

FlowHOI: Flow-based Semantics-Grounded Generation of Hand-Object Interactions for Dexterous Robot Manipulation

Huajian Zeng¹, Lingyun Chen², Jiaqi Yang¹, Yuantai Zhang¹, Fan Shi³, Peidong Liu⁴, Xingxing Zuo^{*,1}

¹MBZUAI, ²TUM, ³NUS, ⁴Westlake University

*Corresponding author

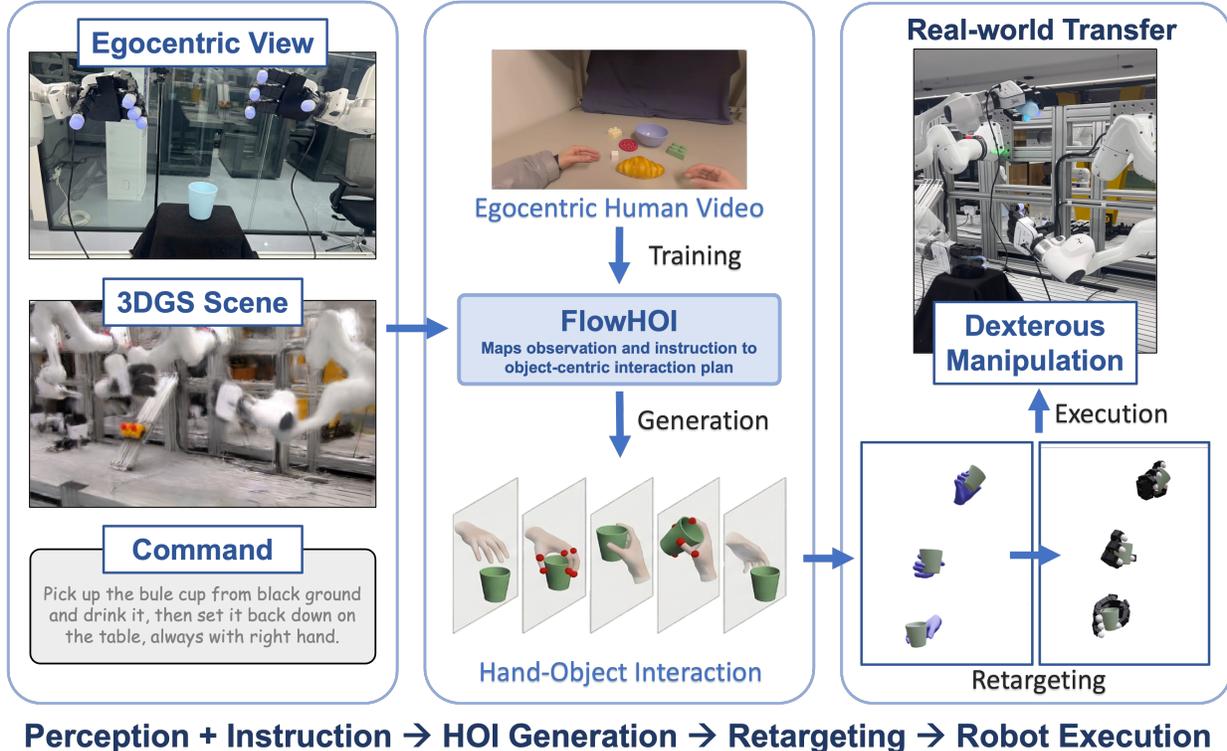


Fig. 1: We present a framework for generating hand-object interaction (HOI) motions conditioned on egocentric observation, text command, and 3D scene context. Our method extracts HOI data from large-scale egocentric videos to learn a grasping prior, then generates manipulation trajectories that respect scene constraints and language instructions. The generated HOI sequence can be retargeted to robot hands for execution in the real world.

Abstract—Recent vision-language-action (VLA) models can generate plausible end-effector motions, yet they often fail in long-horizon, contact-rich tasks because the underlying hand-object interaction (HOI) structure is not explicitly represented. An embodiment-agnostic interaction representation that captures this structure would make manipulation behaviors easier to validate and transfer across robots. We propose FlowHOI, a two-stage flow-matching framework that generates semantically grounded, temporally coherent HOI sequences, comprising hand poses, object poses, and hand-object contact states, conditioned on an egocentric observation, a language instruction, and a 3D Gaussian splatting (3DGS) scene reconstruction. We decouple geometry-centric grasping from semantics-centric manipulation, conditioning the latter on compact 3D scene tokens and a motion-text alignment loss to semantically ground the generated interactions in both the physical scene layout and the language

instruction. To address the scarcity of high-fidelity HOI supervision, we introduce a reconstruction pipeline that recovers aligned hand-object trajectories and meshes from large-scale egocentric videos, yielding an HOI prior for robust generation. Across the GRAB and HOT3D benchmarks, FlowHOI achieves the highest action recognition accuracy and a $1.7\times$ higher physics simulation success rate than the strongest diffusion-based baseline, while delivering a $40\times$ inference speedup. We further demonstrate real-robot execution on four dexterous manipulation tasks, illustrating the feasibility of retargeting generated HOI representations to real-robot execution pipelines. Code and dataset will be released upon paper acceptance.

I. INTRODUCTION

Robotic manipulation in everyday household environments is fundamentally interaction-centric [27, 73]: task success

depends on how the robot establishes and maintains interaction with a target object under clutter, contact constraints, and task semantics. Common manipulation tasks such as opening a lid, pouring from a container, or placing an object require complex hand-object interactions that unfold over time [9]. These interactions are not fully characterized by the robot’s end-effector trajectory alone, but by the underlying context of interaction: where contact occurs given the surrounding scene geometry, how stable contact is achieved and preserved, how the object’s pose or state evolves under interaction, and whether this evolution is semantically consistent with the intended language instruction and the scene affordances.

Recent vision-language-action (VLA) models [49, 10, 101], fine-tuned from large VLMs [4, 80, 6, 59], generate plausible end-effector trajectories but struggle with contact-rich, long-horizon dexterous tasks [110]. This motivates a semantically grounded, embodiment-agnostic intermediate representation [76, 102] that explicitly encodes how contact is established and transitioned to induce language-specified state changes [41], decoupling manipulation from robot-specific control and facilitating transfer across embodiments.

In this work, we propose FlowHOI, a two-stage flow-matching framework for producing semantically grounded hand-object interaction (HOI) sequences. Given an initial egocentric observation, a language instruction, and a 3D Gaussian splatting (3DGS) scene reconstruction, FlowHOI generates temporally coherent and physically plausible HOI sequences comprising hand poses, object poses, and hand-object contact states, all anchored in the observed scene context and aligned with the language instruction. The resulting HOI representation provides a natural interface for enforcing geometric and physical constraints and can be retargeted to downstream robotic dexterous manipulation, leading to improved physical plausibility and robustness [38, 67].

Building such a model raises three central challenges: **(i) Geometric stability vs. semantic grounding.** Long-horizon interactions must maintain contact stability while producing motions that are semantically grounded in both the language instruction and the surrounding 3D scene; naively entangling these two objectives causes the model to compromise between them, resulting in contact drift or semantically inconsistent motions. Inspired by how humans first establish a stable grasp before manipulating objects [43], we decompose generation into a geometry-centric *Grasping* stage and a semantics-centric *Manipulation* stage. The Grasping stage leverages a pretrained grasping prior to produce contact-stable initializations; the Manipulation stage conditions on compact 3D scene tokens extracted from the reconstructed scene and employs a motion-text alignment loss, grounding the generated object state changes in both the physical scene layout and the language instruction. **(ii) Inference efficiency.** Diffusion-based HOI generators [18, 54, 36, 94] require tens to hundreds of denoising steps, taking 3–7 s per sequence, which is prohibitive for real-time downstream planning. We adopt conditional flow matching, reducing inference to 0.16 s per sequence (up to 40× speedup) while maintaining temporally coherent and

physically plausible generation. **(iii) Data scarcity.** High-fidelity HOI supervision is scarce because hand-object interactions are high-dimensional, frequently occluded, and tightly coupled with contact dynamics [25, 104]. We address this by introducing a reconstruction pipeline that recovers aligned hand-object trajectories and meshes from large-scale egocentric videos [37]. The resulting dataset enables the learning of robust HOI priors with strong generalization across objects and tasks.

To the best of our knowledge, FlowHOI is the *first* method to formulate HOI generation as a unified, conditional flow-matching process in a semantics-grounded way. In summary, our contributions are:

- We introduce a two-stage HOI generation framework that decouples geometry-centric grasping from semantics-centric manipulation and employs flow matching for efficient generation, achieving up to 40× speedup over prior diffusion-based methods.
- We design a reconstruction pipeline to extract large-scale, high-fidelity HOI data from egocentric videos, enabling the learning of a robust HOI prior that improves generalization across objects and tasks.
- On GRAB and HOT3D benchmarks, FlowHOI achieves the highest action-recognition accuracy and a 1.7× higher physics-simulation success rate (55.96% vs. 33.03%) over the strongest baseline, while reducing interpenetration volume by up to 21%. We further validate the physical feasibility of generated HOI via retargeting and demonstrate effectiveness on real-world dexterous manipulation tasks.

II. RELATED WORK

Hand-Object Interaction Generation. Synthesizing articulated hand-object motions under realistic contact has been supported by both mocap- and vision-based datasets [50, 95, 64, 23, 5]. Early kinematic methods leverage contact-aware priors, optimization, or implicit representations to reduce interpenetration [35, 29, 55, 44], often with latent-variable models in canonical object spaces [41, 96, 106, 109]. Physics-based RL approaches ensure dynamic feasibility but scale poorly across objects and suffer from sim-to-real gaps [11, 90]. Recent diffusion methods [26, 62, 13] improve temporal coherence, yet DiffH2O [18] still suffers from physical artifacts, semantic inconsistencies, and slow inference, and LatentHOI [54] remains limited in interaction length and data scale. We instead propose a flow matching framework that explicitly models distinct grasping and manipulation phases, achieving faster generation with improved physical plausibility and semantic alignment.

Robot Learning from Human Videos. Learning manipulation from human videos [71, 47, 75] is scalable but faces an embodiment gap. Egocentric datasets [5, 20, 81] partially bridge this gap, while VLA models [49, 10, 101], language-conditioned policies [3, 34], and video generation [7, 99, 72, 28] map perception to actions but lack explicit interaction structure. We instead acquire a robot-agnostic HOI sequence

prior as a motion script that can be retargeted to different embodiments.

Human Motion Synthesis. Diffusion-based methods [98, 45, 16] generate realistic motions from text [32, 83, 107] or action labels [31, 82]. Extending to HOI requires jointly modeling contact-constrained trajectories with dedicated datasets [52, 8, 66] that remain limited in scale [69, 32], and existing methods [103, 53, 79, 100] primarily target full-body motion. Our work generates fine-grained hand-object interactions with a data pipeline that recovers HOI motions from egocentric videos [37] to address data scarcity.

III. PRELIMINARY: FLOW MATCHING

In this section, we briefly review the flow matching framework for generative modeling [61]. Let $q(\mathbf{x})$ denote the unknown data distribution over $\mathbf{x} \in \mathbb{R}^d$ and let $p_0(\mathbf{x}) = \mathcal{N}(\mathbf{0}, \mathbf{I})$ be a simple prior. Flow matching learns a time-dependent vector field $\mathbf{v}(\mathbf{x}, \tau)$ that transports samples from p_0 to a target distribution $p_1 \approx q$ along a continuous probability path $\{p_\tau\}_{\tau \in [0,1]}$. This transport is defined by the neural ODE [15]:

$$\frac{d}{d\tau} \phi_\tau(\mathbf{x}) = \mathbf{v}(\phi_\tau(\mathbf{x}), \tau), \quad \phi_0(\mathbf{x}) = \mathbf{x}, \quad (1)$$

where ϕ_τ is the flow map. Since directly matching the marginal vector field of p_τ is generally intractable, *conditional flow matching* (CFM) [61] instead constructs a tractable conditional path $p_\tau(\mathbf{x} | \mathbf{x}_1)$ for each data sample $\mathbf{x}_1 \sim q$. A standard choice is the optimal-transport path with linear interpolation:

$$\mathbf{x}_\tau = (1 - (1 - \sigma_{\min})\tau)\mathbf{x}_0 + \tau\mathbf{x}_1, \quad \mathbf{x}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad \tau \sim \mathcal{U}[0, 1], \quad (2)$$

where $\sigma_{\min} > 0$ is a small constant that controls the residual stochasticity at the end of the flow. Under this path, the conditional vector field admits a closed form:

$$\mathbf{u}_\tau(\mathbf{x}_\tau | \mathbf{x}_1) = \frac{\mathbf{x}_1 - (1 - \sigma_{\min})\mathbf{x}_\tau}{1 - (1 - \sigma_{\min})\tau}. \quad (3)$$

We parameterize the vector field with a neural network $\mathbf{v}_\theta(\mathbf{x}, \tau, \mathbf{c})$, optionally conditioned on side information \mathbf{c} . The flow matching objective regresses the network vector field to this target:

$$\mathcal{L}_{\text{FM}} = \mathbb{E}_{\mathbf{x}_1, \mathbf{x}_0, \tau} \left[\left\| \mathbf{v}_\theta(\mathbf{x}_\tau, \tau, \mathbf{c}) - \mathbf{u}_\tau(\mathbf{x}_\tau | \mathbf{x}_1) \right\|_2^2 \right]. \quad (4)$$

At inference time, sampling starts from $\mathbf{x}(0) = \mathbf{x}_0 \sim p_0$ and integrates Eq. (1) to $\tau = 1$ using a numerical ODE solver. Euler discretization with K steps yields:

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \Delta\tau \mathbf{v}_\theta(\mathbf{x}_k, \tau_k, \mathbf{c}), \quad \Delta\tau = \frac{1}{K}, \quad \tau_k = \frac{k}{K}. \quad (5)$$

IV. METHODOLOGY

In this section, we present our flow-based HOI motion generation conditioned on a single-frame initial egocentric observation, text command, and 3D scene context. As illustrated in Fig. 2, our framework consists of two stages: a *Grasping* stage that generates approach-and-grasp motions from a pretrained prior fine-tuned on reconstructed egocentric

HOI data, and a *Manipulation* stage that produces subsequent interaction motions conditioned on 3D scene context and language instructions.

A. Problem Formulation

We address HOI motion generation in practical scenarios. Given the first egocentric observation I , from which we extract the initial hand-object state \mathbf{x}_{init} and object geometry \mathcal{M} , an action description T , and a 3DGS scene representation \mathcal{G} , our goal is to generate a temporally coherent hand-object interaction motion over N frames. We denote the object pose at frame t as \mathbf{O}_t , and the left and right hand states as \mathbf{H}_t^l and \mathbf{H}_t^r , respectively. The full HOI trajectory is defined as

$$\mathbf{x} = \{(\mathbf{H}_t^l, \mathbf{H}_t^r, \mathbf{O}_t)\}_{t=0}^{N-1} \in \mathbb{R}^{N \times D}, \quad (6)$$

where D denotes the per-frame feature dimension of the hand and object representations. Given a dataset of M paired interactions $\{(\mathbf{x}^{(i)}, \mathbf{c}^{(i)})\}_{i=1}^M$, our objective is to learn a conditional generative model $p_\theta(\mathbf{x} | \mathbf{c})$, where \mathbf{c} is the conditioning information composed of encoded \mathbf{x}_{init} , \mathcal{M} , T , and \mathcal{G} .

