

LaMP: Learning Vision-Language-Action Policy with 3D Scene Flow as Latent Motion Prior

Xinkai Wang^{1,3}, Chenyi Wang³, Yifu Xu^{2,3}, Mingzhe Ye², Fucheng Zhang³, Jialin Tian^{2,3}, Xinyu Zhan^{2,3}, Lifeng Zhu¹, Cewu Lu^{2,3,4}, and Lixin Yang^{✉2,3}

¹ Southeast University, China ² Shanghai Jiao Tong University, China

³ Shanghai Innovation Institute, China ⁴ Noematrix, China

xinkaiwang@sii.edu.cn, siriussyang@sjtu.edu.cn




Fig. 1: LaMP: a Vision-Language-Action Model with Latent Motion Prior.

LaMP introduces a dense 3D motion prior between VLM perception and action generation. The Motion Expert provides geometric foresight via one-step denoised features, which guide the Action Expert through gated cross-attention for physically grounded manipulation. LaMP achieves superior performance across real-world tasks and simulation benchmarks, significantly outperforming prior VLAs.

Abstract. We introduce **LaMP**, a dual-expert Vision-Language-Action framework that embeds dense 3D scene flow as a latent motion prior for robotic manipulation. Existing VLA models regress actions directly from 2D semantic visual features, forcing them to learn complex 3D physical interactions implicitly. This implicit learning strategy degrades under unfamiliar spatial dynamics. LaMP addresses this limitation by aligning a flow-matching *Motion Expert* with a policy-predicting *Action Expert* through gated cross-attention. Specifically, the Motion Expert generates a one-step partially denoised 3D scene flow, and its hidden states condition the Action Expert without full multi-step reconstruction. We evaluate LaMP on the LIBERO, LIBERO-Plus, and SimplerEnv-WidowX

✉ Corresponding author.

simulation benchmarks as well as real-world experiments. LaMP consistently outperforms evaluated VLA baselines across LIBERO, LIBERO-Plus, and SimplerEnv-WidowX benchmarks, achieving the highest reported average success rates under the same training budgets. On LIBERO-Plus OOD perturbations, LaMP shows improved robustness with an average 9.7% gain over the strongest prior baseline. Our project page is available at <https://summerwxx.github.io/lamp-project-page/>.

Keywords: Robotic Manipulation · 3D Scene Flow · Vision-Language-Action Model

1 Introduction

Humans rarely execute motor tasks by directly mapping raw perception to low-level motion commands. Before grasping an object or closing a drawer, people implicitly form a world model of the necessary motion dynamics: which entities interact, where the interaction occurs, and how the movement unfolds. This latent motion planning bridges the gap between high-level intentions and low-level control [15, 43, 47]. Current VLA models lack an explicit intermediate representation that captures these 3D dynamic relationships, forcing them to infer contact-aware motion implicitly from action labels alone.

Existing Vision-Language-Action (VLA) models aim to equip robots with the ability to execute natural language instructions in visually rich environments [1, 13, 17, 19, 20, 34, 35]. By mapping visual and linguistic inputs to low-level motor commands, recent VLA models such as π_0 [2], $\pi_{0.5}$ [17], Gr00t [1], UniVLA [4], and FlowVLA [51] have achieved strong generalization across tasks and embodiments. However, a fundamental representational mismatch persists: VLM features are largely semantic and 2D-centric, whereas robot manipulation requires reasoning about explicit 3D dynamics for precise control. This mismatch forces models to learn complex 3D physical interactions implicitly from action supervision alone, often leading to brittle execution under unfamiliar spatial dynamics. For example, purely 2D-centric policies frequently fail on tasks requiring precise depth control (such as inserting pegs into holes or stacking objects with tight clearance) when camera viewpoints or object positions differ from training, as reported in LIBERO-Plus camera and layout perturbation benchmarks [12]. We ask: *Can dense 3D scene flow serve as a latent motion prior that explicitly grounds VLA policy learning in physical geometry?*

Prior works attempt to bridge the 2D-to-3D gap through various intermediate modalities. FlowVLA and TraceVLA use optical flow and visual traces for temporal reasoning [48, 51], but 2D flow lacks explicit depth information required for contact-rich manipulation. $\mathcal{F}_1$ and FLARE inject visual foresight into VLA control loops [30, 49], yet they predict pixel-space futures that conflate appearance changes with actionable kinematics. Point-cloud diffusion policies operate directly in 3D environments [40, 45], but they process geometric features as static inputs rather than modeling continuous temporal dynamics. Even methods investigating 3D traces or flows for world modeling [22, 50] generally function as

standalone pipelines trained per task, without access to pretrained VLMs for semantic understanding or cross-task transfer. Consequently, these approaches struggle to provide the continuous, scene-level spatial geometry required for robust physical interactions within a unified, language-conditioned policy.

We argue that dense 3D scene flow is not merely an observable byproduct of actions, but a fundamental latent motion prior that should actively guide policy learning. Building on this insight, we propose **LaMP** (**Latent Motion Prior**), a dual-expert VLA framework (Fig. 1). LaMP aligns a flow-matching Motion Expert with a policy-predicting Action Expert through gated cross-attention. Instead of fully reconstructing 3D scene flow (which incurs prohibitive latency), LaMP conditions action prediction on one-step partially denoised hidden motion states, inspired by predictive latent conditioning in VPP [14] and mimic-video [32]. This one-step strategy avoids the computational burden of multi-step denoising while retaining task-relevant dynamics. The gated cross-attention mechanism starts from near-zero injection and adaptively increases motion guidance only when beneficial, preventing representational collapse of pretrained VLM features while injecting 3D dynamics cues only when spatial precision is required. This design yields actionable geometric foresight without the computational expense of full multi-step reconstruction, enabling efficient robotic deployment.

LaMP achieves the highest average success rates across LIBERO [29], LIBERO-Plus [12], and SimplerEnv-WidowX [25] benchmarks, outperforming general VLAs and geometry-aware models under identical data budgets.

In summary, our contributions are as follows:

- We propose **LaMP**, a dual-expert VLA framework that leverages dense 3D scene flow as a latent motion prior, bridging the gap between 2D semantic VLM features and 3D physical dynamics.
- We introduce a gated motion guidance mechanism that extracts partial denoised motion states and adaptively injects geometric cues into VLM representations via gated cross-attention, preventing representational collapse while maintaining computational efficiency.
- Extensive experiments on LIBERO, LIBERO-Plus, SimplerEnv-WidowX, and real-world tasks confirm that LaMP consistently outperforms existing VLA baselines across all benchmarks.

