with 20 novice users was performed to demonstrate easier human supervision compared to a state-of-the-art heuristic control policy. Seamless supervision of colonoscopy procedures would enable endoscopists to focus on the medical decision rather than on the control of FE.

I. INTRODUCTION

Colonoscopy remains the gold standard for the diagnosis and treatment of colorectal cancer (CRC), which is the third most common malignancy worldwide [1]. During a routine procedure, a Flexible Endoscope (FE) is firstly inserted from the rectum to the caecum and then retracted to detect possible early-stage CRC lesions. Early-stage detection of CRC can improve the survival rate by 90% [1]. FE-based procedures are complex due to non-intuitive mapping between the endoscope tip and the control steering knobs, which requires a long and extensive training process [2]. Consequently, these procedures increase the risks of tissue stretching and perforation, which are the prime cause of patient discomfort, and pain [3]. Moreover, work-related musculoskeletal injuries have been reported among endoscopists due to awkward body posturing [2]. The shortage of adequately trained endoscopists with respect to the increasing clinical




Fig. 1. Deep Visuomotor Control (DVC) flow diagram. The environment provides a state observation S_t . The DVC agent uses the state input to generate an action a_t that is applied to the environment. During the training, DVC learns a policy π_ϕ to perform autonomous navigation. In the evaluation phase, the supervising clinicians can override DVC decisions through action $a_{t'}$.

demand of colonoscopy procedures can lead to potential loss of human lives.

To overcome these limitations, current research efforts are developing navigation systems using robotised FE, such as magnetic actuated FEs [4, 5, 3]. Robotised FE enables the introduction of automation technologies by adding autonomous navigation, which is the most time-consuming step of colonoscopy [5]. This will allow the endoscopists to focus on the clinical aspect of the procedure rather than the control of FE, potentially improving the overall procedure outcome and reducing the training time [3].

During the navigation phase, the clinician mainly uses visual feedback from the FE camera to advance through the lumen [6]. A common gesture observed by endoscopists during the procedure is to centralise the direction of the endoscope towards the lumen centre. Prior works on endoscopic navigation have built rule-based controllers to replicate this gesture by reducing the distance between the image centre and the detected lumen centre [7]. These algorithms fail in situations when the tip of the endoscope approaches close to the colon wall. Such situations occur due to the highly deformable nature of the colon and the variable mobility introduced by patient movements, peristalsis and breathing which makes lumen detection not trivial. These situations require human interventions to correct the motion direction, or they can be handled by adaptive exploration methods, as proposed in this work.

Rule-based controllers are being progressively replaced by data-driven approaches such as Deep Reinforcement Learn-

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie (grant agreement No. 813782 “ATLAS”)

* Corresponding author: Ameya Pore (email: ameya.pore@univr.it)

¹ Department of Computer Science, University of Verona, Italy

² Department of Enginyeria de Sistemes, Automatica i Informatica Industrial, Universitat Politècnica de Catalunya, Barcelona, Spain

³ The BioRobotics Institute, Scuola Superiore Sant’Anna, Pisa, Italy

⁴ Department of Surgical Sciences, University of Torino, Turin, Italy

ing (DRL), since they provide some degree of adaptability [8, 9, 10]. However, the application of DRL in learning surgical task has been limited to low-dimensional physical state features such as robot kinematic data, which are widely accepted to be sample-efficient and trivial to learn [8, 9, 10]. This paper proposes an image-based DRL approach for FE control (Fig. 1) focussing on learning the navigation task by devising an end-to-end policy to map the raw endoscopic images to the control signal of the endoscope, referred henceforth as Deep Visuomotor Control (DVC). We primarily evaluate DVC through a user study with 20 expert endoscopists who perform the navigation task in a realistic virtual simulator.

The introduction of autonomous navigation can improve clinical practice, relieving clinicians from demanding cognitive and physical tasks. However, in safety-critical areas, such as medical robotics, it is highly desirable to maintain human supervision to address ethical and legal concerns. Hence, it is essential to consider human-in-the-loop for DVC deployment in realistic surgical scenarios. Therefore, we conducted a second user study with 20 novice participants to demonstrate that non-expert users can easily supervise autonomous navigation, and DVC reduces the need for human intervention compared to a state-of-the-art method.

This work presents an initial study towards generating adaptive control for colonoscopy by proposing a DVC policy for autonomous navigation and providing its performance evaluation with expert endoscopists.

