

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 GeoMoLa: GEOMETRY-AWARE MOTION LATENTS FOR LEARNING ROBUST MANIPULATION POLICIES

Anonymous authors

Paper under double-blind review

## ABSTRACT

Learning motion latents for robotic manipulation heavily relies on extracting motion patterns from visual sequences, yet effective action abstractions require understanding three-dimensional geometric transformations. Here, we introduce *GeoMoLa* (Geometry-Aware Motion Latents), which learns discrete motion latent codes by predicting how point clouds evolve during manipulation rather than reconstructing visual observations. This four-dimensional objective – spatial geometry changing through time – forces latent representations to encode actual physical motion rather than appearance patterns. *GeoMoLa* achieves state-of-the-art performance using only single-view RGB-D input, while existing methods require multi-view reconstruction, succeeding across diverse manipulation benchmarks. Our ablations reveal that geometric prediction is the key to driving performance, quantitatively validating that manipulation depends on spatial understanding. Furthermore, the learned codes exhibit effective motion abstraction: applying them to novel scenes produces physically consistent transformations regardless of visual context. Our real-world experiments also confirm this robustness capability, achieving robust manipulation with minimal demonstrations in cluttered environments where geometric reasoning determines success. Thus, we demonstrate that effective motion latents for robot control can better emerge from understanding motion through its three-dimensional effects rather than pixel-level patterns.

## 1 INTRODUCTION

Robot manipulation requires learning reusable motion patterns – motion latents (Bruce et al., 2024; Parker-Holder et al., 2024; Ball et al., 2025) – that abstract complex continuous movements into discrete, transferable skills. Current methods mainly learn these motion latents from sequences of two-dimensional images, missing the three-dimensional geometric structure that fundamentally determines manipulation success. A grasping action, for instance, depends not only on visual appearance but also on precise spatial relationships, approach angles, and the continuous evolution of three-dimensional configurations over time.

Therefore, learning motion latents without access to this underlying spatiotemporal geometry may produce representations that fail to generalize across different viewpoints, object poses, or spatial arrangements. Furthermore, this representational gap could create cascading failures in real-world deployment. Robots may not recognize that occluded objects maintain their geometric relationships despite visual changes, and small spatial errors compound across action sequences without understanding of three-dimensional workspace dynamics. Solving this essential representation issue in motion latent learning is necessary for robots to understand manipulation through spatial relationships and physical transformations rather than pixel patterns.

The *core* challenge lies in jointly modeling spatial geometry and temporal dynamics without prohibitive computational cost. Existing methods capture either spatial structure through static three-dimensional representations (Ke et al., 2024; Ze et al., 2024) or temporal patterns through two-dimensional video (Ye et al., 2024; Chen et al., 2024), but not both. Three-dimensional approaches process frozen point clouds without modeling evolution; diffusion policies generate trajectories from fixed scene features; video-based latent learning operates in image space without depth.

Our *key* insight is that effective motion latents must encode geometric transformations in three-dimensional space over time, not static scenes or visual motion. In this paper, we propose *GeoMoLa*

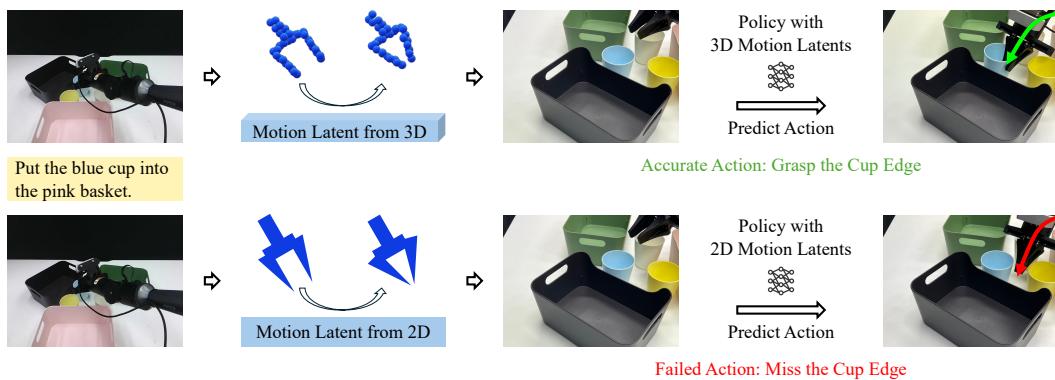


Figure 1: **Policies trained with 2D and 3D motion latents.** When encountering novel and cluttered scenes (left) that differ from the training distribution, the policy trained with 3D motion latents (top) demonstrates superior task performance with more robust control, e.g., enhanced reaching and grasping accuracy.

(Geometry-Aware Motion Latents), which learns **3D discrete motion latent codes (i.e., motion latent)** by predicting how **3D point clouds or point maps evolve during manipulation**. By training latent representations to forecast future geometric states rather than reconstruct current observations, we ensure these codes capture the causal relationship between actions and their spatial effects, creating motion primitives grounded in physical transformation rather than visual appearance.

Our approach demonstrates that explicitly modeling four-dimensional dynamics through geometry-aware motion latents significantly improves manipulation performance. We achieve state-of-the-art results using only single-view RGB-D input, while our learned latent codes produce consistent geometric transformations across different scenes, validating their transferability. Ablation studies reveal that three-dimensional geometric prediction is critical while visual appearance contributes minimally, suggesting rethinking representation learning for manipulation. In real-world experiments with limited demonstrations, our method excels particularly in cluttered environments where spatial reasoning determines success. While focused on rigid-body manipulation, our framework establishes four-dimensional geometric prediction as an effective objective for motion latent learning for robotics.

In summary, we make the following contributions:

- We introduce the first framework that explicitly models robot manipulation as continuous four-dimensional processes while learning **geometry-aware motion latents** through self-supervised prediction of future three-dimensional point cloud evolution.
- We demonstrate through comprehensive ablations that geometric structure in the motion latents is significantly more important than visual appearance for manipulation success.
- We achieve state-of-the-art performance across multiple benchmarks while showing that learned motion latents transfer consistently across different scenes and configurations.
- We show superior real-world performance with limited demonstrations, particularly excelling in cluttered scenarios requiring precise spatial reasoning.

## 2 RELATED WORK

**Motion latents.** Learning motion latent representations has proven effective across diverse applications. GENIE (Bruce et al., 2024; Parker-Holder et al., 2024; Ball et al., 2025) maps user inputs to latent spaces for generating interactive environments, while ILPO (Edwards et al., 2019) employs motion latents for pretraining video game policies. Recent work has explored deriving motion latents directly from observations: LAPA (Ye et al., 2024) and Moto (Chen et al., 2024) extract motion latents from raw inputs to leverage unlabeled data at scale. However, these observation-based approaches overlook the inherent 4D spatiotemporal structure of robotic actions. In contrast, we derive geometry-aware motion latents from 4D trajectories that explicitly capture spatial deformations and temporal dynamics, enabling more robust transfer to robotic manipulation tasks.

**Diffusion models in robotics.** Diffusion models have emerged as powerful tools for robotic manipulation, particularly for trajectory generation and action prediction. ChainedDiffuser (Xian et al., 2023) replaces traditional motion planners with a trajectory diffusion model that conditions on 3D scene

108 features and predicted keyposes from Act3D (Gervet et al., 2023) to generate linking trajectories.  
 109 Building on this, 3D Diffuser Actor (Ke et al., 2024) tackles the more challenging task of jointly  
 110 predicting the next keyposes and linking trajectories, while 3D Diffusion Policy (Ze et al., 2024)  
 111 combines 3D representations with diffusion objectives. We also evaluate against recent diffusion-  
 112 based methods (Bu et al., 2024; Black et al., 2023) and demonstrate superior performance. Unlike  
 113 these approaches that primarily focus on static 3D representations, our method explicitly models 4D  
 114 dynamics through motion latents, enabling better temporal reasoning and generalization.

115 **2D and 3D scene representations for robot manipulation.** 3D scene-to-action policies address  
 116 this limitation through explicit geometric reasoning: C2F-ARM (James et al., 2021) and PerAct  
 117 (Shridhar et al., 2022) voxelize workspaces but face computational scaling challenges; Act3D (Gervet  
 118 et al., 2023) avoids voxelization by sampling and featurizing 3D points through cross-attention; and  
 119 RVT (Goyal et al., 2023) reprojects RGB-D inputs to multiple views before lifting predictions to  
 120 3D. End-to-end image-to-action models like RT-1 (Brohan et al., 2022), RT-2 (Brohan et al., 2023),  
 121 GATO (Reed et al., 2022), BC-Z (Jang et al., 2022), RT-X (Padalkar et al., 2023), Octo (Team et al.,  
 122 2024), and InstructRL (Liu et al., 2022) directly predict 6-DoF poses from 2D images but require  
 123 thousands of demonstrations to implicitly learn 3D geometry. While these methods improve upon  
 124 2D approaches through explicit 3D representations, they still treat manipulation as static spatial  
 125 reasoning. Our approach advances beyond static 3D by modeling actions as continuous 4D processes,  
 126 capturing how spatial configurations evolve over time.

### 127 3 METHOD

128 **Overview.** We aim to learn robotic manipulation policies that map RGB-D observations and task  
 129 instructions to executable actions. Our *key* contribution is learning discrete motion latents – *abstract*  
 130 *motion concepts that guide high-level policy planning* – from 4D spatiotemporal data (3D pointmaps  
 131 over time) rather than 2D video sequences. Motion latents are essential for abstracting reusable  
 132 motion patterns and enabling task generalization, yet most existing methods learn motion latents  
 133 directly from 2D observations, missing crucial 3D geometric information (depth relationships, spatial  
 134 arrangements, and object poses) that fundamentally determine manipulation feasibility.

135 Our geometry-aware motion latents address this limitation by encoding abstract motion concepts  
 136 grounded in 3D geometry. This approach provides three advantages: better generalization across  
 137 manipulation tasks through abstract motion primitives, improved interpretability by encoding 3D  
 138 geometric transformations, and enhanced performance in cluttered environments requiring precise  
 139 spatial reasoning. Our framework comprises two components: a self-supervised pipeline that  
 140 discovers geometry-aware motion latents by predicting future 3D observations from demonstrations,  
 141 and a diffusion-based model that leverages these 3D action latents to generate actions.

#### 142 3.1 PROBLEM FORMULATION

143 We consider a dataset of robotic manipulation demonstrations  $\mathcal{D} = \{(\{\mathbf{o}_i^t, \mathbf{a}_i^t\}_{t=1}^{T_i}, l_i)\}_{i=1}^N$ , where  
 144 each demonstration contains a sequence of observation-action pairs with a natural language instruction  
 145  $l_i$ . Each observation  $\mathbf{o}_i^t = (\mathbf{I}_i^t, \mathbf{D}_i^t)$  consists of an RGB image  $\mathbf{I}_i^t \in \mathbb{R}^{H \times W \times 3}$  and a depth map  
 146  $\mathbf{D}_i^t \in \mathbb{R}^{H \times W}$ . Each action  $\mathbf{a}_i^t = (\mathbf{p}_i^t, \mathbf{r}_i^t, g_i^t)$  specifies the end-effector position  $\mathbf{p}_i^t \in \mathbb{R}^3$ , rotation  
 147  $\mathbf{r}_i^t \in \mathbb{R}^6$ , and gripper state  $g_i^t \in \{0, 1\}$ . We employ the 6D rotation representation proposed by  
 148 Zhou et al. (2018) to circumvent the discontinuity issues inherent in quaternion representations. Our  
 149 objective is to learn a policy  $\pi$  that maps the observation history and task instruction to an action  
 150 chunk:  $\hat{\mathbf{a}}^{t:t+h-1} = \pi(\mathbf{o}^t, l)$ , where  $h$  is the prediction horizon. The challenge lies in capturing  
 151 both geometric manipulation constraints (e.g., collision avoidance, grasp stability) and semantic  
 152 language intent. We address this through a two-stage method: First, we learn a dictionary of geometry-  
 153 aware motion latents (i.e., motion latents) that capture reusable motion patterns from demonstrations  
 154 (Sec. 3.2). Second, we leverage these learned motion latents to guide a diffusion-based policy for  
 155 generating executable trajectories (Sec. 3.3). We now detail each component.