Data Representation. Following prior work [18], we use a compact canonical representation that couples hands and objects while remaining robust to global placement. At each frame t , the object pose is represented as translation and rotation:

$$\mathbf{O}_t = (\mathbf{p}_t^o, \mathbf{r}_t^o), \quad \mathbf{p}_t^o \in \mathbb{R}^3, \quad \mathbf{r}_t^o \in \mathbb{R}^6, \quad (7)$$

where \mathbf{r}_t^o denotes the continuous 6D rotation representation [111]. Each hand $h \in \{l, r\}$ is parameterized in MANO [92] space as

$$\mathbf{H}_t^h = (\tilde{\mathbf{p}}_t^h, \mathbf{r}_t^h, \boldsymbol{\theta}_t^h, \mathbf{s}_t^h) \in \mathbb{R}^{54}, \quad (8)$$

with $\tilde{\mathbf{p}}_t^h \in \mathbb{R}^3$ the hand root translation, $\mathbf{r}_t^h \in \mathbb{R}^6$ the global hand orientation, $\boldsymbol{\theta}_t^h \in \mathbb{R}^{24}$ the MANO pose coefficients in PCA space, and $\mathbf{s}_t^h \in \mathbb{R}^{21}$ the per-joint signed distance (SD) vectors to the object surface. Specifically, let $\mathbf{j}_k^{h,t} \in \mathbb{R}^3$ denote the 3D position of joint $k \in \{1, \dots, 21\}$ for hand h at frame t , and let \mathcal{M}_t be the posed object mesh (obtained by transforming \mathcal{M} with \mathbf{O}_t). Each entry of \mathbf{s}_t^h is defined as

$$\mathbf{s}_{t,k}^h = \mathbf{j}_k^{h,t} - \Pi_{\mathcal{M}_t}(\mathbf{j}_k^{h,t}), \quad (9)$$

where $\Pi_{\mathcal{M}_t}(\cdot)$ returns the closest point on the object surface.

To align with the two-stage pipeline, we anchor global translations at the transition frame $t_g = N_g - 1$, where N_g denotes the number of grasping-stage frames. Specifically, we express hand root translations \mathbf{p}_t^h relative to the object position $\mathbf{p}_{t_g}^o$:

$$\tilde{\mathbf{p}}_t^h = \mathbf{p}_t^h - \mathbf{p}_{t_g}^o, \quad (10)$$

This reduces variance across scenes and objects while keeping the interaction dynamics in a consistent reference frame.

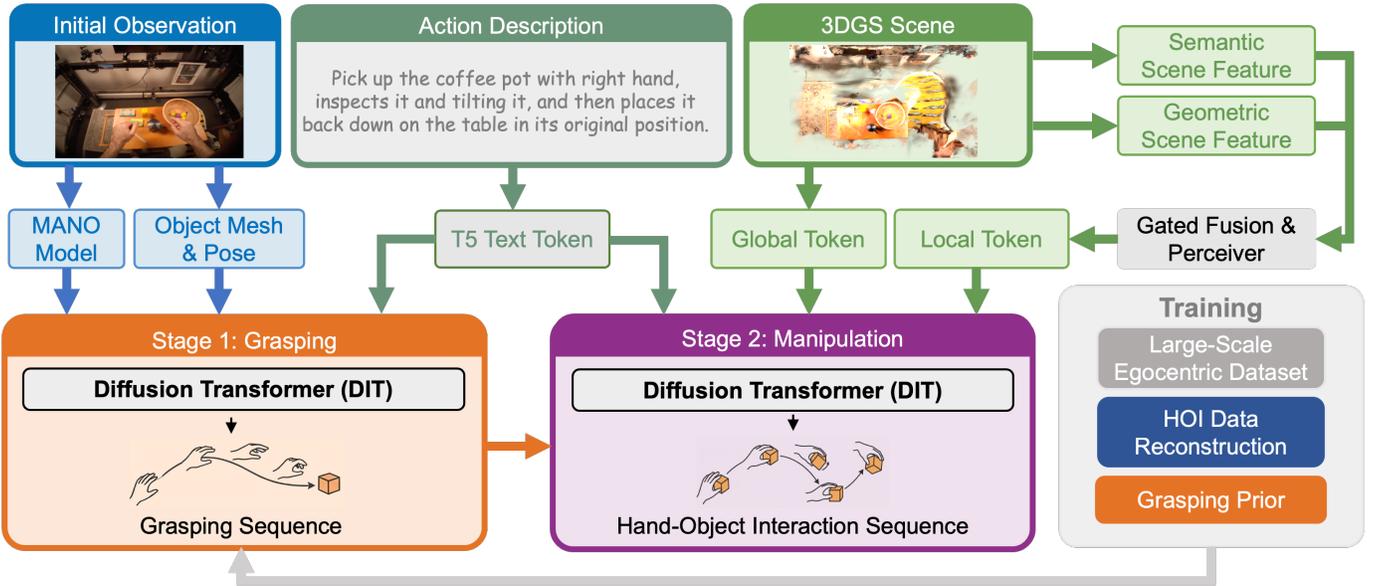


Fig. 2: **Overview of our framework.** Given an egocentric observation, text command, and 3D scene context, our method generates hand-object interaction motions through a two-stage pipeline: (1) a grasping stage that generates hand motion to approach and grasp the object, fine-tuned by reconstructed high-fidelity hand-object interaction data from large-scale egocentric videos, and (2) a manipulation stage that generates the subsequent interaction conditioned on scene and language.

B. Hand-Object Data Reconstruction from Egocentric Videos

Existing HOI datasets [95, 14] are orders of magnitude smaller than full-body MoCap collections [69] due to severe self-occlusion and tightly coupled contact dynamics in egocentric views. To bridge this gap, we build a reconstruction pipeline (Fig. 3) that converts raw egocentric videos from EgoDex [37], which provides RGB streams, camera parameters, text descriptions, and tracked hand keypoints but no object annotations, into high-fidelity HOI training data used exclusively to pretrain our grasping prior (Sec. IV-C).

The pipeline consists of three steps. **(1) Transition detection:** we smooth wrist trajectories and identify the grasp-to-manipulation transition via local speed minima and direction change. **(2) Object reconstruction:** we segment the target object with SAM3 [12], estimate metric depth with DepthAnything3 [60], and reconstruct a mesh with SAM3D [17] from pre-transition frames where the object is static. **(3) Hand-object alignment:** we fit MANO meshes via inverse kinematics, optimize an object translation offset to satisfy fingerpad contact and non-penetration constraints at the transition frame, and propagate the alignment to all frames. Full algorithmic details, loss formulations, and hyperparameters are provided in the supplementary material.

C. Two-Stage Hand-Object Interaction Generation

We naturally decompose HOI into two distinct phases: *Grasping* and *Manipulation*. Following the temporal decoupling strategy in DiffH2O [18], we also adopt a two-stage generation pipeline that explicitly models the two phases using specialized modules. This design allows each stage to effectively leverage the most relevant phase-specific conditioning signals: geometry and reachability for grasping, and action semantics and scene context for manipulation.

Initial Processing. Given the first egocentric observation \mathbf{I} , we estimate the initial hand pose using an off-the-shelf hand tracker [85] and reconstruct the target object mesh together with its 6D pose in the initial camera frame via SAM3D [17]. The estimated hand and object states are transformed into a common world coordinate frame, defined by the first camera’s extrinsics, yielding the initial state $\mathbf{x}_{\text{init}} = (\mathbf{O}_0, \mathbf{H}_0^l, \mathbf{H}_0^r)$. The object geometry \mathcal{M} is encoded using a Basis Point Set (BPS) representation [86], while the action description T is encoded by a frozen T5-Large [89] text encoder.

To incorporate rich scene context, we reconstruct a 3D scene representation using 3D Gaussian Splatting (3DGS) from the egocentric video [30, 68], with moving hands and objects masked out during reconstruction. From the reconstructed 3D scene, we sample N_s 3D points (Gaussian centroids) via Farthest Point Sampling (FPS) [21], denoted as $\mathbf{X} \in \mathbb{R}^{N_s \times 3}$. For the i -th scene point, we extract two complementary feature modalities: (i) a geometric embedding $\mathbf{e}_i \in \mathbb{R}^{d_e}$ capturing local spatial structure via Concerto [108], and (ii) a semantic embedding $\mathbf{u}_i \in \mathbb{R}^{d_u}$ obtained from the language-aligned scene representation of SceneSplat [58]. Stacking both feature types across all sampled scene points yields

$$\mathbf{E} \in \mathbb{R}^{N_s \times d_e}, \quad \mathbf{U} \in \mathbb{R}^{N_s \times d_u}. \quad (11)$$

We further fuse semantic and geometric cues using a linear-complexity gated fusion mechanism adapted from [39]. Specifically, both modalities are linearly projected into a shared latent space of dimension d_h with:

$$\tilde{\mathbf{E}} = \Phi_e(\mathbf{E}), \quad \tilde{\mathbf{U}} = \Phi_u(\mathbf{U}), \quad (12)$$

and combined via a learnable channel-wise gate $\alpha \in \mathbb{R}^{d_h}$ (α is a learnable vector):

$$\mathbf{F} = \tilde{\mathbf{E}} + \sigma(\alpha) \odot \tilde{\mathbf{U}}, \quad (13)$$

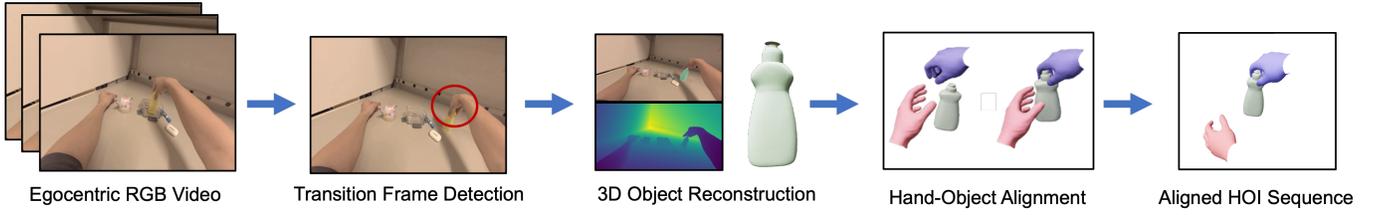


Fig. 3: **Hand-object data reconstruction pipeline.** Given an egocentric RGB video, we detect the grasp-to-manipulation transition frame from wrist motion cues, reconstruct the 3D object mesh from pre-transition frames via segmentation and metric depth estimation, and align the MANO hand mesh with the object under contact and non-penetration constraints to produce an aligned HOI sequence. See supplementary material for the detailed pipeline.

where $\sigma(\cdot)$ denotes the sigmoid function and \odot element-wise multiplication. This formulation enables adaptive modulation of semantic information and geometric structure while preserving linear memory and compute complexity.

To explicitly encode spatial layout, we concatenate a Fourier positional encoding $\gamma(\mathbf{X}) \in \mathbb{R}^{d_h}$ of the 3D point coordinates and apply a final projection to get the hybrid scene tokens:

$$\mathbf{P} = \Psi([\mathbf{F}, \gamma(\mathbf{X})]) \in \mathbb{R}^{N_s \times d_s}, \quad (14)$$

where d_s denotes the dimension of the scene token. Although \mathbf{P} has encoded dense local scene context, in the transformer backbone each per-frame feature of the noisy trajectory serves as a motion token, and the scene tokens are injected via cross-attention. Directly attending to all N_s scene tokens for each of the N motion tokens would incur a cost of $\mathcal{O}(N \cdot N_s)$. To address this, we employ a Perceiver bottleneck [40] to compress the hybrid scene tokens into a compact set of latent tokens:

$$\mathbf{S}_{\text{local}} = \text{Perceiver}(\mathbf{P}; \mathbf{Q}_0) \in \mathbb{R}^{L \times d_s}, \quad (15)$$

where $\mathbf{Q}_0 \in \mathbb{R}^{L \times d_s}$ denotes learnable latent queries and $L \ll N_s$. This bottleneck reduces the per-layer attention complexity to $\mathcal{O}(L \cdot N_s + N \cdot L)$ while preserving interaction-relevant local geometric and semantic structure.

In addition to the local context, a global scene token $\mathbf{S}_{\text{global}} \in \mathbb{R}^{1 \times d_s}$ is used to encode the coarse layout of the scene. Following prior scene-aware motion generation practice [42], we voxelize the reconstructed 3D scene into a coarse occupancy grid $\mathbf{V} \in \{0, 1\}^{H_z \times H_x \times H_y}$ via uniform grid sampling and encode the voxelized 3D scene via a vision transformer (ViT):

$$\mathbf{S}_{\text{global}} = \text{ViT}(\mathbf{V}) \in \mathbb{R}^{1 \times d_s}. \quad (16)$$

The global and local representations play complementary roles: $\mathbf{S}_{\text{global}}$ provides a holistic structural prior, encouraging globally consistent motion generation and avoiding catastrophic collisions, while $\mathbf{S}_{\text{local}}$ captures fine-grained geometric and semantic constraints tied to specific interaction regions.

Grasping Stage. In the grasping stage, the target object is static and we generate only the hand motion that approaches the object and establishes contact. We adopt an x -prediction variant of conditional flow matching for improved temporal stability [56]; details and ablations are provided in the supplementary material.

Since the object is static during grasping, only the hand state is generated. Let $\mathbf{x}^g = (\mathbf{H}^l, \mathbf{H}^r) \in \mathbb{R}^{N_g \times D_h}$ denote the grasping-stage hand states, where D_h is the per-frame hand feature dimension. The conditioning signal for grasping motion generation is defined as:

$$\mathbf{c}^g = \{ \text{BPS}(\mathcal{M}); \text{T5}(T_g); \mathbf{x}_{\text{init}} \}, \quad (17)$$

where T_g is a grasp-focused sub-instruction extracted from the full action instruction T using an MLLM [2], and $\text{T5}(\cdot)$ is the T5-Large [89] text encoder. This explicitly removes manipulation-related semantics from the grasping conditioner, preventing interference from manipulation during grasping. \mathbf{x}_{init} provides the initial hand state and is concatenated to the noisy motion token at each sampling step to anchor the generation.

