

ViTac-Tracing: Visual-Tactile Imitation Learning of Deformable Object Tracing

Yongqiang Zhao^{1*}, Haining Luo^{2*}, Yupeng Wang¹, Emmanouil Spyarakos Papastavridis¹,
Yiannis Demiris², Shan Luo¹

Abstract—Deformable objects often appear in unstructured configurations. Tracing deformable objects helps bringing them into extended states and facilitating the downstream manipulation tasks. Due to the requirements for object-specific modeling or sim-to-real transfer, existing tracing methods either lack generalizability across different categories of deformable objects or struggle to complete tasks reliably in the real world. To address this, we propose a novel visual-tactile imitation learning method to achieve one-dimensional (1D) and two-dimensional (2D) deformable object tracing with a unified model. Our method is designed from both local and global perspectives based on visual and tactile sensing. Locally, we introduce a weighted loss that emphasizes actions maintaining contact near the center of the tactile image, improving fine-grained adjustment. Globally, we propose a tracing task loss that helps the policy to regulate task progression. On the hardware side, to compensate for the limited features extracted from visual information, we integrate tactile sensing into a low-cost teleoperation system considering both the teleoperator and the robot. Extensive ablation and comparative experiments on diverse 1D and 2D deformable objects demonstrate the effectiveness of our approach, achieving an average success rate of 80% on seen objects and 65% on unseen objects. Demos, code and datasets are available at <https://sites.google.com/view/vitac-tracing>.

I. INTRODUCTION

Deformable object manipulation has gained increasing attention in recent years due to its numerous real-world applications, such as cable management, cloth handling, and assistive dressing [1]. A common characteristic of deformable objects is that they commonly appear in unstructured configurations, e.g. a folded-over shoelace or a crumpled towel, where their geometry and task-relevant features are difficult to observe directly (Fig. 1). Tracing transforms deformable objects into an extended configuration by following their edge with fingers from one end to the other [2], [3].

We consider deformable object tracing in two categories: one-dimensional (1D) linear objects (e.g., cables, ropes) and two-dimensional (2D) planar objects (e.g., towels, clothes). Although structurally distinct, these objects exhibit comparable physical properties [4], suggesting the possibility of tracing both categories with a single unified model. Previous work has so far addressed only one category, constrained by the limitations of their methods. Model-based controllers

¹Robot Perception Laboratory, Department of Engineering, King's College London, Strand, London, WC2R 2LS, United Kingdom, {yongqiang.zhao, yupeng.l.wang, emmanouil.spyarakos, shan.luo}@kcl.ac.uk.

²Personal Robotics Laboratory, Department of Electrical and Electronic Engineering, Imperial College London, London SW7 2AZ, United Kingdom {haining.luo18, y.demiris}@imperial.ac.uk.

* Yongqiang Zhao and Haining Luo contributed equally to this work.




Fig. 1. (a) 1D deformable object tracing. (b) 2D deformable object tracing. Using our proposed method, an ABB YuMi traces the objects by sliding a gripper along the 1D deformable object or an edge of the 2D deformable object and transforming them from the unstructured configurations on the left to the extended states on the right.

were among the first explored for the tracing task [2], [5]. However, the requirement for accurate modeling of the object states and dynamics makes it difficult to construct a single controller for multiple objects. This is further compounded by the effectively infinite Degrees of Freedom (DoF) of deformable objects. Reinforcement Learning (RL) provides an alternative for learning unified tracing policies across multiple object types. However, it typically requires carefully designed reward functions and accurate deformable object modeling in simulation, and often suffers sim-to-real gaps [6], [7]. In contrast, Imitation Learning (IL) methods leverage expert demonstrations to accelerate policy learning and avoids the need for explicit modeling or sim-to-real transfer. In this work, we propose an IL approach to trace both 1D and 2D deformable objects within a unified model.

During tracing, maintaining the contact region near the center of the fingertips helps minimize the risk of object slippage [2]. However, due to occlusion by the fingers, visual information alone is typically insufficient to capture grasping conditions. A common solution is to incorporate tactile sensing to provide the missing local information [2]. To date, most budget-efficient teleoperation systems do not include tactile sensing capabilities [8], [9], [10].

In this work, we utilize a dual-arm ABB YuMi and a bespoke 3D-printed kinematic model for teleoperation. Tactile sensing is incorporated by integrating a vision-based

tactile sensor into the gripper finger. Previous studies have shown that sufficient visual and haptic feedback enhances the quality of demonstration data [11]. To this end, we provide the teleoperator with multi-modal feedback through real-time visual and tactile image streaming, complemented by vibration alerts for singularities.

With the visual and tactile information, we propose an imitation learning method for deformable object tracing. This work marks the first step towards a unified model for tracing both 1D and 2D deformable objects. Our method is designed from both the local and global perspectives of the task. To improve fine-grained action adjustment and reduce the risk of dropping the object, we propose a local center loss that prioritizes actions that center the object in the tactile image. Since various deformable manipulation tasks require tracing stopping at vision-detected accurate locations, e.g., stopping when inserting a cable into a clip during cable routing [12], we further design a global task loss to regulate the task progression. Experiments validate the effectiveness of individual components and demonstrate the generalizability of our method, achieving an average success rate of 80% on seen objects and 65% on unseen objects.

In summary, the contributions of this work are as follows:

- We propose a novel visual-tactile imitation learning framework for deformable object tracing, enabling real robots to trace on diverse 1D and 2D deformable objects through a unified policy;
- We introduce a budget-efficient visual-tactile teleoperation system with multi-modal feedback to enrich perception of both the robot and the teleoperator;
- Extensive experiments validate the effectiveness of the individual components of our method and demonstrate the performance on seen objects as well as the generalizability to unseen objects.