2 Related Work

2.1 General Vision-Language-Action Models

VLA models combine pretrained VLMs with action generation, enabling instruction-following at scale [1, 2, 13, 16, 17, 20, 35]. These approaches directly map observations to actions via end-to-end architectures, encoding scene dynamics implicitly within the policy latent space. While achieving strong generalization on standard benchmarks, this implicit encoding often leads to compounding execution

errors during long-horizon tasks. Diffusion-based visuomotor policies [8] generate multimodal action sequences through iterative denoising, and point-cloud variants [40, 45] mitigate spatial ambiguity by operating in 3D, yet they process geometric features as static inputs rather than modeling continuous temporal dynamics. 3D-VLA [47] incorporates 3D scene understanding through point-cloud inputs, but does not follow the SimplifiedEnv/LIBERO evaluation protocols. We therefore treat it as related work rather than a direct end-to-end baseline.

2.2 World Model for Visuomotor Policy

To address the limitation of reactive policies, world model approaches predict future states before action generation. Pixel-space methods [3, 5, 10, 46] offer high expressivity but dedicate substantial capacity to reconstructing visual appearance, distracting optimization from learning actionable kinematics. Motion-space methods provide a more control-aligned abstraction by operating directly on geometric dynamics. $\mathcal{F}_1$ [30] and mimic-video [32] condition action synthesis on predicted future visual representations. 3DFlowAction [50], 3D FDP [31], NovaFlow [23], and TraceGen [22] further ground this prediction in 3D space, using predicted flow trajectories or traces as structured motion priors [27]. While 3D spatial priors improve physical grounding over pixel-space methods, these approaches share a common limitation: they function as standalone pipelines trained per task, without access to pretrained VLMs for semantic understanding or cross-task transfer. LaMP takes a different position. Rather than treating 3D flow prediction as a standalone world model, we integrate it as an internal latent prior within a pretrained VLA framework, enabling embodiment-agnostic knowledge sharing and language-conditioned control with a single unified policy. Unlike prior 3D-aware VLAs [36, 37] that inject *current-frame* 3D geometry, LaMP integrates a *future, chunk-level* dense 3D scene-flow prior within a pretrained VLA framework. Prior 3D flow methods [22, 31, 50] operate as standalone pipelines without VLM semantic pretraining. We found naive fusion of 3D features into VLM layers causes representational collapse, necessitating our gated cross-attention design.

2.3 Intermediate Representations for Manipulation

An alternative strategy introduces auxiliary supervision on current-frame representations rather than predicting futures [44]. FlowVLA [51] and TraceVLA [48] use 2D optical flow and trace supervision to enhance spatial reasoning. Qdepth-VLA, MolmoACT [21, 26] employs depth prediction as auxiliary supervision. UniVLA, villaX [4, 7] learns latent action representations for cross-task transfer. These approaches improve geometric awareness but operate in 2D pixel space or on static features [41], lacking dense temporal-spatial dynamics. In contrast, LaMP uses dense scene-level 3D flow as a latent motion prior within the VLA framework, capturing both spatial geometry and temporal dynamics for contact-aware control.




Fig. 2: Overview of LaMP. (a) Motion Pre-training: The Motion Expert learns 3D scene flow prediction conditioned on VLM features. (b) Action Post-Training: Frozen Motion Expert provides one-step denoised features that fuse with VLM representations via Gated Cross-Attention. (c) Data Curation: 1.6M observation-language-motion triplets from diverse robot embodiments.

3 Method

3.1 Problem Formulation

We consider the problem of predicting an action chunk $\mathbf{A}_{t:t+H}$ given observation O_t and instruction l . Let $\mathbf{z}_t = \phi_{vl}(O_t, l)$ denote the latent representation produced by a pre-trained vision-language backbone. A standard VLA directly models $P(\mathbf{A}_{t:t+H} | \mathbf{z}_t)$.

Rather than learning a direct perception-to-action mapping, we introduce a latent 3D motion representation $\mathbf{M}_t$ as an explicit intermediate variable. The joint distribution of actions and motion, conditioned on the current perception, factorizes naturally by the chain rule as:

$$P(\mathbf{A}_{t:t+H}, \mathbf{M}_t | \mathbf{z}_t) = \underbrace{P(\mathbf{A}_{t:t+H} | \mathbf{z}_t, \mathbf{M}_t)}_{\text{Action Expert}} \underbrace{P(\mathbf{M}_t | \mathbf{z}_t)}_{\text{Motion Expert}} \quad (1)$$

This factorization does not assume conditional independence; the Action Expert conditions on both perceptual features $\mathbf{z}_t$ and motion prior $\mathbf{M}_t$, allowing the model to leverage both semantic and dynamic information jointly. The Motion Expert models $P(\mathbf{M}_t | \mathbf{z}_t)$ via conditional flow matching, learning to generate 3D scene flow that captures the expected scene dynamics. The Action Expert then models $P(\mathbf{A}_{t:t+H} | \mathbf{z}_t, \mathbf{M}_t)$, predicting robot actions conditioned on both the perceptual features and the motion prior.

The overall pipeline is illustrated in Figure 2. Section 3.2 describes the Motion Expert, Section 3.3 introduces the Motion Guidance module and the Action Expert, and Section 3.4 details the two-stage training and inference procedures.

3.2 Motion Expert: Learning a Latent 3D Motion Prior

Existing intermediate representations in VLAs, such as optical flow [11, 51] and visual traces [33, 48], either remain in 2D image space or capture only sparse point trajectories, limiting their spatial consistency for manipulation tasks.

3D Scene Flow Representation. We represent the latent motion as a *dense 3D scene flow* defined over a uniform spatial grid. Specifically, we sample a $K_h \times K_w$ grid of keypoints on the current image plane and track each point across T future timesteps. For each keypoint at each timestep, we predict the displacement increment $\Delta \mathbf{p} = (\Delta u, \Delta v, \Delta d)$, where $(\Delta u, \Delta v)$ denotes the 2D displacement in image coordinates and Δd denotes the change in depth. The full motion representation is $\mathbf{M}_t \in \mathbb{R}^{K \times T \times 3}$, where $K = K_h \times K_w = 400$ and $T = 32$. Compared to 2D flow or sparse traces, this representation preserves scene-level 3D geometry while remaining embodiment-agnostic, providing a strong inductive bias for cross-embodiment generalization. All trajectories are transformed into a unified reference camera frame cam_{ref} to compensate for camera motion. This screen-aligned (u, v, d) representation maintains compatibility with 2D trajectory formats, enabling joint training across diverse data sources. The dense correspondences are produced offline by the TraceForge labeler [22], which tracks each keypoint over $T=32$ future steps. Our 1.6M training triplets use TraceForge-provided depth, falling back to sensor depth where available (*e.g.*, DROID [18]). Because depth enters only at this offline labeling stage, the depth source is interchangeable and no depth sensor is required at inference.

The Motion Expert adopts a CogVideoX-style [42] 3D transformer that operates in motion-token space. We apply spatial patchification with patch size 2×2 , grouping each 2×2 block of keypoints into a single token, resulting in 10×10 spatial tokens per timestep. The model is conditioned on the last-layer VLM hidden features $\mathbf{z}_t$ via Adaptive LayerNorm.

