

Contact-Grounded Policy: Dexterous Visuotactile Policy with Generative Contact Grounding

Zhengtong Xu¹, Yeping Wang³, Ben Abbatematteo², Jom Preechayasomboon²
 Sonny Chan², Nick Colonnese², Amirhossein H. Memar²

¹Purdue University ²Meta Reality Labs Research ³University of Wisconsin–Madison

<https://contact-grounded-policy.github.io/>




Fig. 1: Contact-Grounded Policy (CGP) enables fine-grained, contact-rich dexterous manipulation by grounding multi-point contacts through predicting the actual robot state and tactile feedback, and by using a learned contact-consistency mapping to convert these predictions into controller-executable targets for a compliance controller. CGP supports both dense tactile arrays and vision-based tactile sensors mounted on the hand. We collect demonstrations via teleoperation in both simulation and on a physical robot, and evaluate CGP across multiple dexterous manipulation tasks.

Abstract—Contact-rich dexterous manipulation with multi-finger hands remains an open challenge in robotics because task success depends on multi-point contacts that continuously evolve and are highly sensitive to object geometry, frictional transitions, and slip. Recently, tactile-informed manipulation policies have shown promise. However, most use tactile signals as additional observations rather than modeling contact state or how their action outputs interact with low-level controller dynamics. We present Contact-Grounded Policy (CGP), a visuotactile policy that grounds multi-point contacts by predicting coupled trajectories of actual robot state and tactile feedback, and using a learned contact-consistency mapping to convert these predictions into executable target robot states for a compliance controller. CGP consists of two components: (i) a conditional diffusion model that forecasts future robot state and tactile feedback in a compressed latent space, and (ii) a learned contact-consistency mapping that converts the predicted robot state-tactile pair into executable targets for a compliance controller, enabling it to realize the intended contacts. We evaluate CGP using a physical four-finger Allegro V5 hand with Digit360 fingertip tactile sensors, and a simulated five-finger Tesollo DG-5F hand with dense whole-hand tactile arrays. Across a range of dexterous tasks

including in-hand manipulation, delicate grasping, and tool use, CGP outperforms visuomotor and visuotactile diffusion-policy baselines.

I. INTRODUCTION

Dexterous manipulation with multi-finger robotic hands remains one of the most challenging frontiers in robotics. Unlike rigid end-effectors or simple grippers, dexterous manipulation with hands requires the robot to continuously regulate rich, high-dimensional contact interactions between multiple fingers and the object. These interactions are distributed across multiple contact points and evolve rapidly as the hand and object move. They are highly nonlinear and partially observable, and their outcomes are sensitive to contact geometry, frictional state transitions, and slip.

Existing learning approaches to dexterous manipulation generally fall into three primary paradigms. First, grasp-centric pipelines focus on generating secure grasp configurations [6, 36, 32, 40]. These grasp-centric pipelines are effective

for rigid pick-and-place tasks. However, once a grasp is established, they typically constrain subsequent finger motion, limiting their applicability to contact-rich behaviors that require continuous reconfiguration and active contact regulation, such as in-hand manipulation and tool use. Second, reinforcement learning (RL) can discover complex contact strategies [18, 22, 38, 15, 35]. However, RL policies often face major hurdles in sim-to-real transfer, particularly with visual and tactile observations, and frequently require extensive reward engineering to achieve delicate manipulation on new tasks.

Imitation learning from human demonstrations is well-suited to contact-rich dexterous manipulation because humans naturally regulate distributed, time-varying contacts through coordinated finger motion and compliant interaction. It avoids complex reward design, scales with additional demonstrations, and can generalize across objects and tasks when the dataset covers diverse contact patterns. Motivated by these advantages, recent imitation-learning methods have shown impressive performance by directly leveraging human demonstrations [30, 9, 34]. However, many proposed visuomotor policies still struggle with complex, contact-rich dexterous manipulation [1]. A core limitation is that they primarily predict kinematic trajectories without explicit contact semantics, making it difficult to represent and control evolving multi-point contacts.

Addressing this limitation requires policies to go beyond kinematics-only prediction. First, they must leverage multi-modal perception, reasoning over tactile feedback alongside vision, to reduce ambiguity under partial observability. Second, policies must go beyond using tactile signals as additional observations and instead model contact state and how action outputs interact with low-level controller dynamics. Concretely, the policy should produce controller-compatible targets that reflect contact evolution and can be faithfully executed by the compliance controller during deployment. Otherwise, learned outputs can become physically inconsistent with the environment, leading to slip, overly stiff interactions, and ultimately unreliable execution. Bridging high-level task intent and low-level contact-rich control therefore requires deploying contact predictions through the compliance controller.

In this paper, we introduce **Contact-Grounded Policy (CGP)**, a supervised policy learning framework for dexterous visuotactile manipulation, as shown in Fig. 1. CGP casts dexterous manipulation as a contact-grounding problem: rather than using tactile signals only as additional observations, CGP grounds evolving multi-point contacts by predicting coupled trajectories of the actual robot state and tactile feedback, and by using a learned contact-consistency mapping to convert these predictions into controller-executable target robot states that account for the dynamics of the compliance controller. This representation produces concrete, contact-consistent targets that enable real-time execution to realize the intended contact evolution.

Our contributions are as follows:

1. Contact-Grounded Policy Framework. We introduce Contact-Grounded Policy (CGP), a visuotactile policy that grounds evolving multi-point contacts for dexterous manip-

ulation. CGP does so by predicting coupled trajectories of the actual robot kinematic state and tactile feedback, and then translating these predictions into executable control targets. Specifically, a learned contact-consistency mapping converts each predicted state-tactile pair into a target robot state that accounts for the compliance controller, enabling real-time realization of the intended contact evolution. Across a range of dexterous tasks including in-hand manipulation, delicate grasping, and tool use, CGP outperforms visuomotor and visuotactile diffusion-policy baselines.

2. Efficient Tactile Prediction for Contact Grounding. We utilize a latent tactile prediction model integrated into CGP that delivers high-fidelity tactile forecasts. For efficiency, we compress tactile observations with a KL-regularized variational autoencoder (VAE) and perform prediction in the resulting compact latent space. This design captures diverse contacts while remaining lightweight at runtime, and we demonstrate its effectiveness across both tactile arrays and vision-based tactile sensors. Ablations show that the KL-regularized latent space stabilizes generation and improves downstream policy performance.

II. RELATED WORK

Tactile-informed policy learning has become a prominent paradigm for contact-rich manipulation [34, 33, 19, 13, 41, 10, 37]. A growing body of work goes beyond direct action regression by introducing tactile/force prediction objectives to shape representations, providing richer supervision when action labels are low-dimensional or ambiguous [9, 14, 20]. However, most approaches treat tactile signals as additional observations or auxiliary prediction targets, rather than modeling contact state or how action outputs interact with low-level controller dynamics.

Predicting dense tactile signals alone does not guarantee that the intended contacts will be physically realized during rollout [9, 14, 20]. In many approaches, tactile prediction is only loosely coupled to the control stack; as a result, predicted contact patterns may not be reproduced by the commands executed by the low-level controller. This can yield disconnected contact awareness: models anticipate tactile patterns yet fail to ground them into controller references.

