PointMapPolicy: Structured Point Cloud Processing for Multi-Modal Imitation Learning

Xiaogang Jia^{1*} Qian Wang¹ Anrui Wang¹ Han A. Wang^{2†} Balázs Gyenes¹ Emiliyan Gospodinov¹ Xinkai Jiang¹ Ge Li¹ Hongyi Zhou¹ Weiran Liao¹ Xi Huang¹ Maximilian Beck³ Moritz Reuss¹ Rudolf Lioutikov¹ Gerhard Neumann¹

¹Karlsruhe Institute of Technology ²Reality Labs, Meta ³Johannes Kepler University Linz

Abstract

Robotic manipulation systems benefit from complementary sensing modalities, where each provides unique environmental information. Point clouds capture detailed geometric structure, while RGB images provide rich semantic context. Current point cloud methods struggle to capture fine-grained detail, especially for complex tasks, which RGB methods lack geometric awareness, which hinders their precision and generalization. We introduce PointMapPolicy, a novel approach that conditions diffusion policies on structured grids of points without downsampling. The resulting data type makes it easier to extract shape and spatial relationships from observations, and can be transformed between reference frames. Yet due to their structure in a regular grid, we enable the use of established computer vision techniques directly to 3D data. Using xLSTM as a backbone, our model efficiently fuses the point maps with RGB data for enhanced multi-modal perception. Through extensive experiments on the RoboCasa, CALVIN benchmarks and real robot evaluations, we demonstrate that our method achieves state-of-the-art performance across diverse manipulation tasks. The overview and demos are available on our project page.

1 Introduction

The advent of diffusion-based Imitation Learning (IL) has allowed robots to carry out complex, long-horizon tasks from raw image observations [1, 2]. RGB images are a common observation modality for diffusion policies due to their ubiquitousness and rich semantic information. However, policies conditioned on only RGB images lack 3D geometric information about the scene. This 3D information is crucial for learning generalizable policies that can act precisely in complex 3D scenes, especially when using multiple camera views [3–6]. An alternative modality is point clouds, unstructured sets of 3D points that preserve geometric shape, distances, and spatial relationships. In addition, points captured from multiple camera views can be transformed into a common reference frame and concatenated, yielding a natural and powerful way to fuse multiple cameras. Although numerous works use point clouds as an input modality [7–9], their irregular structure limits the network architectures that can be used with them. In contrast, RGB images are on a regular grid and can be processed using convolutional operators, but are susceptible to changes in perspective and lighting.

Current 3D processing approaches face fundamental limitations that create a critical gap between 3D geometric information and existing 2D vision architectures. Downsampling-based methods [7],

^{*}Correspondence to jia266163@gmail.com

[†]This work is not related to Han A. Wang's position at Meta.

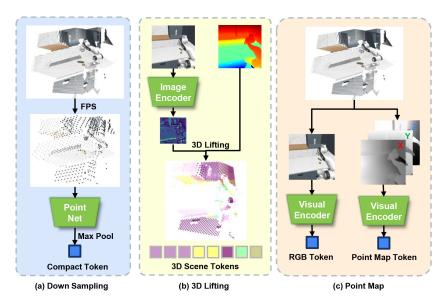


Figure 1: Different approaches for point cloud processing: (a) Downsampling-based methods use Furthest Point Sampling (FPS) to reduce dense point clouds to sparse representations, which PointNet then processes into compact tokens. Some variants employ FPS+KNN to generate structured point patches. (b) Feature-lifting approaches first extract 2D features from images, then project these features into 3D space, creating semantically rich 3D points. (c) Our point map method structures the point cloud as a 2D grid with the same dimensions as corresponding images, enabling direct application of efficient visual encoders to each modality independently.

as shown in Figure 1(a), suffer from an inherent information-fidelity tradeoff: they must dramatically reduce point density through techniques like Farthest Point Sampling (FPS)[10] to remain computationally tractable, inevitably discarding fine-grained geometric details essential for precise manipulation tasks. Feature-lifting approaches [11], as shown in Figure 1(b), face equally problematic limitations as they aggregate 2D features through depth averaging and 3D transformations, introducing information loss while struggling to maintain spatial structure during the lifting process. In this paper, we take inspiration from recent advances from the computer vision community in stereo reconstruction [12, 13] to propose using point maps, as shown in Figure 1(c). Point maps encode 3D information in a regular, 2D grid of Cartesian coordinates. This results in a structured data type that can be used with standard architectures such as ResNet [14], ViT [15], or ConvNeXt [16]. This obviates the need for steps like K-Nearest Neighors (KNN) and Farthest Point Sampling (FPS) [10], which are computationally expensive operations common to point cloud methods [17, 18, 6]. At the same time, because they are geometrically grounded, point maps from multiple views can be transformed into the same reference frame, increasing robustness to perturbations in camera perspective.

We integrate point maps into a standard diffusion-based imitation learning framework based on EDM [19] to demonstrate their effectiveness as a drop-in replacement for RGB images or point clouds. We validate the effectiveness of point map observations on two challenging benchmarks: RoboCasa [20] and CALVIN [21]. These benchmarks feature language-conditioned tasks and diverse scenes, and require spatial reasoning and long-term planning. Across both benchmarks, point map-based policies outperform baselines using RGB, depth maps, or point clouds, demonstrating superior learning efficiency and generalization. Our method is computationally efficient in training and inference, sometimes by an order of magnitude.

Contributions: Our contributions are the following: 1) we propose PointMapPolicy (PMP), a method for diffusion-based imitation learning on point maps, a powerful observation modality that has never been used in diffusion imitation learning; 2) we achieve state of the art results among policies trained from scratch on the CALVIN benchmark [21], and outperform other observation modalities on RoboCasa [20]; 3) we present systematic ablations of point cloud processing methods, vision backbones (e.g. ResNet [14], ViT [15], ConvNeXt [16]), and paradigms for fusing color and geometry information.

2 Related Work

2D Visual Representations for Imitation Learning. Recent imitation learning approaches [1, 22–25] rely predominantly on 2D visual representations such as RGB images or videos. Such representations are widely utilized due to their capacity to capture rich textural and semantic information, as well as their accessibility through low-cost cameras. However, 2D image modalities have inherent limitations: they contain 3D information only implicitly, are vulnerable to viewpoint and lighting changes and occlusions, and typically underperform in tasks requiring detailed spatial reasoning and geometric alignment [3–6].

3D Visual Representations for Imitation Learning. To overcome these limitations, a growing amount of research incorporates explictly 3D representations such as depth maps, point clouds, or voxels. Voxel-based methods like C2F-ARM [26] and Perceiver-Actor [27] voxelize point clouds and use a 3D-convolutional network for action prediction, but require high voxel resolution for precision tasks, resulting in high memory consumption and slow training. DP3 [7] encodes sparse point clouds using FPS, followed by a lightweight MLP to produce a compact embedding vector of the observation. While efficient, this approach discards local geometric structure that can be critical for fine-grained tasks. In contrast, 3D Diffuser Actor [11] computes tokens by lifting 2D image features into 3D space by using averaged depth information and camera parameters, and applies FPS after the first cross-attention layer. FPV-Net [28] fuses RGB and point cloud modalities by injecting global and local image features into a point cloud encoder using adaptive normalization layers, but is still limited by the disadvantages of both modalities.

Multi-View Representation. Complementary work, such as Robot Vision Transformer (RVT) [29], avoids working directly with raw point clouds by proposing a novel multi-view representation. This approach re-renders the point cloud from a set of orthographic virtual cameras, deriving a 7-channel point map (RGBD + XYZ) from each view. RVT-2 [30] improves this approach for high-precision tasks by introducing a multi-stage inference pipeline: it first identifies a region-of-interest, truncates the observation to this area of interest, and then runs policy inference. However, neither of these methods use action diffusion, instead relying on key-frame based manipulation with a motion planner [31]. Furthermore, geometric and color information are fused naively at the channel level, whereas we investigate more sophisticated techniques for fusion.