II. RELATED WORKS

A. Deformable Object Tracing

Deformable object manipulation has drawn increasing attention in robotics [4]. Progress has been made on tasks such as cable routing [13], knot tying [14], rope untying [15], [16], and cloth folding [17], [18]. A common facilitating step in these tasks is transforming the objects into extended or unfolded configurations. For example, spreading a T-shirt into a flat state can simplify the identification of keypoints or landmarks [19]. Object tracing provides a means of transforming deformable objects into such states, and has been employed in some works on cloth unfolding [5] and cable insertion [2].

In this work, we address the tracing task for both 1D and 2D deformable objects, with the aim of recovering fully extended configurations from unstructured, crumpled states. Owing to the effectively infinite number of DoF in deformable objects, tracing demands fine-grained adjustments to maintain stable gripper-object contact throughout the process. In recent years, several studies have attempted to tackle the tracing task for 1D linear objects (some referred to

as “object following”) and for 2D planar objects (also known as “object sliding”).

Traditional control-based methods were among the first to be applied to the tracing task. For 1D objects, She et al. [2], [3] designed grip and pose controllers for cable tracing using real-time tactile feedback. Later, Liu et al. [20] introduced a dynamic contrastive Koopman teleoperator for transparent tube tracing. For 2D objects, Zhang et al. [5] employed a deep Model Predictive Control (MPC) method to iteratively servo garments from raw visual and tactile inputs. Although these approaches demonstrate promising results, they are tailored to specific object types and struggle to generalize across different geometries due to their reliance on accurate modeling of object dynamics and states.

Reinforcement Learning (RL) methods have also been applied to the tracing task [6], [21], [22], [23]. Pecyna et al. [6] achieved 1D object tracing in simulation using a model-free RL approach with visual-tactile fusion. Later, Sun et al. [7] extended the method to real 1D cable following using the Soft Actor-Critic (SAC) algorithm. Tailored for high-precision manipulation, Yu et al. [24] applied a similar strategy to 1D needle threading. For 2D objects, Zhou et al. [25] trained an RL policy in a latent action space for towel tracing. Although model-free RL methods do not directly require object dynamics, these methods depend on either accurate deformable object models in simulation or carefully designed reward functions, which leads to sim-to-real gaps and hinders the development of the RL environment for tracing both 1D and 2D objects.

As a highly effective robot manipulation method [26], Imitation Learning (IL) offers another option for the tracing task. By leveraging expert demonstrations, it does not require accurate deformable object modeling and is not affected by sim-to-real gaps. This makes it an attractive choice for training a unified model capable of tracing multiple types of objects. In this study, we introduce a local center loss and a global task loss to an imitation learning framework, enabling us to address both 1D and 2D deformable object tracing.

B. Vision-Tactile Imitation Learning

In recent years, vision has become the primary modality used in IL [8], [27]. While visual sensing is well-suited for capturing global features, it is susceptible to occlusion by the manipulated object or surrounding items. Tactile sensing complements vision by providing rich local feature information.

Most existing teleoperation systems lack fine-grained tactile feedback [11]. Several recent studies have combined vision and tactile sensing for policy learning with IL. Xue et al. [11] developed a teleoperation system via Augmented Reality (AR) for demonstration collection and proposed a visual-tactile imitation learning method for contact-rich manipulation skills. Meanwhile, Gu et al. [28] integrated a vision-based tactile sensor into the ALOHA system [8] to accomplish precise tasks such as zip-tie insertion and Velcro fastening. In this work, we take a similar approach

and integrate a vision-based tactile sensor into a bimanual teleoperation system.

In IL, collecting expert demonstrations is a crucial step. Existing data collection systems generally rely only on visual feedback [8], [9], [10]. While equipping robots with tactile sensing is straightforward, providing haptic feedback to human operators is equally important. Recent studies [28], [11], employ real-time tactile image monitoring for this purpose, while commercial devices such as Touch Haptic Device [29] can also be used to provide vibration-based feedback. Sufficient visual and haptic feedback to the teleoperator can enhance the consistency and quality of demonstration data [11]. To enhance our data quality, we provide multi-modal feedback to the teleoperator with real-time visual–tactile monitoring and vibration feedback.

III. PROBLEM STATEMENT

We define deformable object tracing as the task of sliding a gripper along a 1D deformable object or an edge of a 2D deformable object, while maintaining continuous contact, until reaching the terminal point. We assume that one end of the object is secured by a second gripper, and the sliding gripper starts near the grasping point and moves towards the free end of the object. The task is illustrated in Fig. 1.

The deformable object tracing problem can be formulated as follows. We model a 1D deformable object, or an edge of a 2D deformable object, as a time-varying spatial curve $\mathcal{C}_t \subset \mathbb{R}^3$ in Cartesian space. The length of \mathcal{C}_t is defined as L . The time steps of this task move from $t = 0$ to $t = T$. Let $p_0 = (x_0, y_0, z_0)$ denote the fixed point in the world frame, $p_t = (x_t, y_t, z_t)$ is the contact point with the moving gripper. o^T represents the gripper’s tactile sensing region.

In the deformable object tracing task, we identify two constraints for successful completion: (1) the object should remain grasped, i.e., $p_t \in \mathcal{C}_t$, meaning the contact point p_t should always lie on the curve; (2) the contact point p_t should stay within the effective tactile sensing region o^T , i.e., $p_t \in o^T$. Since the tactile sensing region covers the fingertip, leaving this region implies the objects has slipped or dropped. Meanwhile, we define two objectives for the task: (1) the distance between the contact point p_t and the fixed point p_0 should gradually converge to the total curve length as t increases, i.e., $\|p_t - p_0\|_2 \rightarrow L$; (2) this distance function between should be monotonically increasing over time, i.e., $\frac{d}{dt}\|p_t - p_0\|_2 \geq 0$.

