

# Supervised Mixture-of-Experts for Surgical Grasping and Retraction

Lorenzo Mazza<sup>†\*1,2,3</sup>, Ariel Rodriguez<sup>†\*1,2,3</sup>, Rayan Younis<sup>4,5</sup>, Martin Lelis<sup>1,2,3</sup>,  
Ortrun Hellig<sup>4,5</sup>, Chenpan Li<sup>1,2</sup>, Sebastian Bodenstedt<sup>1,2,3,4</sup>, Martin Wagner<sup>4,5</sup>, Stefanie Speidel<sup>1,2,3,4</sup>

<sup>†</sup>Equal contribution

<sup>1</sup>Department of Translational Surgical Oncology, NCT/UCC Dresden,

Faculty of Medicine and University Hospital Carl Gustav Carus, TUD, Dresden, Germany

<sup>2</sup>National Center for Tumor Diseases (NCT), NCT/UCC Dresden, a partnership between DKFZ,  
Faculty of Medicine and University Hospital Carl Gustav Carus, TUD, and HZDR, Dresden, Germany.

<sup>3</sup>Faculty of Computer Science, TUD, Dresden, Germany.

<sup>4</sup>The Center for Tactile Internet with Human-in-the-Loop (CeTI), TUD, Dresden, Germany

<sup>5</sup>Department of Visceral, Thoracic and Vascular Surgery,

Faculty of Medicine and University Hospital Carl Gustav Carus, TUD, Dresden, Germany

\*Corresponding authors: lorenzo.mazza@nct-dresden.de, ariel.rodriguezjimenez@nct-dresden.de




Fig. 1: Collaborative autonomous bowel grasping and retraction policy roll-outs across phantom, *ex vivo*, and preliminary *in vivo* porcine settings. The surgeon-controlled gripper (bottom instrument) indicates the target bowel segment, marked by the red rhomboid, while the robot assistant gripper (top instrument) autonomously performs grasping, retraction, and tension maintenance.

**Abstract**—Imitation learning has achieved remarkable success in robotic manipulation, yet its application to surgical robotics remains challenging due to data scarcity, constrained workspaces, and the need for an exceptional level of safety and predictability. We present a supervised Mixture-of-Experts (MoE) architecture designed for phase-structured surgical manipulation tasks, which can be added on top of any autonomous policy. Unlike prior surgical robot learning approaches that rely on multi-camera setups or thousands of demonstrations, we show that a lightweight action decoder policy like Action Chunking Transformer (ACT) can learn complex, long-horizon manipulation from less than 150 demonstrations using solely stereo endoscopic images, when equipped with our architecture. We evaluate our approach on the collaborative surgical task of bowel grasping and retraction, where a robot assistant interprets visual cues from a human surgeon, executes targeted grasping on deformable tissue, and performs sustained retraction. We benchmark our method against state-of-the-art Vision-Language-Action (VLA) models and the standard ACT baseline. Our results show that generalist VLAs fail to acquire the task entirely, even under standard in-distribution conditions. Furthermore, while standard ACT achieves moderate success in-distribution, adopting a supervised MoE architecture significantly boosts its performance, yielding higher success rates in-distribution and demonstrating superior robustness in out-of-distribution scenarios, including novel grasp locations, reduced illumination, and partial occlusions. Notably, it generalizes to unseen testing viewpoints and also transfers zero-shot to *ex vivo* porcine tissue without additional training, offering a promising pathway toward *in vivo* deployment. To support this statement, we present qualitative preliminary results of policy roll-outs during *in vivo* porcine surgery. These results demonstrate that supervised MoE architectures provide a data-efficient approach for learning multi-step dexterous manipulation in visually constrained environments. Code and dataset are available.<sup>1</sup>

## I. INTRODUCTION

Imitation learning (IL) has shown remarkable results in learning manipulation tasks through generative modeling approaches [35, 6]. Recently, Vision-Language-Action (VLA) models [3, 25, 20, 2, 28, 31, 19] have achieved impressive performance by leveraging large-scale datasets [5, 26, 16], demonstrating that foundation models can enable generalist robot policies across diverse tasks and embodiments.

Minimally-invasive surgery (MIS) stands out as a particularly impactful application domain for autonomous manipulation. Staff shortages are expected to worsen relative to the growing surgical treatment needs of our ageing society worldwide [27], creating an urgent need for autonomous surgical assistance. Robot policies show great potential to address this shortcoming by enabling intraoperative autonomous assistance [23, 33], yet several challenges limit the direct adoption of general-purpose IL approaches. Demonstration data is scarce due to ethical and regulatory constraints, inability to repeat procedures purely for data collection, and the prohibitive costs of operating room time and expert surgeon involvement. Data quality is further compromised by noise, due to occlusions, limited control over recording conditions, and the inherent variability of scenes in surgical procedures. Furthermore, workspace constraints preclude multi-view camera or depth

sensor setups, tissue deformation adds complex dynamics not present in rigid object manipulation, and the proximity to delicate anatomical structures demands an exceptional level of safety and predictability. Surgical policies must additionally satisfy strict deployment requirements: lightweight architectures that fit on compact hardware with low inference latency to enable real-time control on resource-constrained systems.

These constraints preclude the use of large pretrained VLA models, which i) require to be trained on large-scale datasets, due to their high parameter count and ii) incur high computational overhead during inference. Recent benchmarks in precision surgical tasks, such as end-to-end suturing [11], confirm that compact, lightweight policies like Action Chunking Transformers (ACT) [35] outperform significantly VLAs when trained on limited surgical data, achieving considerably higher success rates and faster inference.

