TactiDex: A Real-World Tactile-Guided Benchmark for Human-Like Dexterous Manipulation

Suting Ni[1,2] Hanbing Zhang[1,2] Zhenyu Wei[1] Guo Chen[1] Chixuan Zhang[1] Ye Shi[1,2] Jingya Wang[1,*]

{nist2024, zhanghb2025, weizhy2024, chenguo2024, zhangchx12024, shiye, wangjingya}@shanghaitech.edu.cn

[1] Shanghaitech University [2] InstAdapt

arXiv:2607.09190v1 [cs.RO] 10 Jul 2026

Abstract

Tactile feedback is fundamental to Hand-Object Interaction (HOI), governing contact formation, force regulation, and stable manipulation, making it essential for achieving true human-like dexterous manipulation. Yet, current human-to-robot dexterous transfer pipelines primarily rely on kinematic trajectories, resulting in motion imitation without physically grounded interaction. To address this, we introduce TactiDex, a real-world tactile-guided benchmark specifically designed to move dexterous manipulation beyond kinematic mimicry toward contact-level human-likeness. TactiDex provides a comprehensive dataset that elegantly aligns whole-hand tactile signals with multi-granularity kinematic and object states, coupled with standardized evaluation metrics. Building upon this data paradigm, we propose a tactile-driven transfer framework that effectively translates human demonstrations into physically plausible robotic execution. We introduce TactiSkill, a framework built upon a novel tri-component tactile reward that innovatively uses tactile signals as structured supervision. This reward unifies guidance, human-like alignment, and contact constraints into a single objective. Through comprehensive experiments on both single and bimanual tasks, we demonstrate that TactiSkill achieves superior performance in manipulation success and physical realism. This work lays a crucial foundation for advancing tactile-aware dexterous manipulation. Our project page at https://tactidex.github.io/.

Keywords

Tactile-Guided Transfer, HOI Dataset, Dexterous Manipulation

1 Introduction

Hand-object interaction (HOI) lies at the core of embodied intelligence, enabling humans to manipulate tools, operate daily objects, and perform complex dexterous skills. A defining characteristic of human dexterity is the continuous regulation of contact through tactile feedback, which governs contact formation, force modulation, and manipulation stability. While vision provides spatial and geometric awareness [13, 50, 51], tactile perception determines how physical interaction is established and maintained at the hand-object interface. Neuroscience and robotics research consistently show that successful manipulation relies heavily on force modulation rather than mere spatial positioning [30, 38, 52, 56]. Consequently, human dexterity is not solely a function of kinematic trajectories, but fundamentally grounded in contact-level physical interaction.

Transferring such human dexterity to robotic systems remains a longstanding challenge. Recent advances in dexterous manipulation have explored human-to-robot skill transfer through imitation learning [2, 8, 26, 39, 47], motion retargeting [24, 40, 49, 53], and reinforcement learning frameworks [48, 54]. These approaches typically align human hand kinematics with robotic embodiments, leveraging visual observations and joint-level motion supervision to reproduce task behaviors. Despite impressive progress, this trajectory-centric paradigm inherently overlooks the tactile dynamics essential for regulating contact formation and force distribution. Consequently, robotic policies often successfully replicate superficial movement patterns, but deviate significantly in contact behavior, interaction stability, and force modulation.

A key factor underlying this gap is the absence of tactile-rich benchmarks and standardized evaluation protocols. Without strictly synchronized tactile-kinematic demonstrations, learning objectives are implicitly forced to prioritize motion matching over physically grounded interaction, leaving a substantial gap in achieving true human-like robotic dexterity.

To address this gap, we introduce **TactiDex**, a real-world tactile-rich HOI dataset for contact-aware human-to-robot dexterous transfer. TactiDex provides synchronized whole-hand tactile sensing, fine-grained hand kinematics, wrist 6D poses, and object 6D trajectories, forming a temporally aligned and contact-rich representation of human manipulation dynamics. Unlike prior datasets that focus primarily on visual observations or joint trajectories [4, 12, 33, 34, 55, 58], TactiDex captures contact interactions between the hand and object through spatially distributed pressure measurements. The dataset spans diverse single-hand and bimanual manipulation tasks, covering varied object geometries and interaction patterns, thereby providing a broad testbed for studying tactile-aware transfer. Beyond data collection, we establish standardized contact-aware evaluation protocols that quantify contact fidelity, force alignment, and interaction stability. By integrating synchronized multi-granularity tactile and spatial states within a unified framework, TactiDex establishes a benchmark tailored to studying tactile-aware human-to-robot transfer.

Building upon this data paradigm, we propose **TactiSkill**, built on the **Tactile-Guided Tri-Component Transfer Framework** that leverages synchronized contact signals to guide robotic policy learning. Directly incorporating raw pressure measurements into reinforcement learning objectives, however, presents non-trivial

* Corresponding author: wangjingya@shanghaitech.edu.cn

challenges. Naïve formulations based solely on contact force matching often lead to unstable optimization, reward exploitation, or physically unrealistic behaviors in simulation [3, 14, 22, 31, 59]. To address this, we introduce a tri-component tactile objective conceptualizing tactile signals as structured supervision, which integrates (1) contact guidance to encourage meaningful and timely contact formation, (2) human-like alignment to match force distribution patterns observed in human demonstrations, and (3) safety constraints to regularize excessive force and promote stable interaction. This structured formulation enables stable optimization while preserving contact-level fidelity, effectively moving dexterous transfer beyond kinematic mimicry toward physically grounded human-likeness.

In summary, our contributions are as follows:

(1) We introduce **TactiDex**, the first high-precision, long-horizon tactile-rich hand–object interaction dataset for dexterous bimanual manipulation, providing temporally synchronized whole-hand tactile signals, fine-grained hand kinematics, and object 6D trajectories, together with tactile-aware evaluation metrics for human-to-robot transfer.

(2) We propose **TactiSkill**, a tactile-guided human-to-robot transfer framework that leverages tactile sensing as a core supervisory signal for dexterous policy learning. To stabilize contact-rich interaction, we design a tri-component tactile reward framework that (i) encourages proactive contact formation, (ii) aligns tactile distributions with human demonstrations to achieve force-level human-likeness, and (iii) constrains excessive contact forces to ensure physically stable manipulation.

(3) Extensive experiments demonstrate that tactile-guided transfer significantly improves geometric robustness and physical fidelity across diverse manipulation tasks, and we further validate the effectiveness of TactiSkill through real-world dexterous hand deployment, showing that policies derived from our dataset can reproduce stable contact-rich manipulation on physical hardware.

2 Related Work

2.1 Hand-Object Interaction Dataset

Numerous datasets facilitate hand-object interaction (HOI) modeling [4, 12, 23, 33, 34, 46, 55, 58]. Motion-capture datasets [4, 12, 44, 58] provide high-precision 3D trajectories for detailed manipulation analysis, while egocentric datasets [19, 33] expand semantic diversity in open-world and instruction-conditioned settings. However, both paradigms primarily infer contact from geometric proximity or visual observation, leaving actual force modulation physically unmeasured. Synthetic benchmarks [18] offer scalable contact signals but rely on simulated geometry, lacking real-world tactile fidelity. Most recently, OPENTOUCH [42] introduced real-world whole-hand tactile sensing with text annotations; yet, it predominantly focuses on egocentric perception tasks (e.g., cross-modal retrieval and grasp classification) rather than dynamic manipulation control. Consequently, as summarized in Table 1, existing HOI datasets largely lack the tactile grounding necessary for physically interactive policy learning. To bridge this gap, we introduce **TactiDex**, a tactile-grounded HOI benchmark that synchronizes whole-hand tactile sensing with multi-granularity kinematic states and establishes standardized, contact-aware evaluation protocols, enabling physically grounded modeling of human dexterity beyond geometric imitation.

2.2 Dexterous Manipulation & Human-to-Robot Transfer

Dexterous manipulation has witnessed substantial progress through model-based control, reinforcement learning, and large-scale policy learning [1, 7, 27, 48, 54]. Concurrently, learning from human demonstrations offers an intuitive pathway for transferring dexterity to robotic hands [39, 47]. Methods such as imitation learning and motion retargeting align trajectories under morphological constraints [24, 40], while recent RL frameworks refine these transferred policies using reference motions and task-specific rewards [6, 9, 26, 32, 35]. Although these methods significantly improve trajectory reproduction, contact-level discrepancies frequently persist. Geometric alignment does not guarantee consistency in force modulation or pressure distribution, often resulting in unstable grasps, excessive force, or physically implausible interactions during delicate, contact-rich tasks. Therefore, while trajectory imitation is a strong starting point, it fails to fully resolve the physical interaction gap. In contrast, our approach introduces tactile-guided supervision that directly constrains contact formation and force distribution, shifting the human-to-robot transfer paradigm from purely geometric trajectory alignment to physically grounded interaction alignment.

2.3 Tactile-Guided Policy Learning

Tactile sensing is increasingly recognized as a critical modality for robust manipulation. Early tactile-guided reinforcement learning (RL) applications predominantly utilize parallel-jaw grippers to solve foundational tasks like peg-in-hole insertions and terminal assemblies [15, 16, 28, 45]. To improve generalizability, subsequent studies extensively explore visuo-tactile fusion and sim-to-real transfer paradigms to address the reality gap during unknown object manipulation [10, 11, 17, 21, 29, 36, 43, 60]. More recently, tactile RL has scaled to highly articulated multi-fingered dexterous hands [20, 22], accompanied by the development of multimodal pretraining benchmarks [31] and attempts to leverage human tactile demonstrations for imitation learning [57]. Despite this rapid progress, existing frameworks largely treat tactile signals merely as auxiliary observation vectors for unconstrained trial-and-error optimization, suffering fundamentally from the absence of synchronized, whole-hand hand-object interaction (HOI) physical datasets. To bridge this gap, our work introduces a true *human-to-robot* tactile transfer paradigm by leveraging the real-world **TactiDex** dataset. Instead of learning contact heuristics from scratch, **TactiSkill** formulates high-fidelity human force distributions not as passive observations, but as active, structural physical priors. By directly supervising the force modulation of the robotic hand, our framework ensures the resulting policies transcend geometric mimicry to achieve physically grounded, human-like dexterity.

3 TactiDex Dataset

TactiDex, as overviewed in Fig. 1, is a large-scale, real-world tactile-grounded dataset designed to capture the physical dynamics of

Table 1: Comparison with existing HOI datasets. High-Prec. Kinematics: Global motion capture vs. vision-based estimation. Tactile: Whole-hand pressure maps. TactiDex uniquely unifies exact kinematics, physical tactile sensing, and text annotations for long-horizon bimanual manipulation.