We train the Motion Expert using conditional flow matching [28]. Given a ground-truth motion sample $\mathbf{M}_t$ and Gaussian noise $\epsilon \sim \mathcal{N}(0, I)$, we construct the interpolated noisy motion at time $\tau \in [0, 1]$:

$$\mathbf{M}^\tau = (1 - \tau) \epsilon + \tau \mathbf{M}_t \quad (2)$$

The Motion Expert v_θ is trained to predict the velocity field $\mathbf{M}_t - \epsilon$ that transports noise to data:

$$\mathcal{L}_{\text{motion}} = \mathbb{E}_{\tau, \mathbf{M}_t, \epsilon} \|v_\theta(\mathbf{M}^\tau, \tau, \mathbf{z}_t) - (\mathbf{M}_t - \epsilon)\|_2^2 \quad (3)$$

One-step denoising features. Inspired by the partial-denoising conditioning used in video-conditioned policy learning (*e.g.*, VPP and mimic-video [14, 32]), the proposed LaMP framework does not require fully denoised motion latents. Given a pure noise input $\mathbf{M}^0 \sim \mathcal{N}(0, I)$ initialized at $\tau = 0$, we simulate a single denoising step out of $N = 10$ total solver steps (reaching $\tau = 0.1$) and extract the intermediate motion hidden states: $\mathbf{z}_m \in \mathbb{R}^{B \times N \times d_m}$. These features retain the task-relevant dynamics while circumventing the computational cost and bias associated with forcing fully clean motion generation prior to action prediction.

3.3 Motion Guidance and Action Expert

Single-layer motion guidance. Existing architectures typically perform deep fusion across multiple VLM layers. To preserve the pretrained VLM representations while still injecting temporal dynamics, we adapt gated cross-attention to fuse one-step partially denoised motion features into the last-layer VLM features. Let $\tilde{\mathbf{h}}_{\text{motion}} = \text{LN}(\mathbf{W}_{\text{proj}} \mathbf{z}_{\text{m}})$, where $\text{LN}(\cdot)$ denotes layer normalization:

$$\mathbf{z}_{\text{guided}} = \mathbf{z} + \sigma(g) \cdot \text{CA}\left(Q = \text{LN}(\mathbf{z}), K = V = \tilde{\mathbf{h}}_{\text{motion}}\right) \quad (4)$$

where CA represents the multi-head cross-attention operation, and g is a learnable scalar gate initialized at zero, shared across all attention heads and spatial positions. We apply the sigmoid activation $\sigma(\cdot)$ to constrain the gate value to $(0, 1)$, allowing the model to start from a weak motion-injection regime and increase the guidance only when it proves beneficial, leading to more stable optimization.

Action expert. The Action Expert receives the guided features $\mathbf{z}_{\text{guided}}$ and predicts continuous actions via flow matching. Let $\mathbf{a}^\tau = (1 - \tau)\boldsymbol{\epsilon}^a + \tau\mathbf{a}$, where $\boldsymbol{\epsilon}^a \sim \mathcal{N}(0, I)$. The action objective is given by

$$\mathcal{L}_{\text{action}} = \mathbb{E}_{\tau, \mathbf{a}, \boldsymbol{\epsilon}^a} \|v_\phi(\mathbf{a}^\tau, \tau, \mathbf{z}_{\text{guided}}) - (\mathbf{a} - \boldsymbol{\epsilon}^a)\|_2^2 \quad (5)$$

Utilizing motion-guided features as the condition ensures that action denoising depends on explicit 3D dynamics rather than solely on 2D implicit semantics.

3.4 Two-Stage Training and Inference

Stage 1: motion-prior pretraining. We train the Motion Expert using the objective $\mathcal{L}_{\text{motion}}$ on dense 3D scene-flow supervision generated via the TraceForge pipeline [22]. Using this method, we construct 1.6M observation-language-motion triplets from the LIBERO [29], BridgeV2 [39], DROID [18], and InternData-A1 [38] datasets.

Stage 2: motion-guided policy learning. We freeze the pretrained Motion Expert and the VLM backbone, and subsequently train the Motion Guidance module and the Action Expert. The optimization in this stage employs solely the objective $\mathcal{L}_{\text{action}}$. Crucially, we obtain the guidance features by performing the exact same partial denoising process used during inference: starting from pure noise $\mathbf{m}^0 \sim \mathcal{N}(0, I)$, we run the Motion Expert for one step of a 10-step ODE solver and extract the resulting hidden state.

Inference. During inference, LaMP performs four sequential steps: (1) encode the RGB observation o_t and the instruction l using the VLM and read the last-layer features; (2) execute the Motion Expert starting from Gaussian noise at $\tau = 0$ for one step (out of $N = 10$) to $\tau = 0.1$ and extract the motion hidden

state $\mathbf{z}_m$; (3) fuse the motion information into the VLM features via a single-layer gated cross-attention; (4) denoise in the action space from $\tau = 0$ to $\tau = 1$ to sample the clean action $\mathbf{a}_t$. The resulting policy operates online and does not require full motion reconstruction during inference.

4 Simulation Experiments

We design our empirical evaluation to answer the following research questions:

- **Q1: Geometric Foresight.** Does explicitly modeling future 3D scene flow as a generative action prior improve performance on long-horizon planning and fine-grained spatial reasoning tasks? (see Tab. 1 and Sec. 4.2)
- **Q2: 3D Geometry vs. 2D Pixels.** Is depth-aware 3D scene flow superior to 2D optical flow for guiding manipulation policies? (see Fig. 3)
- **Q3: Out-of-Distribution Robustness.** Does the 3D motion prior facilitate robust zero-shot generalization under visual perturbations and cross-domain transfer? (see Tab. 1 and Tab. 2)
- **Q4: Fusion Strategy.** Does gated cross-attention outperform simpler feature fusion strategies for integrating the motion prior into the action policy? (see Fig. 4)

4.1 Evaluation Setup

We evaluate our method on two standard simulation environments and three public benchmarks. **LIBERO** and **LIBERO-Plus** [12, 29] assesses language-conditioned manipulation, spatial reasoning, and long-horizon planning under object and layout variations. **SimplerEnv-WidowX** [25] serves as a high-fidelity proxy for real-world robustness, evaluating policies under realistic visual mismatches.

Training Recipe. Following our decoupled two-stage paradigm, the Motion Expert is pretrained on TraceForge-generated [22] dense 3D scene-flow supervision derived from a mixture of datasets. Subsequently, the Motion Expert and the Qwen3-VL-4B-Instruct backbone are frozen, and we train the Action Expert together with the motion-guidance module. Unless otherwise specified, LaMP employs the one-step partial denoising strategy for motion guidance during both training and inference. Our training pipeline is built upon the starVLA framework [9] and all experiments are conducted on 16 NVIDIA H100 GPUs. Detailed hyperparameters, such as the learning rate and batch size, are provided in the supplementary material. At inference, the Motion Expert adds only $1.35\times$ latency and $1.15\times$ memory while sustaining real-time control (see the supplementary material).