#### 156 3.2 LEARNING MOTION LATENTS

157 We propose a self-supervised framework that discovers motion patterns by learning a dictionary of  
 158 discrete 3D latent codes, referred to as **motion latents**, encoded from the current observation and  
 159 task instructions, and trained to predict future 3D observations. We predict the future 3D geometry  
 160 to encourage the motion latents to capture the underlying motions driving scene evolution. We use  
 161 discrete rather than continuous representations because they naturally cluster similar motions into  
 reusable primitives and provide interpretable, composable action abstractions.

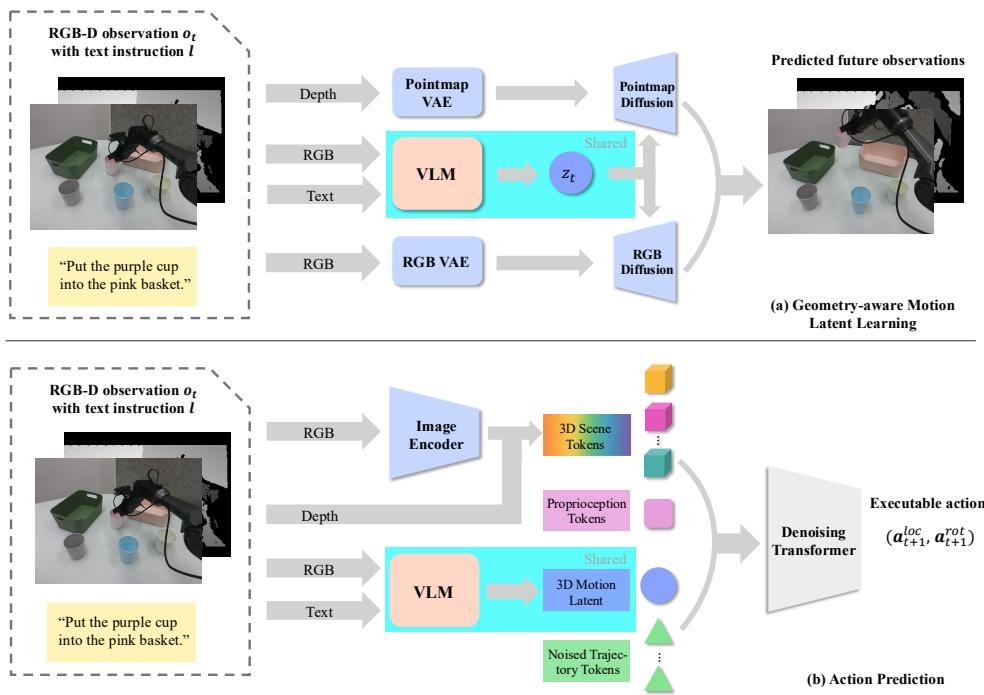


Figure 2: **GeoMoLa framework.** (a) **Geometry-Aware Motion Latent Learning:** RGB-D observations and language instructions are encoded into discrete motion latents via VQ-VAE, trained by predicting future pointmaps and RGB images. This self-supervised objective ensures latent codes capture 4D dynamics (3D geometry over time). (b) **Latent-Conditioned Action Prediction:** The previously trained motion latent encoder and the codebook are applied to guide a 3D denoising transformer to generate 6-DoF trajectories through iterative refinement, using 3D-aware attention mechanisms to accommodate geometric constraints.

### 3.2.1 ENCODING OBSERVATIONS INTO 3D MOTION LATENT

Given a current observation  $\mathbf{o}^t = (\mathbf{I}^t, \mathbf{D}^t)$  and task instruction  $l$ , we first extract a continuous representation using a vision-language encoder  $\phi^{\text{vlm}}$  based on Mini-GPT (Zhu et al., 2023), which provides strong visual-linguistic grounding:  $\mathbf{f}^t = \phi^{\text{vlm}}(\mathbf{o}^t, l)$ . We then discretize this representation using Vector Quantization (VQ-VAE) (van den Oord et al., 2017), where we map  $\mathbf{f}^t$  to a sequence of  $n_s$  discrete codes from a learned codebook  $\mathcal{C} = \{\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_K\}$  with vocabulary size  $K$ :

$$\mathbf{z}^t = \text{VQ}(\mathbf{f}^t) = [z_1^t, z_2^t, \dots, z_{n_s}^t], \quad (1)$$

where each  $z_j^t \in \{1, 2, \dots, K\}$  is an index into the codebook. We employ NSVQ (Vali & Bäckström, 2022) for stable training, addressing gradient collapse issues common in VQ-VAE.

### 3.2.2 LEARNING VIA FUTURE 3D PREDICTION

The core challenge in learning motion latents is ensuring they encode motion semantics rather than visual appearance. By requiring latent codes to predict future geometric transformations, we force them to capture the causal relationship between actions and their effects in 3D space. Specifically, we train the latent codes  $\mathbf{z}^t$  through a conditional prediction task: given the current 3D scene and a latent code, predict how the scene evolves over future timesteps. We use 3D pointmaps to ensure the codes learn geometric transformations rather than pixel-level appearance changes.

**Pointmap Representation.** We convert each RGB-D observation  $\mathbf{o}^t$  to a pointmap  $\mathbf{P}^t \in \mathbb{R}^{H \times W \times 3}$  by back-projecting pixels to 3D coordinates using camera intrinsics and extrinsics.

**Conditional Diffusion for Future Prediction.** We then employ a video latent diffusion architecture (Blattmann et al., 2023) adapted for pointmaps. The model consists of two components:

(1) A 3D-aware VAE with encoder  $\psi^{\text{enc}}$  and decoder  $\psi^{\text{dec}}$  that translates pointmaps to and from a latent space:  $\mathbf{h}_{\text{pm}}^{t'} = \psi^{\text{enc}}(\mathbf{P}^{t'})$ ,  $\hat{\mathbf{P}}^{t'} = \psi^{\text{dec}}(\mathbf{h}_{\text{pm}}^{t'})$ , where  $\mathbf{h}_{\text{pm}}^{t'} \in \mathbb{R}^d$  denotes the latent encoding at timestep  $t'$ . The VAE is initialized from a pre-trained RGB VAE (Blattmann et al., 2023) and fine-tuned on pointmaps through reconstruction. See Appendix D for details.

(2) A diffusion model  $\psi^{\text{diff}}$  that generates future latents conditioned on history and the motion latent  $\mathbf{z}^t$ , following DDPM (Ho et al., 2020):

$$\hat{\mathbf{h}}_{\text{pm}}^{t+1:t+w} = \psi^{\text{diff}}(\mathbf{h}_{\text{pm}}^{t-w+1:t}, \mathbf{z}^t; \phi), \quad (2)$$

where  $w$  is the observation window size. During training, we minimize the denoising objective:  $\mathcal{L}_{\text{diff}}^{\text{pm}} = \mathbb{E}_{k, \epsilon} [\|\epsilon - \epsilon_{\phi}(\mathbf{h}_k, k, \mathbf{h}_{\text{pm}}^{t-w+1:t}, \mathbf{z}^t)\|^2]$ , where  $\mathbf{h}_k = \sqrt{\alpha_k} \cdot \mathbf{h}_{\text{pm}}^{t+1:t+w} + \sqrt{1 - \alpha_k} \cdot \epsilon$  is the noised latent at diffusion step  $k$ ,  $\epsilon \sim \mathcal{N}(0, I)$ , and  $\alpha_k$  follows the noise schedule.

We jointly train an RGB prediction branch using the same architecture and motion latent  $\mathbf{z}^t$  to capture correlated appearance changes, with denoising objective  $\mathcal{L}_{\text{diff}}^{\text{rgb}}$ . The combined training objective is:  $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{diff}}^{\text{pm}} + \mathcal{L}_{\text{diff}}^{\text{rgb}} + \mathcal{L}_{\text{vq}}$ , where  $\mathcal{L}_{\text{vq}} = \|\text{sg}[\mathbf{f}^t] - \mathbf{c}\|_2^2 + \beta \|\mathbf{f}^t - \text{sg}[\mathbf{c}]\|_2^2$  is the VQ-VAE loss with stop-gradient operator  $\text{sg}[\cdot]$  and commitment coefficient  $\beta$ . Through this objective, latent codes encoding similar motions converge to similar discrete values, creating a learned vocabulary of reusable, geometry-aware motion primitives that capture the essential dynamics of state transitions.

### 3.3 ACTION PREDICTION WITH MOTION LATENTS

**Having learned motion latents through self-supervised future prediction.** We now describe how to use learned geometry-aware motion latents to generate executable robot trajectories. The motion latent  $\mathbf{z}^t$  serves as a bridge between high-level task understanding and low-level control, abstracting the essential motion while filtering out irrelevant visual details.

Our proposed action prediction model is a conditional diffusion policy that generates 6-DoF end-effector trajectories given: (1) the current scene observation, (2) the learned motion latent code, and (3) the robot’s proprioceptive state. By conditioning on geometry-aware motion latents rather than raw language instructions or pixels, the policy benefits from motion priors learned across the entire dataset, improving both sample efficiency and generalization.

#### 3.3.1 OVERALL PIPELINE

Specifically, we employ a 3D denoising transformer  $\epsilon_{\theta}$  based on Ke et al. (2024) to generate executable trajectories through conditional diffusion. The model iteratively refines noisy action sequences into precise robot motions. Given the current observation  $\mathbf{o}^t$ , motion latent  $\mathbf{z}^t$ , and proprioceptive state  $\mathbf{c}^t$ , the model generates an action chunk  $\mathbf{a}^{t:t+h-1}$  over a horizon  $h$  via DDPM, where each action  $\mathbf{a}^{t+k} = (\mathbf{p}^{t+k}, \mathbf{r}^{t+k}, \mathbf{g}^{t+k})$  specifies the end-effector position, rotation, and gripper state.

**Input Tokenization.** The transformer processes four types of tokens (illustrated in Fig. 2): (1) **Motion Latent Tokens**: The discrete motion latent  $\mathbf{z}^t = [z_1^t, \dots, z_{n_s}^t]$  from Sec. 3.2 is embedded via learned embeddings  $\mathbf{E}_z \in \mathbb{R}^{n_s \times d_{\text{code}}}$ , providing high-level motion guidance. (2) **Trajectory Tokens**: Each noisy action is encoded through an MLP to produce token  $\mathbf{t}_{\text{traj}}^k \in \mathbb{R}^{d_{\text{model}}}$ . We further concatenate the 3D position  $\mathbf{p}^{t+k}$  as positional information to maintain spatial grounding. (3) **Scene Tokens**: We extract visual features  $\mathbf{F} \in \mathbb{R}^{H \times W \times 3}$  using a frozen CLIP-ResNet50 (Radford et al., 2021) encoder. Each feature  $\mathbf{F}_{ij}$  at spatial location  $(i, j)$  is lifted to 3D position  $\mathbf{q}_{ij}$  using the depth map and camera intrinsic matrix. This produces  $H \times W$  scene tokens  $\{(\mathbf{F}_{ij}, \mathbf{q}_{ij})\}$  combining appearance and 3D position. (4) **Proprioception Token**: The robot state  $\mathbf{c}^t$  (joint angles, end-effector pose) is encoded as  $\mathbf{t}_{\text{prop}} = \text{MLP}(\mathbf{c}^t) + \text{PosEmbed}(\mathbf{p}_{\text{ee}}^t)$ , where  $\mathbf{p}_{\text{ee}}^t$  is the current end-effector position and PosEmbed denotes the positional embedding.