Dataset	Frame	Subj	Obj	Real-World	Bimanual	High-Prec. Kinematics	Tactile	Text Annot.	Long-horizon
ObMan [18]	154K	20	3K						
ContactPose [51]	2.99M	50	25	✓	✓	✓			
HOI4D [33]	3M	9	1000	✓					
DexYCB [4]	582K	10	20	✓		✓			
ARCTIC [12]	2.1M	19	11	✓	✓	✓			
OakInk2 [58]	4.01M	9	75	✓	✓	✓		✓	✓
GRAB [44]	1.62M	10	51	✓	✓	✓			
OPENTOUCH [42]	0.54M	-	800	✓			✓	✓	
TactiDex (Ours)	5.1M	10	49	✓	✓	✓	✓	✓	✓

Figure 1: Overview of the TactiDex Data Collection and Annotation Framework. (Left) Our hardware setup integrates a high-precision motion capture system with a novel dual-glove design—an inner capture glove for kinematics and an outer tactile glove for whole-hand pressure mapping, strictly synchronized via an eSync module. Manipulated objects are digitally reconstructed via 3D scanning. (Center) A user performs complex, contact-rich interactions (e.g., bimanual grinding) within the capture volume. (Right) The pipeline yields strictly synchronized multi-modal data streams, including tactile signals, hand kinematics, and hand-object 6D poses, enriched with comprehensive language descriptions and hierarchical task phase annotations.

human dexterous manipulation. It synchronizes whole-hand tactile sensing, articulated hand kinematics, wrist 6D poses, and object 6D trajectories within a unified motion-capture environment. The dataset comprises **49 calibrated objects** and **757 interaction sequences**, spanning from fundamental pick-and-place behaviors to complex functional operations and object-object bimanual interactions. By providing structured interaction annotations and task-level textual descriptions alongside multimodal sensory streams, TactiDex establishes a physically measurable and semantically grounded benchmark for contact-aware human-to-robot transfer.

3.1 Hardware Setup

To develop the TactiDex dataset, we built a multimodal data collection system for tactile-grounded hand–object interaction that synchronously records hand contact information, hand kinematic states, and accurate object 6D poses in a global coordinate frame. The system consists of three hardware components: a full-hand tactile glove that captures pressure distributions, a hand-tracking glove that acquires articulated joint states, and an OptiTrack motion capture system that provides global 6-DoF poses for both the object and the hand.

3.1.1 Tactile Measurement. The tactile gloves measure pressure via piezoresistive changes in flexible sensors. The tactile array comprises 162 sensing elements densely distributed across the fingertips and palm, with a force resolution of up to 0.01 N and a sampling rate of 17 Hz. We further redesigned the structure of tactile gloves and wearing configuration so that it can be worn externally over the motion capture glove. This modification enables stable acquisition of tactile signals from the whole hand without degrading the kinematic tracking quality of the mocap glove, while also minimizing interference with natural manipulation, improving reliability and consistency for long-duration data collection.

3.1.2 Motion Capture. Hand kinematics are captured using a motion-tracking glove, where IMU measurements are fused with optical tracking data to estimate high-fidelity articulated joint states. Global 6-DoF poses for both the wrist and the objects are provided by an OptiTrack system equipped with 8 PX13W cameras operating at 120 fps within a 1.2 m × 1.8 m workspace. As both hand and object poses are processed within the same motion-capture environment, they are inherently aligned within a unified spatial coordinate system and temporal base.

3.2 Data Collection and Processing

Synchronized Multimodal Capture. To ensure high-fidelity data acquisition, we established a strict collection protocol. Participants perform manipulation tasks within the calibrated multi-camera optical workspace. The hardware-synchronized pipelines ensure frame-level temporal alignment between the 120 Hz kinematic/pose

streams and the tactile pressure data, enabling precise temporal modeling of contact formation and evolution.

Object Modeling and Calibration. The 49 objects in our dataset were meticulously selected to cover diverse geometries, scales, and physical properties. Object meshes are reconstructed using a hybrid pipeline that integrates high-resolution mobile 3D scanning applications with SAM3D-based reconstruction methods [5]. To overcome the inherent feature-matching failures caused by reflective, transparent, or highly symmetric surfaces, we applied specialized non-reflective textured tapes and temporary fiducial markers during the scanning phase. This rigorous treatment ensured robust visual feature extraction, yielding high-fidelity triangular meshes regardless of challenging material properties. Following reconstruction, each object is dynamically calibrated within the OptiTrack system to ensure its 6D trajectory is spatially consistent with the hand's global coordinate frame.

Task Coverage and Data Scope. We recorded 757 interaction sequences encompassing single-hand and bimanual manipulation. Task complexity ranges from basic physical interactions (e.g., pipkplace, inspect) to tool use and functional operations (e.g., click, cut, shake). Each sequence pairs the synchronized multimodal sensory data with structured annotations, including interaction-phase timestamps, task identifiers, and natural-language descriptions, providing a comprehensive foundation for multimodal policy learning.

Optimization and Processing. Following OMOMO [25], we adopt the MANO model [41] for kinematic representation. We independently fit MANO parameters to the raw 3D motion-capture keypoints via L-BFGS optimization. To ensure anatomical and temporal consistency, we share shape parameters across all frames for each subject, and apply both an ℓ_2 pose regularizer and a second-order temporal smoothness penalty to suppress high-frequency jitter.

While kinematic fitting captures accurate free-space trajectories, residual errors often cause physically implausible "floating" or penetration during object interaction. To resolve this, we introduce a two-stage *tactile-constrained post-optimization*. In the first stage, we identify genuine contact intervals by gating geometric proximity with measured tactile pressure, effectively filtering out visually deceptive near-contacts. Within each interval, we select a stable reference grasp frame and optimize its spatial alignment using tactile-modulated per-finger weights: active fingers (exceeding the pressure threshold) are strongly attracted to the object surface, while inactive fingers are down-weighted to prevent phantom contacts. In the second stage, this refined grasp is propagated across the contact interval via temporal inverse kinematics, preserving original motions in non-contact regions. Finally, a collision-aware adjustment optimizes finger articulation against the object's Signed Distance Field (SDF) [37] to eliminate residual penetrations. This pipeline yields the final TactiDex annotations, ensuring that interactions are both kinematically precise and physically grounded by real-world tactile evidence.

4 Methodology

As illustrated in Fig. 2, the core objective of our framework is to leverage the TactiDex dataset to train dexterous hands in simulation, enabling them to replicate both the spatiotemporal kinematics and the dynamic contact patterns of human demonstrations.

4.1 Overview & Problem Formulation

To achieve physically grounded transfer, we model the contact-aware dexterous manipulation as a constrained Markov Decision Process (MDP), denoted by the tuple $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma \rangle$. Unlike traditional kinematic-centric setups, we explicitly incorporate high-fidelity tactile signals as Target Contact Features into the state space to enforce physical realism:

- **State Space $\mathcal{S}$:** At timestep t, the state $\mathbf{s}_t \in \mathcal{S}$ includes standard proprioception (joint positions $\mathbf{q}_t$, velocities $\dot{\mathbf{q}}_t$) and object states (pose $\mathbf{p}_t$). Crucially, we introduce a Target Tactile Reference denoted as $\mathbf{F}_t^{\text{human}}$ derived from the TactiDex, which guides the robot on the exact normal force each finger should apply at the current moment.
- **Action Space $\mathcal{A}$:** The action $\mathbf{a}_t = \Delta\mathbf{q}_t$ is defined as the residual joint position commands, serving to fine-tune the base kinematic trajectory.

The optimization objective is to find an optimal policy $\pi^*(\mathbf{a}_t|\mathbf{s}_t)$ that minimizes kinematic tracking errors while maximizing force consistency with the human demonstration.

4.2 Tactile-Guided Tri-Component Human-to-Robot Transfer

To effectively transfer human demonstrations to morphologically distinct robotic hands, we propose **TactiSkill**, a novel policy learning method built upon the **TactiDex** dataset. As depicted in the central block of Fig. 2, TactiSkill operates through our proposed **Tactile-Guided Tri-Component Transfer Framework**.

Building upon recent residual reinforcement learning architectures designed for morphological retargeting [26], our final policy π^* is composed of a frozen kinematic imitation policy π_{base} and a learnable residual policy π_{res}. However, instead of acting merely as a geometric compensator to fix trajectory errors, TactiSkill fundamentally re-purposes π_{res} as a *Force Modulator*. By actively inducing and regulating simulated contact forces, π_{res} bridges the gap between geometric motion mimicry and true physical interaction through a tightly integrated three-step pipeline.

(1) Finger-wise Sensor-to-Sim Data Alignment.
The TactiDex dataset captures raw sensor signals (high-frequency ADC readings), whereas the physics simulation engine computes contact forces in Newtons based on rigid-body dynamics. To establish a unified physical baseline for supervision, we construct a finger-wise, non-linear calibration mapping function $\mathcal{M} : \mathbb{R}^{\text{raw}} \to \mathbb{R}^N$. This mapping precisely converts the raw human demonstration signals into aligned normal forces consistent with the simulation environment.

(2) Asymmetric Observation with Privileged Tactile Information.
To enhance training stability and sample efficiency, TactiSkill employs an asymmetric Actor-Critic architecture. The *Actor* only receives the target tactile reference $\mathbf{F}_t^{\text{human}}$, mimicking the real-world deployment constraint where the robot possesses a desired force reference but cannot directly foresee the resulting contact dynamics. Conversely, the *Critic* receives *privileged information*—specifically, the real-time simulated fingertip contact force $\mathbf{F}_t^{\text{sim}}$. This enables the Critic to accurately evaluate the value gradient in the latent space,