Related work moves toward more executable contact representations by predicting control-relevant interaction targets, such as adaptive compliance references, rather than treating tactile or force prediction as an isolated modeling objective [12, 31, 5]. However, many such formulations rely on simplified contact abstractions (e.g., a single arm end-effector force vector or a single force vector per gripper finger), which limits their extensibility to distributed contact points across a multi-finger hand. Similarly, [39, 3] predict sparse fingertip forces executed by a low-level controller, whereas real dexterous interaction involves dense, multi-patch contacts across the hand with high-DoF finger motion. Table I summarizes representative policy paradigms with executable contact modeling.

Method Category	Executable Contact Modeling	Multi-Finger Hand	Extensible to Distributed Contacts
Adaptive Compliance Policies [12, 31, 5]	✓	✗	✗
Sparse Fingertip Force Policies [39, 3]	✓	✓	✗
Contact-Grounded Policy (Ours)	✓	✓	✓

TABLE I: Comparison of representative policy paradigms with executable contact modeling.




Fig. 2: Schematic of contact grounding using a 3-DoF revolute finger, illustrating the actual robot state, target robot state, and the resulting contact patches. We assume that each joint is controlled by a low-level proportional-derivative (PD) controller, which can be viewed as a virtual spring-damper that maps the tracking error between target and actual joint angles to motor torques, enabling compliant motion.

III. METHOD

This section presents CGP, a supervised policy learning framework for dexterous visuotactile manipulation that casts dexterous manipulation as contact grounding. CGP grounds high-level task intent into contact-consistent control targets by predicting (i) future actual robot states, (ii) the corresponding expected tactile feedback, and (iii) executable target robot states for the compliance controller. Coupled prediction produces a structured description of the intended contact evolution.

A. Problem Setup

We consider an arm-hand manipulation system equipped with tactile sensing on the hand (e.g., fingertip or full-hand coverage). At each time step t , the robot receives multi-modal observations $\mathbf{o}_t = \{\mathbf{i}_t, \mathbf{u}_t, \mathbf{x}_t\}$. Here, the vision observation \mathbf{i}_t stacks N_i RGB views (e.g., agent-view and wrist-view) and is represented as $\mathbf{i}_t \in \mathbb{R}^{N_i \times H_i \times W_i \times 3}$. The tactile observation \mathbf{u}_t can take two forms: (i) tactile RGB images from N_u sensors, represented as $\mathbf{u}_t \in \mathbb{R}^{N_u \times H_u \times W_u \times 3}$, or (ii) tactile arrays with N_{tac} tactile units and per-unit feature dimension d_{tac} , represented as $\mathbf{u}_t \in \mathbb{R}^{N_{\text{tac}} \times d_{\text{tac}}}$. The robot state \mathbf{x}_t consists of the end-effector pose and hand joint angles, where the end-effector translation is in \mathbb{R}^3 and the rotation is parameterized by a unit quaternion in \mathbb{R}^4 ; overall, we represent $\mathbf{x}_t \in \mathbb{R}^{(3+4+n_h)}$, where n_h is the number of hand joints.

The policy interfaces with the robot through a low-level tracking controller by outputting a target robot state \mathbf{a}_t , which serves as the controller reference. For example, \mathbf{a}_t specifies an operational space impedance controller setpoint for the arm and joint space targets for the hand proportional-derivative (PD) controller. We parameterize the target translation in \mathbb{R}^3 and represent the target rotation with the 6D continuous rotation representation in \mathbb{R}^6 (rot6D), which facilitates stable

action regression [42, 4]; thus, $\mathbf{a}_t \in \mathbb{R}^{(3+6+n_h)}$. The low-level controller converts the target state \mathbf{a}_t into commanded joint torques.

B. Contact Grounding

Contact-rich dexterous manipulation involves distributed, time-varying contacts over a high-DoF hand, so modeling contact via designed parameterizations (e.g., contact locations, modes, or fixed contact points) often fails to capture the full variability and combinatorial complexity of evolving multi-point contact. CGP instead represents contact through coupled state and tactile signals that are directly measurable and can be mapped to controller-executable targets.

A key observation is that, under a specific tactile sensor and compliance controller setup, contact can be represented by the triplet $(\mathbf{x}_t, \mathbf{u}_t, \mathbf{a}_t)$ without explicitly specifying contact locations or modes, where \mathbf{x}_t is the actual robot state, \mathbf{u}_t is tactile feedback, and \mathbf{a}_t is the compliance controller reference (i.e., the target robot state). Under closed-loop tracking, interaction manifests as a discrepancy between \mathbf{a}_t and \mathbf{x}_t , while the resulting contact outcome is reflected in \mathbf{u}_t (Fig. 2). This representation naturally adapts to tactile coverage: full-hand sensing captures distributed, multi-patch contacts, whereas dense fingertip tactile sensing (e.g., vision-based tactile images) can be sufficient when interaction is primarily localized to the fingertips.

Motivated by this coupling, we learn a controller-specific contact-consistency mapping: given $(\mathbf{x}_t, \mathbf{u}_t)$, it predicts a target state \mathbf{a}_t whose reference would reproduce the observed interaction under the compliance controller

$$\mathbf{a}_t = \mathcal{M}_\phi(\mathbf{x}_t, \mathbf{u}_t). \quad (1)$$

Eq. (1) avoids explicitly modeling contact locations/modes or system dynamics; instead, it learns a compact, implicit, and setup-dependent mapping. Because this mapping is tied to a specific embodiment, sensing configuration, and compliance controller dynamics, it can be learned in a purely data-driven manner, while remaining flexible to distributed, evolving multi-point contacts that are difficult to capture with hand-designed contact parameterizations or explicit dynamics models.

C. Policy Pipeline Overview

1) *Architecture*: Fig. 3 summarizes the CGP pipeline. We denote a short history of observations by $\mathbf{O}_t = \{\mathbf{o}_{t-T_o+1}, \dots, \mathbf{o}_t\}$. Here, T_o denotes the observation horizon. CGP consists of two coupled components that implement contact grounding in a structured manner. The first component is a




Fig. 3: Overview of Contact-Grounded Policy (CGP). CGP grounds multi-point contacts by predicting coupled trajectories of actual robot state and tactile feedback, and using a learned contact-consistency mapping to convert these predictions into executable target robot states for a compliance controller.

contact-consistency mapping \mathcal{M}_ϕ , which infers the controller reference consistent with the current interaction, as shown in Eq. (1). The second component is a conditional trajectory generator π_θ that predicts how the interaction should evolve over a future horizon. Conditioned on \mathbf{O}_t , it samples coupled trajectories of future actual robot state and tactile feedback,

$$(\hat{\mathbf{X}}_t, \hat{\mathbf{U}}_t) \sim \pi_\theta(\cdot | \mathbf{O}_t),$$

where $\hat{\mathbf{X}}_t = \{\hat{\mathbf{x}}_{t+1}, \dots, \hat{\mathbf{x}}_{t+T}\}$ and $\hat{\mathbf{U}}_t = \{\hat{\mathbf{u}}_{t+1}, \dots, \hat{\mathbf{u}}_{t+T}\}$. Here, T denotes the prediction horizon. In our implementation, π_θ follows a diffusion-policy-style design and is parameterized by a U-Net denoiser [4].

