IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

Tactile-Driven Dexterous In-Hand Writing via Extrinsic Contact Sensing

Can Zhao [1], Lingzi Xie[2], Bidan Huang [2], Shuai Wang[2], and Daolin Ma [1]

***Abstract*—Dexterous in-hand manipulation, especially involving interactions between grasped objects and external environments, remains a formidable challenge in robotics. This study tackles the complexities of in-hand manipulation under extrinsic contact through a representative three-finger handwriting task. We propose a hybrid arm-hand coordination framework that combines reinforcement learning with compliance control, offering both flexibility and robustness. Leveraging tactile sensors embedded in each finger, our tactile-driven estimation model dynamically predicts in-hand object pose and external contact, eliminating the need for fixed contact states. The proposed framework is first validated in simulation, where it successfully executes diverse writing tasks with accurate contact sensing. Sim-to-Real transfer is achieved through systematic calibration of finger joints and tactile sensors, supported by domain randomization. Real-world experiments further demonstrate the system's adaptability to writing tools with varying physical properties—such as radius, length, mass, and friction—while maintaining stability across different trajectories. Also see https://inhandwriting.github.io/. This work advances robotic manipulation capabilities in unstructured environments.**

***Index Terms*—In-Hand manipulation, force and tactile sensing, reinforcement learning**

I. Introduction

DEXTEROUS in-hand manipulation with extrinsic contact is vital for numerous real-world robotic applications, such as tool usage (e.g., fastening screws with a screwdriver) and component assembly (e.g., aligning and fitting gears in mechanical assemblies). These tasks demand simultaneously managing in-hand dynamics while maintaining stable contact between the manipulated object and the external environment. Although notable advancements have been made in learning-based approaches that utilize vision, proprioception, and tactile sensing for in-hand manipulation tasks [1–4], most state-of-the-art methods assume that the manipulated object remains isolated from extrinsic contact. However, interactions with the environment are often inevitable in practical scenarios, and these external contacts frequently play a critical role in task success [5, 6]. Recent studies have begun exploring extrinsic contact sensing and control, but the prevalent reliance on fixed gripper-object contact constrains the flexibility and practicality of these methods [7, 8].

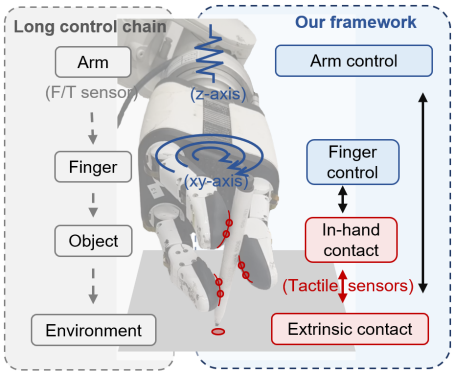

Fig. 1. Overview of our framework (right), which utilizes tactile-based feedback for direct and precise contact sensing, in contrast to the indirect wrist-mounted force/torque sensing (left).

A representative task of in-hand manipulation with extrinsic contact is **In-hand writing**, where a robotic hand dynamically controls a tool (e.g., a pen) while interacting with an external surface. During writing, the robotic hand must coordinate multi-finger joints to maintain a stable grasp on the writing tool and follow reference trajectories (**In-hand contact control**), while adjusting the tool's contact with the writing surface (**Extrinsic contact control**). Successfully achieving this dual-level control introduces three major challenges: i) estimating and controlling both in-hand and extrinsic contact states simultaneously; ii) ensuring a secure grasp on the tool while dynamically adapting to variations in tool properties and external disturbances; and iii) coordinating multi-finger and wrist movements to balance stability and adaptability.

To address these challenges, we propose a tactile-driven framework that synergizes in-hand and extrinsic contact control through a hybrid arm-hand coordination strategy. Our approach leverages rich tactile feedback from fingertip sensors on a three-finger robotic hand to predict continuous object states and external contact conditions. As described in Fig. 1, fingertip tactile sensing, which directly interacts with the manipulated object, provides more localized and detailed contact dynamics than wrist-mounted force/torque sensors. This immediate feedback is essential for maintaining fine-grained contact control during dexterous in-hand manipulation tasks, improving both precision and robustness. By analyzing these tactile signals, we extract key features that enable reinforcement learning (RL) to effectively regulate in-hand contact and trajectory tracking. Our framework further integrates RL-

Manuscript received: March 28, 2025; Revised June 2, 2025; Accepted July 7, 2025. This paper was recommended for publication by J. Kober upon evaluation of the Associate Editor and Reviewers' comments. This work of C. Zhao was supported by the Tencent Robotics X, and by the National Natural Science Foundation of China (No. 12272220). This work was done when C. Zhao and L. Xie were interns at Tencent Robotics X. *(Corresponding author: Bidan Huang)*.

[1]C. Zhao, D. Ma are with the School of Ocean & Civil Engineering, Shanghai Jiao Tong University, Shanghai 200240, China (e-mail: can.zhxx@sjtu.edu.cn, daolinma@sjtu.edu.cn).

[2]L. Xie, B. Huang, and S. Wang are with Tencent Robotics X, Shenzhen 518054, China (e-mail: lingzixie, bidanhuang, shawnshwang@tencent.com).

Digital Object Identifier (DOI): see top of this page.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

based finger motion control with compliant position-force wrist adjustments, ensuring stable yet adaptive contact with external surfaces.

The main contributions of this work are as follows:

- **Tactile-driven Estimation**: Develop a tactile-driven model that predicts object states, in-hand and extrinsic contacts without visual input, eliminating reliance on fixed gripper-object contact.
- **In-hand & Extrinsic Contact Control**: Propose a hybrid arm-hand coordination framework that combines RL-based finger control with compliant wrist adjustments, enabling seamless control of in-hand and external contacts.
- **Real-world Demonstration**: Demonstrate a three-finger robotic hand system capable of stable writing in both simulated and real-world settings. The system adapts to tools with varying properties (e.g., radius, length, mass, and friction) while accurately executing different trajectories.

This work highlights the potential of tactile-driven in-hand manipulation with extrinsic contact, paving the way for more dexterous robotic applications in unstructured environments.