The content of this paper is organised as follows: Sec. II describes the related works, while implemented methods are explained in Sec. III. In Sec. IV, we elaborate on the experimental evaluation. Finally, the results and conclusions are discussed in Sec. V and Sec. VI respectively.

II. RELATED WORKS

The advantages of autonomous navigation in colonoscopy have prompted several studies in this field. In [11], a screw-type endoscope is developed, and motion adjustment is demonstrated using reinforcement learning. This study uses robot kinematics variables as state inputs; however, navigation through the straight segments was slow, and navigation through bends proved awkward due to the robot's size. Several studies have focused on magnetic guided endoscopes [4, 5, 3]. In [4], navigation based on predefined trajectories is presented; hence extending this method to complex non-linear trajectories is challenging. Heuristic path planning algorithms are used in [3] to generate a feasible path in a colon model. This approach employs force-based real-time sensing to guide navigation. Force-based sensing is still not widely available in existing FE devices; moreover, the interpretation of robotic actions without scene visualisation is challenging, hence not suitable for human supervision. In [5], a static perception model is developed, which extracts the centre of the lumen from raw image. The control of endoscope position and orientation is imparted by a proportional controller that aligns the endoscopic image with the centre of the lumen. Similar rule-based controllers have

been previously developed in [7]; however, they require significant manual tasking for non-linear components such as analytically computing image jacobian, and interaction matrix [12]. Moreover, lumen detection could be unstable and prone to errors due to the dynamic nature of the colon and its sharp bends. Learning end-to-end visuomotor representations for direct control using DRL overcomes these limitations without separately designing perception and control models and offers the ability to improve model parameters while training [12].

Some studies have proposed frameworks for training DRL policies to automate surgical tasks [8, 9, 13, 14] such as manipulation of rigid and deformable objects. These studies use simplified environments designed explicitly for robot-assisted surgery to learn the instrument control. Recently, [15] proposed a DRL method for optimising the endoscopic camera viewpoint. These studies use low-dimensional state information for training DRL algorithms, such as kinematic values of the robot, the position of target etc. [8, 9, 13, 15]. In a real colonoscopy scenario, it is challenging to accurately capture the endoscope kinematics due to limits on the sensing capabilities [3], and intra-operative guidance is solely based on visual feedback.




Fig. 2. Colon models used in the experimental phase. (From left to right) ranked in increasing complexity order, C_0 , C_1 , C_2 and C_3 models. The model complexity is characterised by the centerline from rectum to caecum, and the number of acute bending, i.e. >90 degree, which is estimated through visual inspection.

III. METHODS

Our objective is to develop end-to-end joint training for perception and control to learn navigation policies that map raw endoscopic image directly to the control signals of the robotised FE (e.g. motor torques). For simplification purposes, we assume that the motion demanded to the robot is carried out immediately without any latency. We develop a realistic colonoscopy simulator with deformable tissue dynamics, described in Sec. III-A. Further, we explain the implementation details of DVC in Sec. III-B.

A. Simulation platform

Colon simulation - The 3D models of the bowel are segmented in a semi-automated way using a public CT colonography dataset [16]. The segmented models are refined, and volumetrical and superficial meshes are generated. Realistic textures are created and applied to the models using the real endoscopy images from KVASIR dataset [16]. Then, the obtained colon models are loaded into *SOFA (Simulation*

Open Framework Architecture), where a realistic mechanical model based on real-time finite element analysis is generated, setting the simulation parameters in order to mimic the colon tissue behaviour [17]. Additionally, the collision detection between the endoscope and the colon is implemented, and physical constraints are included to realistically restrain the colon deformations. *Unity3D* is adopted for creating high quality and realistic visual rendering including additional visual effects such as reflection on the organ surface or vignetting the peripheral darkening of the endoscopy image [16].




Fig. 3. Representation of the local frame at the endoscope tip. The X-Y plane of the camera is parallel to the image frame, while the z-axis represents the direction of insertion. The green region represents the detected lumen centre.

Endoscope simulation - We assume a scenario close to a magnetically guided FE where external magnets control the motion of the magnetic tip while the tether follows the tip passively [4]. Hence, in this preliminary simulator version, we neglect the effect of the endoscope tether due to multiple collision points with the colon model that could lead to simulation instability. The endoscope tip is modelled as a rigid capsule with weight, length and diameter of 20g, 36mm and 14mm, respectively [18]. An angular drag of 4 rad/sec² is added to account for the frictional resistance [18]. The endoscope tip embeds a camera and has four degrees of freedom for the motion as shown in Fig. 3, i.e. translation (insertion/retraction), roll, bending in two perpendicular directions (pitch/yaw).