Following conditional flow matching (Eq. 2), we sample \mathbf{x}_τ^g , $\tau \in [0, 1]$ along the linear interpolation path between Gaussian noise $\mathbf{x}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and the ground-truth grasping sequence \mathbf{x}_1^g . Instead of directly regressing the velocity \mathbf{u}_τ , we train a network $f_\theta^g(\cdot)$ to predict the clean target \mathbf{x}_1^g :

$$\mathcal{L}_{\text{flow}}^g = \mathbb{E}_{\tau, \mathbf{x}_0, \mathbf{x}_1^g} \left[\|f_\theta^g(\mathbf{x}_\tau^g, \tau, \mathbf{c}^g) - \mathbf{x}_1^g\|_2^2 \right]. \quad (18)$$

The velocity field is then derived via Eq. (3):

$$\mathbf{v}_\theta^g(\mathbf{x}_\tau^g, \tau, \mathbf{c}^g) = \frac{f_\theta^g(\mathbf{x}_\tau^g, \tau, \mathbf{c}^g) - (1 - \sigma_{\min})\mathbf{x}_\tau^g}{1 - (1 - \sigma_{\min})\tau}. \quad (19)$$

To enhance the controllability of motion generation under text instructions and ensure semantic grounding, we additionally introduce a contrastive alignment loss inspired by TMR [84]:

$$\mathcal{L}_{\text{align}} = \frac{1}{2} (\mathcal{L}_{\text{t2m}} + \mathcal{L}_{\text{m2t}}), \quad (20)$$

where \mathcal{L}_{t2m} and \mathcal{L}_{m2t} are symmetric InfoNCE [74] losses that align text and motion embeddings in a shared latent space. The total training loss is:

$$\mathcal{L}_{\text{grasp}} = \mathcal{L}_{\text{flow}}^g + \lambda_{\text{align}} \mathcal{L}_{\text{align}}, \quad (21)$$

with $\lambda_{\text{align}} = 0.1$. At inference, we integrate the derived vector field from $\tau = 0$ to $\tau = 1$ using Euler integration to obtain the denoised grasping trajectory, which terminates at the transition frame t_g .

Manipulation Stage. Unlike grasping, manipulation requires reasoning over longer horizons, where the object must be

moved in a task-consistent manner while preserving the established grasp and respecting scene constraints. To achieve consistency between the two stages, in our manipulation generator, we generate a complete HOI sequence $\mathbf{x}^m \in \mathbb{R}^{N \times D}$, encompassing both the grasping and manipulation stages. The generation process is conditioned on the previously generated grasping trajectory to ensure consistency.

The grasping trajectory $\mathbf{x}_{[0,t_g]}^g$ is treated as a known prefix, and only the post-grasp motion is modeled stochastically. The manipulation generator is conditioned on object geometry \mathcal{M} , language instruction T , scene context ($\mathbf{S}_{\text{local}}, \mathbf{S}_{\text{global}}$), and the grasp transition state $\mathbf{x}_{t_g}^g$ (the terminated state along the previously generated grasping trajectory \mathbf{x}^g):

$$\mathbf{c}^m = \{ \text{BPS}(\mathcal{M}); \text{T5}(T); \mathbf{S}_{\text{local}}; \mathbf{S}_{\text{global}}; \mathbf{x}_{t_g}^g \}. \quad (22)$$

A temporal mask $\mathbf{M} \in \{0, 1\}^N$ separates the fixed grasping segment ($\mathbf{M}_t = 0$ for $t \leq t_g$) from the future manipulation segment ($\mathbf{M}_t = 1$ for $t > t_g$). During inference, at each Euler integration step with flow time τ , the grasping prefix $\mathbf{x}_{[0,t_g]}^g$ is softly inpainted into the noisy trajectory by replacing the grasping segment with the previously generated result:

$$\bar{\mathbf{x}}_\tau^m = \mathbf{M} \odot \mathbf{x}_\tau^m + (1 - \mathbf{M}) \odot \mathbf{x}_{[0,t_g]}^g, \quad \tau < 0.9, \quad (23)$$

so that the manipulation network always observes the correct grasping context while generating future frames. The inpainting is disabled for $\tau \geq 0.9$ to allow the model to refine the full sequence without interference in the final integration steps. During training, the model is optimized with a masked objective applied only to the manipulation portion.

To enforce the continuity between the grasping and manipulation phases, the transition state $\mathbf{x}_{t_g}^g$ is imposed as a hard constraint. For each ODE-based sampling in flow matching, we explicitly clamp the transition state:

$$\bar{\mathbf{x}}_\tau^m [t_g] \leftarrow \mathbf{x}_{t_g}^g. \quad (24)$$

The combination of the subsequence soft inpainting and hard constraint ensures smooth and physically consistent generation of the complete HOI sequence, while preserving the generation model’s ability to jointly reason over the entire interaction. Our subsequence inpainting strategy is conceptually related to the inpainting mechanism used in DiffH2O [18]. However, while DiffH2O formulates inpainting at the level of discrete diffusion steps, we realize this idea within a continuous-time conditional flow matching framework and enforce transition consistency through ODE-level hard constraints.

The manipulation network f_θ^m is also trained with x_1 -prediction using a masked loss:

$$\mathcal{L}_{\text{flow}}^m = \mathbb{E}_{\tau, \mathbf{x}_0, \mathbf{x}_1^m} \left[\|\mathbf{M} \odot (f_\theta^m(\bar{\mathbf{x}}_\tau^m, \tau, \mathbf{c}^m) - \mathbf{x}_1^m)\|_2^2 \right], \quad (25)$$

where the loss is applied only to the unknown manipulation portion of the sequence to prevent trivial copying of the inpainted grasping motion. Similarly, the total training loss for the manipulation stage is:

$$\mathcal{L}_{\text{manip}} = \mathcal{L}_{\text{flow}}^m + \lambda_{\text{align}} \mathcal{L}_{\text{align}}, \quad (26)$$

where $\mathcal{L}_{\text{align}}$ is the contrastive alignment loss (Eq. 20).

A. Datasets

We evaluate our method on two datasets, namely HOT3D [5] and GRAB [95]. We pretrain our grasping model with our curated EgoDex [37] dataset as described in Sec. IV-B; we solely use it for training the grasping prior model and exclude it in our evaluations. GRAB [95] provides high-fidelity hand-object motion capture and is used for quantitative evaluation without scene context. HOT3D [5] offers real-world egocentric recordings with accurate hand and object pose annotations and reconstructed 3D scenes, enabling scene-conditioned evaluation and generalization to unseen objects. Note that both GRAB and HOT3D datasets do not provide 3D scene reconstructions or natural language action descriptions; details of our reconstruction and annotation procedures are provided in the supplementary material.

B. Implementation Details

Our model is implemented in PyTorch [77] and trained on an RTX 6000 Ada GPU (48GB) with a total batch size of 64. For DiffH2O [18] and LatentHOI [54], we use the official implementations and retrain all these baseline models on GRAB and HOT3D using the same data splits and evaluation protocols as our method. To ensure a fair comparison, we replace the CLIP [88] text encoder with a more capable T5 [89] encoder for all compared baseline methods. We adopt conditional flow matching [61] with optimal transport (OT) paths and train the model using x -prediction. Inference is performed using 50-step Euler integration. Additional implementation details are provided in the supplementary material.

C. Evaluation Metrics

Aligned with prior works [18, 54, 26, 97], we evaluate the physical interaction quality, motion quality, and realizable physical feasibility for generated HOI sequences. Detailed metric definitions are provided in the supplementary material.

Physical Interaction Quality. We measure interpenetration volume (IV) and interpenetration depth (ID) to quantify geometric violations between hand and object meshes, and report contact ratio (CR) to characterize sustained contact. Interpenetration volume per contact unit (IVU) is used as a normalized diagnostic metric.

Motion Quality. Semantic correctness is evaluated using action recognition accuracy (AR). Motion diversity is measured by sample diversity (SD) across repeated generations and overall diversity (OD) over the full test set.

Realizable Physical Feasibility. In prior work [54], physical plausibility (Phy) is assessed using heuristic criteria that consider only sustained hand-object contact and whether the object remains above the ground. We also evaluate with this metric. However, it is insufficient to reveal the true physical feasibility of generated human-object interactions. To address this limitation, we further evaluate the realizable physical feasibility of generated HOI sequences in a physics-based simulation environment using Isaac Gym [70].

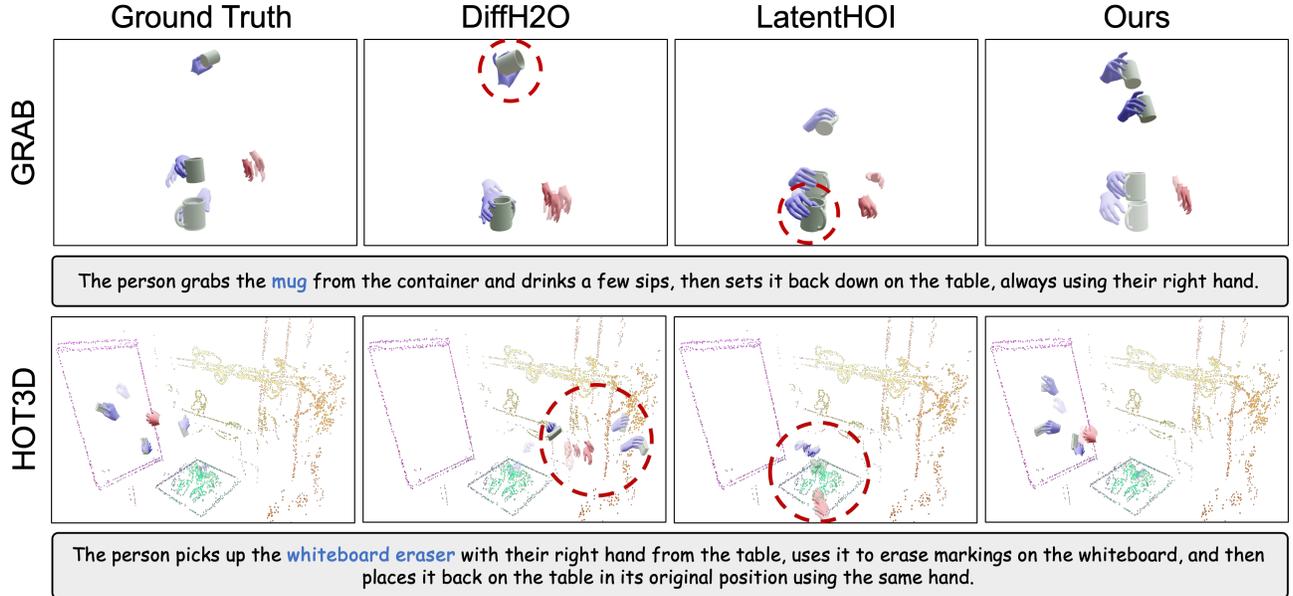


Fig. 4: Qualitative comparison of HOI generation. We compare our method with DiffH2O [18] and LatentHOI [54] against ground truth (GT). Top row: results on GRAB dataset. Bottom row: results on the HOT3D dataset in a 3D scene context. Our method generates more natural grasping poses and physically plausible manipulations that better align with the input action instructions and comply with the surrounding 3D scene layout. Best seen in the supplementary video.

Dataset	Method	Physical Interaction Quality				Motion Quality			Realizable Physical Feasibility			
		IV [cm^3] (\downarrow)	ID [cm] (\downarrow)	CR [%] (\uparrow)	IVU (\downarrow)	AR (\uparrow)	SD [m] (\uparrow)	OD [m] (\uparrow)	Phy [%] (\uparrow)	SR [%] (\uparrow)	HT [s] (\uparrow)	Time [s] (\downarrow)
GRAB	Real Mocap	6.49	0.43	6.62	0.09	1.00	-	0.18	95.56	77.06	2.40	-
	DiffH2O	13.11	1.32	8.14	0.16	0.87	0.12	0.13	89.44	33.03	0.35	6.34
	LatentHOI	13.76	0.99	9.18	0.14	0.78	0.11	0.14	96.40	28.44	0.29	3.57
	Ours	10.93	1.28	6.85	0.13	0.95	0.13	0.16	90.90	55.96	1.50	0.16
HOT3D	Real Mocap	2.88	0.63	1.69	0.15	1.00	-	0.23	16.80	6.67	0.08	-
	DiffH2O	3.25	0.66	2.07	0.18	0.71	0.15	0.23	16.47	4.00	0.03	6.54
	LatentHOI	1.54	0.44	1.59	0.21	0.65	0.12	0.18	14.36	0.00	0.00	3.21
	Ours	3.36	0.72	2.14	0.19	0.78	0.10	0.20	17.15	5.33	0.03	0.16

TABLE I: Quantitative evaluation on GRAB and HOT3D datasets. Best results are in **bold**.

Specifically, we first retarget the generated HOI sequences to the Allegro Hand via inverse kinematics [48]. The retargeted motions are then executed in Isaac Gym using a physics-based tracking controller to generate robot joint actions [63]. We report the success rate (SR) and holding time (HT), during which objects are stably held in the hand, as quantitative measures of realizable physical feasibility. The inference time for HOI generation is also reported.

D. HOI Generation Comparison

We compare FlowHOI with DiffH2O [18] and LatentHOI [54] on GRAB and HOT3D (Table I), focusing on long-horizon contact consistency, semantic alignment, and inference efficiency.

Contact consistency over long horizons. On GRAB, FlowHOI achieves the lowest IV and IVU, indicating that contact geometry remains consistent over extended interactions with less error accumulation. On HOT3D, FlowHOI attains the highest contact ratio while keeping penetration metrics comparable to baselines, suggesting that it maintains sustained, task-relevant contact even under real-world reconstruction noise.

Semantically grounded generation without sacrificing di-

versity. FlowHOI consistently achieves the highest action recognition accuracy, indicating that the generated motions are well grounded in the language instruction and the observed scene context. This improvement stems from conditioning the manipulation stage on compact 3D scene tokens and a motion-text alignment loss, which explicitly anchors the generated object state changes in both the physical scene layout and the language instruction. Importantly, this enhanced semantic grounding does not collapse motion diversity: both sample-level and overall diversity remain comparable to prior methods, with the highest overall diversity observed on GRAB.

Realizable physical feasibility and efficiency. In physics simulation, FlowHOI achieves the best SR on GRAB (55.96% vs. 28.44% for LatentHOI and 33.03% for DiffH2O) and the longest execution duration of 1.5 s, confirming that the generated trajectories remain stable and executable after retargeting. Although LatentHOI achieves the highest heuristic Phy score on GRAB, this metric does not capture physical stability under dynamics, as reflected by its substantially lower SR. By adopting flow matching, inference requires only 0.16 s per sequence, achieving up to 40 \times speedup over diffusion-based baselines (DiffH2O: 6.34 s, LatentHOI: 3.57 s). Qualitative results in Fig. 4 further corroborate these findings, showing

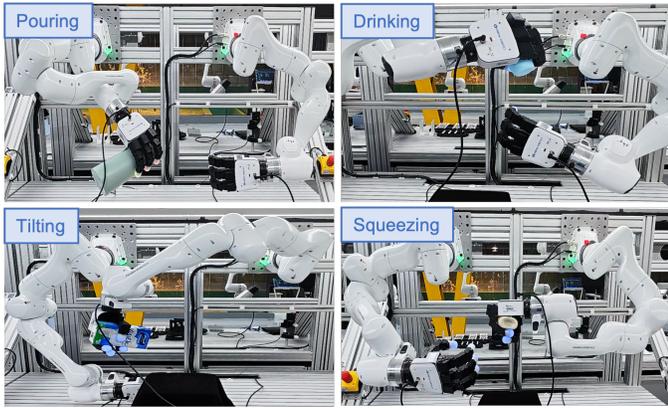


Fig. 5: Showcase of real-world robot applications. We retarget our generated HOI sequence to a Franka Panda arm with Allegro Hand for four contact-rich manipulation tasks: pouring, drinking, tilting, and squeezing. The robot successfully executes contact-rich interactions guided by our HOI sequence.

more realistic grasp configurations and smoother manipulation trajectories, particularly for long-horizon interactions.

