Many Senses, One World: Learning Hierarchical Multimodal World Representations for Bimanual Dexterous Manipulation

Jin Yang, Youlin Wang, Ping Wei,
National Key Laboratory of Human-Machine Hybrid Augmented Intelligence
Institute of Artificial Intelligence and Robotics, Xi'an Jiaotong University
{jin.yang, pingwei}@xjtu.edu.cn

Abstract

Bimanual dexterous manipulation relies on multimodal inputs that capture visual context and non-visual physical interaction cues for precise and coordinated control. To address the challenge of effectively integrating these modalities for policy learning, we propose HiMoDex, a two-stage framework centered on hierarchical multimodal world-state representations. In Stage 1, we pretrain a hierarchical representation network on the collected dataset of 1.19 million synchronized multimodal time steps. The network integrates visual context with tactile, force-torque, and audio cues into unified world-state embeddings. In Stage 2, we freeze the pretrained network and use its embeddings to condition a shared world-action transformer. The transformer jointly learns action generation and action-conditioned future-state prediction. Experiments across five real-world tasks show that HiMoDex achieves an average success rate of 78.7%, representing an absolute improvement of 32.0% over the strongest evaluated baseline. Extensive ablations validate the benefits of hierarchical representation learning and two-stage training. Failure analysis helps explain performance differences across variants. Code, data, and demos will be released at website.

1. Introduction

Bimanual dexterous manipulation is essential for robots to perform real-world tasks, enabling coordinated actions, tool use, and fine manipulation [41]. Building on advances in vision-language models and world models, vision-language-action (VLA) models [4, 5, 23] and world-action models (WAM) [28, 51] have made progress in robotic manipulation. Early methods primarily focused on parallel-jaw grippers, while recent methods have extended these models to bimanual systems with multi-fingered hands [34]. However, these methods face challenges in contact-rich tasks requiring precise control, due to limited physical interaction information from vision alone [49].

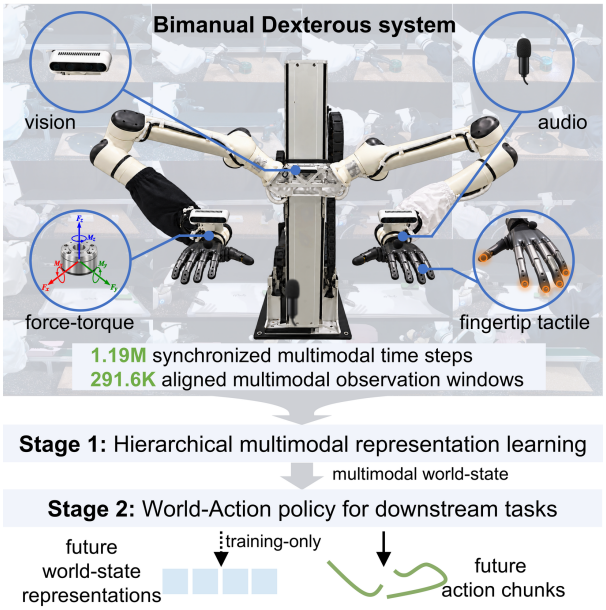


Figure 1. Our bimanual dexterous system integrates multi-view vision, fingertip tactile, wrist force-torque, and audio. We collect 1.19M synchronized multimodal time steps, organized into 291.6K aligned observation windows for multimodal world-state representation learning (Stage 1). The frozen pretrained encoders are then used for downstream policy learning (Stage 2).

Tactile, force-torque, and audio feedback provide complementary cues about physical interaction. Tactile sensing captures hand-object contact, while force sensing measures interaction forces and torques. Audio can reveal impacts, friction, and events such as switch activation [33]. Previous methods have incorporated subsets of these modalities into vision-based policies to improve manipulation reliability [25, 27, 49]. However, jointly exploiting tactile, force-torque, and audio feedback for bimanual dexterous manipulation remains underexplored. A key challenge is to integrate heterogeneous interaction cues with visual context into a unified world-state representation for policy learning.

To this end, we build a bimanual dexterous system with multi-view vision, fingertip tactile sensing, wrist force-torque sensing, and audio. We introduce **HiMoDex**, a two-stage multimodal world-action learning framework for bimanual **Dex**terous manipulation, with **Hi**erarchical **M**ultim**o**dal world-state representation learning as its core, as shown in Fig. 1. In Stage 1, we learn world-state representations from aligned multimodal observation histories through hierarchical encoding and fusion. Multi-view visual histories are encoded into visual context embeddings, while tactile, force-torque, and audio histories are encoded and fused into physical interaction embeddings. A world-state encoder then integrates these intermediate embeddings into unified world-state embeddings. We train this representation network using three complementary objectives: cross-representation consistency, signal reconstruction and change prediction, and task-state contrastive learning.

In Stage 2, we freeze this representation network and use its world-state embeddings for policy learning. The policy combines these embeddings with proprioceptive histories and uses task instructions to guide action generation. A shared world-action transformer jointly learns action generation and action-conditioned future world-state prediction. The future prediction objective provides auxiliary supervision for policy learning. At deployment, only the action generation pathway is retained, avoiding the overhead of future world-state prediction [28, 51].

We evaluate HiMoDex on five real-world tasks covering bimanual coordination, precise manipulation, coordinated tool use, and sustained contact. HiMoDex achieves an average success rate of 78.7%, an absolute improvement of 32.0% over the strongest evaluated baseline. We further conduct extensive ablations to assess the contributions of key components. Task-specific failure statistics further characterize the performance differences across variants. Our main contributions are as follows:

1. We propose a hierarchical multimodal world-state representation learning method that integrates visual context with tactile, force-torque, and audio interaction cues into unified world-state embeddings.
2. We introduce HiMoDex, a two-stage framework using frozen pretrained encoders to learn world-action policies for bimanual dexterous manipulation.
3. HiMoDex outperforms the strongest baseline by 32.0 percentage points in average success rate. Ablations and failure analysis clarify component contributions and performance differences.

2. Related Work

Multimodal Representation Learning for Robotic Manipulation. Multimodal representation learning integrates heterogeneous observations to support robot decision-making and control. Some studies use language supervision to learn representations that support manipulation policies [14, 21, 36, 47]. Specifically, FuSe [20] and OmniVTLA [8] align visual and tactile observations with language descriptions to establish semantic correspondences across modalities. Beyond language supervision, paired multimodal observations provide self-supervised learning signals. Visuo-tactile pretraining learns cross-modal correspondences from spatially paired observations [22]. Masked reconstruction [24] and self-distillation [17] learn robust multimodal representations, while predicting contact events or future observations [16, 30]. HiMoDex introduces a hierarchical representation learning approach that integrates visual context and non-visual interaction cues into a unified world-state representation. This representation is shared by action learning and future-state prediction.

Bimanual Dexterous Manipulation. Compared with bimanual systems using parallel-jaw grippers, those equipped with multifingered hands enable more versatile grasping. Their higher-dimensional action spaces make coordinated control more challenging [41]. To address this challenge, existing methods learn visuomotor policies from expert demonstrations [7, 19, 34, 44]. However, visual observations provide limited information about physical interactions. Recent studies therefore incorporate tactile sensing to complement visual observations [12, 26, 31, 45]. For example, HATO [27] incorporates fingertip tactile feedback into diffusion policies to improve robustness in bimanual block stacking. VTAO-BiManip [42] uses visual-tactile-action pretraining for bimanual bottle-cap unscrewing. In addition to tactile feedback, MoDE-VLA [43] incorporates force observations into a VLA model for contact-rich bimanual dexterous manipulation. In this work, HiMoDex fuses visual, tactile, force, and audio feedback for coordinated control across diverse dexterous manipulation tasks.