Similarly, Surgical Robotics Transformer (SRT) [17] and its hierarchical variant SRT-H [18] have demonstrated that lightweight action transformer policies can learn dexterous multi-step surgical manipulation tasks from visual observations, proving that long-horizon precision tasks can be learned in a data-driven manner. However, key challenges remain unsolved. First, these approaches rely on multi-camera setups — including wrist-mounted cameras — to ensure robust 3D scene understanding, configurations that are often infeasible in MIS settings where only a single endoscopic view of the scene is available. Second, they still require extensive demonstration datasets: for instance, SRT-H’s gallbladder clipping and cutting required approximately 16,000 demonstrations to achieve reliable performance. To address these shortcomings, we propose a supervised Mixture-of-Experts (MoE) extension to action transformer policies. MoE architectures offer a promising framework for modeling multi-step, long-horizon surgical tasks by employing specialized sub-networks (experts) that handle different aspects of the task space [13, 14]. In robotics, MoE has shown promise for learning diverse skills and handling multi-modal action distributions [29, 32]. The key insight is that complex tasks can be decomposed into simpler sub-components, each handled by a dedicated expert, with a gating mechanism determining which experts to activate based on the current context. This decomposition is particularly relevant for surgical tasks, where phase transitions are often well-defined and observable [24, 10]. However, training MoEs end-to-end is notoriously unstable, often suffering from mode collapse or expert underutilization where the gating mechanism fails to effectively distribute task dynamics [36]. We overcome these optimization challenges by exploiting the ordered phase structure of surgical sub-tasks, supervising explicitly the gating network with phase labels, ensuring stable convergence and clear functional specialization for each expert.

Assistant tissue manipulation tasks in MIS mostly consist of tissue grasping and retraction, generally performed under the guidance of the operating surgeon [8, 7]. In MIS for gastroin-

<sup>1</sup> Available at <https://surgical-moe-project.github.io/rss-paper/>.

testinal cancer treatment, surgical assistants must manipulate the small bowel to enable the surgeon to perform anastomoses on the jejunum and ileum [22]. This highlights the need to automate the assistant’s role in small bowel manipulation, in particular bowel grasping and retraction, while maintaining coordinated interaction with the operating surgeon.

In summary, our major contributions are the following:

- We propose a novel supervised Mixture of Experts (MoE) architecture designed for phase-structured surgical tasks that can be integrated into any kind of action transformer policies. We apply our architecture to a lightweight policy such as ACT and show that it can learn from significantly fewer demonstrations than prior work, and relying solely on endoscopic visual feedback — without wrist cameras or multi-view setups — for practical deployment in clinical MIS environments.
- We introduce a novel surgeon-robot collaboration task in laparoscopic bowel retraction, where a human surgeon provides high-level visual cues via a laparoscopic instrument, and the robot executes precise grasping, pulling, and sustained retraction actions. This cooperative paradigm emphasizes human-robot teamwork to enhance efficiency in MIS, where the robot serves as an intelligent assistant handling secondary but crucial tasks, i.e. maintains tissue retraction and tension, while the surgeon focus on critical actions — such as performing suturing or bowel anastomosis.
- We empirically validate the limitations of current state-of-the-art VLAs in the surgical domain. Our results reinforce recent findings that generalist foundation models fail to acquire high-precision surgical policies in data-scarce regimes, establishing the necessity of specialized, lightweight architectures for robust surgical automation.
- We demonstrate two key prerequisite for *in vivo* translation of autonomous surgical policies: (i) viewpoint invariance, showing that training with randomized camera angles enables the policy to generalize to unseen viewpoints without explicit 3D representations; and (ii) zero-shot transfer, where the policy achieves an 80% success rate on *ex vivo* porcine tissue despite being trained solely on phantom data. These results validate the system’s robustness to the geometric variations and visual domain shifts inherent in dynamic clinical environments.

Additionally, we release our code and dataset to facilitate reproducibility and future research in surgical robotics <sup>2</sup>.

## II. METHODS

### A. Hardware and Experimental Setup

We develop an experimental setup using the OpenHELP open-body phantom [15], as shown in figure 2. We use two UR5e industrial robotic arms: one arm remains static and is equipped with a stereo TIPCAM1 S 3D endoscope (Karl Storz SE & Co. KG) to provide visual feedback, while the other is equipped with a mechatronic interface [30] that allows the




Fig. 2: Experimental setup using the OpenHELP open-body phantom, showing the phantom with two robotic arms, one holding the laparoscope and one holding the surgical instrument. The abdominal wall cover is removed for visibility purposes.

attachment of a laparoscopic surgical bowel grasper, allowing controlled opening and closing of the gripper. The latter robot moves while maintaining the remote-center-of-motion (RCM) constraint and is operated via a joystick-based control interface.