B. Deep Visuomotor control

DRL background - The colon navigation problem is formalised into a Markov Decision Process (MDP) represented by a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{P}, \gamma, T)$, where \mathcal{S} denotes the state space, \mathcal{A} is the action space, \mathcal{P} is the transition probability distribution, \mathcal{R} is the reward space, $\gamma \in [0, 1]$ is the discount factor and T is the time horizon per episode. At each timestep t , the environment produces a state observation $s_t \in \mathcal{S}$. The agent then generates an action $a_t \in \mathcal{A}$ according to a policy $a_t \sim \pi(s_t)$, and applies it to the environment to receive a reward $r_t \in \mathcal{R}$ [19]. As a consequence, the agent transitions to a new state s_{t+1} sampled from the transition function $p(s_{t+1}|s_t, a_t)$, $p \in \mathcal{P}$ or terminates the episode at state s_T .

Learning algorithm - The agent's goal is to learn a stochastic behaviour policy π parametrised by ϕ , $\pi_\phi : \mathcal{S} \rightarrow$

$\mathcal{P}(\mathcal{A})$ to maximise the expected future discounted reward $E[\sum_{i=0}^{T-1} \gamma^i r_i]$. We chose PPO [19] as a consolidated DRL algorithm due to overall returns in terms of wall-clock training time and hyper-parameter tuning. It is out of the scope of this work to propose a novel DRL method, while the main goal is to perform a user study to evaluate the performance of image-based DRL in colonoscopy navigation. In the training session, the length of each episode is set as 10k iteration steps, $\gamma = 0.99$, and the batch size and the learning rate hyperparameters are 64 and 3e-4, respectively. The PPO clip ratio was 0.2 with 4 mini-batches per epoch and 4 epochs per iteration. Each training lasted for 1.5 million iteration steps, which was the time taken for the reward function to converge (Fig. 5).

Action space - The preliminary manual control of the endoscope revealed that if the endoscope is directed against the colon wall, especially at the sharp turns, the lumen is not visible. Hence, it is critical to avoid the translation of the endoscope in such scenarios. Therefore, we develop an action strategy, where a translation motion with a constant velocity of $v_{end} = 10\text{mm/sec}$ is carried out only when the lumen is detected. The action space consists of discrete angular rotation values in the three degrees of freedom at the endoscope tip, $\delta\theta_j = \alpha$, $\alpha \in \{0, -1, +1\}$ in the j^{th} spatial dimension. In the tip local reference frame, $j \in x, y, z$ corresponds to the orientation alignment in the horizontal and vertical directions in the image plane and the endoscope roll, respectively (Fig. 3). In cases when the lumen is not visible, the translation velocity of the endoscope is set to zero, and the agent carries out orientation changes to detect the lumen.

Observation space and policy - The sensory input to the DVC agent is composed of a downscaled endoscopic image. The RGB images rendered by the endoscopic camera (1024x1024 pixels) is downscaled to 128x128 pixels. The policy π_ϕ is represented by a CNN architecture, consisting of two convolutional layers (Fig. 1) for encoding visual scene representations. The network details are publicly available on the project website ¹. The output of the convolutional layers are fed into a combination of fully connected layers and Long Short-Term Memory (LSTM) layer to represent time-dependent behaviour, each with 128 rectified units, followed by the linear connections to the output logits π_t for each action a_t and values estimate V_t . A softmax function transforms the logits to action probabilities. The complete network is trained end-to-end to acquire task-specific visual features.

Reward function - The goal of navigation is to reach the end of the colon without any significant complications. Visuomotor control should be able to track the colon during the whole procedure. Successful tracking requires the lumen centre P_L to be close to the image centre P_c . Hence, a dense reward $r_t(s_t, a_t)$ is designed as follows:

$$r_t(s_t, a_t) = \begin{cases} C(1 - (||P_L - P_c||_2 / D_{max})), & L = 1 \\ -1, & L = 0 \end{cases} \quad (1)$$