Diffusion-Policy Backbones. Due to the non-Markovian nature of human demonstrations, where successful decision-making often depends on histories of past observations and actions, early work used RNN-based architectures [32], but struggled with vanishing gradients and limited scalability. This led to the adoption of Transformer-based architectures, which offer global attention and parallelism, enabling superior performance in tasks requiring long-horizon reasoning [33–35], becoming the standard backbone for many methods [7, 11, 28, 22]. However, Transformers are computationally intensive and scale quadratically with the sequence length, which limits the number of tokens that can be used to encode the observation.

To mitigate these challenges, recent works [36, 37] explore State Space Models (SSMs) like Mamba [38], achieving linear-time complexity and improved sample efficiency, particularly in low-data regimes. Additionally, recent recurrent architectures such as xLSTM [39] provide an appealing balance, maintaining the temporal modeling strengths of traditional RNNs while introducing architectural innovations that improve gradient flow and expressiveness. Despite being less expressive than self-attention, xLSTM significantly reduces compute and memory costs, making it well-suited for real-time or resource-constrained applications. X-IL [37] systematically compares different architectural parts, and finds that xLSTM performs competitively with Transformers in multi-modal imitation learning. Building on these insights, PointMapPolicy adopts xLSTM as its diffusion backbone, balancing temporal modeling capability with efficient training and inference.

3 Method

3.1 Problem Formulation

Imitation Learning aims to learn a policy from expert demonstrations. Given a dataset of expert trajectories $\mathcal{D}_{\tau} = \{\tau_i\}_{i=1}^N$, where each trajectory $\tau_i = ((\mathbf{s}_1, \mathbf{a}_1), (\mathbf{s}_2, \mathbf{a}_2), \dots, (\mathbf{s}_K, \mathbf{a}_K))$. The objective is to learn a policy $\pi(\bar{\mathbf{a}}|\mathbf{s})$ that maps observations \mathbf{s} to a sequence of actions $\bar{\mathbf{a}} = (\mathbf{a}_k, \mathbf{a}_{k+1}, \dots, \mathbf{a}_{k+H})$. Predicting sequences of actions, i.e. action chunking, allows for more temporally consistent action

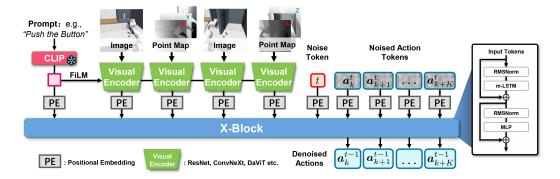


Figure 2: **Overview of PMP.** PMP integrates multiple modalities: language instructions encoded by a pretrained CLIP model, images processed by pretrained visual encoders, and point maps processed by visual encoders trained from scratch. Leveraging x-LSTM as its backbone architecture, PMP efficiently fuses these multi-modal inputs to generate denoised actions.

prediction [40]. Each observation s contains multi-view RGB-D images and language instruction for the current trajectory.

3.2 Score-based Diffusion Policy

Our approach employs the EDM framework for continuous-time action diffusion [19, 41] to generate actions. Diffusion models are generative models that learn to generate new samples through learning to reverse a Gaussian Perturbation process. In PointMapPolicy, we apply a score-based diffusion model to formulate the policy representation $\pi_{\theta}(\bar{a}|s)$. This perturbation and its inverse process can be expressed through a Stochastic Differential Equation (SDE):

$$d\bar{\mathbf{a}} = (\beta_t \sigma_t - \dot{\sigma}_t) \sigma_t \nabla_a \log p_t(\bar{\mathbf{a}}|\mathbf{s}) dt + \sqrt{2\beta_t} \sigma_t d\omega_t, \tag{1}$$

where β_t determines the noise injection rate, $d\omega_t$ represents infinitesimal Gaussian noise, and $p_t(\bar{a}|s)$ denotes the score function of the diffusion process. It guides samples away from high-density regions during the forward process. To learn this score, we train a neural network D_{θ} via score matching [42]:

$$\mathcal{L}_{SM} = \mathbb{E}_{\sigma, \bar{\boldsymbol{a}}, \epsilon} \left[\alpha(\sigma_t) | D_{\theta}(\bar{\boldsymbol{a}} + \epsilon, \mathbf{s}, \sigma_t) - \bar{\boldsymbol{a}}|_2^2 \right], \tag{2}$$

where $D_{\theta}(\bar{a} + \epsilon, s, \sigma_t)$ represents our trainable neural architecture.

After training, we can generate new action sequences beginning with Gaussian noise by iteratively denoising the action sequence with a numerical Ordinary Differential Equation (ODE) solver. Our approach utilizes the DDIM-solver, a specialized numerical ODE-solver for diffusion models [43] that enables efficient action denoising in a minimal number of steps. Across all experiments, our method uses 4 denoising steps.

3.3 Observation Tokenization

We are given an observation s_k in step k as well as a textual language instruction z_{lang} . The language instruction is first tokenized using a pretrained CLIP text encoder [44] to generate language embeddings. For RGB inputs, we use Film-ResNet [45] with pretrained ImageNet weights to generate visual embeddings from the observation s_k .

We define that a Point Map $X \in \mathbb{R}^{H \times W \times 3}$ is a dense 2D field of 3D points that establishes a one-to-one mapping between image pixels and 3D scene points. For an RGB image I of resolution $H \times W$, the corresponding Point Map X satisfies $I_{i,j} \leftrightarrow X_{i,j}$ for all pixel coordinates $(i,j) \in \{1 \dots H\} \times \{1 \dots W\}$, where each pixel intensity $I_{i,j}$ corresponds to a 3D point $X_{i,j} \in \mathbb{R}^3$ in world coordinates.

We convert each depth map $D \in \mathbb{R}^{H \times W}$ to a structured point map representation:

$$\mathbf{M}_t = \phi(D, K_{\text{int}}^{-1}), \quad \mathbf{M}_t \in \mathbb{R}^{H \times W \times C}$$
(3)

where $K_{\rm int}$ are the camera intrinsic parameters obtained through calibration and ϕ is a depth unprojection operation. The result is a multi-channel point map with the same spatial dimensions as the input depth map, where the channel dimension C is typically 3. Points beyond a maximum depth and below a minimum depth are masked out. Point maps from all cameras are transformed into a common world reference frame using the extrinsic parameters of the camera $K_{\rm ext}$.

3.4 PointMapPolicy

PointMapPolicy uses EDM-based action diffusion for decision making and conditions on the multimodal observation tokens generated from RGB and point map modalities. We present and explore multiple paradigms for fusing RGB and geometric data at various stages of processing. We also describe PMP-xyz, a variant with tokens from only the point map modality, for tasks that do not condition on color information.

Fusion of image and point maps. A key advantage of point maps is their ability to provide both geometric and visual embeddings for each camera view, enabling straightforward multimodal fusion. We investigate both early and late fusion approaches. For early fusion *PMP-6ch*, we concatenate point maps with RGB values, creating six-channel inputs (XYZ + RGB). For late fusion, we first tokenize image and point map modalities from each view with separate



Figure 3: Fusion methods. From left to right: *Add, Cat*, and *Attn*.

encoders. Then we explore three methods to fuse encoded tokens, as illustrated in Figure 3: 1) *Add*, element-wise addition of tokens from both modalities, resulting in one token per view; 2) *Cat*, concatenation of tokens from all modalities and views; and 3) *Attn*, using a four-layer transformer module to process tokens using cross-attention to generate fused class tokens for each view. As shown in our ablation studies, we find *Cat* to slightly outperform other late fusion methods, so we choose this for PMP. An overview of PMP with *Cat* fusion is illustrated in Figure 2.

Backbones. Given the multi-modal tokens from Section 3.3, a learnable positional embedding is added to each token. PMP uses a decoder-only backbone from X-IL [37] with x-LSTM as the core computational unit. All tokens are concatenated as inputs to the X-Block, which is the diffusion score network D_{θ} . While Transformers dominate most imitation learning policies, X-IL demonstrated that the recent recurrent architecture xLSTM excels in robot learning tasks. The core computational element within X-Block is the m-LSTM layer, which serves an analogous function to self-attention in Transformer architectures. The denoised action tokens produced by the X-Block are then used to guide the robot's behavior, resulting in a policy that effectively leverages both the geometric precision of point maps and the rich semantic understanding from RGB images.