**Attention Mechanisms.** Then, the transformer employs a two-stage attention strategy to integrate spatial, temporal, and task information: (1) **Self-Attention with 3D Positional Encoding**: We first apply self-attention across all trajectory, scene, and proprioception tokens. To encode spatial relationships, we use the rotary positional embeddings (Su et al., 2021). (2) **Cross-Attention to Motion Latents**: After self-attention, we apply cross-attention from trajectory, scene and proprioception tokens to the motion latent embeddings. This mechanism allows the geometry-aware motion patterns

270 encoded in  $\mathbf{z}^t$  to guide trajectory generation. Unlike conditioning on raw language, the discrete motion  
 271 latents provide structured, geometrically consistent motion priors based on real-time observation.  
 272

273 Finally, we apply MLPs to predict the noise added to the sequence of 3D translations  $\epsilon_{\theta,p}$  and rotations  
 274  $\epsilon_{\theta,r}$ , as well as the gripper state  $\hat{g}$ , with the final trajectory tokens from the transformer outputs. This  
 275 progressively refines the action estimate based on scene geometry and learned motion patterns.  
 276

### 277 3.3.2 TRAINING OBJECTIVE AND INFERENCE PROCESS FOR ACTION PREDICTION

278 **To learn**, we train the denoising transformer  $\epsilon_\theta$  to predict the noise added to the ground-truth action  
 279 chunk. Given a clean action chunk  $\mathbf{a}^{t:t+h-1}$  from demonstrations, we sample noise  $\epsilon \sim \mathcal{N}(0, I)$  at dif-  
 280 fusion step  $i$ , then create the noisy action chunk:  $\mathbf{a}_i = \sqrt{\bar{\alpha}_i} \cdot \mathbf{a}^{t:t+h-1} + \sqrt{1 - \bar{\alpha}_i} \cdot \epsilon$ . The model learns  
 281 to predict  $\epsilon$  given the noisy action chunk and conditioning:  $\mathcal{L}_\theta = \mathbb{E}_{i,\epsilon} [\|\epsilon - \epsilon_\theta(\mathbf{a}_i, i, \mathbf{o}^t, \mathbf{z}^t, \mathbf{c}^t)\|]$ .  
 282 Since actions have components with different scales and properties, we use component-specific  
 283 losses:  
 284

$$\|\epsilon - \epsilon_\theta(\cdot)\| = \lambda_p \|\epsilon_p - \epsilon_{\theta,p}(\cdot)\|_1 + \lambda_r \|\epsilon_r - \epsilon_{\theta,r}(\cdot)\|_1 + \lambda_g \cdot \text{BCE}(g, \hat{g}) \quad (3)$$

285 where subscripts  $p, r, g$  denote position, rotation, and gripper components, respectively. We use L1  
 286 loss for continuous values (which is more robust to outliers in trajectory data) and binary cross-entropy  
 287 (BCE) for the discrete gripper state. The weights  $\{\lambda_p, \lambda_r, \lambda_g\}$  are determined with tuning.  
 288

289 **For inference**, we sample a noisy action chunk  $\mathbf{a}_N \sim \mathcal{N}(0, I)$  and iteratively denoise it following  
 290 DDPM (Ho et al., 2020) using the learned model and the current observation’s motion latent. The  
 291 final denoised action chunk contains executable 6-DoF poses and gripper commands that can be  
 292 directly sent to the robot controller. The geometry-aware motion latent  $\mathbf{z}^t$  ensures that the generated  
 293 trajectory respects both the task semantics and 3D spatial constraints learned from demonstrations.  
 294

## 295 4 EXPERIMENTS

296 We evaluate *GeoMoLa* through comprehensive experiments designed to answer three key questions:  
 297 (1) Does 4D spatiotemporal modeling improve manipulation performance compared to 2D/3D base-  
 298 lines? (2) Are the geometry-aware motion latents interpretable and transferable across tasks? (3) Does  
 299 our approach generalize effectively to real-world scenarios with occlusion and clutter? We conduct  
 300 experiments on two simulation benchmarks (RLBench (James et al., 2019) and CALVIN (Mees et al.,  
 301 2021)) and real-world manipulation tasks using the ALOHA robot (Fu et al., 2024), with extensive  
 302 ablations to validate our design choices.  
 303

### 304 4.1 RLBENCH EVALUATION

305 **Benchmark description.** RLBench is built on top of the CoppeliaSim simulator (Rohmer et al.,  
 306 2013), using a Franka Panda robot to interact with the environment. Our model and all baselines  
 307 are trained to predict the next end-effector keypose rather than the entire trajectory. To execute the  
 308 predicted keypose, we use RLBench’s built-in BiRRT motion planner to generate a feasible trajectory.  
 309 For evaluation, we select a suite of 10 challenging language-conditioned manipulation tasks, including  
 310 166 variations. These variations vary in several types, like position, shape, and color. We use the  
 311 front-view RGB-D camera as input to comply with practical deployment conditions. Performance is  
 312 measured by task completion success rate, defined as the proportion of execution trajectories that  
 313 satisfy the language-specified goal conditions. GNFactor (Ze et al., 2023) and ManiGaussian (Lu  
 314 et al., 2024) require an additional 19 views for 3D reconstruction during training. All other models  
 315 receive only the front-view observation as input during inference.  
 316

317 **Baselines.** We compare against three types of methods: (i) *3D representation-based policies*:  
 318 Act3D (Gervet et al., 2023) voxelizes the workspace and predicts 3D action maps; 3D Diffuser  
 319 Actor (Ke et al., 2024) uses 3D feature fields with diffusion-based trajectory generation; GNFactor (Ze  
 320 et al., 2023) leverages neural radiance fields (Mildenhall et al., 2020) for scene understanding.  
 321 RVT2 (Goyal et al., 2024) **reconstructs scene point cloud for better pose estimation**. (ii) *4D dynamic*  
 322 *framework*: ManiGaussian (Lu et al., 2024) employs dynamic 3D Gaussian Splatting (Kerbl et al.,  
 323 2023) for scene-level spatiotemporal dynamics. (iii) *motion latent methods*: SkillDiffuser (Liang  
 324 et al., 2023) learns discrete skills from 2D observations without explicit 3D geometric grounding.  
 325

324  
 325 **Table 1: Task success rates on RLBench.** *GeoMoLa* achieves the highest average performance  
 326 (80.1%) and ranks first on 9 out of 10 tasks. Results averaged over 5 seeds.

Method	Close jar	Open drawer	Sweep to dustpan	Turn tap	Meat off grill	Stack blocks	Slide block	Put in drawer	Drag stick	Push buttons	Avg.
GNFactor (Ze et al., 2023)	25.3	76.0	28.0	50.7	57.3	4.0	20.0	0.0	37.3	18.7	31.7
ManiGaussian (Lu et al., 2024)	28.0	76.0	64.0	56.0	60.0	<b>12.0</b>	24.0	16.0	92.0	20.0	44.8
Act3D (Gervet et al., 2023)	52.0	84.0	80.0	64.0	66.7	0.0	<b>100.0</b>	54.7	86.7	64.0	65.3
SkillDiffuser (Liang et al., 2023)	64.2	81.0	96.6	70.6	72.1	4.0	87.0	89.2	95.6	83.8	74.4
3D Diffuser Actor (Ke et al., 2024)	66.4	85.6	<b>98.4</b>	75.2	76.0	4.0	87.2	<b>94.4</b>	<b>98.4</b>	84.0	77.0
<b>RVT2</b> (Goyal et al., 2024)	<b>67.0</b>	<b>88.2</b>	<b>96.2</b>	<b>76.0</b>	<b>79.4</b>	<b>12.0</b>	<b>85.1</b>	<b>93.8</b>	<b>97.8</b>	<b>86.2</b>	<b>78.1</b>
<i>GeoMoLa</i> (Ours)	<b>69.4</b>	<b>85.9</b>	<b>98.4</b>	<b>81.4</b>	<b>79.0</b>	<b>12.0</b>	<b>90.5</b>	<b>95.0</b>	<b>98.2</b>	<b>92.1</b>	<b>80.1</b>

327

328

329

330

331

332

333

334

335

336

337

338

339

340

341

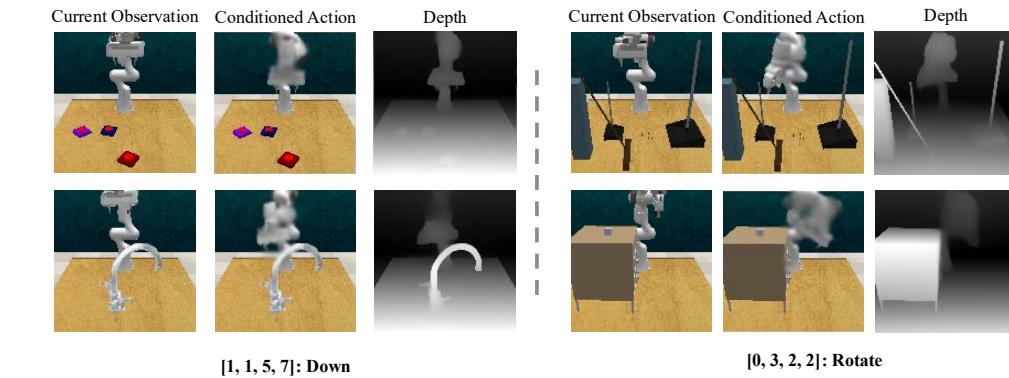
342

343

344

345

346



347 **Figure 3: Cross-scenario motion latent consistency.** Visualization of predicted future observations  
 348 when conditioning the same latent code on different initial scenes. Left: Latent code [1, 1, 5, 7]  
 349 consistently produces downward motion. Right: Latent code [0, 3, 2, 2] consistently generates  
 350 rotational motion. Consistent RGB and depth predictions demonstrate that geometry-aware motion  
 351 latents capture transferable 3D motion primitives independent of scene content (Please check the  
 352 synthesized end-effectors).

353  
 354 **Quantitative Results.** As shown in Tab. 1, our method achieves the highest overall performance  
 355 (80.1%), ranking first on 9 out of 10 tasks. Unlike prior works that condition their action prediction  
 356 models solely on static 3D feature fields, our approach explicitly models dynamics by learning motion  
 357 latents from state transitions. While SkillDiffuser also employs latent skill learning, our method  
 358 further leverages 3D geometric changes along trajectories, enabling more faithful representation of  
 359 scene structure and, consequently, more precise and reliable action generation.

360 **Interpretability of Geometry-Aware Motion Latents.** To validate that our learned motion latents  
 361 encode semantically meaningful motion primitives rather than task-specific behaviors, we conduct  
 362 cross-scenario generalization experiments. We extract latent codes from successful trajectories and  
 363 apply them to different scenes, measuring the consistency of the resulting motion patterns.

364 Fig. 3 demonstrates this cross-scenario transfer capability. When latent code [1, 1, 5, 7] is applied  
 365 to diverse initial configurations, it consistently generates downward motion in the predicted future  
 366 observations. Similarly, code [0, 3, 2, 2] reliably produces rotational motion regardless of scene  
 367 content. This semantic consistency validates that our 4D dynamics pretraining successfully distills  
 368 reusable motion primitives from the continuous space of robot actions.