II. Related work

A. In-hand Manipulation

Dexterous in-hand manipulation has been a focal point in robotics for decades [9, 10]. Early approaches were primarily model-based, relying on analytical planning for precise object control [11, 12]. While advancement in hardware and algorithm design have significantly enhanced the dexterity of these methods [13–15], their performance is often constrained by stringent assumptions about system dynamics and contact interactions. Learning-based approaches have emerged as a powerful alternative, offering increased generalization and dexterity. Many of these approaches rely heavily on vision-based systems for object tracking and control [16–18]. However, vision alone often struggles in highly dynamic scenes, especially when occlusions occur or fine-grained local contact information is critical. In such scenarios, tactile sensing has proven invaluable, providing detailed insights into interactions [19]. Recent work has increasingly explored the combination of proprioception and tactile sensing for in-hand manipulation tasks, focusing on operations such as object reorientation [1–4, 20, 21]. One notable related study is a stick-steering task presented by [22], which utilizes tactile feedback to compute the contact center at each fingertip. In contrast to this, our work goes beyond in-hand manipulation by considering extrinsic contacts between the grasped object and the environment.

B. Extrinsic Contact Sensing and Control

Extrinsic contacts are crucial in tasks that involve interactions between a grasped object and its environment, such as insertion, tool use, and other contact-heavy applications. Recent advancements have explored various methods for sensing and controlling these external contacts. The work [23] proposes a framework for localizing extrinsic contacts with different geometries (points, lines, patches), using visuo-tactile feedback and rigid-body constraints. Kim et al. [7] further use factor graphs to simultaneously estimate and control interactions between the gripper, object and environment by fusing tactile and kinematic measurements. Higuera et al. [24] innovatively combine tactile feedback, proprioception, and implicit geometry representations to estimate diverse object shapes and extrinsic contact configurations. Meanwhile, Kim et al. [25] introduce a vision-based method to localize extrinsic contact using only RGB-D data, bypassing the need for tactile or force-torque sensors. Models like neural deforming contact fields (NDCF) [26] combine visuo-tactile data to track contact patches and deformations, with subsequent iteration NDCF-v2 [27] achieving successful sim-to-real transfer. However, most approaches assume rigid or non-slip gripper-object contact. In contrast, our approach leverages multi-finger tactile sensing to dynamically estimate both in-hand and extrinsic contacts, accommodating various contact conditions.

Effective extrinsic contact control is equally critical. Doshi et al. [28] develop a contact configuration estimator coupled with a mode-dependent controller for manipulating unknown objects, extended by Taylor et al. [29] with factor graph-based estimation for multi-contact regulation. Shirai et al. [30] introduce trajectory optimization for tool-based pivoting, while Oller et al. [8] model compliance for non-prehensile manipulation using deformable tactile sensors. In this work, a RL policy is utilized to continuously track extrinsic contact positions, integrated with a complaint position-force control strategy to maintain stable and adaptive extrinsic contact.

C. Robotic Writing

Robotic handwriting, particularly in-hand writing with multi-finger robotic hands, is an underexplored task due to its complexity [31]. Conventional approaches often rely on robotic arms to generate writing motions, such as studies focused on trajectory planning and manipulation for welding and assembly tasks [32, 33]. While available for above applications, these methods differ significantly from flexible human writing, where finger movements contribute the majority of the motion, and the wrist is typically anchored to the writing surface [34]. Recent advances have introduced learning-based techniques for robotic handwriting, enabling the control of pens and similar tools [16, 35]. However, most of these methods are limited, with some operating only in simulation or generating trajectories without meaningful interaction with the writing surface, which is crucial for capturing realistic contact dynamics. Unlike previous works that either rely solely on arm movements or remain confined to simulation, our framework integrates fine finger motion with compliant wrist control, ensuring seamless arm-hand coordination. Furthermore, it tackles sim-to-real transfer challenges systematically, advancing the capabilities of robotic handwriting.

III. Preliminaries

This paper aims to develop a tactile-driven three-finger robotic hand system capable of in-hand writing while maintaining stable extrinsic contacts. The problem is formulated based on contact dynamics, with the state space defined for both in-hand and extrinsic contact modeling.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

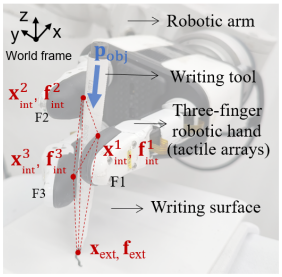

State	Description
$\mathbf{x}_{wt} \in \mathbb{R}^3$	Wrist Pos.
$\mathbf{x}_{fg} \in \mathbb{R}^6$	Finger Joint Pos.
$\mathbf{p}_{obj} \in \mathbb{R}^3$	Writing Tool Pose
$\mathbf{x}_{tg} \in \mathbb{R}^2$	Target Pos.
$\mathbf{v}_{tg} \in \mathbb{R}^2$	Target Velocity
$\mathbf{x}_{int} \in \mathbb{R}^{3\times3}$	In-hand Contact Pos.
$\mathbf{f}_{int} \in \mathbb{R}^3$	In-hand Contact Force
$\mathbf{x}_{ext} \in \mathbb{R}^3$	Extrinsic Contact Pos.
$\mathbf{f}_{ext} \in \mathbb{R}$	Extrinsic Contact Force

Fig. 2. System model and state definition of the robotic in-hand writing system, featuring a three-finger hand on a robotic arm with curved tactile sensor arrays. The writing tool interacts with a fixed horizontal plane.

A. System Modeling

Without relying on visual feedback, the robotic in-hand writing system is modeled as shown in Fig. 2. The system comprises a three-finger robotic hand mounted on a robotic arm, featuring eight fully actuated joints. Each finger has two actuated joints, while fingers 2 and 3 (F2 and F3) include base rotation joints capable of independently spreading up to 90 degrees. Tactile feedback is provided by curved sensor arrays on each fingertip [36], consisting of 128 piezoresistive taxels embedded in a 1-mm silicone layer, measuring normal forces. This design allows for relative slip and brief partial contact loss, relying on multi-point contact and compliant curved fingertip surfaces. The writing tool is modeled as a slender cylindrical object with a small pen tip, interacting with a fixed horizontal plane that serves as a stable extrinsic contact surface. Its interaction with the external surface is simplified to a single point of contact, under the assumption that only the vertical force component is explicitly estimated. Tangential forces are treated as unmodeled disturbances and are implicitly handled through domain randomization in friction parameters.

B. State Space Definition

The system's state space consists of several key variables that govern the robotic in-hand writing process, as listed in Fig. 2. These variables encompass the robotic joint configurations, with the wrist position $\mathbf{x}_{wt}$ and finger joint states $\mathbf{x}_{fg}$, which control the hand's movement and positioning. The in-hand contact positions and forces ($\mathbf{x}_{int}$, $\mathbf{f}_{int}$) are provided by the tactile sensors on the fingertips, offering feedback for precise control. The writing tool's pose is represented by a unit vector $\mathbf{p}_{obj}$, while the initial object pose is denoted as $\mathbf{p}_{prior}$. The target trajectory position and velocity ($\mathbf{x}_{tg}$, $\mathbf{v}_{tg}$) ensure the tool moves smoothly and accurately on the writing surface. Moreover, the extrinsic contacts ($\mathbf{x}_{ext}$, $\mathbf{f}_{ext}$) are estimated based on the in-hand contact positions and the writing surface's height, providing critical information about the interaction between the tool and the external environment. All variables are ultimately expressed in the world frame, except for the finger joint positions $\mathbf{x}_{fg}$, which are normalized and dimensionless.

