Uni-Hand: Universal Hand Motion Forecasting in Egocentric Views

Junyi Ma, Wentao Bao, Jingyi Xu, Guanzhong Sun, Yu Zheng, Erhang Zhang, Xieyuanli Chen, and Hesheng Wang, *Senior Member, IEEE*

***Abstract*—Forecasting how human hands move in egocentric views is critical for applications like augmented reality, human-robot policy transfer, and service/assistive technologies. Recently, several hand trajectory prediction (HTP) methods have been developed to generate future possible hand waypoints, which still suffer from insufficient prediction targets, inherent modality gaps, entangled hand-head motion, and limited validation in downstream tasks. To address these limitations, we present Uni-Hand, a universal hand motion forecasting framework considering multi-modal input, multi-dimensional and multi-target prediction patterns, and multi-task affordances for downstream applications. We harmonize multiple modalities by vision-language fusion, global context incorporation, and task-aware text embedding injection, to forecast hand waypoints in both 2D and 3D spaces. A novel dual-branch diffusion is proposed to concurrently predict human head and hand movements, capturing their motion synergy in egocentric vision. By introducing target indicators, the prediction model can forecast the specific joint waypoints of the wrist or the fingers, besides the widely studied hand center points. In addition, we enable Uni-Hand to additionally predict hand-object interaction states (contact/separation) to facilitate downstream tasks better. To incorporate comprehensive downstream task evaluations in the literature, we build novel benchmarks to assess the real-world applicability of hand motion forecasting algorithms. The experimental results on multiple publicly available datasets and our newly proposed benchmarks demonstrate that Uni-Hand achieves the state-of-the-art performance in multi-dimensional and multi-target hand motion forecasting. Extensive validation in multiple downstream tasks also presents its impressive human-robot policy transfer to enable robotic manipulation, and effective feature enhancement for action anticipation/recognition.**

***Index Terms*—Hand motion forecasting, hand trajectory prediction, egocentric vision, diffusion models.**

I. Introduction

PREDICTING human behaviors over future time horizons is critical for applications such as augmented reality and human-robot interaction. As an important interface for human-environment interaction, the human hand offers valuable motion knowledge that facilitates various downstream tasks in embodied intelligence [1], [2], [3], [4]. However, despite significant progress in other human behavior anticipation tasks like predicting gaze locations [5], [6], [7], [8], and interaction regions [9], [10], [11], forecasting fine-grained human hand motion still remains challenging. To align better with hand movement patterns and potential intentions from a human perspective, some egocentric-vision-only hand motion forecasting (HMF) approaches [12], [13], [14], [15] have been developed in recent years. They receive past human egocentric (first-person) observations and directly predict future hand trajectories in the specified reference frame. However, these methods exhibit significant limitations that hinder their further performance improvement and effectiveness in possible downstream applications:

- *Insufficient prediction targets:* The existing hand motion forecasting paradigms exclusively concentrate on the generation of future hand centers (of the potential bounding boxes) as target motion waypoints [14], [15], [16]. Despite its necessity for providing hand movement trends, these methods basically overlook how different joints of the fingers and wrist move in 2D and 3D spaces. Besides, they do not involve locating the precise timings when the hand will make contact with and separate from target objects. Therefore, there is a lack of sufficient motion knowledge in future time horizons for practical deployments, such as human-robot manipulation policy transfer.
- *Modality gaps:* Existing HMF paradigms predominantly operate on 2D egocentric video inputs, without explicitly incorporating metric-scale 3D spatial observations. While some prior works have shown that temporal modeling, pose estimation, and learned priors can partially infer potential 3D dynamics/structural information from 2D inputs [16], [17], [18], [19], relying only on 2D visual cues without access to raw 3D perception may introduce geometric ambiguity in fine-grained 3D motion forecasting and physically grounded interactions. In

Received 24 June 2025; revised 21 April 2026; accepted 9 May 2026. Date of publication 13 May 2026; date of current version 4 September 2026. This work was supported in part by the National Key R&D Program of China under Grant 2024YFB4708900, in part by the Explorers Program of Shanghai (Basic Research Funding) under Grant 25TS1415100, in part by the Shanghai Municipal Special Program for Basic Research on General AI Foundation Models under Grant 2025SHZDZX025G01, and in part by the Natural Science Foundation of China under Grant 62225309, Grant U24A20278, and Grant 62361166632. Recommended for acceptance by K. I. Kim. *(Corresponding author: Hesheng Wang.)*

Junyi Ma, Yu Zheng, Erhang Zhang, and Hesheng Wang are with the School of Automation and Intelligent Sensing, Shanghai Jiao Tong University, Shanghai 200240, ChinaState Key Laboratory of Avionics Integration and Aviation System-of-Systems Synthesis, Shanghai Key Laboratory of Navigation and Location Based Services, Shanghai 200240, China (e-mail: wanghesheng@sjtu.edu.cn).

Wentao Bao is with Michigan State University, East Lansing, MI 48824 USA.

Jingyi Xu is with the Department of Electronic Engineering, Shanghai Jiao Tong University, Shanghai 200240, China.

Guanzhong Sun is with the China University of Mining and Technology, Xuzhou 221008, China.

Xieyuanli Chen is with the College of Intelligence Science and Technology, National University of Defense Technology, Changsha 410073, China.

We will release our code, pretrained models of Uni-Hand, and novel benchmarks at the project page: https://irmvlab.github.io/unihand.github.io.

This article has supplementary downloadable material available at https://doi.org/10.1109/TPAMI.2026.3692504, provided by the authors.

Digital Object Identifier 10.1109/TPAMI.2026.3692504

0162-8828 © 2026 IEEE. All rights reserved, including rights for text and data mining, and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

Authorized licensed use limited to: ShanghaiTech University. Downloaded on October 19,2026 at 15:54:43 UTC from IEEE Xplore. Restrictions apply.

addition, the absence of textual modality interfaces prevents explicitly incorporating known human intentions or instructions. It also exacerbates the discrepancy between the HMF model optimization and possible deployments in language-conditioned downstream tasks.

- *Lack of prediction-level head motion conditioning:* Human hand motion is closely linked to whole-body kinematics, where the head movement is its most straightforward concurrent motion to provide visual observations. Understanding how the head moves benefits hand motion forecasting. While prior works [15], [16], [18] encode *past* headset camera egomotion into hand state transition to couple head and hand features, they do not explicitly forecast *future* head motion to explicitly condition hand motion prediction. This impairs the prediction models' ability to comprehend future hand motion patterns in the egocentric view of incoming hand-head cooperation.
- *Limited validation in downstream tasks:* The existing HMF frameworks aim to reduce the errors between predicted waypoints and ground-truth human demonstrations. While this objective is well-justified, their capability to support downstream tasks should be further analyzed in this literature. Transferring human actions to robotic manipulation has captured substantial attention nowadays owing to the high accessibility of human demonstrations [1], [20]. In line with this trend, evaluating HMF models via robotic manipulation is particularly natural, as the predicted human hand trajectories from videos can be directly mapped to robot end-effector action planning, providing an interpretable HMF assessment for physical plausibility. Therefore, exploring the transferability of HMF models to robotic manipulation tasks warrants significant research attention. In addition, predicted hand motion contains rich human behavior semantics, which are strongly correlated with action categories. In contrast to designing task-specific ad-hoc decoders [12], [13], how predicted hand motion features directly enhance the existing framework of downstream action category anticipation and recognition needs to be extensively explored.

To address the above limitations, we propose a universal hand motion forecasting framework, dubbed *Uni-Hand*. As illustrated in Fig. 1, Uni-Hand absorbs multi-modal inputs to achieve multi-dimension/-target hand motion prediction, and effectively provides multi-task affordances for downstream applications. **To extend the prediction targets**, Uni-Hand introduces target indicators in denoising diffusion to forecast hand centers or specified joints in both 2D and 3D space. It also integrates an additional interaction state decoder to locate the timings of hand-object contact and separation events. We therefore present a more flexible HMF solution compared to the existing paradigms [13], [15], [18] that only attend to hand center prediction. **To narrow the modality gaps**, we incorporate multiple modalities as the inputs including 2D RGB images, 3D point clouds, past hand waypoints, and text prompts. This not only enhances the awareness of 3D structure information in our framework, but also provides a textual modality interface for integrating possible language instructions in real-world robotic applications. **To address the lack of prediction-level head motion conditioning**, we design a dual-branch diffusion to concurrently forecast head and hand movements, explicitly capturing their synergy for more accurate HMF. Ultimately, **for extensive validation in downstream tasks**, we deploy our framework on multiple real-world

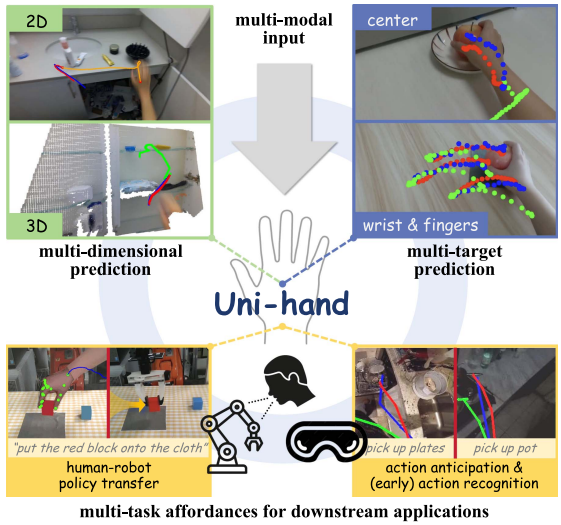


Fig. 1. Uni-Hand is a universal hand motion forecasting framework which facilitates multi-dimensional and multi-target predictions with multi-modal input. It also enables multi-task affordances for downstream applications.

robotic tasks, validating its capability to facilitate atomic skills, language-conditioned implementations, and long-horizon tasks. We also demonstrate its affordable motion features to improve the off-the-shelf framework for conventional action anticipation, early action recognition, and action recognition tasks.

In summary, the **main contributions** of this paper are four-fold:

- *Universal HMF framework:* We present Uni-Hand, a universal hand motion forecasting framework, featuring a versatile solution for multi-dimension/-target HMF and multi-task affordances for downstream applications.
- *Dual-branch diffusion:* We propose a novel dual-branch diffusion, which employs multi-modal information as input to concurrently predict future headset camera egomotion and hand movements in egocentric views.
- *Hybrid Mamba-Transformer denoising:* A novel hybrid Mamba-Transformer module is designed for denoising hand motion latents. It combines the strong temporal modeling ability of Mamba and the global context-awareness of Transformer to harmonize multi-modal features. We also inject text embeddings into the hybrid module to facilitate language-conditioned tasks.
- *HMF benchmark:* To the best of our knowledge, this is the first HMF research considering the assessment of hand forecasting algorithms' effectiveness in enabling comprehensive downstream tasks, including robotic manipulation, action anticipation, early action recognition, and action recognition. The extensive experiments on multiple publicly available datasets and our newly proposed benchmarks demonstrate that Uni-Hand achieves SOTA performance on 2D/3D hand motion forecasting, and can effectively adapt to the downstream tasks.

This work is an extension of the preliminary version [21], in which we only explored how to harmonize multiple input modalities for hand center prediction and achieve decoupling hand-head motion forecasting. In this study, we substantially expand in the following ways:

Authorized licensed use limited to: ShanghaiTech University. Downloaded on October 19,2026 at 15:54:43 UTC from IEEE Xplore. Restrictions apply.

- We extend the original 3D prediction paradigm to a universal framework for multi-dimension/-target HMF and multi-task affordances for downstream applications.
- We incorporate additional prediction targets besides hand centers, encompassing multiple hand joints and interaction states (contact/separation).
- A new textual modality interface (text embedding injection) is designed to enhance our model's awareness of specific task instructions.
- A comprehensive extension of experiments on public datasets and our new benchmarks presents the model's foundation capability to facilitate downstream tasks including real-world robotic manipulation, action anticipation, early action recognition, and action recognition.

This paper is organized as follows. Section II reviews related works in literature. Section III details Uni-Hand's architecture and training/inference schemes. Section IV showcases the extensive experimental results quantitatively and qualitatively. Section V highlights our algorithmic novelty and provides our insights. Ultimately, Section VI concludes the paper.

II. Related Work

A. Hand Trajectory Prediction in Egocentric Vision

Hand trajectory prediction (HTP) is a key component in understanding future hand motion patterns. Given sequential egocentric observations, HTP aims to forecast future possible hand waypoints in the 2D image plane or 3D space. Here, the hand waypoints denote the bounding box centers of detected hands. Motivated by the need for scaling capabilities, some 2D HTP approaches have been pioneeringly developed based on a large volume of human activity videos. For example, OCT [13] uses a Transformer [22] based encoder-decoder architecture to jointly predict hand waypoints and object hotspots. In contrast, Diff-IP2D [15] replaces the Transformer blocks in OCT with egomotion-aware diffusion models, achieving better 2D HTP performance. Similar to Diff-IP2D, EMAG [18] also innovatively encodes homography matrices of camera egomotion into 2D trajectory prediction with the cross-attention mechanism. These two works make initial efforts to integrate headset camera egomotion awareness into the 2D HTP pipeline. To further enhance the interpretability regarding fused camera egomotion, MADiff [16] incorporates Mamba [23] to model sequential hand state transition, based on a new motion-driven selective scan. Its powerful egomotion-aware Mamba block is out-of-the-box with respect to 2D predictive tasks in first-person views. More recently, HandsOnVLM [24] introduces a VLM scheme to transform 2D HTP into a VQA-style task. To further validate the applicability in real-world scenes, Bahl et al. [25] build an affordance model and deploy it for robot data collection and reward-free exploration.

Although 2D HTP is more scalable due to the availability of internet-scale human videos [16], 3D HTP yields more human-intent-aligned results. For instance, USST [14] receives image prompts to directly generate future 3D hand waypoints by autoregressive Transformers. EgoH4 [26] extends to body pose estimation, predicting locations of invisible hands. In contrast to EgoH4 focusing on whole body movements, our recent MMTwin [21] attends to the most critical body part, i.e., the head and hand, in hand-object interaction, forecasting their future motion concurrently. Compared to egomotion-aware 2D HTP [15], [16], [18] that only encodes past head motion, MMTwin builds an additional diffusion branch to predict future head motion latents.

Despite the significant progress in the above-mentioned 2D and 3D hand trajectory prediction, these methods basically only determine future centers of potential hand bounding boxes. However, the hand center constitutes an excessively coarse-grained prediction target, incapable of representing hand pose variations. This hinders fine-grained understanding of human behavior in hand-object interaction. Besides, these methods are agnostic to the potential timings of hand-object contact and separation. Therefore, they cannot be easily transferred to downstream tasks such as robot manipulation, since trajectories can only determine "where to move" rather than "when to grasp". In contrast, Uni-Hand proposed in this work can forecast future 2D and 3D positions of multiple hand joints besides center points, leading to awareness of hand pose changes in egocentric views. We introduce target indicators to let the denoising diffusion model distinguish different prediction targets. Moreover, we develop an additional decoder to forecast future interaction states, which can determine the hand-object contact/separation periods. Thus, Uni-Hand extends the conventional HTP literature to more informative hand motion forecasting by enriching prediction targets, and can practically facilitate multiple downstream applications.