E. Showcase of Real-world Applications

We further evaluate the physical feasibility of executing generated HOI sequences on a real-world dexterous manipulation platform, consisting of two Franka Emika Panda robotic arms [33], each equipped with an Allegro Hand v5 [1]. We consider four contact-rich household manipulation tasks: drinking from a cup, pouring liquid between containers of different sizes, tilting a container, and squeezing dressing. Perception inputs include egocentric RGB observations, reconstructed 3D scene representations, and a natural-language instruction (see Fig. 1). The initial MANO hand pose corresponding to the Allegro hands required by our model is obtained from reading and retargeting the robot hand proprioceptive state via an off-the-shelf kinematics-related retargeting solver [48]. Our generated HOI hand poses are then retargeted to the joint space of Allegro Hand using the same retargeting solver [48]. These retargeted reference actions of Allegro Hand are further refined by an existing dexterous robot hand motion tracker [63] to produce refined robot-executable joint actions, which are finally executed on a real robot using a standard joint impedance controller. Across the showcased tasks, the generated HOI trajectories can be consistently retargeted and successfully executed on the robot, producing stable contact-rich interactions that qualitatively match the intended object-centric behaviors (see Fig. 5).

F. Ablation Study

Effect of Pretraining Grasping Model. Table II compares grasping models with and without pretraining on large-scale egocentric data (Sec. IV-B). We report grasp error (GE), defined as the distance between the generated hand pose and the ground-truth grasp at the end of the grasping stage. Pretraining substantially reduces GE, indicating more accurate and stable grasp initialization. Improvements in penetration-related metrics further indicate that the pretrained grasping prior enhances contact quality in the grasping stage.

Model	IV [cm^3] (\downarrow)	ID [cm] (\downarrow)	CR [%] (\uparrow)	IVU (\downarrow)	GE [m] (\downarrow)
w/o prior	10.40	1.31	6.78	0.13	0.10
w/ prior	10.93	1.28	6.85	0.13	0.06

TABLE II: Effect of pretraining with large-scale egocentric HOI data on GRAB. GE: grasp error at the end of grasping stage.

Method	AR (\uparrow)	Method	AR (\uparrow)	FDE (\downarrow)
CLIP [88]	0.89	w/o scene	0.73	0.33
CLIP + Align. Loss	0.88	PointNet++ [87]	0.59	0.44
T5 [89]	0.92	Concerto [108]	0.75	0.26
T5 + Align. Loss	0.94	SceneSplat [58]	0.73	0.30
		Ours	0.78	0.24

(a) Text Encoder

(b) Scene Encoder

TABLE III: Ablation study on text encoder (a) and scene encoder (b). FDE: final displacement error of object pose.

Semantic grounding via text and scene conditioning. Results are summarized in Table III. For text encoding, T5 encoder consistently outperforms CLIP-based variants in action recognition accuracy. Incorporating the motion-text alignment loss (Eq. (20)) further improves the performance, indicating stronger correspondence between generated motions and language instructions. For scene encoding, removing scene information leads to degraded action recognition accuracy and large final displacement error (FDE). Geometry-only or semantics-only scene representations provide partial improvements, while our fused representation achieves the best performance on both metrics. This demonstrates that jointly modeling spatial geometric constraints and scene semantics is critical for accurate object motion and semantically consistent manipulation.

VI. CONCLUSION

We presented FlowHOI, a two-stage flow-matching framework that generates semantically grounded HOI sequences conditioned on an egocentric observation, a language instruction, and a 3DGS scene reconstruction. By decoupling geometry-centric grasping from semantics-centric manipulation with 3D scene tokens and a motion-text alignment loss, FlowHOI grounds interactions in both the physical scene and the language instruction. A reconstruction pipeline recovering HOI trajectories from egocentric videos further provides a robust prior for generalization. On GRAB and HOT3D, FlowHOI achieves the highest action-recognition accuracy, a $1.7\times$ higher physics-simulation success rate, up to 21% less interpenetration, and a $40\times$ inference speedup over the strongest baseline. Real-robot experiments on four dexterous tasks further validate retargeting to real-world execution.

This work also has several limitations. Our framework assumes accurate initial hand and object state estimation and degrades under heavy occlusion or unreliable reconstruction. The generated trajectories are kinematic and contact-consistent but rely on downstream controllers for dynamics and compliance. Extending to mobile manipulation and learning interaction priors from large-scale exocentric videos are promising future directions.

REFERENCES

- [1] Allegro hand v5. <https://www.allegrohand.com/>, accessed 2026. Wonik Robotics.
- [2] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- [3] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, et al. Do as i can, not as i say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022.
- [4] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35:23716–23736, 2022.
- [5] Prithviraj Banerjee, Sindi Shkodrani, Pierre Moulon, Shreyas Hampali, Shangchen Han, Fan Zhang, Linguang Zhang, Jade Fountain, Edward Miller, Selen Basol, et al. Hot3d: Hand and object tracking in 3d from egocentric multi-view videos. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 7061–7071, 2025.
- [6] Lucas Beyer, Andreas Steiner, André Susano Pinto, Alexander Kolesnikov, Xiao Wang, Daniel Salz, Maxim Neumann, Ibrahim Alabdulmohsin, Michael Tschannen, Emanuele Bugliarello, et al. Paligemma: A versatile 3b vlm for transfer. *arXiv preprint arXiv:2407.07726*, 2024.
- [7] Homanga Bharadhwaj, Debidatta Dwibedi, Abhinav Gupta, Shubham Tulsiani, Carl Doersch, Ted Xiao, Dhruv Shah, Fei Xia, Dorsa Sadigh, and Sean Kirmani. Gen2act: Human video generation in novel scenarios enables generalizable robot manipulation. *arXiv preprint arXiv:2409.16283*, 2024.
- [8] Bharat Lal Bhatnagar, Xianghui Xie, Ilya Petrov, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. Behave: Dataset and method for tracking human object interactions. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, jun 2022.
- [9] Antonio Bicchi and Vijay Kumar. Robotic grasping and contact: A review. In *Proceedings 2000 ICRA. Millennium conference. IEEE international conference on robotics and automation. Symposia proceedings (Cat. No. 00CH37065)*, volume 1, pages 348–353. IEEE, 2000.
- [10] Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Lucy Xiaoyang Shi, James Tanner, Quan Vuong, Anna Walling, Haohuan Wang, and Ury Zhilinsky. π_0 : A vision-language-action flow model for general robot control, 2026. URL <https://arxiv.org/abs/2410.24164>.
- [11] Jona Braun, Sammy Christen, Muhammed Kocabas, Emre Aksan, and Otmar Hilliges. Physically plausible full-body hand-object interaction synthesis. In *2024 International Conference on 3D Vision (3DV)*, pages 464–473. IEEE, 2024.
- [12] Nicolas Carion, Laura Gustafson, Yuan-Ting Hu, Shoubhik Debnath, Ronghang Hu, Didac Suris, Chaitanya Ryali, Kalyan Vasudev Alwala, Haitham Khedr, Andrew Huang, et al. Sam 3: Segment anything with concepts. *arXiv preprint arXiv:2511.16719*, 2025.
- [13] Junuk Cha, Jihyeon Kim, Jae Shin Yoon, and Seungryul Baek. Text2hoi: Text-guided 3d motion generation for hand-object interaction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1577–1585, 2024.
- [14] 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. Dexycb: A benchmark for capturing hand grasping of objects. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9044–9053, 2021.
- [15] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.
- [16] Xin Chen, Biao Jiang, Wen Liu, Zilong Huang, Bin Fu, Tao Chen, and Gang Yu. Executing your commands via motion diffusion in latent space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 18000–18010, 2023.
- [17] Xingyu Chen, Fu-Jen Chu, Pierre Gleize, Kevin J Liang, Alexander Sax, Hao Tang, Weiyao Wang, Michelle Guo, Thibaut Hardin, Xiang Li, et al. Sam 3d: 3dfy anything in images. *arXiv preprint arXiv:2511.16624*, 2025.
- [18] Sammy Christen, Shreyas Hampali, Fadime Sener, Edoardo Remelli, Tomas Hodan, Eric Sauser, Shugao Ma, and Bugra Tekin. Diffh2o: Diffusion-based synthesis of hand-object interactions from textual descriptions. In *SIGGRAPH Asia 2024 Conference Papers*, pages 1–11, 2024.
- [19] Moo K Chung. Gaussian kernel smoothing. *arXiv preprint arXiv:2007.09539*, 2020.
- [20] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. The epic-kitchens dataset: Collection, challenges and baselines. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(11): 4125–4141, 2020.

- [21] Yuval Eldar, Michael Lindenbaum, Moshe Porat, and Yehoshua Y Zeevi. The farthest point strategy for progressive image sampling. *IEEE transactions on image processing*, 6(9):1305–1315, 1997.
- [22] Jakob Engel, Kiran Somasundaram, Michael Goesele, Albert Sun, Alexander Gamino, Andrew Turner, Arjang Talattof, Arnie Yuan, Bilal Souti, Brigid Meredith, et al. Project aria: A new tool for egocentric multi-modal ai research. *arXiv preprint arXiv:2308.13561*, 2023.
- [23] Zicong Fan, Omid Taheri, Dimitrios Tzionas, Muhammed Kocabas, Manuel Kaufmann, Michael J Black, and Otmar Hilliges. Arctic: A dataset for dexterous bimanual hand-object manipulation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12943–12954, 2023.
- [24] Tamar Flash and Neville Hogan. The coordination of arm movements: an experimentally confirmed mathematical model. *Journal of neuroscience*, 5(7):1688–1703, 1985.
- [25] Hongming Fu, Wenjia Wang, Xiaozhen Qiao, Shuo Yang, Zheng Liu, and Bo Zhao. Egograsp: World-space hand-object interaction estimation from egocentric videos. *arXiv preprint arXiv:2601.01050*, 2026.
- [26] Anindita Ghosh, Rishabh Dabral, Vladislav Golyanik, Christian Theobalt, and Philipp Slusallek. Imos: Intent-driven full-body motion synthesis for human-object interactions. In *Computer Graphics Forum*, volume 42, pages 1–12. Wiley Online Library, 2023.
- [27] James J. Gibson. *The Ecological Approach to Visual Perception: Classic Edition*. Houghton Mifflin, 1979.
- [28] Raktim Gautam Goswami, Amir Bar, David Fan, Tsung-Yen Yang, Gaoyue Zhou, Prashanth Krishnamurthy, Michael Rabbat, Farshad Khorrani, and Yann LeCun. World models can leverage human videos for dexterous manipulation. *arXiv preprint arXiv:2512.13644*, 2025.
- [29] Patrick Grady, Chengcheng Tang, Christopher D Twigg, Minh Vo, Samarth Brahmhatt, and Charles C Kemp. Contactopt: Optimizing contact to improve grasps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1471–1481, 2021.
- [30] Qiao Gu, Zhaoyang Lv, Duncan Frost, Simon Green, Julian Straub, and Chris Sweeney. Egolifter: Open-world 3d segmentation for egocentric perception. In *European Conference on Computer Vision*, pages 382–400. Springer, 2024.
- [31] Chuan Guo, Xinxin Zuo, Sen Wang, Shihao Zou, Qingyao Sun, Annan Deng, Minglun Gong, and Li Cheng. Action2motion: Conditioned generation of 3d human motions. In *Proceedings of the 28th ACM international conference on multimedia*, pages 2021–2029, 2020.
- [32] Chuan Guo, Shihao Zou, Xinxin Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. Generating diverse and natural 3d human motions from text. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5152–5161, 2022.
- [33] Sami Haddadin. The franka emika robot: A standard platform in robotics research. *IEEE Robotics & Automation Magazine*, 2024.
- [34] Ce Hao, Kelvin Lin, Siyuan Luo, and Harold Soh. Language-guided manipulation with diffusion policies and constrained inpainting. *arXiv preprint arXiv:2406.09767*, 2024.
- [35] Mohamed Hassan, Yunrong Guo, Tingwu Wang, Michael Black, Sanja Fidler, and Xue Bin Peng. Synthesizing physical character-scene interactions. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–9, 2023.
- [36] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [37] Ryan Hoque, Peide Huang, David J Yoon, Mouli Siva-purapu, and Jian Zhang. Egodex: Learning dexterous manipulation from large-scale egocentric video. *arXiv preprint arXiv:2505.11709*, 2025.
- [38] Cheng-Chun Hsu, Bowen Wen, Jie Xu, Yashraj Narang, Xiaolong Wang, Yuke Zhu, Joydeep Biswas, and Stan Birchfield. Spot: Se (3) pose trajectory diffusion for object-centric manipulation. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4853–4860. IEEE, 2025.
- [39] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141, 2018.
- [40] Andrew Jaegle, Felix Gimeno, Andy Brock, Oriol Vinyals, Andrew Zisserman, and Joao Carreira. Perceiver: General perception with iterative attention. In *International conference on machine learning*, pages 4651–4664. PMLR, 2021.
- [41] Hanwen Jiang, Shaowei Liu, Jiashun Wang, and Xiaolong Wang. Hand-object contact consistency reasoning for human grasps generation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 11107–11116, 2021.
- [42] Nan Jiang, Zhiyuan Zhang, Hongjie Li, Xiaoxuan Ma, Zan Wang, Yixin Chen, Tengyu Liu, Yixin Zhu, and Siyuan Huang. Scaling up dynamic human-scene interaction modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1737–1747, 2024.
- [43] Roland S Johansson and J Randall Flanagan. Coding and use of tactile signals from the fingertips in object manipulation tasks. *Nature Reviews Neuroscience*, 10(5):345–359, 2009.
- [44] Korrawe Karunratanakul, Jinlong Yang, Yan Zhang, Michael J Black, Krikamol Muandet, and Siyu Tang. Grasping field: Learning implicit representations for