LIBERO Benchmark. We evaluate LaMP on all four standard suites: LIBERO-Spatial, LIBERO-Object, LIBERO-Goal, and LIBERO-Long, comprising 10 tasks

Table 1: Comparison of different methods on the LIBERO and SimplerEnv-WidowX benchmarks. We report the task success rate for each suite and the average. **Bold** denotes the best performance, and underline denotes the second best.

Method	LIBERO					SimplerEnv-WidowX						
	Spatial	Object	Goal	Long	Avg	Stack Block	Put Carrot	Put Spoon	Put Eggplant	Avg		
General VLA												
OpenVLA [20]	84.7	88.4	79.2	53.7	76.5	0.0	0.0	4.2	12.5	4.2		
OpenVLA-OFT [19]	97.6	98.4	<u>97.9</u>	94.5	97.1	—	—	—	—	—		
π_0 [2]	96.8	98.8	95.8	85.2	94.2	16.7	0.0	29.1	62.5	40.1		
$\pi_{0.5}$ [17]	<u>98.8</u>	98.2	98.0	92.4	96.9	44.7	64.7	49.3	69.7	57.1		
GR00T N1 [1]	94.4	97.6	93.0	90.6	93.9	16.7	45.8	62.5	20.8	49.5		
Latent-Action VLA												
UniVLA [4]	96.5	96.8	95.6	92.0	95.2	29.2	62.5	83.3	100.0	68.7		
villa-X [7]	97.5	97.0	91.5	74.5	90.1	61.3	46.3	77.9	64.6	62.5		
Video-Based VLA												
mimic-video [32]	94.2	96.8	90.6	—	93.9	29.2	54.2	41.7	100.0	56.3		
WorldVLA [5]	87.6	96.2	83.4	60.0	81.8	—	—	—	—	—		
$\mathcal{F}_1$ [30]	98.2	97.8	95.4	91.3	95.7	50.0	70.8	50.0	66.7	72.9		
2D Flow/Trace-Guided VLA												
FlowVLA [51]	93.2	95.0	91.6	72.6	88.1	<u>62.5</u>	62.5	70.8	100.0	<u>74.0</u>		
TraceVLA [48]	84.6	85.2	75.1	54.1	75.8	16.6	16.6	12.5	65.0	27.7		
3D-Aware VLA												
GeoVLA [37]	98.4	<u>99.0</u>	96.6	<u>96.6</u>	<u>97.7</u>	—	—	—	—	—		
SpatialVLA [36]	88.2	89.9	78.6	55.5	78.1	29.2	25.0	16.7	100.0	42.7		
LaMP												
w/o motion	99.4	99.8	97.4	96.7	98.3	75.0	<u>66.7</u>	<u>79.1</u>	<u>95.8</u>	79.2		
	95.8	98.9	96.6	78.2	92.4	25.0	45.8	66.7	87.5	56.3		

per suite with 500 expert demonstrations each, while training a single policy across all suites. We compare LaMP against a diverse set of VLA architectures: General VLAs (OpenVLA [20], OpenVLA-OFT [19], π_0 [2], $\pi_{0.5}$ [17], GR00T N1 [1]), latent-action VLAs (UniVLA [4], villa-X [7]), video-based VLAs (mimic-video [32], WorldVLA [5], $\mathcal{F}_1$ [30]), and 2D flow/trace-guided VLAs (FlowVLA [51], TraceVLA [48]) and report the task success rate for each suite and the average across all tasks. Each task is evaluated over 50 episodes.

As reported in Table 1, LaMP consistently outperforms strong baselines on LIBERO. In particular, LaMP achieves the best average success rate (98.3%) and the best long-horizon performance on LIBERO-Long (96.7%), exceeding the strongest prior baseline $\pi_{0.5}$ (96.9% average, 92.4% on Long). These results confirm that explicit 3D motion modeling reduces error accumulation over long manipulation horizons, directly addressing **Q1**.

LIBERO-Plus Benchmark (Zero-Shot OOD). All methods in Table 2 are trained exclusively on the original LIBERO demonstrations and evaluated zero-shot on the seven perturbation dimensions of LIBERO-Plus [12], without any additional training data. LaMP achieves the highest average success rate (79.3%), outperforming the strongest prior baseline OpenVLA-OFT (69.6%) by 9.7 percentage points. The gains concentrate on perturbations that alter the visual appearance or spatial layout of the scene: LaMP reaches 97.4% on Background and 95.3% on Light, surpassing all baselines by clear margins. On Robot perturbation, which changes the arm’s kinematic configuration, LaMP attains

Table 2: LIBERO-Plus zero-shot OOD evaluation. All models are trained on LIBERO and evaluated zero-shot on seven perturbation dimensions without additional training data.

Method	LIBERO-Plus							
	Camera	Robot	Language	Light	Background	Noise	Layout	Avg
UniVLA [4]	1.8	46.2	69.6	69.0	81.0	21.2	31.9	42.9
OpenVLA [20]	0.8	3.5	23	8.1	34.8	15.2	28.5	15.6
OpenVLA-OFT [19]	56.4	31.9	79.5	88.7	93.3	75.8	74.2	69.6
π_0 [2]	13.8	6.0	58.8	85.0	81.4	79.0	68.9	53.6
π_0 -Fast [34]	65.1	21.6	61.0	73.2	73.2	74.4	68.8	61.6
WorldVLA [5]	0.1	27.9	41.6	43.7	17.1	10.9	38.0	25.0
LaMP	64.5	69.6	88.2	95.3	97.4	76.9	73.8	79.3
w/o <i>motion</i>	46.7	56.0	82.5	95.3	95.4	69.3	71.0	71.6

69.6% while the next best method (π_0 -Fast) reaches only 21.6%. This robustness stems from our camera-frame motion prior: by reasoning in $(\Delta u, \Delta v, \Delta d)$ space rather than joint space, the learned dynamics transfer across visual and kinematic variations without task-specific adaptation. Ablating the Motion Expert (w/o *motion expert*) reduces the average from 79.3% to 71.6%, with the largest drops on Camera (-17.8) and Robot (-13.6), confirming that the motion prior is the primary source of OOD robustness rather than the VLM backbone alone, directly addressing **Q3**.

SimplerEnv-WidowX Benchmark. Following the SimplerEnv protocol [25], we evaluate policies trained on BridgeV2 real-world data without any fine-tuning on the simulated WidowX embodiment. LaMP achieves **79.2%** average success rate, outperforming the strongest baseline FlowVLA (74.0%) by **5.2%** and $\pi_{0.5}$ (57.1%) by **22.1%** (Table 1). Notably, on Stack Block, the most challenging task where all prior methods struggle, LaMP reaches **75.0%** while the second-best FlowVLA achieves only 62.5% and $\pi_{0.5}$ drops to 44.7%. LaMP also clearly surpasses recent 3D/depth-aware VLAs in this cross-domain setting, *e.g.*, SpatialVLA (42.7%) and RoboVLMs (38.0%) [24, 36].