### B. Data Collection

TABLE I: Task phase segmentation for bowel grasping and retraction.

| Phase | Name             | Description                                                                        |
|-------|------------------|------------------------------------------------------------------------------------|
| 1     | Idle             | Robotic instrument visible; surgeon’s instrument may or may not be present.        |
| 2     | Approach & Grasp | Surgeon indicates target; robot opens grasper, moves to target, and grasps tissue. |
| 3     | Hold             | Robot holds bowel stationary until surgeon grasps opposite end.                    |
| 4     | Retract          | Robot retracts bowel until visually straight and free of folds.                    |
| 5     | Maintain Tension | Robot maintains tension with no further movement while surgeon operates.           |

We record a total of 120 episodes with a fixed viewpoint of the scene from the endoscope. We refer to this dataset as the *fixed-viewpoint dataset*. During each trial, we record stereo image pairs from the endoscope, the binary state of the mechatronic interface (gripper open/closed), and the three-dimensional position of the instrument tip in the camera coordinate system. We segment the bowel retraction task into  $H = 5$  phases as shown in table I, with transitions primarily triggered by the surgeon’s actions. The phases in the dataset are automatically labeled using the gripper state and the magnitude of movement of the robot instrument as a proxy for the transitions. To improve the robustness of downstream policies, we introduce variability across the training demonstrations, including different starting points for the assistant tool position, different grasping locations indicated by the surgeon and slight movements of the bowel across the phantom scene. This ensures a diverse set of trajectories and visual features in the dataset. To evaluate generalization to

<sup>2</sup>Available at <https://surgical-moe-project.github.io/rss-paper/>.






Fig. 3: Overview of the supervised Mixture-of-Experts extension to ACT. Stereo endoscopic images are processed by the ACT backbone to produce latent visual features and action tokens. A phase-aware gating network, supervised using surgical phase labels during training, routes features to specialized experts corresponding to different task phases (idle, grasping, holding, retraction, and tension maintenance). The outputs of the action and gripper experts are combined through phase-weighted aggregation to generate the action chunk.

viewpoint variation, we additionally collect 50 episodes with randomized endoscopic camera angles. We refer to this dataset as the *random-viewpoint dataset*.

### C. Observation and Action Space

At time  $t$ , the observation space of the policy consists of the state  $s_t = (I_t^{\text{left}}, I_t^{\text{right}})$  containing **only the stereo endoscopic image pair at the current time**. Although available, we explicitly exclude proprioceptive data to derive a vision-only policy. This design decouples performance from kinematic sensor noise and calibration drift, and makes it independent of the underlying surgical hardware quality [17]. The action space, similarly to [35], comprises chunks of  $k$  continuous actions  $a_{t:t+k}$ , where each  $a_{t+i} \in \mathbb{R}^3$  represents the delta movement of the instrument tip in the Cartesian space of the camera coordinate system, plus chunks of binary gripper actions  $g_{t:t+k}$ , where each  $g_{t+i} \in \{0, 1\}$ . We additionally denote ground truth phase label  $h \in H$  at time  $t$  with  $\psi_t = h$ . For action notations, we indicate ground truth using symbol  $a^*$  and predictions using  $\hat{a}$ .

### D. Policy

We propose a supervised Mixture-of-Experts (MoE) architecture that is modular and can be integrated into any transformer-based action chunking policy to leverage explicit

task phase structure. Here, we apply it to the lightweight Action Chunking Transformer (ACT) [35], chosen for its efficiency, low latency, and strong performance on data-limited surgical tasks [11]. The base follows ACT’s variational framework: at training time, given a state observation  $s_t$  and a chunk of size  $k$  of ground-truth training actions  $a_{t:t+k}^*$ , a posterior encoder  $q_\phi(z|a_{t:t+k}^*)$  infers latent  $z$ , which is concatenated with visual features and processed through a transformer encoder-decoder with parameters  $\theta$  to reconstruct action chunk  $\hat{a}_{t:t+k}$ . At inference, the posterior encoder is discarded and  $z$  is set to the mean of a Gaussian prior  $p(z) \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ . We extend this with a Phase-Aware MoE block comprising  $H$  parallel experts (one per phase), where each phase expert models the action distribution of a specific task phase, conditioned additionally on  $z$  and  $s_t$ . The block is composed as following:

- 1) **Action Phase-Experts:**  $H$  action heads, where expert  $h$  outputs location parameters  $\mu_{h,t:t+k}(z, s_t, \psi = h) \in \mathbb{R}^{k \times d}$  for a  $d$ -dimensional action space over a chunk of length  $k$ .
- 2) **Gripper Phase-Experts:**  $H$  gripper heads, where expert  $h$  outputs logits  $\nu_{h,t:t+k}(z, s_t, \psi = \psi_h) \in \mathbb{R}^k$ , parameterizing Bernoulli distributions with  $p_h(g_{t+j} = 1 | z, s_t, \psi = h) = \sigma(\nu_{h,t+j})$  for each  $j^{\text{th}}$  gripper action of the chunk.



- 3) **Gating Network:** A phase classifier that models the categorical distributions  $\pi_{t:t+k}(z, s_t)$ , where  $\pi_{h,t+j} = p(\psi_{t+j} = h \mid z, s_t)$  and  $\sum_{h=1}^H \pi_{h,t+j} = 1$ .

Final predictions are phase-weighted mixtures:

$$\hat{a}_{t+j} = \sum_{h=1}^H \pi_{h,t+j} \cdot \mu_{h,t+j}$$

$$\hat{g}_{t+j} = \sum_{h=1}^H \pi_{h,t+j} \cdot \sigma(\nu_{h,t+j})$$

for each  $(t+j)_{th}$  action of the chunk. The architecture is illustrated in figure 3.