- human grasps. In *2020 International Conference on 3D Vision (3DV)*, pages 333–344. IEEE, 2020.
- [45] Korrawe Karunratanakul, Konpat Preechakul, Supasorn Suwajanakorn, and Siyu Tang. Guided motion diffusion for controllable human motion synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2151–2162, 2023.
- [46] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023.
- [47] Alexander Khazatsky, Karl Pertsch, Suraj Nair, Ashwin Balakrishna, Sudeep Dasari, Siddharth Karamcheti, Soroush Nasiriany, Mohan Kumar Srirama, Lawrence Yunliang Chen, Kirsty Ellis, et al. Droid: A large-scale in-the-wild robot manipulation dataset. *arXiv preprint arXiv:2403.12945*, 2024.
- [48] Chung Min Kim, Brent Yi, Hongsuk Choi, Yi Ma, Ken Goldberg, and Angjoo Kanazawa. Pyroki: A modular toolkit for robot kinematic optimization. *arXiv preprint arXiv:2505.03728*, 2025.
- [49] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, et al. Openvla: An open-source vision-language-action model. *arXiv preprint arXiv:2406.09246*, 2024.
- [50] Taein Kwon, Bugra Tekin, Jan Stühmer, Federica Bogo, and Marc Pollefeys. H2o: Two hands manipulating objects for first person interaction recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10138–10148, 2021.
- [51] Xiaolei Lang, Laijian Li, Chenming Wu, Chen Zhao, Lina Liu, Yong Liu, Jiajun Lv, and Xingxing Zuo. Gaussian-lic: Real-time photo-realistic slam with gaussian splatting and lidar-inertial-camera fusion. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 8500–8507. IEEE, 2025.
- [52] Jiaman Li, Jiajun Wu, and C Karen Liu. Object motion guided human motion synthesis. *ACM Transactions on Graphics (TOG)*, 42(6):1–11, 2023.
- [53] Jiaman Li, Alexander Clegg, Roozbeh Mottaghi, Jiajun Wu, Xavier Puig, and C Karen Liu. Controllable human-object interaction synthesis. In *European Conference on Computer Vision*, pages 54–72. Springer, 2024.
- [54] Muchen Li, Sammy Christen, Chengde Wan, Yujun Cai, Renjie Liao, Leonid Sigal, and Shugao Ma. Latentthoi: On the generalizable hand object motion generation with latent hand diffusion. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 17416–17425, 2025.
- [55] Quanzhou Li, Jingbo Wang, Chen Change Loy, and Bo Dai. Task-oriented human-object interactions generation with implicit neural representations. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3035–3044, 2024.
- [56] Tianhong Li and Kaiming He. Back to basics: Let denoising generative models denoise. *arXiv preprint arXiv:2511.13720*, 2025.
- [57] Yang Li, Si Si, Gang Li, Cho-Jui Hsieh, and Samy Bengio. Learnable fourier features for multi-dimensional spatial positional encoding. *Advances in Neural Information Processing Systems*, 34:15816–15829, 2021.
- [58] Yue Li, Qi Ma, Runyi Yang, Huapeng Li, Mengjiao Ma, Bin Ren, Nikola Popovic, Nicu Sebe, Ender Konukoglu, Theo Gevers, et al. Scenesplat: Gaussian splatting-based scene understanding with vision-language pretraining. *arXiv preprint arXiv:2503.18052*, 2025.
- [59] Zhiqi Li, Guo Chen, Shilong Liu, Shihao Wang, Vibashan VS, Yishen Ji, Shiyi Lan, Hao Zhang, Yilin Zhao, Subhashree Radhakrishnan, et al. Eagle 2: Building post-training data strategies from scratch for frontier vision-language models. *arXiv preprint arXiv:2501.14818*, 2025.
- [60] Haotong Lin, Sili Chen, Junhao Liew, Donny Y Chen, Zhenyu Li, Guang Shi, Jiashi Feng, and Bingyi Kang. Depth anything 3: Recovering the visual space from any views. *arXiv preprint arXiv:2511.10647*, 2025.
- [61] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.
- [62] Xueyi Liu and Li Yi. Geneoh diffusion: Towards generalizable hand-object interaction denoising via denoising diffusion. *arXiv preprint arXiv:2402.14810*, 2024.
- [63] Xueyi Liu, Jianibieke Adalibieke, Qianwei Han, Yuzhe Qin, and Li Yi. Dextrack: Towards generalizable neural tracking control for dexterous manipulation from human references. *arXiv preprint arXiv:2502.09614*, 2025.
- [64] Yunze Liu, Yun Liu, Che Jiang, Kangbo Lyu, Weikang Wan, Hao Shen, Boqiang Liang, Zhoujie Fu, He Wang, and Li Yi. Hoi4d: A 4d egocentric dataset for category-level human-object interaction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21013–21022, 2022.
- [65] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [66] Jiaxin Lu, Chun-Hao Paul Huang, Uttaran Bhattacharya, Qixing Huang, and Yi Zhou. Humoto: A 4d dataset of mocap human object interactions. *arXiv preprint arXiv:2504.10414*, 2025.
- [67] Tyler Ga Wei Lum, Olivia Y Lee, C Karen Liu, and Jeannette Bohg. Crossing the human-robot embodiment gap with sim-to-real rl using one human demonstration. *arXiv preprint arXiv:2504.12609*, 2025.
- [68] Zhaoyang Lv, Maurizio Monge, Ka Chen, Yufeng Zhu, Michael Goesele, Jakob Engel, Zhao Dong, and Richard Newcombe. Photoreal scene reconstruction from an egocentric device. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers*, pages 1–11,

- 2025.
- [69] Naureen Mahmood, Nima Ghorbani, Nikolaus F Troje, Gerard Pons-Moll, and Michael J Black. Amass: Archive of motion capture as surface shapes. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5442–5451, 2019.
- [70] Viktor Makoviychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, et al. Isaac gym: High performance gpu-based physics simulation for robot learning. *arXiv preprint arXiv:2108.10470*, 2021.
- [71] Ajay Mandlekar, Jonathan Booher, Max Spero, Albert Tung, Anchit Gupta, Yuke Zhu, Animesh Garg, Silvio Savarese, and Li Fei-Fei. Scaling robot supervision to hundreds of hours with roboturk: Robotic manipulation dataset through human reasoning and dexterity. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1048–1055. IEEE, 2019.
- [72] Jiageng Mao, Sicheng He, Hao-Ning Wu, Yang You, Shuyang Sun, Zhicheng Wang, Yanan Bao, Huizhong Chen, Leonidas Guibas, Vitor Guizilini, et al. Robot learning from a physical world model. *arXiv preprint arXiv:2511.07416*, 2025.
- [73] Matthew T Mason. *Mechanics of robotic manipulation*. MIT press, 2001.
- [74] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [75] Abby O’Neill, Abdul Rehman, Abhiram Maddukuri, Abhishek Gupta, Abhishek Padalkar, Abraham Lee, Acorn Pooley, Agrim Gupta, Ajay Mandlekar, Ajinkya Jain, et al. Open x-embodiment: Robotic learning datasets and rt-x models: Open x-embodiment collaboration 0. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6892–6903. IEEE, 2024.
- [76] Mingjie Pan, Jiayao Zhang, Tianshu Wu, Yinghao Zhao, Wenlong Gao, and Hao Dong. Omnimanip: Towards general robotic manipulation via object-centric interaction primitives as spatial constraints. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 17359–17369, 2025.
- [77] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- [78] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4195–4205, 2023.
- [79] Xiaogang Peng, Yiming Xie, Zizhao Wu, Varun Jampani, Deqing Sun, and Huaizu Jiang. Hoi-diff: Text-driven synthesis of 3d human-object interactions using diffusion models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 2878–2888, 2025.
- [80] Zhiliang Peng, Wenhui Wang, Li Dong, Yaru Hao, Shaohan Huang, Shuming Ma, and Furu Wei. Kosmos-2: Grounding multimodal large language models to the world. *arXiv preprint arXiv:2306.14824*, 2023.
- [81] Toby Perrett, Ahmad Darkhalil, Saptarshi Sinha, Omar Emara, Sam Pollard, Kranti Kumar Parida, Kaiting Liu, Prajwal Gatti, Siddhant Bansal, Kevin Flanagan, et al. Hd-epic: A highly-detailed egocentric video dataset. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 23901–23913, 2025.
- [82] Mathis Petrovich, Michael J Black, and Gül Varol. Action-conditioned 3d human motion synthesis with transformer vae. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10985–10995, 2021.
- [83] Mathis Petrovich, Michael J Black, and Gül Varol. Temos: Generating diverse human motions from textual descriptions. In *European Conference on Computer Vision*, pages 480–497. Springer, 2022.
- [84] Mathis Petrovich, Michael J Black, and Gül Varol. Tmr: Text-to-motion retrieval using contrastive 3d human motion synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9488–9497, 2023.
- [85] Rolandos Alexandros Potamias, Jinglei Zhang, Jiankang Deng, and Stefanos Zafeiriou. Wilor: End-to-end 3d hand localization and reconstruction in-the-wild. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 12242–12254, 2025.
- [86] Sergey Prokudin, Christoph Lassner, and Javier Romero. Efficient learning on point clouds with basis point sets. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4332–4341, 2019.
- [87] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017.
- [88] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmlR, 2021.
- [89] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67, 2020.
- [90] Aravind Rajeswaran, Vikash Kumar, Abhishek Gupta, Giulia Vezzani, John Schulman, Emanuel Todorov, and Sergey Levine. Learning complex dexterous manipula-

- tion with deep reinforcement learning and demonstrations. *arXiv preprint arXiv:1709.10087*, 2017.
- [91] Santhosh K Ramakrishnan, Aaron Gokaslan, Erik Wijmans, Oleksandr Maksymets, Alex Clegg, John Turner, Eric Undersander, Wojciech Galuba, Andrew Westbury, Angel X Chang, et al. Habitat-matterport 3d dataset (hm3d): 1000 large-scale 3d environments for embodied ai. *arXiv preprint arXiv:2109.08238*, 2021.
- [92] Javier Romero, Dimitrios Tzionas, and Michael J. Black. Embodied hands: Modeling and capturing hands and bodies together. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)*, 36(6), November 2017.
- [93] SkylandX. MetaCam: 3D scanner for spatial intelligence. <https://skylandx.com/>, accessed 2026.
- [94] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.
- [95] Omid Taheri, Nima Ghorbani, Michael J Black, and Dimitrios Tzionas. Grab: A dataset of whole-body human grasping of objects. In *European conference on computer vision*, pages 581–600. Springer, 2020.
- [96] Omid Taheri, Vasileios Choutas, Michael J Black, and Dimitrios Tzionas. Goal: Generating 4d whole-body motion for hand-object grasping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13263–13273, 2022.
- [97] Purva Tendulkar, Dídac Surís, and Carl Vondrick. Flex: Full-body grasping without full-body grasps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21179–21189, 2023.
- [98] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Daniel Cohen-Or, and Amit H Bermano. Human motion diffusion model. *arXiv preprint arXiv:2209.14916*, 2022.
- [99] Boyang Wang, Nikhil Sridhar, Chao Feng, Mark Van der Merwe, Adam Fishman, Nima Fazeli, and Jeong Joon Park. This&that: Language-gesture controlled video generation for robot planning. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 12842–12849. IEEE, 2025.
- [100] Yinhuai Wang, Jing Lin, Ailing Zeng, Zhengyi Luo, Jian Zhang, and Lei Zhang. Physshoi: Physics-based imitation of dynamic human-object interaction. *arXiv preprint arXiv:2312.04393*, 2023.
- [101] Yuqi Wang, Xinghang Li, Wenxuan Wang, Junbo Zhang, Yingyan Li, Yuntao Chen, Xinlong Wang, and Zhaoxiang Zhang. Unified vision-language-action model. *arXiv preprint arXiv:2506.19850*, 2025.
- [102] Rongtao Xu, Jian Zhang, Minghao Guo, Youpeng Wen, Haoting Yang, Min Lin, Jianzheng Huang, Zhe Li, Kaidong Zhang, Liqiong Wang, et al. A0: An affordance-aware hierarchical model for general robotic manipulation. *arXiv preprint arXiv:2504.12636*, 2025.
- [103] Sirui Xu, Zhengyuan Li, Yu-Xiong Wang, and Liang-Yan Gui. Interdiff: Generating 3d human-object interactions with physics-informed diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14928–14940, 2023.
- [104] Yue Xu, Yong-Lu Li, Zhemin Huang, Michael Xu Liu, Cewu Lu, Yu-Wing Tai, and Chi-Keung Tang. Egopca: A new framework for egocentric hand-object interaction understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5273–5284, 2023.
- [105] Chandan Yeshwanth, Yueh-Cheng Liu, Matthias Nießner, and Angela Dai. Scannet++: A high-fidelity dataset of 3d indoor scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12–22, 2023.
- [106] He Zhang, Yuting Ye, Takaaki Shiratori, and Taku Komura. Manipnet: neural manipulation synthesis with a hand-object spatial representation. *ACM Transactions on Graphics (ToG)*, 40(4):1–14, 2021.
- [107] Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. Motiondiffuse: Text-driven human motion generation with diffusion model. *IEEE transactions on pattern analysis and machine intelligence*, 46(6):4115–4128, 2024.
- [108] Yujia Zhang, Xiaoyang Wu, Yixing Lao, Chengyao Wang, Zhuotao Tian, Naiyan Wang, and Hengshuang Zhao. Concerto: Joint 2d-3d self-supervised learning emerges spatial representations. *arXiv preprint arXiv:2510.23607*, 2025.
- [109] Juntian Zheng, Qingyuan Zheng, Lixing Fang, Yun Liu, and Li Yi. Cams: Canonicalized manipulation spaces for category-level functional hand-object manipulation synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 585–594, 2023.
- [110] Yifan Zhong, Xuchuan Huang, Ruochong Li, Ceyao Zhang, Zhang Chen, Tianrui Guan, Fanlian Zeng, Ka Num Lui, Yuyao Ye, Yitao Liang, et al. Dex-graspv1a: A vision-language-action framework towards general dexterous grasping. *arXiv preprint arXiv:2502.20900*, 2025.
- [111] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5745–5753, 2019.

FlowHOI: Flow-based Semantics-Grounded Generation of Hand-Object Interactions for Dexterous Robot Manipulation

Supplementary Material

This supplementary document provides additional details and results that complement the main paper. We begin with implementation details of the model architecture and training configurations in Sec. A, followed by an ablation study on the flow matching prediction target in Sec. B. We then describe the hand-object data reconstruction pipeline and dataset preprocessing in Sec. C and Sec. D, respectively. Next, we elaborate on the physics simulation evaluation protocol in Sec. E and provide details on the real-world experiment setup in Sec. F. We then analyze the scene representation in Sec. G. Full evaluation metric definitions are provided in Sec. H. We then present additional qualitative results in Sec. I. Finally, we discuss failure cases in Sec. J. We also provide a supplementary video with animated results.

A. IMPLEMENTATION DETAILS

This section provides additional implementation details beyond Sec. V-B of the main paper.

Both models use a DiT [78] backbone with pre-norm. The manipulation model uses $d=512$, $d_{ff}=1024$; the grasping model uses $d=256$, $d_{ff}=512$. Both use $L=8$ layers and dropout 0.1. Text is encoded by a frozen T5-Large [89], projected to d and injected via cross-attention. Object geometry uses BPS [86] encoding, projected to 256-dim and concatenated per-frame.

The scene is represented by up to 25,000 points sampled via Farthest Point Sampling (FPS) [21]. Geometric features from Concerto [108] $d_e=1536$ and semantic features from SceneSplat [58] $d_u=768$ are fused via bidirectional cross-attention with $d_h=512$ for manipulation, 256 for grasping. 3D coordinates use Fourier positional encoding [57], projected to 64-dim. Fused features are compressed by a Perceiver [40] bottleneck into $K=256$ scene tokens of dimension d . For global scene encoding, we voxelize the scene at $48 \times 48 \times 24$ resolution and process it with a ViT, injected via AdaLN [78].

We use AdamW [65] with lr 10^{-4} , weight decay 0.01, and linear annealing. The grasping model trains for 500K steps; the manipulation model for 200K steps. We adopt logit-normal time sampling and classifier-free guidance with drop probability 0.1 during training and scale 2.5 at inference with 50-step Euler integration.