4 Simulation Experiments

We conduct experiments on two simulation benchmarks RoboCasa [20] and CALVIN [21]. We aim to answer the following questions: Q1) How does PointMapPolicy compare to state-of-the-art 2D and 3D imitation learning policies? Q2) How do the fusion methods perform compared to other modalities? Q3) How does point map representation compare to other point cloud processing methods? Q4) Can current vision encoders effectively extract the geometric and semantic information from point maps required for robust decision-making?

RoboCasa: The RoboCasa benchmark [20] is a large-scale simulation framework designed to evaluate IL agents across a wide range of household manipulation tasks. Built on a physically realistic environment with rich visual rendering, RoboCasa supports task diversity, long-horizon behaviors, and fine-grained physical interactions, making it a compelling testbed for assessing both generalization and behavior diversity in policy learning. We use the RoboCasa benchmark to assess whether our proposed point map representation can enable effective learning and generalization across manipulation tasks of increasing complexity, object count, and behavioral variation.

CALVIN: The CALVIN benchmark [21] provides a large-scale framework for evaluating language-conditioned IL policies in visually rich, long-horizon manipulation tasks. The benchmark contains 34 distinct manipulation tasks such as button-pressing, drawer-opening, object-picking, and pushing.

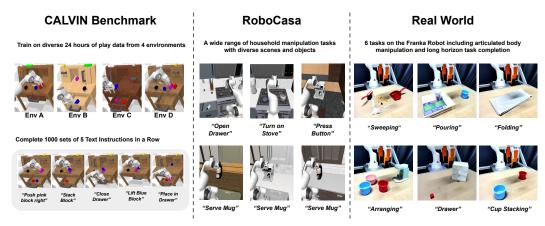


Figure 4: Overview of Simulation and Real World Experiments used to test PointMapPolicy. From left to right: CALVIN Benchmark [21], RoboCasa [20], and our Real World Setup.

Category	Task	BC	GR00T-N1	DP3	3DA	RGB	Depth	PMP-6ch	PMP-xyz	PMP
	PnPCounterToMicrowave	2.0	0.0	$4.0_{\pm 1.6}$	0.0	10.0±4.3	3.3 ± 0.9	$9.3\pm_{0.9}$	13.3 ± 3.4	10.7 ± 3.8
Pick and Place	PnPCounterToSink	2.0	1.0	0.7 ± 0.9	0.0	5.3±1.9	4.7 ± 0.9	$8.7{\scriptstyle\pm0.9}$	$6.7{\pm}2.5$	6.7 ± 3.4
rick and riace	PnPMicrowaveToCounter	2.0	0.0	4.0 ± 2.8	0.0	10.7±3.8	8.0 ± 1.6	12.0 ± 1.6	$16.0{\scriptstyle\pm1.6}$	$16.0{\pm}6.5$
	PnPSinkToCounter	8.0	5.9	1.3 ± 0.9	0.0	14.7±0.9	$3.3{\scriptstyle\pm0.9}$	9.3 ± 4.7	$8.0{\scriptstyle\pm1.6}$	$16.7{\scriptstyle\pm5.0}$
Open/Close Drawers	OpenDrawer	42.0	42.2	46.0 ± 3.3	18.0	44.7±7.7	56.7 ± 0.9	40.0 ± 3.3	$60.0{\scriptstyle\pm4.3}$	56.0 ± 8.2
Open/Close Drawers	CloseDrawer	80.0	96.1	$60.0{\scriptstyle\pm1.6}$	80.0	84.0±7.1	$92.0{\scriptstyle\pm0.0}$	$75.3{\scriptstyle\pm3.4}$	$96.0{\scriptstyle\pm1.6}$	$91.3{\pm}3.8$
Twisting Knobs	TurnOnStove	32.0	25.5	24.7 ± 4.1	18.0	18.7±1.9	22.7 ± 0.9	35.3 ± 2.5	$43.3{\scriptstyle\pm5.7}$	41.3 ± 5.0
	TurnOffStove	4.0	15.7	$7.3{\scriptstyle\pm1.9}$	8.0	13.3±0.9	$14.0{\scriptstyle\pm1.6}$	$16.7{\scriptstyle\pm2.5}$	$20.0{\pm}3.3$	18.0 ± 0.0
	TurnOnSinkFaucet	38.0	59.8	42.0±3.3	26.0	64.0±7.1	$78.7{\scriptstyle\pm2.5}$	64.7 ± 4.1	76.7 ± 4.1	66.7±8.1
Turning Levers	TurnOffSinkFaucet	50.0	67.7	$42.0 {\pm} 4.9$	44.0	63.3±7.7	$76.0{\scriptstyle\pm5.9}$	73.3 ± 9.6	$82.0{\scriptstyle\pm1.6}$	$66.7{\scriptstyle\pm9.4}$
	TurnSinkSpout	54.0	42.2	$58.7{\pm}6.8$	28.0	50.0±4.9	$\textbf{76.0} {\pm} 4.9$	$69.3 {\pm} 3.8$	$76.0 {\scriptstyle\pm1.6}$	48.7 ± 7.4
	CoffeePressButton	48.0	56.9	14.7 ± 0.9	8.0	70.7±13.7	84.0 ± 4.3	76.7 ± 6.6	82.7 ± 5.0	92.0 ± 3.3
Pressing Buttons	TurnOnMicrowave	62.0	73.5	$39.3 {\pm} 7.4$	34.0	48.0±4.9	44.0 ± 0.0	$64.7{\scriptstyle\pm5.2}$	$49.3{\scriptstyle\pm5.0}$	$64.7{\scriptstyle\pm3.4}$
	TurnOffMicrowave	70.0	57.8	$62.7{\scriptstyle\pm5.7}$	30.0	69.3±9.0	$68.0{\scriptstyle\pm7.1}$	$75.3{\scriptstyle\pm2.5}$	$70.0{\scriptstyle\pm4.9}$	$84.0{\pm}6.5$
Insertion	CoffeeServeMug	22.0	34.3	21.3±0.9	0.0	60.0±2.8	57.3±6.2	48.7±7.7	69.3±6.6	49.3±3.4
HISCHIOH	CoffeeSetupMug	0.0	2.0	4.0 ± 2.8	2.0	16.0±2.8	15.3 ± 0.9	10.7 ± 0.9	$16.7{\pm}3.8$	$26.7{\pm}3.4$
Average Success Rate		32.25	36.28	27.04	18.50	40.16	44.00	43.12	49.12	47.22

Table 1: Success rate (%) for each task in RoboCasa [20]. The models were trained for 50 epochs with 50 human demonstrations per task and evaluated with 50 episodes for each task. The bold numbers highlight the best achieved success rate for that task among all the models.

Each rollout consists of a sequence of 5 language instructions, and the agent must complete one task before proceeding to the next. Policies are evaluated on 1,000 such instruction chains per seed, and success is measured by the average number of correctly completed tasks in each sequence.

Experimental Setup: For RoboCasa, each model was trained for 50 epochs using three random seeds, with performance measured at the 30th, 40th, and 50th checkpoints, selecting the best result. To ensure fair comparison, all models across different modalities use identical backbone parameters. For the CALVIN benchmark, models were trained for 25 epochs, with the best success rate reported from the 10th, 15th, 20th, and 25th checkpoints.

Baselines: For RoboCasa, we benchmark against Behavioral Cloning (BC) [20], GR00T-N1 [46], 3D Diffusion Policy (DP3) [7], and 3D Diffuser Actor (3DA) [11]. Note that GR00T-N1 results use 100 demonstrations, while all other methods use 50 human demonstrations. To systematically evaluate the effectiveness of our representation, we further compare other against image-based baselines using only RGB data (RGB), and only depth data (Depth). We then compare PMP against multiple variants introduced in Section 3.4: *PMP-6ch* directly uses 6-channel point maps as inputs, and *PMP-xyz* only uses xyz coordinates as inputs. All five methods share the same architectures and parameters for fair comparison. Details can be found in Appendix 6.