This factorization is intentional. Rather than directly mapping observations to controller references, CGP first predicts interaction targets in future actual robot state and tactile feedback that describe the desired contact evolution, and then converts them into executable target robot states via \mathcal{M}_ϕ . This separation allows the trajectory generator to model contact evolution as paired state-tactile trajectories, while the mapping realizes the intended contacts by producing low-level controller references.

2) *Residual Mapping*: We formulate the contact-consistency mapping \mathcal{M}_ϕ in residual form. Rather than regressing the target state directly, it outputs an offset from the current actual state, which anchors learning, improves conditioning, and yields more robust targets under the compliance controller.

3) *Inference and Execution*: At test time, CGP samples future actual robot state and tactile feedback trajectories from π_θ and converts them into target robot states step-wise,

$$\hat{\mathbf{a}}_{t+k} = \mathcal{M}_\phi(\hat{\mathbf{x}}_{t+k}, \hat{\mathbf{u}}_{t+k}), \quad k = 1, \dots, T.$$

The compliance controller tracks the next target robot state, and the policy replans in a receding-horizon manner.

D. Latent Tactile Generation

Raw tactile observations for dexterous hands can be high-dimensional, especially for vision-based tactile sensors and dense tactile arrays, making long-horizon generation computationally expensive. To enable efficient real-time generation

while preserving contact variability, CGP generates tactile trajectories in a learned latent space.

1) *Tactile Compression with VAE*: We train a variational autoencoder (VAE) [16] to compress raw tactile observations \mathbf{u}_t into a compact latent representation, $\mathbf{h}_t = E(\mathbf{u}_t)$, $\hat{\mathbf{u}}_t = G(\mathbf{h}_t)$, where E and G denote the encoder and decoder, respectively. A compact, well-behaved latent space is critical for stable and efficient diffusion-based generation; therefore, we use KL regularization to encourage learning such a compact latent representation [24]. The VAE backbone can be adapted to different tactile modalities (e.g., tactile arrays or tactile images) while sharing a common training objective and pipeline. In all cases, tactile observations are mapped to a fixed-dimensional latent $\mathbf{h}_t \in \mathbb{R}^M$.

2) *Coupled Diffusion over State and Tactile Latent*: After compression, tactile trajectories are generated in latent space together with actual robot states. We define the coupled future trajectory as $\mathbf{Y}_t = [\mathbf{x}_{t+1:t+T}, \mathbf{h}_{t+1:t+T}]$, which concatenates future actual robot states and tactile latent states over the horizon.

We model the conditional distribution of \mathbf{Y}_t given the observation history \mathbf{O}_t using a diffusion model. Following standard DDPM/DDIM training [11, 26], we sample a diffusion step $j \in \{1, \dots, J\}$ and perturb the clean trajectory \mathbf{Y}_t^0 with Gaussian noise, $\mathbf{Y}_t^j = \alpha_j \mathbf{Y}_t^0 + \sigma_j \epsilon$, $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, where $\{\alpha_j, \sigma_j\}$ are fixed noise-schedule coefficients. A U-Net denoiser π_θ is trained to predict the injected noise conditioned on \mathbf{O}_t ,

$$\mathcal{L}_{\text{diff}}(\theta) = \mathbb{E}_{(\mathbf{O}_t, \mathbf{Y}_t^0), \epsilon, j} \left[\|\epsilon - \pi_\theta(\mathbf{O}_t, \mathbf{Y}_t^j, j)\|^2 \right].$$

E. Implementation Details

This section describes key implementation choices in CGP.

1) *Tactile Encoders and Decoders*: CGP employs tactile encoders across multiple components, including diffusion conditioning, contact-consistency mapping, and tactile compression, while decoders are used only for tactile compression. Since these components are trained with different objectives and at different stages, encoder weights are not shared, though all encoders follow the same architectural template.

We use different tactile modalities in simulation and on real hardware: dense tactile arrays in simulation and vision-based

tactile sensors (RGB tactile images) on the physical system. We validate the proposed pipeline on both sensor types.

In simulation, tactile sensing is provided as a dense tactile array $\mathbf{u}_t \in \mathbb{R}^{N_{\text{tac}} \times d_{\text{tac}}}$ (e.g., $d_{\text{tac}}=3$ for 3D force vectors). We use a 1D ResNet-style architecture that applies convolutions along the tactile-unit dimension N_{tac} , treating d_{tac} as channels. For decoding, the latent code is projected to a compact 1D feature map and progressively upsampled with transposed convolutions and residual refinement, followed by a final convolution to reconstruct tactile features.

On real hardware, tactile sensing is provided by N_u vision-based tactile sensors, each producing an RGB image. We adopt a per-sensor design with shared 2D ResNet-style encoder-decoder weights, processing each tactile image independently to produce consistent fingertip representations. The encoder consists of a convolutional stem followed by strided residual blocks and global pooling, while the decoder mirrors this structure with progressive upsampling and residual refinement. A final convolutional layer reconstructs the RGB tactile image.

2) *Visual Encoder and Diffusion*: For visual encoding and diffusion, we follow the architecture of diffusion policy [4], using a U-Net-based conditional diffusion model and DDIM sampling [26] for efficient inference.

In simulation, visual features, tactile features, and the low-dimensional robot state are concatenated and injected into the diffusion U-Net via FiLM conditioning [21]. On real hardware, each tactile image is encoded using the shared tactile encoder (Sec. III-E1), and the resulting per-sensor features are aggregated via cross-sensor self-attention. The aggregated tactile feature is then concatenated with the visual feature and robot state, and applied through FiLM conditioning throughout the diffusion backbone.

3) *Contact-Consistency Mapping*: The contact-consistency mapping is a lightweight network. For tactile arrays, we prioritize tactile fidelity by decoding the diffusion-predicted tactile latent to raw tactile, re-encoding it, and concatenating the resulting feature with the actual robot state to predict the target state using an MLP. This design preserves detailed contact information for target prediction.

For vision-based tactile sensors, we use a more lightweight design for real-time deployment. Reconstructing tactile images over the prediction horizon is substantially more expensive, so we instead concatenate the diffusion-predicted tactile latent states with the actual robot states and feed them directly to an MLP head. This greatly reduces inference cost while retaining sufficient tactile information for contact-consistent execution.

We use different designs in simulation and on real hardware because the tactile modalities and runtime constraints differ: array decoding is lightweight and preserves high-frequency contact details, whereas image reconstruction is compute-heavy and can become a bottleneck for real-time inference.

IV. SYSTEM OVERVIEW

A. Simulation Setup

For the simulation environment, we integrate a real-time finite element solver with Unreal Engine via a native plugin.




Fig. 4: Teleoperation pipeline. We use a Meta Quest 3 headset for VR-based hand tracking in simulation and an OptiTrack system for mocap-based hand tracking on the real robot. Despite different tracking front-ends, both settings share the same retargeting and controller-stack architecture.

The solver uses reduced-order models with data-driven hyper-reduction to simulate contact among rigid and compliant objects at interactive rates, providing rich per-contact state for dexterous manipulation [27, 2, 44, 25]. Unreal Engine handles rendering and I/O; the physics step is advanced natively each frame. The simulated robot system consists of a UR5 arm under an operational space impedance controller [28, 43] and a Tesollo DG-5F five-finger 20-DoF hand under a joint space PD controller. To capture rich multi-contact interactions, we instrument the inner surfaces of all fingers and the palm with a dense tactile array of 748 sensing points, each reporting a 3D force vector, yielding tactile observations in $\mathbb{R}^{748 \times 3}$. Visual observations are provided by two RGB cameras, including an agent-view camera and a wrist-mounted camera.