### E. Training Procedure

Following the conditional variational framework of ACT [35], training optimizes a variational lower bound (ELBO) on the log-likelihood of demonstration trajectories [34], where the phase labels  $\psi_{t:t+k}^*$  are observed during training. This supervised approach leverages privileged phase information to guide expert specialization during training. By assuming Laplace-distributed action errors, Bernoulli gripper states, and categorical phase distributions, we obtain the following training objective (1), which comprises four components: i) the action reconstruction loss (L1) trains the weighted mixture of experts to match demonstration actions and provides robustness against outliers in human demonstrations [35]; ii) the phase cross-entropy loss (CE) directly supervises the gating network, treating phase prediction as an auxiliary task, guiding the MoE experts to specialize per task phase; iii) the gripper binary cross-entropy loss (BCE) trains the weighted Bernoulli mixture for discrete gripper actions; iv) the KL term regularizes the learned amortized posterior encoder [21, 4, 12], that we denote as  $q_\phi$ .

$$\begin{aligned} \mathcal{L}(\theta, \phi) = & \alpha \sum_{j=0}^{k-1} \|\hat{a}_{t+j} - a_{t+j}^*\|_1 + \gamma \sum_{j=0}^{k-1} \text{CE}(\pi_{t+j}, \psi_{t+j}^*) \\ & + \delta \sum_{j=0}^{k-1} \text{BCE}(\hat{g}_{t+j}, g_{t+j}^*) + \beta \cdot D_{\text{KL}}(q_\phi \| p(z)) \quad (1) \end{aligned}$$

As benchmarks, we train ACT, SmolVLA and  $\pi_{0.5}$  with standard hyperparameters, using the open-source LeRobot codebase [1]. As SmolVLA and  $\pi_{0.5}$  include language instruction and proprioceptive state in their input space, we extend  $s_t$  with a fixed language instruction and a padding proprioceptive vector of zeros. Table II summarizes the evaluated models, while additional implementation details are provided in the supplementary material. We train and deploy each policy on a single RTX5000 Nvidia GPU, with the exception of  $\pi_{0.5}$  which required one A100 Nvidia GPU for training. The training time of SmolVLA was 14 hours, 8 hours for  $\pi_{0.5}$  and 3 hours for ACT and ACT+MoE.

TABLE II: Comparison of model variants, parameter counts, and training frameworks.

| Policy        | Variant | Parameters | Codebase |
|---------------|---------|------------|----------|
| $\pi_{0.5}$   | Base    | 4 B        | LeRobot  |
| SmolVLA       | Small   | 0.24 B     | LeRobot  |
| ACT+MoE(Ours) | –       | 53.3 M     | Ours     |
| ACT           | –       | 52 M       | Ours     |

## III. EXPERIMENTS

To evaluate the success rate of the learned policies on the robotic platform, two trained medical students and one surgical resident review each policy roll-out in a single-blinded process and label the final frame as either a success or a failure. A success in the final frame is defined as tissue grasped by two graspers and retracted with sufficient tension on the bowel segment. The final outcome of each roll-out is then determined by majority voting. We first evaluate the policy trained on the fixed viewpoint dataset using environment conditions as close as possible to the training scene and the same camera angle. We denote these tests as *in-distribution*. Consequently, we assess generalization capabilities by testing four out-of-distribution conditions: i) grasping bowel sections not seen during training, ii) operating under severely reduced scene illumination, iii) handling partial occlusions from phantom fat, and iv) using a slightly different camera angle compared to the training data. We denote these policy roll-outs as *out-of-distribution*. For the in-distribution evaluation, the endoscopic field of view was divided into 12 angular sectors. The robotic instrument was initialized from 20 distinct starting positions distributed across these sectors. Each starting position was paired with one of eight distinct grasping points on the bowel, resulting in 20 unique starting position and grasping point combinations per policy. For the out-of-distribution evaluation, we performed five trials for each perturbation type, sampling starting positions from the same 20-position set used in-distribution. The novel-grasp-location perturbation differed only in the target grasp locations: the surgeon indicated targets on a third bowel segment at image coordinates not observed during training. This setting directly tests whether the policy follows the surgeon’s visual cue rather than memorizing a fixed spatial prior. The best-performing policy (based on success rate) is tested zero-shot on *ex vivo* porcine bowel, replacing the phantom bowel in the experimental setup. We conduct 15 trials in the *ex vivo* configuration to evaluate whether the learned policies transfer to real tissue without additional training. Additionally we train again the best performing policy on the random viewpoint dataset, and test it on unseen camera angles, to verify its robustness to the 3D scene in prevision of *in vivo* conditions.

## IV. RESULTS AND DISCUSSION

### A. In-Distribution Roll-Outs

We report the success rates of the policies for in-distribution experiments in table III.

| Policy           | In-Distribution |               |               |               |                     |
|------------------|-----------------|---------------|---------------|---------------|---------------------|
|                  | Reaching        | Grasping      | Retracting    | End-To-End    | Inference Frequency |
| ACT              | 16/20           | 13/20         | 12/20         | 10/20         | 27 Hz               |
| $\pi_{0.5}$      | 2/20            | 4/20          | 0/20          | 0/20          | 10 Hz (compiled)    |
| SmolVLA          | 5/20            | 2/20          | 0/20          | 0/20          | 3.3 Hz              |
| ACT + MoE (Ours) | <b>20/20</b>    | <b>20/20*</b> | <b>19/20*</b> | <b>17/20*</b> | 27 Hz               |

TABLE III: Success rates (divided per sub-task) of policies trained on the fixed viewpoint dataset. In-distribution roll-outs inside the phantom environment.