This strong sim-to-sim transfer despite the visual domain gap stems from our 3D scene-flow prior. By reasoning in 3D camera coordinate space rather than pixel space, LaMP acquires a geometrically grounded representation that is robust to visual appearance changes between real and simulated environments, further addressing **Q3**. LaMP also generalizes to a bimanual embodiment, reaching 83.9/81.8 average success on 10 RoboTwin 2.0 tasks [6] under clean/randomized settings (see the supplementary material).

4.2 Ablation Studies

To systematically address our RQs and isolate the contributions of our motion prior and architectural choices, we evaluate four model variants on SimplerEnv-WidowX benchmark (see Fig. 3 and Fig. 4):

1. **No-Motion.** The Motion Expert is entirely removed, reducing the model to a purely reactive VLA without generative foresight.
2. **2D-Flow Prior.** The dense 3D scene flow prior is replaced with a 2D optical-flow prior by masking out the depth dimension during training while keeping u and v unchanged.
3. **Add.** The denoised motion hidden states are directly added to the corresponding action tokens at each layer, without any learnable gating.
4. **Concat+MLP.** The motion hidden states and action tokens are concatenated along the feature dimension and projected back via a two-layer MLP.

Effectiveness of Geometric Foresight.

Comparing LaMP with the No-Motion variant reveals severe performance degradation across both benchmarks. On LIBERO (See Tab. 1), removing the Motion Expert reduces the average success rate from 98.3% to 92.4% (−5.9%), with the largest drop on LIBERO-Long (96.7% → 78.2%, −18.5%). On SimplerEnv-WidowX (See Fig. 3), the degradation is even more pronounced: the average drops from 79.2% to 56.3% (−22.9%), with Stack Block plummeting by 50.0% (75.0% → 25.0%). This substantial gap confirms that explicit visual foresight is critical for long-horizon planning and bridging the real-to-sim domain gap. Without generating future geometric states, the policy struggles with goal alignment in complex multi-stage scenes and fails to transfer skills learned from BridgeV2 demonstrations to unseen visual configurations in simulation, directly addressing **Q1**.




Fig. 3: Ablation study on the SimplerEnv-WidowX benchmark. Success rates of LaMP are compared against variants without motion priors and with 2D flow priors across four manipulation tasks.

3D Geometry vs. 2D Pixels. Compared to the 2D-Flow variant, our 3D prior yields consistent gains across all four tasks. While 2D flow provides temporal cues, it lacks physical grounding for precise spatial manipulation. The 3D scene-flow prior improves Stack Block by 16.7% (58.3% → 75.0%) and Put Spoon by 8.3% (70.8% → 79.1%). Notably, on Put Eggplant both methods achieve comparable performance (91.7% vs. 95.8%), suggesting that tasks requiring less fine-grained 3D reasoning can be solved with 2D temporal cues alone. This result confirms that depth-aware geometric representations are indispensable for contact-rich manipulation tasks with stringent spatial constraints, directly addressing **Q2**.

Gated Fusion vs. Alternative Fusion Strategies. To investigate how the motion prior is best integrated into the action policy, we compare our gated cross-attention module against the Add and Concat+MLP variants described above on the SimplerEnv-WidowX benchmark.




Fig. 5: Real-world experiment platform and task overview. (a) **In-domain tasks:** Pick-and-Place, Deformable manipulation, and Long-horizon tasks. (b) **top: Hardware setup:** Flexiv Rizon 4 arm with Robotiq 2F-85 gripper and Intel D415 camera. **bottom: OOD test conditions:** unseen layout, object, and background. (c) **Visualizations of motion foresight.** The predicted motion trajectories are overlaid on the current observation. The color gradient from blue to red indicates the temporal progression of the predicted motion.

As shown in Fig. 4, gated cross-attention consistently outperforms both alternatives across all four tasks. The most striking gap appears on Stack Block, where the **Add** variant achieves only 33.3% compared to 75.0% for our gated module (-41.7%), and **Concat+MLP** reaches 45.8% (-29.2%). This collapse indicates that unweighted motion injection harms performance when the predicted flow conflicts with visual observations. On simpler tasks such as Put Eggplant, the performance gap narrows (**Add**: 91.7%, **Concat+MLP**: 83.3%, **Gated**: 95.8%), confirming that selective gating is most critical when spatial precision is required. These results show that the gated mechanism is essential for reliably exploiting the motion prior across tasks of varying difficulty, directly addressing Q4.




Fig. 4: Ablation study on fusion strategies. Success rates of Gated Cross-Attention are compared against Concat-MLP and Add variants across four manipulation tasks.

5 Real-World Experiments

5.1 Setup

Hardware Platform. We conduct real-world experiments using a 7-DoF Flexiv Rizon 4 robotic arm equipped with a 1-DoF Robotiq 2F-85 gripper. A single Intel

RealSense D415 RGB-D camera is mounted to provide global visual observations. The action space consists of a 6-DoF end-effector pose (relative to the current pose) and a 1-DoF gripper open/close state. All devices are connected to a workstation equipped with an NVIDIA RTX 4090 GPU for model inference and real-time control. Robot demonstrations are collected via haptic teleoperation.

5.2 Evaluation

In-Domain Tasks. We design three representative manipulation tasks spanning rigid-body manipulation, deformable-object folding, and long-horizon multi-step planning, as illustrated in Fig. 5 (left).

Pick-and-Place (Stack Cup) is a rigid-body 6-DoF pick-and-place task in which the robot grasps a blue cup and stacks it into a green bowl. The task proceeds in two stages: (I) approach the cup from the side and grasp it, and (II) approach the bowl and release the cup inside. Successful execution requires precise 6-DoF positioning to avoid collisions with the bowl rim, thereby evaluating the spatial reasoning capability of the policy for contact-aware trajectory planning.

Fold Towel requires the robot to fold a blue towel into a neat triangle through sequential corner manipulation. The task consists of four stages: (I) pick corner 1, (II) fold corner 1 to the opposite edge, (III) pick corner 2, and (IV) fold corner 2 to complete the triangle. The task is considered successful if the resulting shape closely approximates a triangle with well-aligned edges. This task evaluates the ability of the policy to handle deformable dynamics and sequential multi-step reasoning.

Making Bread is a long-horizon multi-step task that requires the robot to serve bread and pour an egg. The task involves three stages: (I) place the bread onto a plate, (II) grasp the handle of the pot and pour the egg onto the bread, and (III) return the pot to its original position. This task evaluates the capability of the policy for long-horizon planning and maintaining task coherence across multiple manipulation primitives.

Out-of-Distribution (OOD) Tasks. To evaluate generalization beyond the training distribution, we construct three conditions unseen in the expert demonstrations, as shown in Fig. 5 (middle).

Unseen Layout: positions and orientations in the workspace are significantly altered from those observed during training, testing the robustness of the policy to spatial configuration changes.