¹<https://github.com/Ameyapores/DVC>

TABLE I
NAVIGATION PARAMETERS USED FOR VALIDATION WITH THEIR DESCRIPTION

	Navigation metrics	Description
1	Time of insertion (TOI)	TOI is measured from the time point where the initial movement of the endoscope is detected to the time point when the caecum is reached.
2	Colon Wall Collision (CWC)	CWC refers to the scenario when the endoscope tip touches the colon wall. The average diameter of the colon models used is 5cm, hence a threshold of $\delta d = 1cm$ is decided to classify the deformation as CWC.
3	Normalised distance	Distance travelled is crucial as multiple backward motions, reversing the direction, can lead to suboptimal trajectories. The distance travelled is measured using the position values of the endoscope tip. This distance is normalised by the centerline distance of the colon model in order to compare among different colon models. Normalised distance above 1 indicates a path distance longer than the centerline path, while a normalised distance below 1 indicates a shorter path than the centerline was followed.
4	Average Lumen Distance (LD)	Lumen centralisation is believed to create smooth insertion trajectories hence the lumen distance in the image plane is recorded at each timepoint. This distance is normalised by the size of the image to get a value in $[0,1]$. Lumen distance value 0 denotes image centre (P_c) coincides with the detected lumen (P_L), and value 1 denotes that the detected lumen is at the farthest point.




Fig. 4. Proposed adaptive threshold segmentation pipeline for lumen detection. a) original RGB frame b) Image mask for the detected lumen (in green) c) distance vector between the image centre P_c and the centroid of the detected darkest regions P_L .

where $D_{max} = 1/2 * (Imagewidth) = 64$, is the normalisation factor which is the maximum distance possible, L represents the lumen detection flag, (1 denotes lumen detected, 0 denotes no lumen detected), the hyperparameter C is chosen as 1. Moreover, the agent is awarded a reward of +10 if the colon end is reached and -10 if it returns to the original starting point, to encourage the agent to move unidirectional towards the caecum. To detect the colon lumen in the endoscope image, we build a threshold segmentation algorithm that runs in real-time at 30fps based on [20], where the image is first segmented to detect the darkest and most distinct region, with the presumption that this area contains the distal lumen with high probability (Fig. 4). The segmentation is performed by converting the RGB image to greyscale and cropping a circular region centred with the image and a diameter equal to the image width to remove the vignette effect on the corner.

IV. EXPERIMENTAL EVALUATION

The experimental goal is to compare the navigation performance of the DVC agents, the baseline method of rule-based control [5] and the endoscopists. Hence, we create a pipeline where the position and orientation values of the endoscope, lumen distance in the image space, colon deformations, and camera image can be recorded within our developed simulator.




Fig. 5. Learning curve of DVC trained on varying complexity of colon, using three colon models. Cumulative reward is normalised in the range $[-1, 1]$. The shaded area spans the range of values obtained when training the agent starting from five different initialisation seeds.

Endoscopist data acquisition - A group of 20 expert endoscopists (more than four years of experience) were asked to make navigation attempts in the simulation scene developed in Sec. III-A. Due to the time constraints and COVID regulations, four colon models were selected considering the opinion of domain experts to represent progressively more complex scenarios (Fig. 2). Each endoscopist was instructed to navigate the colon models from the rectum to the caecum using a PlayStation (Sony Interactive Entertainment, USA) joystick device. The model C_0 , which conforms with the shape and size of the average human colon, was used to familiarise the endoscopists with the controls before initialising the trials. The trials started with endoscopist attempts on the C_1 colon, followed by randomised attempts on C_2 and C_3 colon. The randomness between C_2 and C_3 colon was introduced to identify performance bias based on the colon model.

Training DVC agents - We conduct three experiments to validate the DVC. First, our aim is to determine the