World-Action Models for Robotic Manipulation. World-action models extend world models by coupling future-state prediction with action generation. Based on their prediction targets, existing methods can be broadly categorized into generative approaches and latent-space approaches. Generative approaches couple future observation prediction with action learning [15, 46, 50]. For example, DreamZero [48] jointly generates future videos and actions using a pretrained video diffusion backbone. Latent-space approaches instead predict or align future representations to avoid costly reconstruction of action-irrelevant observation details [10, 13, 35]. JEPA-WAM [28] couples action learning with prediction of a joint representation of current and future observations. Beyond the prediction target, the deployment efficiency of WAMs also depends on whether future prediction is performed at inference. Fast-WAM [51] demonstrates that training-time future prediction improves policy performance under action-only inference [10, 28]. In this work, HiMoDex adopts a latent-space approach, jointly

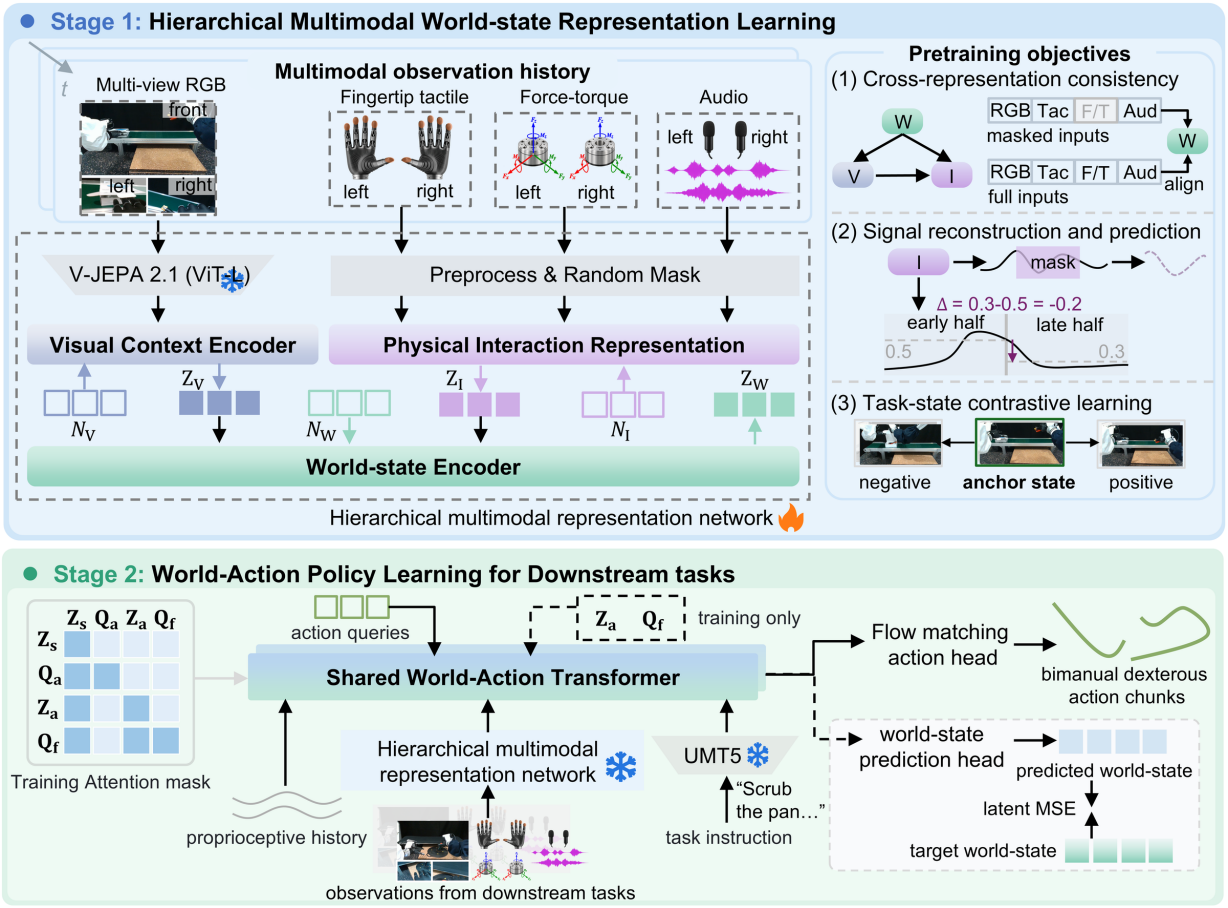


Figure 2. Overview of the two-stage training framework of HiMoDex. **Stage 1** pretrains a hierarchical multimodal representation network using multi-view RGB, fingertip tactile, force-torque, and audio observation histories. **Stage 2** trains a shared world-action transformer for downstream tasks, combining flow-matching action generation with future world-state prediction.

learning action generation and future world-state prediction through a shared world-action transformer.

3. Methodology

3.1. Problem Formulation.

We learn a bimanual dexterous manipulation policy $\pi_\phi(\mathrm{A}_t \mid l, \mathcal{O}_t, \mathrm{P}_t)$ from expert demonstrations, conditioned on a natural-language task instruction l, multimodal observation history $\mathcal{O}_t$, and proprioceptive history P_t. Each action chunk $\mathrm{A}_t = [a_t, \ldots, a_{t+H-1}]$ contains H consecutive actions. $\mathcal{O}_t$ comprises multi-view vision, fingertip tactile, tool force-torque, and audio signals. All data are synchronized using timestamps. Visual, tactile, force-torque, and proprioceptive histories contain $T = 16$ consecutive steps ending at t on the 30 Hz front-view camera timeline, spanning approximately 0.5 s [37]. Following ManiWAV [33], audio is a continuous two-second segment ending at t.

3.2. Architecture Overview.

As shown in Fig. 2, HiMoDex follows two-stage training. In Stage 1 (S1), we pretrain a hierarchical multimodal representation network $\mathcal{F}_\theta$ to encode temporally aligned observation histories into unified world-state embeddings. We maintain an online network $\mathcal{F}_\theta$ and a target network $\mathcal{F}_{\bar{\theta}}$, whose parameters are updated via an exponential moving average (EMA) of the online parameters. The online network is optimized using three pretraining objectives: cross-representation consistency, signal reconstruction and prediction, and task-state contrastive learning. In Stage 2 (S2), we freeze the EMA target network $\mathcal{F}_{\bar{\theta}}$ from S1 and use it to encode observation histories from downstream tasks into world-state embeddings. A shared world-action transformer uses these embeddings to jointly learn action generation and action-conditioned future world-state prediction.

3.3. S1: Hierarchical Multimodal Representation.

In S1, $\mathcal{F}_\theta$ encodes visual and non-visual observation histories into visual context embeddings Z_V and interaction embeddings Z_I, respectively. A world-state encoder then integrates these embeddings into unified world-state embeddings Z_W. This hierarchical representation captures both visual context and physical interactions.

Visual Context Encoding. Multi-view visual signals are acquired from one front-view camera and two wrist-mounted cameras. Each RGB frame is resized to 384×384 pixels and normalized. Each camera sequence is then independently encoded by a frozen V-JEPA 2.1 ViT-L backbone [2, 38]. The extracted features $\mathrm{X}_t^{\mathrm{v}}$ are spatially averaged [39] and temporally interpolated to T observation steps. A trainable linear layer projects the aligned features to dimension D. The visual context encoder then uses N_{V} learnable visual context queries to integrate these multi-view temporal features [18]. It consists of two cross-attention blocks followed by four self-attention blocks. The resulting visual context embeddings are $\mathrm{Z}_{\mathrm{V}} \in \mathbb{R}^{T \times N_{\mathrm{V}} \times D}$.