B. Diffusion Models in Hand Motion Analysis

In recent years, hand motion analysis has attracted much attention due to the informational richness of hand-centric data for downstream tasks. While object-agnostic hand mesh recovery [27], [28], [29], [30] and object-aware 3D HOI reconstruction [31], [32], [33], [34] have witnessed significant advances, synthesizing reasonable hand poses around target objects and analyzing hand dynamics along the time axis are more challenging. With the advent of diffusion models [35], [36], some works such as HOIDiffusion [37], Gears [38], and DiffH2O [39] seamlessly integrate the powerful generation ability of diffusion models to generate hand meshes from scratch. However, they basically underestimate how the hand approaches the target object. More recently, PEAR [40] considers analyzing the hand motion trend during the interaction process. In contrast to PEAR, Tang et al. [41] develop a prompt-based future-driven diffusion model for predicting future 3D coordinates of multiple hand joints. Despite being more fine-grained, this method does not attend to wide-range wrist movements. Conversely, Diff-IP2D [15] and MADiff [16] both utilize DiffuSeq [42]-style models to forecast future wide-range hand movements, but lose detailed joint motion as mentioned before.

This work also adapts the advanced generation ability of diffusion models, to achieve multi-dimensional and multi-target future hand motion forecasting. Unlike the existing diffusion-based schemes, we propose a novel hybrid Mamba-Transformer architecture as the denoising model to harmonize multi-modal input. Besides, we introduce camera egomotion analysis within the dual-branch diffusion to enable more comprehensive hand motion understanding.

C. Downstream Tasks of Hand Motion Forecasting

As hands are fundamental to physical interaction in human daily life, they have driven growing research interest in human-robot policy transfer for manipulation tasks. Traditional imitation learning [43], [44] techniques rely on costly teleoperation

Authorized licensed use limited to: ShanghaiTech University. Downloaded on October 19,2026 at 15:54:43 UTC from IEEE Xplore. Restrictions apply.

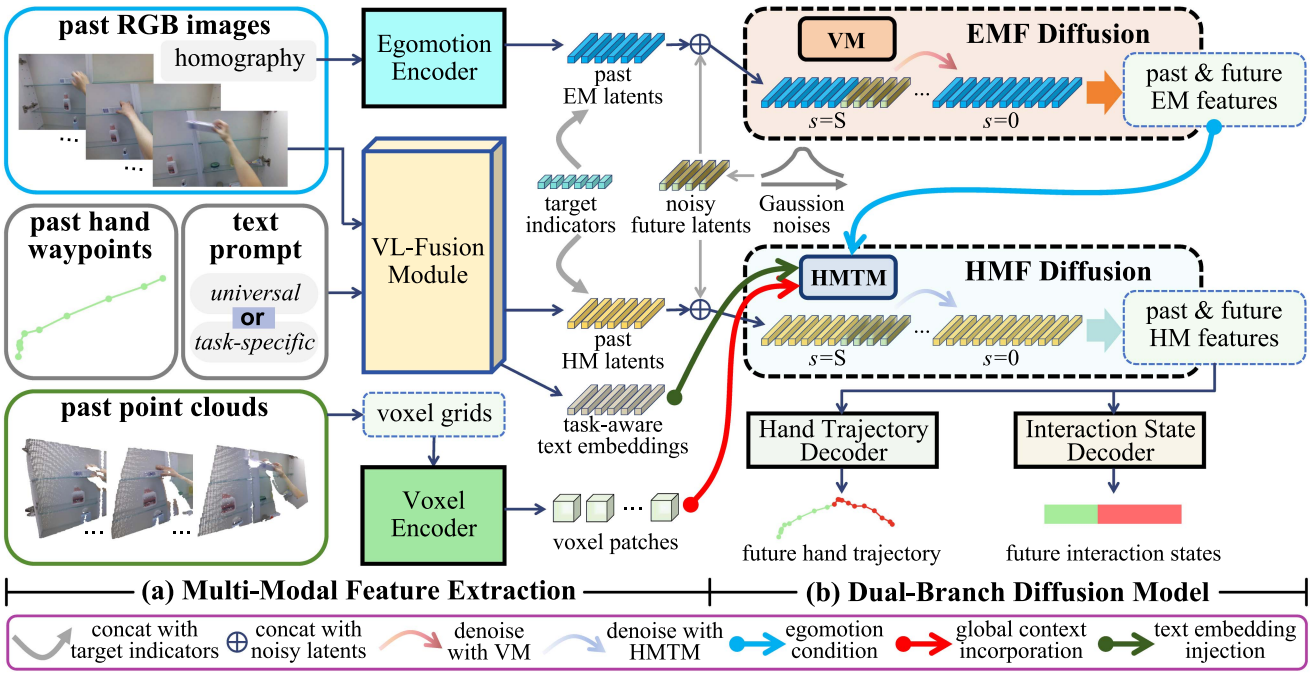


Fig. 2. System overview of Uni-Hand. Uni-Hand (a) converts multi-modal input into latent feature spaces, and (b) decouples predictions of future egomotion latents (EM latents) and hand motion latents (HM latents) by a novel dual diffusion. The vanilla Mamba (VM) is used for denoising in the ego-motion-forecasting diffusion (EMF diffusion). We further design a new denoising model in hand-motion-forecasting diffusion (HMF diffusion) with a hybrid Mamba-Transformer module (HMTM). The predicted HM latents are ultimately decoded to future hand trajectories and interaction states.

and lack cross-embodiment generalization. In contrast, widely available human videos already contain rich demonstrations. They are robot-agnostic and offer transferable knowledge for manipulation. Therefore, directly capturing human hand motion data from egocentric or exocentric videos to guide robot policy generation is a promising path towards embodied intelligence. Compared to producing pretrained visual representations from large-scale human videos [45], [46], [47], [48], extracting explicit affordances related to human hands enables more straightforward human-robot policy transfer. For example, OKAMI [49] retargets the body motions and hand poses separately from a single video imitation to humanoid robot actions. Following OKAMI using only one video demonstration, YOTO [50] proposes a proliferating strategy for one-shot teaching, enabling more demonstrations for the bimanual diffusion policy. Instead of relying entirely on human video data, Motion Tracks [51] combines robot and human video demonstrations to train a motion prediction network. Compared to these imitation learning schemes, ManipTrans [20] models the transfer process with reinforcement learning. Wen et al. [52] design a two-stage framework for any-point trajectory modeling, which guides policy learning with predicted future point trajectories. More recently, Qiu et al. [1] unify the state-action space of humanoid embodiments and human hands for action prediction.

Inspired by these human-robot transfer strategies, in this work, we expand the existing HMF assessment paradigm [12], [13], [15] by evaluating Uni-Hand's capability to support robot manipulation tasks. That is, we directly use pretrained Uni-Hand to generate executable trajectories and gripper actions for a real robot. The success rates of multiple manipulation tasks highlight its potential for real-world downstream deployments. In addition, we also present whether Uni-Hand can benefit the other three critical downstream tasks in egocentric vision, i.e., action anticipation [53], [54], [55], [56], early action recognition [56], [57], [58], and action recognition [56], [59], [60], [61], [62]. The improvement of anticipation and recognition accuracy by introducing hand motion features shows Uni-Hand's applicability in possible wearable assistive technologies.

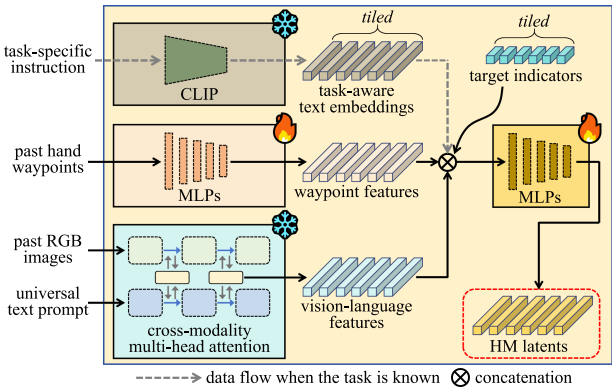


Fig. 3. Architecture of the VL-fusion module. It generates HM latents for the following HMF diffusion by fusing vision-language features, waypoint features, and task-aware text embeddings.

III. Proposed Method

A. System Overview

Here we first provide the overall inference pipeline of Uni-Hand in Fig. 2. Given a sequence of past egocentric observations $\mathcal{O} = \{O_t\}^{0}_{t=-N_p+1}$, past waypoints of the hand center or a specific joint $\mathcal{H}_p = \{H_t\}^{0}_{t=-N_p+1}$ ($H_t \in \mathbb{R}^2$ or $\mathbb{R}^3$), and a text

Authorized licensed use limited to: ShanghaiTech University. Downloaded on October 19,2026 at 15:54:43 UTC from IEEE Xplore. Restrictions apply.

prompt L, our proposed Uni-Hand aims to predict future hand trajectories $\mathcal{H}_{\mathrm{f}} = \{H_t\}_{t=1}^{N_{\mathrm{f}}} (H_t \in \mathbb{R}^2 \text{ or } \mathbb{R}^3)$, as well as interaction states $\mathcal{G}^{\mathrm{f}} = \{G_t\}_{t=1}^{N_{\mathrm{f}}} (G_t \in \{0, 1\})$ if needed. N_{p} and N_{f} denote the number of frames in the past and future time horizons. Each observation O_t consists of an egocentric 2D RGB image $I_t \in \mathbb{R}^{c \times h \times w}$ and 3D point clouds $D_t \in \mathbb{R}^{n \times 3}$. The temporal interaction state $G_t = 0$ indicates that the hand and target object are physically separated, while $G_t = 1$ represents that they are in contact. Following the previous works [13], [14], [15], we forecast the future hand waypoints on a fixed reference frame (e.g., the image plane and the camera coordinate system of the first frame $O_{-N_{\mathrm{p}}+1}$ for 2D and 3D HMF tasks, respectively). As can be noted, our framework absorbs multi-modal information to enhance the prediction model's perceptual capability of environments and tasks, compared to the current baselines [12], [13], [15], [18] only using image inputs. Besides, UniHand can also predict waypoints for multiple joints (e.g., $\mathcal{H}_{\mathrm{f}}^{j0}$, $\mathcal{H}_{\mathrm{f}}^{j4}$, and $\mathcal{H}_{\mathrm{f}}^{j8}$) in parallel since the data from different joints can be packed into a single batch.

Specifically, we first extract features/latents from the input RGB images, hand waypoints, and text prompt (Fig. 2(a), Section III-B). The VL-fusion module generates hand motion latents (HM latents) by fusing vision-language features, waypoint features, task-aware text embeddings, and target indicators. The egomotion encoder computes sequential homography matrices of the input image sequence as headset camera egomotion, and then encodes them to egomotion latents (EM latents). To integrate 3D structure features, the voxel encoder converts the past point clouds to voxel patches. Subsequently, a dual-branch diffusion model (Fig. 2(b), Section III-C), encompassing hand-motion-forecasting diffusion (HMF diffusion) and ego-motion-forecasting diffusion (EMF diffusion), is tailored for concurrently predicting future HM latents and EM latents conditioned on their past counterparts. The EMF diffusion uses the vanilla Mamba to efficiently denoise the noisy future latents. In contrast, the HMF diffusion adopts our hybrid Mamba-Transformer module as the denoising model, which consists of stacked egomotion-aware Mamba blocks and structure-aware/task-aware Transformer. It harmonizes hand motion modeling for egocentric vision, text embedding injection for task-aware optimization, and global context incorporation for 3D structure awareness. Ultimately, the denoised future HM latents are decoded to future hand trajectories and interaction states. We train the components of our Uni-Hand in an end-to-end manner with multiple loss functions (Section III-D).

B. Multi-Modal Feature Extraction

1) Egomotion Encoder: Since our framework considers modeling head motion in hand motion forecasting, we first calculate sequential homography $\mathcal{M}_{\mathrm{p}} = \{M_t\}_{t=-N_{\mathrm{p}}+1}^{0} (M_t \in \mathbb{R}^{3\times3})$ from I_t to represent past headset camera egomotion following [15], [18]. M_t denotes the homography matrix between t th frame and the first frame $(t = -N_{\mathrm{p}} + 1)$ estimated by SIFT descriptors [63] with RANSAC [64]. Subsequently, the egomotion encoder, composed of MLPs, encodes egomotion homography $\mathcal{M}$ to past egomotion features $F_{\mathrm{p}}^{\mathrm{em}} \in \mathbb{R}^{N_{\mathrm{p}} \times f}$ as latents for the following EMF diffusion. Note that the ground-truth (GT) future EM latents $F_{\mathrm{f}}^{\mathrm{em}} \in \mathbb{R}^{N_{\mathrm{f}} \times f}$ are obtained by encoding the GT homography matrices $\mathcal{M}_{\mathrm{f}} = \{M_t\}_{t=1}^{N_{\mathrm{f}}}$ computed from future image frames using the same egomotion encoder as for the past frames. They only exist in the training process for supervision and will be replaced by sampled noises $F_{\mathrm{noise}}^{\mathrm{em}} \in \mathbb{R}^{N_{\mathrm{f}} \times f}$ to enable denoising-based inference in EMF diffusion as shown in Fig. 2(b).

2) VL-Fusion Module: For hand motion forecasting, we develop the VL-fusion module to generate HM latents by fusing multiple modalities, including RGB images, past waypoints, and text prompts. As shown in Fig. 3, here we support two types of text prompting for different HMF applications, the universal prompt and the task-specific instruction. When the target hand motion is task-agnostic, we advocate using the universal prompt "*hand*" to generate vision-language fusion features $X^{\mathrm{vl}} \in \mathbb{R}^{(N_{\mathrm{p}}+\ell)\times x}$ by the pretrained GLIP [65] following [16]. x is the feature channel dimension. ℓ equals N_{f} during training, and is set to 0 during inference since future HM latents will be replaced by sampled noises in the inference process of our HMF diffusion. The universal text prompt helps to semi-implicitly capture hand poses and hand-scenario relationships within each image through GLIP's deepest cross-modality multi-head attention. Conversely, once the task being performed by the hand is known, an additional task-specific instruction (e.g., "*push the block onto the cloth*") is absorbed by the VL-fusion module to generate the text embedding $X^{\mathrm{task}} \in \mathbb{R}^{x}$ with CLIP [66] for task-awareness in hand motion forecasting.

Moreover, we use MLPs to encode past hand waypoints to waypoint features $X^{\mathrm{wp}} \in \mathbb{R}^{(N_{\mathrm{p}}+\ell)\times x}$, which are further concatenated with X^{vl} and tiled X^{task}. As showcased in Fig. 3, the additional target indicators are further concatenated with them to designate the prediction target. The target indicators consist of the tiled one-hot embeddings $X^{\mathrm{tar}} \in \mathbb{R}^{e}$, where e denotes the total number of prediction targets. Ultimately, the feature combos are fused by MLPs to produce the HM latents $F_{\mathrm{p}}^{\mathrm{hm}} \in \mathbb{R}^{N_{\mathrm{p}} \times f}$ and $F_{\mathrm{f}}^{\mathrm{hm}} \in \mathbb{R}^{N_{\mathrm{f}} \times f}$ for the following HMF diffusion. f represents the channel dimension of HM latents. $F_{\mathrm{f}}^{\mathrm{hm}}$ only exists in the training stage for reconstruction supervision since noisy future latents $F_{\mathrm{noise}}^{\mathrm{hm}} \in \mathbb{R}^{N_{\mathrm{f}} \times f}$ is concatenated to $F_{\mathrm{p}}^{\mathrm{hm}}$ in the inference stage of the HMF diffusion. By seamlessly integrating the target indicators into the framework, Uni-Hand identifies HM latents of different hand centers and joints, and achieves multi-target forecasting with the following HMF diffusion.