IV. Method

Our method consists of three core components, as shown in Fig. 4. First, the state estimation module (Sec. IV-A) uses tactile sensor feedback to accurately estimate the writing tool's pose and contact states. Second, the hybrid hand-arm coordination algorithm (Sec. IV-B) integrates: i) an RL policy that enables the three-finger robotic hand to track target trajectories by managing the writing tool's pose and force transmission while maintaining stable object configurations; and ii) a compliance control strategy that ensures stable external contact forces through robotic wrist adjustments. Finally, sim-to-real transfer techniques (Sec. IV-C) ensure the effective deployment of the entire pipeline in real-world environments.

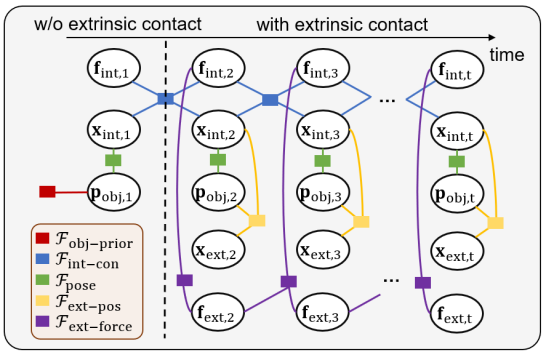

Fig. 3. Tactile-driven factor graph, including state variables (object pose, contact positions and forces) and measurement constraints (initial pose prior, in-hand and extrinsic contact data).

A. In-hand and Extrinsic Contact Estimation

In this section, we develop a factor graph-based state estimation framework to infer the real-time pose of the writing tool ($\mathbf{p}_{obj}$), intrinsic contact states ($\mathbf{x}_{int}$, $\mathbf{f}_{int}$), and extrinsic contact conditions ($\mathbf{x}_{ext}$, $\mathbf{f}_{ext}$) by fusing tactile and proprioceptive measurements, without visual input. The estimation process is formalized by defining the optimization variables and posing the problem as:

$$\Theta = \{\mathbf{p}_{obj}, \mathbf{x}_{int}, \mathbf{f}_{int}, \mathbf{x}_{ext}, \mathbf{f}_{ext}\}, \tag{1}$$

$$\Theta^* = \arg\max_{\Theta} P(\mathbf{p}_{obj}, \mathbf{x}_{int}, \mathbf{f}_{int}, \mathbf{x}_{ext}, \mathbf{f}_{ext}, \mathbf{p}_{obj,0}) \tag{2}$$

$$\equiv \min_{\Theta} \sum_{t=1}^{T} \sum_{i} \|\mathcal{F}_{i,t}(\Theta_t)\|^2_{\Sigma_{i,t}},$$

where the maximum a posteriori (MAP) estimation reduces to minimizing a nonlinear least-squares objective under Gaussian noise. Here, T denotes the total number of time steps, $\mathcal{F}_{i,t}$ represents the i-th residual at time step t, and $\Sigma_{i,t}$ is the corresponding covariance matrix.

As shown in Fig. 3, circles represent the variables (unknown states), while cubes denote factors (constraints or measurements). Below are the key factors included in the graph:

- **Object Prior Factor**: Ensures that the writing tool starts from a known initial pose,

$$\mathcal{F}_{\text{object-prior}} = \|\mathbf{p}_{obj,t} - \mathbf{p}_{prior}\|^2_{\Sigma_{\text{pose-prior}}}.$$

- **Object Pose Estimation Factor**: Estimates the object's pose based on the geometric relationship between in-hand contact positions $\mathbf{x}_{int}$,

$$\mathcal{F}_{pose,t} = \left\| \mathbf{p}_{obj,t} - \frac{(\mathbf{x}^2_{int,t} - \mathbf{x}^1_{int,t}) \times (\mathbf{x}^3_{int,t} - \mathbf{x}^1_{int,t})}{\|\mathbf{x}^2_{int,t} - \mathbf{x}^1_{int,t}\| \|\mathbf{x}^3_{int,t} - \mathbf{x}^1_{int,t}\|} \right\|^2_{\Sigma_{pose}},$$

where $\times$ denotes the cross product.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

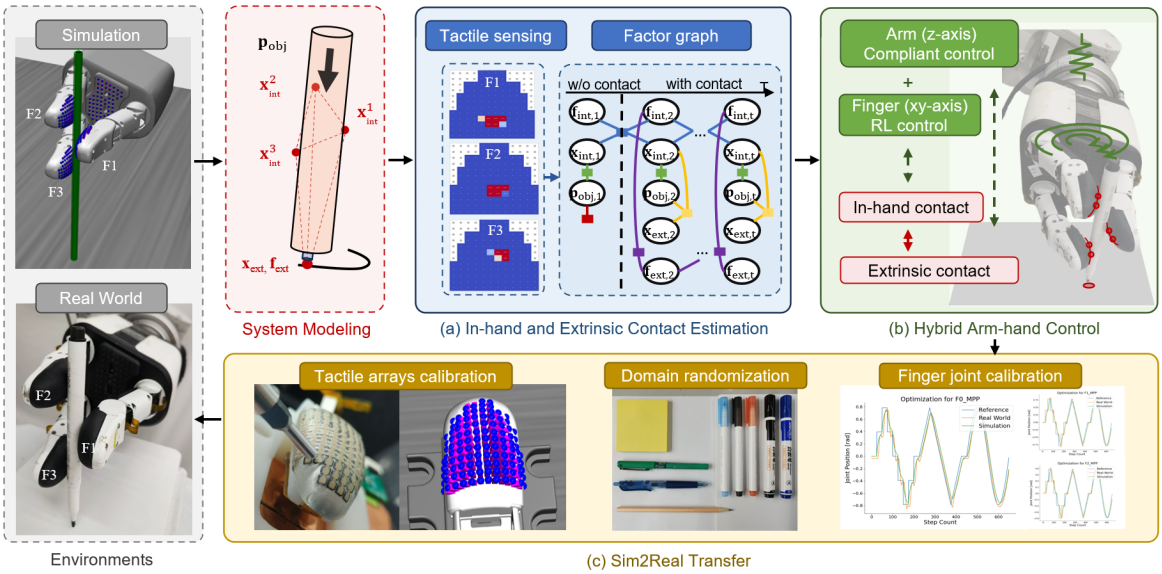

Fig. 4. Overall pipeline. (a) In-hand and extrinsic contact estimation via factor graph; (b) hybrid arm-hand coordination control framework that combines fingers' RL policy with wrist's compliant control; (c) sim-to-real transfer for real-world in-hand writing, including system identification and domain randomization.