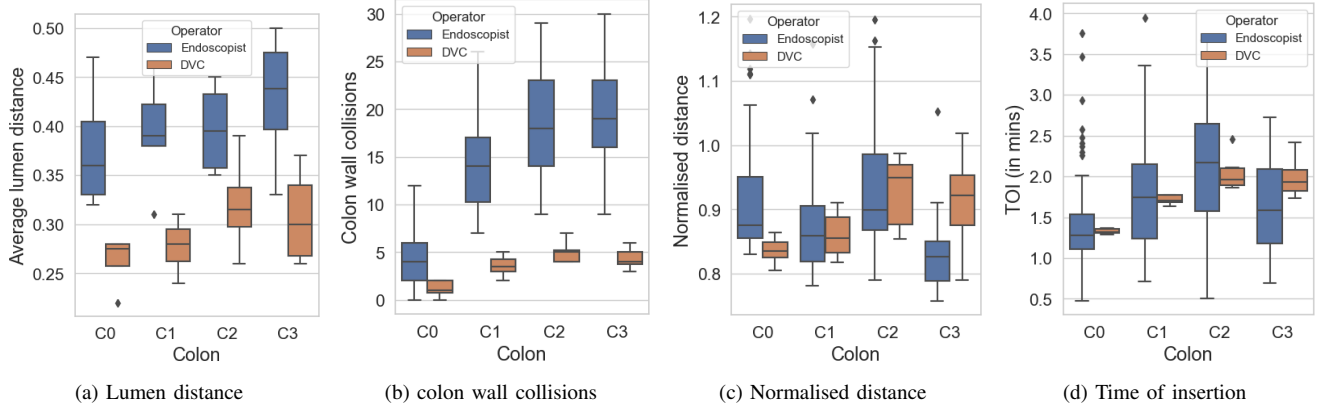


Fig. 6. Navigation performance comparison plots between DVC_{C_0} and endoscopists. Several parameters are plotted a) Average lumen distance, b) Colon wall collisions, c) Normalised distance, d) Time of insertion.

sample efficiency of training on different levels of colon complexity. Hence, we train DVC agents separately using the same models employed during the endoscopist experiment. Second, to establish a comparative analysis between the DVC and endoscopists, a similar experimental workflow was followed as in the endoscopist experiments, where the DVC was only trained on C_0 (DVC_{C_0}) and tested on C_1 , C_2 and C_3 colons. Third, the DVC was trained on the C_0 model followed by training on the C_1 ($DVC_{C_0+C_1}$) to test if training on a complex colon after a simple one improves performance. To keep the overall iteration steps for DVC training at 1.5 million, the training on C_0 was terminated after 1 million iteration steps and loaded back to train on C_1 for 500k iteration steps (Table. II).

Supervision - 20 novice participants (no endoscopy experience) were asked to supervise the performance of the rule-based controller agent and the DVC agent. The experimental workflow consisted of three trials: the participants attempted to navigate C_1 followed by randomised attempts on C_2 and C_3 colon model, similar to the endoscopists data acquisition. C_0 was used for initial training. Each trial was characterised by

- 1) Manual control: Participants were instructed to exclusively control the endoscope using a joystick during the entire duration of the procedure.
- 2) Rule-based baseline [5]: A proportional controller is generated for orientation control that aligns the image center (P_c) to the detected lumen (P_L) as follows:

$$\delta\theta = \beta \begin{bmatrix} P_{Lx} - P_{cx} \\ P_{Ly} - P_{cy} \end{bmatrix} \quad (2)$$

We refer to the distance between P_L and P_c as Lumen Distance (LD).

- 3) DVC: A fully trained DVC_{C_0} was deployed.

In control strategy 2, the rule-based controller indicates the requirement of manual supervision when the lumen centre is not detected. In control strategy 3, the agent is given ($\Delta_t = 50$) iteration steps to search for lumen when the lumen is not detected. After Δ_t steps, the DVC notifies the requirement of

human supervision, and manual control is activated. In both the control strategies, the user has an override option to take control when unsafe behaviour is encountered, e.g. collision with the colon wall or direction of motion reversed. Once the manual control is active, the participants can navigate the endoscope safely and give back the control to the DVC or rule-based controller. During each attempt, the number of interventions by the participant was recorded. A low-resolution (128x128 pixels) image is displayed to facilitate interpretability of machine decisions, however users have the option to change to high resolution (1024x1024 pixel) display. After all the trials, users were asked to complete a NASA Task Load Index (TLX) questionnaire [21], to score human-perceived workload.

Data Analysis - Four different parameters are used to score the navigation performance. Time of Insertion (TOI) and the number of Colon Wall Collisions (CWC) are qualitative assessment measures for colonoscopy procedures [22], while average LD and the normalised distance travelled are the two metrics devised in this study to measure the accuracy of the trajectories. The details of each parameter are elaborated in Table. I. When the user or DVC reversed its direction of motion and returned to the rectum or perforated heavily to destabilise the colon simulation, it was considered a failed navigation attempt.

V. RESULTS AND DISCUSSION

The learning curves when DVC is trained on different levels of colon complexity are presented in Fig. 5. C_0 represents a simplistic model; hence, the DVC agent reaches high reward values in relatively fewer steps than in other colon models. A high reward indicates that the agent successfully learns to complete the navigation task. Whereas C_2 represents high complexity, the agent requires 1.2 million steps for high-reward convergence. The C_1 training curve lies between C_0 and C_2 . This suggests that the training time is related to colon complexity. However, note that DVC_{C_0} can navigate other complex colon models, i.e. it acquires task-specific features that can generalise to other colon models (Table. II).