3) Voxel Encoder: Point Cloud Preprocessing: Uni-Hand advances 3D structure-awareness by absorbing point cloud observations for more reasonable HMF compared to the existing methods which only receive RGB images. The headset RGB-D camera captures a depth image aligned with its RGB counterpart I_t, which is directly transformed to the input point cloud $D_t \in O_t$ for Uni-Hand. To address the practical constraints of computational efficiency and memory requirements, we convert the holistic point cloud into a more compact voxelized form, as shown in Fig. 4. Specifically, we first exploit MobileSAM [67] to remove depth points projected to arms for each D_t. This is because arm movements generate artifact points hindering global 3D representation accuracy during the following multi-frame aggregation. Subsequently, all the purified point clouds within the past egocentric observations $\mathcal{O}$ are transformed into a unified global coordinate system by visual odometry. Then we aggregate them into voxel grids to avoid disturbance of the unordered point cloud structure, and further reduce the data volume.

Voxel Encoding: After the above point cloud preprocessing, the voxel encoder in Uni-Hand encodes the dense 3D voxels into a sparser representation $X^{\mathrm{vox}} \in \mathbb{R}^{N_{\mathrm{vox}} \times f}$ through 3D convolutions. That is, each input point cloud corresponds to N_{vox}

Authorized licensed use limited to: ShanghaiTech University. Downloaded on October 19,2026 at 15:54:43 UTC from IEEE Xplore. Restrictions apply.

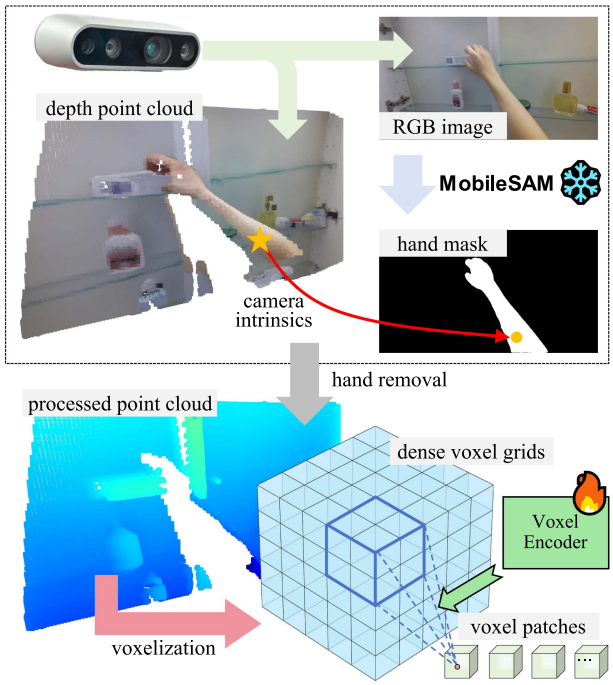


Fig. 4. Hand removal for purified point clouds. We regard the voxel patches encoded by the voxel encoder as 3D global context for the denoising process in the HMF diffusion.

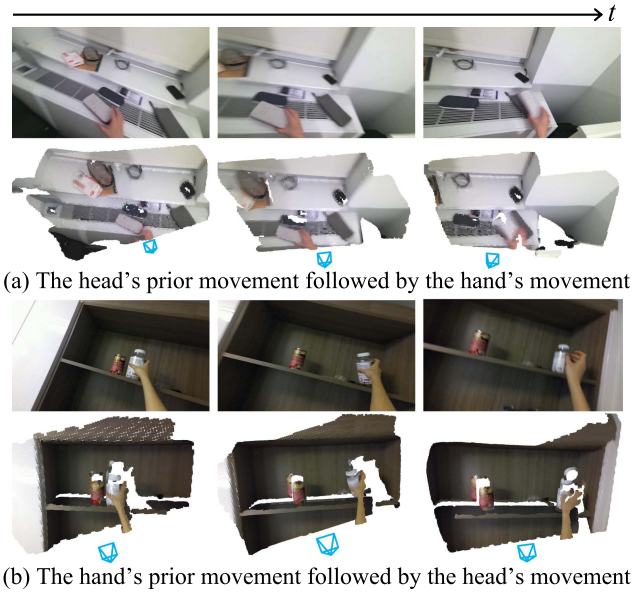


Fig. 5. Examples of head movement (corresponding to camera egomotion) and hand movement entangled during the hand-object interaction process in egocentric views in the EgoPAT3D dataset [68]. Here we present the RGB images and point clouds, as well as camera poses to clarify the hand-head motion trend.

voxel patches with the same channel dimension f as HM latents. Notably, while these voxel patches share dimensional alignment with HM latents, we do not integrate them into HM latents. Instead, they are maintained as separate global contexts of 3D interaction environments for the denoising model in the following HMF diffusion. This stems from the fact that the unified global representation of aggregated point clouds has not been time-varying compared to sequential egocentric observations. Merging it into time-varying HM latents is not required.

C. Dual-Branch Diffusion Model

1) Hand-Head Motion Coordination: Our proposed dual-brach diffusion model is aimed at capturing the synergy between hand movements and headset camera egomotion, and concurrently predicting future HM latents and EM latents. The hand-head motion coordination within the future interaction process is reflected in three aspects: 1) Hand movements often follow head motion, as the head's prior motion provides key visual cues for trajectory planning (Fig. 5(a)), 2) head movements may follow hand actions, as subconscious and faster hand motions can lead to head responses (Fig. 5(b)), and 3) humans keep hand movements within egocentric views to ensure accurate target contact Fig. 5(a) and (b). We argue that predicting hand motion agnostic to future head motion does not align with real human behavior planning. Instead, explicitly decoupling the entangled hand-head movements helps prediction models to better understand synergy motion patterns and potential intentions of interaction. In light of this, we propose the dual-branch diffusion model, encompassing the EMF diffusion and HMF diffusion to predict future EM latents (headset camera egomotion) and HM latents (hand motion) concurrently.

2) EMF Diffusion: The EMF diffusion is developed to denoise future EM latents, thus obtaining future head motion trend for the following hand motion forecasting. As shown in Fig. 2(b), the EMF diffusion leverages Mamba [23] as the denoising model to convert the noisy future EM latents $F_{\text{noise}}^{\text{em}}$ to $\hat{F}_{\text{f}}^{\text{em}} \in \mathbb{R}^{N_{\text{f}} \times f}$ conditioned on the past EM latents F_{p}^{em}. The predicted $\hat{F}_{\text{f}}^{\text{em}}$ implicitly represents future head motion trend during the interaction process. Since head motion exhibits simpler dynamics than hand motion, and head motion forecasting must be sufficiently efficient to provide conditioning signals for hand motion forecasting, we adopt the vanilla Mamba blocks for efficient denoising in our proposed EMF diffusion. Besides, we omit the process of decoding the future EM latents into specific homography matrices. This eliminates the ambiguity inherent in selecting supervision signals for alternative homography parameterizations [69].

3) HMF Diffusion: With the predicted camera egomotion trend, the HMF diffusion aims to forecast future HM latents, which will be transformed into the explicit future hand trajectories and interaction states by the devised decoders. As shown in Fig. 2(b), the HMF diffusion receives past HM latents F_{p}^{hm} to predict future counterparts $\hat{F}_{\text{f}}^{\text{hm}}$, conditioned on F_{p}^{em} and $\hat{F}_{\text{f}}^{\text{em}}$ predicted by the EMF diffusion. Here we propose a novel hybrid Mamba-Transformer module (HMTM) as the denoising model. As showcased in Fig. 6(a), HMTM comprises three key components: 1) stacked egomotion-aware Mamba (EAM) blocks adapted from MADiff [16], which employ Motion-Driven Selective Scan (MDSS) to seamlessly incorporate egomotion homography features into Mamba's state transitions, 2) our newly designed structure-aware Transformer (SAT), which absorbs voxel patches as 3D global context by multi-head corss attention, and 3) task-aware Transformer (TAT), where we inject text embeddings of task instructions into the HMF denoising process.

Authorized licensed use limited to: ShanghaiTech University. Downloaded on October 19,2026 at 15:54:43 UTC from IEEE Xplore. Restrictions apply.

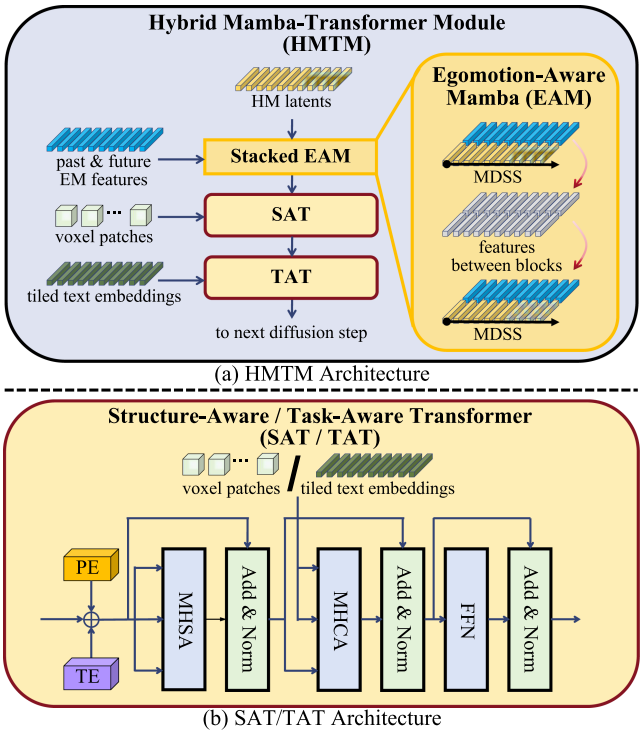


Fig. 6. Architecture of the hybrid Mamba-Transformer module. It denoises HM latents with stacked EAM, SAT, and TAT blocks.

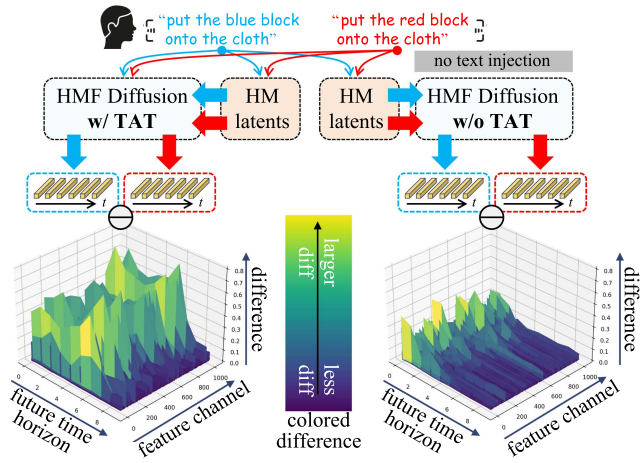


Fig. 7. Difference matrices of denoised future HM latents between two text instructions, "put the blue block onto the square cloth" and "put the red block onto the square cloth". They are visualized as height fields. This illustration reflects the effectiveness of our devised TAT to enhance the awareness of specific tasks in hand feature prediction.

Stacked Egomotion-Aware Mamba (EAM): Concretely, we first concatenate $F_{\mathrm{p}}^{\mathrm{em}}$ with the predicted $\hat{F}_{\mathrm{f}}^{\mathrm{em}}$ to $\hat{F}_{\mathrm{pf}}^{\mathrm{em}} \in \mathbb{R}^{(N_{\mathrm{p}}+N_{\mathrm{f}})\times f}$ along the time dimension, leading to the holistic feature sequence of headset camera egomotion in past and future time horizons. Similarly, we concatenate $F_{\mathrm{p}}^{\mathrm{hm}}$ with the sampled noise $F_{\mathrm{noise}}^{\mathrm{hm}}$ to $F_{\mathrm{pf}}^{\mathrm{hm}} \in \mathbb{R}^{(N_{\mathrm{p}}+N_{\mathrm{f}})\times f}$. Then we implement MDSS in EAM for each denoising step of the HMF diffusion, to denoise the future part of $F_{\mathrm{pf}}^{\mathrm{hm}}$ conditioned on the holistic sequential EM latents $\hat{F}_{\mathrm{pf}}^{\mathrm{em}}$. Therefore, we integrate egomotion-aware Mamba's powerful ability of temporal feature modeling, as well as head motion awareness into Uni-Hand. It can narrow the visual gap between the current camera view and the prediction view, and adaptively retain critical causality across temporal sequences, as demonstrated in the previous work [16].

Structure-Aware Transformer (SAT): The output of the stacked EMA blocks is further processed by the structure-aware Transformer in HMTM. SAT explicitly absorbs 3D global contexts of interaction environments for each denoising step of the HMF diffusion. As illustrated in Fig. 6(b), after the standard positional/temporal encoding (PE/TE) [15], [42], we sequentially implement multi-head self-attention (MHSA) on HM latents $F_{\mathrm{pf}}^{\mathrm{hm}}$, and multi-head cross-attention (MHCA) between the output of MHSA and voxel patches X^{vox}. Specifically, for MHCA input, we use the HM latents from the preceding MHSA and Add&Norm as the query, and let the voxel patches X^{vox} be the key and value. Inspired by human perception of 3D spatial relationships with stereo vision, Uni-Hand incorporates voxel patches as environmental global context through MHCA in each denoising step. This enables more plausible hand motion forecasting thanks to the potential layout understanding and collision avoidance from 3D structure awareness.

Task-Aware Transformer (TAT): Following the devised SAT, the task-aware Transformer injects specific task information into hand motion forecasting. It facilitates the model's awareness of specific human intentions or downstream task instructions. As shown in Fig. 6(b), TAT has the same architecture as SAT. However, it performs multi-head cross-attention between the output of its MHSA and the tiled text embeddings X^{task} of the task instruction. We found that only merging the task-aware text embeddings into HM latents (introduced in Section III-B2) is not enough to facilitate generating future hand waypoints that comply with the input task instruction. The reason could be that the task specificity inherent in HM latents is continuously weakened as the network depth increases and progressive denoising operations are performed in the HMF diffusion. We thus need to continuously inject original task information by MHCA to enhance the model's task discrimination capability. Fig. 7 illustrates the different matrices of denoised future HM latents between two text instructions, "put the blue block onto the square cloth" and "put the red block onto the square cloth". They are computed by $|\hat{F}_{\mathrm{f}}^{\mathrm{hm,blue}} - \hat{F}_{\mathrm{f}}^{\mathrm{hm,red}}|$ to present the feature difference from the two generally similar text instructions with different key targets (blue block vs. red block). As can be seen, after introducing the task-aware Transformer, besides only fusing text embeddings into HM latents, the feature discrepancies between the predicted future HM latents are significantly pronounced. That is, TAT helps the prediction model to distinguish different interaction tasks for more reasonable hand motion forecasting. Its positive effect becomes more obvious in later timings, which will directly decide how well the predicted hand motion interacts with the target object. We will further ablate the task-aware text embedding injection in Section IV-F. Note that we deploy TAT in Uni-Hand only when the task being performed by human hands is known to us. The output latents of SAT will be directly fed to the next diffusion step in task-agnostic HMF.