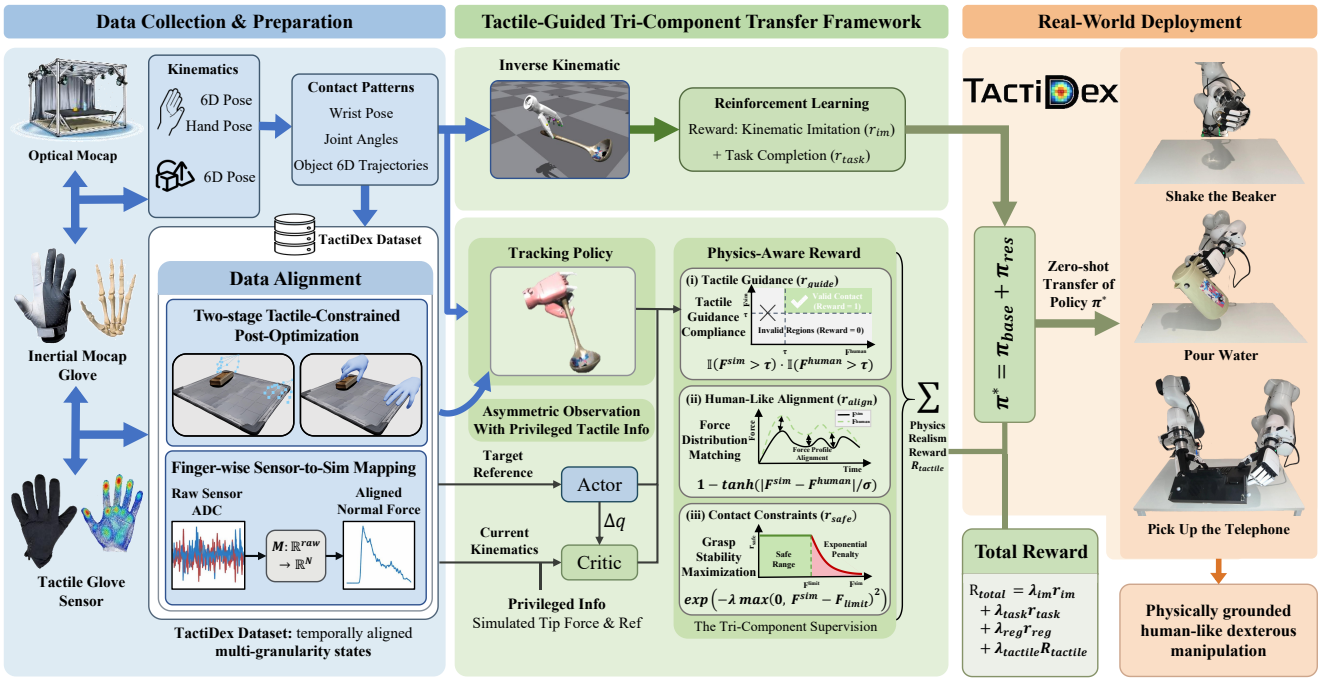

Figure 2: Overview of the TactiSkill Pipeline. (Left) Data Collection & Preparation: Raw multi-modal signals from human demonstrations undergo a two-stage tactile-constrained post-optimization and sensor-to-sim mapping, yielding the physically aligned TactiDex dataset. (Center) Tactile-Guided Tri-Component Transfer Framework: We train a residual tracking policy using an asymmetric Actor-Critic architecture, where the Critic receives privileged simulated contact forces. A physics-aware tri-component reward (Tactile Guidance, Human-Like Alignment, and Contact Constraints) enforces force-level human-likeness. (Right) Real-World Deployment: The composite policy (π^*) enables zero-shot sim-to-real transfer, achieving physically grounded, human-like dexterous manipulation on actual robotic hardware across diverse tasks.

efficiently guiding the Actor to adjust joint poses to approximate the target force distribution.

(3) Physics-Aware Reward: The Tri-Component Supervision. To enforce force-level human-likeness, the composite tactile reward R_{tactile} regulates the normal contact force F_i of the i-th finger (out of N_f fingers). It is formulated as a sum of three distinct components, *Tactile Guidance* r_{guide}, *Human-Like Alignment* r_{align} and *Contact Constraints* r_{safe} balanced by weights w_g, w_a, w_s.

$$R_{\text{tactile}} = w_g \cdot r_{\text{guide}} + w_a \cdot r_{\text{align}} + w_s \cdot r_{\text{safe}}. \tag{1}$$

i. Tactile Guidance. Overcomes the common "air grasping" issue by providing a binary contact bonus. It acts as a "contact trigger," rewarding the agent only when it successfully replicates an active contact event initiated by the human demonstrator, where τ is a minimal contact threshold:

$$r_{\text{guide}} = \frac{1}{N_f} \sum_{i=1}^{N_f} \mathbb{I}\left(F_i^{\text{sim}} > \tau\right) \cdot \mathbb{I}\left(F_i^{\text{human}} > \tau\right). \tag{2}$$

ii. Human-Like Alignment. Merely touching the object is insufficient; the robot must replicate the force modulation patterns. We employ a robust Tanh-based distance metric to minimize the discrepancy between simulated and human forces. Unlike Mean Squared Error, this formulation provides normalized gradients and preserves the force distribution across fingers while resisting sensor outliers (σ is a scaling factor):

$$r_{\text{align}} = \frac{1}{N_f} \sum_{i=1}^{N_f} \left(1 - \tanh\left(\frac{|F_i^{\text{sim}} - F_i^{\text{human}}|}{\sigma}\right)\right). \tag{3}$$

iii. Contact Constraints. Prevents the policy from exploiting the physics engine to generate unrealistically high forces. We define a dynamic safety limit $F_{\text{limit}} = F_i^{\text{human}} + \delta$ (where δ is a safety margin). An exponential penalty is applied only when the simulated force exceeds this limit, ensuring a physically stable and safe Sim-to-Real transfer:

$$r_{\text{safe}} = \exp\left(-\lambda \sum_{i=1}^{N_f} \max\left(0, F_i^{\text{sim}} - F_{\text{limit}}\right)^2\right). \tag{4}$$

Total Reward Formulation. To ensure the policy converges to a kinematically accurate and physically plausible solution, the final reward function at each timestep t combines our novel tactile supervision with standard kinematic constraints:

$$R_{total} = \underbrace{\lambda_{im} r_{im} + \lambda_{task} r_{task} + \lambda_{reg} r_{reg}}_{\text{Kinematic \& Task Constraints}} + \underbrace{\lambda_{tactile} R_{tactile}}_{\text{Physical Realism}}, \tag{5}$$

While r_{im} (pose tracking), r_{task} (object trajectory), and r_{reg} (energy efficiency) guide the robot to the correct spatiotemporal configuration, the R_{tactile} term acts as a critical dynamic regularizer, guaranteeing that the interaction forces remain human-like and physically stable.

5 Experiments

5.1 Data Split and Evaluation Metrics

To rigorously evaluate the effectiveness of our tactile-guided transfer framework, we select 73 representative interaction sequences from the **TactiDex** dataset. This evaluation split encompasses both single-hand and bimanual scenarios, covering diverse object geometries, materials, and varying task complexities. Based on the evaluation dimensions shown in Table 2, we comprehensively assess the transferred policies across four distinct aspects:

1) Kinematic and Object Tracking. Following established practices [26], we evaluate geometric alignment using Object Tracking Error in Translation (**OTE_t**) and Rotation (**OTE_r**). To evaluate how closely the robotic hand replicates human kinematics, we compute the Mean Per-Joint Position Error (**MPJPE**) and specifically the Mean Per-Fingertip Position Error (**TipErr**), as fingertips are the primary medium for interaction.

2) Tactile and Physical Fidelity. To benchmark physical contact-level human-likeness—a core contribution of TactiSkill—we introduce a suite of fine-grained tactile metrics. Let $f_{t,i}^{\text{sim}}$ and $f_{t,i}^{\text{gt}}$ denote the normal force magnitude on the i-th finger (out of N_f total fingers) at timestep t (out of T total frames).

- **Mean Tactile Force Error (MTFE).** Computes the mean absolute difference between simulated forces and the aligned human ground truth across all fingers and temporal frames:

$$\text{MTFE} = \frac{1}{T \cdot N_f} \sum_{t=1}^{T} \sum_{i=1}^{N_f} \left| f_{t,i}^{\text{sim}} - f_{t,i}^{\text{gt}} \cdot \right| \tag{6}$$

- **Contact F1-Score:** We binarize the continuous force signals into contact states $C_{t,i} = \mathbb{I}(f_{t,i} > \tau_c)$ using a minimal contact threshold τ_c. The F1-Score evaluates the spatial-temporal precision and recall of the simulated contact states against the ground truth, measuring whether the robot engages the correct fingers at the optimal time.

3) Safety and Reliability. To explicitly penalize unstable physical spikes and momentary destructive penetrations, we monitor the peak transient force discrepancy $\text{MaxErr} = \max_{t,i} |f_{t,i}^{\text{sim}} - f_{t,i}^{\text{gt}}|$ and introduce two episode-level safety metrics:

- **PeakSafe@3N:** Calculates the fraction of episodes where the maximum instantaneous force error stays strictly below a safety threshold of 3N (MaxErr < 3 N), indicating stable and safe execution without destructive force spikes.
- **SafeTac@3N:** A strictly combined metric requiring the episode to simultaneously achieve tactile-aware success *and* satisfy the peak safety constraint (MaxErr < 3 N).

4) Task Success Rate (SR). We evaluate task completion at two distinct levels of strictness. First, the **Kinematic Success Rate (SR_kin)** adopts the criteria from [26]: a trial is successful if $\text{OTE}_\text{r} \leq 30°$, $\text{OTE}_\text{t} \leq 3$ cm, MPJPE ≤ 8 cm, and TipErr ≤ 6 cm. For bimanual tasks, both hands must simultaneously meet these spatial conditions. Second, to explicitly emphasize physical realism, we employ a stricter **Tactile-Aware Success Rate (SR_tac)**. A trial achieves tactile success only if it satisfies the kinematic criteria *and* demonstrates high physical fidelity, specifically requiring MTFE ≤ 3.0 N and Contact F1 ≥ 0.3.

5.2 Main Results and Ablative Analysis

To isolate the contributions of our tactile-guided supervision, we evaluate our **TactiSkill** framework against a kinematic baseline (**ManipTrans** [26]) and conduct an ablation study removing core tactile components: (1) **w/o Bonus**, (2) **w/o Align**, and (3) **w/o Safety** (Table 2).

The Insufficiency of Pure Kinematic Imitation. The purely kinematic baseline (**Base**) overfits to geometric alignment at the expense of physical realism. While achieving a seemingly adequate SR_kin (72.91%) and OTE_t (1.1947 cm), this performance is deceptive. By ignoring contact mechanics, the baseline frequently accomplishes goals through "air grasping" or mesh penetration. Consequently, its tactile-aware success rate (SR_tac) drops precipitously to 39.35%, and its safety metrics remain low (Contact F1 = 0.5569, SafeTac@3N = 35.84%). This unequivocally indicates that kinematic imitation alone is insufficient for robust tactile manipulation.

The Synergistic Effect of Tactile Guidance. In contrast, the full **TactiSkill** framework demonstrates that tactile supervision *synergizes* with kinematic tracking to act as a powerful physical regularizer. By enforcing valid contacts, TactiSkill prevents collapsing into invalid geometric optima, attaining the highest SR_kin (81.95%) alongside the lowest tracking errors (OTE_t = 0.9577 cm, OTE_r = 6.1974°). Most importantly, TactiSkill vastly outperforms all baselines in physical fidelity. It secures the highest Contact F1-score (0.7384) and boosts SR_tac to 64.64%. The superior episode-level safety metrics, particularly PeakSafe@3N (72.80%) and SafeTac@3N (53.57%), validate that our method yields highly reliable and human-like manipulation behavior without destructive force spikes. Specifically, this leap in tactile fidelity highlights a critical shift from passive trajectory following to active force regulation.

Ablation on Contact Initiation and Alignment. The ablation variants validate our tri-component design. Removing the contact bonus (**w/o Bonus**) cripples the incentive to initiate stable grasps, causing the Contact F1-score to plummet to 0.5425, and PeakSafe@3N to drop to 57.00%. Omitting force alignment (**w/o Align**) leads to spatially awkward grasps. Failing to match the ground-truth force distribution means the policy cannot utilize anatomically optimal contact patches. This misaligned application severely degrades dynamic stability, reducing Contact F1-score (0.5134) and resulting in a subpar tactile-aware success rate (SR_tac = 42.20%).