369 **Future Observation Prediction Quality.** Beyond semantic consistency, accurate future prediction  
 370 is crucial for action planning. We evaluate the quality of predicted observations using both per-  
 371 ceptual and geometric metrics. Fig. 4 demonstrates our superior performance compared to another  
 372 reconstruction-based 4D dynamic method. Despite ManiGaussian’s access to 19 additional camera  
 373 views for the training of dynamic Gaussian Splatting, our pointmap approach produces more accurate  
 374 predictions with sharper object boundaries. The pointmap representation naturally preserves scene  
 375 structure during the diffusion process and is much simpler for learning the geometry information,  
 376 while 3D Gaussian deformation often introduces artifacts at occlusion boundaries due to data scarcity.  
 377 This prediction fidelity directly impacts action generation – accurate future state prediction enables  
 better trajectory planning, particularly for tasks involving precise object interactions.

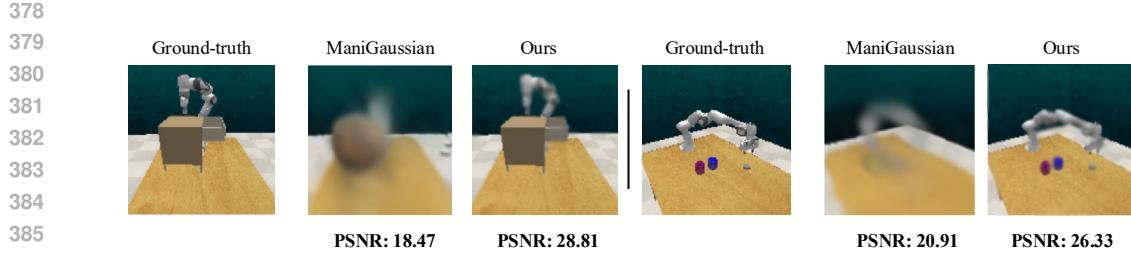


Figure 4: **Qualitative comparison of future observation prediction.** ManiGaussian requires 19 additional training views and produces blurred predictions with geometric inconsistencies (see distorted object boundaries). *GeoMoLa* generates sharper, geometrically consistent predictions using only single-view training. We report PSNR for quantitative comparison (higher is better).

Table 2: **Long-horizon task chaining on CALVIN.** Methods are evaluated on completing sequences of 1-5 tasks in the unseen environment D. *GeoMoLa* achieves the longest average task chain (3.58).

Method	Training Data	Task Completion Rate (%)					Avg. Length
		1 Task	2 Tasks	3 Tasks	4 Tasks	5 Tasks	
DP3 (Ze et al., 2024)	Language-annotated	28.7	2.7	0.0	0.0	0.0	0.31
GR-1 (Wu et al., 2023)	Language-annotated	85.4	71.2	59.6	49.7	40.1	3.06
SuSIE (Black et al., 2023)	All play data	87.0	69.0	49.0	38.0	26.0	2.69
RVT2 (Goyal et al., 2024)	Language-annotated	90.4	76.8	61.6	50.2	39.8	3.23
3D Diffuser Actor (Ke et al., 2024)	Language-annotated	93.8	80.3	66.2	53.3	41.2	3.35
Clover (Bu et al., 2024)	Language-annotated	96.0	83.5	70.8	57.5	45.4	3.53
<i>GeoMoLa</i> (Ours)	Language-annotated	95.4	84.0	72.4	59.0	46.4	3.58

## 4.2 CALVIN EVALUATION

**Benchmark Description.** CALVIN (Mees et al., 2021) evaluates long-horizon task execution in PyBullet, requiring sequential completion of language-specified sub-tasks. It includes 34 distinct tasks across 4 environments (A, B, C, D) with varying textures and object positions. Each environment contains a Franka Emika Panda robot, desk, sliding door, drawer, LED button, light switch, and colored blocks. CALVIN provides 24 hours of teleoperated unstructured play data, 35% of which are annotated with language descriptions (18k trajectory videos). Each instruction chain includes five language instructions that need to be executed sequentially. We evaluate on the challenging zero-shot generalization setting: training on environments A, B, C and testing on the unseen environment D.

**Baselines.** We compare against: (i) *3D-based methods*: DP3 (Ze et al., 2024) encodes RGB-D into 3D features for diffusion trajectory prediction (we add language conditioning following (Ke et al., 2024)); *3D Diffuser Actor* and *RVT2* that previously mentioned. (ii) *Video-pretrained models*: GR-1 (Wu et al., 2023) leverages large-scale internet video pretraining; SuSIE (Black et al., 2023) uses all available play data including unannotated sequences. (iii) *Closed-loop methods*: Clover (Bu et al., 2024) encodes RGB-D inputs, introducing a closed-loop error correction mechanism to achieve better video diffusion policy.

**Results Analysis.** Tab. 2 shows *GeoMoLa* achieves the highest average task sequence length of 3.58, completing 46.4% of 5-task chains compared to 41.2% for 3D Diffuser Actor. The performance gap widens with sequence length – while Clover’s closed-loop correction provides only a 0.6% advantage on completion of the first subtask, *GeoMoLa* outperforms other methods when completing the entire long-horizon subtask sequences. This trend validates our hypothesis: explicit 4D dynamics modeling reduces compounding errors that accumulate over long horizons. *While RVT2 utilizes RGB-D input, it obtains supervision only from ground-truth poses. Our results demonstrate that learning motion patterns enables more accurate action prediction without requiring additional views..* The stark failure of DP3 (0.31 average length) despite using 3D representations highlights that static geometric features alone are insufficient for sequential manipulation. GR-1’s internet-scale pretraining provides strong single-task performance (85.4%) but degrades rapidly in multi-task scenarios, suggesting that generic video understanding does not transfer directly to precise robotic control.

432

433

434

435

436

437

438

439

440

441

442

## 443 4.3 ABLATION STUDIES

444 To validate our architectural choices, we conduct ablation studies examining the contribution of  
445 each modality in the 4D dynamics learning phase. We design two baselines that pretrain the motion  
446 latent space using only a single modality – either 2D image prediction or pointmap prediction – by  
447 removing the pointmap or RGB prediction branch from our framework, respectively.448 **Importance of Geometric vs. Appearance Modeling.** Tab. 3 reveals a critical insight: removing  
449 pointmap prediction causes a substantial performance drop (CALVIN: -0.20 avg. length; RLBench:  
450 -2.1%), while removing RGB prediction has minimal impact (-0.08 avg. length; -0.7%). This asym-  
451 metry demonstrates that explicit 3D geometry is fundamental for learning transferable manipulation  
452 primitives, while appearance primarily provides auxiliary context.453 **The role of geometry in different motion types.** By analysing the task-specific results on RLBench,  
454 we find that in rotation-heavy tasks (e.g., “sweep to dustpan”, “open drawer”), removing point map  
455 prediction reduced success rates by nearly 9% on average, while in translation-heavy tasks (e.g.,  
456 “close jar”, “push buttons”), the drop was less than 1%. This suggests that 3D geometry is especially  
457 important for tasks involving fine rotational control and complex object interactions, aligning with our  
458 motivation that geometric awareness supports more precise gripper state estimation and interaction  
459 reasoning. Thus, the key value of our approach lies in providing consistent robustness in geometrically  
460 complex settings, rather than a uniform boost across all tasks.

## 461 4.4 REAL-WORLD VALIDATION

462 **Experimental Setup.** We collected real-world demonstration data using the ALOHA robot platform.  
463 RGB-D observations are captured via an Intel RealSense camera at 640×480 resolution from a front  
464 view and subsequently downsampled to 256×256 for processing. During inference, target gripper  
465 poses are executed using the MoveIt package in ROS (Coleman et al., 2014).466 **Data Collection and Training.** We consider six distinct tasks, with 20 demonstration trajectories  
467 recorded for each task. To ensure diversity within each task, variations in object quantities, positions,  
468 colors, and other attributes were intentionally introduced across different trajectories. The collected  
469 demonstration data were subsequently used to train our model. During the testing phase, the model  
470 was evaluated on the same set of tasks under zero-shot transfer conditions, where object configurations,  
471 spatial layouts, and visual properties such as color were systematically altered to assess generalization  
472 beyond the training demonstrations. We evaluated 10 episodes for each task and reported the success  
473 rate. Models are trained from scratch on this limited data to evaluate sample efficiency in real-world  
474 settings. Detailed task information and visualization of our settings could be found in Appendix. C.  
475476 **Table 4: Real-world manipulation performance.** Success rates over 10 trials per task. *GeoMoLa*  
477 shows consistent improvements, particularly in cluttered scenarios (Clean cup) and precise manipula-  
478 tion (Stack cubes). † indicates tasks with significant occlusion during execution.

Method	Clean cup†	Stack cups	Put cups on shelf	Stack cubes	Place dish	Place cube	Average
SkillDiffuser (Liang et al., 2023)	30.0%	0.0%	20.0%	20.0%	20.0%	40.0%	21.7%
3D Diffuser Actor (Ke et al., 2024)	50.0%	10.0%	40.0%	20.0%	40.0%	80.0%	40.0%
<i>GeoMoLa</i> w/o Pointmap	20.0%	10.0%	30.0%	20.0%	50.0%	80.0%	35.0%
<i>GeoMoLa</i> (Ours)	60.0%	30.0%	40.0%	50.0%	50.0%	90.0%	53.3%

479 **Results and Analysis.** Tab. 4 shows *GeoMoLa* achieving a 53.3% average success rate, a 13.3%  
480 improvement over 3D Diffuser Actor. In particular, we observe consistent gains in tasks such as  
481 “Clean cup”, which involves significant occlusion, as well as “Stack cups” and “Stack cubes”, which

486 require precise manipulation. These improvements highlight the benefits of our approach in modeling  
487 dynamics and scene geometry, enabling more accurate action prediction. More detailed visualization  
488 discussion can be found in Fig. 1 and the Appendix C. The results demonstrate the effectiveness of  
489 our method for real-world robotic manipulation in diverse and previously unseen settings.  
490

## 491 5 CONCLUSION 492

493 This work demonstrates that motion latent learning for robotic manipulation benefits significantly  
494 from grounding in four-dimensional geometric transformations rather than visual sequences. Our  
495 ablation studies provide quantitative evidence: removing geometric prediction degrades motion latent  
496 quality substantially while visual prediction contributes minimally to performance. The success  
497 with limited real-world demonstrations indicates that geometry-aware latent representations naturally  
498 capture manipulation-relevant motion primitives without requiring extensive datasets. Future work  
499 could extend this motion latent framework to deformable objects and investigate hierarchical planning  
500 where high-level policies compose learned geometric primitives. Our results suggest that effective  
501 motion latent representations for robotics should encode how objects move through three-dimensional  
502 space over time rather than how they appear visually, providing another perspective on learning  
503 reusable manipulation skills from unlabeled demonstrations.  
504  
505  
506  
507  
508  
509  
510  
511  
512  
513  
514  
515  
516  
517  
518  
519  
520  
521  
522  
523  
524  
525  
526  
527  
528  
529  
530  
531  
532  
533  
534  
535  
536  
537  
538  
539

540 **6 ETHICS STATEMENT**  
541

542 This work strictly adheres to established research ethics guidelines. Our research focuses on de-  
 543 veloping geometry-aware motion learning for robotic manipulation and does not involve human  
 544 subjects, animal experiments, or raise concerns related to privacy or security. All experiments were  
 545 conducted in simulation environments using synthetic data and self-collected demonstrations from  
 546 consenting researchers, with no personal data, biometric information, or sensitive content involved.  
 547 The geometry-aware motion learning framework is designed for research purposes in controlled  
 548 robotic manipulation tasks and presents minimal risks of misuse. We have carefully considered the  
 549 broader impacts of our work and believe it contributes positively to the advancement of reliable and  
 550 sample-efficient robotic systems. The improved generalization capabilities and reduced data require-  
 551 ments demonstrated by our method could benefit applications in assistive robotics, healthcare, elder  
 552 care, and accessibility when properly validated in real-world settings. We acknowledge that advances  
 553 in robotic automation may have economic implications for certain job categories. However, our work  
 554 focuses specifically on benign manipulation tasks such as grasping, stacking, and placing objects, with  
 555 the goal of enhancing human-robot collaboration rather than replacement. The framework is designed  
 556 to augment human capabilities in controlled environments rather than substitute for human judgment  
 557 and creativity. While any advancement in robotic capabilities could potentially be misapplied, our  
 558 research addresses fundamental manipulation learning without developing applications for harmful  
 559 purposes. All authors have thoroughly reviewed this work and acknowledge compliance with research  
 560 ethics standards.