After the last denoising diffusion step in the HMF diffusion, we obtain the denoised HM latents $\hat{F}_{\mathrm{pf}}^{\mathrm{hm}}$. As shown in Fig. 2(b), the future part $\hat{F}_{\mathrm{f}}^{\mathrm{hm}} \in \mathbb{R}^{N_{\mathrm{f}}\times f}$ of $\hat{F}_{\mathrm{pf}}^{\mathrm{hm}}$ is ultimately decoded by

Authorized licensed use limited to: ShanghaiTech University. Downloaded on October 19,2026 at 15:54:43 UTC from IEEE Xplore. Restrictions apply.

the MLP-based hand trajectory decoder and interaction state decoder to future 2D/3D hand waypoints $\hat{\mathcal{H}}_{\mathrm{f}} = \{\hat{H}_t\}_{t=1}^{N_{\mathrm{f}}}$ and future interaction states $\hat{\mathcal{G}}_{\mathrm{f}} = \{\hat{G}_t\}_{t=1}^{N_{\mathrm{f}}}$. The denoised HM features, predicted waypoints and interaction states, can also be regarded as affordances for multiple downstream tasks, which will be validated in Section IV-E.

D. Training and Inference

1) Partial Noising and Denoising: We exploit partial noising/denoising proposed by Gong et al. [42] for the training and inference stages of both EMF diffusion and HMF diffusion. Specifically, we anchor the past latents $F_{\mathrm{p}}^{\mathrm{em}}$ and $F_{\mathrm{p}}^{\mathrm{hm}}$ in forward and reverse steps. In practice, although the EM latents and HM latents from the past time horizons are modified by noise corruption and denoising operations, they are manually overridden with their initial values (obtained via Multi-Modal Feature Extraction) after each diffusion step.

2) Loss Functions: We train Uni-Hand end-to-end, using six losses: 1) $\mathcal{L}_{\mathrm{VLB}}^{\mathrm{em}}$ for recovering future EM latents in the EMF diffusion, 2) $\mathcal{L}_{\mathrm{VLB}}^{\mathrm{hm}}$ for recovering future HM latents in the HMF diffusion, 3) trajectory displacement loss $\mathcal{L}_{\mathrm{dis}}$, 4) trajectory angle loss $\mathcal{L}_{\mathrm{angle}}$, 5) regularization term $\mathcal{L}_{\mathrm{reg}}$, and 6) interaction state loss $\mathcal{L}_{\mathrm{state}}$. We detail these losses as follows.

Losses for EMF Diffusion: To generate reasonable motion trends of the headset camera by the EMF diffusion, we set $\mathcal{L}_{\mathrm{VLB}}^{\mathrm{em}}$ for recovering future EM latents following the diffusion-related loss used by [15]:

$$\mathcal{L}_{\mathrm{VLB}}^{\mathrm{em}} = \sum_{s=2}^{S} ||\mathbf{z}_0^{\mathrm{em}} - f_{\mathrm{EMF}}(\mathbf{z}_s^{\mathrm{em}}, s)||^2 + ||F_{\mathrm{f}}^{\mathrm{em}} - \hat{F}_{\mathrm{f}}^{\mathrm{em}}||^2, \quad (1)$$

where s is the index of each diffusion step, and S is the total diffusion steps. $\mathbf{z}_s^{\mathrm{em}}$ is the latent after the s-th denoising step. f_{EMF} denotes the function of the denoising model in the EMF diffusion. $F_{\mathrm{f}}^{\mathrm{em}}$ is the GT future EM latents from the observations in future time horizons, and $\hat{F}_{\mathrm{f}}^{\mathrm{em}}$ is the counterparts finally predicted by the EMF diffusion as mentioned before.

Losses for HMF Diffusion: The diffusion-related loss $\mathcal{L}_{\mathrm{VLB}}^{\mathrm{hm}}$ supervising the predicted future HM latents in the HMF diffusion is similarly defined as:

$$\mathcal{L}_{\mathrm{VLB}}^{\mathrm{hm}} = \sum_{s=2}^{S} ||\mathbf{z}_0^{\mathrm{hm}} - f_{\mathrm{HMF}}(\mathbf{z}_s^{\mathrm{hm}}, s)||^2 + ||F_{\mathrm{f}}^{\mathrm{hm}} - \hat{F}_{\mathrm{f}}^{\mathrm{hm}}||^2, \quad (2)$$

where $\mathbf{z}_s^{\mathrm{hm}}$ indicates the latent after the s-th denoising step. f_{HMF} represents the function of the denoising model in the HMF diffusion. $F_{\mathrm{f}}^{\mathrm{hm}}$ is the GT future HM latents encoded from the observations in future time horizons. $\hat{F}_{\mathrm{f}}^{\mathrm{hm}}$ is the counterparts predicted by the HMF diffusion.

In addition, we utilize the trajectory displacement loss $\mathcal{L}_{\mathrm{dis}}$ and the trajectory angle loss $\mathcal{L}_{\mathrm{angle}}$ to optimize the hand waypoints decoded by the hand trajectory decoder:

$$\mathcal{L}_{\mathrm{dis}} = \frac{1}{N_{\mathrm{f}}} \sum_{t=1}^{N_{\mathrm{f}}} D_{\mathrm{dis}}(H_t, \hat{H}_t), \quad (3)$$

$$\mathcal{L}_{\mathrm{angle}} = \frac{1}{N_{\mathrm{f}}} \sum_{t=0}^{N_{\mathrm{f}}-1} D_{\cos}(H_{t+1} - H_t, \hat{H}_{t+1} - \hat{H}_t), \quad (4)$$

where $D_{\mathrm{dis}}(\cdot)$ denotes the Euclidean distance between predicted hand waypoints and GT ones. $\hat{H}_t$ is the output hand waypoints of the hand trajectory decoder, while H_t is its GT counterpart. $D_{\cos}(\cdot)$ represents the cosine similarity between two input motion vectors.

Following the previous works [15], [16], we also introduce the regularization term $\mathcal{L}_{\mathrm{reg}}$ to penalize the difference between the pseudo trajectory predictions $\tilde{H}_t$ and H_t:

$$\mathcal{L}_{\mathrm{reg}} = \frac{1}{N_{\mathrm{f}}} \sum_{t=1}^{N_{\mathrm{f}}} D_{\mathrm{dis}}(H_t, \tilde{H}_t), \quad (5)$$

where $\tilde{H}_t$ is obtained by directly feeding $F_{\mathrm{f}}^{\mathrm{hm}}$ to the hand trajectory decoder rather than the predicted $\hat{F}_{\mathrm{f}}^{\mathrm{hm}}$. This regularization helps to improves training stability and optimization performance.

To optimize interaction state forecasting, we utilize the interaction state loss $\mathcal{L}_{\mathrm{int}}$ defined as:

$$\mathcal{L}_{\mathrm{int}} = \mathrm{CE}\left(\mathcal{G}_{\mathrm{f}}, \hat{\mathcal{G}}_{\mathrm{f}}\right), \quad (6)$$

where $\mathrm{CE}(\cdot)$ is the binary cross-entropy loss function. $\hat{\mathcal{G}}_{\mathrm{f}}$ is the predicted interaction states decoded by the interaction state decoder, and $\mathcal{G}_{\mathrm{f}}$ is its GT label. Through the supervision of the interaction state loss, the timings of hand-object contact will be predicted with probabilities closer to 1.

Total Loss Function: The total loss function to train our proposed Uni-Hand is the weighted sum of all the above losses, which is expressed as:

$$\begin{aligned}\mathcal{L}_{\mathrm{total}} = {} & \lambda_{\mathrm{VLB}}^{\mathrm{em}}\mathcal{L}_{\mathrm{VLB}}^{\mathrm{em}} + \lambda_{\mathrm{VLB}}^{\mathrm{hm}}\mathcal{L}_{\mathrm{VLB}}^{\mathrm{hm}} + \lambda_{\mathrm{dis}}\mathcal{L}_{\mathrm{dis}} + \lambda_{\mathrm{angle}}\mathcal{L}_{\mathrm{angle}} \\ & + \lambda_{\mathrm{reg}}\mathcal{L}_{\mathrm{reg}} + \lambda_{\mathrm{int}}\mathcal{L}_{\mathrm{int}}, \end{aligned} \quad (7)$$

where $\lambda_{\mathrm{VLB}}^{\mathrm{em}}$, $\lambda_{\mathrm{VLB}}^{\mathrm{hm}}$, λ_{dis}, λ_{angle}, λ_{reg}, and λ_{int} are weights to balance the trade-off between these losses.

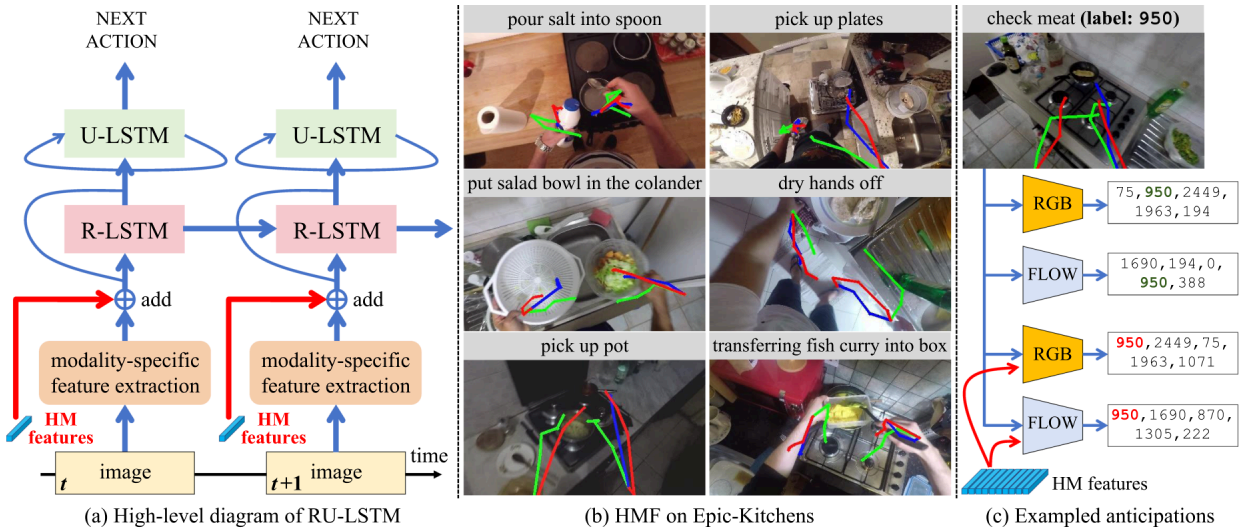
IV. Experimental Results

A. Benchmarks

We evaluate our proposed Uni-Hand and the existing baselines on five publicly available datasets including EgoPAT3D-DT [14], [68], H2O-PT [14], [70], HOT3D-Clips [71], Ego4D [72], and Epic-Kitchens-55 [73], and two self-recorded datasets including the Cup-Apple-Box-Hand (CABH) benchmark and Hand-ALOHA-Transfer (HAT) benchmark. In the HAT benchmark, we conduct human-robot transfer experiments on a real robot platform, ALOHA [74], to present how Uni-Hand effectively supports manipulation tasks in real-world setups.

1) Public Datasets: We follow the setups of the prior works [14], [21] to organize EgoPAT3D-DT, H2O-PT, and HOT3D-Clips datasets. Please refer to Supp. Mat., Section A for more details. For Ego4D, we attend to the 5-frame annotated clip with unconstrained head motions for each prediction in its future hand prediction benchmark, and let all the evaluated approaches receive the first three frames (three frames preceding the pre-condition frame) as input to predict the hand locations in the subsequent two frames (pre-condition frame and contact frame). To further demonstrate how our Uni-Hand supports the downstream action anticipation, early action recognition, and action recognition tasks, we organize the Epic-Kitchens-55 dataset following the splits of OCT [13]. The evaluation paradigm for the downstream action anticipation task is illustrated in Fig. 8. As shown in Fig. 8(a), we select the widely-used and off-the-shelf action anticipation framework, RU-LSTM [56], to present the feature enhancement from our HMF method. Concretely, we

Authorized licensed use limited to: ShanghaiTech University. Downloaded on October 19,2026 at 15:54:43 UTC from IEEE Xplore. Restrictions apply.



(a) High-level diagram of RU-LSTM (b) HMF on Epic-Kitchens (c) Exampled anticipations

Fig. 8. Illustration of testing Uni-Hand in the downstream action anticipation task. In the egocentric images of (b) and (c), the red, green, and blue lines denote hand center waypoints predicted by Uni-Hand, OCT [13], and GT labels, respectively.

first train Uni-Hand with annotated hand center waypoints in Epic-Kitchens-55, leading to a good HMF capability as depicted in Fig. 8(b). Afterwards, when implementing downstream action anticipation, we directly add the HM features (latents recovered by HMF diffusion of Uni-Hand) to the modality-specific features from the vanilla RU-LSTM. The combined features are further processed by the following rolling-unrolling operations through R-LSTM and U-LSTM. Then we optimize RU-LSTM for the action anticipation task, by minimizing the loss between the predicted action categories (linearly transformed from U-LSTM outputs) and the ground-truth labels. Note that we consider the modality-specific branches of RU-LSTM, i.e., the RGB branch and the flow branch, to demonstrate that our HMF method can enhance both RGB and flow features for higher anticipation accuracy. We do not use the object branch of RU-LSTM due to the mismatch in channel dimensions between HMF latents and object features. The feature enhancement operations for early action recognition and action recognition tasks are similar to action anticipation, while we only consider the ANTICIPATION period of RU-LSTM and discard its ENCODING period. Please refer more details about RU-LSTM structures to the original paper [56].

2) Cup-Apple-Box-Hand Benchmark: To further demonstrate that our method has the potential to scale up with low-cost devices for data collection, we used headset RealSense D435i to collect 1200 egocentric videos for three real-world tasks, i.e., *place the cup on the coaster* (Task 1), *put the apple on the plate* (Task 2), and *place the box on the shelf* (Task 3), to build the CABH benchmark (see Fig. 9). For each task, 350 video clips are used for training with the other 50 clips for evaluation. Each clip has around 5 seconds, with the first 50% regarded as the past sequences and the latter 50% as the future ones. To validate that Uni-Hand can achieve plausible hand motion forecasting using only egocentric vision without expensive MoCap systems for GT annotations, we utilize the visual grounding model, Grounded SAM [75], and the hand mesh recovery approach, HaMeR [27], to label 2D hand centers and joints on videos. Their 3D positions can be further recovered with depth observations of the RGB-D camera. To demonstrate that the choice of prediction canvas does not constrain the superiority of Uni-Hand, we directly predict future hand waypoints with each future frame as the canvas.