The Paradox of the Safety Penalty. Interestingly, removing the safety penalty (**w/o Safety**) yields the lowest mean errors (MPJPE = 5.0327 cm, MTFE = 0.0921 N). However, these seemingly superior *average* metrics are highly deceptive. Without an explicit penalty for transient force spikes, the unconstrained policy learns to game the average metric, permitting occasional, highly destructive force impulses. This is evidenced by the degradation in peak safety: compared to the full model, PeakSafe@3N drops to 68.70%, and

Table 2: Quantitative comparison of TactiSkill against the kinematic baseline and ablation variants over 73 sequences. SR_{kin} denotes kinematic success, while SR_{tac} enforces strict tactile fidelity criteria. PeakSafe@3N measures the fraction of episodes whose peak force error stays below 3N, and SafeTac@3N further requires tactile-aware success under the same safety constraint. Best results are bolded.

Method	Success Rate		Object Tracking		Hand Kinematics		Tactile Fidelity		Safety & Reliability	
	SR_{kin} ↑	SR_{tac} ↑	OTE-t (cm) ↓	OTE-r (°) ↓	MPJPE (cm) ↓	TipErr (cm) ↓	MTFE (N) ↓	Contact F1 ↑	PeakSafe@3N ↑	SafeTac@3N ↑
Kinematic Base [26]	0.7291	0.3935	1.1947	7.4509	5.0990	5.1833	0.1491	0.5569	0.6366	0.3584
TactiSkill (w/o Bonus)	0.7687	0.4193	1.2031	8.2507	5.1743	5.1746	0.1529	0.5425	0.5700	0.3568
TactiSkill (w/o Align)	0.7637	0.4220	1.0828	6.9556	5.2308	5.2643	0.1412	0.5134	0.5988	0.3441
TactiSkill (w/o Safety)	0.7705	0.4124	1.0265	6.3860	**5.0327**	5.2398	**0.0921**	0.5686	0.6870	0.3872
TactiSkill (Full)	**0.8195**	**0.6464**	**0.9577**	**6.1974**	5.0915	**5.0793**	0.1236	**0.7384**	**0.7280**	**0.5357**

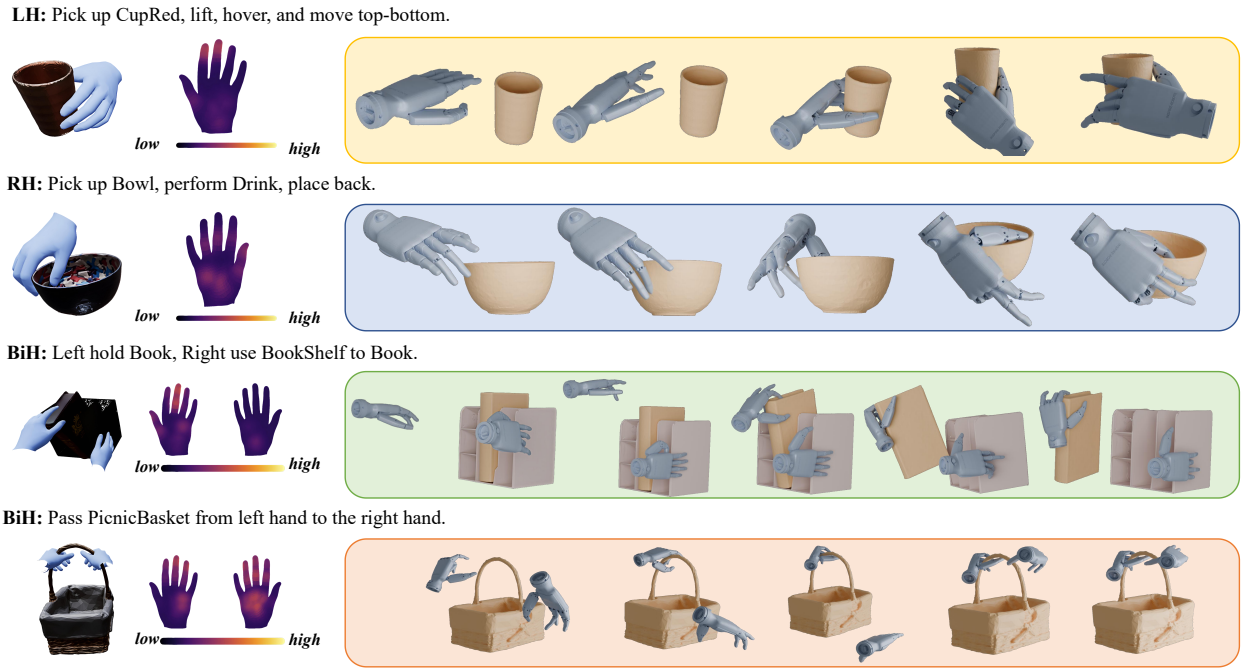

Figure 3: Qualitative Results of TactiSkill. Transfer sequences for single-hand (LH/RH) and bimanual (BiH) tasks. Left: Optimized human HOI reference and whole-hand tactile prior. Right (colored panels): Five-frame execution rollout of the simulated dexterous hand. Guided by our tactile supervision, the policy establishes physically grounded contact, effectively avoiding unnatural hovering or penetrations.

functional safe success (SafeTac@3N) falls significantly to 38.72%. This highlights our core philosophy: the safety term is crucial not for minimizing isolated local average errors, but for suppressing momentary destructive behaviors, thereby achieving genuinely reliable, safety-aware manipulation at the task level.

5.3 Qualitative Analysis

Figure 3 showcases the qualitative transfer results of our **TactiSkill** framework evaluated on four representative sequences from the **TactiDex** dataset. For each single-hand and bimanual task, we present the hand-object interaction (HOI) reference alongside a five-frame rollout of the simulated execution.

Guided by our tactile-aware supervision, the dexterous hands successfully establish flush and physically stable contact, effectively avoiding unnatural penetrations or hovering. Specifically, in dynamic single-hand tasks (Rows 1-2), our policy seamlessly adapts its multi-fingered posture to diverse object geometries, maintaining secure grasps during complex translations and simulated actions. Furthermore, in demanding bimanual scenarios (Rows 3-4), TactiSkill exhibits exceptional coordination. Whether performing asymmetrical manipulation or dynamic spatial handovers, the framework

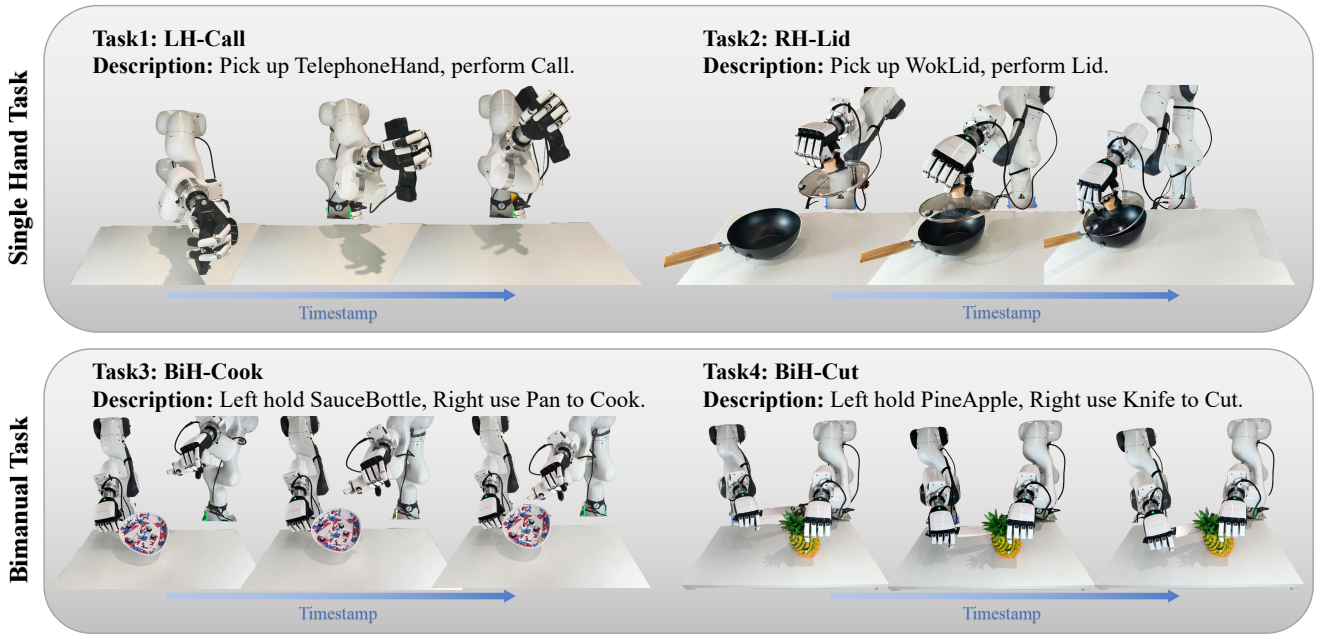

Figure 4: Real-world deployment on our bimanual Franka–Inspire platform. Left: hardware setup with dual Franka arms and Inspire dexterous hands. Right: representative contact-rich execution snapshots.

accurately aligns respective contact patches and regulates interactive forces without dropping the object. Ultimately, this high-fidelity sequential execution underscores the efficacy of our physics-aware tri-component reward in long-horizon environments.

5.4 Real-World Deployment

We deploy our policy on a bimanual platform with two 7-DoF Franka arms and two Inspire dexterous hands. Since the Inspire hand exposes only 6 independent actuators, we retarget the simulated trajectory $\mathbf{q}^{\text{sim}}_{1:T}$ to an actuation trajectory $\tilde{\mathbf{q}}_{1:T} \in \mathbb{R}^6$.

A coarse anchor mapping $\mathbf{q}^{\text{anchor}}_t = \Psi(\mathbf{q}^{\text{sim}}_t)$ provides initialization, and the actuator configuration is expanded to the full kinematic chain via $\Gamma(\cdot)$. The retargeting objective is

$$\mathcal{L} = \mathbb{E}_t \left[\sum_f \|\mathbf{x}^{\text{sim}}_{t,f} - \Phi_f(\Gamma(\tilde{\mathbf{q}}_t))\|_2 + \lambda_a \|\tilde{\mathbf{q}}_t - \mathbf{q}^{\text{anchor}}_t\|_2^2 + \lambda_s \|\tilde{\mathbf{q}}_{t+1} - \tilde{\mathbf{q}}_t\|_2^2 \right] \cdot \tag{7}$$

Arm motions are executed by solving inverse kinematics to track the retargeted wrist pose. Tactile sensing is available on the Inspire hands but is not used as feedback in the current deployment.