561 **7 REPRODUCIBILITY STATEMENT**  
562

563 We have taken extensive measures to ensure the reproducibility of our work on the GeoMoLa frame-  
 564 work. The complete architectural details of our geometry-aware latent action learning pipeline and  
 565 diffusion-based action prediction model are described in Sec. 3 and Appendix D. All experimental  
 566 configurations, including task specifications for RLBench and CALVIN benchmarks, batch size, learn-  
 567 ing rate, training epochs, and diffusion model parameters are detailed in Sec. 4 and Appendix B, C, D.  
 568 The geometry-aware encoding formulation, latent action space construction, and diffusion-based  
 569 prediction procedure are precisely specified with corresponding equations and algorithms provided  
 570 in Sec. 3 and Appendix D. Our experimental evaluation builds upon established publicly available  
 571 benchmarks (RLBench and CALVIN) using standard evaluation protocols, with all success metrics  
 572 and evaluation criteria clearly defined for direct comparability with existing work. Implementation  
 573 details include the geometry encoder architecture, latent space dimensionality, diffusion model config-  
 574 uration, and specific hyperparameters for all baseline methods documented in Sec. 3 and Appendix D.  
 575 Our ablation studies examining the contributions of each component are systematically presented  
 576 with quantitative results. For real-world validation, we provide detailed collection procedures for our  
 577 demonstration dataset, including robot setup, task specifications, and data preprocessing steps. Upon  
 578 publication, we will make available the complete implementation of GeoMoLa, including trained  
 579 models, configuration files, example usage scripts, and our collected demonstration dataset with  
 580 comprehensive documentation to facilitate reproduction of all reported results and enable extension  
 581 of our work.

582 **REFERENCES**  
583

584 Philip J. Ball, Jakob Bauer, Frank Belletti, Bethanie Brownfield, Ariel Ephrat, Shlomi Fruchter,  
 585 Agrim Gupta, Kristian Holsheimer, Aleksander Holynski, Jiri Hron, Christos Kaplanis, Marjorie  
 586 Limont, Matt McGill, Yanko Oliveira, Jack Parker-Holder, Frank Perbet, Guy Scully, Jeremy Shar,  
 587 Stephen Spencer, Omer Tov, Ruben Villegas, Emma Wang, Jessica Yung, Cip Baetu, Jordi Berbel,  
 588 David Bridson, Jake Bruce, Gavin Buttimore, Sarah Chakera, Bilva Chandra, Paul Collins, Alex  
 589 Cullum, Bogdan Damoc, Vibha Dasagi, Maxime Gazeau, Charles Gbadamosi, Woohyun Han,  
 590 Ed Hirst, Ashyana Kachra, Lucie Kerley, Kristian Kjems, Eva Knoepfel, Vika Koriakin, Jessica  
 591 Lo, Cong Lu, Zeb Mehring, Alex Moufarek, Henna Nandwani, Valeria Oliveira, Fabio Pardo, Jane  
 592 Park, Andrew Pierson, Ben Poole, Helen Ran, Tim Salimans, Manuel Sanchez, Igor Saprykin,  
 593 Amy Shen, Sailesh Sidhwani, Duncan Smith, Joe Stanton, Hamish Tomlinson, Dimple Vijaykumar,  
 Luyu Wang, Piers Wingfield, Nat Wong, Keyang Xu, Christopher Yew, Nick Young, Vadim Zubov,  
 Douglas Eck, Dumitru Erhan, Koray Kavukcuoglu, Demis Hassabis, Zoubin Gharamani, Raia

594 Hadsell, Aäron van den Oord, Inbar Mosseri, Adrian Bolton, Satinder Singh, and Tim Rocktäschel.  
 595 Genie 3: A new frontier for world models. 2025.

596

597 Kevin Black, Mitsuhiro Nakamoto, Pranav Atreya, Homer Rich Walke, Chelsea Finn, Aviral Kumar,  
 598 and Sergey Levine. Zero-shot robotic manipulation with pretrained image-editing diffusion models.  
 599 *ArXiv*, abs/2310.10639, 2023. URL <https://api.semanticscholar.org/CorpusID:264172455>.

600

601 A. Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, and Do-  
 602 minik Lorenz. Stable video diffusion: Scaling latent video diffusion models to large datasets.  
 603 *ArXiv*, abs/2311.15127, 2023. URL <https://api.semanticscholar.org/CorpusID:265312551>.

604

605 Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn,  
 606 Keerthana Gopalakrishnan, Karol Hausman, Alexander Herzog, Jasmine Hsu, Julian Ibarz, Brian  
 607 Ichter, Alex Irpan, Tomas Jackson, Sally Jesmonth, Nikhil J. Joshi, Ryan C. Julian, Dmitry  
 608 Kalashnikov, Yuheng Kuang, Isabel Leal, Kuang-Huei Lee, Sergey Levine, Yao Lu, Utsav Malla,  
 609 Deeksha Manjunath, Igor Mordatch, Ofir Nachum, Carolina Parada, Jodilyn Peralta, Emily Perez,  
 610 Karl Pertsch, Jornell Quiambao, Kanishka Rao, Michael S. Ryoo, Grecia Salazar, Pannag R.  
 611 Sanketi, Kevin Sayed, Jaspiar Singh, Sumedh Anand Sontakke, Austin Stone, Clayton Tan, Huong  
 612 Tran, Vincent Vanhoucke, Steve Vega, Quan Ho Vuong, F. Xia, Ted Xiao, Peng Xu, Sichun Xu,  
 613 Tianhe Yu, and Brianna Zitkovich. Rt-1: Robotics transformer for real-world control at scale.  
 614 *ArXiv*, abs/2212.06817, 2022. URL <https://api.semanticscholar.org/CorpusID:254591260>.

615

616 Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Krzysztof Choromanski, Tianli  
 617 Ding, Danny Driess, Kumar Avinava Dubey, Chelsea Finn, Peter R. Florence, Chuyuan Fu,  
 618 Montse Gonzalez Arenas, Keerthana Gopalakrishnan, Kehang Han, Karol Hausman, Alexander  
 619 Herzog, Jasmine Hsu, Brian Ichter, Alex Irpan, Nikhil J. Joshi, Ryan C. Julian, Dmitry Kalashnikov,  
 620 Yuheng Kuang, Isabel Leal, Sergey Levine, Henryk Michalewski, Igor Mordatch, Karl Pertsch, Kan-  
 621 ishka Rao, Krista Reymann, Michael S. Ryoo, Grecia Salazar, Pannag R. Sanketi, Pierre Sermanet,  
 622 Jaspiar Singh, Anikait Singh, Radu Soricu, Huong Tran, Vincent Vanhoucke, Quan Ho Vuong,  
 623 Ayzaan Wahid, Stefan Welker, Paul Wohlhart, Ted Xiao, Tianhe Yu, and Brianna Zitkovich. Rt-2:  
 624 Vision-language-action models transfer web knowledge to robotic control. *ArXiv*, abs/2307.15818,  
 625 2023. URL <https://api.semanticscholar.org/CorpusID:260293142>.

626

627 Jake Bruce, Michael D Dennis, Ashley Edwards, Jack Parker-Holder, Yuge Shi, Edward Hughes,  
 628 Matthew Lai, Aditi Mavalankar, Richie Steigerwald, Chris Apps, et al. Genie: Generative  
 629 interactive environments. In *Forty-first International Conference on Machine Learning*, 2024.

630

631 Qingwen Bu, Jia Zeng, Li Chen, Yanchao Yang, Guyue Zhou, Junchi Yan, Ping Luo, Heming Cui,  
 632 Yi Ma, and Hongyang Li. Closed-loop visuomotor control with generative expectation for robotic  
 633 manipulation. *ArXiv*, abs/2409.09016, 2024. URL <https://api.semanticscholar.org/CorpusID:272653959>.

634

635 Yi Chen, Yuying Ge, Weiliang Tang, Yizhuo Li, Yixiao Ge, Mingyu Ding, Ying Shan, and Xihui Liu.  
 Moto: Latent motion token as the bridging language for learning robot manipulation from videos.  
 636 2024. URL <https://api.semanticscholar.org/CorpusID:277151378>.

637

638 D.M. Coleman, Ioan Alexandru Sucan, Sachin Chitta, and Nikolaus Correll. Reducing the barrier  
 639 to entry of complex robotic software: a moveit! case study. *ArXiv*, abs/1404.3785, 2014. URL  
<https://api.semanticscholar.org/CorpusID:13939653>.

640

641 Ashley Edwards, Himanshu Sahni, Yannick Schroeder, and Charles Isbell. Imitating latent policies  
 642 from observation. In *International conference on machine learning*, pp. 1755–1763. PMLR, 2019.

643

644 Zipeng Fu, Tony Zhao, and Chelsea Finn. Mobile aloha: Learning bimanual mobile manipulation  
 645 with low-cost whole-body teleoperation. *ArXiv*, abs/2401.02117, 2024. URL <https://api.semanticscholar.org/CorpusID:266755740>.

646

647 Théophile Gervet, Zhou Xian, Nikolaos Gkanatsios, and Katerina Fragkiadaki. Act3d: 3d feature  
 648 field transformers for multi-task robotic manipulation. In *Conference on Robot Learning*, 2023.  
 649 URL <https://api.semanticscholar.org/CorpusID:259308821>.

648 Ankit Goyal, Jie Xu, Yijie Guo, Valts Blukis, Yu-Wei Chao, and Dieter Fox. Rvt: Robotic view  
 649 transformer for 3d object manipulation. *ArXiv*, abs/2306.14896, 2023. URL <https://api.semanticscholar.org/CorpusID:259262273>.

650

651 Ankit Goyal, Valts Blukis, Jie Xu, Yijie Guo, Yu-Wei Chao, and Dieter Fox. Rvt2: Learning precise  
 652 manipulation from few demonstrations. *RSS*, 2024.

653

654 Jonathan Ho, Ajay Jain, and P. Abbeel. Denoising diffusion probabilistic models. *ArXiv*,  
 655 abs/2006.11239, 2020. URL <https://api.semanticscholar.org/CorpusID:219955663>.

656

657 Stephen James, Z. Ma, David Rovick Arrojo, and Andrew J. Davison. Rlbench: The robot learning  
 658 benchmark & learning environment. *IEEE Robotics and Automation Letters*, 5:3019–3026, 2019.  
 659 URL <https://api.semanticscholar.org/CorpusID:202889132>.

660

661 Stephen James, Kentaro Wada, Tristan Laidlow, and Andrew J. Davison. Coarse-to-fine q-attention:  
 662 Efficient learning for visual robotic manipulation via discretisation. *2022 IEEE/CVF Conference  
 663 on Computer Vision and Pattern Recognition (CVPR)*, pp. 13729–13738, 2021. URL <https://api.semanticscholar.org/CorpusID:235606348>.

664

665 Eric Jang, Alex Irpan, Mohi Khansari, Daniel Kappler, Frederik Ebert, Corey Lynch, Sergey  
 666 Levine, and Chelsea Finn. Bc-z: Zero-shot task generalization with robotic imitation learning.  
 667 *ArXiv*, abs/2202.02005, 2022. URL <https://api.semanticscholar.org/CorpusID:237257594>.