- **In-hand Contact Constraint Factor**: Constrains in-hand contact positions and forces based on tactile data. In case of temporary local detachment, the in-hand contact is assumed to remain the same as in the previous timestep,

$$\mathcal{F}_{\text{int-con},t} = (\|\mathbf{x}_{\text{int},t} - \mathbf{x}_{\text{int},t-1}\|^2 + \|\mathbf{f}_{\text{int},t} - \mathbf{f}_{\text{int},t-1}\|^2)_{\Sigma_{\text{consist}}} + \left\|\mathbf{x}_{\text{int},t} - \frac{\sum_j w_{j,t}\mathbf{x}_{j,t}}{\sum_j w_{j,t}}\right\|^2_{\Sigma_{\text{int-pos}}} + \|\mathbf{f}_{\text{int},t} - \sum_j \mathbf{f}_{j,t}\|^2_{\Sigma_{\text{int-force}}},$$

  where $\mathbf{w}_j$, $\mathbf{x}_j$ and $\mathbf{f}_j$ represents the weight factor based on taxel pressure, the coordinates measured by the tactile taxels, and the force measured by each tactile taxel.
- **Extrinsic Contact Point Constraint Factor**: Constrains the extrinsic contact point to remain on the writing surface,

$$\mathcal{F}_{\text{ext-pos},t} = \|\mathbf{x}_{\text{ext},t} - (\mathbf{x}_{\text{cm},t} + \mathbf{p}_{\text{obj},t} \cdot \Delta_t)\|^2_{\Sigma_{\text{ext-pos}}},$$

  where $\mathbf{x}_{\text{cm}} = \frac{1}{3}\sum_{i=1}^{3} \mathbf{x}^i_{\text{int}}$ denotes the centroid of three in-hand contact positions in the global coordinate system, $\mathbf{p}_{\text{obj}}$ is the object's pose, and Δ is the vertical distance between the in-hand contact center and the external surface A.
- **External Contact Force Balance Factor**: Maintains force equilibrium between the tool and its environment, ensuring that the estimated extrinsic force—modeled only in the normal direction—remains stable on the writing surface.

$$\mathcal{F}_{\text{ext-force},t} = \|\mathbf{f}_{\text{ext},t} - (\mathbf{f}_{\text{g}} - (\mathbf{f}_{\text{int},t})_z)\|^2_{\Sigma_{\text{force-balance}}} + \|\mathbf{f}_{\text{ext},t} - \mathbf{f}_{\text{ext},t-1}\|^2_{\Sigma_{\text{ext-consist}}},$$

  where $\mathbf{f}_{\text{g}} = m \cdot \mathbf{g}$ is the gravitational force acting on the object, and $(\mathbf{f}_{\text{int},t})_z$ represents the sum of the vertical components of the in-hand contact forces.

To solve the optimization problem in Eq. (2), the Incremental Smoothing and Mapping (iSAM) algorithm [37] from the GTSAM library is adopted. Although the estimation objective spans all time steps $t = 1, \ldots, T$, iSAM optimization is performed incrementally, updating the factor graph as new measurements arrive. Starting from an initial estimate $\mathbf{x}_0 = \{\mathbf{p}_{\text{obj},0}, \mathbf{x}_{\text{int},0}, \mathbf{f}_{\text{int},0}, \mathbf{x}_{\text{ext},0}, \mathbf{f}_{\text{ext},0}\}$, it adds new variables and factors while maintaining sparsity through periodic reordering and relinearization. This enables real-time state estimation aligned with tactile and proprioceptive inputs.

B. Hybrid Arm-hand Control Policy

To achieve stable and adaptive writing, we propose a hybrid arm-hand control framework that combines reinforcement learning (RL)-based in-hand control with compliant wrist adjustments. This framework effectively balances dexterity and stability by decoupling complex finger movements from wrist control, ensuring consistent extrinsic contact.

1) Reinforcement learning policy: The finger motion control is formulated as a finite-horizon, discounted Markov decision process (MDP) in which a robotic hand (agent) selects actions to interact with a stochastic environment over time. The goal is to optimize the control policy using a deep reinforcement learning (RL) algorithm. Specifically, Proximal Policy Optimization (PPO) is trained for 8e6 timesteps across 12 parallel environments on a desktop with an Intel i5-12600KF CPU and NVIDIA RTX 3060 GPU. The policy uses an MLP with two hidden layers (256 units each), ReLU activations, and separate heads for policy and value outputs.

The observation space is designed to capture essential task-related information, focusing on: i) **In-hand contact features**—fingertip contact positions $\mathbf{x}_{\text{int}}$ and forces $\mathbf{f}_{\text{int}}$ from tactile sensors help the system to detect subtle changes in object orientation or position, ensuring stability and control even during slight slips or partial contact loss; ii) **Object pose** $\mathbf{p}_{\text{obj}}$—referenced to the initial pose $\mathbf{p}_{\text{prior}}$, links hand motion to pen-tip trajectory on the writing surface; iii) **Extrinsic contact features** $\mathbf{x}_{\text{ext}}$ and $\mathbf{f}_{\text{ext}}$—reflect the interaction dynamics between the writing tool and external surface, allowing real-time slip compensation and stable contact. These observation features not only help the RL agent maintain precise pen control and consistent contact with the surface while adapting

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

to task changes, but also support generalization to new writing scenarios and broader robotic manipulation tasks. The action space specifies the desired displacements of six finger joints. The reward function is designed to encourage the agent to manipulate the writing tool effectively while maintaining stable contact. The overall reward function is defined as:

$$r = r_{\text{pos}} + r_{\text{height}} + r_{\text{time}} + r_{\text{smooth}} + r_{\text{con}} + r_{\text{ext}}, \tag{3}$$

where each term contributes to optimizing different aspects of task performance. The position reward r_{pos} encourages the pen tip to approach the target position, and r_{height} promotes alignment with the predefined writing plane. The time reward r_{time} maintains task persistence, whereas r_{smooth} penalizes abrupt changes in action, ensuring smoother finger movements. Moreover, the in-hand contact reward r_{con} encourages stable grasping by maintaining appropriate fingertip contact forces, and the extrinsic contact reward r_{ext} penalizes large deviations in external forces, preventing slips and maintaining stability.