IV. METHODOLOGY

A. Overview

The proposed visual-tactile imitation learning framework is illustrated in Fig. 3. Demonstration data is collected using a visual-tactile teleoperation system equipped with real-time visual/tactile monitoring and vibration feedback, as shown in Fig. 2. We then use the collected dataset to train the proposed action-chunking-based policy with the local center loss, the Kullback–Leibler (KL) divergence loss, and the global task loss. The trained policy is subsequently deployed on a real robot.




Fig. 2. Visual-tactile teleoperation system for collecting demonstrations. On the robot side, a top-view camera is installed, and a tactile sensor is mounted on the follower robot’s gripper. On the teleoperator side, visual and tactile images are monitored in real time, and vibration motors are mounted on the leader robot’s gripper.

B. Visual-Tactile Teleoperation System

1) **Robot Side:** As discussed in Section II-B, tactile sensors can provide rich information about grasping conditions. A vision-based tactile sensor, derived from GelSight Wedge [30], is mounted on the gripper. For the visual input, a ZED 2 stereo camera provides a top-down view of the Cartesian space (Fig. 2). The robot is controlled in joint position mode, with the End-Effector (EE) velocity limited to 400 mm/s. The robot platform is further equipped with an Nvidia Jetson Orin running the camera driver and the robot’s ROS driver. All drivers are containerized in Docker on ROS Noetic.

2) **Teleoperator Side:** Section II-B also highlighted the importance of haptic feedback. The absence of such feedback can limit the teleoperator’s awareness of grasping conditions. To address this, we enhance the teleoperator’s perception by adding a display screen to stream visual and tactile image from the robot side in real time. Meanwhile, as our tracing policy is trained with EE poses rather than joint angles as in [30] (see Section V-B for the rationale behind this choice), singular configurations become a concern due to the non-unique mapping between EE poses to joint configurations. To address this, we provide haptic feedback through vibration motors (DAOKAI DC 5V Mini) mounted on the leader grippers, which are triggered when the EE approaches singular configurations. To calculate the proximity to singularity, we adopt Yoshikawa Manipulability Index [31]

$$w(q) = \sqrt{\det(J(q)J(q)^T)}, \quad (1)$$

where q is the joint states of the robot, and $J(\cdot)$ is the Jacobian matrix. Higher $w(q)$ indicates better dexterity, meaning the robot can move in multiple directions from the current configuration. The motors are triggered when the index falls below $\lambda_w \max(w)$, where $\lambda_w = 0.2$ is selected via grid search in our experiments. All three feedback devices are shown in Fig. 2.

Using this visual-tactile teleoperation system, we collect kinematic information o^K , visual information o^V , and tactile

information o^T from the robot side, while ground-truth action information a is collected from the teleoperator side. We define the observation at time step t as $o_t = \{o_t^K, o_t^V, o_t^T\}$. A collected demonstration episode \mathcal{D} is thus represented as $\mathcal{D} = \{(o_t, a_t)\}_{t=0}^T$.

C. Tracing Policy Learning

We train the tracing policy via transformer-based policy network, as illustrated in Fig. 3. Specifically, we adopt the Action Chunking Transformer (ACT) algorithm [8], where the policy $\pi_\theta(\hat{a}_{t:t+k}|o_t)$ takes an observation o_t as input and predicts the next k sequential actions \hat{a} . ACT optimizes a reconstruction loss $\mathcal{L}_{reconst}$ to reconstruct actions from the inputs and a regularization loss \mathcal{L}_{reg} to regularize the Transformer encoder toward a Gaussian prior.

$$\mathcal{L}_{reconst} = \text{MAE}(\hat{a}_{t:t+k}, a_{t:t+k}), \quad (2)$$

$$\mathcal{L}_{reg} = D_{KL}(q_\phi(z|a_{t:t+k}, \bar{o}_t) || \mathcal{N}(0, I)), \quad (3)$$

where MAE is the mean absolute error, i.e., L1 loss. $q_\phi(z|a_{t:t+k}, \bar{o}_t)$ denotes the Transformer encoder, z is a style variable, and \bar{o}_t represents o_t without image observations. We refer the reader to [8] for the technical details of ACT.

In addition to the aforementioned losses, we introduce a local center loss to enhance tactile-feedback-driven adjustments, and a global task loss to improve task progression.

1) Local Center Loss: As noted in Section III, a key constraint in the tracing task is maintaining contact within the gripper’s tactile sensing region, i.e., $p_t \in o^T$. Real-time tactile monitoring allows teleoperators to perceive gripper-object contact during demonstration collection. However, owing to the high-DoF dynamics of deformable objects, contacts near the edges of the tactile sensing region are prone to drop from the robot gripper or slip to the imperceptible area, as illustrated in Fig. 4. The ideal contact location lies at the center of the sensing region. To encourage this, we introduce a center loss that prioritizes the actions leading to contact points moving closer to the sensor center.

We first localize the contact point p_t via classical image processing, since the high-resolution tactile texture is highly salient, enabling robust contact-point extraction, as shown in Fig. 6. A contact mask is generated through grayscale conversion, thresholding, and Gaussian filtering, followed by contour extraction. The largest contour is then fitted with an ellipse (preferred) or characterized via Principal Component Analysis (PCA), yielding the pixel-coordinate contact point $p_t^{\text{tac}} = (u_t^{\text{tac}}, v_t^{\text{tac}})$. Let the sensing-region center be $c = (u_c, v_c)$, we introduce a weight factor

$$w_t = \exp\left(-\frac{\|p_t^{\text{tac}} - c\|}{c}\right). \quad (4)$$

We combine the original reconstruction loss with the proposed weight to obtain the local center loss

$$\mathcal{L}_{center} = w_{t:t+k} \cdot \text{MAE}(\hat{a}_{t:t+k}, a_{t:t+k}). \quad (5)$$

2) Global Task Loss: The definition of tracing discussed in Section III illustrates the constraint of maintaining object–grripper contact and the goal of reaching the full curve length. While the center loss addresses the contact constraint, it does not assist the policy in achieving the overall task goal. In many manipulation tasks, precise termination is critical, e.g., tracing to the corner of a towel to enable folding, or stopping before a cable slides off from the gripper in insertion tasks. We therefore introduce a global task loss to regulate the task progression.