Each frame is a 117-dim vector encoding bimanual wrist poses, hand articulations in 24D PCA, per-joint signed distances to the object, and object pose. The full sequence length is $N=200$ frames with $N_g=50$ for grasping and 150 for manipulation.

Contrastive Alignment Loss. We employ a TMR-inspired [84] symmetric InfoNCE loss to align motion and text embeddings in a shared 512-dim latent space. The motion encoder is a 4-layer Transformer encoder with a learnable [CLS] token. The text encoder applies mean

pooling over frozen T5-Large [89] embeddings and projects them to the same 512-dim space via a two-layer MLP. The temperature τ is initialized to 0.7, learnable, and clamped to $[0.01, 100]$. The alignment loss is weighted by $\lambda_{\text{align}}=0.1$ and directly added to the flow matching loss. These alignment components are only used during training and discarded at inference.

B. FLOW MATCHING PREDICTION TARGET ANALYSIS

Different from standard flow matching that predicts the velocity field \mathbf{v} , we propose to directly predict the clean data \mathbf{x}_1 as the model output. Recent studies [56] have conducted extensive analyses on the choice of prediction targets and consistently advocate for directly predicting the clean data (x -prediction), rather than noise (ϵ -prediction) or flow velocity (v -prediction). These works are motivated by the *manifold assumption*, which posits that natural data, such as images, concentrate on a low-dimensional manifold embedded in a high-dimensional ambient space, while quantities such as noise ϵ or flow velocity $v = x - \epsilon$ are inherently off-manifold and distributed across the full ambient space.

In our problem, although the interaction sequence is represented in a high-dimensional space, physically and semantically plausible hand-object motions are highly constrained (e.g., kinematics, joint limits, contact consistency, and temporal coherence), suggesting a similarly thin *interaction manifold*. Under velocity-based training, small high-frequency errors in \mathbf{v}_θ are directly injected into the ODE integration and can accumulate across steps, manifesting as temporal jitter [56]. In contrast, x -prediction repeatedly estimates an on-manifold state and then derives the vector field from it, which empirically improves temporal stability. To validate this hypothesis, we conduct an ablation study comparing the two prediction targets (x -prediction and v -prediction) on the GRAB [95] dataset. We evaluate temporal smoothness using two complementary metrics (see Sec. H for formal definitions): (1) *jerk*, the third derivative of object position/rotation with respect to time, and (2) *hand acceleration*, which measures the second-order temporal stability of the generated hand motion, decomposed into global wrist acceleration (Acc_g) and local finger acceleration (Acc_l) for both positional and rotational components. As shown in Tab. IV, x -prediction achieves lower angular jerk and consistently lower local hand acceleration (both positional and rotational), indicating smoother finger articulation. The global wrist metrics are comparable between the two targets, suggesting that the primary benefit of x -prediction lies in stabilizing fine-grained hand dynamics rather than coarse wrist trajectories. Fig. 6 further visualizes the generated trajectories, where v -prediction exhibits noticeable high-frequency jitter whereas x -prediction produces temporally smooth results.

Prediction Target	Jerk (Pos.) [m/s^3] (\downarrow)	Jerk (Ang.) [rad/s^3] (\downarrow)	Acc _g (Pos.) [m/s^2] (\downarrow)	Acc _l (Pos.) [m/s^2] (\downarrow)	Acc _g (Rot.) [rad/s^2] (\downarrow)	Acc _l (Rot.) [rad/s^2] (\downarrow)
<i>v</i> -prediction	4.9721	137.9271	0.2523	0.1232	0.9567	3.2414
<i>x</i> -prediction (Ours)	0.0438	0.4071	0.0204	0.0055	0.0835	0.0339

TABLE IV: Ablation study on flow matching prediction targets on GRAB [95]. *x*-prediction yields lower angular jerk and hand acceleration while maintaining competitive interaction quality. Acc_g/Acc_l: global/local acceleration (see Sec. H).

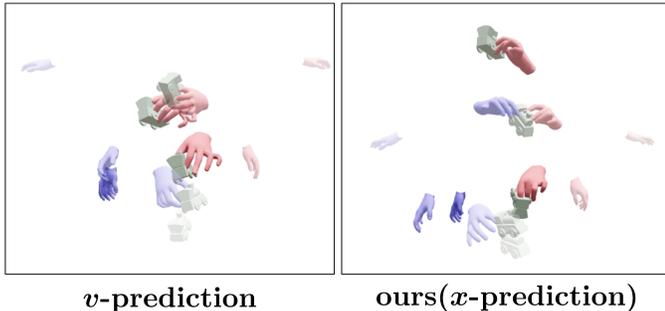


Fig. 6: Qualitative comparison of prediction targets. Our *x*-prediction produces temporally smooth and stable hand trajectories, while *v*-prediction exhibits noticeable spatial jitter and inconsistent grasping poses across frames.

C. HAND-OBJECT DATA RECONSTRUCTION DETAILS

The goal of this pipeline is to reconstruct high-fidelity hand-object interaction data from large-scale egocentric videos [37], which provides large-scale training data for pretraining the grasping model and improves generalization across objects and tasks. This section provides the full algorithmic details, loss formulations, and hyperparameters of the reconstruction pipeline summarized in Sec. IV-B of the main paper. An overview of the pipeline is illustrated in Fig. 7.

A. Step 1: Transition Frame Detection

To separate the grasp phase from the manipulation phase, we detect a transition frame using wrist motion cues. We first smooth the wrist trajectory with a Gaussian kernel [19] ($\sigma=2.0$) and compute wrist speed as the displacement between consecutive frames. Candidate transitions are identified at time steps where the speed reaches a local minimum within a 1-second window and the orientation changes by more than 30° in the window.

B. Step 2: 3D Object Reconstruction

We use an LLM (GPT-4o [2]) to extract the label of the object of interest from the given text prompt. The object is stable during the grasping phase, allowing us to reconstruct its 3D geometry from the video start frame to the transition frame. We uniformly sample eight frames from this interval and apply SAM3 [12] to segment the target object using the specified label. For each sampled frame, the highest-confidence mask is selected. In cases where SAM3 fails to produce a valid segmentation, we examine nearby frames within the same interval to recover a suitable mask.

Since SAM3D [17] reconstructs geometry from a single RGB image but does not recover metric scale, we first estimate metric depth for the initial frame using DepthAnything3 [60] and back-project it into a 3D point map using the known camera parameters. We then apply the segmentation mask to

extract the object point map, which provides the metric reference. Given this metrically scaled point map and the initial RGB frame, we reconstruct the object mesh with SAM3D. Given the pose of the RGB image, we further transform the reconstructed metric mesh from SAM3D into the transition frame’s world coordinates using the known camera extrinsics.

C. Step 3: Hand-Object Alignment

Despite sequential video observations from multiple cameras in EgoDex [37], misalignment persists between the reconstructed object and the tracked hand keypoints in the world frame. We align the hand and object based on two hypotheses: (i) the fingerpads should be in contact (at least three fingers including the thumb) with the object surface at the moment of grasp completion, and (ii) no deep penetration should occur.

At the transition frame between the grasping and manipulation stages, we first fit the MANO hand mesh from the tracked hand keypoints using optimization-based inverse kinematics (IK) [48], which allows us to extract the fingerpad vertices from the mesh. We then estimate an optimal object translation offset $\Delta p^o \in \mathbb{R}^3$ that minimizes a weighted fingerpad-to-object surface discrepancy at the transition frame by gradient descent. If the resulting grasp does not satisfy the contact hypotheses, we further fine-tune the MANO parameters correction $\Delta \theta$ with:

$$\min_{\Delta p^o, \Delta \theta} \mathcal{L}_{\text{dist}} + \lambda_{\text{pen}} \mathcal{L}_{\text{pen}} + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}}, \quad (\text{S1})$$

where $\Delta p^o \in \mathbb{R}^3$ is the object translation offset, $\Delta \theta \in \mathbb{R}^{24}$ is the MANO parameters correction in PCA space, $\mathcal{L}_{\text{dist}}$ penalizes MANO fingerpad to object surface distances, \mathcal{L}_{pen} penalizes negative signed distances of fingerpad vertices, and \mathcal{L}_{reg} regularizes deviation from the IK-fitted pose.

Finally, we propagate both the estimated object translation offset Δp^o and the optimized MANO pose correction $\Delta \theta$ at the transition frame to all frames prior to the transition, resulting in aligned hand-object data for the entire grasping phase.

D. DATASET PREPROCESSING DETAILS

This section details our preprocessing and annotation procedures for each dataset used in Sec. V-A of the main paper.

A. GRAB

GRAB [95] is a comprehensive full-body grasping and manipulation dataset containing 1,335 sequences of human interactions with 51 everyday objects. Each sequence provides high-quality 3D motion capture data for both the human body and hands, together with accurate object meshes and trajectories. Although GRAB does not include scene information, we include it in our evaluation due to the high fidelity of its hand-object interaction data and the complexity of the manipulation actions it contains.

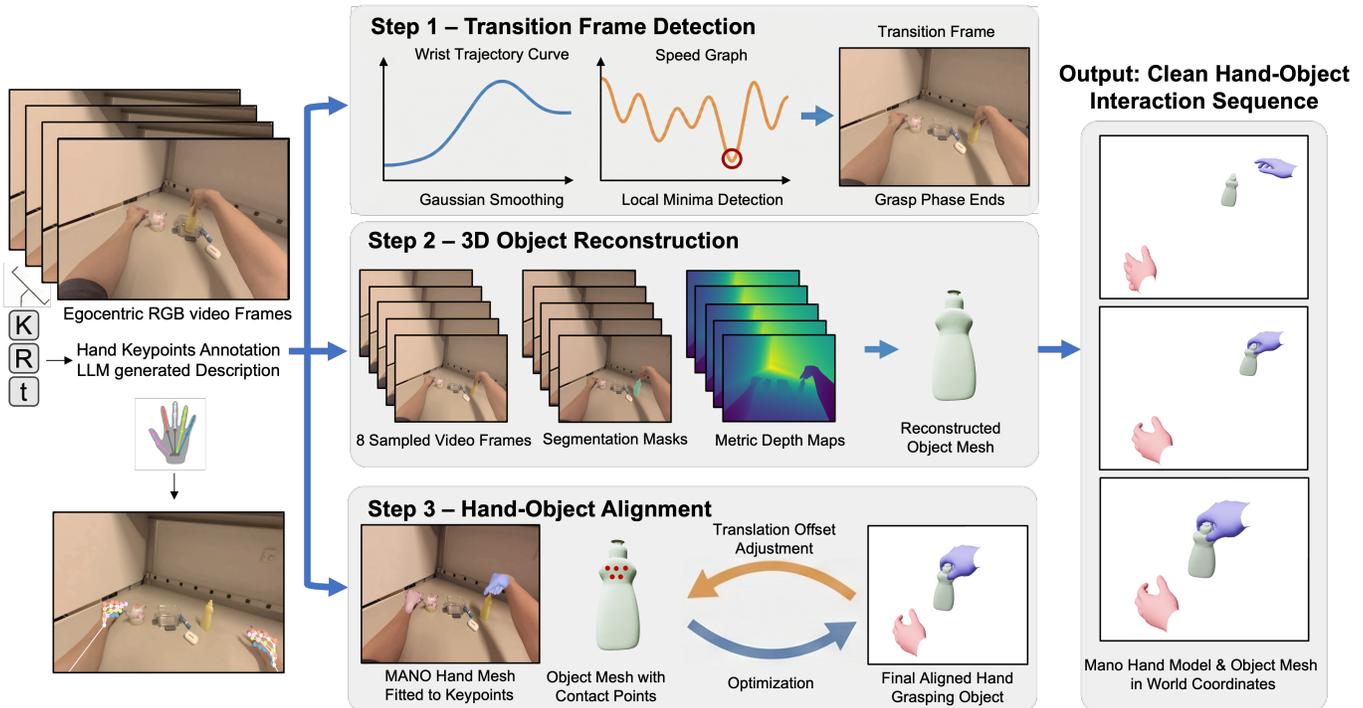


Fig. 7: **Detailed hand-object data reconstruction pipeline.** Given an egocentric RGB video with tracked hand keypoints and camera parameters, we reconstruct HOI data in three steps: (1) *Transition frame detection*: we smooth wrist trajectories and identify the grasp-to-manipulation transition via local speed minima and direction change; (2) *3D object reconstruction*: we segment the target object with SAM3 [12], estimate metric depth with DepthAnything3 [60], and reconstruct a mesh with SAM3D [17] using pre-transition frames where the object is static; (3) *Hand-object alignment*: we fit MANO hand meshes via inverse kinematics and optimize an object translation offset to satisfy fingerpad contact and non-penetration constraints at the transition frame, then propagate the alignment to all frames.

Following [18], we focus exclusively on the hand-object interaction components and discard full-body motion information. We adopt the dataset split (1,125 train / 210 test) and textual action descriptions introduced in [18]. The raw motion capture data is downsampled from ~ 120 Hz to 20 Hz.

B. HOT3D

For training our semantically grounded and scene-aware manipulation model, there is currently no dataset that simultaneously satisfies all of the following requirements: (i) high-quality 3D annotations for hand-object interactions [14, 95], (ii) high-fidelity real-world scene capture comparable to large-scale indoor datasets [105, 91], and (iii) explicit action-level semantic descriptions.

HOT3D [5] is close to meeting these requirements. The dataset provides accurate 3D hand and object pose annotations synchronized with egocentric videos, enabling detailed analysis of fine-grained manipulation activities.

In our experiment, we use the publicly available training split of HOT3D, as the ground-truth annotations for the official test set have not been released and are primarily intended for benchmark evaluation. The training split consists of 136 recordings collected from nine subjects, with each recording lasting approximately two minutes and containing multiple hand-object interactions.

To obtain scene representations, we reconstruct each record-

ing using 3DGS [46, 68] from the recorded egocentric videos captured by the Project Aria glasses [22]. We then automatically extract hand-object interaction clips by detecting object motion and temporally segmenting the sequences accordingly. Specifically, each interaction clip is divided into a grasping stage and a manipulation stage. The grasping stage is defined as the interval from the end of the previous interaction to the moment of grasp completion. The manipulation stage covers the remaining period where the hand actively manipulates the object.

Following this procedure, we obtain 1,802 interaction clips in total, split into 1,441 for training and 361 for testing, where the test set contains interactions involving unseen objects. Finally, we use OpenAI GPT-4o [2] to generate natural-language descriptions for each extracted interaction clip, including the objects involved and the actions performed.

E. PHYSICS SIMULATION ANALYSIS

This section provides additional details on the physics simulation experiments described in Sec. V-C of the main paper.

We evaluate the realizable physical feasibility of generated HOI sequences in Isaac Gym [70]. The evaluation pipeline consists of three steps: (1) retargeting the generated MANO [92] hand trajectories to the Allegro Hand [1] via inverse kinematics [48], (2) executing the retargeted motions

using a physics-based tracking controller [63] that produces robot joint torques, and (3) evaluating whether the object remains stably grasped and the task is completed. We describe each step in detail below.