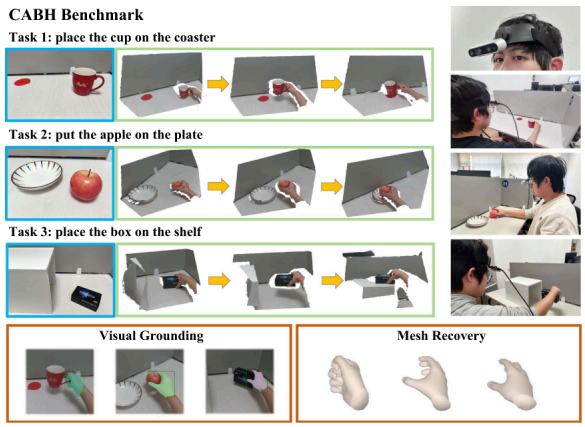


Fig. 9. Our self-collected CABH benchmark includes three hand-object interaction tasks to efficiently evaluate HMF.

3) Hand-ALOHA-Transfer Benchmark: Task Definition: We also propose the Hand-ALOHA-Transfer benchmark by recording hand motion with the top camera mounted on ALOHA [74], as shown in Fig. 10. The top camera, RealSense LiDAR Camera L515, captures egocentric RGB images and point clouds (depth) for hand motion forecasting and end-effector action prediction. In the HAT benchmark, we introduce **five** manipulation tasks:

- *Push Task (♥):* push the block onto the cloth. 400 hand motion videos in total for training Uni-Hand. 10 trials to evaluate real-world performance.
- *Pick-and-Place Task (♦):* put the block onto the cloth. 400 hand motion videos in total for training Uni-Hand. 10 trials to evaluate real-world performance.
- *Harder Pick-and-Place Task (♣):* stack the blue block onto the red block. 400 hand motion videos in total for training Uni-Hand. 10 trials to evaluate real-world performance.

Authorized licensed use limited to: ShanghaiTech University. Downloaded on October 19,2026 at 15:54:43 UTC from IEEE Xplore. Restrictions apply.

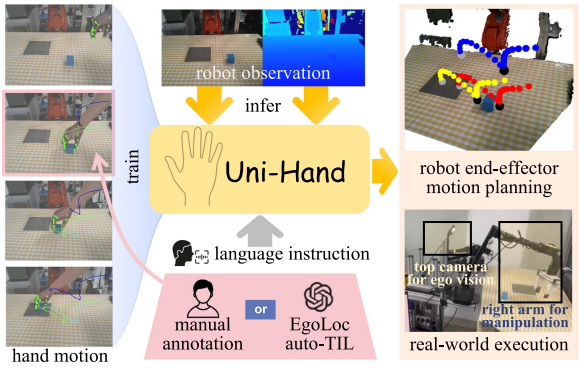


Fig. 10. Our scheme to deploy Uni-Hand to real-world robotic manipulation tasks.

- *Language-Conditioned Pick-and-Place Task (♠):* put the {*blue*, *red*} block onto the cloth. 800 hand motion videos for training, 400 videos per category (*blue*, *red*). 10 trials to evaluate real-world performance, 5 trials per category.
- *Long-Horizon Pick-and-Place Task (⋆):* put the block, the toy, and the banana onto the plate. 800 hand motion videos in total for training Uni-Hand. 10 trials to evaluate real-world performance.

These multiple manipulation tasks are used to evaluate both HMF performance and human-robot policy transfer ability.

Deployment Scheme: As illustrated in Fig. 10, we train Uni-Hand with recorded hand motion data, enabling it to generate complete executable actions for the robot end-effector in a look-then-move manner. Concretely, the predicted 3D trajectories are used for tracking, and the predicted interaction states help the robot to decide when to open and close its gripper. The GT 3D hand waypoints are captured by HaMeR [27]. To label the GT interaction states, we introduce two approaches: (1) manual annotation, and (2) automated temporal interaction localization (auto-TIL) using EgoLoc [76]. In both the training and inference stages with the HAT benchmark, only the first observation frame (i.e., $N_{\mathrm{p}} = 1$, agnostic to hands) is used as the egocentric visual inputs. The prediction model directly predicts the global motion with complete waypoints and interaction states, instead of step-by-step reasoning conditioned on iterative observations and states. This is because the robot manipulates autonomously without visual access to human hands, depending entirely on environmental observations. We should not encode observations involving human hands in the training stage. Our HAT benchmark thus evaluates the HMF methods' ability to generate the holistic trajectory and interaction states from scratch without the constraints of observed motion (See Supp. Mat., Sec. S for more details). Besides, for all the robotic manipulation tasks, we curate the mapping from human hand grasp poses to robot grasp poses, which is detailed in Supp. Mat., Sec. C.

B. Uni-Hand Configurations

For vision-language features from images and universal text prompt shown in Fig. 3, we extract the outputs of the deepest cross-modality multi-head attention module (X-MHA) in GLIP [65], which are further processed by the channel adapter proposed by [16]. Besides, in Fig. 3, the channel dimension of all the features before concatenation is set as 1024, and the counterpart of the target indicators is determined by the number of prediction targets (e.g., 5 if predicting five hand joints). To generate 3D global context shown in Fig. 4, we convert input point clouds into voxel grids of dimensions $20 \times 20 \times 20$ with a resolution of 0.05 m. Each voxel patch has a size of 27×1024. For both EMF and HMF diffusion, we set the channel dimension of the latent features to 1024. The total number of diffusion steps is set to 1000, while the EMF diffusion takes only one step to predict EM latents for high efficiency, and the HMF diffusion takes 100 steps to predict future HM latents. We adaptively adjust the input and output dimension of the MLPs for encoding and decoding hand trajectories according to the predictive space (2D/3D). Notably, for hand-ALOHA motion transfer, the camera homography is an identity matrix since the robot base is fixed in our HAT benchmark. Besides, we empirically found that the HMF diffusion can generate more stable trajectories for the robot end-effector when increasing the denoising steps to 1000. There is a trade-off between inference efficiency and planning accuracy. The hybrid pattern of Mamba and Transformer in our devised hybrid Mamba-Transformer module is designed according to the ablation study in Supp. Mat., Section E. We train Uni-Hand using AdamW optimizer [77] with a learning rate of 5e-5 for 1 K epochs on EgoPAT3D-DT, H2O-PT, Ego4D, and our self-recorded benchmarks, and with a learning rate of 5e-6 for 2 K epochs on the HOT3D-Clips dataset.

C. Evaluation on Multi-Dimensional Hand Trajectory Prediction

We first validate our claim that the proposed Uni-Hand predicts plausible hand center waypoints in both 2D and 3D dimensions. Predicting hand centers constitutes the most fundamental capability required by the HMF method. The evaluation metrics including averaged displacement error (ADE) and final displacement error (FDE) are reported following the previous works [14], [15], [24]. The prediction canvas [15] is selected as the first frame of each video clip. That is, all the 3D waypoints are expressed relative to the first camera coordinate, and all the 2D waypoints are transformed to the first image plane. The evaluation in the 3D space follows the absolute scale in meters. For the evaluation in 2D space, we normalize the predicted waypoint pixel coordinates by the image size. We separately forecast left- and right-hand motion. We compare our Uni-Hand with the existing baselines including Constant Velocity Hand (CVH) [15], OCT [13], USST [14], S-Mamba [78], Diff-IP2D [15], EER (based on Diff-IP2D) [17], EMAG [18], and MADiff [16]. We modify S-Mamba [78] originally designed for general time series forecasting into our diffusion-based paradigm to predict HM tokens. We additionally replace the 2D input and output waypoints, as well as the corresponding encoders and decoders with 3D counterparts in Diff-IP2D, EER, EMAG, and MADiff since they were originally developed for 2D forecasting tasks, obtaining the baselines Diff-IP3D, EER3D, EMAG3D, and MADiff3D.

Here we present the center prediction performance of our proposed Uni-Hand and the baselines on the publicly available datasets and our CABH benchmark. As Table I∼ Table IV show, Uni-Hand produces the lowest prediction errors on most metrics compared to the SOTA baselines. Although some baselines show competitive performance in 2D and 3D respectively, they show specialized and lack consistent superiority across both spaces due to a lack of multi-modal fusion. In contrast, our universal

Authorized licensed use limited to: ShanghaiTech University. Downloaded on October 19,2026 at 15:54:43 UTC from IEEE Xplore. Restrictions apply.

TABLE I
COMPARISON OF PERFORMANCE ON 3D HAND TRAJECTORY PREDICTION ON THE PUBLICLY AVAILABLE DATASETS. BEST AND SECONDARY RESULTS ARE VIEWED IN **BOLD BLACK** AND **GREEN** COLORS RESPECTIVELY.

Approach	EgoPAT3D-DT (seen)		EgoPAT3D-DT (unseen)		H2O-PT		HOT3D-Clips	
	ADE↓	FDE↓	ADE↓	FDE↓	ADE↓	FDE↓	ADE↓	FDE↓
CVH [15]	1.100	1.278	0.952	1.018	0.146	0.148	1.273	1.358
OCT* [13]	0.370	0.524	0.309	0.397	0.103	0.126	0.188	0.215
USST* [14]	**0.183**	**0.341**	**0.120**	**0.185**	**0.031**	**0.052**	0.123	0.157
S-Mamba [78]	0.185	0.355	0.138	0.207	0.038	0.074	**0.117**	**0.132**
Diff-IP3D [15]	0.199	0.377	0.156	0.229	0.049	0.081	0.147	0.164
EMAG3D [18]	0.207	0.365	0.166	0.209	0.062	0.085	0.141	0.158
EER3D [17]	0.195	0.377	0.158	0.233	0.047	0.075	0.150	0.162
MADiff3D [16]	**0.183**	0.363	0.139	0.224	0.032	0.059	0.120	0.147
Uni-Hand (ours)	**0.170**	**0.336**	**0.118**	**0.189**	**0.030**	**0.050**	**0.104**	**0.131**

* The baselines are re-evaluated according to the erratum: https://github.com/oppo-us-research/USST/commit/beebdb963a702b08de3a4cf8d1ac9924b544abc4.

TABLE II
COMPARISON OF PERFORMANCE ON 2D HAND TRAJECTORY PREDICTION ON THE PUBLICLY AVAILABLE DATASETS. BEST AND SECONDARY RESULTS ARE VIEWED IN **BOLD BLACK** AND **GREEN** COLORS RESPECTIVELY.

Approach	EgoPAT3D-DT (seen)		EgoPAT3D-DT (unseen)		H2O-PT		HOT3D-Clips	
	ADE↓	FDE↓	ADE↓	FDE↓	ADE↓	FDE↓	ADE↓	FDE↓
CVH [15]	0.180	0.230	0.188	0.221	0.206	0.208	0.437	0.444
OCT* [13]	0.108	0.122	0.091	0.147	0.387	0.381	0.190	0.240
USST* [14]	0.082	0.118	0.060	0.087	0.040	**0.068**	0.133	0.172
S-Mamba [78]	0.084	0.141	0.071	0.118	0.051	0.094	0.130	**0.160**
Diff-IP2D [15]	0.080	0.130	0.066	0.087	0.042	0.074	0.169	0.208
EMAG [18]	0.095	0.143	0.092	0.139	0.054	0.093	0.155	0.194
EER [17]	0.072	0.121	0.056	**0.080**	0.039	0.069	0.166	0.204
MADiff [16]	**0.065**	**0.105**	**0.054**	**0.086**	**0.039**	**0.068**	**0.124**	**0.160**
Uni-Hand (ours)	**0.064**	**0.109**	**0.052**	0.087	**0.036**	**0.065**	**0.118**	**0.156**

TABLE III
COMPARISON OF PERFORMANCE ON HAND TRAJECTORY PREDICTION ON THE EGO4D DATASET IN 2D SPACE. BEST AND SECONDARY RESULTS ARE VIEWED IN **BOLD BLACK** AND **GREEN** COLORS.

Approach	ADE↓	FDE↓
USST [14]	0.196	0.208
EMAG [18]	0.221	0.228
Diff-IP2D [15]	0.187	0.208
S-Mamba [78]	**0.182**	**0.201**
Uni-Hand (ours)	**0.176**	**0.193**

TABLE IV
COMPARISON OF PERFORMANCE ON HAND TRAJECTORY PREDICTION ON CABH IN 3D AND 2D SPACES WITH AUTOMATED ANNOTATIONS. BEST AND SECONDARY RESULTS ARE VIEWED IN **BOLD BLACK** AND **GREEN** COLORS.

Approach	Task 1		Task 2		Task 3	
	3D metrics					
	ADE↓	FDE↓	ADE↓	FDE↓	ADE↓	FDE↓
USST [14]	0.102	0.125	0.109	0.128	0.103	0.130
Diff-IP3D [15]	0.053	0.057	0.055	0.078	0.064	0.080
S-Mamba [78]	**0.045**	**0.055**	**0.050**	**0.072**	**0.058**	**0.061**
Uni-Hand (ours)	**0.041**	**0.052**	**0.044**	**0.061**	**0.047**	**0.053**
	2D metrics					
USST [14]	0.115	0.155	0.127	0.168	0.124	0.161
Diff-IP2D [15]	0.084	**0.100**	0.094	0.117	0.109	0.126
S-Mamba [78]	**0.080**	0.102	**0.081**	**0.107**	**0.102**	**0.108**
Uni-Hand (ours)	**0.066**	**0.086**	**0.077**	**0.099**	**0.094**	**0.095**

framework accommodates both 3D structure information and 2D camera egomotion, achieving consistently comparable performance in both 2D and 3D spaces. Besides, Uni-Hand consistently outperforms the baseline methods on Ego4D, which indicates that Uni-Hand can more effectively forecast semantically driven hand motion under limited observation (only 3-frame input) and highly dynamic head motion, rather than only extrapolating short-term kinematic hand patterns. Figs. 11 and 12 visualize the predicted hand center waypoints, where Uni-Hand's predictions have better stability and directionality in multiple dimensions. This suggests that Uni-Hand better captures hand motion patterns and potential human intentions during interaction by harmonizing multi-modal information. Moreover, HOT3D-Clips has small data volumes to supervise prediction models, where our Uni-Hand consistently outperforms the other baselines. This implies that our method better accommodates situations with limited training data.

D. Evaluation on Multi-Target Hand Trajectory Prediction

We further validate Uni-Hand's capability to forecast plausible future waypoints of multiple hand joints rather than centers. We do not involve EgoPAT3D-DT here due to a lack of joint annotations. For H2O-PT and HOT3D-Clips which have labeled GT 3D joint positions, we predict trajectories of five joints: one on the wrist and two each on the thumb and index finger (see Supp. Mat., Section B). These prediction targets are selected according to the recent work [51] considering potential downstream tasks. Moreover, we provide a data extension of our CABH benchmark (CABH-E) compared to our preliminary version [21]. As shown in the first row of Fig. 13, we record "place the apple into the square area from multiple directions" to demonstrate that Uni-Hand captures multimodal behavior [43] without generating jittery

Authorized licensed use limited to: ShanghaiTech University. Downloaded on October 19,2026 at 15:54:43 UTC from IEEE Xplore. Restrictions apply.