\* indicates statistically significant improvement of ACT + MoE over the ACT baseline (two-sided Fisher’s exact test,  $p < 0.05$ ; exact  $p$ -values: Grasping  $p = 0.008$ , Retracting  $p = 0.020$ , End-To-End  $p = 0.041$ ). Improvements of ACT + MoE over both  $\pi_{0.5}$  and SmolVLA are highly significant ( $p < 10^{-7}$ ) on all reported metrics.

Our benchmarking reveals a significant performance gap between generalist VLA models and specialized action transformers. Both VLA baselines fail to complete the task end-to-end. SmolVLA proves unable to model the trajectory dynamics, producing erratic and dangerous actions against the target anatomy. While  $\pi_{0.5}$  demonstrates a slight improvement in grasping capabilities, it suffers from severe temporal incoherence; the model frequently violates task phase constraints, initiating retraction motions before securing the grasp or anticipating the surgeon’s handover prematurely. Consequently, it achieves a 0% end-to-end success rate. We argue that both models fail to model the task due to the limited amount of training data available compared to the actual number of parameters that they need to optimize for, shown in Table II.

In contrast, the standard ACT baseline demonstrates reasonable competency, achieving a 50% success rate. However, it lacks fine-grained dexterity, frequently resulting in tissue slippage or imprecise end-effector positioning during critical phase transitions. Incorporating our Supervised MoE architecture yields a significant performance gain. The specialized experts enable precise phase handling, boosting the grasping success rate from 60% (ACT) to 85% and the overall end-to-end success rate from 50% to 85%. This represents a 70% *relative improvement* over the standard ACT baseline, confirming that explicit expert supervision significantly enhances policy robustness and dexterity in data-constrained regimes. In terms of computational efficiency, our approach maintains the real-time applicability of the base architecture. ACT + MoE policy operates at 27 Hz, incurring negligible inference overhead compared to standard ACT. Conversely, VLA baselines exhibit significantly higher latency—with  $\pi_{0.5}$  running at 10 Hz and SmolVLA at 3.3 Hz—rendering them impractical for the high-frequency control loops required in surgical automation.

### B. Out-of-Distribution Roll-Outs

Given the complete failure of VLA models in the standard setting, we exclude them from further evaluation. We restrict the out-of-distribution (OOD) analysis to a comparison between the standard ACT baseline and our MoE-augmented policy. The results are detailed in Table IV.

| Policy           | Out-of-Distribution |              |              |              |
|------------------|---------------------|--------------|--------------|--------------|
|                  | Reaching            | Grasping     | Retracting   | End-To-End   |
| ACT              | 16/20               | 13/20        | 6/20         | 6/20         |
| ACT + MoE (Ours) | <b>19/20</b>        | <b>16/20</b> | <b>13/20</b> | <b>13/20</b> |

TABLE IV: Success rates (divided per sub-task) of policies trained on the fixed viewpoint dataset. OOD roll-outs inside the phantom environment.

two-sided Fisher’s exact test  $p$ -values: Reaching  $p = 0.342$ , Grasping  $p = 0.480$ , Retracting  $p = 0.056$ , End-To-End  $p = 0.056$ .

### C. Ex Vivo Zero-Shot Roll-Outs

Motivated by the superior out-of-distribution performance of our supervised MoE-ACT, we select it for the next two tests: i) zero-shot testing on *ex vivo* porcine bowel and ii) retraining on the random-viewpoint dataset and testing on unseen viewpoints. The model achieves an **80%** success rate on *ex vivo* porcine bowel (**12/15**) (Figure 1, central panel), validating its generalization capabilities. Of the three failures, two were due to grasping two bowels at the same time — but still completing the retraction — and only one consisted of a complete failure.

### D. Random Viewpoint Roll-Outs

We further evaluate the ability of MoE-ACT to generalize to diverse camera viewpoints. We train the policy on the joint combination of the fixed and random-viewpoint datasets, to simulate the more realistic viewing conditions of clinical *in vivo* procedures, where precise camera positioning cannot be fixed or known a priori.

After retraining, our policy roll-outs achieve an 82% success rate (18/22) on unseen testing viewpoints, demonstrating robust performance and implicit 3D scene understanding. In figure 4 we show representative examples of the diverse set of camera angles used during testing, with considerable variations in both zoom and orientation levels. These variations more closely approximate the random viewing conditions that autonomous policies should be robust against when translated to real surgical scenarios.




Fig. 4: Roll-outs of our policy trained on the random-viewpoint dataset generalize on unseen camera viewpoints, showing robust performance across zoom and orientation changes. Examples show initial (left) and final (right) frames of the roll-outs. Human-controlled surgeon instrument on the bottom, robot-controlled assistant instrument on the top.




Fig. 5: Ablation on the amount of training data demonstrations plotted against the policy success rates for In-Distribution Roll-Outs

### E. Qualitative Analysis and Ablation Studies

We first analyze the qualitative behavior of the policies to understand the performance gap observed in the main benchmarks. We find that the ACT baseline frequently exhibits superficial grasping behavior, characterized by insufficient tissue purchase. As illustrated in Fig. 8, the baseline’s end-effector trajectories often fail to achieve the necessary approach depth compared to our MoE-augmented policy. This results in weak engagement with the phantom bowel tissue, leading to frequent slippage during the sustained retraction phase.