2) Compliant Position-Force Control: To maintain consistent contact between the pen tip and the writing surface, we design a compliant position-force control mechanism for the wrist. Inspired by human strategies—where the wrist moves the pen tip toward the paper, adjusts the applied force upon contact, and then uses the fingers for writing—our approach mirrors this process. The wrist control consists of three modes: position control, force control, and smooth transition. Initially, the wrist moves toward the paper using position control:

$$u_{\text{pos}} = K_{p,w}(z_{\text{target}} - z_{\text{wrist}}) + K_{d,w}\dot{z}_{\text{wrist}}, \tag{4}$$

where z_{target}, z_{wrist} and $\dot{z}_{\text{wrist}}$ denote the target height, current height, and current velocity of the wrist, respectively. $K_{p,w}$ and $K_{d,w}$ are the proportional and derivative gains.

Once the pen tip establishes stable contact with the surface, the system fully switches to force control:

$$u_{\text{force}} = K_{f,w}(f_{\text{desired}} - \mathbf{f}_{\text{ext}}), \tag{5}$$

where f_{desired} and $\mathbf{f}_{\text{ext}}$ represent the desired and actual contact force, respectively, and $K_{f,w}$ is the force control gain.

To ensure a rapid yet smooth transition from position to force control, we introduce a transition mode triggered when the detected contact force exceeds a predefined threshold f_{th} for N control cycles, which minimizes oscillations caused by the paper's stiffness. During the i_{th} cycle of the smooth transition, the position command increment in the vertical direction is given by: $u_{\text{pos}}(i) = (1 - i/N) \cdot u_{\text{pos}} + i/N \cdot u_{\text{force}}$, where i is initialized to 1 and increases to N control cycles.

The wrist motion primarily occurs along the z-direction, as the fingers have limited freedom in this axis. However, controlling the wrist in the vertical direction is essential for adjusting the external contact force. This hybrid control strategy dynamically adapts to changing external forces and writing trajectories. By decoupling in-hand control from wrist adjustments, it ensures smooth transitions and enhances overall system performance in real-world applications.

C. Sim-to-Real Transfer

In this study, we first train the RL agent in a simulated Mujoco environment that closely replicates the real robot setup, including tactile sensors, as shown in Fig. 4. To enable effective sim-to-real transfer, it is essential to address critical factors affecting robot performance: accurate modeling of tactile signals, system identification for sensor and actuator alignment, and domain randomization to enhance generalization.

1) Tactile Signal Modeling: To enhance sim-to-real consistency, we simulate tactile signals in MuJoCo with both spatial and temporal fidelity. Each fingertip tactile sensor is modeled as a distributed array of taxels, replicating the layout and dead-zone characteristics of real-world tactile sensors. Contact detection is threshold-based, with mean positions and vector-summed normal forces computed from active taxels. To emulate real-world noise, we apply dropout smoothing (zero-hold), signal rounding, and gain scaling informed by empirical calibration. These heuristics capture common tactile limitations such as jitter, dead zones, and quantization.

2) System identification: To minimize discrepancies in policy transfer, system identification on the real hardware are performed to ensure quantitative consistency between simulated and real responses. Specifically, we calibrate the force–voltage mapping of the tactile arrays, which serves as the empirical basis for gain scaling and noise modeling in simulation. Each taxel's voltage output v_{tac} is mapped to the applied force F_{tac} using a calibration platform that applies known normal forces. The sensor response is modeled using a nonlinear function with a dead zone, expressed as: $F_{\tilde{\text{tac}}} = a \cdot \frac{b v_{\tilde{\text{tax}}}}{1 - b v_{\tilde{\text{tax}}}} + c$. The fitting parameters are identified using the Levenberg-Marquardt (LM) algorithm, achieving an average fitting error of less than 5%.

Simulated finger joint responses are also calibrated to match real-world counterparts. Each joint operates under position control in simulation, governed by the position control gain $K_p^{(sim)}$ and joint damping $K_d^{(sim)}$. A reference joint angle sequence is applied with a control frequency of 250 Hz, using step, triangular, and sinusoidal signals. Optimization is performed using CMA-ES [38], formulated as: $\min_{K_p^{(sim)}, K_d^{(sim)}} \sum_{t=0}^{T} \|\theta_{\text{sim},t} - \theta_{\text{real},t}\|^2, \quad \text{s.t.} K_p^{(sim)}, K_d^{(sim)} > 0$. The final average gains are $K_p^{(sim)} = 33.90$ and $K_d^{(sim)} = 2.63$, with the MPP joints exhibiting slightly higher values due to their more complex dynamics. These calibration and optimization processes enhance the fidelity of the simulated joint responses and improve alignment with real-world behavior.

3) Domain Randomazation: Domain randomization improves system robustness by introducing variability in object properties and environmental conditions. Key parameters—such as the pen's mass, geometry, initial grasping position, and friction coefficients—are randomly sampled from broad ranges to capture real-world uncertainty. To further promote policy generalization, the training process includes diverse target trajectories—including circular, straight, heart-shaped, figure-eight, and spiral paths—as illustrated in Fig. 5 (a). Exposure to these variations during training enables the learned policy to remain resilient under unpredictable changes, improving the reliability of sim-to-real transfer.

V. Evaluation

This section evaluates the proposed framework in both simulated and real-world environments, demonstrating accu-

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

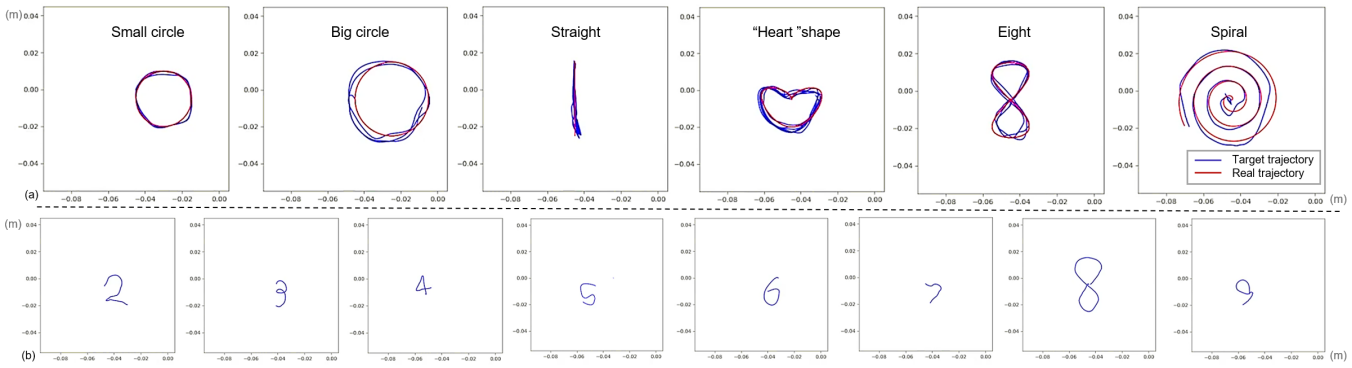

Fig. 5. In-hand writing performance in simulation. (a) Trained trajectories, including circles, straight lines, "heart" shapes, eights, and a spiral (from left to right), with blue indicating the actual trajectories and red representing the target trajectories. (b) Unseen trajectories, showcasing various digits.