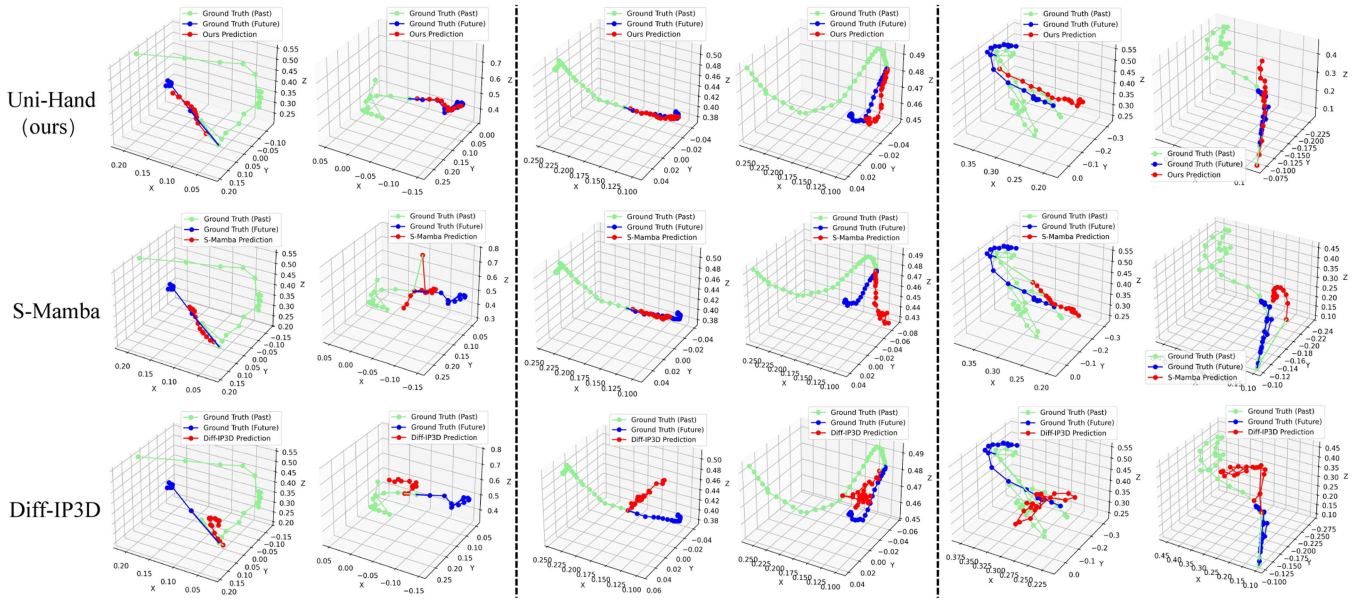


Fig. 11. Visualization of predicted hand trajectories in the 3D space (left: EgoPAT3D-DT; middle: H2O-PT; right: HOT3D-Clips). We show the holistic sequence including observed past hand waypoints (green), ground-truth future ones (blue), and predicted future counterparts (red) by our Uni-Hand and two SOTA baselines.

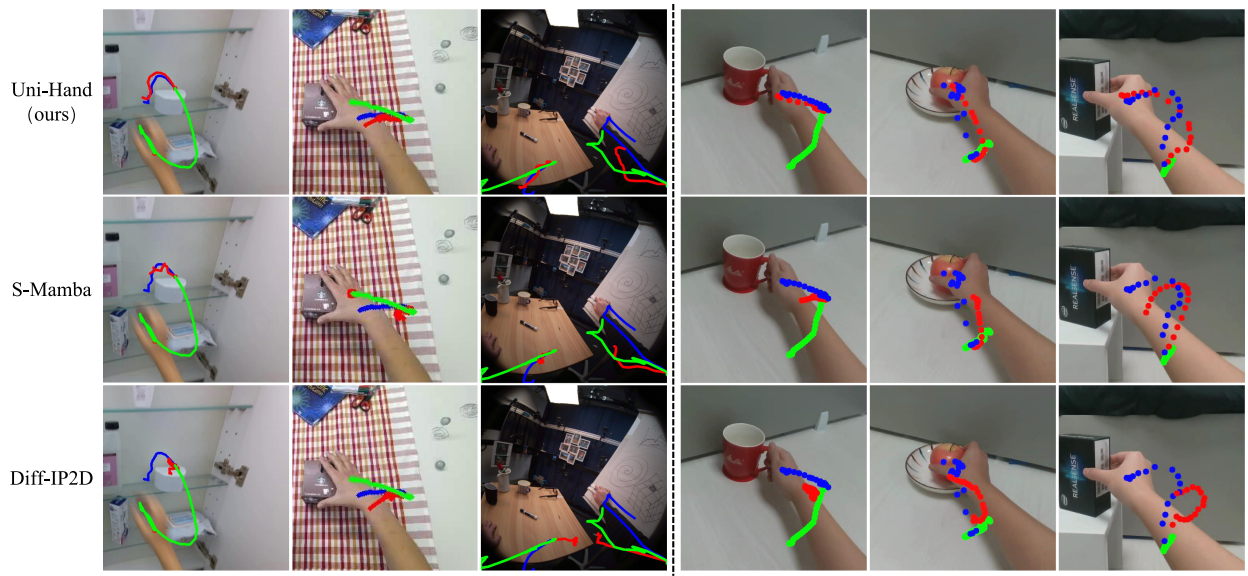


Fig. 12. Visualization of 2D hand center forecasting (left: EgoPAT3D-DT, H2O-PT, and HOT3D-Clips; right: CABH). We show the holistic sequence including observed past joint waypoints (green), ground-truth future ones (blue), and predicted future counterparts (red) by our Uni-Hand and two SOTA baselines. Here we showcase the predictions on the first frame of each input video clip for EgoPAT3D-DT, H2O-PT, and HOT3D-Clips, while presenting the last frame for our CABH.

actions from multiple directions. For CABH-E, we also generate 3D GT positions by HaMeR with depth observations, and forecast trajectories of three joints, which are on the wrist, the thumb tip, and the middle fingertip (see Fig. 13). This target selection follows the setups from CoM [79], which emphasizes using the middle fingertips instead of the index ones. We thus showcase the performance on forecasting joint waypoints with different target combinations. The final displacement errors are averaged among the target joints. Note that the SOTA baselines in this experiment are implemented separately on different joints since they do not support multi-target prediction originally.

Table V shows that Uni-hand consistently outperforms the baselines on the three datasets. Especially for CABH-E, Uni-hand reduces ADE and FDE of the second-best baseline by 0.055 and 0.064, respectively. We also visualize the predicted joints' positions projected to the last frame of each input video in Fig. 13. Uni-Hand generally generates plausible joint waypoints extending in correct directions, with the final positions all successfully meeting the requirement of placing the apple in the target area. These results demonstrate that our universal framework achieves good multi-target hand forecasting, extending the scope of the existing hand trajectory prediction that only attends to hand centers.

Authorized licensed use limited to: ShanghaiTech University. Downloaded on October 19,2026 at 15:54:43 UTC from IEEE Xplore. Restrictions apply.

TABLE V
COMPARISON OF PERFORMANCE ON MULTI-JOINT WAYPOINT PREDICTION IN 3D SPACE ON THE H2O-PT, HOT3D-CLIPS, AND CABH-E. BEST AND SECONDARY RESULTS ARE VIEWED IN **BOLD BLACK** AND **GREEN** COLORS RESPECTIVELY.

Approach	H2O-PT		HOT3D-Clips		CABH-E	
	ADE↓	FDE↓	ADE↓	FDE↓	ADE↓	FDE↓
USST* [14]	0.102	0.123	0.162	0.180	0.105	0.122
Diff-IP3D [15]	0.152	0.160	0.200	0.204	0.093	0.110
S-Mamba [78]	0.046	0.087	0.200	0.229	0.099	0.130
Uni-Hand w/o X^{tar}	**0.038**	**0.076**	**0.121**	**0.151**	**0.047**	**0.056**
Uni-Hand (ours)	**0.036**	**0.066**	**0.118**	**0.148**	**0.038**	**0.046**

TABLE VI
PERFORMANCE COMPARISON ON OUR PROPOSED HAT BENCHMARK. WE PRESENT THE ERRORS OF HAND TRAJECTORY PREDICTION AND INTERACTION STATE PREDICTION ON THE VALIDATION SET (20 OUT OF THE TRAINING SAMPLES), AS WELL AS THE SUCCESS RATES (SR) OF 10 TEST TRIALS FOR EACH MANIPULATION TASK.

Tasks	S-Mamba [78]				Uni-Hand (ours)			
	ADE↓	FDE↓	MAE↓	SR↑	ADE↓	FDE↓	MAE↓	SR↑
♥	0.038	0.049	-	40%	0.022	0.018	1.3	100%
♦	0.053	0.051	-	20%	0.035	0.012	1.9	90%
♣	0.061	0.040	-	0%	0.040	0.016	0.8	80%
♠	0.075	0.076	-	0%	0.054	0.028	1.8	80%
★	0.100	0.109	-	0%	0.059	0.071	1.9	20%

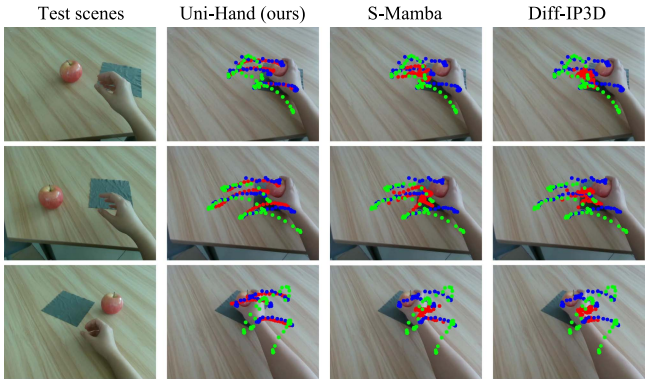


Fig. 13. Visualization of multi-target prediction on our CABH-E. We show the holistic sequence including observed past joint waypoints (green), ground-truth future ones (blue), and predicted future counterparts (red) by our Uni-Hand and two SOTA baselines. Here we showcase the predictions on the last frame of each input video clip.

Table V also ablates the contribution of our devised target indicators. The baseline Uni-Hand w/o X^{tar} separately forecasts different joints without using target indicators. The results indicate that although Uni-Hand w/o X^{tar} has outperformed the other SOTA baselines, Uni-Hand can better identify which joint it currently attends to, and predicts more accurate joint waypoints. Therefore, the target indicators are demonstrated as key components in our universal HMF framework to achieve multi-target forecasting. Please refer to Supp. Mat., Section I for more ablations on target indicator formulations and conditioning mechanisms.

E. Evaluation on Downstream Tasks

In this section, we evaluate the claim that Uni-Hand enables multi-task affordances for downstream applications.

1) Human-Robot Manipulation Policy Transfer: This experiment is conducted on our proposed Hand-ALOHA-Transfer benchmark to validate Uni-Hand's ability to facilitate multiple downstream robotic manipulation tasks. The five tasks and prediction configurations in this benchmark have been detailed in Section IV-A3. We activate the task-aware Transformer and inject the text embeddings of task instructions into it to enhance the model's adherence to specific task setups. Here, the GT interaction states are labeled manually by default. More experimental details and results with EgoLoc annotations will be introduced in Supp. Mat., Section K. Table VI first shows 3D ADE and FDE metrics of Uni-Hand and the SOTA baseline S-Mamba on the validation sets of the real-world scenes. Uni-Hand generates end-effector trajectories that align better with human motion patterns. We also observed lower prediction errors around hand-object contact timings (see Supp. Mat., Section P), which suggests that hand movements hold lower uncertainties at these timings, leading to more stable supervision signals to optimize Uni-Hand. This also facilitates plausible action planning despite relatively higher errors at other timings (e.g., the initial stage). Besides, we also present the Mean Absolute Error (MAE) of the contact-separation/separation-contact transition timings within the predicted interaction states on these tasks. Since the baseline [78] is proposed without the ability for interaction state prediction, we do not report its MAE and directly use GT labels as its pseudo predictions for the following real-world execution. Table VI shows that Uni-Hand predicts accurate hand-object contact/separation states for reasonable action planning. Then, we directly transfer the prediction models to the real-world deployment, and let the robot end-effector follow the generated trajectories and grasp actions.

As shown in Table VI, for 10 test trials of each designated manipulation task, Uni-Hand exhibits significant superiority with higher success rates than S-Mamba. The baseline is infeasible with 0% SR for the harder pick-and-place, language-conditioned pick-and-place, and long-horizon pick-and-place tasks. Fig. 14 shows how the end-effector accurately follows Uni-Hand's hand motion forecasting results for these tasks. This demonstrates that Uni-Hand provides the required motion knowledge for

Authorized licensed use limited to: ShanghaiTech University. Downloaded on October 19,2026 at 15:54:43 UTC from IEEE Xplore. Restrictions apply.

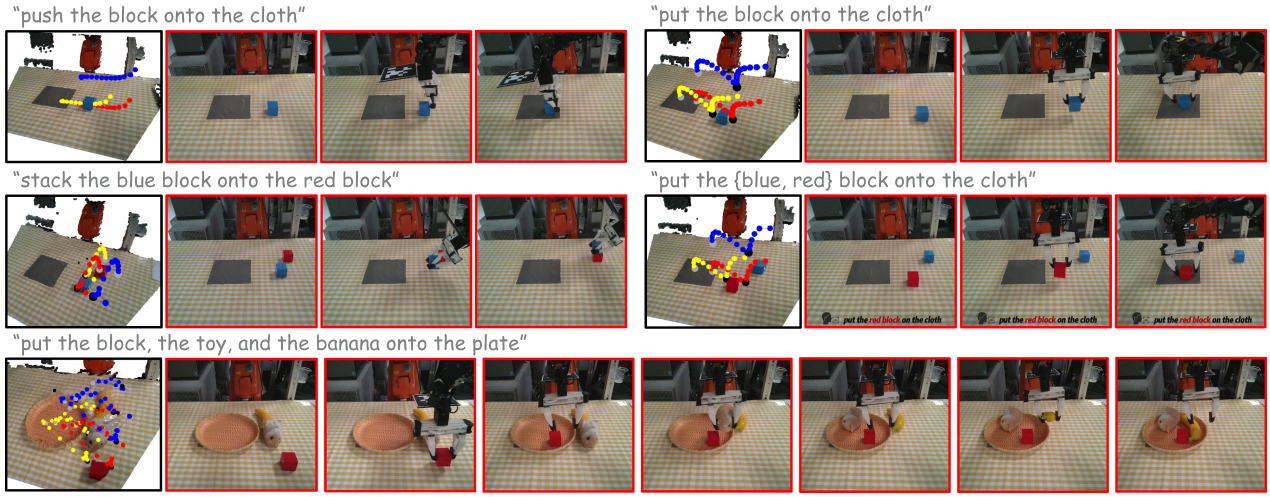


Fig. 14. Real-world test on ALOHA. For each manipulation task, we illustrate the trajectories generated from scratch by Uni-Hand and how the gripper follows the action planning. The blue dots denote the predicted hand wrist positions, which spawn the yellow and red dots corresponding to the index fingertip and the thumb tip. The black and grey dots denote the pick and place timings of the gripper, respectively.