Retargeting. Given the generated MANO hand trajectories, we first convert the MANO parameters to 21 keypoint positions per hand at each frame via forward kinematics. Since the Allegro Hand is approximately $1.6\times$ larger than the human hand, we uniformly scale the MANO keypoints accordingly to match the robot hand size. We then solve for the 16-DOF Allegro joint angles and the 6-DOF wrist pose via inverse kinematics using pyroki [48], a JAX-based least-squares IK solver. The optimization minimizes a weighted combination of cost terms: (i) a *local alignment* cost that preserves relative inter-joint distances and angles to maintain hand structure, (ii) a *global alignment* cost that matches the robot link positions to the target keypoints, (iii) *wrist position and rotation* costs that anchor the robot wrist to the MANO wrist pose using translation error and $SO(3)$ geodesic distance respectively, and (iv) *temporal smoothness* costs on both the joint angles and the root trajectory to suppress jitter. Joint angle limits from the Allegro URDF are enforced as hard constraints. Since the Allegro Hand has four fingers while MANO models five, we approximate the pinky targets by offsetting the ring finger keypoints laterally.

Tracking. We use DexTrack [63], a reinforcement learning-based tracking controller, to execute the retargeted trajectories under physics simulation. The controller is trained with Proximal Policy Optimization (PPO) in Isaac Gym [70] across thousands of parallel environments. The policy observes the current hand joint positions and velocities, fingertip states, object pose, and the residual between the current state and the reference trajectory at the current and future timesteps. It outputs cumulative residual position targets for the 16 actuated hand joints, which are tracked by low-level PD controllers. The reward function encourages (i) minimizing the distance between the fingertips and the object surface, (ii) tracking the reference hand pose with separate coefficients for global translation, wrist rotation, and finger joint angles, (iii) matching the target object position and orientation, and (iv) smooth joint velocity profiles. To improve robustness, we apply domain randomization over hand joint stiffness and damping, hand and object masses, friction coefficients, and observation and action noise during training.

Evaluation. As shown in Fig. 8, although the kinematic trajectories of both the human hand (a) and the retargeted Allegro hand (b) appear visually plausible, executing them under physics simulation (c) can reveal failures: the object may detach or slip from the grasp due to insufficient contact forces, incorrect friction modeling, or dynamically unstable grasp configurations. This demonstrates that kinematic feasibility alone does not guarantee physically stable interactions, motivating the use of physics simulation as a more rigorous evaluation protocol.

For each method, we run 10 simulation trials per test sequence and report the average success rate and holding time.

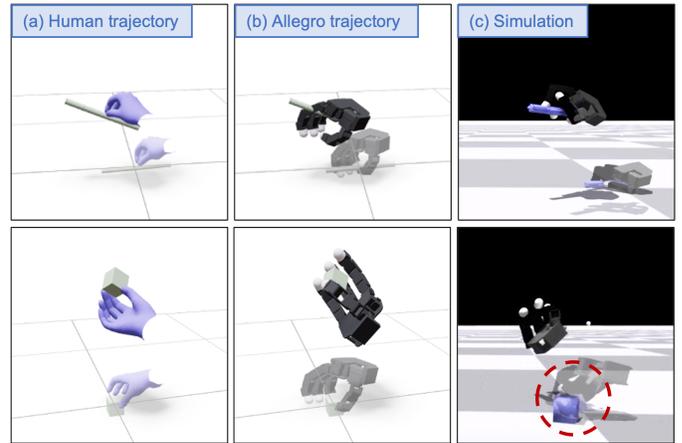


Fig. 8: Comparison of HOI trajectories: (a) kinematic human hand trajectory, (b) retargeted Allegro hand trajectory, and (c) physics simulation in Isaac Gym [70]. Although the kinematic trajectories in (a) and (b) appear plausible, executing them under physics simulation reveals failures: the object detaches and slips from the grasp (bottom row, red circle), indicating that kinematically feasible motions do not guarantee physically stable interactions.

An episode is considered successful if the object remains within 5 cm of its target position while maintaining finger contact throughout the trajectory. All simulation parameters follow the default settings from DexTrack [63].

F. REAL-WORLD EXPERIMENT DETAILS

This section provides additional details on the real-world dexterous manipulation experiments showcased in the main paper.

Hardware Setup. Our platform consists of two Franka Emika Panda robotic arms [33], each equipped with an Allegro Hand v5 [1]. The 3D scene is reconstructed offline prior to each experiment: we capture a multi-view video sequence using a MetaCam [93] scanner and reconstruct the scene via Gaussian-LIC [51], yielding a 3DGS representation consistent with our training pipeline.

Retargeting Pipeline. Given the generated MANO hand trajectories, we decompose the retargeting into two stages: (1) *Arm retargeting*: the generated MANO wrist 6D pose at each frame is used as the end-effector target for the Franka arm. We solve the 7-DOF joint angles using the Franka built-in inverse kinematics solver. (2) *Hand retargeting*: the MANO finger articulations are retargeted to the 16-DOF Allegro Hand joint angles using pyroki [48], following the same procedure described in Sec. E. The retargeted trajectories are further refined offline by the DexTrack [63] sim-to-real policy, which was trained entirely in Isaac Gym [70] and deployed without additional fine-tuning on the real robot.

Execution. The refined joint trajectories are executed in an open-loop fashion via a standard joint impedance controller on both the Franka arms and the Allegro hands at 1 kHz. Since the execution is open-loop, no real-time object state feedback is required during task execution.

Evaluation. We qualitatively evaluate the system on four contact-rich household manipulation tasks: drinking from a cup, pouring liquid between containers of different sizes, tilting a container, and squeezing dressing. Across all tasks, the generated HOI trajectories are consistently retargeted and successfully executed, producing stable contact-rich interactions that match the intended behaviors. We refer the reader to Fig. 5 in the main paper and the supplementary video for visual results.

G. SCENE REPRESENTATION DETAILS AND ANALYSIS

This section provides additional details on the scene representation described in Sec. IV-C of the main paper.

Fig. 9 visualizes the geometric features (Concerto [108]) and semantic features (SceneSplat [58]) of two example scenes in top-down view, colored by PCA of the respective feature spaces. The geometric features (left column) primarily encode spatial properties such as surface normals, curvature, and local point distributions. As highlighted by the black dashed circles, geometric features can distinguish objects on a table from the table surface itself based on their different spatial structures. The semantic features (right column) encode language-aligned semantics, grouping points by object category or functional role. Points on the same object share coherent representations regardless of spatial location, while semantically different objects are clearly separated (black dashed circles). This complementarity motivates our bidirectional fusion: geometric features provide precise spatial grounding for localization and collision avoidance, while semantic features enable the model to identify task-relevant objects for text-conditioned generation.

H. EVALUATION METRIC DETAILS

This section provides detailed definitions and implementation details of the evaluation metrics reported in Sec. V-C of the main paper.

A. Physical Interaction Quality

Let \mathcal{M}_h^t and \mathcal{M}_o^t denote the hand and object meshes at frame t , respectively. For a hand mesh vertex \mathbf{v} , we denote by $d(\mathbf{v}, \mathcal{M}_o^t)$ the signed distance to the object surface, where negative values indicate penetration. All physical interaction metrics (IV, ID, CR, IVU) are computed independently for each hand and then averaged across hands.

Interpenetration Volume (IV). Interpenetration volume measures the volumetric intersection between the hand and object meshes:

$$IV_t = \text{Vol}(\mathcal{M}_h^t \cap \mathcal{M}_o^t), \quad (\text{S2})$$

and is averaged over time:

$$IV = \frac{1}{T} \sum_{t=1}^T IV_t. \quad (\text{S3})$$

Interpenetration Depth (ID). Interpenetration depth is defined as the mean penetration distance across all penetrating hand vertices:

$$ID_t = \frac{1}{|\mathcal{V}_{\text{pen}}^t|} \sum_{\mathbf{v} \in \mathcal{V}_{\text{pen}}^t} \max(0, -d(\mathbf{v}, \mathcal{M}_o^t)), \quad (\text{S4})$$

where $\mathcal{V}_{\text{pen}}^t = \{\mathbf{v} \in \mathcal{V}_h \mid d(\mathbf{v}, \mathcal{M}_o^t) < 0\}$ is the set of penetrating vertices at frame t . ID is temporally averaged over all frames.

Contact Ratio (CR). The contact ratio quantifies the proportion of hand vertices that are close to the object surface:

$$\text{CR}_t = \frac{1}{|\mathcal{V}_h|} \sum_{\mathbf{v} \in \mathcal{V}_h} \mathbb{I}(d(\mathbf{v}, \mathcal{M}_o^t) \leq \delta), \quad (\text{S5})$$

where δ is a fixed distance threshold (set to 5 mm in all experiments). The final CR score is obtained by averaging over interaction frames (frames where hand-object contact occurs).

Interpenetration Volume per Contact Unit (IVU). Since low penetration can be trivially achieved by avoiding contact, we additionally report interpenetration volume normalized by the contact area at each frame. For each contact frame t (where $\text{CR}_t > 0$), we compute the per-frame IVU as:

$$\text{IVU}_t = \frac{IV_t}{A_{\text{contact}}^t}, \quad A_{\text{contact}}^t = \text{CR}_t \cdot \text{Area}(\mathcal{M}_h), \quad (\text{S6})$$

and the final IVU is the mean over all contact frames:

$$\text{IVU} = \frac{1}{|\mathcal{T}_c|} \sum_{t \in \mathcal{T}_c} \text{IVU}_t, \quad (\text{S7})$$

where $\mathcal{T}_c = \{t \mid \text{CR}_t > 0\}$ is the set of frames with non-zero contact.

B. Motion Quality

Action Recognition Accuracy (AR). To evaluate semantic correctness, we apply a pretrained action recognition classifier f_{act} to generated motion sequences. Given a generated motion $\hat{\mathbf{x}}_i$ with target action label y_i , AR is computed as:

$$\text{AR} = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(\arg \max f_{\text{act}}(\hat{\mathbf{x}}_i) = y_i). \quad (\text{S8})$$

Sample Diversity (SD). For each conditioning input, we generate K motion samples and compute the mean pairwise ℓ_2 distance between their hand joint trajectories, normalized by the number of frames:

$$\text{SD} = \frac{2}{K(K-1)} \sum_{1 \leq i < j \leq K} \frac{1}{T'} \left\| \mathbf{h}^{(i)} - \mathbf{h}^{(j)} \right\|_2, \quad (\text{S9})$$

where $\mathbf{h}^{(i)} \in \mathbb{R}^{2J \cdot 3 \cdot T'}$ is the flattened joint position trajectory of the i -th sample. We use $K = 4$ in our experiments.

Overall Diversity (OD). Overall diversity is computed as the mean pairwise ℓ_2 distance between hand joint trajectories across all generated test samples, normalized by the number of frames:

$$\text{OD} = \frac{2}{S(S-1)} \sum_{1 \leq i < j \leq S} \frac{1}{T'} \left\| \mathbf{h}_i - \mathbf{h}_j \right\|_2, \quad (\text{S10})$$

where S is the total number of generated motions.

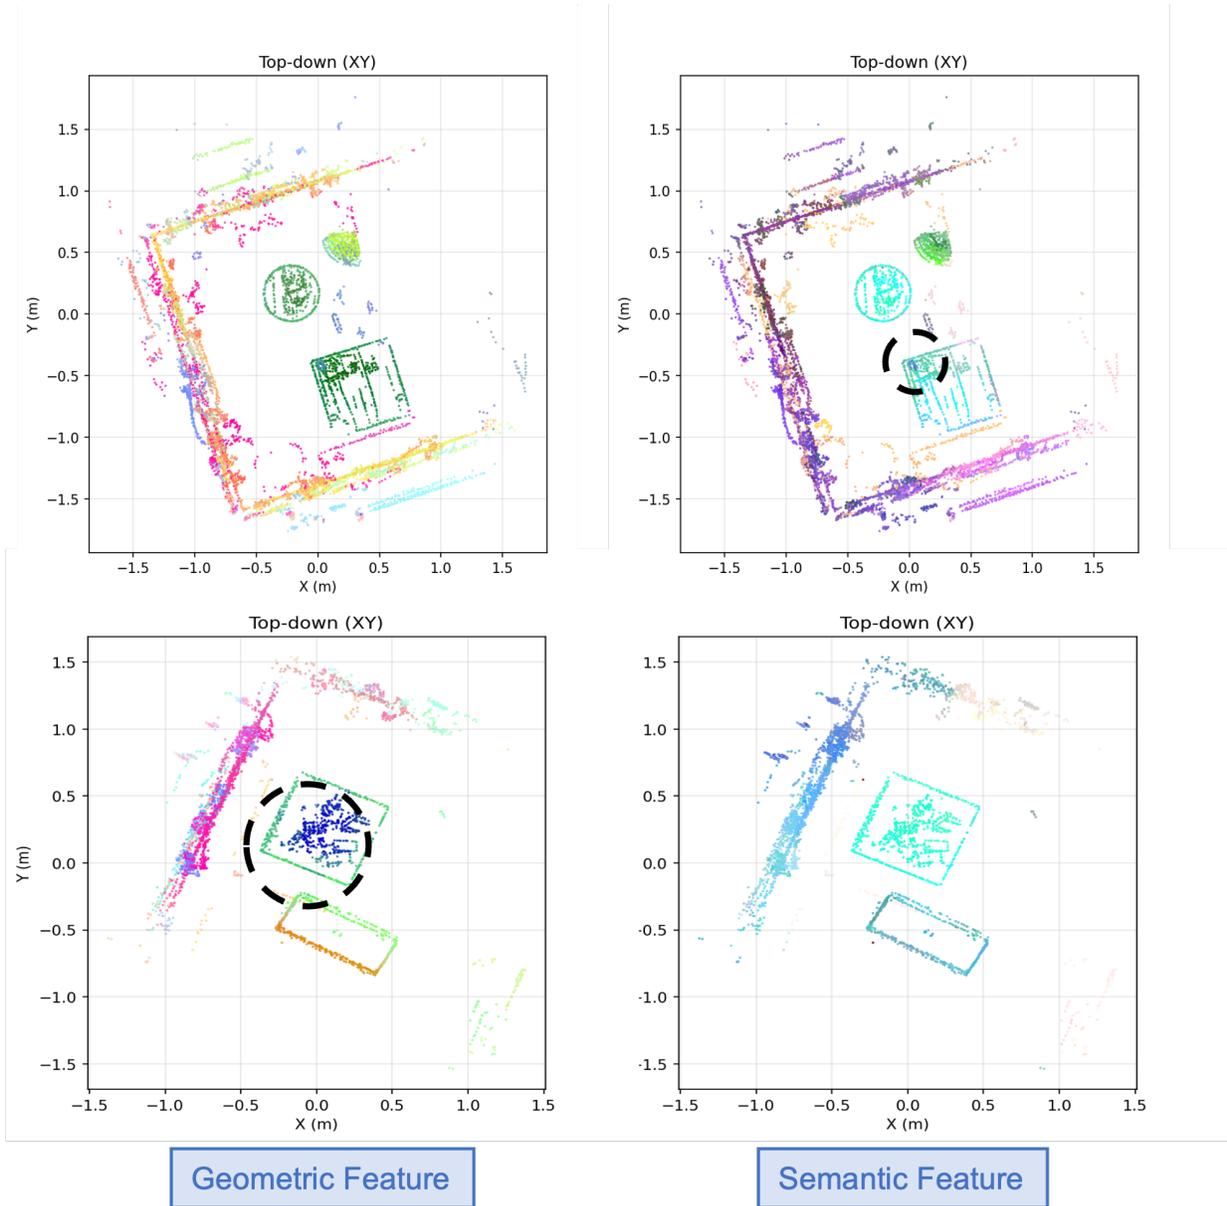


Fig. 9: **Scene feature analysis.** Top-down views of two scenes comparing geometric features (Concerto [108], left) and semantic features (SceneSplat [58], right), colored by PCA of the respective feature spaces. Black dashed circles highlight where geometric features successfully distinguish objects from the supporting surface (bottom-left) and where semantic features successfully discriminate semantically distinct regions (top-right). This complementarity motivates our bidirectional fusion.