Physical Interaction Representation. We encode tactile, force-torque, and audio histories with modality-specific Intra-Inter Sensor Transformers and fuse their outputs into interaction embeddings via Gated Interaction Fusion.

Non-visual signal preprocessing. Tactile signals are collected from the five fingertips of each BrainCo Revo2 hand. We normalize the normal force and tangential force magnitude and convert the tangential direction to its sine and cosine components. Force-torque signals are collected from two end-effector sensors, each measuring three force and three torque components. Modality-specific multilayer perceptron (MLP) networks project these measurements into features $\mathrm{X}_t^{\mathrm{tac}} \in \mathbb{R}^{T \times 10 \times D}$ and $\mathrm{X}_t^{\mathrm{f-t}} \in \mathbb{R}^{T \times 2 \times D}$. Audio from two hand-mounted microphones is converted to log-mel spectrograms [33]. A convolutional projection and cross-attention with learnable queries [18] produce audio features $\mathrm{X}_t^{\mathrm{aud}} \in \mathbb{R}^{T \times 2 \times D}$ aligned with the T observation steps. These are augmented with learnable embeddings for modality identity, hand identity, and temporal position.

Intra-Inter Sensor Transformer. Physical interaction cues are reflected in both the temporal evolution of individual sensor signals and the variation in responses across sensing locations. To jointly model these temporal and spatial dependencies, we design the Intra-Inter Sensor Transformer (I^2ST) to learn spatiotemporal representations. Separate I^2ST parameters are used for different modalities.

I^2ST comprises three sequential modules. The intra-sensor encoding module uses two 1D ConvNeXt [32] blocks to capture intra-sensor local temporal dynamics. Each block combines a depthwise temporal convolution with a channel-wise MLP. The inter-sensor attention module then applies self-attention across sensors of the same modality at each time step to model relationships among their responses. Finally, the temporal attention module applies self-attention over the full observation window to capture long-range dependencies within each sensor stream. Finally, I^2ST produces modality-specific features $\mathrm{Z}_{\mathrm{tac}}$, $\mathrm{Z}_{\mathrm{f-t}}$, and $\mathrm{Z}_{\mathrm{aud}}$.

Algorithm 1 Gated Interaction Fusion

Input: Available modalities $\mathcal{M} \subseteq \{\mathrm{tac}, \mathrm{f-t}, \mathrm{aud}\}$; features $\{\mathrm{Z}_m\}_{m \in \mathcal{M}}$; learnable interaction tokens $\mathrm{I} \in \mathbb{R}^{N_{\mathrm{I}} \times D}$; fusion depth L
Output: $\mathrm{Z}_{\mathrm{I}} \in \mathbb{R}^{T \times N_{\mathrm{I}} \times D}$
Initialization: $\mathrm{Z}_m^{(0)} \leftarrow \mathrm{Z}_m$, $m \in \mathcal{M}$; broadcast I over T steps to obtain $\mathrm{I}^{(0)}$
for $\ell = 0, \ldots, L-1$ **do**
  if $\mathcal{M} = \varnothing$ **then**
    $\mathrm{I}^{(\ell+1)} \leftarrow \mathrm{I}^{(\ell)}$
  else
    for $m \in \mathcal{M}$ **do**
      $\mathrm{U}_m \leftarrow E_m^{(\ell)}\big(\mathrm{Concat}(\mathrm{Z}_m^{(\ell)}, \mathrm{I}^{(\ell)})\big)$
      $(\mathrm{Z}_m^{(\ell+1)}, \widetilde{\mathrm{I}}_m) \leftarrow \mathrm{Split}(\mathrm{U}_m)$
      $s_m \leftarrow \mathrm{Linear}_m^{(\ell)}\big(\mathrm{RMSNorm}(\widetilde{\mathrm{I}}_m)\big)$
    end for
    $(\alpha_m)_{m \in \mathcal{M}} \leftarrow \mathrm{Softmax}_{\mathcal{M}}\big((s_m)_{m \in \mathcal{M}}\big)$
    $\mathrm{I}^{(\ell+1)} \leftarrow \sum_{m \in \mathcal{M}} \alpha_m \odot \widetilde{\mathrm{I}}_m$
  end if
end for
Flatten $\mathrm{I}^{(L)}$ into $T \times N_{\mathrm{I}}$, apply two self-attention layers, and obtain Z_{I}

Gated Interaction Fusion. Different modalities provide complementary interaction cues. For example, tactile and force-torque signals can indicate that an electric brush is held, while audio can indicate whether it is powered on. We therefore introduce the Gated Interaction Fusion module to integrate these cues into a unified interaction representation.

Algorithm 1 summarizes the fusion procedure. Shared interaction tokens exchange information with the modality-specific features $\mathrm{Z}_{\mathrm{tac}}$, $\mathrm{Z}_{\mathrm{f-t}}$, and $\mathrm{Z}_{\mathrm{aud}}$ over L rounds. For each modality, a separate module $E_m^{(\ell)}$ updates its features and generates candidate interaction tokens at each observation step. Parameters are separate across modalities and fusion rounds. A softmax gate across available modalities combines these candidates at each time step and interaction-token position. Gate projections are initialized to zero to ensure uniform initial weights.

World-state Encoder. At each observation step, we concatenate N_{W} learnable world-state queries with the visual context embeddings Z_{V} and interaction embeddings Z_{I}. A Transformer module jointly updates these tokens. The outputs at the world-state query positions are then arranged into a sequence of length $T \times N_{\mathrm{W}}$. A second Transformer module models dependencies across time and query positions, producing world-state embeddings $\mathrm{Z}_{\mathrm{W}} \in \mathbb{R}^{T \times N_{\mathrm{W}} \times D}$.

Pretraining objectives. During pretraining, $\mathcal{F}_\theta$ processes observation histories with randomly dropped non-visual modalities, while $\mathcal{F}_{\bar{\theta}}$ receives the same histories without masking. We denote the resulting online and target embeddings by Z_u and $\overline{\mathrm{Z}}_u$, respectively, where $u \in \{\mathrm{V}, \mathrm{I}, \mathrm{W}\}$. We optimize $\mathcal{F}_\theta$ using the following three objectives.

(1) Cross-representation consistency. Following latent-space prediction with EMA targets [1], we predict full-observation target embeddings from online embeddings along three directions: $\mathcal{R} = \{\mathrm{W} \to \mathrm{V}, \mathrm{W} \to \mathrm{I}, \mathrm{V} \to \mathrm{I}\}$. The first two encourage Z_W to retain visual context and physical interaction information. The third encourages Z_V to capture visual cues relevant to physical interactions. Each direction uses an independent predictor $f_{r\to g}$. At each time step, cross-attention updates N_g learnable queries using Z_r as keys and values. The resulting $T \times N_g$ tokens are then processed by self-attention [18]. The prediction matches the shape of its target $\overline{\mathrm{Z}}_g \in \mathbb{R}^{T\times N_g \times D}$. All EMA targets are stop-gradient. In addition, to improve the robustness of world-state representations, we align online world-state embeddings Z_W with their full-observation EMA targets $\overline{\mathrm{Z}}_\mathrm{W}$ under non-visual modality dropout. The combined loss is

$$\begin{aligned}\mathcal{L}_{\text{cross}} = &\sum_{(r\to g)\in\mathcal{R}} \lambda_{r\to g}\, \mathbb{E}\big[\ell_{\text{reg}}\big(f_{r\to g}(\mathrm{Z}_r), \overline{\mathrm{Z}}_g\big)\big] \\ &+ \lambda_{\text{con}}\, \mathbb{E}_{\text{drop}}\big[\ell_{\text{reg}}\big(\mathrm{Z}_\mathrm{W}, \overline{\mathrm{Z}}_\mathrm{W}\big)\big]\,,\end{aligned} \tag{1}$$

where ℓ_{reg} is the mean-squared error. $\mathbb{E}_{\text{drop}}$ averages over windows with at least one dropped non-visual modality. We set $\lambda_{\mathrm{W}\to\mathrm{V}} = \lambda_{\mathrm{W}\to\mathrm{I}} = 0.5$ and $\lambda_{\mathrm{V}\to\mathrm{I}} = \lambda_{\text{con}} = 1.0$.