668

669 Tsung-Wei Ke, Nikolaos Gkanatsios, and Katerina Fragkiadaki. 3d diffuser actor: Policy dif-  
 670 fusion with 3d scene representations. *ArXiv*, abs/2402.10885, 2024. URL <https://api.semanticscholar.org/CorpusID:267740176>.

671

672 Bernhard Kerbl, Georgios Kopanas, Thomas Leimkuehler, and George Drettakis. 3d gaussian  
 673 splatting for real-time radiance field rendering. *ACM Transactions on Graphics (TOG)*, 42:1 – 14,  
 674 2023. URL <https://api.semanticscholar.org/CorpusID:259267917>.

675

676 Zhixuan Liang, Yao Mu, Hengbo Ma, Masayoshi Tomizuka, Mingyu Ding, and Ping Luo. Skilldif-  
 677 fuser: Interpretable hierarchical planning via skill abstractions in diffusion-based task execution.  
 678 *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 16467–  
 679 16476, 2023. URL <https://api.semanticscholar.org/CorpusID:266361979>.

680

681 Hao Liu, Lisa Lee, Kimin Lee, and P. Abbeel. Instruction-following agents with jointly  
 682 pre-trained vision-language models. *ArXiv*, abs/2210.13431, 2022. URL <https://api.semanticscholar.org/CorpusID:253098249>.

683

684 Guanxing Lu, Shiyi Zhang, Ziwei Wang, Changliu Liu, Jiwen Lu, and Yansong Tang. Manigaus-  
 685 sian: Dynamic gaussian splatting for multi-task robotic manipulation. In *European Conference  
 686 on Computer Vision*, 2024. URL <https://api.semanticscholar.org/CorpusID:268379077>.

687

688 Oier Mees, Lukás Hermann, Erick Rosete-Beas, and Wolfram Burgard. Calvin: A benchmark for  
 689 language-conditioned policy learning for long-horizon robot manipulation tasks. *IEEE Robotics  
 690 and Automation Letters*, 7:7327–7334, 2021. URL <https://api.semanticscholar.org/CorpusID:244908821>.

691

692 Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi,  
 693 and Ren Ng. Nerf. *Communications of the ACM*, 65:99 – 106, 2020. URL <https://api.semanticscholar.org/CorpusID:213175590>.

694

695 Abhishek Padalkar, Acorn Pooley, Ajinkya Jain, Alex Bewley, Alex Herzog, Alex Irpan, Alexan-  
 696 der Khazatsky, Anant Rai, Anikait Singh, Anthony Brohan, Antonin Raffin, Ayzaan Wahid,  
 697 Ben Burgess-Limerick, Beomjoon Kim, Bernhard Schölkopf, Brian Ichter, Cewu Lu, Charles  
 698 Xu, Chelsea Finn, Chenfeng Xu, Cheng Chi, Chenguang Huang, Christine Chan, Chuer Pan,  
 699 Chuyuan Fu, Coline Devin, Danny Driess, Deepak Pathak, Dhruv Shah, Dieter Büchler, Dmitry  
 700 Kalashnikov, Dorsa Sadigh, Edward Johns, Federico Ceola, Fei Xia, Freek Stulp, Gaoyue Zhou,  
 701 Gaurav S. Sukhatme, Gautam Salhotra, Ge Yan, Giulio Schiavi, Hao Su, Haoshu Fang, Haochen

702 Shi, Heni Ben Amor, Henrik I Christensen, Hiroki Furuta, Homer Rich Walke, Hongjie Fang, Igor  
 703 Mordatch, Ilija Radosavovic, Isabel Leal, Jacky Liang, Jaehyung Kim, Jan Schneider, Jasmine Hsu,  
 704 Jeannette Bohg, Jeff Bingham, Jiajun Wu, Jialin Wu, Jianlan Luo, Jiayuan Gu, Jie Tan, Jihoon Oh,  
 705 Jitendra Malik, Jonathan Tompson, Jonathan Yang, Joseph J. Lim, João Silvério, Junhyek Han,  
 706 Kanishka Rao, Karl Pertsch, Karol Hausman, Keegan Go, Keerthana Gopalakrishnan, Ken Gold-  
 707 berg, Kendra Byrne, Kenneth Oslund, Kento Kawaharazuka, Kevin Zhang, Keyvan Majd, Krishan  
 708 Rana, Krishna Parasuram Srinivasan, Lawrence Yunliang Chen, Lerrel Pinto, Liam Tan, Lionel  
 709 Ott, Lisa Lee, Masayoshi Tomizuka, Maximilian Du, Michael Ahn, Mingtong Zhang, Mingyu  
 710 Ding, Mohan Kumar Srirama, Mohit Sharma, Moo Jin Kim, Muhammad Zubair Irshad, Naoaki  
 711 Kanazawa, Nicklas Hansen, Nicolas Manfred Otto Heess, Nikhil J. Joshi, Niko Suenderhauf,  
 712 Norman Di Palo, Nur Muhammad Mahi Shafiullah, Oier Mees, Oliver Kroemer, Pannag R. Sanketi,  
 713 Paul Wohlhart, Peng Xu, Pierre Sermanet, Priya Sundaresan, Quan Ho Vuong, Rafael Rafailov, Ran  
 714 Tian, Ria Doshi, Russell Mendonca, Rutav Shah, Ryan Hoque, Ryan C. Julian, Samuel Bustamante,  
 715 Sean Kirmani, Sergey Levine, Sherry Moore, Shikhar Bahl, Shivin Dass, Shuran Song, Sichun Xu,  
 716 Siddhant Haldar, Simeon Adebola, Simon Guist, Soroush Nasiriany, Stefan Schaal, Stefan Welker,  
 717 Stephen Tian, Sudeep Dasari, Suneel Belkhale, Takayuki Osa, Tatsuya Harada, Tatsuya Mat-  
 718 sushima, Ted Xiao, Tianhe Yu, Tianli Ding, Todor Davchev, Tony Zhao, Travis Armstrong, Trevor  
 719 Darrell, Vidhi Jain, Vincent Vanhoucke, Wei Zhan, Wenxuan Zhou, Wolfram Burgard, Xi Chen,  
 720 Xiaolong Wang, Xinghao Zhu, Xuanlin Li, Yao Lu, Yevgen Chebotar, Yifan Zhou, Yifeng Zhu,  
 721 Ying Xu, Yixuan Wang, Yonatan Bisk, Yoonyoung Cho, Youngwoon Lee, Yuchen Cui, Yueh hua  
 722 Wu, Yujin Tang, Yuke Zhu, Yunzhu Li, Yusuke Iwasawa, Yutaka Matsuo, Zhuo Xu, and Zichen Jeff  
 723 Cui. Open x-embodiment: Robotic learning datasets and rt-x models. *ArXiv*, abs/2310.08864,  
 2023. URL <https://api.semanticscholar.org/CorpusID:263626099>.

724 Jack Parker-Holder, Philip Ball, Jake Bruce, Vibhavari Dasagi, Kristian Holsheimer, Chris-  
 725 tos Kaplanis, Alexandre Mouferek, Guy Scully, Jeremy Shar, Jimmy Shi, Stephen Spencer,  
 726 Jessica Yung, Michael Dennis, Sultan Kenjeyev, Shangbang Long, Vlad Mnih, Harris  
 727 Chan, Maxime Gazeau, Bonnie Li, Fabio Pardo, Luyu Wang, Lei Zhang, Frederic Besse,  
 728 Tim Harley, Anna Mitenkova, Jane Wang, Jeff Clune, Demis Hassabis, Raia Hadsell,  
 729 Adrian Bolton, Satinder Singh, and Tim Rocktäschel. Genie 2: A large-scale foun-  
 730 dation world model. 2024. URL <https://deepmind.google/discover/blog/genie-2-a-large-scale-foundation-world-model/>.

731 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,  
 732 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever.  
 733 Learning transferable visual models from natural language supervision. In *International Conference*  
 734 *on Machine Learning*, 2021. URL <https://api.semanticscholar.org/CorpusID:231591445>.

735 Scott Reed, Konrad Zolna, Emilio Parisotto, Sergio Gomez Colmenarejo, Alexander Novikov, Gabriel  
 736 Barth-Maron, Mai Giménez, Yury Sulsky, Jackie Kay, Jost Tobias Springenberg, Tom Eccles, Jake  
 737 Bruce, Ali Razavi, Ashley D. Edwards, Nicolas Manfred Otto Heess, Yutian Chen, Raia Hadsell,  
 738 Oriol Vinyals, Mahyar Bordbar, and Nando de Freitas. A generalist agent. *ArXiv*, abs/2205.06175,  
 739 2022. URL <https://api.semanticscholar.org/CorpusID:248722148>.

740 Eric Rohmer, Surya P. N. Singh, and Marc Freese. V-rep: A versatile and scalable robot simulation  
 741 framework. *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp.  
 742 1321–1326, 2013. URL <https://api.semanticscholar.org/CorpusID:960339>.

743 Mohit Shridhar, Lucas Manuelli, and Dieter Fox. Perceiver-actor: A multi-task transformer for robotic  
 744 manipulation. *ArXiv*, abs/2209.05451, 2022. URL <https://api.semanticscholar.org/CorpusID:252199474>.

745 Jianlin Su, Yu Lu, Shengfeng Pan, Bo Wen, and Yunfeng Liu. Roformer: Enhanced trans-  
 746 former with rotary position embedding. *ArXiv*, abs/2104.09864, 2021. URL <https://api.semanticscholar.org/CorpusID:233307138>.

747 Octo Model Team, Dibya Ghosh, Homer Rich Walke, Karl Pertsch, Kevin Black, Oier Mees,  
 748 Sudeep Dasari, Joey Hejna, Tobias Kreiman, Charles Xu, Jianlan Luo, You Liang Tan, Pan-  
 749 ntag R. Sanketi, Quan Vuong, Ted Xiao, Dorsa Sadigh, Chelsea Finn, and Sergey Levine.  
 750 Octo: An open-source generalist robot policy. *ArXiv*, abs/2405.12213, 2024. URL <https://api.semanticscholar.org/CorpusID:266379116>.

756 Mohammad Hassan Vali and Tom Bäckström. Nsvq: Noise substitution in vector quantization for  
 757 machine learning. *IEEE Access*, 10:13598–13610, 2022.

758

759 Aäron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning.  
 760 In *Neural Information Processing Systems*, 2017. URL <https://api.semanticscholar.org/CorpusID:20282961>.

761

762 Hongtao Wu, Ya Jing, Chi-Hou Cheang, Guangzeng Chen, Jiafeng Xu, Xinghang Li, Minghuan  
 763 Liu, Hang Li, and Tao Kong. Unleashing large-scale video generative pre-training for visual  
 764 robot manipulation. *ArXiv*, abs/2312.13139, 2023. URL <https://api.semanticscholar.org/CorpusID:266374724>.

765

766 Zhou Xian, Nikolaos Gkanatsios, Theophile Gervet, Tsung-Wei Ke, and Katerina Fragkiadaki.  
 767 Chaineddiffuser: Unifying trajectory diffusion and keypose prediction for robotic manipulation.  
 768 In *7th Annual Conference on Robot Learning*, 2023. URL <https://openreview.net/forum?id=W0zgY2mBTA8>.

769

770 Seonghyeon Ye, Joel Jang, Byeongguk Jeon, Se June Joo, Jianwei Yang, Baolin Peng, Ajay Mandlekar,  
 771 Reuben Tan, Yu-Wei Chao, Bill Yuchen Lin, Lars Lidén, Kimin Lee, Jianfeng Gao, Luke S. Zettlemoyer,  
 772 Dieter Fox, and Minjoon Seo. Latent action pretraining from videos. *ArXiv*, abs/2410.11758,  
 773 2024. URL <https://api.semanticscholar.org/CorpusID:273351190>.