C. Realizable Physical Feasibility

Physical Plausibility (Phy). Following LatentHOI [54], physical plausibility is a heuristic, per-sequence binary score. A sequence is deemed physically plausible if (i) at least one hand joint remains within a signed distance of 5 mm to the object surface for every frame during the interaction phase, and (ii) the object stays above the ground plane throughout the sequence. The final Phy score is the percentage of test sequences satisfying both criteria. Note that this metric only checks kinematic proximity and does not account for dynamic stability; as discussed in Sec. E, a high Phy score does not necessarily imply physically realizable interactions.

Success Rate (SR). For simulation evaluation, a trial is

considered successful if the object is lifted above the ground plane by at least 5 cm while maintaining contact with the hand. The success rate is the fraction of such sequences:

$$\text{SR} = \frac{1}{N} \sum_{n=1}^N s_n, \quad (\text{S11})$$

where $s_n \in \{0, 1\}$ indicates whether the object in the n -th trial is lifted off the ground (> 5 cm) with sustained hand-object contact.

Holding Time (HT). Holding time measures the duration during which objects are stably held in the hand and is

averaged over successful trials:

$$\text{HT} = \frac{\sum_{n=1}^N s_n \tau_n}{\sum_{n=1}^N s_n}, \quad (\text{S12})$$

where τ_n denotes the duration during which the object is stably held in the n -th trial.

D. Ablation-Specific Metrics

Grasp Error (GE). Grasp error measures the mean per-joint Euclidean distance between the generated and ground-truth hand poses at the transition frame t_g :

$$\text{GE} = \frac{1}{N_{\text{test}}} \sum_{i=1}^{N_{\text{test}}} \frac{1}{|\mathcal{J}^{(i)}|} \sum_{j \in \mathcal{J}^{(i)}} \left\| \hat{\mathbf{j}}_{t_g, j}^{(i)} - \mathbf{j}_{t_g, j}^{(i)} \right\|_2, \quad (\text{S13})$$

where $\hat{\mathbf{j}}_{t_g, j}^{(i)}, \mathbf{j}_{t_g, j}^{(i)} \in \mathbb{R}^3$ are the generated and ground-truth 3D positions of the j -th MANO joint at the transition frame, and $\mathcal{J}^{(i)}$ is the set of joints belonging to the grasping hand(s), determined by proximity to the object at the ground-truth transition frame.

Final Displacement Error (FDE). Final displacement error measures the Euclidean distance between the generated and ground-truth object poses at the end of the manipulation stage:

$$\text{FDE} = \frac{1}{N_{\text{test}}} \sum_{i=1}^{N_{\text{test}}} \left\| \hat{\mathbf{O}}_{N-1}^{(i)} - \mathbf{O}_{N-1}^{(i)} \right\|_2, \quad (\text{S14})$$

where $\hat{\mathbf{O}}_{N-1}^{(i)}$ and $\mathbf{O}_{N-1}^{(i)}$ denote the generated and ground-truth object poses at the final frame.

Jerk. Jerk measures the smoothness of generated object motion as the third derivative of position with respect to time, i.e., the rate of change of acceleration. Human motion naturally follows a *minimum jerk* principle [24], and lower jerk values indicate smoother, more natural trajectories. Given the object trajectory $\mathbf{p}^t \in \mathbb{R}^3$ over T uniformly sampled frames, we compute the positional jerk via finite differences:

$$\text{Jerk}_{\text{pos}} = \frac{1}{T} \sum_{t=1}^T \left\| \frac{d^3 \mathbf{p}^t}{dt^3} \right\|_2, \quad (\text{S15})$$

where derivatives are approximated using central differences (`np.gradient`). We analogously define the angular jerk Jerk_{ang} over the object rotation trajectory $\boldsymbol{\theta}^t \in \mathbb{R}^3$ (axis-angle representation).

Hand Acceleration. To assess the stability of generated hand motion, we compute acceleration metrics that separately evaluate global wrist movement and local finger articulation. Given the MANO joint positions $\mathbf{J}^t \in \mathbb{R}^{J \times 3}$ at frame t , we decompose them into wrist positions $\mathbf{p}_w^t \in \mathbb{R}^3$ and finger positions $\mathbf{p}_f^t \in \mathbb{R}^{(J-1) \times 3}$. Local finger positions are computed relative to the wrist: $\tilde{\mathbf{p}}_f^t = \mathbf{p}_f^t - \mathbf{p}_w^t$, isolating finger articulation from global wrist movement. We then compute acceleration via second-order finite differences and report the mean absolute acceleration:

$$\text{Acc}_g^{\text{pos}} = \frac{1}{T} \sum_t \left| \frac{d^2 \mathbf{p}_w^t}{dt^2} \right|, \quad \text{Acc}_l^{\text{pos}} = \frac{1}{T} \sum_t \left| \frac{d^2 \tilde{\mathbf{p}}_f^t}{dt^2} \right|, \quad (\text{S16})$$

where $|\cdot|$ denotes element-wise absolute value and the result is averaged over all spatial components and both hands. The rotational counterparts $\text{Acc}_g^{\text{rot}}$ and $\text{Acc}_l^{\text{rot}}$ are defined analogously over the wrist and finger axis-angle rotations, respectively. Separating global and local components ensures that wrist trajectory smoothness and finger articulation stability are evaluated independently.

I. ADDITIONAL QUALITATIVE RESULTS

We provide additional qualitative results to complement the main paper. Fig. 10 provides qualitative comparisons with DiffH2O [18] and LatentHOI [54] on GRAB [95]. Fig. 11 presents scene-conditioned results on HOT3D [5], demonstrating that the generated trajectories adapt to the surrounding scene layout. **We strongly encourage the reader to view the supplementary video for animated results.**

J. FAILURE CASES

We identify two representative failure modes of our framework (see Fig. 12):

(a) Object Size Mismatch. When the object is too large or too small relative to the hand, the generated grasp may not fully enclose the object, resulting in the object appearing to float in the air without stable contact. A potential improvement is to incorporate explicit object size conditioning or adaptive grasp aperture prediction based on object geometry.

(b) Incorrect Contact Establishment. The hand and object may fail to establish proper contact during the grasping phase, leading to physically implausible configurations where fingers do not properly wrap around the object surface. Future work could incorporate contact-aware losses or physics-based refinement during inference to ensure proper finger-object contact. Additionally, improving the quality of training data with more accurate contact annotations could also help mitigate this issue.

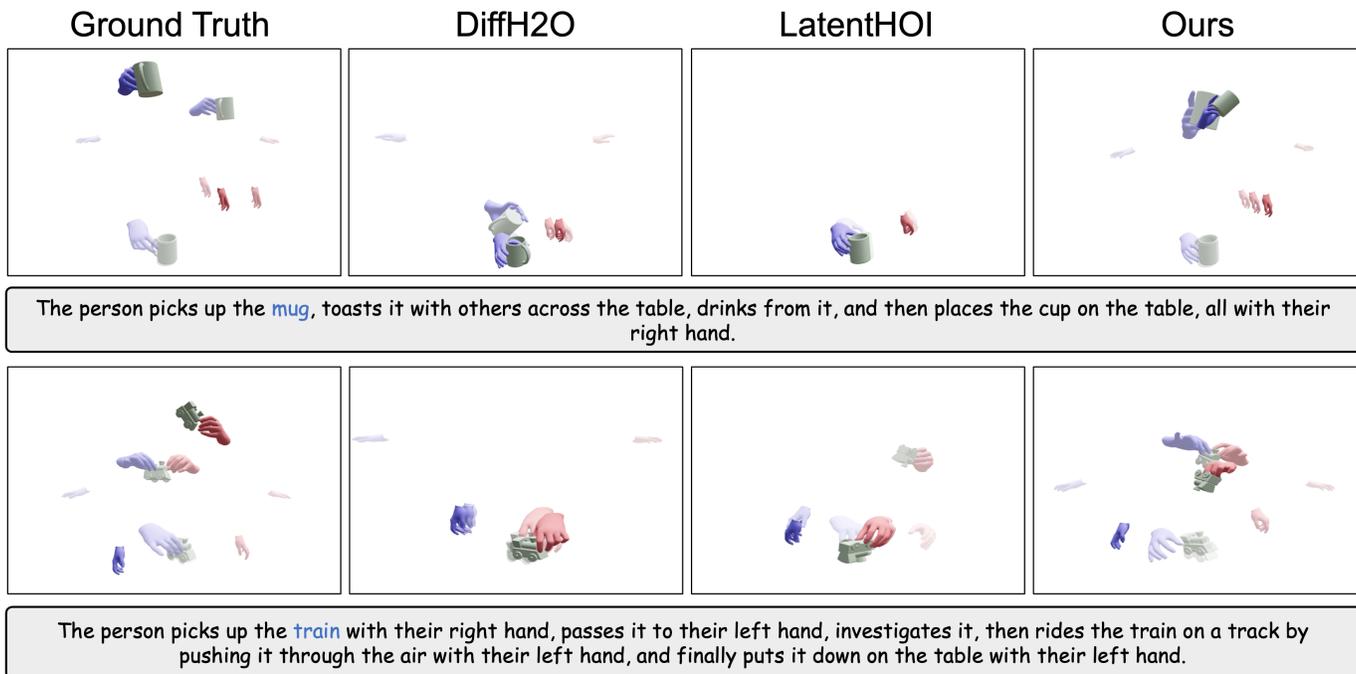


Fig. 10: **Qualitative comparison on GRAB.** Comparison with DiffH2O [18] and LatentHOI [54] on diverse actions. Our method produces more physically plausible hand-object interactions with accurate grasping poses and smooth bimanual coordination.

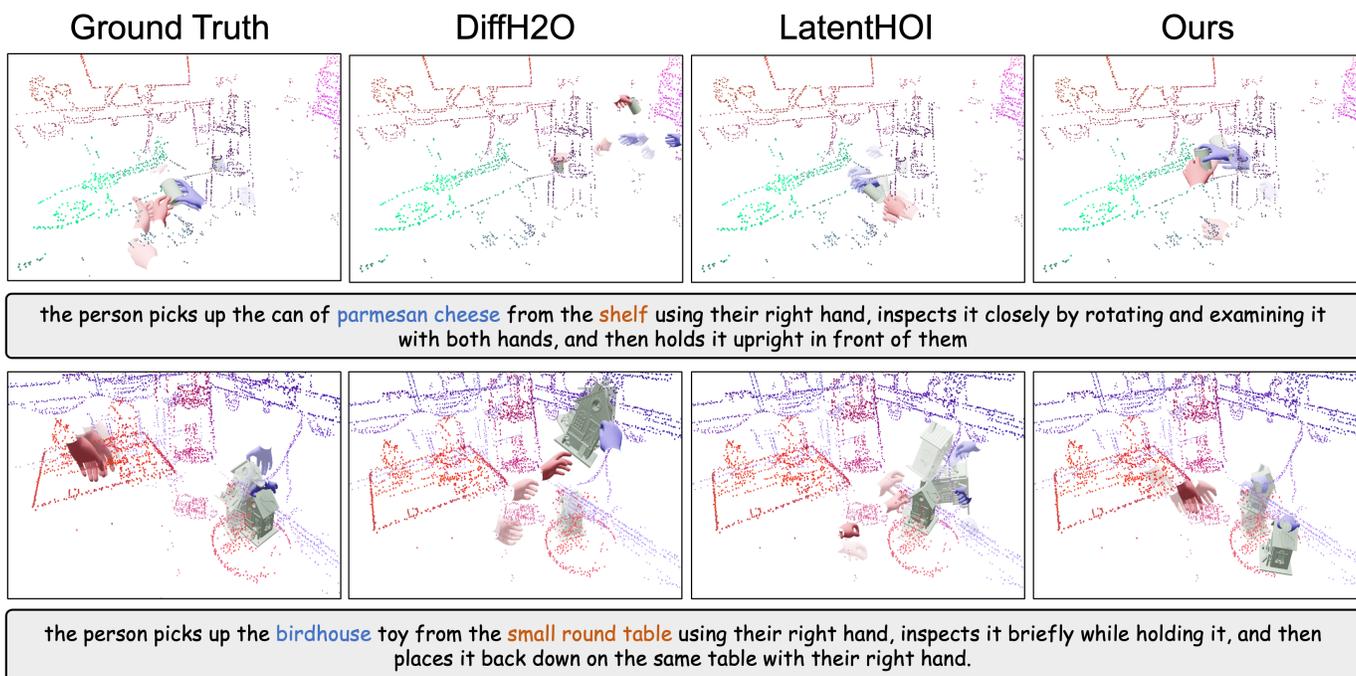


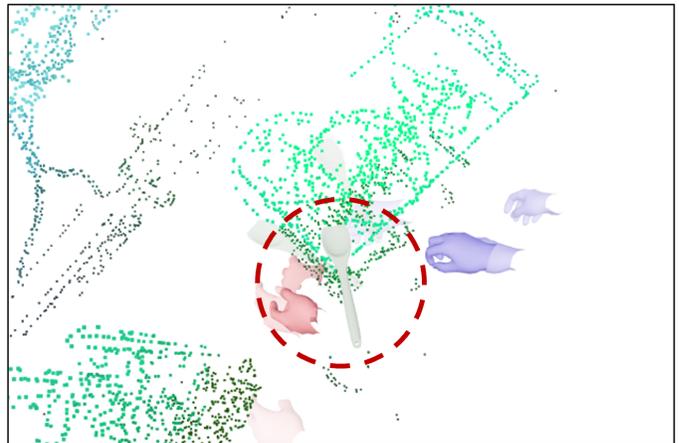
Fig. 11: **Additional qualitative results on HOT3D with scene context.** The generated hand trajectories adapt to the surrounding scene layout, correctly interacting with objects on shelves and tables.

(a) Object Size Mismatch



The person picks up the **small pyramid** from the high table using their right hand, passes it to their left hand, then hands it to someone else on their left side with their left hand, and finally puts it back on the table using their right hand.

(b) Incorrect Contact Establishment



The person picks up the **wooden spoon** from the **table** with their left hand, holds it while interacting with a bowl of colorful objects using their right hand, and then places the spoon onto a shelf.

Fig. 12: **Representative failure cases.** Red dashed circles highlight the erroneous regions. (a) *Object size mismatch*: the small pyramid floats away from the hand due to a mismatch between object size and generated grasp aperture. (b) *Incorrect contact establishment*: the hand fails to properly wrap around the wooden spoon, resulting in physically implausible finger-object contact. Left: GRAB [95]; right: HOT3D [5] with scene context.