Demonstrations are collected via VR teleoperation using a Meta Quest 3 headset (Fig. 4). The tracked hand pose is mapped to the robot end-effector pose to define the arm target [8]. For the fingers, we retarget the operator’s tracked fingertip positions to the DG-5F by scaling them to its kinematic range, and use RelaxedIK [23, 29] to solve inverse kinematics and produce finger joint targets, as shown in Fig. 4.

B. Real-Robot Setup

On real hardware, we use a Franka Panda arm under an operational space impedance controller [28, 43] and an Allegro V5 four-finger 16-DoF hand under a joint space PD controller. Tactile sensing is provided by four Digit360 sensors [17] mounted on the fingertips. We use tactile images as the tactile modality. Visual sensing follows the same two-view design as in simulation.

Real-robot demonstrations are collected using an OptiTrack motion-capture system and an instrumented glove for hand tracking [7], as shown in Fig. 4. We map the tracked human hand pose to the robot end-effector pose to define the arm target. For hand teleoperation, we perform fingertip-based retargeting: we map glove-estimated fingertip positions into the Allegro hand workspace, ignore the human pinky, and solve inverse kinematics with RelaxedIK [23, 29] to obtain finger joint targets, as shown in Fig. 4.

Method	Sim			Real	
	In-Hand Box Flipping 60 demos	Fragile Egg Grasping 100 demos	Dish Wiping 100 demos	Jar Opening 45 demos	Real In-Hand Box Flipping 90 demos
Contact-Grounded Policy	66.0%	74.8%	58.4%	93.3%	80.0%
Visuotactile DP	58.0%	70.0%	43.6%	66.7%	60.0%
Visuomotor DP	53.2%	53.2%	42.4%	73.3%	60.0%

TABLE II: Success rate over 5 challenging contact-rich, dexterous manipulation tasks. DP is short for diffusion policy. For the simulation tasks (left three), we report the mean success rate over the last 5 checkpoints, evaluated with 250 rollouts in total. For the real-world tasks, we report the success rate over 15 consecutive rollouts.




Fig. 5: Snapshots of Contact-Grounded Policy (CGP) rollouts on three simulated tasks, showing time-aligned predicted and observed tactile feedback. At each inference step, the diffusion model predicts the next 16 steps of tactile feedback and actual states, which are mapped to target states and executed for 8 steps before the next inference. Predicted tactile is time-aligned with subsequent observations after execution, and the close match indicates that CGP executes contact-grounded targets and realizes the predicted contact evolution. Full rollout videos are provided in the supplementary material.

V. EXPERIMENTS

We evaluate CGP through three complementary sets of experiments, each targeting a distinct component of the framework. First, we evaluate full policy rollouts on both simulated and real-world contact-rich dexterous manipulation tasks to




Fig. 6: Hand configuration predictions by the contact-consistency mapping for unseen grasps (Section V-B). This figure provides high-level evidence that contact can be consistently represented through this mapping in a way that generalizes across diverse contact configurations.

assess end-to-end performance. Second, we study hand configuration prediction to directly validate the contact-consistency mapping in isolation. In this experiment, we highlight the necessity of several key design choices in the mapping, and provide high-level evidence that contact can be consistently represented through this mapping in a way that generalizes across diverse contact configurations. Finally, we analyze tactile reconstruction and compression to validate the design of the latent tactile representation used for real-time prediction, and quantify how alternative design choices impact policy performance.

A. Performance on Contact-Rich Dexterous Manipulation

We evaluate CGP in full closed-loop execution on three simulated tasks and two real-robot tasks involving contact-rich,

Input Modality	Tactile Encoder	Abs. Mode ↓	Residual Mode ↓
State + Tactile	ResNet1D	8.80 ± 0.24	5.94 ± 0.20
	MLP	12.50 ± 0.32	8.33 ± 0.32
	Transformer	14.39 ± 0.38	9.58 ± 0.48
State Only	-	16.05 ± 0.39	10.64 ± 0.38
Tactile Only	ResNet1D	35.93 ± 0.89	12.15 ± 0.20
	MLP	36.86 ± 0.25	12.72 ± 0.25
	Transformer	43.11 ± 0.91	14.62 ± 0.43

TABLE III: Hand configuration prediction results (Section V-B). MAE denotes mean absolute error. Values are reported in joint-angle space as MAE ($\times 10^{-3}$ rad), shown as mean \pm std over 3 seeds. Model initializations and the train/validation splits are determined by the random seed. “Abs.” stands for absolute.




Fig. 7: Ablation results of KL regularization for tactile compression. MAE stands for mean absolute error. “Recon.” stands for reconstruction.

dexterous manipulation. As baselines, we compare against a standard visuomotor diffusion policy that uses only visual observations, and a visuotactile diffusion policy that follows the same formulation but additionally conditions on tactile observations. To ensure a fair comparison, we implement both visuomotor and visuotactile diffusion-policy baselines with identical visual and tactile encoders and the same U-Net backbone as CGP. All tasks require coordinated finger motion and force modulation while maintaining stable, continuously evolving contacts under the compliance controller. Task details are provided in the appendix and the supplementary video.

Table II reports success rates across all tasks. CGP consistently outperforms both visuomotor and visuotactile diffusion-policy baselines. The improvements are pronounced in tasks with sustained or delicate contact, such as dish wiping, in-hand box flipping, and jar opening, where contact evolution plays a central role.

Fig. 5 visualizes snapshots of rollouts and provides direct evidence of effective contact grounding. At each replanning step, CGP predicts a short-horizon tactile trajectory, maps it into contact-consistent targets, and executes these targets under compliant tracking. To verify that predicted contacts are actually realized during execution, we time-align tactile frames predicted at earlier replanning steps with tactile feedback observed at the corresponding future time steps. The aligned predicted and observed tactile signals exhibit alignment across




Fig. 8: Inference time comparison. Average time over 50 runs for visuomotor diffusion-policy baseline, visuotactile diffusion-policy baseline, and Contact-Grounded Policy (CGP) using 8-step DDIM denoising. To guarantee fairness, shared network modules utilize the same architecture and size across all methods. For architectural details, please refer to Section III-E.

contact regions and force patterns, despite the stochasticity of multi-contact interactions. This close correspondence indicates that CGP does not merely forecast plausible tactile outcomes, but generates interaction targets that the robot can reliably execute to reproduce the predicted contact evolution. Overall, the visualization confirms that tactile prediction in CGP is tightly coupled to execution through contact-consistent targets, rather than being treated as an auxiliary objective decoupled from control.

Moreover, Fig. 5 compares two overlays: aligning only the hand target/actual states (with a fixed hand-skeleton base) versus aligning the end-effector (which determines the hand base). The two overlays are similar for in-hand box flipping and egg grasping, suggesting that most interactions are hand-dominated; in contrast, dish wiping shows a clear target-actual pose offset when the sponge contacts the plate (Fig. 5, third row), indicating arm-driven force application and demonstrating that our formulation generalizes across both hand and arm-driven contact interactions.