Assuming the object edge is locally straight during the task, where one end is fixed and the segment between the fixed point p_0 and the contact point p_t is typically tensioned straight, the traced length can be inferred from p_0 and p_t , and the completion index I is computed from the traced length and the overall length. We firstly localize p_t^{tac} from the tactile image and convert it to the gripper frame (coincident with the sensor frame)

$$p_t^{\text{grripper}} = \left(\frac{u_t^{\text{tac}} - u_c}{p2m}, \frac{v_t^{\text{tac}} - v_c}{p2m}, 0\right), \quad (6)$$

where $p2m$ is the pixel-to-meter scale.

Then, we can calculate p_t through p_t^{tac} using coordinate transformation

$$p_t = T_{\text{grripper}}^{\text{world}} p_t^{\text{grripper}}, \quad (7)$$

where $T_{\text{grripper}}^{\text{world}}$ is the transform matrix from the gripper frame to the world frame.

Following that, the ground-truth completion index I of task is able to be labeled on the collected demonstrations as

$$I = \min(\max(\frac{\|p_t - p_0\|_2}{\|p_T - p_0\|_2}, 0), 1). \quad (8)$$

As shown in Fig. 3, we extend the policy network with a completion index prediction branch, which predicts completion squence $\hat{I}_{t:t+k}$ alongside the action sequence. This branch is optimized using

$$\mathcal{L}_{task} = \text{MSE}(\hat{I}_{t:t+k}, I_{t:t+k}), \quad (9)$$

where MSE denotes the mean squared error (L2 loss).

Therefore, the overall loss \mathcal{L} of the proposed network is

$$\mathcal{L} = \mathcal{L}_{center} + \lambda_{reg} \mathcal{L}_{reg} + \lambda_{task} \mathcal{L}_{task}, \quad (10)$$

where $\lambda_{reg} = 100$ and $\lambda_{task} = 100$ are selected via grid search.

D. Task Evaluation Metrics

Following [6], [5], [10], we identify the experiment results and use their rates as evaluation metrics

- **Success:** the gripper follows the object to the terminal end and maintains grasp there (any point within the final 5% of length).
- **Robot collision:** the robot collides with the object or itself and cannot recover from the current state.
- **Early stopping:** the gripper does not reach the terminal end but does not drop the object.

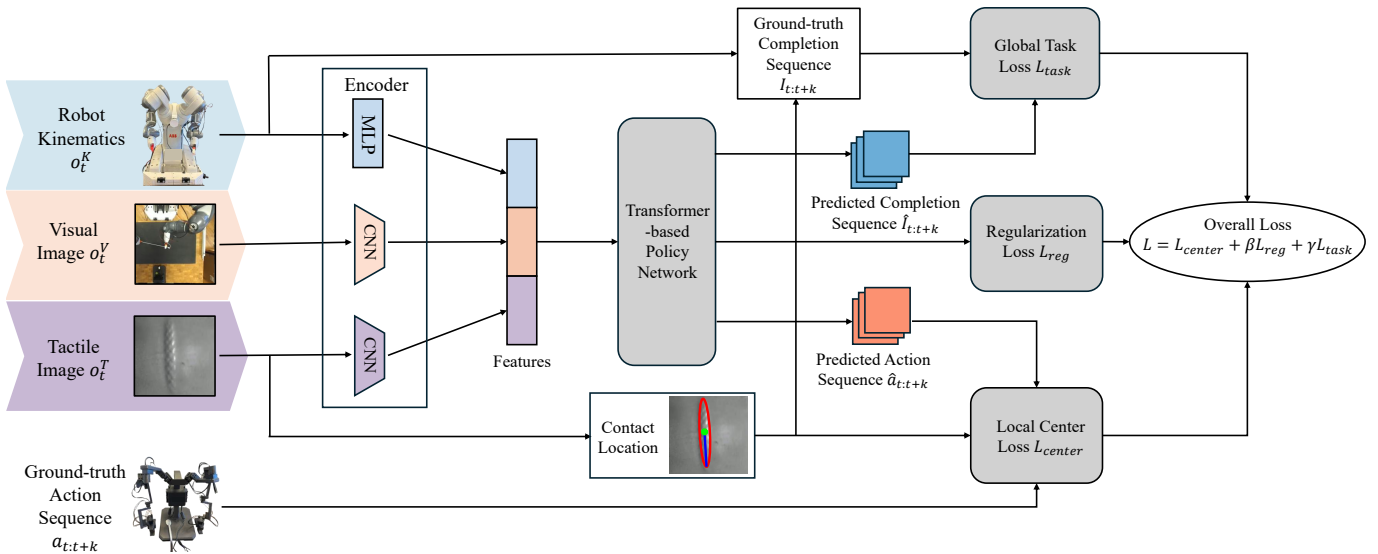


Fig. 3. Overview of the proposed tracing policy learning framework. The inputs include robot kinematics o_t^K , visual image o_t^V , and tactile image o_t^T collected from the follower robot, while the ground truth consists of action sequence $a_{t:t+k}$ recorded from the leader robot. Input features are first extracted using a Multilayer Perceptron (MLP) and Convolutional Neural Networks (CNNs). These features are then concatenated and fed into a Transformer-based policy network, which is trained using a combination of three loss functions: local center loss, global task loss, and regularization loss.