Unseen Object: Object instances are replaced with unseen counterparts that are not present in the training demonstrations. Specifically, for *Pick-and-Place*, the blue cup is replaced with a red cup, or the green bowl is replaced with a white bowl, testing generalization to unseen colors. For *Fold Towel*, evaluation is conducted using an unseen towel of a different color. For *Making Bread*, the plate used for serving is replaced with a bowl, which requires the policy to adapt to a different container geometry.

Unseen Background: The workspace surface is changed (*e.g.*, replacing the wooden tabletop with a green mat), introducing background variations not covered by the expert demonstrations and evaluating the visual robustness of the policy.

Baselines. We compare LaMP against the following baselines under identical data budgets: (1) π_0 [2], a mainstream flow-matching-based VLA that directly maps RGB observations to actions; and (2) 3D Flow Diffusion Policy (3D FDP) [31], a diffusion-based policy that operates on 3D flow representations for spatially-aware action generation. Since no official implementation is available, we re-implement 3D FDP based on the original paper.




Fig. 6: Real-world task evaluations. LaMP outperforms π_0 and 3D FDP across all task categories.

Training Protocol. We collect 160 expert demonstrations across three tasks on the target robot, totaling approximately one hour of teleoperation data. Following the warm-up protocol in [22], we fine-tune LaMP with 10 demonstrations per task to translate the embodiment-agnostic motion prior into the action space of the target robot, without requiring any robotic action data for pretraining. We then train a single unified policy on all three tasks using the full dataset, as does π_0 . Unlike these VLAs, 3D FDP requires training a separate model for each individual task. This one-policy multi-task design enables cross-task knowledge transfer while preserving the embodiment-agnostic geometric reasoning inherited from the pretrained motion prior.

Evaluation Protocol. For each in-domain task, we perform 10 trials per method with fixed initial layouts. For each OOD condition, we perform 5 trials per method per task. Unseen Layout tests randomized object positions and orientations, while Unseen Object and Unseen Background test training-distribution-matched layouts. We evaluate all methods under identical hardware and lighting conditions to ensure fair comparison. We report the average task progress score (see Fig. 5) as the primary evaluation metric. The metric definition is provided in the supplementary material.

5.3 Results

LaMP consistently outperforms both baselines across all task categories (Fig. 6). On Pick-and-Place, LaMP achieves **80%** average task progress, surpassing π_0 (70%) by 10 percentage points and 3D FDP (53%) by 27 percentage points. The gains are most pronounced on Deformable manipulation, where LaMP reaches

50% while both π_0 and 3D FDP plateau at 40%. This result confirms that the motion prior provides a stronger inductive bias for deformable dynamics.

On Long-Horizon tasks, LaMP maintains **80%** progress compared to 70% for π_0 and 60% for 3D FDP, demonstrating superior multi-step planning coherence. Under Out-of-Domain conditions, LaMP achieves **62.5%** while π_0 drops to 52.5% and 3D FDP collapses to 26.8%. This margin confirms that camera-frame geometric reasoning is more resilient to visual distribution shifts than pixel-level representations. Notably, 3D FDP fails under OOD conditions, dropping 26.2 points from its in-domain Pick-and-Place performance, whereas LaMP degrades by only 17.5 points. These results confirm that the embodiment-agnostic motion prior learned from diverse video data transfers effectively to unseen robots without requiring large-scale robotic demonstrations for pretraining.

5.4 Motion Foresight Visualizations

We visualize the Motion Expert’s predictive foresight to confirm that it captures physically meaningful dynamics rather than serving as opaque regularization. As shown in Fig. 5 (right), the predicted 3D trajectories anticipate the arm’s approach and manipulation *before* execution, spanning rigid grasping, deformable folding, and multi-step long-horizon tasks. Note that the Action Expert only conditions on 1-step partially denoised hidden states ($\tau = 0.1$); the full ODE rollout from $\tau = 0$ to $\tau = 1.0$ is performed solely for visualization.

6 Limitations and Future Work

LaMP demonstrates that dense 3D scene flow serves as an effective latent motion prior within pretrained VLA frameworks, with particular gains in long-horizon planning and cross-embodiment transfer. Several limitations remain. Our Motion Expert operates at a fixed spatial resolution and temporal horizon, so adaptive-resolution and longer-horizon motion representations are a natural extension. It also depends on TraceForge-generated 3D scene flow for supervision, and learning motion priors directly from unlabeled in-the-wild video would further improve scalability. Finally, motion is injected only at the last VLM layer to avoid the representational collapse observed with multi-layer fusion, and both a full multi-layer ablation and broader real-world, multi-embodiment evaluation are left for future work.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (No. 62506232); Shanghai Committee of Science and Technology (Yangfan: No. 24YF2722000; No.24511103200); Science and Technology Major Project of Jiangsu Province (No. BG2024041).

Lixin Yang (corresponding author) is affiliated with the School of Artificial Intelligence, Shanghai Jiao Tong University and serving as a project advisor at the Shanghai Innovation Institute.