For CALVIN, we primarily compare our approach against models without robot-specific pretraining, though we include all results for reference. DP3, 3DA, and CLOVER [47] are selected as representa-

Train→Test	Method	PrT	Action Type	No. Instructions in a Row (1000 chains)			Avg. Len.		
				1	2	3	4	5	
	RoboFlamingo [50]	√	Cont.	82.4%	61.9%	46.6%	33.1%	23.5%	2.47
	SuSIE [51]	✓	Diffusion	87.0%	69.0%	49.0%	38.0%	26.0%	2.69
	GR-1 [52]	✓	Cont.	85.4%	71.2%	59.6%	49.7%	40.1%	3.06
	OpenVLA [23]	✓	Discrete	91.3%	77.8%	62.0%	52.1%	43.5%	3.27
	RoboDual [53]	✓	Diffusion	94.4%	82.7%	72.1%	62.4%	54.4%	3.66
	Seer [49]	✓	Cont.	94.4%	87.2%	79.9%	72.2%	64.3%	3.98
	MoDE [24]	✓	Diffusion	96.2%	88.9%	81.1%	71.8%	63.5%	4.01
	Seer-Large [49]	✓	Cont.	96.3%	91.6%	86.1%	80.3%	74.0%	4.28
$ABC \rightarrow D$	DP3 [7]	×	Diffusion	28.7%	2.7%	0.0%	0.0%	0.0%	0.31
	MDT [48]	×	Diffusion	63.1%	42.9%	24.7%	15.1%	9.1%	1.55
	3DA [11]	×	Diffusion	92.2%	78.7%	63.9%	51.2%	41.2%	3.27
	MoDE (scratch) [24]	×	Diffusion	91.5%	79.2%	67.3%	55.8%	45.3%	3.39
	CLOVER [47]	×	Diffusion	96.0%	83.5%	70.8%	57.5%	45.4%	3.53
	Seer (scratch) [49]	×	Cont.	93.0%	82.4%	72.3%	62.6%	53.3%	3.64
	Seer-Large (scratch) [49]	×	Cont.	92.7%	84.6%	76.1%	68.9%	60.3%	3.83
	RGB	×	Diffusion	89.9%	75.4%	60.8%	49.8%	39.1%	3.15
	PMP-xyz (ours)	×	Diffusion	73.0%	51.9%	37.0%	24.5%	16.1%	2.03
	PMP (ours)	×	Diffusion	96.1%	88.6%	80.5%	72.3%	63.6%	4.01

Table 2: Evaluation results on the CALVIN benchmark under ABC→D. All results report the average rollout length averaged over 1000 instruction chains.

tive policies using RGB-D inputs. MDT [48], MoDE [24], and Seer [49] are selected as RGB-based policies. Further details of these baselines can be found in Appendix A.2.

4.1 Main Results

RoboCasa. We present the main results in Table 1. PMP-xyz demonstrates significant advantages over prior 3D baselines DP3 and 3DA, achieving an average success rate of 49.12%—nearly 20% higher. It also outperforms 2D baselines BC and GR00T-N1 by approximately 13%. The above results address **Q1**. Our cross-modality evaluation using consistent architectures reveals that incorporating 3D information consistently improves performance in RoboCasa. Specifically, PMP-xyz shows a 6% improvement over the Depth-only model, highlighting the value of structured point maps. While PMP (47.22%) outperforms PMP-6ch, demonstrating the benefits of late fusion, it still falls 2% short of PMP-xyz. This pattern suggests that most RoboCasa tasks favor geometric information, likely due to the diversity of objects and scenes.

CALVIN. On the CALVIN benchmark, PMP achieves a score of 4.01, outperforming all other models trained from scratch and many models that leverage pretrained data, as shown in Table 2. Our method even outperforms Seer-Large (scratch) which scores 3.83 despite using 24 Transformer layers compared to our smaller model using only 10 x-LSTM blocks. This answers **Q1** in the affirmative.

PMP-xyz performs poorly with an average rollout length of only 2.03, while the RGB-only model achieves a respectable score of 3.15. This performance disparity stems from CALVIN's heavy reliance on color information for task execution, with many instructions explicitly referencing colors (e.g., "red block", "pink block"). This finding highlights a key limitation of purely geometric representations in color-dependent scenarios.

These contrasting results between RoboCasa and CALVIN benchmarks underscore the complementary nature of geometric and visual information. While PMP-xyz excels in geometry-heavy tasks (RoboCasa), it struggles with color-dependent tasks (CALVIN). This demonstrates that multimodal fusion approaches like PMP provide the most robust and versatile performance across diverse task domains by adaptively leveraging the most relevant modality for each scenario, addressing **Q2**.

4.2 Point Cloud Encoding

While our main results demonstrate the effectiveness of PMP compared to other methods, these comparisons involve different policy backbone architectures. To isolate the contribution of our point cloud encoding approach, we conduct a controlled ablation study where we fix the policy backbone (X-Block) and systematically vary only the point cloud encoder.

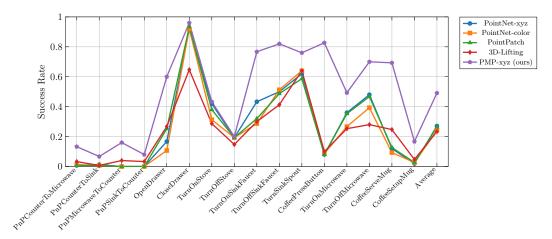


Figure 5: Ablation study comparing point cloud encoders with fixed X-Block policy backbone. Our PMP-xyz method substantially outperforms baseline encoders across all manipulation tasks, demonstrating that our improvements arise from the point cloud encoding approach.

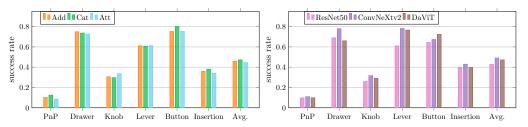


Figure 6: *Left*: Performance comparison of various fusion methods between point maps and images. *Right*: Performance comparison of different visual encoders for point map processing.

We conducted controlled experiments on RoboCasa using identical xLSTM backbones with different point cloud processing encoders: 1) PointNet-xyz: Following DP3 [7], we gather point clouds from 3 camera views and use Furthest Point Sampling (FPS) to downsample to 1024 points, then apply MLP with maxpooling to create a compact 3D token. 2) PointNet-color: Same process as PointNet-xyz but using colored points with XYZRGB information. 3) PointPatch: We use FPS to sample 256 center points, apply k-Nearest Neighbors to create 256 point patches with 32 points each, tokenize each patch using MLP with maxpooling, then process the resulting tokens with a transformer to generate compact 3D representations. 4) 3D-Lifting: We extract CLIP features (frozen) from each camera view and lift the 2D features into 3D space, then use a transformer to process the lifted tokens. The 3D tokens are then passed to the diffusion policy with an identical X-Block backbone.

Figure 5 presents the success rates over 16 RoboCasa tasks. Our PMP-xyz achieves an average success rate of 49.12%, substantially outperforming all baselines. The consistent improvements demonstrate that the point maps approach effectively captures the spatial understanding necessary for robotic manipulation, independent of the downstream policy architecture, addressing **Q3**.

4.3 Ablation Study

We additionally conduct three ablations across the 6 categories of RoboCasa:

Fusion of Images and Point Maps. One key advantage of point maps over traditional point cloud processing methods is their structural similarity to RGB images from corresponding views. This alignment enables direct fusion of visual representations with point cloud data on a per-view basis. We evaluate three fusion strategies—*Add*, *Cat*, and *Attn*—which are described in Figure 3. The comparative performance of these fusion strategies is presented in Figure 6. Although the performance differences are modest, *Cat* consistently emerges as the most effective fusion approach.