To evaluate the data efficiency of our approach, we performed an ablation study by training both the baseline and




Fig. 6: Confusion matrix of the MoE gating network on the validation dataset, demonstrating the effectiveness of the auxiliary phase classification task.




Fig. 7: Expert utilization rates on the validation dataset. The activation frequency aligns with the distribution of task phases in the training data.

our MoE policy on progressively smaller subsets of the phantom dataset: 100% (120 episodes), 50% (60 episodes), and 25% (30 episodes), with a consistent 10% validation split. This analysis aims to investigate learning robustness under conditions of extreme data scarcity and identify the minimum dataset size required for learning multi-step surgical tasks. The results, shown in Fig. 5, show that in the lowest data regime (25%), both policies perform identically (45% success rate), suggesting that 30 demonstrations represent a lower bound where data scarcity bottlenecks performance regardless of architecture. However, as data availability increases, a significant divergence emerges: while the standard ACT baseline plateaus at 50% success, our MoE policy effectively leverages the additional data, scaling from 60% to 85% success.

We further extend our analysis by examining the learned feature representations of the vision encoder of our MoE policy during the inference roll-outs. We construct saliency







Fig. 8: Top: trajectories of ACT and our ACT + MoE decoder from the same starting position to the same target grasping point, showing approach and retraction of the instrument. Bottom: frames of ACT (top) and ACT + MoE (bottom) highlight the difference in grasping depth. The MoE policy visibly demonstrates a deeper and more secure grasp, a consistent pattern observed during policy roll-outs. Robot-controlled instrument is grasping, human-controlled instrument is pointing.

maps using an adaptation of AblationCAM [9] for regression tasks, where we ablate regions of the visual feature maps and measure the resulting change in the predicted action norm, identifying which visual regions most strongly influence the policy’s action magnitude. As shown in Fig. 9, the resulting saliency maps reveal distinct patterns across task phases: early in episodes, high saliency concentrates on the robotic instrument’s position; when the surgeon provides visual cues, saliency shifts to the indicated grasping target; during retraction, the policy exhibits high sensitivity to the bowel segment between the robot’s and surgeon’s grasping points. This phase-dependent shift in salient regions suggests that the policy has learned to extract task-relevant visual features for each phase exploiting successfully its MoE block.

We analyze the performance of the phase classifier serving as the gating network for the action experts. The validation

set confusion matrix, presented in Fig. 6, demonstrates high classification accuracy, confirming that surgical phase classification is effectively learned as an auxiliary task. This explicit supervision enables the policy to correctly route states to their specialized experts. Furthermore, we examine the expert utilization rates in Fig. 7. We observe that the expert activation distribution closely mirrors the phase frequency of the training dataset. This validates that our method ensures balanced expert specialization, effectively preventing the policy from suffering from mode collapse or expert underutilization—common failure modes in unsupervised mixture-of-experts training.

Finally, we show preliminary qualitative results of the supervised MoE policy during *in vivo* porcine surgery<sup>3</sup>. In Fig. 10, we show two examples of successful policy rollouts in the aforementioned setting. We plan to extensively evaluate and present quantitative results in an *in vivo* porcine environment in our future work.

## V. CONCLUSION

In this work, we presented a Supervised Mixture-of-Experts architecture that enables lightweight action transformer policies to perform multi-step surgical manipulation tasks from limited data. Additionally, we demonstrated that a purely vision-based policy can generalize to unseen camera viewpoints if trained with sufficient geometric variability. Our experiments show that by incorporating random viewpoints during training, the policy maintains high success rates on unseen camera angles at test time, effectively preventing overfitting to a static 3D scene configuration. This capability, combined with successful zero-shot transfer to *ex vivo* tissue, underscores the potential of our method for real-world surgical assistance where the endoscopic view is not known a priori.

Limitations of our current framework include the reliance on manual phase supervision to guide the MoE gating network. Future work will investigate unsupervised learning methods to discover latent task skills implicitly from demonstration data. Moreover, we plan to equip our policies with real-time depth vision derived from stereoscopic endoscope feed to further enhance 3D spatial understanding and robustness for *in vivo* deployment. This is motivated by the observed behaviour of the policy in the preliminary *in vivo* experiments, where the lack of depth is an evident factor for the failure of the task in the grasping phase. We include examples of this in the supplementary material.

In conclusion, this work has shown that Imitation Learning coupled with a Mixture-of-Experts architecture can automate multi-step minimally-invasive surgical tasks using only endoscopic vision, building toward the goal of deploying learned policies *in vivo*.

## ACKNOWLEDGMENTS

The authors would like to acknowledge the support of the Federal Ministry of Research, Technology, and Space

<sup>3</sup>All procedures were approved by the local state authority (TVV43/2023, Saxony, Germany), and conducted in accordance with institutional ethical standards for animal experimentation and the registered protocol.




Fig. 9: AblationCAM heatmaps: the policy vision encoder focuses first on the robot instrument (top left), then on the surgeon instrument (top right), and at the end on stretching the bowel (bottom)




Fig. 10: Qualitative examples of two roll-outs of MoE-ACT policy during *in vivo* porcine surgery. Rows show task phases: reach (top), grasp (middle) and retract (bottom). Robot-controlled instrument on the top, human-controlled instrument on the right.

of Technology. In addition, this work has been partially supported by the Federal Ministry of Research, Technology, and Space (BMFTR) within the Robotics Institute Germany (RIG), grant agreement no. 16ME1005.