6 Conclusion and Discussion

Conclusion In this work, we introduce **TactiDex**, a tactile-rich benchmark designed to advance dexterous manipulation from purely kinematic imitation toward physically grounded, human-like interaction. TactiDex provides synchronized object trajectories, articulated hand motion, and temporally aligned tactile signals, enabling contact-aware evaluation and learning. Building upon this benchmark, we propose **TactiSkill** under the **Tactile-Guided Tri-Component Transfer Framework**, which explicitly regulates contact formation, force magnitude, and temporal alignment. Through comprehensive quantitative and qualitative experiments, we demonstrate that incorporating structured tactile supervision significantly improves object stability, contact fidelity, and manipulation success over purely kinematic transfer. Overall, our results show that tactile guidance is not merely auxiliary, but essential for bridging the gap between motion imitation and physically realistic dexterous manipulation.

Limitations & Future Work Discrepancies in sensor noise, resolution, and physical compliance between simulation and real hardware can degrade the accuracy of learned force interactions when deploying policies directly. To mitigate this, future work will explore on-robot fine-tuning with real tactile feedback, enabling closed-loop adaptation. More broadly, integrating vision, proprioception, and whole-hand tactile sensing within a unified, contact-centric learning framework represents a promising direction toward achieving robust, real-world dexterous manipulation.

Acknowledgments

This work was supported by National Natural Science Foundation of China (62406195, 62303319), HPC Platform of ShanghaiTech University, Core Facility Platform of Computer Science and Communication of ShanghaiTech University, and Key Laboratory of Intelligent Perception and Human-Machine Collaboration (Shanghaitech University), Ministry of Education, and Shanghai Engineering Research Center of Intelligent Vision and Imaging.

References

[1] Sridhar Pandian Arunachalam, Sneha Silwal, Ben Evans, and Lerrel Pinto. 2023. Dexterous imitation made easy: A learning-based framework for efficient dexterous manipulation. In *2023 ieee international conference on robotics and automation (icra)*. IEEE, 5954–5961.

[2] Homanga Bharadhwaj, Jay Vakil, Mohit Sharma, Abhinav Gupta, Shubham Tulsiani, and Vikash Kumar. 2024. Roboagent: Generalization and efficiency in robot manipulation via semantic augmentations and action chunking. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 4788–4795.

[3] Bibit Bianchini, Mathew Halm, and Michael Posa. 2023. Simultaneous learning of contact and continuous dynamics. In *Conference on Robot Learning*. PMLR, 3966–3978.

[4] Yu-Wei Chao, Wei Yang, Yu Xiang, Pavlo Molchanov, Ankur Handa, Jonathan Tremblay, Yashraj S Narang, Karl Van Wyk, Umar Iqbal, Stan Birchfield, et al. 2021. Dexycb: A benchmark for capturing hand grasping of objects. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 9044–9053.

[5] Xingyu Chen, Fu-Jen Chu, Pierre Gleize, Kevin J Liang, Alexander Sax, Hao Tang, Weiyao Wang, Michelle Guo, Thibaut Hardin, Xiang Li, et al. 2025. Sam 3d: 3dfy anything in images. *arXiv preprint arXiv:2511.16624* (2025).

[6] Yuanpei Chen, Yiran Geng, Fangwei Zhong, Jiaming Ji, Jiechuang Jiang, Zongqing Lu, Hao Dong, and Yaodong Yang. 2023. Bi-dexhands: Towards human-level bimanual dexterous manipulation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 46, 5 (2023), 2804–2818.

[7] Yuanpei Chen, Chen Wang, Yaodong Yang, and C Karen Liu. 2024. Object-centric dexterous manipulation from human motion data. *arXiv preprint arXiv:2411.04005* (2024).

[8] Xuxin Cheng, Jialong Li, Shiqi Yang, Ge Yang, and Xiaolong Wang. 2024. Open-television: Teleoperation with immersive active visual feedback. *arXiv preprint arXiv:2407.01512* (2024).

[9] Sudeep Dasari, Abhinav Gupta, and Vikash Kumar. 2023. Learning dexterous manipulation from exemplar object trajectories and pre-grasps. In *2023 IEEE international conference on robotics and automation (ICRA)*. IEEE, 3889–3896.

[10] Zihan Ding, Nathan F Lepora, and Edward Johns. 2020. Sim-to-real transfer for optical tactile sensing. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 1639–1645.

[11] Zihan Ding, Ya-Yen Tsai, Wang Wei Lee, and Bidan Huang. 2021. Sim-to-real transfer for robotic manipulation with tactile sensory. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 6778–6785.

[12] Zicong Fan, Omid Taheri, Dimitrios Tzionas, Muhammed Kocabas, Manuel Kaufmann, Michael J Black, and Otmar Hilliges. 2023. ARCTIC: A dataset for dexterous bimanual hand-object manipulation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 12943–12954.

[13] Hao-Shu Fang, Chenxi Wang, Hongjie Fang, Minghao Gou, Jirong Liu, Hengxu Yan, Wenhai Liu, Yichen Xie, and Cewu Lu. 2023. Anygrasp: Robust and efficient grasp perception in spatial and temporal domains. *IEEE Transactions on Robotics* 39, 5 (2023), 3929–3945.

[14] Harrison Field, Max Yang, Yijiong Lin, Efi Psomopoulou, David Barton, and Nathan F Lepora. 2025. Text2Touch: Tactile In-Hand Manipulation with LLM-Designed Reward Functions. *arXiv preprint arXiv:2509.07445* (2025).

[15] Letian Fu, Huang Huang, Lars Berscheid, Hui Li, Ken Goldberg, and Sachin Chitta. 2023. Safe self-supervised learning in real of visuo-tactile feedback policies for industrial insertion. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 10380–10386.

[16] Beining Han, Abhishek Joshi, and Jia Deng. 2025. Zero-shot Sim2Real Transfer for Magnet-Based Tactile Sensor on Insertion Tasks. *arXiv preprint arXiv:2505.02915* (2025).

[17] Johanna Hansen, Francois Hogan, Dmitriy Rivkin, David Meger, Michael Jenkin, and Gregory Dudek. 2022. Visuotactile-rl: Learning multimodal manipulation policies with deep reinforcement learning. In *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 8298–8304.

[18] Yana Hasson, Gul Varol, Dimitrios Tzionas, Igor Kalevatykh, Michael J Black, Ivan Laptev, and Cordelia Schmid. 2019. Learning joint reconstruction of hands and manipulated objects. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 11807–11816.

[19] Ryan Hoque, Peide Huang, David J Yoon, Mouli Sivapurapu, and Jian Zhang. 2025. Egodex: Learning dexterous manipulation from large-scale egocentric video. *arXiv preprint arXiv:2505.11709* (2025).

[20] Wenbin Hu, Bidan Huang, Wang Wei Lee, Sicheng Yang, Yu Zheng, and Zhibin Li. 2025. Dexterous in-hand manipulation of slender cylindrical objects through deep reinforcement learning with tactile sensing. *Robotics and Autonomous Systems* 186 (2025), 104904.

[21] Binghao Huang, Yixuan Wang, Xinyi Yang, Yiyue Luo, and Yunzhu Li. 2024. 3d-vitac: Learning fine-grained manipulation with visuo-tactile sensing. *arXiv preprint arXiv:2410.24091* (2024).

[22] Yitaek Kim, Casper Hewson Rask, and Christoffer Sloth. 2025. Tac2Motion: Contact-Aware Reinforcement Learning with Tactile Feedback for Robotic Hand Manipulation. *arXiv preprint arXiv:2509.17812* (2025).

[23] Taein Kwon, Bugra Tekin, Jan Stühmer, Federica Bogo, and Marc Pollefeys. 2021. H2o: Two hands manipulating objects for first person interaction recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*. 10138–10148.

[24] Arjun S Lakshmipathy, Jessica K Hodgins, and Nancy S Pollard. 2025. Kinematic motion retargeting for contact-rich anthropomorphic manipulations. *ACM Transactions on Graphics* 44, 2 (2025), 1–20.

[25] Jiaman Li, Jiajun Wu, and C Karen Liu. 2023. Object motion guided human motion synthesis. *ACM Transactions on Graphics (TOG)* 42, 6 (2023), 1–11.

[26] Kailin Li, Puhao Li, Tengyu Liu, Yuyang Li, and Siyuan Huang. 2025. Maniptrans: Efficient dexterous bimanual manipulation transfer via residual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6991–7003.

[27] Sizhe Li, Zhiao Huang, Tao Chen, Tao Du, Hao Su, Joshua B Tenenbaum, and Chuang Gan. 2023. Dexdeform: Dexterous deformable object manipulation with human demonstrations and differentiable physics. *arXiv preprint arXiv:2304.03223* (2023).

[28] Yuchao Li, Ziqi Jin, Jin Liu, and Daolin Ma. 2025. Visuo-tactile feedback policies for terminal assembly facilitated by reinforcement learning. *Frontiers in Robotics and AI* 12 (2025), 1660244.

[29] Wenyu Liang, Fen Fang, Cihan Acar, Wei Qi Toh, Ying Sun, Qianli Xu, and Yan Wu. 2022. Visuo-tactile manipulation planning using reinforcement learning with affordance representation. *arXiv preprint arXiv:2207.06608* (2022).

[30] Yijiong Lin, Alex Church, Max Yang, Haoran Li, John Lloyd, Dandan Zhang, and Nathan F Lepora. 2023. Bi-touch: Bimanual tactile manipulation with sim-to-real deep reinforcement learning. *IEEE Robotics and Automation Letters* 8, 9 (2023), 5472–5479.

[31] Qingtao Liu, Yu Cui, Zhengnan Sun, Gaofeng Li, Jiming Chen, and Qi Ye. 2025. VTDexmanip: A dataset and benchmark for visual-tactile pretraining and dexterous manipulation with reinforcement learning. In *The Thirteenth International Conference on Learning Representations*.

[32] Xueyi Liu, Kangbo Lyu, Jieqiong Zhang, Tao Du, and Li Yi. 2024. Parameterized quasi-physical simulators for dexterous manipulations transfer. In *European Conference on Computer Vision*. Springer, 164–182.

[33] Yunze Liu, Yun Liu, Che Jiang, Kangbo Lyu, Weikang Wan, Hao Shen, Boqiang Liang, Zhoujie Fu, He Wang, and Li Yi. 2022. Hoi4d: A 4d egocentric dataset for category-level human-object interaction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 21013–21022.

[34] Yun Liu, Haolin Yang, Xu Si, Ling Liu, Zipeng Li, Yuxiang Zhang, Yebin Liu, and Li Yi. 2024. Taco: Benchmarking generalizable bimanual tool-action-object understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 21740–21751.

[35] Viktor Makoviychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, et al. 2021. Isaac gym: High performance gpu-based physics simulation for robot learning. *arXiv preprint arXiv:2108.10470* (2021).

[36] Elle Miller, Trevor McInroe, David Abel, Oisin Mac Aodha, and Sethu Vijayakumar. 2025. Enhancing Tactile-based Reinforcement Learning for Robotic Control. *arXiv preprint arXiv:2510.21609* (2025).