TABLE I
QUANTITATIVE RESULTS OF SIMULATED IN-HAND WRITING

Shape	Circle		Straight	"Heart"	Eight	Spiral
	small	big				
Trajectory perimeter /m	0.094	0.126	0.004	0.093	0.122	0.631
x_{err} /$10^{-3}\cdot$ m	0.087	0.845	0.063	0.613	3.624	3.227
y_{err} /$10^{-3}\cdot$ m	0.074	0.515	1.726	1.162	3.039	3.247

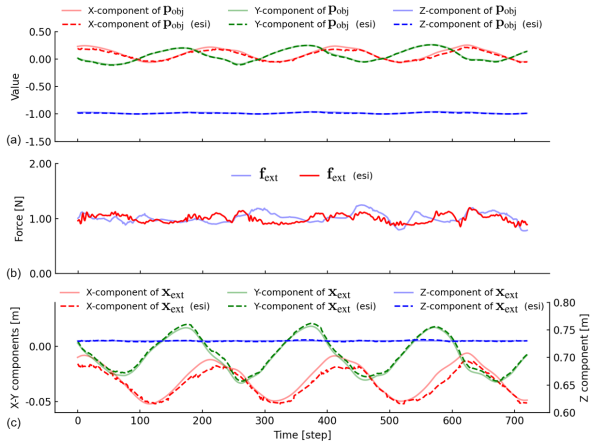

Fig. 6. Tactile estimation pipeline evaluation, including $\mathbf{p}_{\text{obj}}$, $\mathbf{f}_{\text{ext}}$, and $\mathbf{x}_{\text{ext}}$.

rate tactile estimation, task generalization, and adaptability to diverse tools and surfaces. Ablation studies emphasize the importance of contact perception and arm-hand collaboration.

A. Simulation Performance

To assess the learning algorithm, we first train RL policies on various reference trajectories in simulation, including circles, straight lines, heart shapes, figure-eights, and a spiral. As shown in Fig. 5 (a), the trained policies successfully execute in-hand writing, with the actual pen tip trajectories (blue) closely following the reference trajectories (red). TABLE I summarizes the average position errors for different tasks. For instance, we analyze the "circle" policy under varying radii. As the radius increases, tracking errors grow due to the limited range of finger motion, affecting control accuracy. While increasing the pen length theoretically extends the trajectory range, it reduces finger control stability. Optimal tracking performance is achieved for trajectories within a 5 cm range. The figure-eight and spiral trajectories pose additional challenges, with the former requiring sharp turns and the latter demanding precise control, especially in the inner loops, leading to larger errors.

To evaluate the generalization capability of the trained policies, we also perform digit-writing tasks in simulation, where the robotic hand automatically selects the most appropriate pre-trained policy. As demonstrated in Fig. 5 (b), the learned policies exhibit strong generalization, effectively executing unseen trajectories and highlighting the system's robustness and adaptability across different writing tasks.

Furthermore, we evaluate the precision and real-time performance of the tactile estimation pipeline, which estimates the object's pose, extrinsic contact forces, and positions during circle-writing tasks. As shown in Fig. 6, the close agreement between estimated values and ground truth demonstrates the effectiveness of the tactile perception module in enhancing contact stability and trajectory accuracy. The factor graph estimator runs in parallel with the RL policy, with per-step updates under 5ms on an Intel i5-12600KF CPU in simulation and around 10ms on a Xeon 6242R in real-world execution—providing low-latency feedback for closed-loop control.

B. Real-World Performance

To assess the real-world performance of the proposed framework, experiments are conducted using a three-finger robotic hand, mounted on a six-axis industrial robotic arm. The framework operates solely on tactile feedback, with the RL policy running at 50Hz, while finger joint control and wrist compliant control are executed at 250Hz. Each writing trial begins after the robot grasps the tool with its three fingertips.

1) Ablation Study: To validate the necessity of explicitly estimating extrinsic contact status and arm-hand coordination control—as opposed to relying solely on an end-to-end policy—we conduct ablation studies on the circle-drawing task. The compliant control is set to maintain a 0.5 N external contact force at the pen tip. We compare our proposed framework against two baseline variants: i) **Raw SP** adapts the "stir policy" from [22], originally trained via RL for mid-air object stirring using fingertip tactile and joint state inputs, and fine-tunes it for the writing task—without wrist control or explicit extrinsic contact estimation. ii) **Compliant SP** extends Raw SP with a compliant wrist controller to passively maintain contact,

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

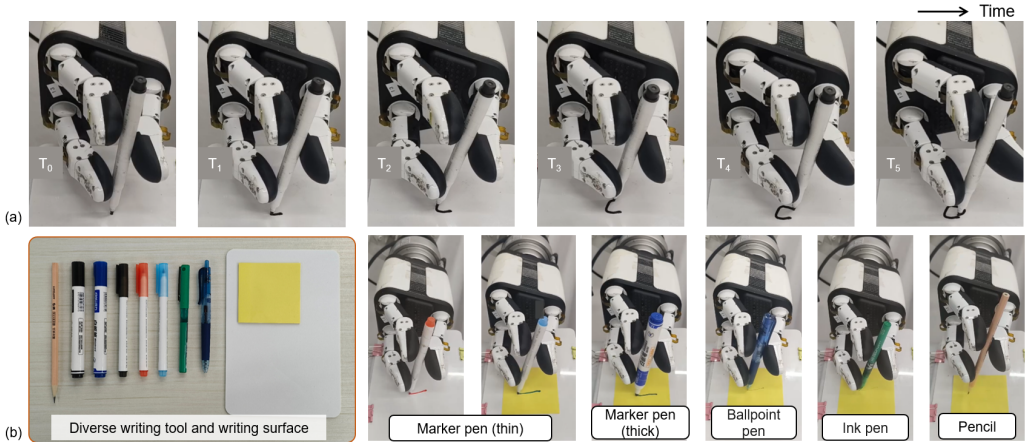

Fig. 7. Overall performance of our method: A three-fingered robotic hand mounted on a Franka arm performs in-hand writing without any visual input (top), demonstrating adaptability across various writing tools and surfaces with different physical properties (bottom).