(2) Signal reconstruction and change prediction. We reconstruct masked tactile and force-torque measurements from online interaction embeddings Z_I. We randomly select one retained tactile or force-torque modality and mask a contiguous measurement segment. We implement masking by replacing the projected modality features in this segment with a learnable mask token. The online network then encodes the masked observations into Z_I. At each time step, we average the N_I embeddings in Z_I and use a modality-specific MLP to reconstruct the masked measurements. The prediction targets comprise normalized and smoothed tactile force magnitudes and end-effector force-torque components. The reconstruction loss $\mathcal{L}_{\text{est}}$ is the smooth L_1 loss at masked positions. Reconstruction is skipped when neither modality is retained.

For each retained modality $m \in \{\mathrm{tac}, \mathrm{f-t}\}$, we also predict the change $\Delta\mathbf{y}_m$ between the two halves of the T-step observation window.

$$\Delta\mathbf{y}_m = \frac{1}{|\mathcal{T}_2|}\sum_{\tau\in\mathcal{T}_2}\mathbf{y}_{m,\tau} - \frac{1}{|\mathcal{T}_1|}\sum_{\tau\in\mathcal{T}_1}\mathbf{y}_{m,\tau}, \tag{2}$$

where $\mathbf{y}_{m,\tau}$ is the target measurement vector. The index sets $\mathcal{T}_1$ and $\mathcal{T}_2$ correspond to the first and second halves, with $|\mathcal{T}_1| = |\mathcal{T}_2| = T/2$. We average Z_I within each half and concatenate the two resulting vectors. A separate modality-specific MLP uses this vector to predict $\widehat{\Delta\mathbf{y}}_m$. We compute the change loss $\mathcal{L}_\Delta$ from the smooth L_1 errors between $\widehat{\Delta\mathbf{y}}_m$ and $\Delta\mathbf{y}_m$. Following [6], we assign larger sample weights to windows with greater signal changes when computing this loss. The combined loss is $\mathcal{L}_{\text{sig}} = \lambda_{\text{est}}\mathcal{L}_{\text{est}} + \lambda_\Delta\mathcal{L}_\Delta$. We set $\lambda_{\text{est}} = \lambda_\Delta = 0.1$.

(3) Task-state contrastive learning. We define task states as manipulation stages within each task and assign a state label to each time step. This objective encourages consistent representations of the same stage across demonstrations while distinguishing different stages. State labels are obtained through an LLM-assisted annotation pipeline based on multimodal cues. These labels supervise the learning of representations that associate each task stage with its characteristic contact patterns.

Within each mini-batch, we construct triplets consisting of an anchor, a positive with the same state label, and a negative with a different state label. For each triplet, we apply a triplet-margin loss [40] with cosine similarity separately to Z_u for $u \in \{\mathrm{V}, \mathrm{I}, \mathrm{W}\}$. The weighted sum of the three losses defines $\mathcal{L}_{\text{state}}$, with each coefficient set to 0.1.

Overall, the pretraining objective of S1 is the sum of $\mathcal{L}_{\text{cross}}$, $\mathcal{L}_{\text{sig}}$, and $\mathcal{L}_{\text{state}}$.

3.4. S2: World-Action Policy Learning.

In S2, we learn a manipulation policy for downstream tasks. We use the frozen EMA network $\mathcal{F}_{\bar{\theta}}$ pretrained in S1 to extract world-state embeddings. A shared world-action transformer is jointly trained for action generation and action-conditioned future world-state prediction.

Shared World-Action Transformer. During training, it receives five groups of inputs. *1) State embeddings* Z_s. We concatenate world-state embeddings from the frozen $\mathcal{F}_{\bar{\theta}}$ with proprioceptive embeddings to form Z_s. Each proprioceptive state in P_t contains 12 arm joint positions, 12 end-effector pose components, and 12 hand joint positions. *2) Action queries* Q_a. The action head uses these H learnable queries to generate $\mathrm{A}_t \in \mathbb{R}^{H\times 36}$. Each action matches the dimensionality of a proprioceptive state. *3) Ground-truth action embeddings* Z_a. An MLP encodes the T ground-truth actions aligned with the future prediction window. With temporal position embeddings added, Z_a conditions future world-state prediction. *4) Future world-state queries* Q_f. These $T \times N_\mathrm{W}$ learnable queries are used to predict future world-state embeddings. Both Z_a and Q_f are used only during training. *5) Task instruction* l. A frozen UMT5 [9] encoder maps l to instruction embeddings. Only the action queries Q_a cross-attend to these embeddings.

Each transformer block applies masked self-attention, instruction cross-attention, and a feed-forward network. Following Fast-WAM [51], we use a structured attention

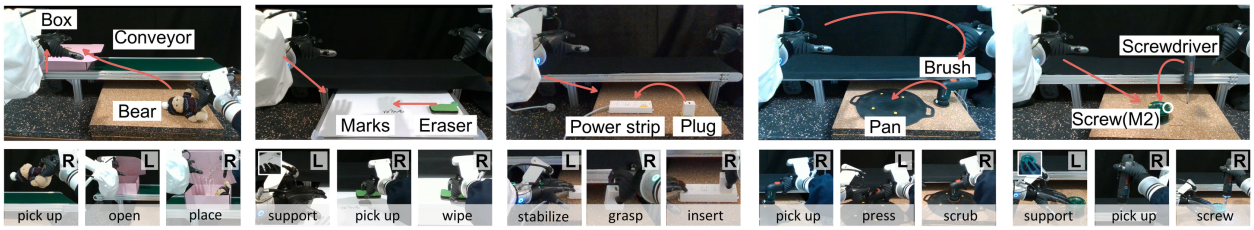


Figure 3. Overview of the five real-world evaluation tasks. (a) Pack Conveyor Box (Pack), (b) Wipe Whiteboard (Wipe), (c) Insert Plug (Insert), (d) Scrub Pan (Scrub), and (e) Tighten Screw (Tighten). These tasks span diverse manipulation skills.