References

1. Bjorck, J., Castaneda, F., Cherniadev, N., Da, X., Ding, R., Fan, L., Fang, Y., Fox, D., Hu, F., Huang, S., Jang, J., Jiang, Z., Kautz, J., Kundalia, K., Lao, L., Li, Z., Lin, Z., Lin, K., Liu, G., Llorente, E., Magne, L., Mandlekar, A., Narayan, A., Nasiriany, S., Reed, S., Tan, Y.L., Wang, G., Wang, Z., Wang, J., Wang, Q., Xiang, J., Xie, Y., Xu, Y., Xu, Z., Ye, S., Yu, Z., Zhang, A., Zhang, H., Zhao, Y., Zheng, R., Zhu, Y.: GR00T N1: An open foundation model for generalist humanoid robots. arXiv preprint arXiv:2503.14734 (2025)
2. Black, K., Brown, N., Driess, D., Esmail, A., Equi, M., Finn, C., Fusai, N., Groom, L., Hausman, K., Ichter, B., Jakubczak, S., Jones, T., Ke, L., Levine, S., Li-Bell, A., Mothukuri, M., Nair, S., Pertsch, K., Shi, L., Smith, L., Tanner, J., Vuong, Q., Walling, A., Wang, H., Zhilinsky, U.: π_0 : A vision-language-action flow model for general robot control. arXiv preprint arXiv:2410.24164 (2024)
3. Black, K., Nakamoto, M., Atreya, P., Walke, H., Finn, C., Kumar, A., Levine, S.: Zero-shot robotic manipulation with pretrained image-editing diffusion models. arXiv preprint arXiv:2310.10639 (2023)
4. Bu, Q., Yang, Y., Cai, J., Gao, S., Ren, G., Yao, M., Luo, P., Li, H.: Learning to act anywhere with task-centric latent actions. In: Proceedings of Robotics: Science and Systems (2025)
5. Cen, J., Yu, C., Yuan, H., Jiang, Y., Huang, S., Guo, J., Li, X., Song, Y., Luo, H., Wang, F., et al.: Worldvla: Towards autoregressive action world model. arXiv preprint arXiv:2506.21539 (2025)
6. Chen, T., Chen, Z., Chen, B., Cai, Z., Liu, Y., Li, Z., Liang, Q., Lin, X., Ge, Y., Gu, Z., et al.: RoboTwin 2.0: A scalable data generator and benchmark with strong domain randomization for robust bimanual robotic manipulation. arXiv preprint arXiv:2506.18088 (2025)
7. Chen, X., Wei, H., Zhang, P., Zhang, C., Wang, K., Guo, Y., Yang, R., Wang, Y., Xiao, X., Zhao, L., et al.: Villa-x: enhancing latent action modeling in vision-language-action models. arXiv preprint arXiv:2507.23682 (2025)
8. Chi, C., Xu, Z., Feng, S., Cousineau, E., Du, Y., Burchfiel, B., Tedrake, R., Song, S.: Diffusion policy: Visuomotor policy learning via action diffusion. The International Journal of Robotics Research **44**(10-11), 1684–1704 (2025)
9. starVLA Contributors: Starvla: A lego-like codebase for vision-language-action model developing (1 2025)
10. Du, Y., Yang, S., Dai, B., Dai, H., Nachum, O., Tenenbaum, J., Schuurmans, D., Abbeel, P.: Learning universal policies via text-guided video generation. Advances in neural information processing systems **36**, 9156–9172 (2023)
11. Fang, Y., Ranasinghe, K., Xue, L., Zhou, H., Tan, J., Xu, R., Heinecke, S., Xiong, C., Savarese, S., Szafir, D., et al.: Robotic vla benefits from joint learning with motion image diffusion. arXiv preprint arXiv:2512.18007 (2025)
12. Fei, S., Wang, S., Shi, J., Dai, Z., Cai, J., Qian, P., Ji, L., He, X., Zhang, S., Fei, Z., et al.: LIBERO-Plus: In-depth robustness analysis of vision-language-action models. arXiv preprint arXiv:2510.13626 (2025)
13. Ghosh, D., Walke, H.R., Pertsch, K., Black, K., Mees, O., Dasari, S., Hejna, J., Kreiman, T., Xu, C., Luo, J., Tan, Y.L., Chen, L.Y., Vuong, Q., Xiao, T., Sanketi, P.R., Sadigh, D., Finn, C., Levine, S.: Octo: An open-source generalist robot policy. In: Proceedings of Robotics: Science and Systems (2024)
14. Hu, Y., Guo, Y., Wang, P., Chen, X., Wang, Y.J., Zhang, J., Sreenath, K., Lu, C., Chen, J.: Video prediction policy: A generalist robot policy with predictive visual

- representations. In: International Conference on Machine Learning. pp. 24328–24346. PMLR (2025)
15. Huang, C.P., Wu, Y.H., Chen, M.H., Wang, Y.C.F., Yang, F.E.: Thinkact: Vision-language-action reasoning via reinforced visual latent planning. arXiv preprint arXiv:2507.16815 (2025)
 16. Huang, T., Li, D., Yang, R., Zhang, Z., Yang, Z., Tang, H.: MobileVLA-R1: Reinforcing vision-language-action for mobile robots. arXiv preprint arXiv:2511.17889 (2025)
 17. Intelligence, P., Black, K., Brown, N., Darpinian, J., Dhabalia, K., Driess, D., Esmail, A., Equi, M., Finn, C., Fusai, N., Galliker, M.Y., Ghosh, D., Groom, L., Hausman, K., Ichter, B., Jakubczak, S., Jones, T., Ke, L., LeBlanc, D., Levine, S., Li-Bell, A., Mothukuri, M., Nair, S., Pertsch, K., Ren, A.Z., Shi, L.X., Smith, L., Springenberg, J.T., Stachowicz, K., Tanner, J., Vuong, Q., Walke, H., Walling, A., Wang, H., Yu, L., Zhilinsky, U.: $\pi_{0.5}$: a vision-language-action model with open-world generalization. arXiv preprint arXiv:2504.16054 (2025)
 18. Khazatsky, A., Pertsch, K., Nair, S., Balakrishna, A., Dasari, S., Karamcheti, S., Nasiriany, S., Srirama, M.K., Chen, L.Y., Ellis, K., et al.: Droid: A large-scale in-the-wild robot manipulation dataset. arXiv preprint arXiv:2403.12945 (2024)
 19. Kim, M.J., Finn, C., Liang, P.: Fine-tuning vision-language-action models: Optimizing speed and success. arXiv preprint arXiv:2502.19645 (2025)
 20. Kim, M.J., Pertsch, K., Karamcheti, S., Xiao, T., Balakrishna, A., Nair, S., Rafailov, R., Foster, E.P., Sanketi, P.R., Vuong, Q., Kollar, T., Burchfiel, B., Tedrake, R., Sadigh, D., Levine, S., Liang, P., Finn, C.: OpenVLA: An open-source vision-language-action model. In: Proceedings of The 8th Conference on Robot Learning. pp. 2679–2713 (2025)
 21. Lee, J., Duan, J., Fang, H., Deng, Y., Liu, S., Li, B., Fang, B., Zhang, J., Wang, Y.R., Lee, S., et al.: MolmoAct: Action reasoning models that can reason in space. arXiv preprint arXiv:2508.07917 (2025)
 22. Lee, S., Jung, Y., Chun, I., Lee, Y.C., Cai, Z., Huang, H., Talreja, A., Dao, T.D., Liang, Y., Huang, J.B., et al.: TraceGen: World modeling in 3d trace space enables learning from cross-embodiment videos. arXiv preprint arXiv:2511.21690 (2025)
 23. Li, H., Sun, L., Hu, Y., Ta, D., Barry, J., Konidaris, G., Fu, J.: Novaflow: Zero-shot manipulation via actionable flow from generated videos. arXiv preprint arXiv:2510.08568 (2025)
 24. Li, X., Li, P., Qian, L., et al.: What matters in building vision-language-action models for generalist robots. *Nature Machine Intelligence* **8**, 158–172 (2026). <https://doi.org/10.1038/s42256-025-01168-7>
 25. Li, X., Hsu, K., Gu, J., Mees, O., Pertsch, K., Walke, H.R., Fu, C., Lunawat, I., Sieh, I., Kirmani, S., et al.: Evaluating real-world robot manipulation policies in simulation. In: Conference on Robot Learning. pp. 3705–3728. PMLR (2024)
 26. Li, Y., Chen, Y., Zhou, M., Li, H., Zhang, Z., Zhao, D.: Qdepth-vla: quantized depth prediction as auxiliary supervision for vision-language-action models. In: Conference on Autonomous Agents and Multiagent Systems (2026)
 27. Lin, B., Xu, Y., Zhan, X., Fang, H., Tian, J., Zhang, F.C., Li, Y.L., Lu, C., Yang, L.: ChronoFlow-Policy: Unifying past-current-future interaction flow in visuomotor policy learning. In: European Conference on Computer Vision (ECCV) (2026)
 28. Lipman, Y., Chen, R.T., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. In: The Eleventh International Conference on Learning Representations (2023)