Vision Encoders for Point Maps. To assess how well existing visual architectures process Point Map representations, we conduct a comparative analysis of three prominent visual encoders: FiLM-ResNet50, ConvNeXt-v2, and DaViT. The results in Figure 6 demonstrate that while all encoders can effectively process point maps, ConvNeXtv2 consistently outperforms the others across all RoboCasa tasks, addressing **Q4**.

Understanding Model Attention Patterns. To uncover where the model attends during action prediction, we apply Grad-CAM++ [54] to highlight the regions most influential for action decisions across different modalities. For detailed visualizations, see Appendix D.

4.4 Computation Resources and Inference Time

For the CALVIN experiments, PMP employs Film-ResNet50 as encoders for both images and point maps, with 8 x-Blocks as backbones (512 latent dimensions), totaling 147M trainable parameters. Training utilizes 4 Nvidia RTX 6000 Ada GPUs with 128 samples per GPU (512 total batch size). Each epoch completes in approximately 13 minutes, allowing full training (25 epochs) in under 6 hours, excluding evaluation time. More details can be found in Appendix E.

Regarding computational efficiency, we conducted inference latency benchmarks for our models using ConvNeXt-v2 encoders on a single Nnidia RTX 5080 GPU (batch size 1). Across 1000 prediction cycles, PMP-xyz demonstrates remarkable efficiency with an average inference time of 2.9 ms, while PMP requires only 3.9 ms, maintaining real-time performance.

5 Real-World Experiments

We evaluate PMP on six challenging real-world robot manipulation tasks: Arranging, Folding, Cup-Stacking, Drawer, Pouring, and Sweeping. An overview of our robot setup is shown in Figure 7. The robot's perception system consists of two RGB-D cameras mounted on the left and right sides of the workspace.

5.1 Real-World Benchmark

Real-world Setup. We evaluate our policies on a 7-DOF Franka Panda robot in six challenging tasks. Visual information is captured by two Orbbec Femto Bolt cameras, positioned to provide left and right views. These sensors provide both RGB and depth images, which are used to generate calibrated 3D point clouds. All RGB and depth images are resized to 180×320 resolution. The robot operates in an 8-dimensional action space, including joint positions and gripper state.

Datasets. For collecting demonstrations, we use a teleoperation system consisting of a leader robot and a follower robot. For each task, we collect varying numbers of language-conditioned trajectories as detailed in Table 3. To ensure robust evaluation, we randomly initialize the object and goal states, introducing significant variation in the objects used. For instance, in the sweep task, the broom can appear in 10 different areas, and the garbage in 4 different areas. In addition, the number, positions, and even categories of trash items are varied in the collection and evaluation.

5.2 Baselines and Metrics

Baseline. To evaluate the effectiveness of the point-map representation, we benchmark our methods against RGB-only policy, sharing with the same backbone. Each method is evaluated over 20 trials per task at training checkpoints 70,000, 80,000, and 90,000, using randomized initial object states to ensure robustness. We report results from the best-performing checkpoint for each method.

Metrics. Given the long-horizon nature and complexity of the tasks, we introduce a structured scoring metric to enable fair and detailed comparisons. Each task is decomposed into multiple stages, with each successfully completed stage contributing 1 point to the overall score. The final task score is the sum of the completed intermediate stages, providing a more granular measure of progress and policy effectiveness. The details of our scoring metrics can be found in Table 5.

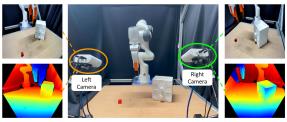


Table 3: The table shows average completed stages. The Max. indicates the total number of stages per task.

Tasks	Demos	Methods with Scores					
Tuono	Per Task	RGB	PMP-xyz	PMP	Max		
Arranging	80	2.05	2.10	2.25	3		
Folding	45	2.1	0.80	2.50	3		
Cup-Stacking	75	1.40	0.45	2.10	3		
Drawer	120	2.00	2.15	2.40	4		
Pouring	80	1.55	1.60	1.80	4		
Sweeping	90	1.80	0.80	2.15	4		

Figure 7: Real world experiments consisting of six tasks. The left figure shows our setup and the Drawer task. The Table shows the average completed stages out of 20 evaluations.

5.3 Real-World Main Results

We evaluate all the methods with 20 rollouts per task. As can be seen in Table 3, our proposed PMP policy consistently outperforms all baselines across all evaluated real-world tasks. Compared to the RGB-only policy, PMP achieves at least a 0.2-point improvement in accumulated scores, demonstrating the effectiveness of fusing both point-map and RGB modalities. Notably, on the Folding task, PMP increases the score from 2.1 to 2.5 using only 45 demonstrations, showcasing strong sample efficiency.

Interestingly, the PMP-xyz also outperforms the RGB-only baseline on several tasks, underscoring the value of spatial structure in guiding action prediction. However, its performance drops significantly in tasks involving deformable objects, such as Folding and Sweeping, where object geometry would change over actions. In these scenarios, the lack of appearance cues leads to coarse and less reliable action predictions. This is especially evident in Cup-Stacking, a task that explicitly requires reasoning about object color, further highlighting the importance of RGB input. Overall, these results validate the effectiveness and generalizability of PMP in handling diverse and challenging manipulation tasks in the real world.

6 Limitation and Future Work

Our current approach has two main limitations. First, simply concatenating the point map and RGB tokens may not optimally leverage the complementary information in each modality. More sophisticated fusion mechanisms could potentially extract richer cross-modal relationships and further improve performance. Second, our point map visual encoders are trained entirely from scratch, which constrains their performance compared to the RGB modality that benefits from ImageNet pretraining. For future work, developing pretraining objectives specifically designed for point map encoders represents a promising direction. Just as vision models benefit substantially from pretraining on large image datasets, establishing similar paradigms for point map representations could dramatically improve performance, enabling more robust geometric feature learning before fine-tuning on specific robotic tasks.

7 Conclusion

We present PointMapPolicy (PMP), a novel diffusion-based imitation learning framework that effectively integrates 3D geometric reasoning with standard vision techniques. By projecting depth pixels into a multi-channel image of XYZ coordinates, PMP leverages existing visual encoders, while an efficient xLSTM-based diffusion network denoises action tokens to generate precise control sequences. Empirical results on RoboCasa and CALVIN demonstrate that PMP not only achieves state-of-the-art performance but also offers significantly faster training and inference. Comprehensive ablations on observation modalities and fusion strategies further highlight the clear advantages of structured point-map representations. Looking forward, we plan to explore large-scale pretraining of point-map models to extend generalization across diverse robotic tasks.

8 Acknowledgments

This work was supported by funding from the pilot program Core Informatics of the Helmholtz Association (HGF). NS and GN were supported by the Carl Zeiss Foundation under the project JuBot (Jung Bleiben mit Robotern). Xiaogang Jia and Xinkai Jiang acknowledge the support from the China Scholarship Council (CSC). The authors also acknowledge support by the state of Baden-Württemberg through the HoreKa supercomputer funded by the Ministry of Science, Research and the Arts Baden-Württemberg, and by the German Federal Ministry of Education and Research.