Fig. 4. Tactile images under different contact locations between the object and gripper. (a) Object grasped near the center of the tactile sensing region; (b) Object grasped near the front edge of the tactile sensing region; (c) Object grasped near the rear edge of the tactile sensing region. When the object is grasped near the edges of the sensing region, it is more likely to slip into an imperceptible area.

- **Over-tracing:** the gripper reaches the final 5% but fails to maintain grasp.
- **Object dropping:** the gripper drops the object before reaching the end.

We introduce two other metrics to further quantify the policy performance:

- **Success time:** the time length the robot spends to complete the tracing task successfully.
- **Completion ratio:** the ratio between the distance the gripper reached (measured this distance from the final contact point to the fixed point) and the deformable object length, i.e., $\|p_T - p_0\|_2/L$.

V. EXPERIMENT RESULTS

A. Experimental and Training Setups

For data collection, we selected two 1D and two 2D deformable objects: a flat shoelace, a braided cable, a face towel, and a microfiber cloth (Fig. 5). We further show the example tactile images of the tested objects in Fig. 6. For each object, 25 demonstrations were collected using our

visual–tactile teleoperation system, resulting in 100 demonstrations in total. During collection, we recorded visual images, tactile images, robot states, and kinematic model states as action commands at 30 Hz. The visual and tactile images were cropped and resized to 480×480 . The 16-dimensional action–state input consisted of either 14 joint states with 2 gripper states (for models trained in joint space) or 2 EE poses with 2 gripper states (for models trained in Cartesian space). We further test our policies on two unseen objects: a 1D synthetic rope and a 2D cotton napkin. In addition, we keep lighting constant, use a black sponge pad as the table surface, and randomize object placement by stacking the 1D and 2D deformable objects during data collection and policy deployment.

We set the chunk size to 60 steps and use separate ResNet18 backbones to extract visual and tactile features. During training, image brightness, contrast, and gamma are randomized. Each model is trained for 15,000 epochs, and the checkpoint with the lowest validation loss is selected. Training is performed on a workstation equipped with an AMD Ryzen Threadripper Pro 5965WX processor, 128 GB of RAM, and an RTX 4090 GPU. Temporal aggregation is disabled, following ACT’s tuning recommendations [10]. During inference, the models run on a desktop computer with an Intel i9 CPU (24 cores), 32 GB of RAM. For models trained in Cartesian space, the output EE poses are converted to joint angles before being executed.

B. Proprioceptive Representation Comparison

The two common representations for proprioception are the joint angles and EE poses. While joint angles provide raw low-level proprioception, end-effector poses represent a higher-level kinematic abstraction. To investigate the effect of proprioception representation on the tracing task, we

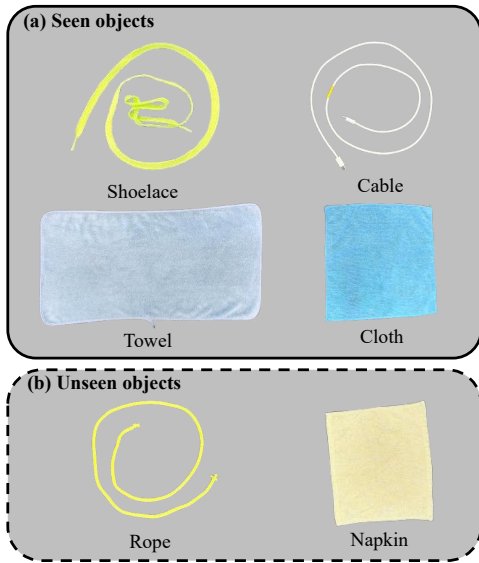


Fig. 5. 1D and 2D deformable objects used in the experiments. (a) Seen objects; (b) Unseen objects. For each seen object, 25 demonstrations were collected. The trained models were tested on both seen and unseen objects.

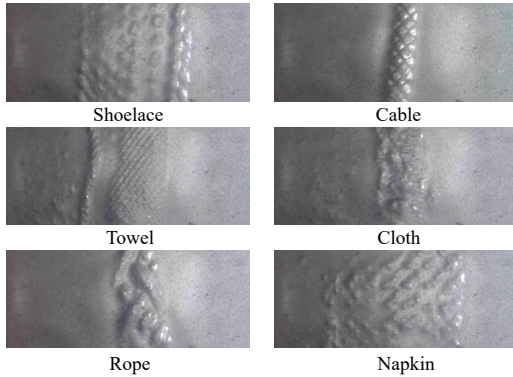


Fig. 6. Tactile images providing high-resolution texture information of the tested 1D deformable objects and the hemmed edges of 2D deformable objects. The folded hems of the 2D objects show textures similar to the 1D objects. The tactile textures of the unseen objects also exhibit resemblance to those of the seen objects.

compared models trained with two different forms of input with all four datasets.

We evaluated both models on the four objects used for data collection, with 10 trials per object, resulting in 40 experiments per model. The accumulated results are shown in Table I. Compared with the joint-space model (last row) achieved a higher success rate of 80%. This trend is also visible in Fig. 7, where the completion ratio of the joint-space model is noticeably lower. These findings suggest that the higher-level abstraction of EE poses is more suitable for the tracing task. Since tracing is inherently defined in task space, providing the policy with EE poses aligns the input with the task objective and avoids redundancy present in joint-angle representations. This reduces ambiguity and enables finer-grained adjustments, which is evident in the lower object-dropping rate. The joint-space model exhibits a much higher frequency of object drops, indicating a weaker ability to adjust object orientation relative to the fingers. Therefore, we adopt EE poses as the proprioceptive




Fig. 7. Success time and completion ratio for the component ablation study. Overall, the tested models showed similar success time. It is observable that all variants demonstrated lower completion ratios, indicating degraded task performance.