(BMFTR) within the research program “Communication Systems: Souverän. Digital. Vernetzt.”, Joint Project 6G-life (Project ID: 16KIS2413K). This work is also supported by the project “Next Generation AI Computing (gAI)” funded by the Bavarian Ministry of Science and the Arts and the Saxon Ministry for Science, Culture, and Tourism. The authors further acknowledge support from the German Research Foundation (DFG, Deutsche Forschungsgemeinschaft) under the Germany’s Excellence Strategy - EXC 2050/2 - Project ID 390696704 - Cluster of Excellence “Centre for Tactile Internet with Human-in-the-Loop” (CeTI) at TUD Dresden University

## REFERENCES

- [1] Anonymous. LeRobot: An Open-Source Library for End-to-End Robot Learning. In *Submitted to The Fourteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=CiZMMAFQR3>. under review.
- [2] Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny Driess, Adnan Esmail, Michael Robert Equi, Chelsea Finn, Niccolo Fusai, Manuel Y. Galliker, Dibya Ghosh, Lachy Groom, Karol Hausman, brian ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Devin LeBlanc, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Allen Z. Ren, Lucy Xiaoyang Shi, Laura Smith, Jost Tobias Springenberg, Kyle Stachowicz, James Tanner, Quan Vuong, Homer Walke, Anna Walling, Haohuan Wang, Lili Yu, and Ury Zhilinsky.  $\pi_{0.5}$ : a Vision-Language-Action Model with Open-World Generalization. In Joseph Lim, Shuran Song, and Hae-Won Park, editors, *Proceedings of The 9th Conference on Robot Learning*, volume 305 of *Proceedings of Machine Learning Research*, pages 17–40. PMLR, 27–30 Sep 2025. URL <https://proceedings.mlr.press/v305/black25a.html>.
- [3] Kevin Black et al.  $\pi_0$ : A Vision-Language-Action Flow Model for General Robot Control. *arXiv preprint arXiv:2410.24164*, 2024.
- [4] David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 2017.
- [5] Anthony Brohan, Noah Brown, Justice Carbajal, et al. RT-1: Robotics Transformer for Real-World Control at Scale. In *Robotics: Science and Systems (RSS)*, 2023.
- [6] Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research*, page 02783649241273668, 2023.
- [7] Anita Chiu, Wilbur B Bowne, Kelley A Sookraj, Michael E Zenilman, Abe Fingerhut, and George S Ferzli. The role of the assistant in laparoscopic surgery: important considerations for the apprentice-in-training. *Surgical innovation*, 15(3):229–236, 2008.
- [8] Perioperative Care Collaborative. The perioperative care collaborative position statement: surgical first assistant. PCC, 2018.
- [9] Saurabh Desai and Harish G Ramaswamy. Ablation-CAM: Visual explanations for deep convolutional network via gradient-free localization. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 983–991, 2020.
- [10] Isabel Funke, Sebastian Bodenstedt, Florian Oehme, Felix von Bechtolsheim, Jürgen Weitz, and Stefanie Speidel. Using 3D convolutional neural networks to learn spatiotemporal features for automatic surgical gesture recognition in video. In *International conference on medical image computing and computer-assisted intervention*, pages 467–475. Springer, 2019.
- [11] Jesse Haworth, Juo-Tung Chen, Nigel Nelson, Ji Woong Kim, Masoud Moghani, Chelsea Finn, and Axel Krieger. SutureBot: A Precision Framework & Benchmark For Autonomous End-to-End Suturing. *arXiv preprint arXiv:2510.20965*, 2025.
- [12] Irina Higgins et al. beta-VAE: Learning basic visual concepts with a constrained variational framework. In *ICLR*, 2017.
- [13] Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87, 1991.
- [14] Michael I Jordan and Robert A Jacobs. Hierarchical mixtures of experts and the EM algorithm. *Neural computation*, 6(2):181–214, 1994.
- [15] H. G. Kenngott, J. J. Wünscher, M. Wagner, A. Preukschas, A. L. Wekerle, P. Neher, S. Suwelack, S. Speidel, F. Nickel, D. Oladokun, L. Maier-Hein, R. Dillmann, H. P. Meinzer, and B. P. Müller-Stich. OpenHELP (Heidelberg laparoscopy phantom): Development of an open-source surgical evaluation and training tool. *Surgical Endoscopy*, 29(11):3338–3347, 2015. doi: 10.1007/s00464-015-4094-0.
- [16] Alexander Khazatsky, Karl Pertsch, Suraj Nair, Ashwin Balakrishna, Sudeep Dasari, Siddharth Karamcheti, Soroush Nasiriany, Mohan Kumar Srirama, Lawrence Yunliang Chen, Kirsty Ellis, et al. Droid: A large-scale in-the-wild robot manipulation dataset. *arXiv preprint arXiv:2403.12945*, 2024.
- [17] Ji Woong Kim, Tony Z Zhao, Samuel Schmidgall, Anton Deguet, Marin Kobilarov, Chelsea Finn, and Axel Krieger. Surgical robot transformer (srt): Imitation learning for surgical tasks. *arXiv preprint arXiv:2407.12998*, 2024.
- [18] Ji Woong Kim, Juo-Tung Chen, Pascal Hansen, Lucy Xiaoyang Shi, Antony Goldenberg, Samuel Schmidgall, Paul Maria Scheikl, Anton Deguet, Brandon M White, De Ru Tsai, et al. SRT-H: A hierarchical framework for autonomous surgery via language-conditioned imitation learning. *Science robotics*, 10(104):eadt5254, 2025.
- [19] Moo Jin Kim, Chelsea Finn, and Percy Liang. Fine-tuning vision-language-action models: Optimizing speed and success. *arXiv preprint arXiv:2502.19645*, 2025.
- [20] Moo Jin Kim et al. OpenVLA: An Open-Source Vision-Language-Action Model. *arXiv preprint arXiv:2406.09246*, 2024.
- [21] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [22] Ioannis T Konstantinidis, Philip Ituarte, Yanghee Woo, Susanne G Warner, Kurt Melstrom, Jae Kim, Gagandeep Singh, Byrne Lee, Yuman Fong, and Laleh G Melstrom. Trends and outcomes of robotic surgery for gastrointestinal (GI) cancers in the USA: maintaining perioperative and oncologic safety. *Surgical Endoscopy*, 34(11):4932–4942, 2020.