TABLE II
COMPARISON BETWEEN DVC_{C_0} AND $DVC_{C_0+C_1}$

	DVC_{C_0}				$DVC_{C_0+C_1}$			
	Average LD	CWC	TOI	Normalised distance	Average LD	CWC	TOI	Normalised distance
C_0	0.27 ± 0.01	0.5 ± 0.25	1.37 ± 0.05	0.84 ± 0.02	0.24 ± 0.02	1 ± 1	1.32 ± 0.03	0.84 ± 0.02
C_1	0.30 ± 0.01	3.3 ± 1.5	1.74 ± 0.04	0.88 ± 0.08	0.25 ± 0.01	3.3 ± 0.5	1.70 ± 0.07	0.82 ± 0.03
C_2	0.36 ± 0.03	5 ± 1	2.22 ± 0.2	0.97 ± 0.03	0.28 ± 0.01	4.6 ± 0.5	1.89 ± 0.03	0.85 ± 0.01
C_3	0.35 ± 0.02	4.6 ± 0.5	2.15 ± 0.23	0.92 ± 0.08	0.29 ± 0.03	3 ± 1	1.78 ± 0.05	0.89 ± 0.09
Mean	0.31 ± 0.04	5.0 ± 1.2	2.20 ± 0.75	0.90 ± 0.04	0.23 ± 0.04	4.3 ± 1.2	1.96 ± 0.59	0.86 ± 0.04

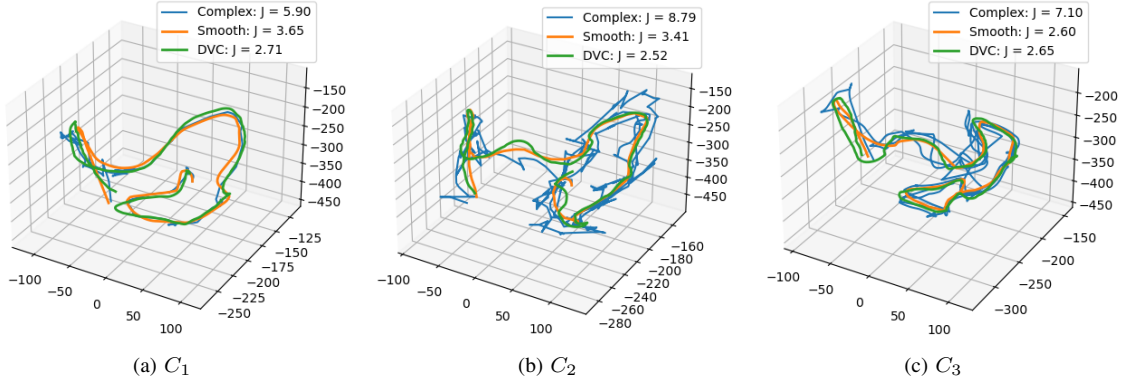


Fig. 7. Trajectory plot of DVC, complex and smoothest endoscopist performance for a) C_1 b) C_2 c) C_3 models respectively.

TABLE III
NASA TASK LOAD INDEX FOR NOVICE USERS.

	Manual control	Rule-based control	DVC
Mental demand	63	33	18
Physical demand	65	38	9
Temporal demand	30	47	17
Performance	25	34	12
Effort	57	38	10
Frustration	42	41	12
Mean workload	47	38	13

Comparative analysis - The performance data of 20 endoscopists was acquired while 10 different DVC agents were trained on the C_0 starting with a different random seed. The realism of the simulation was confirmed and validated by expert clinicians. Additionally, all the users positively evaluated the joystick used to navigate the endoscope inside the colon as intuitive, user-friendly and easy to learn. Fig. 6 shows the comparison of the average LD, the number of CWC, the completion time and the normalised distance travelled. There is a significant difference in the average LD and the number of CWC between the endoscopists and the DVC. DVC shows precise tip centralisation and less number of CWC compared to endoscopists. One of the reasons for this difference is that clinicians tend to push the colon wall at acute bends of colon junctions (see supplementary video). This is a gesture sometimes clinicians follow due to the rigid constraints of

the clinically available FEs. Whereas DVC is trained on reward feedback to minimise LD, it stays centralised to avoid contact with the wall. For the normalised distance and TOI, a substantial difference is not noted. There is more variance observed in the performance of the endoscopists. Some endoscopists followed a convoluted trajectory that increased the normalised distance and time of insertion, while others followed smoother trajectories that resulted in the lower normalised distance and TOI. Fig. 7 shows the most complex and smoothest trajectories demonstrated by the endoscopists and the trajectory executed by DVC_{C_0} for the C_1 , C_2 and C_3 colons. The smoothness of a trajectory is estimated using a jerk index J (cm/sec^3) which characterises the average rate of change of acceleration in a movement [23]. Human operators tend to show wide variance in performing optimal trajectories, while DVC performance stays in the average range.