TABLE I
ABLATION EXPERIMENTAL RESULTS TO VALIDATE THE EFFECTIVENESS OF INDIVIDUAL COMPONENTS.

Methods	Success rate (Wilson 95% CI)	Robot collision	Early stopping	Over-tracing	Object dropping
Joint Space	70.0% [54.6, 81.9]	1/40	4/40	2/40	5/40
w/o Vision	65.0% [49.5, 77.9]	4/40	2/40	8/40	0/40
w/o Tactile	60.0% [44.6, 73.7]	2/40	5/40	1/40	8/40
w/o Center Loss	65.0% [49.5, 77.9]	4/40	1/40	0/40	9/40
w/o Task Loss	67.5% [52.0, 79.9]	3/40	3/40	7/40	0/40
Ours	80.0% [65.2, 89.5]	2/40	2/40	3/40	1/40

representation for the remainder of this work.

C. Component Ablation Study

To evaluate the contribution of different sensing modalities and objective terms, we performed an ablation study. Specifically, we compared the proposed system against variants trained without tactile input, without visual input, without the proposed center loss, and without the proposed task loss. Each model was tested in 10 trials per object, and the results were aggregated (Table I). Overall, the results show that both visual and tactile information, as well as the two proposed losses, contributed to the success of the experiments. The success time across all tested models did not exhibit noticeable differences.

1) *Ablation on Sensing Modalities*: We first assess the effect of removing tactile or visual input to determine the role of each sensing modality. Without visual information, the robot is more prone to over-tracing, reflecting its reduced ability to control task termination. In contrast, without tactile information, the robot achieves lower completion ratio (Fig. 7) and is more likely to drop the object, indicating

TABLE II

EXPERIMENTAL RESULTS USING DIFFERENT TRAINING DATASETS.

Dataset types	Tested objects	Success rate (Wilson 95% CI)	Robot collision	Early stopping	Over-tracing	Object dropping
1D	Shoelace	80.0% [49.0, 94.3]	1/10	0/10	1/10	0/10
	Cable	90.0% [59.6, 98.2]	1/10	0/10	0/10	0/10
	Towel	60.0% [31.3, 83.2]	2/10	1/10	1/10	0/10
	Cloth	80.0% [49.0, 94.3]	0/10	1/10	1/10	0/10
2D	Towel	70.0% [39.7, 89.2]	1/10	1/10	0/10	1/10
	Cloth	70.0% [39.7, 89.2]	0/10	1/10	1/10	1/10
Ours	Shoelace	90.0% [59.6, 98.2]	0/10	1/10	0/10	0/10
	Cable	80.0% [49.0, 94.3]	1/10	1/10	0/10	0/10
	Towel	70.0% [39.7, 89.2]	1/10	0/10	1/10	1/10
	Cloth	80.0% [49.0, 94.3]	0/10	0/10	0/10	2/10

weaker fine-grained adjustment of the tracing direction. These results suggest that the model relies more on visual information for task progression, while tactile information plays a key role in maintaining stable contact between the object and the gripper.

2) *Ablation on Objective Losses*: We also examine the impact of excluding the center loss and task loss to evaluate their contribution to policy performance. Without the center loss, the policy exhibits noticeably higher rates of object dropping. This suggests that the regulation introduced by the center loss helps the model learn adjustment actions more effectively. Although our demonstration data consist only of successful episodes, teleoperator actions are not always ideal, i.e., they do not consistently move the contact region towards the fingertip center. The center loss weighs actions and prioritizes those keep the object near the center of the tactile image, thereby improving the model’s learning of tracing direction adjustments. Without the task loss, failure modes show a higher rate of over-tracing and slightly more cases of early stopping. This indicates that including the completion index as an output enables the model to better track task progression and recognize the task endpoint.

The dominant failure modes are: (1) robot collision due to excessive contact force and friction/entanglement; (2) early stopping/over-tracking under visual occlusion near the object end; and (3) object dropping due to weak tactile feedback and an unstable grasp. The towel shows a lower success rate mainly because its longer length amplifies accumulated tracking errors and interaction-induced disturbances.

D. Effect of Multi-Object Training

The previous experiments have validated the design choices of our method. In this section, we investigate whether unifying deformable object tracing with a single model improves or harms task performance. We train models with individual object datasets (top four rows in Table II), with grouped 1D and 2D datasets (middle section of Table II), and with all object datasets combined (bottom section of Table II). Each model is then tested on its corresponding objects, e.g., the model trained on the shoelace dataset is

TABLE III

EXPERIMENTAL RESULTS ON UNSEEN OBJECTS.

Unseen objects	Success rate (Wilson 95% CI)	Robot collision	Early stopping	Over-tracing	Object dropping
Rope	70.0% [48.1, 85.5]	0/20	4/20	0/20	2/20
Napkin	60.0% [38.7, 78.1]	2/20	0/20	4/20	2/20

tested on the shoelace. Overall, the success rate is higher for 1D objects and the results indicate similar performance between unified and individual models.

For the three models tested on the shoelace, the success rate ranges between 80% and 90%, and similar results are observed for the cable. For the towel, the success rate falls between 60% and 70%, while for the cloth it is between 70% and 80%. We attribute this difference to the intrinsic properties of the objects. Unlike 1D objects, 2D objects have fabric parts that dangle outside the fingertips (Fig. 1). Due to gravity, they are more likely to slide out through the front opening of the gripper. This is evident from the object-dropping rate, which was zero for 1D objects. Furthermore, the towel is larger and heavier than the cloth, making it more susceptible to the effects of gravity. From the perspective of multi-object training, we do not observe any noticeable differences between models trained on multiple datasets and those trained on a single dataset.