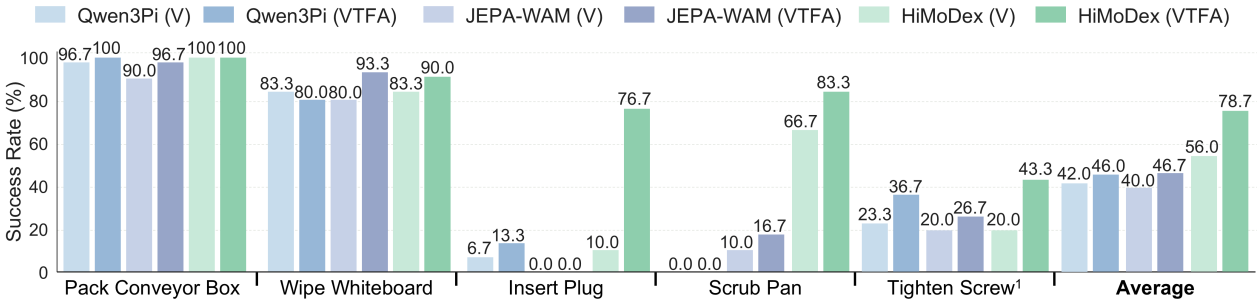


Figure 4. Performance comparison across five real-world bimanual dexterous manipulation tasks. Each method is evaluated over 30 trials per task, and HiMoDex achieves the highest average success rate. No method completes the particularly challenging Tighten Screw task (as described in Sec. 4.1). Tighten Screw[1] indicates the success rate of grasping the screwdriver and bringing it near the screw. The average is computed over the full-task success rates of the first four tasks and the subgoal success rate for Tighten Screw[1].

mask (Fig. 2). Action queries attend to Z_s and one another, while future world-state queries additionally access Z_a. Attention is bidirectional within each token group. This mask prevents ground-truth actions and future prediction features from entering the action generation branch.

Training Objective. We train the action head with flow matching [29], conditioned on the transformer outputs at Q_a. Given noisy actions and flow time, the head predicts velocities along linear paths from Gaussian noise to expert action chunks A_t. The action loss $\mathcal{L}_{\text{FM}}$ is the mean-squared error between predicted and target velocities.

For future world-state prediction, an MLP maps the transformer outputs at Q_f to predicted embeddings. The frozen EMA network encodes the future observation window to provide the targets. The world prediction loss $\mathcal{L}_{\text{W}}$ is the mean-squared error between predicted and target embeddings. Both losses train the shared transformer through $\mathcal{L}_{\text{policy}} = \mathcal{L}_{\text{FM}} + \lambda_{\text{W}}\mathcal{L}_{\text{W}}$, with $\lambda_{\text{W}} = 1.0$. At inference, we retain only the action generation branch.

More details of S1 and S2 are provided in the Appendix.

4. Experiments

4.1. Experimental Setups

Stage 1: Multimodal Pretraining Data. We collected 1,187,768 time steps of synchronized multimodal observations using our robotic platform. These yielded 291,578 aligned observation windows of $T = 16$ steps. We used 287,718 windows for pretraining and reserved 3,860 for representation evaluation. More details are in the Appendix.

Stage 2: Downstream Tasks. We collected 770 expert demonstrations across five bimanual dexterous manipulation tasks to train the WAM policy. These tasks assess a range of manipulation capabilities. We varied initial object positions and orientations across demonstrations to increase data diversity. Despite some scenario overlap with S1, S2 demonstrations were collected independently with different environmental conditions. They were excluded from S1 pretraining. Fig. 3 shows the five task settings.

Task 1: Pack Conveyor Box. The right hand grasps a toy bear while the left hand opens the lid of a box on a conveyor belt. The right hand then places the toy inside. This task evaluates basic bimanual coordination.

Task 2: Wipe Whiteboard. The robot grasps an eraser and wipes marks from the whiteboard. This task evaluates the ability to sustain controlled contact.

Task 3: Insert Plug. The left hand stabilizes a power strip. The right hand grasps a power plug and inserts it into the designated slot. This task evaluates the ability to perform precise, contact-rich manipulation.

Task 4: Scrub Pan. The right hand grasps an electric brush and the left hand presses its activation button. After

Table 1. Ablation results over 10 trials per configuration and task. Success rates (SR, %). F1-F3 report task-specific failure-type counts (0-10). Darker shades indicate higher counts. S1 and S2 in the group labels denote ablations targeting Stage 1 and Stage 2, respectively.

Group	ID	Configuration	Pack			Wipe				Insert			Scrub				Tighten[1]			Avg.↑
			SR↑	F1	F2	SR↑	F1	F2	F3	SR↑	F1	F2	SR↑	F1	F2	F3	SR↑	F1	F2	
I	01	HiMoDex(VTFA)	**100**			80			2	**70**		3	**80**	2			40	6		**74**
II-S1	02	w/o $\mathcal{L}_{\text{sig}}$	**100**			**90**			1	10	6	3	0		10		20	8		44
	03	w/o $\mathcal{L}_{\text{state}}$	**100**			80			2	0		10	10	9			30	5	2	44
	04	w/o $(\mathcal{L}_{\text{sig}} + \mathcal{L}_{\text{state}})$	70	1	2	70			3	30		7	10			9	10	5	4	38
III-S1	05	Flat	90		1	70	2		1	40	2	4	70	3			20	5	3	58
	06	H-concat	90		1	**90**			1	0	6	4	40	3	3		20	8		48
IV-S1	08	w/o gates	90		1	80			2	20		8	70	3			40	6		60
V-S2	09	V	**100**			70			3	10		9	70	2		1	20	8		54
	10	VT	**100**			70			3	**70**		3	70	2		1	40	6		70
	11	VTF	**100**			**90**			1	60		4	60	4			**50**	5		72
VI-S2	12	w/o $\mathcal{L}_{\text{W}}$	**100**			80			2	20		8	60			4	20	8		56
VII	13	End-to-end	**100**			50		3	2	30		7	40	4		2	40	6		52

Failure count: 0 1 2 3 4 5 6 7 8 9 10

Failure types. *Pack:* F1, insecure object grasp; F2, failure to grasp the object. *Wipe:* F1, failure to grasp the eraser; F2, insecure eraser grasp; F3, residual marks after wiping. *Insert:* F1, failure to grasp the plug; F2, failed insertion. *Scrub:* F1, insufficient pressing force to activate the switch; F2, off-target pressing that fails to activate the switch; F3, insufficient scrubbing strokes. *Tighten:* F1: failure to grasp the screwdriver; F2: failure to approach the screw.

activation, the brush makes a sound. The right hand then scrubs the bottom of the pan. This task evaluates coordinated bimanual tool use.

Task 5: Tighten Screw. We introduce this task as the most challenging evaluation setting. It probes the limits of the model's capabilities. First, the front-view camera is 0.6 m from the object, making the M2 screw difficult to observe clearly, as shown in Fig. 3(e). The left arm must therefore stabilize the workpiece while positioning its wrist-mounted camera for a closer view. Second, grasping is more challenging than in *Scrub Pan*, as accidental contact can topple the unsecured, upright electric screwdriver. Third, the right hand must precisely align the screwdriver with the screw head and tighten the 5 mm-long screw. This near-instantaneous tightening phase is difficult to learn.

Evaluation Metrics and Baselines. We evaluate policies by task success rate, using 30 trials per task for main comparisons and 10 for ablations. Detailed success criteria are provided in the Appindex. We compare HiMoDex with Qwen3Pi and JEPA-WAM [28] as representative VLA and WAM baselines, respectively. Qwen3Pi is implemented using StarVLA [11] with a Qwen3-VL [3] backbone and a flow-matching action head. Its design incorporates choices from π_0 [5] and $\pi_{0.5}$ [4]. For each framework, we evaluate vision-only (V) and full-modality (VTFA) variants. The latter uses visual, tactile, force-torque, and audio inputs.

4.2. Main Results

As shown in Fig. 4, HiMoDex (VTFA) achieves an average success rate of 78.7% across five tasks, outperforming all baseline configurations. It achieves absolute gains of 32.7% and 32.0% in success rate over the VTFA variants of Qwen3Pi and JEPA-WAM, respectively. The advantage is pronounced on Insert Plug and Scrub Pan. HiMoDex achieves success rates of 76.7% and 83.3%, whereas the strongest baseline results are only 13.3% and 16.7%, respectively. These results highlight its effectiveness in manipulation tasks involving precise alignment, sustained contact, and tool use. From V to VTFA, HiMoDex's average success rate rises from 56.0% to 78.7%. This gain of 22.7% exceeds the corresponding gains of 4.0% for Qwen3Pi and 6.7% for JEPA-WAM. HiMoDex improves substantially on Insert Plug, Scrub Pan, and Tighten Screw[1]. These differences suggest that the benefits of additional modalities depend on how effectively a model uses them.