The result of splitting the training into two colon models $DVC_{C_0+C_1}$ and evaluating on other colon models are shown in Table. II. There is an improvement in the lumen detection performance for $DVC_{C_0+C_1}$ in comparison to DVC_{C_0} . DVC_{C_0} reaches high rewards at 500k iteration steps; hence there is no additional feedback to improve the performance. We speculate that the agent reaches suboptimal local minima, while, when the DVC trained on C_0 is loaded to train on C_1 , it encounters acute bends, which offers the potential to maximise the cumulative reward. There is no considerable improvement on other navigation parameters, i.e. CWC, TOI

and normalised distance.

Supervision - The human interventions are divided into two parts. First, where the user overrides the control due to unsafe behaviour and second, where the system demands human supervision. The average human intervention required for rule-based baseline was 5 ± 1.8 for human override and 2.5 ± 1.5 when the system demanded human control, while for DVC, the average number of human interventions are 0.1 ± 0.5 for human override and 0.05 ± 0.2 when the system demanded human control. This difference is attributed to DVC's adaptability to search for new insertion directions when the lumen is not easily detected, whereas the rule-based controller lacks this ability. The NASA-TLX for each control strategy is shown in Table. III. Regarding ease of use, participants found manual control and rule-based controller more demanding in all task load categories, while a substantial workload reduction is observed for DVC.

VI. CONCLUSIONS

Prior works on autonomous colonoscopy navigation use heuristic control policies that fail to adapt to situations where detecting lumen is not straightforward and requires frequent human intervention. We propose a DVC method that learns a mapping between the endoscopic images and the control signal. Motion data from 20 endoscopists was acquired and compared to DVC control. Our performance evaluation shows an equivalent performance in terms of the time of insertion and the distance travelled. However, DVC reduces the number of colon wall collisions and shows efficient lumen tracking, improving safety. Furthermore, we conducted a second novice user study to demonstrate that supervision of DVC control significantly reduces the user workload with overall performance comparable to expert endoscopists.

Our future work will demonstrate the formal validation of the proposed virtual simulator. Moreover, the proposed strategy will be tested on the real robotic FE system. Hence, some simplified assumptions made in this simulation study need to be reconsidered in the learning loop. For example, when controlling a real robotic FE it is common that the required movements are not as smooth as in the simulator and they are not performed immediately due to friction and the conformation of the colon walls.