- [23] Yonghao Long, Anran Lin, Derek Hang Chun Kwok, Lin Zhang, Zhenya Yang, Kejian Shi, Lei Song, Jiawei Fu, Hongbin Lin, Wang Wei, et al. Surgical embodied intelligence for generalized task autonomy in laparoscopic robot-assisted surgery. *Science Robotics*, 10(104): eadt3093, 2025.
- [24] Lena Maier-Hein, Swaroop S Vedula, Stefanie Speidel, Nassir Navab, Ron Kikinis, Adrian Park, Matthias Eisenmann, Hubertus Feussner, Germain Forestier, Stamatia Giannarou, et al. Surgical data science for next-generation interventions. *Nature Biomedical Engineering*, 1(9):691–696, 2017.
- [25] Octo Model Team et al. Octo: An Open-Source Generalist Robot Policy. *arXiv preprint arXiv:2405.12213*, 2024.
- [26] Open X-Embodiment Collaboration. Open X-Embodiment: Robotic Learning Datasets and RT-X Models. *arXiv preprint arXiv:2310.08864*, 2024.
- [27] Sathira Kasun Perera, Susannah Jacob, Brooke E Wilson, Jacques Ferlay, Freddie Bray, Richard Sullivan, and Michael Barton. Global demand for cancer surgery and an estimate of the optimal surgical and anaesthesia workforce between 2018 and 2040: a population-based modelling study. *The Lancet Oncology*, 22(2):182–189, 2021.
- [28] Karl Pertsch, Kyle Stachowicz, Brian Ichter, Danny Driess, Suraj Nair, Quan Vuong, Oier Mees, Chelsea Finn, and Sergey Levine. Fast: Efficient action tokenization for vision-language-action models. *arXiv preprint arXiv:2501.09747*, 2025.
- [29] Moritz Reuss, Jyothish Pari, Pulkit Agrawal, and Rudolf Lioutikov. Efficient diffusion transformer policies with mixture of expert denoisers for multitask learning. *arXiv preprint arXiv:2412.12953*, 2024.
- [30] Alexander Schüßler, Christian Kunz, Rayan Younis, Benjamin Alt, Jamie Paik, Martin Wagner, and Franziska Mathis-Ullrich. Semi-Autonomous Robotic Assistance for Gallbladder Retraction in Surgery. *IEEE Robotics and Automation Letters*, 2025.
- [31] Mustafa Shukor, Dana Aubakirova, Francesco Capuano, Pepijn Kooijmans, Steven Palma, Adil Zouitine, Michel Aractingi, Caroline Pascal, Martino Russi, Andres Marafioti, et al. Smolvla: A vision-language-action model for affordable and efficient robotics. *arXiv preprint arXiv:2506.01844*, 2025.
- [32] Wenxuan Song, Han Zhao, Pengxiang Ding, Can Cui, Shangke Lyu, Yaning Fan, and Donglin Wang. Germ: A generalist robotic model with mixture-of-experts for quadruped robot. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 11879–11886. IEEE, 2024.
- [33] R Younis, A Yamlahi, S Bodenstedt, PM Scheikl, A Kisilenko, M Daum, A Schulze, PA Wise, F Nickel, F Mathis-Ullrich, et al. A surgical activity model of laparoscopic cholecystectomy for co-operation with collaborative robots. *Surgical Endoscopy*, 38(8):4316–4328, 2024.
- [34] Lantao Yu, Tianhe Yu, Jiaming Song, Willie Neiswanger, and Stefano Ermon. Offline imitation learning with sub-optimal demonstrations via relaxed distribution matching. In *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence and Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence and Thirteenth Symposium on Educational Advances in Artificial Intelligence*, AAAI’23/IAAI’23/EAAI’23. AAAI Press, 2023. ISBN 978-1-57735-880-0. doi: 10.1609/aaai.v37i9.26305. URL <https://doi.org/10.1609/aaai.v37i9.26305>.
- [35] Tony Z. Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware. In *Proceedings of Robotics: Science and Systems*, Daegu, Republic of Korea, July 2023. doi: 10.15607/RSS.2023.XIX.016.
- [36] Hongyi Zhou, Denis Blessing, Ge Li, Onur Celik, Xiaogang Jia, Gerhard Neumann, and Rudolf Lioutikov. Variational distillation of diffusion policies into mixture of experts. *Advances in Neural Information Processing Systems*, 37:12739–12766, 2024.