774

775 Yanjie Ze, Ge Yan, Yueh-Hua Wu, Annabella Macaluso, Yuying Ge, Jianglong Ye, Nicklas Hansen,  
 776 Li Erran Li, and X. Wang. Gnfactor: Multi-task real robot learning with generalizable neural  
 777 feature fields. *ArXiv*, abs/2308.16891, 2023. URL <https://api.semanticscholar.org/CorpusID:261396262>.

778

779 Yanjie Ze, Gu Zhang, Kangning Zhang, Chenyuan Hu, Muhan Wang, and Huazhe Xu. 3d diffu-  
 780 sion policy. *ArXiv*, abs/2403.03954, 2024. URL <https://api.semanticscholar.org/CorpusID:268253298>.

781

782 Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation  
 783 representations in neural networks. *2019 IEEE/CVF Conference on Computer Vision and Pattern  
 784 Recognition (CVPR)*, pp. 5738–5746, 2018. URL <https://api.semanticscholar.org/CorpusID:56178817>.

785

786 Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigpt-4: Enhancing  
 787 vision-language understanding with advanced large language models. *ArXiv*, abs/2304.10592,  
 788 2023. URL <https://api.semanticscholar.org/CorpusID:258291930>.

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809

810 A THE USE OF LARGE LANGUAGE MODELS  
811812 This work utilized large language models as supplementary tools to enhance writing quality, including  
813 improving clarity, maintaining consistency between sections, and refining adherence to academic  
814 writing conventions. The core research concepts, methodological approaches, and findings represent  
815 original contributions by the authors.  
816817 B SIMULATION EXPERIMENTS  
818819 B.1 DATASET COMPOSITION  
820821 **RLBench.** RLBench is a comprehensive and large-scale benchmark and learning environment  
822 designed to advance research in vision-guided robotic manipulation. The platform is tailored to  
823 support a variety of research areas, including reinforcement learning, imitation learning, multi-task  
824 learning, and few-shot learning.  
825826 In our specific experimental setup, we utilize a subset of 10 manipulation tasks from the RLBench  
827 environment to evaluate the multi-task capabilities of our agents. These tasks are selected to cover a  
828 diverse range of challenges, involving different objects, objectives, and required skills. The variations  
829 within each task are designed to test an agent’s ability to understand and adapt to changes in color,  
830 placement, size, and object category.  
831832 Tab. 5 shows the details of the composition of the RLBench dataset.  
833

Task	Variation Type	# of Variations	Avg. Keyframes	Language Description Example
close jar	color	20	6.0	“close the — jar”
meat off grill	category	2	5.0	“take the — off the grill”
open drawer	placement	3	3.0	“open the — drawer”
sweep to dustpan	size	2	4.6	“sweep dirt to the — dustpan”
turn tap	placement	2	2.0	“turn — tap”
slide block	color	4	4.7	“slide the block to — target”
put in drawer	placement	3	12.0	“put the item in the — drawer”
drag stick	color	20	6.0	“use the stick to drag the cube onto the — — target”
push buttons	color	50	3.8	“push the — button, [then the — button]”
stack blocks	color, count	60	14.6	“stack — — blocks”

842 Table 5: Dataset composition of 10 manipulation tasks in RLBench [James et al. \(2019\)](#).  
843844  
845 **CALVIN.** A key evaluation protocol within CALVIN is the “ABC→D” setup, which is specifically  
846 designed to test an agent’s ability to generalize to a new, unseen environment. This setup is considered  
847 one of the most challenging evaluations in the benchmark.  
848849 The visualization of the Calvin setting is demonstrated in Fig. 5. This zero-shot generalization task,  
850 where the agent must apply learned skills to a completely new setting, is crucial for assessing the  
851 robustness and adaptability of the control policy. The ABC→D setup measures how well a policy  
852 can transfer its understanding of language and manipulation to an unfamiliar setting.  
853854 B.2 ADDITIONAL VISUALIZATION OF FUTURE OBSERVATIONS PREDICTION  
855856 Besides the visualization shown in Fig. 4 on RLBench, we also provide the future observations  
857 prediction results on Calvin. Fig. 6 illustrates the generated future observations conditioned on the  
858 current scene and the inferred latent action  $z^t$  for the “Open the drawer” task. Our model produces  
859 highly accurate and temporally consistent predictions in both RGB and depth modalities. The  
860 generated frames faithfully capture the geometric displacement of the drawer and the manipulator  
861 trajectory over multiple timesteps, showing smooth and physically plausible motion progression.  
862 Importantly, the predicted depth maps remain well aligned with the RGB predictions, indicating that  
863 the model preserves 3D scene structure rather than merely hallucinating pixel-level appearance. This  
864 coherence highlights that the latent action representation encodes the causal effect of actions in 3D  
865 space, enabling the diffusion model to generate consistent future trajectories.  
866

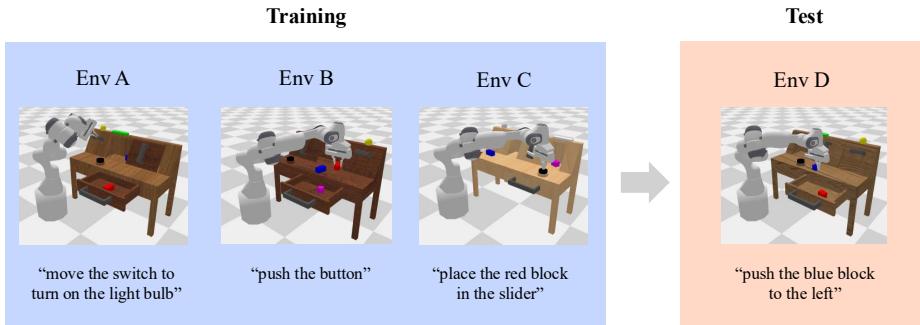


Figure 5: Illustration of the four different environments in CALVIN (Mees et al., 2021).

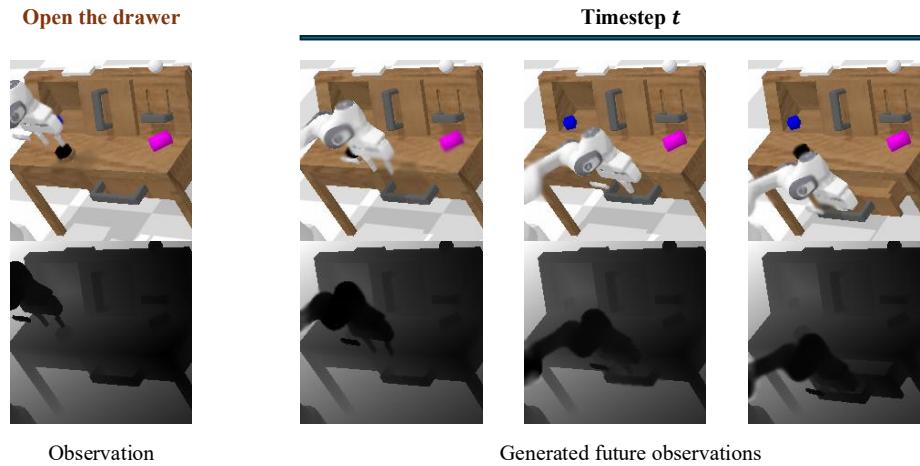


Figure 6: Future observation prediction on Calvin (Mees et al., 2021).

## C REAL-WORLD EXPERIMENTS

### C.1 DATASET COMPOSITION

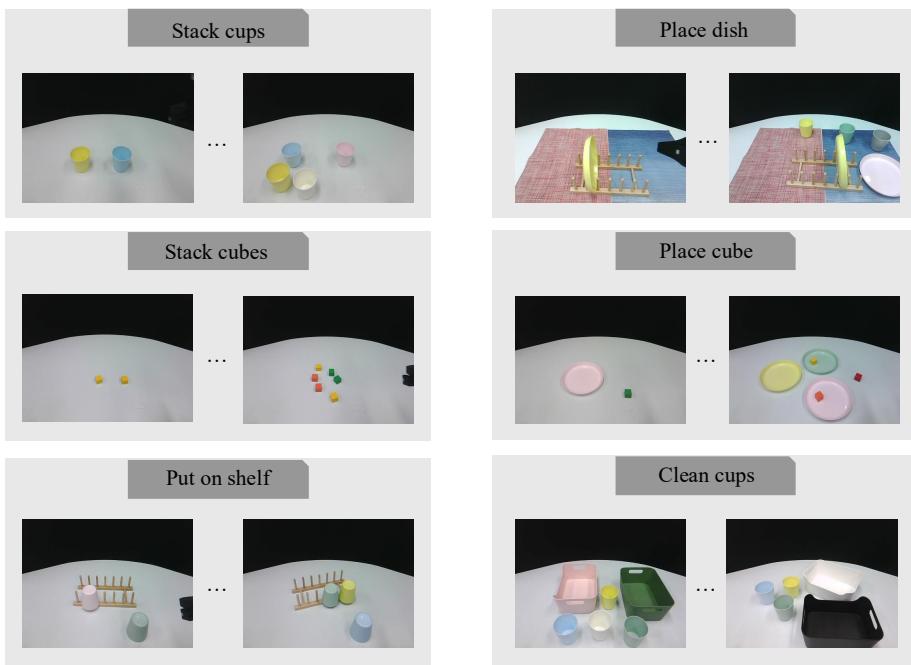
Task	Variation Type	Language Description Example
Place dish	color, count	“Place the — dish on the — tablecloth”
Clean cups	color, count, placement	“Put the — cup into the — basket”
Stack cups	color, count, placement	“Stack the — cup on the — cup”
Stack cubes	color, count	“Stack the — cubes”
Put cups on shelf	placement	“Put the — cup on the shelf next to the — cup”
Place cube	color, count	“Place the — block on the — plate”

Table 6: Dataset composition of 6 manipulation tasks in real robot experiments.

We evaluate our approach on six real-world manipulation tasks, each incorporating controlled variations in object placement, color, and count. This design introduces perceptual and spatial diversity, challenging the agent to generalize across visually distinct yet semantically similar scenarios. The visualization of our settings is shown in Fig. 7.

### C.2 MORE EXPERIMENTAL RESULTS

**Real-World Experiments.** We deploy our trained policy on a real Franka Panda manipulator and evaluate it on six long-horizon tabletop manipulation tasks: (1) “Put the blue cup on the shelf next to

Figure 7: **Visualization of the real-world experiment setup.**

the green cup”, (2) “Stack the yellow cup on the green cup”, (3) “Place the green cube on the pink plate”, (4) “Stack the orange cubes”, (5) “Place the blue dish on the blue tablecloth”, and (6) “Put the yellow cup into the white basket”. As shown in Fig. 8, our policy successfully completes all six tasks with smooth and collision-free trajectories, demonstrating strong sim-to-real transfer. The robot consistently executes precise grasps, object placements, and stacking behaviors, even in cluttered and visually diverse scenes. These results confirm that the learned latent actions generalize to real-world execution and maintain their semantic meaning outside of the simulation domain.

**Case study of precise action control.** Fig. 9 presents a qualitative comparison of action execution in a real-world cluttered scene, highlighting the advantage of our 4D modeling approach over a 3D static baseline. The top row (3D Diffuser Actor) demonstrates failure: although the robot successfully grasps and lifts the target cup, it misaligns during placement due to a lack of temporal dynamics and spatial reasoning — resulting in an unstable or incorrect stack. In contrast, our method (bottom row) leverages 4D spatio-temporal modeling to predict not only where but when and how to act, enabling precise control throughout the motion trajectory. As shown, our agent successfully stacks the cup with stable alignment, even under visual occlusion and object clutter. This illustrates that 4D-aware policy learning is critical for achieving reliable, fine-grained manipulation in dynamic physical environments — a capability absent in purely 3D state-based models.