References

- [1] Lingdong Chi, Aravind Rajeswaran Gupta, Silvio Savarese, Li Fei-Fei, and Nicholas Rhinehart. Diffusion policy: Visuomotor policy learning via action diffusion. In *International Conference on Learning Representations (ICLR)*, 2023.
- [2] Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, et al. π0: A vision-language-action flow model for general robot control, 2024. *URL https://arxiv. org/abs/2410.24164*.
- [3] Xiaogang Jia, Denis Blessing, Xinkai Jiang, Moritz Reuss, Atalay Donat, Rudolf Lioutikov, and Gerhard Neumann. Towards diverse behaviors: A benchmark for imitation learning with human demonstrations. *arXiv preprint arXiv:2402.14606*, 2024.
- [4] Haoyi Zhu, Yating Wang, Di Huang, Weicai Ye, Wanli Ouyang, and Tong He. Point cloud matters: Rethinking the impact of different observation spaces on robot learning. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024. URL https://openreview.net/forum?id=zgSnSZ0Re6.
- [5] Skand Peri, Iain Lee, Chanho Kim, Li Fuxin, Tucker Hermans, and Stefan Lee. Point cloud models improve visual robustness in robotic learners, 2024. URL https://arxiv.org/abs/2404.18926.
- [6] Balazs Gyenes, Nikolai Franke, Philipp Becker, and Gerhard Neumann. PointpatchRL masked reconstruction improves reinforcement learning on point clouds. In 8th Annual Conference on Robot Learning, 2024. URL https://openreview.net/forum?id=3jNEz3kUS1.
- [7] Yanjie Ze, Gu Zhang, Kangning Zhang, Chenyuan Hu, Muhan Wang, and Huazhe Xu. 3d diffusion policy: Generalizable visuomotor policy learning via simple 3d representations. *arXiv* preprint arXiv:2403.03954, 2024.
- [8] Eugenio Chisari, Nick Heppert, Max Argus, Tim Welschehold, Thomas Brox, and Abhinav Valada. Learning robotic manipulation policies from point clouds with conditional flow matching. In 8th Annual Conference on Robot Learning, 2024. URL https://openreview.net/forum?id=vtEn8NJWlz.
- [9] Zhou Xian, Nikolaos Gkanatsios, Theophile Gervet, Tsung-Wei Ke, and Katerina Fragkiadaki. Chaineddiffuser: Unifying trajectory diffusion and keypose prediction for robotic manipulation. In 7th Annual Conference on Robot Learning, 2023. URL https://openreview.net/forum?id=W0zgY2mBTA8.
- [10] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017.
- [11] Tsung-Wei Ke, Nikolaos Gkanatsios, and Katerina Fragkiadaki. 3d diffuser actor: Policy diffusion with 3d scene representations. *arXiv preprint arXiv:2402.10885*, 2024.
- [12] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20697–20709, June 2024.
- [13] Vincent Leroy, Yohann Cabon, and Jerome Revaud. Grounding image matching in 3d with mast3r, 2024.

- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. pages 770–778, 2016. URL https://openaccess.thecvf.com/content_cvpr_2016/html/He_Deep_Residual_Learning_CVPR_2016_paper.html.
- [15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=YicbFdNTTy.
- [16] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11976–11986, June 2022.
- [17] Yatian Pang, Wenxiao Wang, Francis E. H. Tay, Wei Liu, Yonghong Tian, and Li Yuan. Masked Autoencoders for Point Cloud Self-supervised Learning. In Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, *European Conference on Computer Vision 2022*, pages 604–621, Cham, 2022. Springer Nature Switzerland. ISBN 978-3-031-20086-1. doi: 10.1007/978-3-031-20086-1_35.
- [18] Guangyan Chen, Meiling Wang, Yi Yang, Kai Yu, Li Yuan, and Yufeng Yue. PointGPT: Autoregressively Generative Pre-training from Point Clouds, May 2023. URL http://arxiv.org/abs/2305.11487. arXiv:2305.11487 [cs].
- [19] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. URL https://openreview.net/forum?id=k7FuTOWMOc7.
- [20] Soroush Nasiriany, Abhiram Maddukuri, Lance Zhang, Adeet Parikh, Aaron Lo, Abhishek Joshi, Ajay Mandlekar, and Yuke Zhu. Robocasa: Large-scale simulation of everyday tasks for generalist robots. arXiv preprint arXiv:2406.02523, 2024.
- [21] Oier Mees, Daniel Merk, Yannik Bukschat, Harsh Lodhi, and Wolfram Burgard. Calvin: A benchmark for language-conditioned policy learning for long-horizon robot manipulation tasks. In *Conference on Robot Learning (CoRL)*, 2022.
- [22] Songming Liu, Lingxuan Wu, Bangguo Li, Hengkai Tan, Huayu Chen, Zhengyi Wang, Ke Xu, Hang Su, and Jun Zhu. RDT-1b: a diffusion foundation model for bimanual manipulation. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=yAzN4tz7oI.
- [23] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan P Foster, Pannag R Sanketi, Quan Vuong, Thomas Kollar, Benjamin Burchfiel, Russ Tedrake, Dorsa Sadigh, Sergey Levine, Percy Liang, and Chelsea Finn. Open-VLA: An open-source vision-language-action model. In 8th Annual Conference on Robot Learning, 2024. URL https://openreview.net/forum?id=ZMnD6QZAE6.
- [24] Moritz Reuss, Jyothish Pari, Pulkit Agrawal, and Rudolf Lioutikov. Efficient diffusion transformer policies with mixture of expert denoisers for multitask learning. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=nDmwloEl3N.
- [25] Octo Model Team, Dibya Ghosh, Homer Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Charles Xu, Jianlan Luo, Tobias Kreiman, You Liang Tan, Pannag Sanketi, Quan Vuong, Ted Xiao, Dorsa Sadigh, Chelsea Finn, and Sergey Levine. Octo: An open-source generalist robot policy. In *Proceedings of Robotics: Science and Systems*, Delft, Netherlands, 2024.
- [26] Stephen James, Kentaro Wada, Tristan Laidlow, and Andrew J. Davison. Coarse-to-Fine Q-attention: Efficient Learning for Visual Robotic Manipulation via Discretisation . In 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 13729–13738, Los Alamitos, CA, USA, June 2022. IEEE Computer Society. doi:

- 10.1109/CVPR52688.2022.01337. URL https://doi.ieeecomputersociety.org/10.1109/CVPR52688.2022.01337.
- [27] Mohit Shridhar, Lucas Manuelli, and Dieter Fox. Perceiver-actor: A multi-task transformer for robotic manipulation. In 6th Annual Conference on Robot Learning, 2022. URL https://openreview.net/forum?id=PS_eCS_WCvD.
- [28] Atalay Donat, Xiaogang Jia, Xi Huang, Aleksandar Taranovic, Denis Blessing, Ge Li, Hongyi Zhou, Hanyi Zhang, Rudolf Lioutikov, and Gerhard Neumann. Towards fusing point cloud and visual representations for imitation learning. In 7th Robot Learning Workshop: Towards Robots with Human-Level Abilities, 2025. URL https://openreview.net/forum?id=5cG7ilWX1V.
- [29] Ankit Goyal, Jie Xu, Yijie Guo, Valts Blukis, Yu-Wei Chao, and Dieter Fox. RVT: Robotic view transformer for 3d object manipulation. In 7th Annual Conference on Robot Learning, 2023. URL https://openreview.net/forum?id=0hPkttoGAf.
- [30] Ankit Goyal, Valts Blukis, Jie Xu, Yijie Guo, Yu-Wei Chao, and Dieter Fox. RVT-2: Learning precise manipulation from few demonstrations. In RSS 2024 Workshop on Geometric and Algebraic Structure in Robot Learning, 2024. URL https://openreview.net/forum?id=T6c4qoc6PG.
- [31] Mohit Shridhar, Lucas Manuelli, and Dieter Fox. Perceiver-actor: A multi-task transformer for robotic manipulation. In *Proceedings of the 6th Conference on Robot Learning (CoRL)*, 2022.
- [32] Ajay Mandlekar, Danfei Xu, Josiah Wong, Soroush Nasiriany, Chen Wang, Rohun Kulkarni, Li Fei-Fei, Silvio Savarese, Yuke Zhu, and Roberto Martín-Martín. What matters in learning from offline human demonstrations for robot manipulation. In 5th Annual Conference on Robot Learning, 2021. URL https://openreview.net/forum?id=JrsfBJtDFdI.
- [33] Nur Muhammad Mahi Shafiullah, Zichen Jeff Cui, Ariuntuya Altanzaya, and Lerrel Pinto. Behavior transformers: Cloning \$k\$ modes with one stone. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. URL https://openreview.net/forum?id=agTr-vRQsa.
- [34] Moritz Reuss, Maximilian Li, Xiaogang Jia, and Rudolf Lioutikov. Goal-conditioned imitation learning using score-based diffusion policies. In *Robotics: Science and Systems*, 2023. URL https://doi.org/10.15607/RSS.2023.XIX.028.
- [35] Homanga Bharadhwaj, Jay Vakil, Mohit Sharma, Abhinav Gupta, Shubham Tulsiani, and Vikash Kumar. Roboagent: Generalization and efficiency in robot manipulation via semantic augmentations and action chunking. In *First Workshop on Out-of-Distribution Generalization in Robotics at CoRL* 2023, 2023. URL https://openreview.net/forum?id=Pt5N3OG5wP.
- [36] Xiaogang Jia, Qian Wang, Atalay Donat, Bowen Xing, Ge Li, Hongyi Zhou, Onur Celik, Denis Blessing, Rudolf Lioutikov, and Gerhard Neumann. Mail: Improving imitation learning with mamba, 2024. URL https://arxiv.org/abs/2406.08234.
- [37] Xiaogang Jia, Atalay Donat, Xi Huang, Xuan Zhao, Denis Blessing, Hongyi Zhou, Hanyi Zhang, Han A. Wang, Qian Wang, Rudolf Lioutikov, and Gerhard Neumann. X-IL: Exploring the design space of imitation learning policies. In 7th Robot Learning Workshop: Towards Robots with Human-Level Abilities, 2025. URL https://openreview.net/forum?id=8w00JsTaFz.
- [38] Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces, 2024. URL https://openreview.net/forum?id=AL1fq05o7H.
- [39] Maximilian Beck, Korbinian Pöppel, Markus Spanring, Andreas Auer, Oleksandra Prudnikova, Michael K Kopp, Günter Klambauer, Johannes Brandstetter, and Sepp Hochreiter. xLSTM: Extended long short-term memory. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://openreview.net/forum?id=ARAxPPIAhq.
- [40] Tony Z Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual manipulation with low-cost hardware. *arXiv preprint arXiv:2304.13705*, 2023.