On Pack Conveyor Box, all variants achieve at least 90.0% success. Additional modalities provide modest gains, suggesting that visual observations support strong performance on this task. On the most challenging task, Tighten Screw[1], HiMoDex (VTFA) achieves 43.3% success, compared with 36.7% for Qwen3Pi (VTFA). Although the performance gap is modest, HiMoDex uses only 327.62M trainable parameters in S2, approximately one-eighth of Qwen3Pi's 2606.04M.

Model Generalization We evaluate lighting, texture,

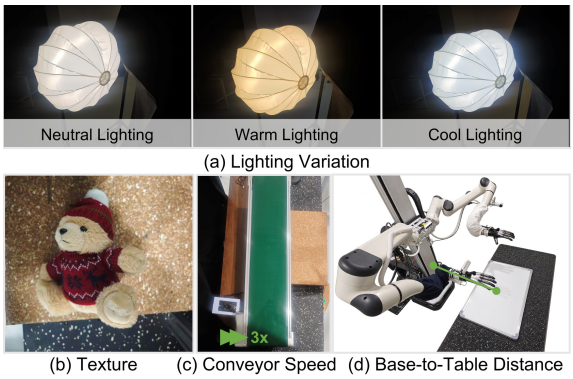


Figure 5. Settings for generalization evaluation. (a) different lighting conditions, (b) changes in object texture, (c) increased conveyor speed, and (d) variations in robot base-to-table distance.

and conveyor-speed variations on Pack Conveyor Box (Fig. 5(a)–(c)) and robot base-to-table distance variations on Wipe Whiteboard (Fig. 5(d)). As shown in Table 2, HiMoDex achieves the highest average success rate of 87.5%, outperforming JEPA-WAM and Qwen3Pi by 7.5 and 10.0 percentage points, respectively. Its high success under lighting and texture changes suggests robustness to visual appearance variations. Under physical variations, HiMoDex outperforms both baselines at altered base-to-table distances and matches the strongest baseline at increased conveyor speed. These results support its ability to handle visual changes while accommodating variations in object motion and workspace geometry.

4.3. Ablation Studies and Analysis

Table 1 presents ablation results. Each configuration is evaluated over 10 trials per task. It also reports task-specific failure types and their frequencies. For a fair comparison, we independently reevaluate HiMoDex (VTFA) and HiMoDex (V) using the same checkpoints as in the main evaluation. It obtains average success rates close to those in Fig. 4.

Effect of S1 Training Objectives. Group II-S1 examines the contributions of $\mathcal{L}_{\text{sig}}$ and $\mathcal{L}_{\text{state}}$. Removing either objective reduces average success from 74% to 44%, while removing both further lowers it to 38% (02-04). These results highlight the importance of both objectives for downstream policy learning. The failure statistics provide further insight into their respective roles. Without $\mathcal{L}_{\text{sig}}$ (02) increases grasping failures on Insert Plug (F1) from zero to six and off-target pressing failures on Scrub Pan (F2) from zero to ten. These positioning errors suggest that $\mathcal{L}_{\text{sig}}$ helps learn representations supporting accurate spatial targeting. In contrast, removing $\mathcal{L}_{\text{state}}$ (03) increases insertion failures on Insert Plug (F2) from three to ten and activation failures caused by insufficient pressing force on Scrub Pan (F1) from two to nine. This pattern suggests that $\mathcal{L}_{\text{state}}$ helps the policy use contact feedback to adjust its actions.

Effect of Hierarchical Fusion. Group III-S1 compares two S1 fusion variants. Flat (05) bypasses the visual context and interaction encoders. It processes modality features and learnable world queries through self-attention to obtain world-state embeddings. H-concat (06) concatenates the visual context and interaction embeddings to form the world representation. Flat reduces average success from 74% to 58% and underperforms the full model on all tasks. It suggests a benefit from forming intermediate representations before constructing world-state embeddings. However, retaining these embeddings alone is still insufficient. H-concat achieves 48% average success, with the largest deficits on Insert Plug and Scrub Pan. On Insert Plug, plug-lifting failures increase from zero to six. These results favor the complete hierarchical design, which combines intermediate representations with a learned world-state encoder.

Table 2. Comparison of generalization performance (%). The four conditions correspond to Fig. 5 (a)-(d).

Method	Lighting	Texture	Speed	Distance	Avg.↑
Qwen3Pi (VTFA)	80	100	70	60	77.5
JEPA-WAM (VTFA)	80	100	60	80	80.0
HiMoDex (VTFA)	90	100	70	90	**87.5**

Modality Contributions. Groups I and V-S2 examine how different modality combinations affect policy performance. As tactile, force, and audio inputs are progressively introduced, the average success rate steadily increases, demonstrating the benefits of multimodal observations. However, additional modalities do not benefit every task. VTF outperforms VTFA on Wipe and Tighten[1] (11 vs. 01). The failure breakdown further suggests that tactile input reduces insertion failures on Insert (09 vs. 10), force reduces residual marks on Wipe (10 vs. 11), and audio reduces switch-activation failures on Scrub (11 vs. 01).

Additional Design Choices. Replacing the learnable gates in Algorithm 1 with equal-weight averaging (08) reduces average success from 74% to 60% (01 vs. 08). This supports the benefit of adaptively weighting modality contributions. In addition, removing $\mathcal{L}_{\text{W}}$ (12) during S2 reduces average success to 56%. This supports the value of future-prediction supervision for policy learning, consistent with the findings of Fast-WAM [51]. Finally, we evaluate an end-to-end variant (13) that jointly trains the representation network and policy in a single stage. This variant retains all representation and policy objectives but achieves only 52% average success. This supports pretraining world-state representations before optimizing the policy.

The Appendix provides more results and details.

5. Conclusion

We propose HiMoDex, a multimodal world-action policy for bimanual dexterous manipulation, learned through two-stage training. In the first stage, we learn hierarchical world-

state representations that integrate visual context with tactile, force-torque, and audio interaction cues. In the second stage, we freeze the pretrained encoders and train a shared world-action transformer for action generation and action-conditioned future-state prediction, retaining only action generation at deployment. Experiments across five real-world task settings show that HiMoDex achieves 78.7% average success, exceeding the strongest baseline by 32.0 percentage points. Ablations validate the benefits of hierarchical multimodal representations and two-stage training. We also contribute a multimodal dataset with 1.19 million synchronized time steps and an expert demonstration dataset for diverse bimanual dexterous manipulation tasks.

LLM Usage Statement

During manuscript preparation, we used Codex to assist with language polishing, grammar correction, and suggestions for improving the clarity, organization, and coherence of the manuscript text. All resulting revisions were reviewed by the authors, who take full responsibility for the accuracy and integrity of the final manuscript.