## D IMPLEMENTATION DETAILS

### D.1 LEARNING GEOMETRY-AWARE LATENT ACTIONS.

To learn geometry-aware latent actions, we employ two different branches of diffusion models to predict the future observations.

Given RGB-D observation  $\mathbf{o}^t$ , we convert it to a pointmap:

$$\mathbf{P}^t = \text{BackProject}(\mathbf{o}^t) \in \mathbb{R}^{H \times W \times 3}.$$

**Latent Encoding.** Pointmaps are encoded by a 3D-aware VAE:

$$\mathbf{h}_{\text{pm}}^{t'} = \psi^{\text{enc}}(\mathbf{P}^{t'}), \quad \hat{\mathbf{P}}^{t'} = \psi^{\text{dec}}(\hat{\mathbf{h}}_{\text{pm}}^{t'}).$$

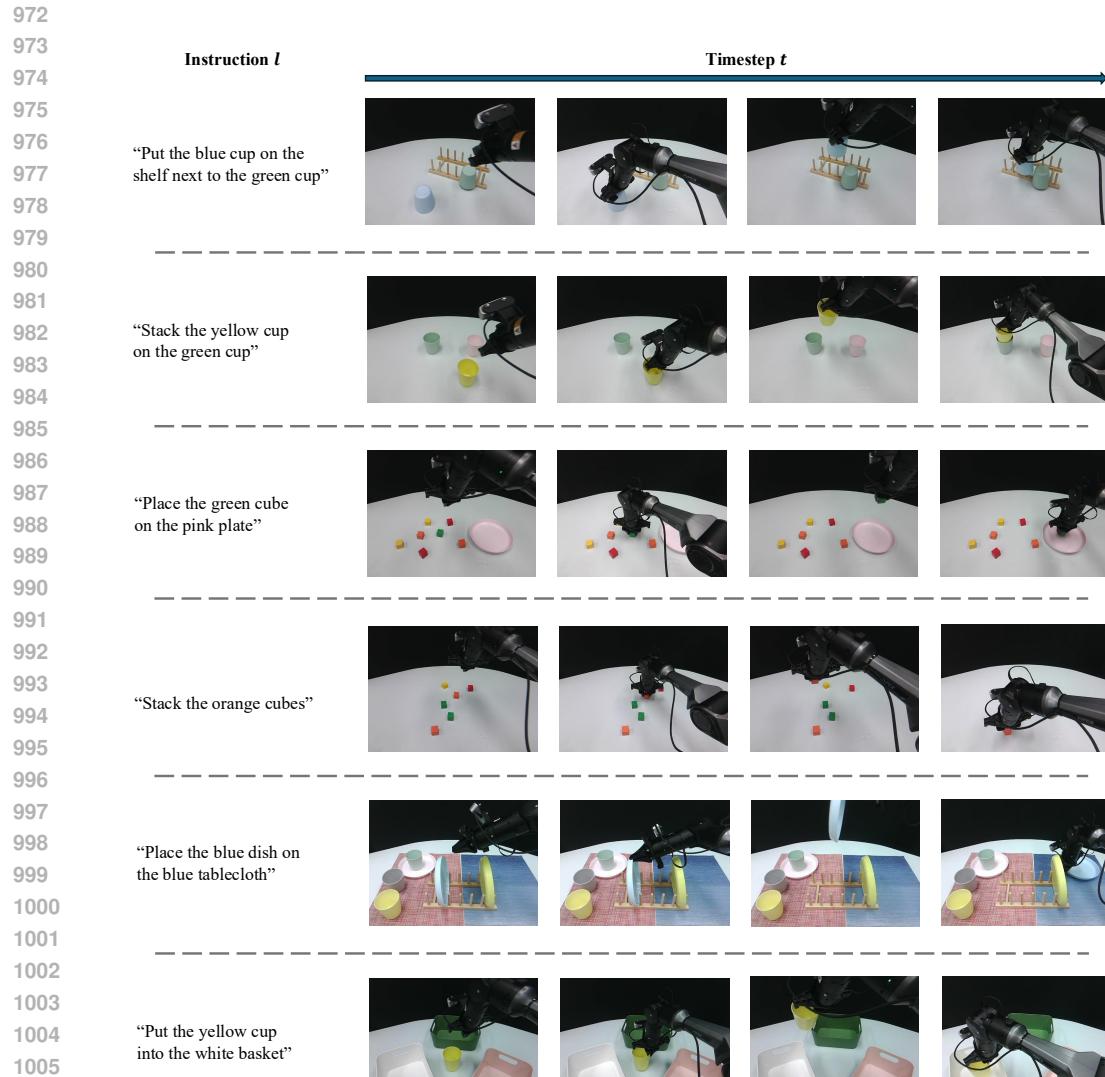


Figure 8: More qualitative results on real-world experiments.

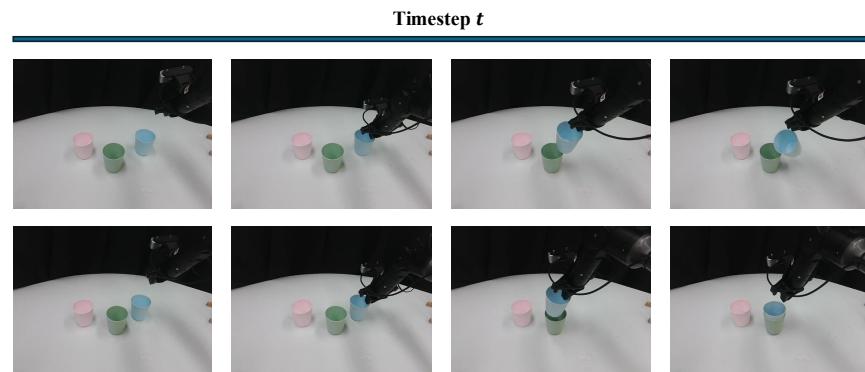


Figure 9: Qualitative comparison in the real-world cluttered scene.

1026 The pointmap VAE is initialized by the RGB VAE. Before the latent action learning, we first finetune  
 1027 it on the video sequences of demonstration for a few epochs.  
 1028

1029 **Latent Action-Conditioned Future Prediction.** We generate future latents with a conditional  
 1030 diffusion model adapted from DDPM Ho et al. (2020) with motion latent  $z^t$ :

$$1031 \hat{\mathbf{h}}_{\text{pm}}^{t+1:t+w} = \psi^{\text{diff}}(\mathbf{h}_{\text{pm}}^{t-w+1:t}, \mathbf{z}^t; \phi).$$

1032 At each diffusion step  $k$ , we add Gaussian noise:  
 1033

$$1034 \mathbf{h}_k^{\text{pm}} = \sqrt{\alpha_k} \cdot \hat{\mathbf{h}}_{\text{pm}}^{t+1:t+w} + \sqrt{1 - \alpha_k} \cdot \epsilon, \quad \epsilon \sim \mathcal{N}(0, I),$$

1035 and minimize the denoising objective:  
 1036

$$\mathcal{L}_{\text{diff}}^{\text{pm}} = \mathbb{E}_{k, \epsilon} [\|\epsilon - \epsilon_{\phi}(\mathbf{h}_k^{\text{pm}}, k, \mathbf{h}_{\text{pm}}^{t-w+1:t}, \mathbf{z}^t)\|^2].$$

1038 **Joint RGB Prediction.** An RGB branch predicts future appearance using the same architecture and  
 1039 conditioning:  
 1040

$$\mathcal{L}_{\text{diff}}^{\text{rgb}} = \mathbb{E}_{k, \epsilon} [\|\epsilon - \epsilon_{\phi}^{\text{rgb}}(\mathbf{h}_k^{\text{rgb}}, k, \mathbf{h}_{\text{rgb}}^{t-w+1:t}, \mathbf{z}^t)\|^2].$$

1042 **Vector-Quantized Latent Regularization.** Latent actions are discretized using VQ-VAE loss:  
 1043

$$\mathcal{L}_{\text{vq}} = \|\text{sg}[\mathbf{f}^t] - \mathbf{c}\|_2^2 + \beta \|\mathbf{f}^t - \text{sg}[\mathbf{c}]\|_2^2.$$

1045 **Overall Objective.** The total training loss is:  
 1046

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{diff}}^{\text{pm}} + \mathcal{L}_{\text{diff}}^{\text{rgb}} + \mathcal{L}_{\text{vq}}.$$

1048 Note that both two VAEs are all frozen during the training process.  
 1049

## 1050 D.2 ACTION PREDICTION WITH GEOMETRY-AWARE LATENT ACTIONS.

1052 The 3D denoising transformer  $\epsilon_{\theta}$  is a multi-layer conditional diffusion model that iteratively refines  
 1053 noisy action sequences into executable robot trajectories. The model takes four types of input tokens:  
 1054 trajectory tokens, scene tokens, proprioception tokens, and latent action tokens.

1055 **Token Embeddings.** We embed each noisy action  $\mathbf{a}^{t+k}$  using a MLP, producing a  $d_{\text{model}} = 256$ -  
 1056 dimensional trajectory token. Scene features are extracted using a frozen CLIP-ResNet50 encoder  
 1057 and lifted to 3D with depth information and camera intrinsics. The proprioceptive state  $\mathbf{c}^t$  is encoded  
 1058 with an MLP and added with the positional embedding of the end-effector pose.

1059 **Transformer Backbone.** The denoising network contains: (i) a multi-head self-attention layer  
 1060 over all trajectory, scene, and proprioception tokens with rotary 3D positional encodings; (ii) a  
 1061 cross-attention layer that attends from all tokens to the latent action embeddings  $z_t$  to inject high-level  
 1062 motion priors; In order to decrease the computational requirements, we subsample a number of visual  
 1063 tokens using Farthest Point Sampling (FPS). The sampled visual tokens, proprioception tokens, and  
 1064 noisy position/rotation tokens attend to each other.

1065 **Output Heads.** The final transformer layer outputs refined trajectory tokens, which are decoded  
 1066 using two independent MLP heads: One predicts the noise of position and 6D rotation, and the other  
 1067 predicts the gripper open/close state.

1068 **3D relative attention.** We formulate the detailed attention as follows:  
 1069

$$1070 a_{q,k} \propto x_q^T M(p_q - p_k) x_k$$

- 1071 •  $a_{q,k}$ : attention weight between query token  $q$  and key token  $k$
- 1072 •  $x_q$ : feature vector of the query token
- 1073 •  $x_k$ : feature vector of the key token
- 1074 •  $p_q$ : 3D position of the query token
- 1075 •  $p_k$ : 3D position of the key token
- 1076 •  $M(p_q - p_k)$ : matrix-valued function that depends only on the relative 3D position between  
 1077 query and key

---

		<b>RLBench</b>	<b>CALVIN</b>
<b>Transformer</b>			
1083	image_size	256	200
1084	embedding_dim	120	192
1085	camera_views	1	2
1086	FPS : % of sampled tokens	20%	33%
1087	diffusion_timestep	100	25
<b>Latent action</b>			
1088	Patch size	16	16
1089	Hidden size	768	768
1090	Codebook size	64	64
1091	Codebook dim	32	32
<b>Training</b>			
1093	batch_size	240	5400
1094	learning_rate	$1e^{-4}$	$3e^{-4}$
1095	weight_decay	$5e^{-4}$	$5e^{-3}$
1096	total_epochs	$1.6e^4$	90
1097	optimizer	Adam	Adam

---

Table 7: Comparison of configurations between RLBench and CALVIN.

### D.3 HYPER-PARAMETERS

The summary of used hyper-parameters for training/evaluating our model is described in Tab. 7.