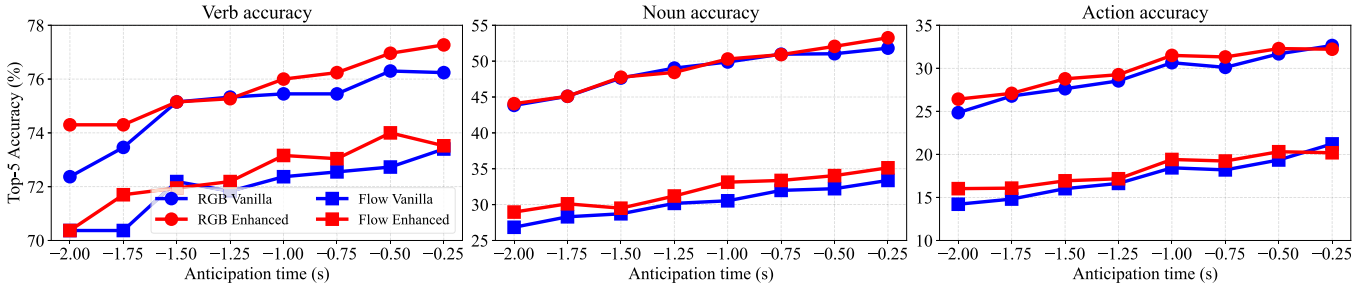


Fig. 15. Performance comparison between RU-LSTM branches enhanced by our Uni-Hand's HM features and the vanilla model on the downstream action anticipation task. We follow the evaluation metrics of the prior work [56].

practical deployments on a real robot platform. It benefits policy learning for atomic skills such as push, pick, and place, and also holds promise for completing complex language-conditioned and long-horizon robotic tasks.

2) Action Anticipation: We have showcased in Fig. 8(b) that our Uni-Hand can predict good hand waypoints after being optimized on Epic-Kitchens. This indicates that Uni-Hand can generate reasonable hand motion features (HM features) related to specific future human activities. Therefore, we explore whether the future HM features can enhance the downstream action anticipation, a task that predicts which category of action will happen. As shown in Fig. 15, after we add the HM features recovered by our HMF diffusion to the vanilla modality-specific features in the RGB and flow branches of RU-LSTM, the Top-5 verb/noun/action accuracy increases at most anticipation times τ_a. Fig. 8(c) also presents an example of Top-5 action categories predicted by RU-LSTM before and after adding HM features. As can be seen, the feature enhancement by our HMF method facilitates a higher probability for the correct action category. These results highlight that the hand motion features predicted by Uni-Hand can be effective affordances encompassing future activity information for the downstream action anticipation.

3) Early Action Recognition and Action Recognition: Different from action anticipation, the downstream early action recognition and action recognition tasks focus on reasoning human activity categories according to partial or complete observations. Following RU-LSTM [56], we simply discard its ENCODING period and only attend to the observations in the ANTICIPATION period. As noted in Table VII, after adding HM features from Uni-Hand, the recognition accuracy of RGB and flow branches in RU-LSTM increases at most observation rates.

TABLE VII
TOP-5 ACTION ACCURACY (%) OF RU-LSTM BRANCHES ON THE DOWNSTREAM EARLY ACTION RECOGNITION AND ACTION RECOGNITION TASKS. WE PRESENT THE PERFORMANCE AT THE OBSERVATION RATES 12.5% $\sim$ 87.5% FOR EARLY ACTION RECOGNITION, AND 100% FOR ACTION RECOGNITION FOLLOWING [56]. BEST RESULTS ARE VIEWED IN **BOLD BLACK**.

Approach	Early action recognition						
	12.5%	25.0%	37.5%	50.0%	62.5%	75.0%	87.5%
RGB (vanilla)	**36.40**	**38.21**	39.18	**40.45**	41.29	41.90	41.66
RGB (enhanced)	36.22	37.85	**39.42**	39.78	**41.60**	**42.26**	**42.32**
Flow (vanilla)	19.77	**24.67**	**27.63**	28.30	29.99	30.35	30.71
Flow (enhanced)	**19.83**	24.43	26.90	**28.54**	**30.29**	**31.02**	**31.74**

Approach	Action recognition (100%)	
	Centers	Multiple joints
RGB (vanilla)	41.78	41.78
RGB (enhanced)	**42.74**	**42.82**
Flow (vanilla)	30.96	30.96
Flow (enhanced)	**31.74**	**31.86**

Authorized licensed use limited to: ShanghaiTech University. Downloaded on October 19,2026 at 15:54:43 UTC from IEEE Xplore. Restrictions apply.

TABLE VIII
ABLATION STUDY ON CAMERA EGOMOTION. SE(3) AS EGOMOTION REPRESENTS THE BASELINE REPLACING THE INPUT CAMERA HOMOGRAPHY WITH 6-DOF POSES. UNI-HAND W/O EMF REPRESENTS THE BASELINE WITHOUT THE EMF DIFFUSION. BEST RESULTS ARE VIEWED IN **BOLD BLACK**.

Approach	EgoPAT3D-DT (seen)		EgoPAT3D-DT (unseen)		H2O-PT	
	ADE↓	FDE↓	ADE↓	FDE↓	ADE↓	FDE↓
SE(3)	0.257	0.446	0.217	0.308	0.032	0.063
w/o EMF	0.186	0.363	0.137	0.231	0.031	0.053
Uni-Hand	**0.170**	**0.336**	**0.118**	**0.189**	**0.030**	**0.050**

Note that we additionally train a Uni-Hand model to achieve multi-joint motion forecasting (the wrist, the thumb tip, and the index fingertip), and enhance RU-LSTM by adding these new HM features. For action recognition with the observation rate of 100%, HM features encompassing multi-joint motion information yield additional performance improvements. This further validates the advantages of multi-target predictions by Uni-Hand. However, compared to the results of action anticipation in Fig. 15, the feature enhancement from Uni-Hand exhibits an overall weaker influence on these two downstream tasks. This suggests that the positive effect of predicted HM features diminishes when partial action observations are available. Nevertheless, Uni-Hand still affords informative HM features for action recognition thanks to a deeper understanding of hand motion patterns.

Note that Uni-Hand is not intended as a standalone action anticipation/recognition model, but as an effective HMF framework whose predicted motion features can be directly plugged into existing anticipation/recognition models. Given that RU-LSTM has already been extensively optimized in the original work [56], consistent performance gains indicate that Uni-Hand provides meaningful future hand affordances, while larger improvements would require redesigning the downstream anticipation architecture.

F. Ablation Studies

Camera egomotion forecasting: We first ablate the headset camera egomotion forecasting on EgoPAT3D-DT and H2O-PT. Concretely, we conduct a baseline by removing the EM Diffusion, and regarding the last past camera homography as the constant egomotion in future time horizons. We present the hand center forecasting performance in 3D space in Table VIII. As can be noted, forecasting future headset camera egomotion improves hand prediction accuracy. This demonstrates that our proposed dual-branch diffusion model in Uni-Hand captures the synergy between head motion and hand motion for a better understanding of the future interaction process. Besides, Fig. 16 visualizes the failure cases of the variant without camera egomotion forecasting in 2D space. As can be seen, explicitly forecasting future camera egomotion provides strong spatiotemporal constraints on hand trajectory prediction, leading to improved directional consistency and temporal stability.

Camera egomotion representation: Despite the ablation studies on homography as camera egomotion provided by the previous works [15], we additionally present 3D hand center forecasting performance when we regard SE(3), i.e., 6-DOF poses from visual odometry, as camera egomotion instead of homography. As shown in Table VIII, replacing the vanilla egomotion homography with SE(3) leads to a significant degradation in HMF performance. This suggests that the homography-based representation is better suited for modeling egomotion in this context. It also aligns with the fact that observed hands are constrained to 2D image planes, and homography matrices more effectively capture camera egomotion variations coupled with hand motions.

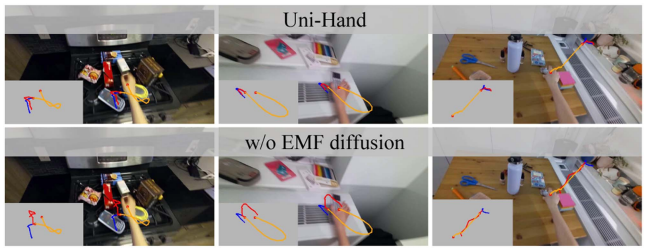


Fig. 16. Failure cases of the Uni-Hand's variant without future egomotion forecasting on EgoPAT3D-DT. The observed past hand waypoints, predicted future hand waypoints, and GT labels are represented in green, red, and blue, respectively, with the first frame of each sequence as the visualization canvas.

TABLE IX
ABLATION STUDY ON MULTIMODAL INPUTS. BEST RESULTS ARE VIEWED IN **BOLD BLACK**.

Input modalities				Seen		Unseen	
waypoint	image	text	point cloud	ADE↓	FDE↓	ADE↓	FDE↓
✓				0.178	0.356	0.124	0.205
✓	✓			0.173	0.350	0.122	0.201
✓	✓	✓		0.171	0.347	0.122	0.200
✓	✓	✓	✓	**0.170**	**0.336**	**0.118**	**0.189**

TABLE X
ABLATION STUDY ON TEXT EMBEDDING INJECTION. BEST RESULTS ARE VIEWED IN **BOLD BLACK**.

Approach	ADE↓	FDE↓	MAE↓	SR↑
Uni-Hand w/o TAT	0.057	0.034	2.0	40%
Uni-Hand	**0.054**	**0.028**	**1.8**	**80%**

Multi-modal inputs: To ablate multiple modalities, we incrementally add past hand waypoints, RGB images, the universal text prompt, and point clouds in model inputs. The experimental results of 3D hand center forecasting on EgoPAT3D-DT shown in Table IX indicate that our proposed framework well harmonizes multi-modal information to achieve strong forecasting performance. The incorporation of both the universal text prompt and point cloud data effectively bridges modality gaps in existing literature.

Text embedding injection: We have shown in Fig. 7 that our text embedding injection in hybrid Mamba-Transformer facilitates the HM features denoised for task-aware HMF. Here we further present the quantitative results on the validation and test sets of the task "put the {*blue*, *red*} block onto the cloth" of our HAT benchmark. We conduct the baseline, Uni-Hand w/o TAT, by removing the task-aware Transformer in the denoising model. Table X indicates that our text embedding injection with TAT reduces prediction errors, while also improving success rates significantly.

Authorized licensed use limited to: ShanghaiTech University. Downloaded on October 19,2026 at 15:54:43 UTC from IEEE Xplore. Restrictions apply.

V. Discussion

Uni-Hand introduces a new algorithmic paradigm for hand motion forecasting that extends the existing approaches in multiple aspects, including modeling structure, prediction targets, conditioning objectives, and output space. At the modeling level, Uni-Hand is the first framework to jointly forecast future head egomotion and hand motion via a coordinated dual-branch diffusion architecture, enabling hand motion prediction to be explicitly conditioned on future egomotion. Beyond purely kinematic forecasting, Uni-Hand expands the output space by jointly predicting hand trajectories and hand-object contact/separation states, reasoning not only about where the hand will move but also when interactions will occur. Moreover, Uni-Hand introduces devised target indicators that enable a single predictive model to forecast heterogeneous hand targets, including the hand center, wrist, and finger joints, within a unified denoising process. The framework further incorporates global 3D scene context through a voxel-grid-based formulation, leading to 3D structure awareness under real-world geometric constraints. Ultimately, Uni-Hand injects task-aware textual embeddings directly into the diffusion denoising process through the dedicated cross-attention mechanism, aligning future hand motion generation with high-level task intentions. As can be noted, these algorithmic contributions redefine hand motion forecasting as a multi-target, interaction-aware, and task-conditioned generative problem, which we hope will inspire future research on egocentric human-object interaction.

VI. Conclusion

We have developed a universal framework namely Uni-Hand to achieve multi-dimensional and multi-target hand motion forecasting in egocentric views. Uni-Hand absorbs multi-modal data and predicts plausible hand trajectories in both 2D and 3D spaces, considering hand-head motion synergy. It attends to the future motion of hand centers and any joints, as well as hand-object interaction states. Extensive experiments on public datasets and our newly proposed benchmarks have demonstrated that our universal framework outperforms the existing SOTA baselines in hand motion forecasting across 2D/3D spaces and multiple targets, and facilitates multi-task affordances for downstream applications, including robotic manipulation, action anticipation, early action recognition, and action recognition. In the future work, we plan to introduce tactile sensing for contact-rich human-robot policy transfer [80].

References

[1] R.-Z. Qiu et al., "Humanoid policy human policy," 2025, *arXiv:2503.13441*.

[2] S. Kareer et al., "Egomimic: Scaling imitation learning via egocentric video," 2024, *arXiv:2410.24221*.

[3] D. Gong, L. Hao, J. Yu, and G. Zuo, "Bionic design of a dexterous anthropomorphic hand actuated by antagonistic PAMs," in *Proc. IEEE Int. Conf. Real-time Comput. Robot.*, pp. 493–498, 2020.

[4] A. Bandini and J. Zariffa, "Analysis of the hands in egocentric vision: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 6, pp. 6846–6866, Jun. 2023.

[5] M. Zhang, K. T. Ma, J. H. Lim, Q. Zhao, and J. Feng, "Deep future gaze: Gaze anticipation on egocentric videos using adversarial networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, pp. 4372–4381, 2017.

[6] B. Lai, M. Liu, F. Ryan, and J. M. Rehg, "In the eye of transformer: Global–local correlation for egocentric gaze estimation and beyond," *Int. J. Comput. Vis.*, vol. 132, no. 3, pp. 854–871, 2024.

[7] P. Li, X. Hou, L. Wei, G. Song, and X. Duan, "Efficient and low-cost deep-learning based gaze estimator for surgical robot control," in *Proc. IEEE Int. Conf. Real-Time Comput. Robot.*, 2018, pp. 58–63.

[8] Y. Li, M. Liu, and J. M. Rehg, "In the eye of beholder: Joint learning of gaze and actions in first person video," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 639–655.

[9] D. Roy, R. Rajendiran, and B. Fernando, "Interaction region visual transformer for egocentric action anticipation," in *Proc. IEEE Winter Conf. Appl. Comput. Vis.*, 2024, pp. 6740–6750.

[10] T. Nagarajan, Y. Li, C. Feichtenhofer, and K. Grauman, "EGO-TOPO: Environment affordances from egocentric video," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 163–172.

[11] L. Mur-Labadia, R. Martinez-Cantin, J. Guerrero, G. M. Farinella, and A. Furnari, "Aff-ttention! affordances and attention models for short-term object interaction anticipation," 2024, *arXiv:2406.01194*.

[12] M. Liu, S. Tang, Y. Li, and J. M. Rehg, "Forecasting human-object interaction: Joint prediction of motor attention and actions in first person video," in *Proc. Eur. Conf. Comput. Vis.*, pp. 704–721, 2020.

[13] S. Liu, S. Tripathi, S. Majumdar, and X. Wang, "Joint hand motion and interaction hotspots prediction from egocentric videos," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 3282–3292.

[14] W. Bao et al., "Uncertainty-aware state space transformer for egocentric 3D hand trajectory forecasting," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2023, pp. 13656–13665.

[15] J. Ma, J. Xu, X. Chen, and H. Wang, "Diff-IP2D: Diffusion-based hand-object interaction prediction on egocentric videos," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2025, pp. 4291–4298,.