References

[1] Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun, and Nicolas Ballas. Self-supervised learning from images with a joint-embedding predictive architecture. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15619–15629. IEEE, 2023. 5

[2] Mido Assran, Adrien Bardes, David Fan, Quentin Garrido, Russell Howes, Matthew Muckley, Ammar Rizvi, Claire Roberts, Koustuv Sinha, Artem Zholus, et al. V-jepa 2: Self-supervised video models enable understanding, prediction and planning. *arXiv preprint arXiv:2506.09985*, 2025. 4

[3] Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, et al. Qwen3-vl technical report. *arXiv preprint arXiv:2511.21631*, 2025. 7

[4] Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny Driess, Adnan Esmail, Michael Robert Equi, Chelsea Finn, Niccolo Fusai, Manuel Y. Galliker, Dibya Ghosh, Lachy Groom, Karol Hausman, brian ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Devin LeBlanc, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Allen Z. Ren, Lucy Xiaoyang Shi, Laura Smith, Jost Tobias Springenberg, Kyle Stachowicz, James Tanner, Quan Vuong, Homer Walke, Anna Walling, Haohuan Wang, Lili Yu, and Ury Zhilinsky. $\pi_{0.5}$: a vision-language-action model with open-world generalization. In *Proceedings of the Conference on Robot Learning*, pages 17–40, 2025. 1, 7

[5] Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Robert Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Lucy Xiaoyang Shi, Laura Smith, James Tanner, Quan Vuong, Anna Walling, Haohuan Wang, and Ury Zhilinsky. π_0: A vision-language-action flow model for general robot control. In *Proceedings of Robotics: Science and Systems*, 2025. 1, 7

[6] Ning Cheng, Jinan Xu, Wanlin Li, Yangzhi Chen, Jing Gao, Yiqun Wang, Kelan Peng, and Wenjuan Han. τ: Learning touch-augmented vision-language-action models from future visual supervision. *arXiv preprint arXiv:2607.24485*, 2026. 5

[7] Xuxin Cheng, Jialong Li, Shiqi Yang, Ge Yang, and Xiaolong Wang. Open-television: Teleoperation with immersive active visual feedback. In *Proceedings of the Conference on Robot Learning*, pages 2729–2749, 2025. 2

[8] Zhengxue Cheng, Yiqian Zhang, Anni Tang, Keyu Wang, Wenkang Zhang, Haoyu Li, Hengdi Zhang, and Li Song. Omnivtla: Vision-tactile-language-action models with semantic-aligned tactile sensing. *IEEE Robotics and Automation Letters*, 2026. 2

[9] Hyung Won Chung, Noah Constant, Xavier Garcia, Adam Roberts, Yi Tay, Sharan Narang, and Orhan Firat. Unimax: Fairer and more effective language sampling for large-scale multilingual pretraining. In *International Conference on Learning Representations*, 2023. 5

[10] Mateo Clémente, George Z Xue, Amin Abyaneh, Liheng Ma, Rui Heng Yang, Tongtong Cao, Zhanguang Zhang, and Yingxue Zhang. La-jepa: Compact robot policies with latent world modeling. *arXiv preprint arXiv:2611.08156*, 2026. 2

[11] StarVLA Community. Starvla: A lego-like codebase for vision-language-action model developing. *arXiv preprint arXiv:2604.05014*, 2026. 7

[12] Runyu Ding, Yuzhe Qin, Jiyue Zhu, Chengzhe Jia, Shiqi Yang, Ruihan Yang, Xiaojuan Qi, and Xiaolong Wang. Bunny-visionpro: Real-time bimanual dexterous teleoperation for imitation learning. In *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 12248–12255. IEEE, 2025. 2

[13] Xueji Fang, Boqiang Duan, Hua Wu, Jingdong Wang, and Guo-Jun Qi. Latent evolving world action model. *arXiv preprint arXiv:2609.27455*, 2026. 2

[14] Ruoxuan Feng, Yuxuan Zhou, Siyu Mei, Dongzhan Zhou, Pengwei Wang, Shaowei Cui, Bin Fang, Guocai Yao, and Di Hu. Anytouch 2: General optical tactile representation learning for dynamic tactile perception. *arXiv preprint arXiv:2602.09617*, 2026. 2

[15] Yanjiang Guo, Yucheng Hu, Jianke Zhang, Yen-Jen Wang, Xiaoyu Chen, Chaochao Lu, and Jianyu Chen. Prediction with action: Visual policy learning via joint denoising process. *Advances in Neural Information Processing Systems*, 37:112386–112410, 2024. 2

[16] Liang Heng, Haoran Geng, Kaifeng Zhang, Pieter Abbeel, and Jitendra Malik. Vitacformer: Learning cross-modal representation for visuo-tactile dexterous manipulation. *arXiv preprint arXiv:2506.15953*, 2025. 2

[17] Carolina Higuera, Akash Sharma, Taosha Fan, Chaithanya Krishna Bodduluri, Byron Boots, Michael Kaess, Mike Lambeta, Tingfan Wu, Zixi Liu, Fran-

cois Robert Hogan, et al. Tactile beyond pixels: Multisensory touch representations for robot manipulation. In *Proceedings of the Conference on Robot Learning*, pages 105–123, 2025. 2

[18] Andrew Jaegle, Felix Gimeno, Andy Brock, Oriol Vinyals, Andrew Zisserman, and Joao Carreira. Perceiver: General perception with iterative attention. In *International conference on machine learning*, pages 4651–4664. PMLR, 2021. 4, 5

[19] Zhenyu Jiang, Yuqi Xie, Kevin Lin, Zhenjia Xu, Weikang Wan, Ajay Mandlekar, Linxi Jim Fan, and Yuke Zhu. Dexmimicgen: Automated data generation for bimanual dexterous manipulation via imitation learning. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 16923–16930, 2025. 2

[20] Joshua Jones, Oier Mees, Carmelo Sferrazza, Kyle Stachowicz, Pieter Abbeel, and Sergey Levine. Beyond sight: Finetuning generalist robot policies with heterogeneous sensors via language grounding. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5961–5968. IEEE, 2025. 2

[21] Gi-Cheon Kang, Junghyun Kim, Kyuhwan Shim, Jun Ki Lee, and Byoung-Tak Zhang. Clip-rt: Learning language-conditioned robotic policies from natural language supervision. *arXiv preprint arXiv:2411.00508*, 2024. 2

[22] Justin Kerr, Huang Huang, Albert Wilcox, Ryan Hoque, Jeffrey Ichnowski, Roberto Calandra, and Ken Goldberg. Self-supervised visuo-tactile pretraining to locate and follow garment features. In *Proceedings of Robotics: Science and Systems*, 2023. 2

[23] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan P Foster, Pannag R Sanketi, Quan Vuong, Thomas Kollar, Benjamin Burchfiel, Russ Tedrake, Dorsa Sadigh, Sergey Levine, Percy Liang, and Chelsea Finn. Openvla: An open-source vision-language-action model. In *Proceedings of the Conference on Robot Learning*, pages 2679–2713, 2025. 1

[24] Rickmer Krohn, Vignesh Prasad, Gabriele Tiboni, and Georgia Chalvatzaki. Self-supervised multisensory pretraining for contact-rich robot reinforcement learning. *IEEE Robotics and Automation Letters*, 11(6):6799–6806, 2026. 2

[25] Hao Li, Yizhi Zhang, Junzhe Zhu, Shaoxiong Wang, Michelle A Lee, Huazhe Xu, Edward Adelson, Li Fei-Fei, Ruohan Gao, and Jiajun Wu. See, hear, and feel: Smart sensory fusion for robotic manipulation. In *Proceedings of the Conference on Robot Learning*, pages 1368–1378, 2023. 1

[26] Xukun Li, Yu Sun, Lei Zhang, Bo-Sheng Huang, Yibo Peng, Yuan Meng, Haojun Jiang, Shaoxuan Xie, Guocai Yao, Alois Knoll, Zhenshan Bing, Xinlong Wang, and Zhenguo Sun. DECO: Decoupled multimodal diffusion transformer for bimanual dexterous manipulation with a plugin tactile adapter. In *Proceedings of the 43rd International Conference on Machine Learning*, pages 70113–70133, 2026. 2