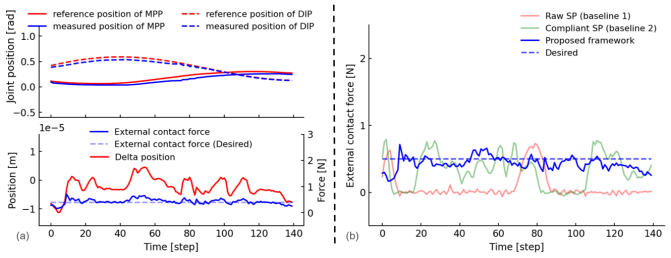

Fig. 8. Ablation study. (a) finger and arm response curves for our framework, (b) external contact forces for policy variations during one full circular arc.

but still without extrinsic contact estimation. iii) **Proposed Framework** (ours) incorporates both explicit contact state estimation and arm-hand coordination control. Each method is evaluated on its ability to maintain stable contact force and produce complete writing patterns over 140 policy inference steps (about one full circular arc).

Results in Fig. 8 (b) indicate that **Raw SP** struggles to maintain stable contact with the paper due to the absence of wrist control. Disturbances from the paper's reactive force often cause the pen tip to lose contact entirely. **Compliant SP** benefits compliant control to improve contact stability, but still lacks robustness under disturbance due to missing external feedback. In contrast, the proposed full framework achieves human-like writing behavior. As demonstrated in Fig. 8 (a), the finger successfully tracks the desired joint trajectories, guiding the pen tip along the target pattern based on the RL policy. Simultaneously, compliant control fine-tunes the wrist to maintain appropriate contact forces, ensuring continuous and stable trajectory tracking.

2) Generalization: To further validate the generalization capability of our framework in real-world scenarios, it is tested with various writing tools and surfaces exhibiting different physical properties (e.g., radii, lengths, masses and friction), as shown in Fig. 7. The results demonstrate that the system can successfully performs writing tasks across a range of tools, surfaces, and trajectories, highlighting its adaptability to diverse real-world conditions. This adaptability underscores the robustness and versatility of the framework in handling a wide array of environments.

VI. Conclusion and Discussions

We propose a hybrid arm-hand coordination framework that integrates reinforcement learning (RL) with compliant force-position control to enable dexterous in-hand writing under extrinsic contact. By equipping each fingertip with tactile sensors, the system dynamically infers external contact forces and positionswithout requiring fixed gripper-object contact. This flexibility allows the robotic hand to handle writing tools with varying physical properties while executing complex trajectories with stability and precision. A key contribution is the successful sim-to-real transfer: through careful calibration of joint parameters and tactile sensors, along with domain randomization, the proposed framework performs reliably in real-world environments. Our experiments demonstrate robust performance in both simulated and physical settings, despite variations in tool characteristics and contact dynamics.

While the framework performs well in controlled settings, there is still room for improvement. i) Complex trajectories requiring high precision are still difficult to execute consistently due to real-world control constraints. The three-fingered hand, though offering a balance of dexterity and simplicity, faces stability challenges for certain paths. This may be improved through additional fingers, better placement, or adaptive grasping. ii) Writing tools with rigid and slippery tips—like pencils—often yield shallow or inconsistent marks, as high stiffness and low friction with the surface make it harder to maintain stable contact. iii) Our approach assumes point contact, whereas real-world interactions often involve line or surface contacts, requiring broader contact modeling. Looking forward, extending the system to accommodate inclined or uneven surfaces will further enhance its adaptability.

Overall, this work lays the foundation for a broad range of robotic manipulation tasks in unstructured and dynamic environments, such as tool usage and assembly, advancing the capabilities of human-like dexterity in robotics.

References

[1] T. Chen, M. Tippur, S. Wu, V. Kumar, E. Adelson, and P. Agrawal, "Visual dexterity: In-hand reorientation of novel and complex object shapes," *Science Robotics*, vol. 8, no. 84, p. eadc9244, 2023.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

[2] H. Qi, A. Kumar, R. Calandra, Y. Ma, and J. Malik, "In-hand object rotation via rapid motor adaptation," in *Conf. Robot. Learn.*, pp. 1722–1732, PMLR, 2023.

[3] J. Wang, Y. Yuan, H. Che, H. Qi, Y. Ma, J. Malik, and X. Wang, "Lessons from learning to spin" pens"," *arXiv preprint arXiv:2407.18902*, 2024.

[4] S. Suresh, H. Qi, T. Wu, T. Fan, L. Pineda, M. Lambeta, J. Malik, M. Kalakrishnan, R. Calandra, *et al.*, "Neuralfeels with neural fields: Visuotactile perception for in-hand manipulation," *Science Robotics*, vol. 9, no. 96, p. eadl0628, 2024.

[5] N. C. Dafle, A. Rodriguez, R. Paolini, B. Tang, S. S. Srinivasa, M. Erdmann, M. T. Mason, I. Lundberg, H. Staab, and T. Fuhlbrigge, "Extrinsic dexterity: In-hand manipulation with external forces," in *IEEE Int. Conf. Robot. Autom.*, pp. 1578–1585, IEEE, 2014.

[6] Y. Shirai, D. K. Jha, A. U. Raghunathan, and D. Romeres, "Robust pivoting: Exploiting frictional stability using bilevel optimization," in *IEEE Int. Conf. Robot. Autom.*, pp. 992–998, IEEE, 2022.

[7] S. Kim, D. K. Jha, D. Romeres, P. Patre, and A. Rodriguez, "Simultaneous tactile estimation and control of extrinsic contact," in *IEEE Int. Conf. Robot. Autom.*, pp. 12563–12569, IEEE, 2023.

[8] M. Oller, D. Berenson, and N. Fazeli, "Tactile-driven nonprehensile object manipulation via extrinsic contact mode control," *arXiv preprint arXiv:2405.18214*, 2024.

[9] A. M. Okamura, N. Smaby, and M. R. Cutkosky, "An overview of dexterous manipulation," in *IEEE Int. Conf. Robot. Autom.*, vol. 1, pp. 255–262, IEEE, 2000.

[10] A. I. Weinberg, A. Shirizly, O. Azulay, and A. Sintov, "Survey of learning-based approaches for robotic in-hand manipulation," *Front. Robot. AI*, vol. 11, p. 1455431, 2024.

[11] L. Han and J. C. Trinkle, "Dextrous manipulation by rolling and finger gaiting," in *IEEE Int. Conf. Robot. Autom.*, vol. 1, pp. 730–735, IEEE, 1998.

[12] R. Platt, A. H. Fagg, and R. A. Grupen, "Manipulation gaits: Sequences of grasp control tasks," in *IEEE Int. Conf. Robot. Autom.*, vol. 1, pp. 801–806, IEEE, 2004.