[16] J. Ma, X. Chen, W. Bao, J. Xu, and H. Wang, "MADiff: Motion-aware mamba diffusion models for hand trajectory prediction on egocentric videos," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 48, vol. 3, pp. 3250–3267, Mar. 2026.

[17] J. Ma, E. Zhang, J. Xu, X. Chen, and H. Wang, "Improved 2D hand trajectory prediction with multi-view consistency," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2025, pp. 20646–20653.

[18] M. Hatano, R. Hachiuma, and H. Saito, "EMAG: Ego-motion aware and generalizable 2D hand forecasting from egocentric videos," in *Proc. Eur. Conf. Comput. Vis. Workshops*, 2024, pp. 119–136.

[19] B. Zhou, Y. Zhan, Z. Zhang, and Z. Lu, "MEgoHand: Multimodal egocentric hand-object interaction motion generation," 2025, *arXiv:2505.16602*.

[20] K. Li, P. Li, T. Liu, Y. Li, and S. Huang, "ManipTrans: Efficient dexterous bimanual manipulation transfer via residual learning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2025, pp. 6991–7003.

[21] J. Ma, W. Bao, J. Xu, G. Sun, X. Chen, and H. Wang, "Novel diffusion models for multimodal 3D hand trajectory prediction," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2025, pp. 2408–2415.

[22] A. Vaswani et al., "Attention is all you need," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 6000–6010.

[23] A. Gu and T. Dao, "Mamba: Linear-time sequence modeling with selective state spaces," 2023, *arXiv:2312.00752*.

[24] C. Bao, J. Xu, X. Wang, A. Gupta, and H. Bharadhwaj, "HandsOnVLM: Vision-language models for hand-object interaction prediction," 2024, *arXiv:2412.13187*.

[25] S. Bahl, R. Mendonca, L. Chen, U. Jain, and D. Pathak, "Affordances from human videos as a versatile representation for robotics," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, pp. 13778–13790, 2023.

[26] M. Hatano, Z. Zhu, H. Saito, and D. Damen, "The invisible EgoHand: 3D hand forecasting through egobody pose estimation," 2025, *arXiv:2504.08654*.

[27] G. Pavlakos, D. Shan, I. Radosavovic, A. Kanazawa, D. Fouhey, and J. Malik, "Reconstructing hands in 3D with transformers," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2024, pp. 9826–9836.

[28] H. Dong et al., "Hamba: Single-view 3D hand reconstruction with graph-guided bi-scanning mamba," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, vol. 37, pp. 2127–2160, 2024.

[29] M. Li, H. Zhang, Y. Zhang, R. Shao, T. Yu, and Y. Liu, "HHMR: Holistic hand mesh recovery by enhancing the multimodal controllability of graph diffusion models," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2024, pp. 645–654.

[30] Y. Tian, H. Zhang, Y. Liu, and L. Wang, "Recovering 3D human mesh from monocular images: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 12, pp. 15406–15425, Dec. 2023.

[31] J. On et al., "BIGS: Bimanual category-agnostic interaction reconstruction from monocular videos via 3D gaussian splatting," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2025, pp. 17437–17447.

Authorized licensed use limited to: ShanghaiTech University. Downloaded on October 19,2026 at 15:54:43 UTC from IEEE Xplore. Restrictions apply.

[32] Y. Ye, A. Gupta, and S. Tulsiani, "What's in your hands? 3D reconstruction of generic objects in hands," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 3895–3905.
[33] Y. Liu et al., "EasyHOI: Unleashing the power of large models for reconstructing hand-object interactions in the wild," 2024, *arXiv:2411.14280*.
[34] Z. Zhu and D. Damen, "Get a grip: Reconstructing hand-object stable grasps in egocentric videos," 2023, *arXiv:2312.15719*.
[35] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2020, pp. 6840–6851.
[36] P. Dhariwal and A. Nichol, "Diffusion models beat GANs on image synthesis," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2021, pp. 8780–8794.
[37] M. Zhang, Y. Fu, Z. Ding, S. Liu, Z. Tu, and X. Wang, "Hoidiffusion: Generating realistic 3D hand-object interaction data," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2024, pp. 8521–8531.
[38] K. Zhou, B. L. Bhatnagar, J. E. Lenssen, and G. Pons-Moll, "Gears: Local geometry-aware hand-object interaction synthesis," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2024, pp. 20634–20643.
[39] S. Christen et al., "DifH2O: Diffusion-based synthesis of hand-object interactions from textual descriptions," in *Proc. SIGGRAPH Asia 2024 Conf. Papers*, 2024, pp. 1–11.
[40] Z. Zhang, H. Luo, W. Zhai, Y. Cao, and Y. Kang, "PEAR: Phrase-based hand-object interaction anticipation," 2024, *arXiv:2407.21510*.
[41] B. Tang, K. Zhang, W. Luo, W. Liu, and H. Li, "Prompting future driven diffusion model for hand motion prediction," in *Proc. Eur. Conf. Comput. Vis.*, 2024, pp. 169–186.
[42] S. Gong, M. Li, J. Feng, Z. Wu, and L. Kong, "DiffuSeq: Sequence to sequence text generation with diffusion models," in *Proc. Int. Conf. Learn. Representations*, 2023, pp. 2–9.
[43] C. Chi et al., "Diffusion policy: Visuomotor policy learning via action diffusion," *Int. J. Robot. Res.*, vol. 44, pp. 1684–1704, 2023.
[44] Y. Ze, G. Zhang, K. Zhang, C. Hu, M. Wang, and H. Xu, "3D diffusion policy: Generalizable visuomotor policy learning via simple 3D representations," 2024, *arXiv:2403.03954*.
[45] S. Nair, A. Rajeswaran, V. Kumar, C. Finn, and A. Gupta, "R3M: A universal visual representation for robot manipulation," 2022, *arXiv:2203.12601*.
[46] A. Majumdar et al., "Where are we in the search for an artificial visual cortex for embodied intelligence?," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2023, Art. no. 31.
[47] T. Xiao, I. Radosavovic, T. Darrell, and J. Malik, "Masked visual pre-training for motor control," 2022, *arXiv:2203.06173*.
[48] S. Pramanick et al., "EgoVLPv2: Egocentric video-language pre-training with fusion in the backbone," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2023, pp. 5285–5297.
[49] J. Li et al., "OKAMI: Teaching humanoid robots manipulation skills through single video imitation," in *Proc. Annu. Conf. Robot Learn.*, 2024, pp. 299–317.
[50] H. Zhou, R. Wang, Y. Tai, Y. Deng, G. Liu, and K. Jia, "You only teach once: Learn one-shot bimanual robotic manipulation from video demonstrations," 2025, *arXiv:2501.14208*.
[51] J. Ren, P. Sundaresan, D. Sadigh, S. Choudhury, and J. Bohg, "Motion tracks: A unified representation for human-robot transfer in few-shot imitation learning," 2025, *arXiv:2501.06994*.
[52] C. Wen et al., "Any-point trajectory modeling for policy learning," 2023, *arXiv:2401.00025*.
[53] H. Mittal, N. Agarwal, S.-Y. Lo, and K. Lee, "Can't make an omelette without breaking some eggs: Plausible action anticipation using large video-language models," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2024, pp. 18580–18590.
[54] Z. Qi, S. Wang, W. Zhang, and Q. Huang, "Uncertainty-boosted robust video activity anticipation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 46, vol. 12, pp. 7775–7792, 2024.
[55] Z. Zhong, D. Schneider, M. Voit, R. Stiefelhagen, and J. Beyerer, "Anticipative feature fusion transformer for multi-modal action anticipation," in *Proc. IEEE Winter Conf. Appl. Comput. Vis.*, 2023, pp. 6068–6077.
[56] A. Furnari and G. M. Farinella, "Rolling-unrolling LSTMs for action anticipation from first-person video," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 43, no. 11, pp. 4021–4036, Nov. 2021.
[57] A. Stergiou and D. Damen, "The wisdom of crowds: Temporal progressive attention for early action prediction," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 14709–14719.
[58] J. Weng, X. Jiang, W.-L. Zheng, and J. Yuan, "Early action recognition with category exclusion using policy-based reinforcement learning," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 30, no. 12, pp. 4626–4638, Dec. 2020.
[59] T. Perrett, A. Masullo, T. Burghardt, M. Mirmehdi, and D. Damen, "Temporal-relational crosstransformers for few-shot action recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 475–484.
[60] Y.-D. Zheng, Z. Liu, T. Lu, and L. Wang, "Dynamic sampling networks for efficient action recognition in videos," *IEEE Trans. Image Process.*, vol. 29, pp. 7970–7983, 2020.
[61] Y. Wang, Z. Gao, Q. Wang, Z. Chen, P. Li, and Q. Hu, "TAMT: Temporal-aware model tuning for cross-domain few-shot action recognition," 2024, *arXiv:2411.19041*.
[62] M. Hatano, R. Hachiuma, R. Fujii, and H. Saito, "Multimodal cross-domain few-shot learning for egocentric action recognition," in *Proc. Eur. Conf. Comput. Vis.*, 2024, pp. 182–199.
[63] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *Int. J. Comput. Vis.*, vol. 60, pp. 91–110, 2004.
[64] M. A. Fischler and R. C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM*, vol. 24, no. 6, pp. 381–395, 1981.
[65] L. Harold et al., "Grounded language-image pre-training," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 10965–10975.
[66] A. Radford et al., "Learning transferable visual models from natural language supervision," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 8748–8763.
[67] C. Zhang et al., "Faster segment anything: Towards lightweight sam for mobile applications," 2023, *arXiv:2306.14289*.
[68] Y. Li et al., "Egocentric prediction of action target in 3D," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 20971–20980.
[69] D. DeTone, T. Malisiewicz, and A. Rabinovich, "Deep image homography estimation," 2016, *arXiv:1606.03798*.
[70] T. Kwon, B. Tekin, J. Stühmer, F. Bogo, and M. Pollefeys, "H2O: Two hands manipulating objects for first person interaction recognition," in *Proc. Int. Conf. Comput. Vis.*, 2021, pp. 10138–10148.
[71] P. Banerjee et al., "Hot3D: Hand and object tracking in 3D from egocentric multi-view videos," 2024, *arXiv:2411.19167*.
[72] K. Grauman et al., "Ego4D: Around the world in 3,000 hours of egocentric video," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 18995–19012.
[73] D. Damen et al., "Scaling egocentric vision: The epic-kitchens dataset," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 720–736.
[74] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, "Learning fine-grained bimanual manipulation with low-cost hardware," 2023, *arXiv:2304.13705*.
[75] T. Ren et al., "Grounded SAM: Assembling open-world models for diverse visual tasks," 2024, *arXiv:2401.14159*.
[76] E. Zhang, J. Ma, Y.-D. Zheng, Y. Zhou, and H. Wang, "Zero-shot temporal interaction localization for egocentric videos," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2025, pp. 19554–19561.
[77] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*.
[78] Z. Wang et al., "Is mamba effective for time series forecasting?," *Neurocomputing*, vol. 619, 2025, Art. no. 129178.
[79] C. Wang et al., "Chain-of-modality: Learning manipulation programs from multimodal human videos with vision-language-models," 2025, *arXiv:2504.13351*.
[80] P. Hao et al., "TLA: Tactile-language-action model for contact-rich manipulation," *Robot Learn.*, vol. 3, no. 1, pp. 2–18, 2026.



Junyi Ma received the bachelor's and master's degrees from the Beijing Institute of Technology, in 2020 and 2023, respectively. He is currently working toward the PhD degree with IRMV Lab, Shanghai Jiao Tong University. He is currently supervised by Prof. Hesheng Wang. His research interests include human motion analysis, human-robot skill transfer, and computer vision.

Authorized licensed use limited to: ShanghaiTech University. Downloaded on October 19,2026 at 15:54:43 UTC from IEEE Xplore. Restrictions apply.



Wentao Bao received the BS and MS degrees from the Wuhan University, China, and the PhD degree from the Computer Science and Engineering Department, Michigan State University, in USA. His research focuses on open-world computer vision, including video activity recognition, prediction, and understanding in 2D/3D visual space. He has more than 20 research works in leading CV/ML journals or conference venues, such as CVPR, ICCV, ECCV, ACM MM, and *IEEE Transactions on Image Processing*.



Jingyi Xu received the bachelor's and master's degrees from the Beijing Institute of Technology, in 2020 and 2023, respectively. She is currently working toward the PhD degree with the Department of Electronic Engineering, Shanghai Jiao Tong University. Her research interests include intelligent robots and autonomous driving.



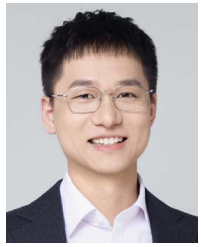
Guanzhong Sun received the bachelor's degree from Nantong University, in 2018, and the master's degree from the China University of Mining and Technology, in 2022. He is currently working toward the PhD degree with the IRMV Lab, China University of Mining and Technology. He is currently supervised by Prof. Hesheng Wang. His research interests include vision-based manipulation and visual servoing.



Yu Zheng received the BEng degree from the Department of Artificial Intelligence, Shanghai Jiao Tong University, China, in 2023. He is currently working toward the PhD degree with Control Science and Engineering, Shanghai Jiao Tong University, China. His current research interests include robot manipulation and computer vision.



Erhang Zhang received the bachelor's degree from Shandong University, in 2025. He is currently working toward the PhD degree with IRMV Lab, Shanghai Jiao Tong University. He is currently supervised by Prof. Hesheng Wang. His research interests include large-scale foundation models and embodied intelligence.



Xieyuanli Chen received the bachelor's degree in electrical engineering and automation from Hunan University, in 2015, the master's degree in robotics from the National University of Defense Technology, in 2017, and the PhD degree from the Photogrammetry and Robotics Laboratory, University of Bonn. He is an associate professor with the National University of Defense Technology, China. He serves as associate editor for IEEE RA-L, ICRA, and IROS. His research interests include SLAM, localization, mapping, and robot perception.



Hesheng Wang (Senior Member, IEEE) received the BEng degree in electrical engineering from the Harbin Institute of Technology, Harbin, China, in 2002, and the MPhil and PhD degrees in automation and computer-aided engineering from the Chinese University of Hong Kong, Hong Kong, in 2004 and 2007, respectively. He is currently a distinguished professor with the School of Automation and Intelligent Sensing, Shanghai Jiao Tong University, Shanghai, China. His current research interests include visual servoing, intelligent robotics, computer vision, and autonomous driving. He is an associate editor of Robotic Intelligence and Automation and the International Journal of Humanoid Robotics, a senior editor of the *IEEE/ASME Transactions on Mechatronics*, an editor-in-chief of Robot Learning. He served as an associate editor of the *IEEE Transactions on Robotics* from 2015 to 2019, an *IEEE Transactions on Automation Science and Engineering* from 2021 to 2023, He was the General Chair of IEEE/RSJ IROS 2025, IEEE ROBIO 2022 and IEEE RCAR 2016.

Authorized licensed use limited to: ShanghaiTech University. Downloaded on October 19,2026 at 15:54:43 UTC from IEEE Xplore. Restrictions apply.