29. Liu, B., Zhu, Y., Gao, C., Feng, Y., Liu, Q., Zhu, Y., Stone, P.: LIBERO: Benchmarking knowledge transfer for lifelong robot learning. *Advances in Neural Information Processing Systems* **36**, 44776–44791 (2023)
30. Lv, Q., Kong, W., Li, H., Zeng, J., Qiu, Z., Qu, D., Song, H., Chen, Q., Deng, X., Pang, J.: F1: A vision-language-action model bridging understanding and generation to actions. *arXiv preprint arXiv:2509.06951* (2025)
31. Noh, S., Nam, D., Kim, K., Lee, G., Yu, Y., Kang, R., Lee, K.: 3d flow diffusion policy: Visuomotor policy learning via generating flow in 3d space. *arXiv preprint arXiv:2509.18676* (2025)
32. Pai, J., Achenbach, L., Montesinos, V., Forrai, B., Mees, O., Nava, E.: Mimic-Video: Video-action models for generalizable robot control beyond VLAs. *arXiv preprint arXiv:2512.15692* (2025)
33. Patratskiy, M.A., Kovalev, A.K., Panov, A.I.: Spatial traces: Enhancing vla models with spatial-temporal understanding. *Optical Memory and Neural Networks* **34**(Suppl 1), S72–S82 (2025)
34. Pertsch, K., Stachowicz, K., Ichter, B., Driess, D., Nair, S., Vuong, Q., Mees, O., Finn, C., Levine, S.: Fast: Efficient action tokenization for vision-language-action models. *arXiv preprint arXiv:2501.09747* (2025)
35. Qu, D., Song, H., Chen, Q., Chen, Z., Gao, X., Ye, X., Lv, Q., Shi, M., Ren, G., Ruan, C., et al.: Eo-1: Interleaved vision-text-action pretraining for general robot control. *arXiv preprint arXiv:2508.21112* (2025)
36. Qu, D., Song, H., Chen, Q., Yao, Y., Ye, X., Ding, Y., Wang, Z., Gu, J., Zhao, B., Wang, D., et al.: SpatialVLA: Exploring spatial representations for visual-language-action model. *arXiv preprint arXiv:2501.15830* (2025)
37. Sun, L., Xie, B., Liu, Y., Shi, H., Wang, T., Cao, J.: GeoVLA: Empowering 3D representations in vision-language-action models. *arXiv preprint arXiv:2508.09071* (2025)
38. Tian, Y., Yang, Y., Xie, Y., Cai, Z., Shi, X., Gao, N., Liu, H., Jiang, X., Qiu, Z., Yuan, F., et al.: Interndata-a1: Pioneering high-fidelity synthetic data for pre-training generalist policy. *arXiv preprint arXiv:2511.16651* (2025)
39. Walke, H.R., Black, K., Zhao, T.Z., Vuong, Q., Zheng, C., Hansen-Estruch, P., He, A.W., Myers, V., Kim, M.J., Du, M., et al.: Bridgedata v2: A dataset for robot learning at scale. In: *Conference on Robot Learning*. pp. 1723–1736 (2023)
40. Wang, C., Fang, H., Fang, H.S., Lu, C.: RISE: 3D perception makes real-world robot imitation simple and effective. In: *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems*. pp. 2870–2877 (2024)
41. Xu, M., Xu, Z., Xu, Y., Chi, C., Wetzstein, G., Veloso, M., Song, S.: Flow as the cross-domain manipulation interface. In: *Proceedings of The 8th Conference on Robot Learning*. pp. 2475–2499 (2024)
42. Yang, Z., Teng, J., Zheng, W., Ding, M., Huang, S., Xu, J., Yang, Y., Hong, W., Zhang, X., Feng, G., et al.: Cogvideox: Text-to-video diffusion models with an expert transformer. In: *The Thirteenth International Conference on Learning Representations* (2025)
43. Yin, S., Yin, K., Chen, W., Liu, Y., Li, G., Lin, L.: Ddp-wm: Disentangled dynamics prediction for efficient world models. *arXiv preprint arXiv:2602.01780* (2026)
44. Zawalski, M., Chen, W., Pertsch, K., Mees, O., Finn, C., Levine, S.: Robotic control via embodied chain-of-thought reasoning. In: *Proceedings of The 8th Conference on Robot Learning*. pp. 3157–3181 (2024)
45. Ze, Y., Zhang, G., Zhang, K., Hu, C., Wang, M., Xu, H.: 3D diffusion policy: Generalizable visuomotor policy learning via simple 3D representations. In: *Proceedings of Robotics: Science and Systems* (2024)

46. Zhang, W., Liu, H., Qi, Z., Wang, Y., Yu, X., Zhang, J., Dong, R., He, J., Wang, H., Zhang, Z., et al.: Dreamvla: A vision-language-action model dreamed with comprehensive world knowledge. In: The Thirty-ninth Annual Conference on Neural Information Processing Systems (2025)
47. Zhen, H., Qiu, X., Chen, P., Yang, J., Yan, X., Du, Y., Hong, Y., Gan, C.: 3d-vla: A 3d vision-language-action generative world model. In: International Conference on Machine Learning. pp. 61229–61245. PMLR (2024)
48. Zheng, R., Liang, Y., Huang, S., Gao, J., Daumé III, H., Kolobov, A., Huang, F., Yang, J.: TraceVLA: Visual trace prompting enhances spatial-temporal awareness for generalist robotic policies. arXiv preprint arXiv:2412.10345 (2024)
49. Zheng, R., Wang, J., Reed, S., Bjorck, J., Fang, Y., Hu, F., Jang, J., Kundalia, K., Lin, Z., Magne, L., Narayan, A., Tan, Y.L., Wang, G., Wang, Q., Xiang, J., Xu, Y., Ye, S., Kautz, J., Huang, F., Zhu, Y., Fan, L.: Flare: Robot learning with implicit world modeling. In: Lim, J., Song, S., Park, H.W. (eds.) Proceedings of The 9th Conference on Robot Learning. Proceedings of Machine Learning Research, vol. 305, pp. 3952–3971. PMLR (27–30 Sep 2025)
50. Zhi, H., Chen, P., Zhou, S., Dong, Y., Wu, Q., Han, L., Tan, M.: 3DFlowAction: Learning cross-embodiment manipulation from 3d flow world model. arXiv preprint arXiv:2506.06199 (2025)
51. Zhong, Z., Yan, H., Li, J., Liu, X., Gong, X., Zhang, T., Song, W., Chen, J., Zheng, X., Wang, H., et al.: Flowvla: Visual chain of thought-based motion reasoning for vision-language-action models. arXiv preprint arXiv:2508.18269 (2025)