- [41] Moritz Reuss, Maximilian Li, Xiaogang Jia, and Rudolf Lioutikov. Goal conditioned imitation learning using score-based diffusion policies. In *Proceedings of Robotics: Science and Systems (RSS)*, 2023.
- [42] Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural Computation*, 23(7):1661–1674, 2011. doi: 10.1162/NECO_a_00142.
- [43] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *ICLR*, 2021.
- [44] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR, 2021.
- [45] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.
- [46] Johan Bjorck, Fernando Castañeda, Nikita Cherniadev, Xingye Da, Runyu Ding, Linxi Fan, Yu Fang, Dieter Fox, Fengyuan Hu, Spencer Huang, et al. Gr00t n1: An open foundation model for generalist humanoid robots. *arXiv preprint arXiv:2503.14734*, 2025.
- [47] Xinyang Bu, Fan Yang, Vikash Kumar, Stephan Mandt, and Dieter Fox. Clover: Closed-loop visual planning for long-horizon manipulation with diffusion policies. *arXiv preprint arXiv:2409.09016*, 2024.
- [48] Moritz Reuss, Ömer Erdinç Yağmurlu, Fabian Wenzel, and Rudolf Lioutikov. Multimodal diffusion transformer: Learning versatile behavior from multimodal goals. In *Robotics: Science and Systems*, 2024.
- [49] Yilun Zhou, Yufei Wu, De-An Huang, Yilun Zhang, Vivek Ganapathi, Yao Lu, and Josh Tenenbaum. Seer: Pretraining vision-language-action models on robot play. *arXiv* preprint *arXiv*:2402.12412, 2024.
- [50] Xinghang Li, Minghuan Liu, Hanbo Zhang, Cunjun Yu, Jie Xu, Hongtao Wu, Chilam Cheang, Ya Jing, Weinan Zhang, Huaping Liu, et al. Vision-language foundation models as effective robot imitators. In *International Conference on Learning Representations*, 2024.
- [51] Andrew Black et al. Zero-shot skill composition with semantic feature fusion. *arXiv preprint arXiv:2310.08573*, 2023.
- [52] Hongtao Wu, Ya Jing, Chilam Cheang, Guangzeng Chen, Jiafeng Xu, Xinghang Li, Minghuan Liu, Hang Li, and Tao Kong. Unleashing large-scale video generative pre-training for visual robot manipulation. In *International Conference on Learning Representations*, 2024.
- [53] Qingwen Bu, Hongyang Li, Li Chen, Jisong Cai, Jia Zeng, Heming Cui, Maoqing Yao, and Yu Qiao. Towards synergistic, generalized, and efficient dual-system for robotic manipulation. *arXiv* preprint arXiv:2410.08001, 2024.
- [54] Aditya Chattopadhay, Anirban Sarkar, Prantik Howlader, and Vineeth N Balasubramanian. Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. In 2018 IEEE Winter Conference on Applications of Computer Vision (WACV), pages 839–847, 2018. doi: 10.1109/WACV.2018.00097.
- [55] Oier Mees, Harsh Lodhi, Daniel Merk, Yannik Bukschat, and Wolfram Burgard. What matters in language conditioned robotic imitation learning over long horizons. *arXiv* preprint *arXiv*:2211.08853, 2022.
- [56] Shuang Zhang, Ilija Radosavovic, Kaiming He, Angel Yao, Dinesh Misra, and Deepak Pathak. Lad: Language as an abstraction for diffusion policy. In *Conference on Robot Learning (CoRL)*, 2022.

- [57] Duy Ha, Ajay Mandlekar, Andy Zeng, Yunzhu Chai, Jacky Liang, Yao Li, Kelvin Wang, Yuan Tian, Xinyang Zeng, Tianhao Xiao, et al. Scaling vision-language models for robot learning with imitation transformers. In *Conference on Robot Learning (CoRL)*, 2023.
- [58] Zichen Liu, Rakibul Hoque, Kevin Zakka, Shuran Song, and Pete Florence. Libero: Benchmarking knowledge transfer for language-instructed robotic manipulation. In *Conference on Robot Learning (CoRL)*, 2023.
- [59] Yufei Wu, Kiran Alwala, Vivek Ganapathi, Sudeep Sharma, Yilun Chang, Yicheng Zhang, Yilun Zhou, et al. Susie: Scaling up instruction-following policies for robot manipulation. *arXiv* preprint arXiv:2402.17552, 2024.
- [60] Yilun Zhou, Yen-Chen Lin, De-An Huang, Yao Lu, and Josh Tenenbaum. Mode: Mixture of diffusion experts for long-horizon manipulation. In *Conference on Robot Learning (CoRL)*, 2024.

NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [NA].
- [NA] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "[No]" or "[NA]" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- Delete this instruction block, but keep the section heading "NeurIPS Paper Checklist",
- Keep the checklist subsection headings, questions/answers and guidelines below.
- Do not modify the questions and only use the provided macros for your answers.

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: Our main contribution is a novel and efficient 3D representation. This is clearly outlined in the abstract and expanded in the introduction and method sections. Claims about performance are verified through experiments on simulation and real robot evaluations.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals
 are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We discuss the limitations of this work in Sec. 6.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: This paper does not include theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We describe the baselines, evaluation details of simulation benchmarks in A, the real world experiment in B, and corresponding hyperparameters to reproduce the experiment results in the Appdendix. C.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
- (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We will open source the codes in the near future once they are cleaned up and anonymity is not a concern anymore. All the experiments we conducted were using open-source datasets. In the experiments section and appendix we provide information to get access to the data.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.

- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We list the hyperparameter with the optimizer for each of the algorithms in Tab. 6, and how they were chosen is in the appendix C.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental
 material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: The experiment of RoboCase provides confidence intervals.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error
 of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: The compute resources are described in Appendix. E

Guidelines:

• The answer NA means that the paper does not include experiments.

- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]
Justification: Yes
Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a
 deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: There is no societal impact of the work performed.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: This paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with
 necessary safeguards to allow for controlled use of the model, for example by requiring
 that users adhere to usage guidelines or restrictions to access the model or implementing
 safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do
 not require this, but we encourage authors to take this into account and make a best
 faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: Yes, we do used pretrained models and open source code base for baselines, which is clearly stated in both experiment section and appendix.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the
 package should be provided. For popular datasets, paperswithcode.com/datasets
 has curated licenses for some datasets. Their licensing guide can help determine the
 license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: We plan to open source the code in the future.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: This paper does not involve crowdsourcing nor research with human subjects. Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: This paper does not involve crowdsourcing nor research with human subjects. Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core method development of this paper does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.