[13] A. S. Morgan, K. Hang, B. Wen, K. Bekris, and A. M. Dollar, "Complex in-hand manipulation via compliance-enabled finger gaiting and multi-modal planning," *IEEE Robot. Autom. Lett.*, vol. 7, no. 2, pp. 4821–4828, 2022.

[14] F. Khadivar and A. Billard, "Adaptive fingers coordination for robust grasp and in-hand manipulation under disturbances and unknown dynamics," *IEEE Trans. Robot.*, vol. 39, no. 5, pp. 3350–3367, 2023.

[15] M. Yu, Y. Jiang, C. Chen, Y. Jia, and X. Li, "Robotic in-hand manipulation for large-range precise object movement: The rgmc champion solution," *arXiv preprint arXiv:2502.07472*, 2025.

[16] A. Nagabandi, K. Konolige, S. Levine, and V. Kumar, "Deep dynamics models for learning dexterous manipulation," in *Conf. Robot. Learn.*, pp. 1101–1112, PMLR, 2020.

[17] T. Pang, H. T. Suh, L. Yang, and R. Tedrake, "Global planning for contact-rich manipulation via local smoothing of quasi-dynamic contact models," *IEEE Trans. Robot.*, 2023.

[18] B. Huang, Y. Chen, T. Wang, Y. Qin, Y. Yang, N. Atanasov, and X. Wang, "Dynamic handover: Throw and catch with bimanual hands," *arXiv preprint arXiv:2309.05655*, 2023.

[19] A. Handa, A. Allshire, V. Makoviychuk, A. Petrenko, R. Singh, J. Liu, D. Makoviichuk, K. Van Wyk, A. Zhurkevich, B. Sundaralingam, *et al.*, "Dextreme: Transfer of agile in-hand manipulation from simulation to reality," in *IEEE Int. Conf. Robot. Autom.*, pp. 5977–5984, IEEE, 2023.

[20] G. Khandate, M. Haas-Heger, and M. Ciocarlie, "On the feasibility of learning finger-gaiting in-hand manipulation with intrinsic sensing," in *IEEE Int. Conf. Robot. Autom.*, pp. 2752–2758, IEEE, 2022.

[21] M. Yang, C. Lu, A. Church, Y. Lin, C. Ford, H. Li, E. Psomopoulou, D. A. Barton, and N. F. Lepora, "Anyrotate: Gravity-invariant in-hand object rotation with sim-to-real touch," *arXiv preprint arXiv:2405.07391*, 2024.

[22] W. Hu, B. Huang, W. W. Lee, S. Yang, Y. Zheng, and Z. Li, "Dexterous in-hand manipulation of slender cylindrical objects through deep reinforcement learning with tactile sensing," *Robotics and Autonomous Systems*, p. 104904, 2024.

[23] D. Ma, S. Dong, and A. Rodriguez, "Extrinsic contact sensing with relative-motion tracking from distributed tactile measurements," in *IEEE Int. Conf. Robot. Autom.*, pp. 11262–11268, IEEE, 2021.

[24] C. Higuera, S. Dong, B. Boots, and M. Mukadam, "Neural contact fields: Tracking extrinsic contact with tactile sensing," in *IEEE Int. Conf. Robot. Autom.*, pp. 12576–12582, IEEE, 2023.

[25] L. Kim, Y. Li, M. Posa, and D. Jayaraman, "Im2contact: Vision-based contact localization without touch or force sensing," in *Conf. Robot. Learn.*, pp. 1533–1546, PMLR, 2023.

[26] M. Van der Merwe, Y. Wi, D. Berenson, and N. Fazeli, "Integrated object deformation and contact patch estimation from visuo-tactile feedback," *arXiv preprint arXiv:2305.14470*, 2023.

[27] C. Higuera, J. Ortiz, H. Qi, L. Pineda, B. Boots, and M. Mukadam, "Perceiving extrinsic contacts from touch improves learning insertion policies," *arXiv preprint arXiv:2309.16652*, 2023.

[28] N. Doshi, O. Taylor, and A. Rodriguez, "Manipulation of unknown objects via contact configuration regulation," in *IEEE Int. Conf. Robot. Autom.*, pp. 2693–2699, IEEE, 2022.

[29] O. Taylor, N. Doshi, and A. Rodriguez, "Object manipulation through contact configuration regulation: multiple and intermittent contacts," in *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, pp. 8735–8743, IEEE, 2023.

[30] Y. Shirai, D. K. Jha, A. U. Raghunathan, and D. Hong, "Tactile tool manipulation," in *IEEE Int. Conf. Robot. Autom.*, pp. 12597–12603, IEEE, 2023.

[31] D. Prattichizzo, L. Meli, and M. Malvezzi, "Digital handwriting with a finger or a stylus: a biomechanical comparison," *IEEE Trans. Haptics*, vol. 8, no. 4, pp. 356–370, 2015.

[32] V. Petrone, E. Ferrentino, and P. Chiacchio, "Time-optimal trajectory planning with interaction with the environment," *IEEE Robot. Autom. Lett.*, vol. 7, no. 4, pp. 10399–10405, 2022.

[33] C. C. Beltran-Hernandez, D. Petit, I. G. Ramirez-Alpizar, and K. Harada, "Variable compliance control for robotic peg-in-hole assembly: A deep-reinforcement-learning approach," *Applied Sciences*, vol. 10, no. 19, p. 6923, 2020.

[34] J. Zhou, X. Chen, U. Chang, J.-T. Lu, C. C. Y. Leung, Y. Chen, Y. Hu, and Z. Wang, "A soft-robotic approach to anthropomorphic robotic hand dexterity," *IEEE Access*, vol. 7, pp. 101483–101495, 2019.

[35] J. T. Grace, P. Chanrungmaneekul, K. Hang, and A. M. Dollar, "Direct self-identification of inverse jacobians for dexterous manipulation through particle filtering," in *IEEE Int. Conf. Robot. Autom.*, pp. 13862–13868, IEEE, 2024.

[36] P. Lu, J. Liang, B. Huang, S. Yang, and W. W. Lee, "Thermoformed electronic skins for conformal tactile sensor arrays," in *IEEE Int. Conf. Robot. Autom.*, pp. 13898–13903, IEEE, 2024.

[37] M. Kaess, A. Ranganathan, and F. Dellaert, "isam: Incremental smoothing and mapping," *IEEE Trans. Robotics*, vol. 24, no. 6, pp. 1365–1378, 2008.

[38] N. Hansen and A. Ostermeier, "Completely derandomized self-adaptation in evolution strategies," *Evolutionary computation*, vol. 9, no. 2, pp. 159–195, 2001.