B. Hand Configuration Prediction

We design a controlled hand configuration prediction task to validate the effectiveness of the contact-consistency mapping. This experiment evaluates the mapping from the actual robot state and tactile feedback to the target robot state, without involving downstream policy learning, and serves two purposes: (i) to justify the necessity of key design choices in the mapping, and (ii) to provide high-level evidence that contact can be consistently represented through this mapping in a way that generalizes across diverse contact configurations.

We collect 150 teleoperated grasping episodes with 4114 frames across 11 objects in simulation. Object details are provided in the appendix. Each episode corresponds to a

Method	Box			Egg			Dish		
	MAE↓	Active MAE↓	KL↓	MAE↓	Active MAE↓	KL↓	MAE↓	Active MAE↓	KL↓
MLP (w/ KL)	1.37	12.24	0.21	0.77	6.70	0.15	1.86	8.14	0.15
MLP (w/o KL)	1.12	9.93	0.59	0.54	4.51	0.40	1.40	6.20	0.42
ResNet1D (w/ KL)	1.26	12.07	0.12	0.69	5.95	0.22	1.54	6.80	0.24
ResNet1D (w/o KL)	0.97	9.91	0.73	0.45	3.92	0.43	1.02	4.49	0.45
Transformer (w/ KL)	1.66	14.88	0.19	1.58	10.58	0.28	3.19	13.57	0.21
Transformer (w/o KL)	1.69	14.60	0.97	1.18	9.08	0.61	3.28	13.57	0.53

TABLE IV: Tactile reconstruction and compression results on the validation set. MAE stands for mean absolute error. MAE averages the per-unit force error over all 748 tactile units, where each unit’s force magnitude is computed as the resultant of the 3-axis forces. Active MAE computes the same error but only over tactile units whose ground-truth force is non-zero (i.e., active contact units). MAE and Active MAE are reported in 10^{-2} N.

unique grasp configuration with diverse contact patterns. To avoid data leakage across correlated frames, we split data at the episode level with an aggressive 1:1 train-validation split.

Table III shows that accurate prediction requires both actual robot state and tactile feedback. Removing either modality substantially increases error, supporting the contact-grounding hypothesis. Tactile-only inputs cannot resolve global hand configuration, while state-only inputs miss contact-induced variations. Together, they capture contact-consistent structure that transfers across objects and contact patterns.

We also ablate mapping designs. A ResNet-style tactile encoder outperforms MLP and Transformer variants, and predicting residual targets reduces error compared to absolute prediction, indicating that contact grounding is best modeled as contact-conditioned corrections around the actual state. These results motivate the final design of the contact-consistency mapping used in CGP. Qualitative results in Fig. 6 further illustrate that, under the ResNet encoder and residual formulation, the model can accurately predict complex and previously unseen target robot state when informative tactile feedback is available. This indicates that the learned mapping generalizes beyond memorized grasp poses and captures transferable, contact-consistent structure across diverse interactions.

C. Tactile Reconstruction, Compression, and Design Impact

We evaluate tactile reconstruction and compression to validate the latent tactile representation used for prediction and to quantify how key design choices affect downstream policy performance. Using the same demonstration dataset as policy training, we split episodes into train and validation with a 1:1 ratio. We size-match encoder and decoder capacity across methods.

Table IV shows that ResNet-based encoders achieve strong reconstruction quality across tasks. While removing KL regularization slightly lowers reconstruction error, it produces a poorly structured latent space with much higher KL divergence. This effect carries over to policy learning. Fig. 7 shows that removing KL consistently degrades rollout performance despite better reconstruction metrics. KL regularization yields a more compact and better-conditioned latent space, which is also reflected by the reduced KL loss in Table IV. This improves the stability of diffusion-based prediction [24] and

the fidelity of long-horizon tactile forecasts. Overall, these results validate KL-regularized latent tactile representations in CGP for reliable tactile prediction and contact-grounded control.

Additional quantitative and qualitative reconstruction results on real vision-based tactile sensors (Digit360) are provided in the appendix, highlighting both the effectiveness of our design and its applicability across tactile sensor types.

D. Time Efficiency

We evaluate the inference efficiency of CGP and compare it with visuomotor and visuotactile diffusion-policy baselines. Fig. 8 reports the average inference time across two tactile sensing modalities, including tactile arrays and Digit360. All inference timings are measured on an NVIDIA A100 80GB GPU. Despite modeling future tactile feedback and contact-consistent targets, CGP achieves inference latency comparable to both visual and visuotactile diffusion-policy baselines.

VI. LIMITATIONS AND FUTURE WORK

A key limitation of CGP is its specificity to both sensing and control. The contact-consistency mapping relies on tactile observations and is learned under a particular compliance controller, so it does not readily transfer across sensor types or controller configurations. As a result, cross-sensor and cross-controller contact grounding remains challenging. In this work, CGP is trained from scratch for each sensor, and changes in sensor type require re-training or adaptation.

A promising direction is cross-sensor and cross-controller co-training. Beyond co-training, conditioning the contact-consistency mapping on controller parameters and robot physical parameters (e.g., impedance gains, update rate, and kinematics/dynamics) is also a promising way to further improve generalization. Together, co-training and additional conditioning would better accommodate deployment shifts and support systematic transfer across sensing and control settings, reducing the need for re-training.

Another limitation is that we do not explicitly study direct sim-to-real transfer. Instead, we evaluate CGP through separate simulated and real-world deployments to isolate the algorithmic benefits of the contact-grounded policy formulation from complex visual, tactile, and dynamics domain shifts. Extending

CGP toward sim-to-real transfer is a promising future direction, potentially through sim-and-real co-training, domain-randomized visual and tactile observations, and adaptation methods that align contact representations across simulated and real interactions.

Finally, we validate CGP under a single-task training and evaluation protocol, where each policy is trained and rolled out on one task. While this demonstrates the benefits of contact grounding, it does not test whether CGP can transfer contact knowledge across tasks with different objectives, objects, and contact regimes. Scaling to broader task distributions will likely require cross-task co-training with more diverse demonstrations and interactions, as well as larger and more expressive architectures that can represent a wider range of contact behaviors.