E. Experiments on Unseen Objects

We evaluate the generalization ability of the learned policy by testing it on previously unseen objects, including an 1D rope and a 2D napkin (Fig. 5). As shown in Table III, although remaining slightly higher for the 1D object, the success rate on unseen objects decreases to overall 65% compared with the 80% for the seen objects. Termination-related failures (early stopping and over-tracing) are also more frequent among the failure modes. This suggests that performance is more affected by differences in the visual appearance of the unseen objects (Fig. 5). This interpretation is also supported by the tactile features shown in Fig. 6: the unseen objects (bottom row) exhibit resemblance to the seen objects. Overall, the proposed model adapts to unseen objects but is less accurate in identifying the task endpoint.

VI. CONCLUSIONS

In this work, we present an IL-based approach for deformable object tracing using a single unified policy with both visual and tactile sensing, implemented in a teleoperation system with multi-modal feedback. To capture both local corrections and global task progression, we introduce a center loss and a task loss. Trained on demonstrations from four 1D and 2D deformable objects, the model achieves an overall success rate of 80%, and ablation studies validate the contribution of each component. Tests on two unseen objects further demonstrate generalizability, achieving a 65% success rate.

Several directions could further improve the current method: (1) combining mechanical intelligence (e.g., specially designed grippers [32], [33] such as V-shaped [34] or

hole-shaped [35]) with machine learning; (2) incorporating more deeply coupled sensor-fusion strategies in policy training [36], [37]; (3) expanding the evaluation with more trials and a larger set of objects; and (4) leveraging the modularity of the proposed components, including the visual–tactile teleoperation system, the local center loss, and the goal task loss, which can be readily integrated into other imitation learning algorithms (e.g., Diffusion Policy [27]) and further validated for general applicability.