[37] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. 2019. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 165–174.

[38] Haozhi Qi, Ashish Kumar, Roberto Calandra, Yi Ma, and Jitendra Malik. 2023. In-hand object rotation via rapid motor adaptation. In *Conference on Robot Learning*. PMLR, 1722–1732.

[39] Yuzhe Qin, Yueh-Hua Wu, Shaowei Liu, Hanwen Jiang, Ruihan Yang, Yang Fu, and Xiaolong Wang. 2022. Dexmv: Imitation learning for dexterous manipulation from human videos. In *European Conference on Computer Vision*. Springer, 570–587.

[40] Yuzhe Qin, Wei Yang, Binghao Huang, Karl Van Wyk, Hao Su, Xiaolong Wang, Yu-Wei Chao, and Dieter Fox. 2023. Anyteleop: A general vision-based dexterous robot arm-hand teleoperation system. *arXiv preprint arXiv:2307.04577* (2023).

[41] Javier Romero, Dimitrios Tzionas, and Michael J Black. 2022. Embodied hands: Modeling and capturing hands and bodies together. *arXiv preprint arXiv:2201.02610* (2022).

[42] Yuxin Ray Song, Jinzhou Li, Rao Fu, Devin Murphy, Kaichen Zhou, Rishi Shiv, Yaqi Li, Haoyu Xiong, Crystal Elaine Owens, Yilun Du, et al. 2025. OPEN-TOUCH: Bringing Full-Hand Touch to Real-World Interaction. *arXiv preprint arXiv:2512.16842* (2025).

[43] Entong Su, Chengzhe Jia, Yuzhe Qin, Wenxuan Zhou, Annabella Macaluso, Binghao Huang, and Xiaolong Wang. 2024. Sim2real manipulation on unknown objects with tactile-based reinforcement learning. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 9234–9241.

[44] Omid Taheri, Nima Ghorbani, Michael J Black, and Dimitrios Tzionas. 2020. GRAB: A dataset of whole-body human grasping of objects. In *European conference on computer vision*. Springer, 581–600.

[45] Jiaxian Tang, Xiaogang Yuan, and Shaodong Li. 2025. Visual–Tactile Fusion and SAC-Based Learning for Robot Peg-in-Hole Assembly in Uncertain Environments. *Machines* 13, 7 (2025), 605.

[46] Jie Tian, Ran Ji, Lingxiao Yang, Suting Ni, Yuexin Ma, Lan Xu, Jingyi Yu, Ye Shi, and Jingya Wang. 2024. Gaze-guided hand-object interaction synthesis: Dataset and method. *arXiv preprint arXiv:2403.16169* (2024).

[47] Chen Wang, Haochen Shi, Weizhuo Wang, Ruohan Zhang, Li Fei-Fei, and C Karen Liu. 2024. Dexcap: Scalable and portable mocap data collection system for dexterous manipulation. *arXiv preprint arXiv:2403.07788* (2024).

[48] Ruicheng Wang, Jialiang Zhang, Jiayi Chen, Yinzhen Xu, Puhao Li, Tengyu Liu, and He Wang. 2022. Dexgraspnet: A large-scale robotic dexterous grasp dataset for general objects based on simulation. *arXiv preprint arXiv:2210.02697* (2022).

[49] Youzhuo Wang, Jiayi Ye, Chuyang Xiao, Yiming Zhong, Heng Tao, Hang Yu, Yumeng Liu, Jingyi Yu, and Yuexin Ma. 2025. DexH2R: A benchmark for dynamic dexterous grasping in human-to-robot handover. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 12702–12712.

[50] Bowen Wen, Jonathan Tremblay, Valts Blukis, Stephen Tyree, Thomas Müller, Alex Evans, Dieter Fox, Jan Kautz, and Stan Birchfield. 2023. Bundlesdf: Neural 6-dof tracking and 3d reconstruction of unknown objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 606–617.

[51] Bowen Wen, Wei Yang, Jan Kautz, and Stan Birchfield. 2024. Foundationpose: Unified 6d pose estimation and tracking of novel objects. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 17868–17879.

[52] Ziwei Xia, Zhen Deng, Bin Fang, Yiyong Yang, and Fuchun Sun. 2022. A review on sensory perception for dexterous robotic manipulation. *International Journal of Advanced Robotic Systems* 19, 2 (2022), 17298806221095974.

[53] Chendong Xin, Mingrui Yu, Yongpeng Jiang, Zhefeng Zhang, and Xiang Li. 2026. Analyzing Key Objectives in Human-to-Robot Retargeting for Dexterous Manipulation. *IEEE Robotics and Automation Practice* (2026).

[54] Yinzhen Xu, Weikang Wan, Jialiang Zhang, Haoran Liu, Zikang Shan, Hao Shen, Ruicheng Wang, Haoran Geng, Yijia Weng, Jiayi Chen, et al. 2023. 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*. 4737–4746.

[55] Lixin Yang, Kailin Li, Xinyu Zhan, Fei Wu, Anran Xu, Liu Liu, and Cewu Lu. 2022. Oakink: A large-scale knowledge repository for understanding hand-object interaction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 20953–20962.

[56] Jessica Yin, Haozhi Qi, Youngsun Wi, Sayantan Kundu, Mike Lambeta, William Yang, Changhao Wang, Tingfan Wu, Jitendra Malik, and Tess Hellebrekers. 2025. OSMO: Open-Source Tactile Glove for Human-to-Robot Skill Transfer. *arXiv preprint arXiv:2512.08920* (2025).

[57] Kelin Yu, Yunhai Han, Qixian Wang, Vaibhav Saxena, Danfei Xu, and Ye Zhao. 2023. Mimictouch: Leveraging multi-modal human tactile demonstrations for contact-rich manipulation. *arXiv preprint arXiv:2310.16917* (2023).

[58] Xinyu Zhan, Lixin Yang, Yifei Zhao, Kangrui Mao, Hanlin Xu, Zenan Lin, Kailin Li, and Cewu Lu. 2024. Oakink2: A dataset of bimanual hands-object manipulation in complex task completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 445–456.

[59] Shenao Zhang, Wanxin Jin, and Zhaoran Wang. 2023. Adaptive barrier smoothing for first-order policy gradient with contact dynamics. In *International Conference on Machine Learning*. PMLR, 41219–41243.

[60] Zifan Zhao, Siddhant Haldar, Jinda Cui, Lerrel Pinto, and Raunaq Bhirangi. 2025. Touch begins where vision ends: Generalizable policies for contact-rich manipulation. *arXiv preprint arXiv:2506.13762* (2025).

A TactiDex Details

A.1 Data Scale and Inventory

The TactiDex dataset features over 700 meticulously recorded interaction sequences, specifically targeting human-to-robot dexterous manipulation. To ensure comprehensive coverage of contact-rich scenarios, we curated a diverse set of 49 everyday objects spanning multiple categories.

As illustrated in Fig. 5, these objects encompass various geometries, sizes, and physical properties. Rather than being limited to simple rigid bodies, our inventory ranges from delicate items (e.g., eggs) and everyday containers to complex tools and matching multi-part accessories (e.g., a grinder with a rod, or a wok with a lid). This strategic selection explicitly supports the execution of intricate, long-horizon object-object interactions.

Furthermore, to provide a holistic view of the interaction diversity, Fig. 6 visualizes the macroscopic distribution of our recorded sequences. Following previous large-scale interaction datasets, we employ a Sankey diagram to map the explicit hierarchical relationships between task divisions (i.e., single-hand versus bimanual tasks), fine-grained action types (e.g., functional simulation, bimanual handovers, object-object interactions), and specific object categories.

A.2 Hardware Setup and Capture System

To capture high-fidelity, synchronized human demonstrations, we constructed a custom multi-modal capture arena that seamlessly integrates global motion capture, fine-grained hand kinematics, and whole-hand tactile sensing.

Optical Motion Capture System. The global tracking arena is powered by an OptiTrack system equipped with 8 PrimeX 13W (PX13W) cameras, operating at a high sampling rate of 120 fps. Prior to data collection, the capture volume is meticulously calibrated using a specialized vendor-provided active wand. This rigorous calibration process achieves an exceptional mean ray error of exactly 0.2 mm, ensuring sub-millimeter precision for recording both hand trajectories and object 6D poses.

Kinematic and Tactile Dual-Glove Setup. For fine-grained finger kinematics, human demonstrators wear MANUS Metagloves Pro, renowned for their high-precision and low-latency joint tracking. To guarantee anatomical accuracy across different subjects, a standardized calibration routine is strictly performed before every single recording session.

Simultaneously, the user wears our high-resolution tactile glove on the exterior to capture continuous pressure maps. Integrating these two state-of-the-art systems presented a unique hardware challenge: tracking the Metagloves within the OptiTrack ecosystem requires a rigid marker backplate. To accommodate the outer tactile glove without damaging its delicate internal wiring or compromising sensor fidelity, we engineered a custom, non-destructive mechanical modification for the backplate integration (as illustrated in Fig. 7).

Object Registration and Migration. Beyond capturing raw trajectories, we ensure all object interactions are grounded in a unified physical space. Once the 3D meshes of the 49 objects are generated, we manually align their scale, translation, and rotation within the multi-view OptiTrack coordinate system. This standardized spatial registration process guarantees that the recorded 6D object poses are inherently aligned with the mesh geometry, allowing the entire dataset to be easily migrated and correctly rendered on any new computational platform or simulation environment.

A.3 Standard Operating Procedure (SOP) and Collection Pipeline

Given the unprecedented scale of the TactiDex dataset, relying on manual recording protocols for 700+ multi-modal sequences would be highly susceptible to inconsistencies and human error.

To systematically standardize our collection process, we designed and deployed a dedicated web-based data collection system. This integrated web platform serves as the central command hub during demonstrations. It dynamically displays the required task instructions to the human subject, tracks the real-time progress of the task inventory, and logs metadata for each recorded trial. By formalizing the entire workflow through this automated SOP interface, we significantly reduced overhead and ensured that task executions remain highly consistent across the entire dataset.

B TactiSkill Implementation

B.1 Network Architecture and State-Action Space

Our TactiSkill framework is implemented using Proximal Policy Optimization (PPO) with an asymmetric Actor-Critic architecture. Both networks are modeled as multi-layer perceptrons (MLPs) featuring hidden dimensions of $[512, 256, 128]$ with ELU activations.

To bridge the sim-to-real gap, the state space is rigorously decoupled. Taking the 12-DoF Inspire hand as our primary single-hand example, the observation space consists of three distinct components (Note: for bimanual tasks, all dimensions are exactly doubled):

1. Proprioceptive State ($S_{prop} \in \mathbb{R}^{46}$): Contains the internal kinematic status of the robotic hand.

- **Joint States (36-dim):** Current joint positions $\mathbf{q} \in \mathbb{R}^{12}$, and their trigonometric encodings $\sin(\mathbf{q}), \cos(\mathbf{q}) \in \mathbb{R}^{24}$ to ensure continuity in rotational space.
- **Wrist Base State (10-dim):** The 6D pose (quaternion) and linear/angular velocities of the wrist base, excluding absolute global positioning to promote spatial invariance.