REFERENCES

- [1] Shan An, Ziyu Meng, Chao Tang, Yuning Zhou, Tengyu Liu, Fangqiang Ding, Shufang Zhang, Yao Mu, Ran Song, Wei Zhang, et al. Dexterous manipulation through imitation learning: A survey. *arXiv preprint arXiv:2504.03515*, 2025.
- [2] Yue Chang, Peter Yichen Chen, Zhecheng Wang, Maurizio M Chiamonte, Kevin Carlberg, and Eitan Grinspun. Licrom: Linear-subspace continuous reduced order modeling with neural fields. In *SIGGRAPH Asia 2023 Conference Papers*, pages 1–12, 2023.
- [3] Claire Chen, Zhongchun Yu, Hojung Choi, Mark Cutkosky, and Jeannette Bohg. Dexforce: Extracting force-informed actions from kinesthetic demonstrations for dexterous manipulation. *IEEE Robotics and Automation Letters*, 10(6):6416–6423, 2025. doi: 10.1109/LRA.2025.3568318.
- [4] Cheng Chi, Siyuan Feng, Yilun Du, Zhenjia Xu, Eric Cousineau, Benjamin Burchfiel, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. In *Proceedings of Robotics: Science and Systems*, 2023.
- [5] Hojung Choi, Yifan Hou, Chuer Pan, Seongheon Hong, Austin Patel, Xiaomeng Xu, Mark R Cutkosky, and Shuran Song. In-the-wild compliant manipulation with umi-ft. *arXiv preprint arXiv:2601.09988*, 2026.
- [6] Hao-Shu Fang, Hengxu Yan, Zhenyu Tang, Hongjie Fang, Chenxi Wang, and Cewu Lu. Anydex-grasp: General dexterous grasping for different hands with human-level learning efficiency. *arXiv preprint arXiv:2502.16420*, 2025.
- [7] Shangchen Han, Beibei Liu, Robert Wang, Yuting Ye, Christopher D Twigg, and Kenrick Kin. Online optical marker-based hand tracking with deep labels. *Acm Transactions on Graphics (TOG)*, 37(4):1–10, 2018.
- [8] Shangchen Han, Po-chen Wu, Yubo Zhang, Beibei Liu, Linguang Zhang, Zheng Wang, Weiguang Si, Peizhao Zhang, Yujun Cai, Tomas Hodan, et al. Umetrack: Unified multi-view end-to-end hand tracking for vr. In *SIGGRAPH Asia 2022 Conference Papers*, pages 1–9, 2022.
- [9] Liang Heng, Haoran Geng, Kaifeng Zhang, Pieter Abbeel, and Jitendra Malik. Vitacformer: Learning cross-modal representation for visuo-tactile dexterous manipulation. *arXiv preprint arXiv:2506.15953*, 2025.
- [10] Carolina Higuera, Akash Sharma, Taosha Fan, Chaithanya Krishna Bodduluri, Byron Boots, Michael Kaess, Mike Lambeta, Tingfan Wu, Zixi Liu, Francois Robert Hogan, et al. Tactile beyond pixels: Multisensory touch representations for robot manipulation. In *Conference on Robot Learning*, pages 105–123. PMLR, 2025.
- [11] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 2020.
- [12] Yifan Hou, Zeyi Liu, Cheng Chi, Eric Cousineau, Naveen Kuppuswamy, Siyuan Feng, Benjamin Burchfiel, and Shuran Song. Adaptive compliance policy: Learning approximate compliance for diffusion guided control. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4829–4836, 2025.
- [13] Binghao Huang, Yixuan Wang, Xinyi Yang, Yiyue Luo, and Yunzhu Li. 3d-vitac: Learning fine-grained manipulation with visuo-tactile sensing. In *8th Annual Conference on Robot Learning*, 2024.
- [14] Zixuan Huang, Huaidian Hou, and Dmitry Berenson. Unified multimodal diffusion forcing for forceful manipulation. *arXiv preprint arXiv:2511.04812*, 2025.
- [15] Gagan Khandate, Siqi Shang, Eric T Chang, Tristan Luca Saidi, Yang Liu, Seth Matthew Dennis, Johnson Adams, and Matei Ciocarlie. Sampling-based exploration for reinforcement learning of dexterous manipulation. In *Proceedings of Robotics: Science and Systems*, 2023.
- [16] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [17] Mike Lambeta, Tingfan Wu, Ali Sengul, Victoria Rose Most, Nolan Black, Kevin Sawyer, Romeo Mercado, Haozhi Qi, Alexander Sohn, Byron Taylor, et al. Digitizing touch with an artificial multimodal fingertip. *arXiv preprint arXiv:2411.02479*, 2024.
- [18] Toru Lin, Zhao-Heng Yin, Haozhi Qi, Pieter Abbeel, and Jitendra Malik. Twisting lids off with two hands. In *Conference on Robot Learning*, 2024.
- [19] Jason Jingzhou Liu, Yulong Li, Kenneth Shaw, Tony Tao, Ruslan Salakhutdinov, and Deepak Pathak. Factr: Force-attending curriculum training for contact-rich policy learning. In *Proceedings of Robotics: Science and Systems*, 2025.
- [20] Zhuoyang Liu, Jiaming Liu, Jiadong Xu, Nuowei Han, Chenyang Gu, Hao Chen, Kaichen Zhou, Renrui Zhang, Kai Chin Hsieh, Kun Wu, et al. Mla: A multisensory language-action model for multimodal understanding and forecasting in robotic manipulation. *arXiv preprint arXiv:2509.26642*, 2025.
- [21] Ethan Perez, Florian Strub, Harm De Vries, Vincent

- Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018.
- [22] Haozhi Qi, Brent Yi, Mike Lambeta, Yi Ma, Roberto Calandra, and Jitendra Malik. From simple to complex skills: The case of in-hand object reorientation. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2025.
- [23] Daniel Rakita, Bilge Mutlu, and Michael Gleicher. Relaxedik: Real-time synthesis of accurate and feasible robot arm motion. In *Robotics: Science and Systems*, volume 14, pages 26–30, 2018.
- [24] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022.
- [25] Cristian Romero, Dan Casas, Maurizio Chiaramonte, and Miguel A Otaduy. Learning contact deformations with general collider descriptors. In *SIGGRAPH Asia 2023 Conference Papers*, pages 1–10, 2023.
- [26] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *Proceedings of International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=St1giarCHLP>.
- [27] Yutian Tao, Maurizio Chiaramonte, and Pablo Fernandez. Interpolated adaptive linear reduced order modeling for deformation dynamics. *arXiv preprint arXiv:2509.25392*, 2025.
- [28] UT-Austin-RPL. deoxys_control. GitHub Repository, 2022. URL https://github.com/UT-Austin-RPL/deoxys_control.
- [29] Yeping Wang, Pragathi Praveena, Daniel Rakita, and Michael Gleicher. Rangedik: An optimization-based robot motion generation method for ranged-goal tasks. *arXiv preprint arXiv:2302.13935*, 2023.
- [30] Mengda Xu, Han Zhang, Yifan Hou, Zhenjia Xu, Linxi Fan, Manuela Veloso, and Shuran Song. DexUMI: Using human hand as the universal manipulation interface for dexterous manipulation. In *9th Annual Conference on Robot Learning*, 2025. URL <https://openreview.net/forum?id=XrgRvBklWu>.
- [31] Xiaomeng Xu, Yifan Hou, Zeyi Liu, and Shuran Song. Compliant residual DAGger: Improving real-world contact-rich manipulation with human corrections. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL <https://openreview.net/forum?id=cjcm5LYVWm>.
- [32] Yinzhen Xu, Weikang Wan, Jialiang Zhang, Haoran Liu, Zikang Shan, Hao Shen, Ruicheng Wang, Haoran Geng, Yijia Weng, Jiayi Chen, et al. Unidexgrasp: Universal robotic dexterous grasping via learning diverse proposal generation and goal-conditioned policy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4737–4746, 2023.
- [33] Zhengtong Xu, Raghava Uppuluri, Xinwei Zhang, Cael Fitch, Philip Glen Crandall, Wan Shou, Dongyi Wang, and Yu She. Unit: Data efficient tactile representation with generalization to unseen objects. *IEEE Robotics and Automation Letters*, 2025.
- [34] Han Xue, Jieji Ren, Wendi Chen, Gu Zhang, Yuan Fang, Guoying Gu, Huazhe Xu, and Cewu Lu. Reactive diffusion policy: Slow-fast visual-tactile policy learning for contact-rich manipulation. In *Proceedings of Robotics: Science and Systems*, 2025.
- [35] Max Yang, Alex Church, Yijiong Lin, Christopher J Ford, Haoran Li, Efi Psomopoulou, David AW Barton, Nathan F Lepora, et al. Anyrotate: Gravity-invariant in-hand object rotation with sim-to-real touch. In *8th Annual Conference on Robot Learning*, 2024.
- [36] Jianglong Ye, Keyi Wang, Chengjing Yuan, Ruihan Yang, Yiquan Li, Jiyue Zhu, Yuzhe Qin, Xueyan Zou, and Xiaolong Wang. Dex1b: Learning with 1b demonstrations for dexterous manipulation. In *Proceedings of Robotics: Science and Systems*, 2025.
- [37] Qi Ye, Qingtao Liu, Siyun Wang, Jiaying Chen, Yu Cui, Ke Jin, Huajin Chen, Xuan Cai, Gaofeng Li, and Jiming Chen. Visual-tactile pretraining and online multitask learning for humanlike manipulation dexterity. *Science Robotics*, 11(110):eady2869, 2026.
- [38] Zhao-Heng Yin, Binghao Huang, Yuzhe Qin, Qifeng Chen, and Xiaolong Wang. Rotating without seeing: Towards in-hand dexterity through touch. In *Proceedings of Robotics: Science and Systems*, 2023.
- [39] Di Zhang, Chengbo Yuan, Chuan Wen, Hai Zhang, Junqiao Zhao, and Yang Gao. Kinedex: Learning tactile-informed visuomotor policies via kinesthetic teaching for dexterous manipulation. In *9th Annual Conference on Robot Learning*, 2025. URL <https://openreview.net/forum?id=GKueYvjqSS>.
- [40] Jialiang Zhang, Haoran Liu, Danshi Li, XinQiang Yu, Haoran Geng, Yufei Ding, Jiayi Chen, and He Wang. Dexgraspnet 2.0: Learning generative dexterous grasping in large-scale synthetic cluttered scenes. In *8th Annual Conference on Robot Learning*, 2024.
- [41] Jialiang Zhao, Naveen Kuppaswamy, Siyuan Feng, Benjamin Burchfiel, and Edward Adelson. Polytouch: A robust multi-modal tactile sensor for contact-rich manipulation using tactile-diffusion policies. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 104–110, 2025. doi: 10.1109/ICRA55743.2025.11128816.
- [42] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5745–5753, 2019.
- [43] Yifeng Zhu and Abhishek Joshi. Viola: Imitation learning for vision-based manipulation with object proposal priors. In *Proceedings of Conference on Robot Learning*, 2022.
- [44] Zeshun Zong, Xuan Li, Minchen Li, Maurizio M Chiara-

monte, Wojciech Matusik, Eitan Grinspun, Kevin Carlberg, Chenfanfu Jiang, and Peter Yichen Chen. Neural stress fields for reduced-order elastoplasticity and fracture. In *SIGGRAPH Asia 2023 Conference Papers*, pages 1–11, 2023.

APPENDIX A ADDITIONAL TASK DETAILS

Table V summarizes the task, training, and inference specifications used in our experiments, including three simulated dexterous manipulation tasks with dense tactile arrays and two real-world tasks with vision-based tactile sensing (Digit360).

In-Hand Box Flipping requires continuously reconfiguring multiple fingers to rotate the object while keeping it securely in-hand, balancing sufficient normal force to prevent slip with controlled rolling as contact points migrate across the hand; the robot initial state and the box position and orientation are randomized at the start of each rollout. The task is considered successful when the box is flipped to an upright standing configuration and remains standing without toppling.

Fragile Egg Grasping stresses delicate force regulation under uncertainty: the hand must lift and stabilize the egg while avoiding excessive compression, relying on timely micro-adjustments to prevent slip; the robot initial state and the egg position and orientation are randomized at the start of each rollout. In simulation, we design a fragile egg asset that shatters into fragments when the contact force applied to the egg exceeds a preset threshold. The task is considered successful when the egg is lifted and held stably for 4 seconds without being dropped or cracked; see the supplementary video for qualitative examples.

Dish Wiping is a sustained contact-rich task that requires maintaining a stable contact patch over a long horizon while coordinating arm motion and compliance, where small deviations in force, orientation, or contact location can quickly accumulate into ineffective wiping; we randomize the robot initial state, the dish position and orientation, the sponge position, and the table height at the start of each rollout. In simulation, we simulate the sponge as a soft body (see Section IV.A *Simulation Setup* in the paper), so it deforms under contact during wiping and better captures compliant contact. The task is considered successful when the dish is wiped clean and the sponge is placed to the side at the end of the rollout; see the supplementary video for qualitative examples.

Jar Opening is a contact-intensive task that requires reliable grasping and continuous rotation under load, maintaining high-friction contact to transmit torque while remaining sensitive to alignment and abrupt contact changes during twisting; the robot initial state and the jar position and orientation are randomized at the start of each rollout. The task is considered successful when the lid is fully unscrewed and falls off the jar.

Real In-Hand Box Flipping transfers the in-hand rotation challenge to real hardware, combining complex multi-contact transitions with execution uncertainty, where small errors in

contact timing or force can lead to drift, slip, or drops; the robot initial state and the box position and orientation are randomized at the start of each rollout. The task is considered successful when the box is flipped to an upright standing configuration and remains standing without toppling.

In simulation, tactile observations are provided as a dense tactile array (Table V), while vision consists of two RGB views (agent-view and wrist-view). On real hardware, tactile observations are RGB images from four fingertip Digit360 sensors, and vision follows the same two-view design.

For all policies, we use the same temporal settings at rollout time: an observation horizon of 2 steps, a prediction horizon of 16 steps, and an execution horizon of 8 steps, following the default configuration of diffusion policy [4].

Fig. 12 provides a real-robot visualization by overlaying the target and actual robot states during Contact-Grounded Policy rollouts. Before contact, the small target-actual mismatch mainly reflects steady-state PD tracking error under gravity and friction. During object interaction, the gap becomes noticeably larger, reflecting the offset produced by the joint-space compliance controller under contact forces. When a finger loses contact, the gap rapidly shrinks, consistent with the removal of contact-induced forces. Together, this visualization provides complementary evidence that Contact-Grounded Policy predicts controller-executable target states, and that the resulting target-actual mismatch tracks contact engagement.

APPENDIX B VAE RECONSTRUCTION EXAMPLES FOR TACTILE ARRAYS

Fig. 9 shows VAE reconstruction examples on the validation set for tactile arrays. As shown in Fig. 9, using a 32-dimensional latent space (Table V), the VAE reconstructs dense whole-hand tactile observations with high fidelity across diverse, unseen contact configurations. The reconstructions preserve key contact structure, including the locations of active units, the extent of contact regions, and relative force magnitudes, supporting the use of this latent representation for downstream tactile prediction and contact grounding.