REFERENCES

- [1] J. Zhu, A. Cherubini, C. Dune, D. Navarro-Alarcon, F. Alambeigi, D. Berenson, F. Ficuciello, K. Harada, J. Kober, X. Li *et al.*, “Challenges and outlook in robotic manipulation of deformable objects,” *IEEE Robotics & Automation Magazine*, vol. 29, no. 3, pp. 67–77, 2022.
- [2] Y. She, S. Wang, S. Dong, N. Sunil, A. Rodriguez, and E. Adelson, “Cable manipulation with a tactile-reactive gripper,” *The International Journal of Robotics Research*, vol. 40, no. 12-14, pp. 1385–1401, 2021.
- [3] N. Sunil, S. Wang, Y. She, E. Adelson, and A. R. Garcia, “Visuotactile affordances for cloth manipulation with local control,” in *Conference on Robot Learning*. PMLR, 2023, pp. 1596–1606.
- [4] H. Yin, A. Varava, and D. Kragic, “Modeling, learning, perception, and control methods for deformable object manipulation,” *Science Robotics*, vol. 6, no. 54, p. eabd8803, 2021.
- [5] F. Zhang and Y. Demiris, “Visual-tactile learning of garment unfolding for robot-assisted dressing,” *IEEE Robotics and Automation Letters*, vol. 8, no. 9, pp. 5512–5519, 2023.
- [6] L. Pecyna, S. Dong, and S. Luo, “Visual-tactile multimodality for following deformable linear objects using reinforcement learning,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 3987–3994.
- [7] C. Sun, B. Duan, K. Qian, and Y. Zhao, “Learning tactilemotor policy for robotic cable following via sim-to-real transfer,” in *International Conference on Intelligent Robotics and Applications*. Springer, 2023, pp. 63–73.
- [8] T. Zhao, V. Kumar, S. Levine, and C. Finn, “Learning fine-grained bimanual manipulation with low-cost hardware,” *Robotics: Science and Systems XIX*, 2023.
- [9] P. Wu, Y. Shentu, Z. Yi, X. Lin, and P. Abbeel, “Gello: A general, low-cost, and intuitive teleoperation framework for robot manipulators,” in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 12 156–12 163.
- [10] H. Luo, R. C. Quesada, F. E. Casado, N. Lingg, and Y. Demiris, “Interface matters: Comparing first and third-person perspective interfaces for bi-manual robot behavioural cloning,” in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 7385–7391.
- [11] H. Xue, J. Ren, W. Chen, G. Zhang, Y. Fang, G. Gu, H. Xu, and C. Lu, “Reactive diffusion policy: Slow-fast visual-tactile policy learning for contact-rich manipulation,” *arXiv preprint arXiv:2503.02881*, 2025.
- [12] J. Luo, C. Xu, X. Geng, G. Feng, K. Fang, L. Tan, S. Schaal, and S. Levine, “Multistage cable routing through hierarchical imitation learning,” *IEEE Transactions on Robotics*, vol. 40, pp. 1476–1491, 2024.
- [13] H. Lee, V. M. Staven, and C. Sloth, “Sim-to-real transfer for cable manipulation with tactile sensing and homotopy-constrained path planning,” in *2024 12th International Conference on Control, Mechatronics and Automation (ICCA)*. IEEE, 2024, pp. 382–387.
- [14] K. Suzuki, M. Kanamura, Y. Suga, H. Mori, and T. Ogata, “In-air knotting of rope using dual-arm robot based on deep learning,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 6724–6731.
- [15] V. Viswanath, K. Shivakumar, J. Kerr, B. Thananjeyan, E. Novoseller, J. Ichnowski, A. Escontrela, M. Laskey, J. E. Gonzalez, and K. Goldberg, “Autonomously untangling long cables,” *arXiv preprint arXiv:2207.07813*, 2022.
- [16] X. Huang, D. Chen, Y. Guo, X. Jiang, and Y. Liu, “Untangling multiple deformable linear objects in unknown quantities with complex backgrounds,” *IEEE Transactions on Automation Science and Engineering*, vol. 21, no. 1, pp. 671–683, 2023.
- [17] A. Doumanoglou, J. Stria, G. Peleka, I. Mariolis, V. Petrik, A. Karagakos, L. Wagner, V. Hlaváč, T.-K. Kim, and S. Malassiotis, “Folding clothes autonomously: A complete pipeline,” *IEEE Transactions on Robotics*, vol. 32, no. 6, pp. 1461–1478, 2016.
- [18] T. Lips, V.-L. De Gussemme, and F. Wyffels, “Learning keypoints for robotic cloth manipulation using synthetic data,” *IEEE Robotics and Automation Letters*, vol. 9, no. 7, pp. 6528–6535, 2024.
- [19] A. Canberk, C. Chi, H. Ha, B. Burchfiel, E. Cousineau, S. Feng, and S. Song, “Cloth funnels: Canonicalized-alignment for multi-purpose garment manipulation,” *arXiv preprint arXiv:2210.09347*, 2022.
- [20] A. Liu, K. Qian, B. Duan, and S. Luo, “Dynamic contrastive koopman operator for tactile servo of deformable linear objects,” *IEEE Transactions on Industrial Electronics*, 2025.
- [21] Y. Zhao, X. Jing, K. Qian, D. F. Gomes, and S. Luo, “Skill generalization of tubular object manipulation with tactile sensing and sim2real learning,” *Robotics and Autonomous Systems*, vol. 160, p. 104321, 2023.
- [22] Y. Zhao, K. Qian, B. Duan, and S. Luo, “Fots: A fast optical tactile simulator for sim2real learning of tactile-motor robot manipulation skills,” *IEEE Robotics and Automation Letters*, vol. 9, no. 6, pp. 5647–5654, 2024.
- [23] Z. Bi, K. Chen, C. Zheng, Y. Li, H. Li, and J. Ma, “Interactive navigation for legged manipulators with learned arm-pushing controller,” in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2025, pp. 9–16.
- [24] Z. Yu, W. Xu, S. Yao, J. Ren, T. Tang, Y. Li, G. Gu, and C. Lu, “Precise robotic needle-threading with tactile perception and reinforcement learning,” in *Conference on Robot Learning*. PMLR, 2023, pp. 3266–3276.
- [25] W. Zhou, S. Bajracharya, and D. Held, “Plas: Latent action space for offline reinforcement learning,” in *Conference on Robot Learning*. PMLR, 2021, pp. 1719–1735.
- [26] B. Fang, S. Jia, D. Guo, M. Xu, S. Wen, and F. Sun, “Survey of imitation learning for robotic manipulation,” *International Journal of Intelligent Robotics and Applications*, vol. 3, no. 4, pp. 362–369, 2019.
- [27] C. Chi, Z. Xu, S. Feng, E. Cousineau, Y. Du, B. Burchfiel, R. Tedrake, and S. Song, “Diffusion policy: Visuomotor policy learning via action diffusion,” *The International Journal of Robotics Research*, p. 02783649241273668, 2023.
- [28] N. Gu, K. Kosuge, and M. Hayashibe, “Tactilealoha: Learning bimanual manipulation with tactile sensing,” *IEEE Robotics and Automation Letters*, 2025.
- [29] I. 3D Systems, “Touch haptic device,” <https://www.3dsystems.com/haptics-devices/touch>, 2025, accessed: 2025-09-02.
- [30] S. Wang, Y. She, B. Romero, and E. Adelson, “Gelsight wedge: Measuring high-resolution 3d contact geometry with a compact robot finger,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 6468–6475.
- [31] T. Yoshikawa, “Manipulability of robotic mechanisms,” *The international journal of Robotics Research*, vol. 4, no. 2, pp. 3–9, 1985.
- [32] Y. Zhou, X. Jiang, D. Chen, Y. Guo, and Y. Liu, “Design of a gripper for cable assembly with integrated in-hand cable manipulation functions,” in *2021 IEEE International Conference on Robotics and Biomimetics (ROBIO)*. IEEE, 2021, pp. 1929–1934.
- [33] Y. Zhou, P. Zhou, S. Wang, and Y. She, “In-hand singulation, scooping, and cable untangling with a 5-dof tactile-reactive gripper,” *Advanced Robotics Research*, p. 202500020, 2025.
- [34] M. Yu, B. Liang, X. Zhang, X. Zhu, L. Sun, C. Wang, S. Song, X. Li, and M. Tomizuka, “In-hand following of deformable linear objects using dexterous fingers with tactile sensing,” in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 13 518–13 524.
- [35] J. Zhu, B. Navarro, R. Passama, P. Fraisse, A. Crosnier, and A. Cherubini, “Robotic manipulation planning for shaping deformable linear objects with environmental contacts,” *IEEE Robotics and Automation Letters*, vol. 5, no. 1, pp. 16–23, 2019.
- [36] Y. Ou, S. Xia, J. Fan, C. Zhou, Y. Huang, P. Shuai, Y. Xue, and Z. Hou, “PI-vap: A tightly coupled self-localization framework for underwater robots using point-line features and visual-acoustic-pressure sensor fusion,” *IEEE/ASME Transactions on Mechatronics*, vol. 30, no. 6, pp. 4116–4128, 2025.
- [37] Z. Wu, Y. Zhao, and S. Luo, “Convitac: Aligning visual-tactile fusion with contrastive representations,” in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2025, pp. 8545–8552.