2. Target Reference ($S_{target} \in \mathbb{R}^{330}$): Provides dense spatial and temporal cues from the human demonstration to guide the imitation process.

- **Tactile & Geometry Prior (133-dim):** A Basis Point Set (BPS) encoding of the object's point cloud (128-dim) coupled with the ground-truth target tactile distances for the fingertips (5-dim).
- **Future Kinematic Trajectory (197-dim):** The delta kinematics (positions, rotations, linear/angular velocities) over the future horizon for the wrist (23-dim), finger joints (135-dim), manipulated object (23-dim), and the relative object-to-joint transformations (16-dim).

3. Privileged Information ($S_{priv} \in \mathbb{R}^{49}$): Accessible exclusively to the Critic during simulation training to accurately estimate the value function.

Figure 5: Object Inventory of TactiDex. Rendered meshes of the 49 diverse everyday objects utilized in our data collection. The collection spans a wide range of geometries, physical scales, and functional categories. Notably, our inventory intentionally includes matching multi-part object pairs (e.g., teacups with lids, grinders with rods) to explicitly facilitate complex, long-horizon bimanual interactions and tool use.

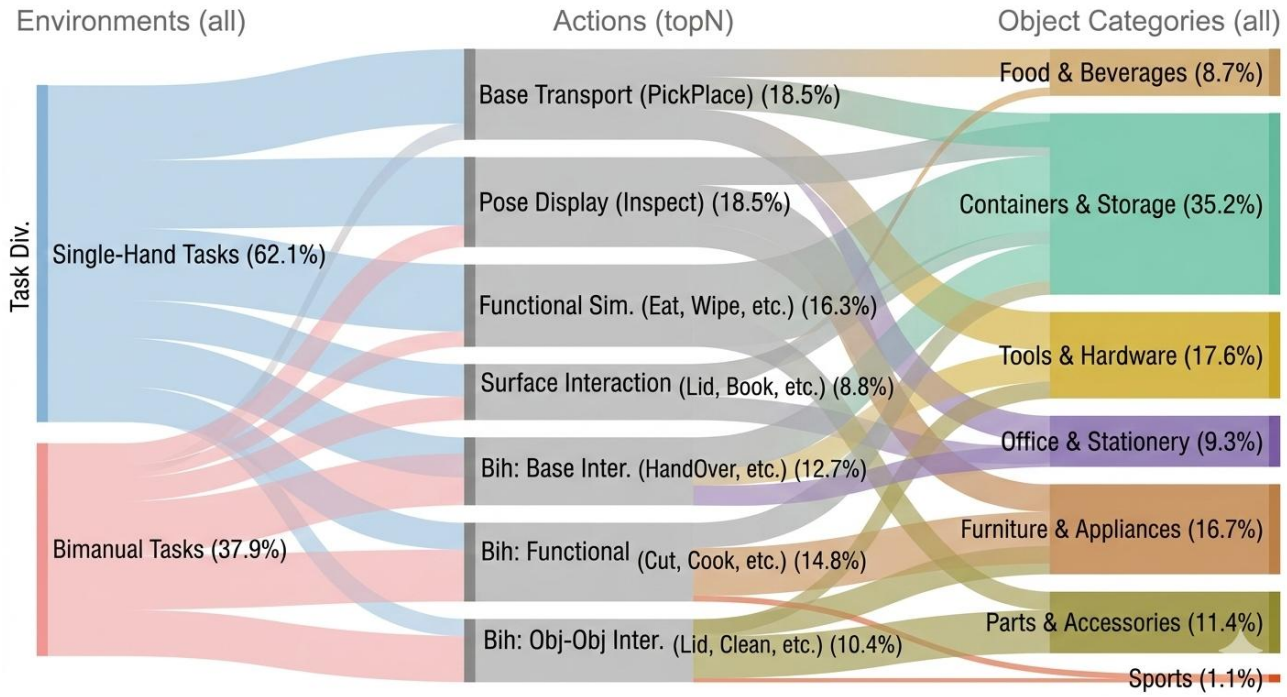

Figure 6: Task and Interaction Distribution. A Sankey diagram illustrating the hierarchical mapping among task actions, semantic grasp types, and object categories across the 700+ sequences in the TactiDex dataset.

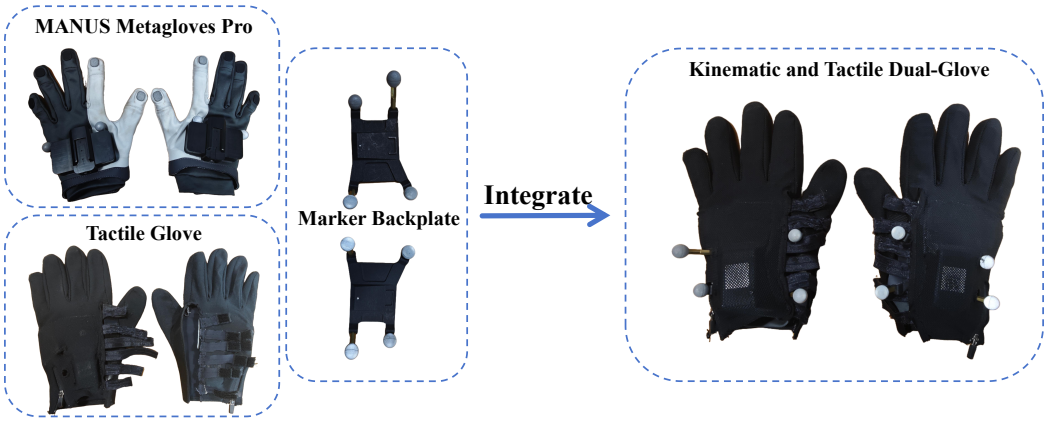

Figure 7: Custom Dual-Glove Hardware Integration. To synchronously capture high-precision kinematics and tactile signals, we developed a customized backplate modification. This design safely routes the tracking markers without interfering with the dense wiring matrix of the outer tactile glove.

- **Exact Object Dynamics (14-dim):** Ground-truth Cartesian position (3), quaternion (4), velocities (6), and object weight (1).
- **High-Fidelity Contact Force (20-dim):** Simulated fingertip contact forces $\mathbf{F}^{sim} \in \mathbb{R}^{20}$, explicitly capturing both the 3D directional force vectors and their scalar magnitudes across all 5 fingertips.
- **Internal Physics (15-dim):** Joint velocities $\dot{\mathbf{q}} \in \mathbb{R}^{12}$ and the object's exact Center of Mass (CoM) coordinates (3-dim).

Network Inputs & Action Space: During training, the Actor network observes only the deployable real-world states: $\mathcal{S}_{actor} = S_{prop} \oplus S_{target} \in \mathbb{R}^{376}$ (or 752 for bimanual). The Critic network receives the full state including privileged dynamics: $\mathcal{S}_{critic} = S_{prop} \oplus S_{target} \oplus S_{priv} \in \mathbb{R}^{425}$ (or 850 for bimanual). The action space $\mathcal{A} \in \mathbb{R}^{12}$ outputs the residual target joint positions, which are executed via a low-level PD controller.

B.2 Hyperparameters and Reward Weights

Table 3 comprehensively lists the hyperparameters used for training the residual policy, including the specific weights for our tactile-guided tri-component reward.

The simulation environment is built upon NVIDIA Isaac Gym with a PhysX backend.

- **Physics Parameters:** The simulation operates at 60 Hz with 2 sub-steps per frame. To ensure stable grasping and realistic interaction dynamics, we set the lateral friction coefficient of both the fingertips and objects to 1.0, and the joint damping to 0.5 Ns/m.
- **Tactile Signal Mapping:** To bridge the gap between human demonstrations and the simulation engine, raw ADC signals from the physical tactile sensors are mapped to simulated normal forces using a linear calibration model: $F_{sim} = k \cdot (ADC_{raw} - ADC_{offset})$, where k is determined via a calibrated force gauge. Within the PhysX engine, this ground-truth force corresponds to the rigid-body contact impulse divided by the simulation time step.

C Real-World Hardware Deployment Details

While Section 5 in the main text outlines our retargeting objective for zero-shot sim-to-real transfer, this section provides the low-level engineering and communication details required to reproduce our bimanual hardware experiments.

C.1 Bimanual System Architecture and Arm Control

Our physical deployment platform consists of two 7-DoF Franka Emika Panda arms and two Inspire dexterous hands. The entire dual-arm control framework is orchestrated using the Deoxys control interface.

To mitigate the sim-to-real global coordinate gap, the Franka arms track relative delta-poses rather than absolute Cartesian coordinates. Before execution, the initial pose of the simulated trajectory is aligned with the physical robot's starting state. At each timestep, the delta position and relative rotation (computed via quaternion transformations) are applied to the physical anchor pose. The arms are then driven by an Operational Space Controller (OSC) operating at a 100 Hz control loop ($\Delta t = 0.01$ s), which converts Cartesian tracking errors into executable joint torques.

C.2 Kinematic Optimization for the Inspire Hand

The simulated hand possesses 12 DoF, whereas the physical Inspire hand is underactuated, driven by only 6 independent linear actuators: Index flexion, Middle flexion, Ring flexion, Pinky flexion, Thumb pitch, and Thumb yaw.

To execute the simulated trajectory $\mathbf{q}_{1:T}^{\text{sim}}$ on the physical hand, we employ an offline optimization-based retargeting pipeline. We utilize *PyTorch Kinematics* to construct the differentiable forward kinematics chain of the Inspire hand. For each trajectory, we initialize a 6-DoF command sequence $\tilde{\mathbf{q}}_{1:T}$ using a heuristic anchor mapping. We then utilize the Adam optimizer (learning rate = 0.01,

Table 3: Key Hyperparameters and Reward Coefficients for TactiSkill Training

Category	Parameter	Value
PPO Algorithm		
	Learning Rate	3×10^{-4}
	Batch Size (Total)	32, 768
	Mini-batch Size	8, 192
	Optimization Epochs	5
	Discount Factor (γ)	0.99
	GAE Parameter (λ)	0.95
Tactile Reward		
	Guidance Weight (w_g)	0.6
	Alignment Weight (w_a)	1.4
	Safety Penalty Weight (w_s)	0.01
	Contact Threshold (τ)	0.3 N
	Safety Force Limit (F_{limit})	40.0 N
Kinematic & Task Reward		
	Object Tracking Weight (Pos / Rot)	5.0 / 3.0
	Fingertip Tracking Weight (Thumb / Others)	0.9 / 0.6∼0.8
	Joint Tracking Weight (Level 1 / Level 2)	0.5 / 0.3
	Energy Penalty & Velocity Constraints	0.05∼0.5

100 iterations) to minimize the following composite loss:

$$\mathcal{L} = W_{\text{pos}}\mathcal{L}_{\text{pos}} + W_{\text{anchor}}\mathcal{L}_{\text{anchor}} + W_{\text{smooth}}\mathcal{L}_{\text{smooth}} \tag{8}$$

Here, $\mathcal{L}_{\text{pos}}$ penalizes the Cartesian distance between the 12-DoF simulated fingertips and the 6-DoF retargeted fingertips. $\mathcal{L}_{\text{anchor}}$ prevents the solution from drifting too far from the anchor prior, and $\mathcal{L}_{\text{smooth}}$ penalizes temporal jitter. Based on empirical tuning, the weights are set to $W_{\text{pos}} = 50.0$, $W_{\text{anchor}} = 20.0$, and $W_{\text{smooth}} = 10.0$.