APPENDIX C TACTILE RECONSTRUCTION RESULTS ON DIGIT360 SENSORS

This section reports additional tactile reconstruction and compression results on real Digit360 sensors, complementing the tactile-array analysis in the main paper. To prevent data leakage across correlated frames, we treat each episode as the atomic unit for dataset splitting, and use a 1:1 train-validation split at the episode level.

Table VI summarizes validation results for Digit360 tactile compression on two real-world tasks (in-hand box flipping and jar opening) under different latent dimensions and with/without KL regularization. Compared with KL-regularized training, removing the KL term yields slightly lower reconstruction error, but incurs a substantially larger latent KL loss, indicating a poorly structured latent distribution




Fig. 9: VAE reconstruction examples on the validation set for tactile arrays. We show ground truth (top) and reconstruction (bottom).




Fig. 10: VAE reconstruction examples on the validation set for Digit360 tactile images. We show ground truth (top) and reconstruction (bottom).




Fig. 11: 11 objects with different shapes and sizes used for the hand configuration prediction experiment.

despite marginally improved reconstruction. We further vary the total latent dimension across 80, 160, and 320 (i.e., 20, 40, and 80 dimensions per sensor for four fingertip sensors). Under KL regularization, reconstruction error remains largely unchanged as the latent dimension increases, while the validation KL loss decreases modestly, suggesting improved latent organization with higher capacity.

For Contact-Grounded Policy, smaller latents are preferred for rollout-time efficiency since tactile prediction is performed

in latent space over a horizon. Balancing reconstruction quality, latent regularity, and runtime efficiency, we adopt an 80-dimensional latent for Digit360 in Contact-Grounded Policy, i.e., 20 dimensions per sensor for the four fingertip sensors (Table V). Figure 10 further visualizes validation reconstructions under this setting: despite the compact per-sensor latent, the VAE reconstructs Digit360 tactile images with clear contact structure, capturing deformation patterns and contact regions that are critical for downstream tactile prediction and contact grounding.

APPENDIX D HAND CONFIGURATION PREDICTION DATASET

We collect 150 teleoperated grasping episodes comprising 4114 frames across 11 distinct objects in simulation, as shown in Fig. 11. Each episode corresponds to a unique grasp configuration, covering diverse contact patterns and finger arrangements. To prevent data leakage across correlated frames, we treat each episode as the atomic unit for dataset splitting. We adopt an aggressive 1:1 train-validation split at the episode level, resulting in a challenging generalization setting for validation.

Task	Domain	Demos	T	T_o	T_a	Frequency	Action	State	Tactile	Vision	Latent	KL Weight
In-Hand Box Flipping	Sim	60	16	2	8	5 Hz 8 DDIM steps	29	27	748×3 (Tactile Array)	2×RGB 180×320	32	1×10^{-5}
Fragile Egg Grasping	Sim	100	16	2	8	5 Hz 8 DDIM steps	29	27	748×3 (Tactile Array)	2×RGB 180×320	32	1×10^{-5}
Dish Wiping	Sim	100	16	2	8	5 Hz 8 DDIM steps	29	27	748×3 (Tactile Array)	2×RGB 180×320	32	1×10^{-4}
Jar Opening	Real	45	16	2	8	5 Hz 8 DDIM steps	25	23	4×RGB 72×72 (Digit360)	2×RGB 240×320	80	5×10^{-5}
Real In-Hand Box Flipping	Real	90	16	2	8	5 Hz 8 DDIM steps	25	23	4×RGB 72×72 (Digit360)	2×RGB 240×320	80	5×10^{-5}

TABLE V: Summaries of the task, training, and inference specifications. Here, T is the prediction horizon, T_o is the temporal length of conditioned observations, and T_a is the execution horizon; **Frequency** reports both the policy rollout frequency and the denoising sampling steps: we rollout the policy at 5 Hz for all tasks and use 8 DDIM steps per inference. The specific inference time is reported in Fig. 8 of the paper. In simulation, the action space has 29 dimensions (arm position + rot6D [4, 42] + 20 hand joints) and the robot state observation has 27 dimensions (arm position + quaternion + 20 hand joints). On real hardware, the action space has 25 dimensions (arm position + rot6D + 16 hand joints) and the robot state observation has 23 dimensions (arm position + quaternion + 16 hand joints). rot6D facilitates stable action regression [4, 42]. **Latent** denotes the tactile latent dimension: 32 for tactile arrays and 80 for Digit360 (20 per sensor, 4 sensors). Vision uses two RGB views (agent and wrist) with resolutions shown in the table. **KL Weight** is the weighting coefficient of the KL term in the training objective.

Latent Dim	Model	Box				Jar			
		MAE ↓	KL Loss ↓	PSNR (dB) ↑	SSIM ↑	MAE ↓	KL Loss ↓	PSNR (dB) ↑	SSIM ↑
80	w/ KL	9.01	0.0694	35.2050	0.9870	8.10	0.0647	35.9445	0.9883
	w/o KL	4.55	0.7781	40.4590	0.9939	3.88	0.6247	42.9845	0.9957
160	w/ KL	9.41	0.0341	35.2768	0.9874	8.40	0.0337	36.5394	0.9891
	w/o KL	4.66	0.7902	40.4800	0.9941	3.98	0.6899	42.9143	0.9957
320	w/ KL	9.02	0.0171	35.2693	0.9870	8.29	0.0155	36.1177	0.9885
	w/o KL	4.56	0.7539	40.5191	0.9941	3.84	0.6787	43.0426	0.9957

TABLE VI: Validation results for tactile compression on real in-hand box flipping and jar opening tasks. MAE denotes mean absolute error (reported in $\times 10^{-3}$), computed on tactile images normalized to $[0, 1]$.

APPENDIX E

ROBUSTNESS OF THE CONTACT-CONSISTENCY MAPPING TO TACTILE NOISE

We evaluate the robustness of the contact-consistency mapping using the hand configuration prediction experiment in Fig. 6 and Table III. To simulate tactile prediction errors, we inject Gaussian noise into the tactile inputs, where the noise standard deviation is specified as a percentage of the dataset-level tactile standard deviation. As shown in Fig. 13, the prediction error remains stable across a wide range of noise levels. This result suggests that the contact-consistency mapping is robust to moderate tactile perturbations, and that errors in tactile prediction do not catastrophically propagate to the mapped robot state. In addition, the model using both tactile and robot state inputs consistently outperforms the tactile-only variant, while residual mapping remains more accurate than absolute mapping. These results further support our design choices for conditioning the mapping on proprioceptive state and predicting residual corrections rather than absolute hand configurations.

● Target State ● Actual State




Fig. 12: Snapshot of real-world Contact-Grounded Policy rollouts, with overlaid target and actual robot states. Before contact, the small target-actual mismatch primarily reflects the steady-state tracking error of the low-level PD controller (e.g., due to gravity and friction; see the first column). During contact, the gap grows as the joint-space compliance controller yields to contact forces; in this phase, the gap reflects the physical interaction. When a finger loses contact, the contact forces disappear and the gap quickly shrinks.




Fig. 13: Robustness of the contact-consistency mapping to tactile noise. Gaussian noise is injected into tactile inputs with different standard deviations, expressed as percentages of the dataset-level tactile standard deviation. The prediction error remains stable across a wide range of noise levels, indicating that the mapping is robust to tactile perturbations.