REFERENCES

- [1] C. Chen, M. Hoffmeister, and H. Brenner, "The toll of not screening for colorectal cancer," *Expert review of gastroenterology & hepatology*, vol. 11, no. 1, pp. 1–3, 2017.
- [2] G. Ciuti, K. Skonieczna-Żydecka, W. Marlicz, V. Iacovacci, H. Liu, D. Stoyanov, A. Arezzo, M. Chiurazzi, E. Toth, H. Thorlacius *et al.*, "Frontiers of robotic colonoscopy: a comprehensive review of robotic colonoscopes and technologies," *Journal of Clinical Medicine*, vol. 9, no. 6, p. 1648, 2020.
- [3] H.-E. Huang, S.-Y. Yen, C.-F. Chu, F.-M. Suk, G.-S. Lien, and C.-W. Liu, "Autonomous navigation of a magnetic colonoscope using force sensing and a heuristic search algorithm," *Scientific reports*, vol. 11, no. 1, pp. 1–15, 2021.
- [4] A. Z. Taddese, P. R. Slawinski, K. L. Obstein, and P. Valdastri, "Non-holonomic closed-loop velocity control of a soft-tethered magnetic capsule endoscope," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2016, pp. 1139–1144.
- [5] J. W. Martin, B. Scaglioni, J. C. Norton, V. Subramanian, A. Arezzo, K. L. Obstein, and P. Valdastri, "Enabling the future of colonoscopy with intelligent and autonomous magnetic manipulation," *Nature machine intelligence*, vol. 2, no. 10, pp. 595–606, 2020.
- [6] 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.
- [7] N. v. d. Stap, C. H. Slump, I. A. Broeders, and F. v. d. Heijden, "Image-based navigation for a robotized flexible endoscope," in *International Workshop on Computer-Assisted and Robotic Endoscopy*. Springer, 2014, pp. 77–87.
- [8] F. Richter, R. K. Orosco, and M. C. Yip, "Open-sourced reinforcement learning environments for surgical robotics," *arXiv preprint arXiv:1903.02090*, 2019.
- [9] E. Tagliabue, A. Pore, D. Dall'Alba, E. Magnabosco, M. Piccinelli, and P. Fiorini, "Soft tissue simulation environment to learn manipulation tasks in autonomous robotic surgery," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 3261–3266.
- [10] A. Pore, E. Tagliabue, M. Piccinelli, D. Dall'Alba, A. Casals, and P. Fiorini, "Learning from demonstrations for autonomous soft-tissue retraction," in *2021 International Symposium on Medical Robotics (ISMR)*. IEEE, 2021, pp. 1–7.
- [11] G. Trovato, M. Shikanai, G. Ukawa, J. Kinoshita, N. Murai, J. Lee, H. Ishii, A. Takanishi, K. Tanoue, S. Ieiri *et al.*, "Development of a colon endoscope robot that adjusts its locomotion through the use of reinforcement learning," *International journal of computer assisted radiology and surgery*, vol. 5, no. 4, pp. 317–325, 2010.
- [12] J. Ibarz, J. Tan, C. Finn, M. Kalakrishnan, P. Pastor, and S. Levine, "How to train your robot with deep reinforcement learning: lessons we have learned," *The International Journal of Robotics Research*, vol. 40, no. 4–5, pp. 698–721, 2021.
- [13] J. Xu, B. Li, B. Lu, Y.-H. Liu, Q. Dou, and P.-A. Heng, "Surrol: An open-source reinforcement learning centered and dvrc compatible platform for surgical robot learning," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 1821–1828.
- [14] P. M. Scheikl, B. Gyenes, T. Davitashvili, R. Younis, A. Schulze, B. P. Müller-Stich, G. Neumann, M. Wagner, and F. Mathis-Ullrich, "Cooperative assistance in robotic surgery through multi-agent reinforcement learning," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 1859–1864.
- [15] Y.-H. Su, K. Huang, and B. Hannaford, "Multicamera 3d viewpoint adjustment for robotic surgery via deep reinforcement learning," *Journal of Medical Robotics Research*, vol. 6, no. 01n02, p. 2140003, 2021.
- [16] K. İncetan, I. O. Celik, A. Obeid, G. I. Gokceler, K. B. Ozyuruk, Y. Almalioglu, R. J. Chen, F. Mahmood, H. Gilbert, N. J. Durr *et al.*, "Vr-caps: A virtual environment for capsule endoscopy," *Medical image analysis*, vol. 70, p. 101990, 2021.
- [17] M. B. Christensen, K. Oberg, and J. C. Wolchok, "Tensile properties of the rectal and sigmoid colon: a comparative analysis of human and porcine tissue," *Springerplus*, vol. 4, no. 1, pp. 1–10, 2015.
- [18] Y. Wang, Y. He, X. Deng, Z. Lei, Y. Chen, and M. Li, "Learning friction model for tethered capsule robot," in *2021 6th International Conference on Robotics and Automation Engineering (ICRAE)*. IEEE, 2021, pp. 280–285.
- [19] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [20] D. Wang, X. Xie, G. Li, Z. Yin, and Z. Wang, "A lumen detection-based intestinal direction vector acquisition method for wireless endoscopy systems," *IEEE Transactions on Biomedical Engineering*, vol. 62, no. 3, pp. 807–819, 2014.
- [21] S. G. Hart and L. E. Staveland, "Development of nasa-tlx (task load index): Results of empirical and theoretical research," in *Advances in psychology*. Elsevier, 1988, vol. 52, pp. 139–183.
- [22] M. F. Kaminski, S. Thomas-Gibson, M. Bugajski, M. Bretthauer, C. J. Rees, E. Dekker, G. Hoff, R. Jover, S. Suchanek, M. Ferlitsch *et al.*, "Performance measures for lower gastrointestinal endoscopy: a european society of gastrointestinal endoscopy (esge) quality improvement initiative," *Endoscopy*, vol. 49, no. 04, pp. 378–397, 2017.
- [23] R. Shadmehr and S. Wise, "Supplementary documents for "computational neurobiology of reaching and pointing,"" 2005.