C.3 Command Mapping and Open-Loop Execution

Once optimized, the 6-DoF joint radians are calibrated and linearly mapped to discrete actuation commands $\in [0, 1000]$. These commands are continuously dispatched to drive the Inspire hand via a 115200 baud serial interface.

Crucially, while the physical Inspire hands are equipped with tactile sensors, our current deployment relies exclusively on open-loop execution (i.e., the tactile signals are not fed back into the policy online). The successful sim-to-real transfer demonstrates a core advantage of our TactiSkill framework: by actively regulating simulated contact forces via the tri-component reward during training, the policy intrinsically learns robust, physics-aware force distributions. Consequently, the resulting zero-shot kinematic trajectories naturally establish stable, physically grounded grasps on the real hardware, successfully overcoming typical geometric domain gaps without requiring fragile online tactile feedback loops.

D Extended Qualitative Results and Video Guide

D.1 Visualization of the TactiDex Dataset

Before evaluating the learned policies, we first visualize the rich, multi-modal human demonstrations captured in our TactiDex dataset. As the foundation of our tactile-guided reinforcement learning framework, the dataset goes beyond standard kinematic tracking by synchronously recording high-resolution contact forces.

As illustrated in Fig. 8, we present representative sequences across diverse manipulation scenarios, including left-hand (LH), right-hand (RH), and long-horizon bimanual (BiH) tasks. Alongside the 3D kinematic hand-object spatial relationships, we visualize the corresponding human contact heat maps at each timestep (mapped to a purple-to-yellow color scale). These dynamic pressure distributions explicitly reveal the underlying physical intent of the demonstrator. For instance, the heat maps clearly distinguish between the localized fingertip forces required to pinch an apple or hold a telephone, versus the broader palm-level pressure required to firmly press a blackboard eraser. By capturing these nuanced, physics-grounded interactions, the TactiDex dataset provides the essential dense tactile priors that our TactiSkill framework leverages during policy training.

D.2 Qualitative Comparison of Interaction Quality

To visually evaluate the impact of our proposed tri-component tactile reward, we provide a localized, frame-by-frame comparison between our *TactiSkill* policy and a standard kinematic-only imitation baseline. As illustrated in Fig. 9, we examine representative left-hand (LH), right-hand (RH), and bimanual (BiH) tasks along their execution timestamps.

Without the explicit guidance of tactile priors, the baseline policy relies entirely on minimizing joint-space discrepancies. This purely kinematic approach frequently suffers from critical physical inconsistencies during complex interactions, as highlighted in the magnified red regions:

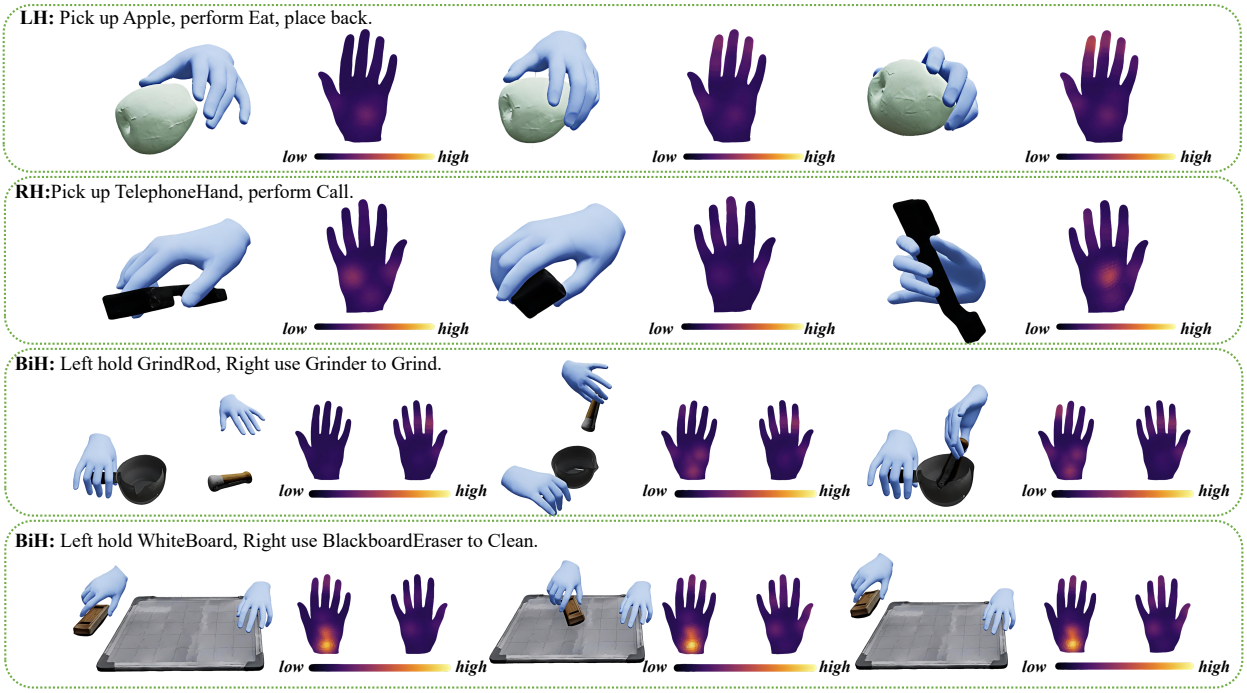

Figure 8: Multi-modal Human Demonstrations from the TactiDex Dataset. Representative sequences for LH (Apple), RH (TelephoneHand), and BiH (Grinder, WhiteBoard) tasks. Each frame pairs the 3D kinematic hand-object interaction with its synchronized tactile pressure heat map (where purple indicates low pressure and yellow indicates high pressure). These sequences demonstrate how our dataset captures the precise, task-specific force distributions required for complex object manipulation.

- **Kinematic-only Baseline:** The policy fails to respect physical contact boundaries. This results in obvious *unnatural hovering* (e.g., fingers leaving large gaps around the telephone handle or paddle) or *severe mesh penetration* (e.g., fingers aggressively clipping through the kettle handle to artificially secure a grasp).
- **TactiSkill (Ours):** By actively regularizing simulated contact forces during the RL process, our policy ensures the robotic fingertips establish *flush, physically grounded contacts* (highlighted in the magnified green regions). Our tactile-guided policy gracefully conforms to the object's surface boundaries, maintaining stable and realistic grasps across all manipulation stages.

D.3 Supplementary Video Guide

The dynamic performance, temporal stability, and zero-shot real-world deployment of our framework are best appreciated in motion. We strongly encourage the reviewers to refer to the accompanying **Supplementary Video** for a comprehensive evaluation of the uncurated policy rollouts and hardware experiments.

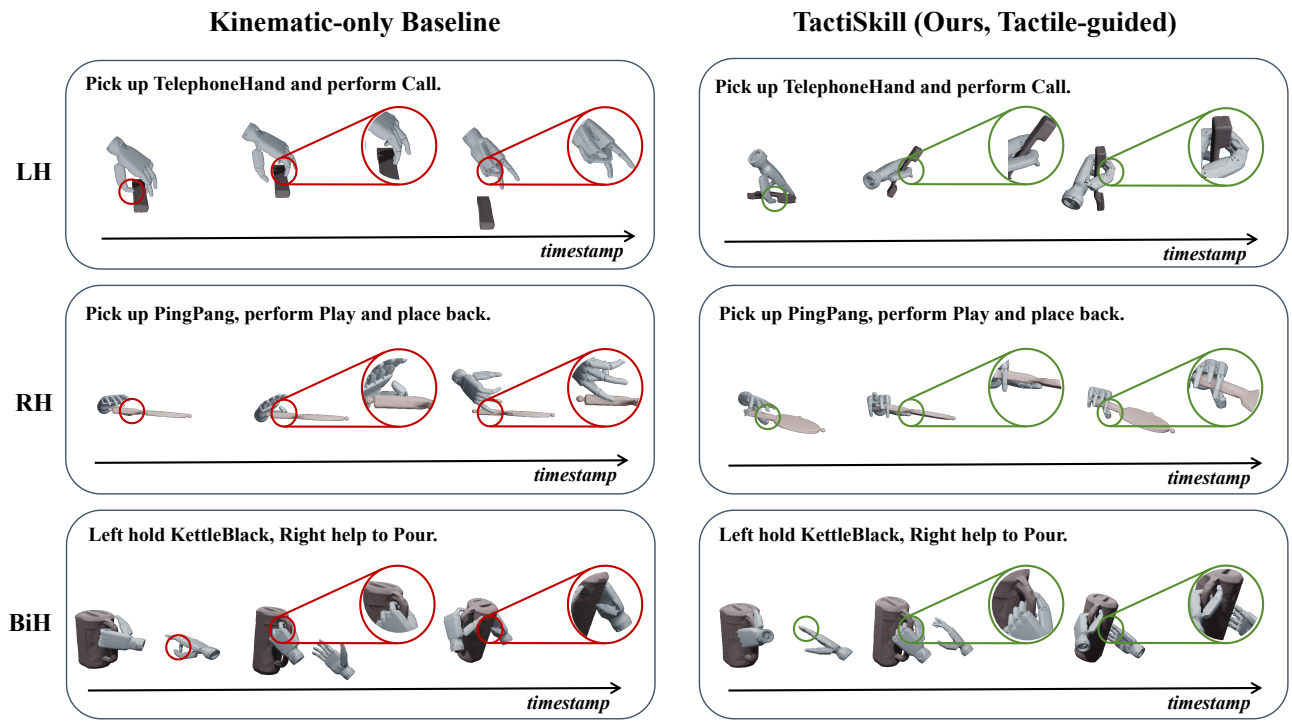

Figure 9: Qualitative Comparison of Contact Quality. Frame-by-frame rollouts of LH (TelephoneHand), RH (PingPang), and BiH (KettleBlack) tasks. (Left) The kinematic-only baseline often results in unnatural hovering or severe mesh penetration, as explicitly highlighted by the red magnified views. (Right) Our TactiSkill policy seamlessly establishes flush, stable contact with diverse object geometries, effectively preventing physical artifacts as shown in the green magnified views.