[27] Toru Lin, Yu Zhang, Qiyang Li, Haozhi Qi, Brent Yi, Sergey Levine, and Jitendra Malik. Learning visuotactile skills with two multifingered hands. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 5637–5643, 2025. 1, 2

[28] Yihan Lin, Jiawei He, Shifeng Bao, Chen Zhao, Yang Li, Xiaobo Wang, Yan Wang, Cheng Chi, and Jing Zhang. Jepa-wam: Learning vision-language-action policies with joint-embedding world modeling. *arXiv preprint arXiv:2608.09381*, 2026. 1, 2, 7

[29] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022. 6

[30] Fangchen Liu, Chuanyu Li, Yihua Qin, Jing Xu, Pieter Abbeel, and Rui Chen. Vitamin: Learning contact-rich tasks through robot-free visuo-tactile manipulation interface. *arXiv preprint arXiv:2504.06156*, 2025. 2

[31] Xiangcheng Liu, Tianhao Wu, Le Zheng, Yidong Wang, Bowen Jiang, Mingjie Pan, Xinlin Ren, Yi Liu, and Jianlan Luo. Star: Sparse tactile representation learning in vision-tactile-language-action models for dexterous manipulation. *arXiv preprint arXiv:2609.12549*, 2026. 2

[32] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *IEEE/CVF conference on computer vision and pattern recognition*, pages 11966–11976. IEEE, 2022. 4

[33] Zeyi Liu, Cheng Chi, Eric Cousineau, Naveen Kuppuswamy, Benjamin Burchfiel, and Shuran Song. Maniwav: Learning robot manipulation from in-the-wild audio-visual data. In *Proceedings of the Conference on Robot Learning*, pages 947–962, 2025. 1, 3, 4

[34] Hao Luo, Ye Wang, Wanpeng Zhang, Sipeng Zheng, Ziheng Xi, Chaoyi Xu, Haiweng Xu, Haoqi Yuan, Chi Zhang, Yiqing Wang, et al. Being-h0.5: Scaling human-centric robot learning for cross-embodiment generalization. *arXiv preprint arXiv:2601.12993*, 2026. 1, 2

[35] Hao Luo, Wanpeng Zhang, Yicheng Feng, Sipeng Zheng, Haiweng Xu, Chaoyi Xu, Ziheng Xi, Yuhui Fu, and Zongqing Lu. Being-h0. 7: A latent world-action model from egocentric videos. *arXiv preprint arXiv:2605.00078*, 2026. 2

[36] Wenxuan Ma, Chaofan Zhang, Yinghao Cai, Guocai Yao, Shaowei Cui, and Shuo Wang. Fg-cltp: Fine-grained contrastive language tactile pretraining for robotic manipulation. *IEEE Robotics and Automation Letters*, 2026. 2

[37] Tomohiro Motoda, Masaki Murooka, Keisuke Shirai, Hanbit Oh, Ryoichi Nakajo, Shotaro Miwa, Roman Mykhailyshyn, Hugo Duarte, and Yukiyasu Domae. Self-supervised anchoring of fingertip sensing to proprioception and proactive actions for robot imitation learning. *arXiv preprint arXiv:2609.29822*, 2026. 3

[38] Lorenzo Mur-Labadia, Matthew Muckley, Amir Bar, Mido Assran, Koustuv Sinha, Mike Rabbat, Yann LeCun, and Nicolas Ballas. V-jepa 2.1: Unlocking dense features in video self-supervised learning. In *European Conference on Computer Vision*, pages 671–689. Springer, 2026. 4

[39] Sharon S Musa, Fereshteh Forghani, Harrish Thasarathan, Sonia Joseph, Matthew Kowal, and Konstantinos G Derpanis. What, where, and how: Probing spatiotemporal representations in video foundation models. *arXiv preprint arXiv:2609.01551*, 2026. 4

[40] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clus-

tering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015. 5

[41] Kenneth Shaw, Yulong Li, Jiahui Yang, Mohan Kumar Srirama, Ray Liu, Haoyu Xiong, Russell Mendonca, and Deepak Pathak. Bimanual dexterity for complex tasks. In *Proceedings of the Conference on Robot Learning*, pages 5166–5183, 2025. 1, 2

[42] Zhengnan Sun, Zhaotai Shi, Jiayin Chen, Qingtao Liu, Yu Cui, Jiming Chen, and Qi Ye. Vtao-bimanip: Masked visual-tactile-action pre-training with object understanding for bimanual dexterous manipulation. In *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3201–3208. IEEE, 2025. 2

[43] Tutian Tang, Xingyu Ji, Wanli Xing, Ce Hao, Wenqiang Xu, Lin Shao, Cewu Lu, Qiaojun Yu, Jiangmiao Pang, and Kaifeng Zhang. Towards human-like manipulation through rl-augmented teleoperation and mixture-of-dexterous-experts vla. *arXiv preprint arXiv:2603.08122*, 2026. 2

[44] Chen Wang, Haochen Shi, Weizhuo Wang, Ruohan Zhang, Li Fei-Fei, and C Karen Liu. Dexcap: Scalable and portable mocap data collection system for dexterous manipulation. In *Proceedings of Robotics: Science and Systems*, 2024. 2

[45] Xinyi Wang, Donghan Li, Zi'Ang Chen, Chong Yu, Chen Xin, Peng Ye, Yingkai Sun, and Tao Chen. Robotacdex: A dexterous visual-tactile-action dataset for humanoid manipulation. *arXiv preprint arXiv:2606.31836*, 2026. 2

[46] Hongtao Wu, Ya Jing, Chilam Cheang, Guangzeng Chen, Jiafeng Xu, Xinghang Li, Minghuan Liu, Hang Li, and Tao Kong. Unleashing large-scale video generative pre-training for visual robot manipulation. In *International Conference on Learning Representations*, pages 10641–10662, 2024. 2

[47] Jingyun Yang, Baiyu Shi, Timothy Yu, Haitian Liu, Alberta Longhini, Weichen Wang, Rika Antonova, Zhenan Bao, and Jeannette Bohg. Sota: Soft tactile skins for dexterous manipulation. *arXiv preprint arXiv:2610.02338*, 2026. 2

[48] Seonghyeon Ye, Yunhao Ge, Kaiyuan Zheng, Shenyuan Gao, Sihyun Yu, George Kurian, Suneel Indupuru, You Liang Tan, Chuning Zhu, Jiannan Xiang, et al. World action models are zero-shot policies. *arXiv preprint arXiv:2602.15922*, 2026. 2

[49] Jiawen Yu, Hairuo Liu, Qiaojun Yu, Jieji Ren, Ce Hao, Haitong Ding, Guangyu Huang, Guofan Huang, Yan Song, Panpan Cai, et al. Forcevla: Enhancing vla models with a force-aware moe for contact-rich manipulation. *Advances in Neural Information Processing Systems*, 38:93409–93439, 2025. 1

[50] Haoran Yuan, Zekai Wang, Boning Shao, Haoran Lu, Trevor Darrell, Ismini Lourentzou, and Wei Zhan. Dextacwam: A visuo-tactile world-action model for dexterous manipulation. *arXiv preprint arXiv:2609.24976*, 2026. 2

[51] Tianyuan Yuan, Zibin Dong, Yicheng Liu, and Hang Zhao. Fast-wam: Do world action models need test-time future imagination? *arXiv preprint arXiv:2603.16666*, 2026. 1, 2, 5, 8