A Simulation Experiment Details

A.1 RoboCasa Benchmark

RoboCasa [20] is a large-scale simulation framework developed to train generalist robots in realistic and diverse home environments, with a particular focus on kitchen scenarios. The benchmark comprises 100 tasks, including 25 atomic tasks with 50 human demonstrations and 75 composite tasks with auto-generated demonstrations. These tasks are centered around eight fundamental robotic skills relevant to real-world home environments: (1) pick-and-place, (2) opening and closing doors, (3) opening and closing drawers, (4) twisting knobs, (5) turning levers, (6) pressing buttons, (7) insertion, and (8) navigation.

To comprehensively evaluate our method, we selected five tasks from the atomic tasks described in Table 4, each representing a distinct skill.

Task Name	Description				
Pick-and-Place Tasks					
PickPlace_Counter_To_Microwave	Pick an object from the counter and place it inside the microwave (door is open)				
PickPlace_Counter_To_Sink	Pick an object from the counter and place it in the sink.				
PickPlace_Microwave_To_Counter	Pick an object from the microwave and place it on the counter (door is open).				
PickPlace_Sink_To_Counter	Pick an object from the sink and place it on the counter next to the sink.				
Drawer Tasks					
Open_Drawer	Open a drawer.				
Close_Drawer	Close a drawer.				
Stove Tasks					
Stove_On	Turn on a specific stove burner by twisting its knob.				
Stove_Off	Turn off a specific stove burner by twisting its knob.				
Sink Tasks					
SinkFaucet_On	Turn on the sink faucet to start water flow.				
SinkFaucet_Off	Turn off the sink faucet to stop water flow.				
Turn_Sink_Spout	Rotate the sink spout.				
Coffee Machine Tasks					
Coffee_Press_Button	Press the button to pour coffee into the mug.				
Coffee_Setup_Mug	Place the mug into the coffee machine's mug holder.				
Coffee_Serve_Mug	Remove the mug from the holder and place it on the counter.				
Microwave Tasks					
Microwave_On	Start the microwave by pressing the start button.				
Microwave_Off	Stop the microwave by pressing the stop button.				

Table 4: RoboCasa task set.

A.2 CALVIN Benchmark

Benchmark Setup. The CALVIN benchmark [21] is a long-horizon manipulation benchmark featuring four visually distinct tabletop environments (A–D), each containing a common set of objects and 34 manipulation tasks. Agents are given natural language instructions describing sequences of up to 5 tasks to be executed in order. The primary evaluation involves completing 1000 such instruction chains in environment D. Agents are scored by the number of tasks successfully completed per chain, with a maximum rollout length of 5.

Evaluation Protocol. We evaluate PointMapPolicy on one standard CALVIN settings: $ABC \rightarrow D$, where the policy is trained on environments A, B, and C, and evaluated zero-shot on D. Only 1% of the play data is paired with language, requiring models to learn from primarily unlabeled data. The $ABC \rightarrow D$ setup tests visual and environmental generalization, while $D \rightarrow D$ emphasizes efficiency in low-resource, language-scarce settings.

Baselines. We compare against a broad set of state-of-the-art language-conditioned policies spanning imitation, diffusion, and foundation-model-based architectures:

• HULC [55]: a hierarchical imitation learning framework that combines a discrete skill planner (learned via VAE) with a low-level controller trained on dense teleoperation data. It does not use language pretraining or frozen encoders.

- LAD [56]: extends HULC by replacing the planner with a high-level diffusion model that learns to generate latent skills conditioned on goals. The policy uses a U-Net-based diffusion process and CLIP-encoded language inputs.
- **Distill-D** [57]: a CLIP-conditioned continuous-time diffusion policy. It distills robot behaviors from demonstrations using dense supervision and large batch sizes, incorporating CLIP goals and continuous action prediction.
- MT-ACT [58]: a transformer-based policy that predicts temporally extended action chunks.
 It encodes entire trajectories using VAE-style representation learning and uses a transformer encoder-decoder for action generation.
- **RoboFlamingo** [50]: based on OpenFlamingo, this model integrates a frozen VLM with a lightweight policy head. It is pretrained on large-scale vision-language data and finetuned on CALVIN using supervised behavior cloning.
- **SuSIE** [59]: a scalable instruction-following diffusion policy. It is pre-trained on curated robot demonstrations and uses an instruction-conditioned denoising process with significant offline finetuning.
- GR-1 [52]: a powerful decoder-only transformer trained on large-scale synthetic video data.
 The model is capable of generating long sequences of actions and is finetuned on CALVIN for grounding.
- CLOVER [47]: a video diffusion planner that predicts intermediate visual goals via video generation and closes the loop using low-level policy feedback. It does not require internetscale pretraining and achieves strong multi-step rollout success.
- **MoDE** [60]: Mixture-of-Diffusion-Experts model with sparse routing. It supports both small (non-pretrained) and large (pretrained) variants. The pretrained variant achieves top performance while maintaining low inference cost.
- Seer / Seer-Large [49]: large-scale transformer models pretrained on 1000+ hours of robot play data. Seer incorporates language, vision, and action streams into a unified transformer and achieves strong generalization, particularly when scaled up.

Results. Table 2 presents results on ABC \rightarrow D settings. In the challenging ABC \rightarrow D generalization setting, PointMapPolicy surpasses earlier methods such as HULC and Distill-D and performs competitively with MoDE and Seer. Unlike these baselines, our method uses no pretraining, no frozen vision-language encoders, and no large-scale robot datasets. These results demonstrate the effectiveness of using grounded spatial goal representations and point-based reactive policies for generalizable, long-horizon language-conditioned control.

A.3 CALVIN Evaluation Details

CALVIN Task Suite. The CALVIN benchmark [21] consists of four environments (A–D) with identical object sets and tasks but differing in visual appearance and lighting. Each episode comprises natural language instructions describing a chain of 5 atomic manipulation goals selected from 34 total tasks. These include pushing, pulling, opening drawers, pressing buttons, etc. The robot receives RGB-D observations from multiple camera views and proprioceptive information. The action space is 7-dimensional, including end-effector translation, rotation, and gripper control.

Language-Conditioned Evaluation. Agents are evaluated on their ability to execute long-horizon natural language commands consisting of 5 consecutive tasks. The evaluation split uses 1000 unique instruction chains in environment D. The model receives a reward of 1 upon successful completion of each task in the chain, and evaluation metrics report the average number of tasks completed per rollout over 3 random seeds.

ABC→**D Setup.** To evaluate generalization, the model is trained on environments A, B, and C, and tested zero-shot on environment D. No data from environment D is used during training. This split emphasizes robustness to domain shifts and visual changes.

Drawer	Stack	Fold	Score
Open the upper/lower drawer	Pick the correct cup	Pick the towel	1
Pick the object and place it into drawer	Stack failed	Fold the towel	2
Close the Drawer	Stack the cups	Fold the towel perfectly	3
-	_	_	4
Pour	Sweep	Arrange	Score
Pick the cup	Pick the broom	Open the mixer machine	1
Pour the contents	Sweep 20% of garbage	Pick the container and place on pad	2
Pour all contents into container	Sweep 50% of garbage	Close the mixer machine	3
Put the cup back	Sweep all garbage	_	4

Table 5: Task score metric details of real robot and evaluation standards.

B Real World Experiment Details

We conducted six real-world experiments on a Franka Panda Robot: Drawer, Stack, Pour, Sweep, Fold, and Arrange.

B.1 Task Metric

Given the complexity and long-horizon nature of the tasks, we decompose each task into several discrete stages. The final score is computed as the total number of successfully completed stages. Details of the scoring metric design are provided in Table 5.

B.2 Task Description

Drawer: In the Drawer task, there is a cabinet with two drawers and two different objects, a cube and a cylinder. The robot must follow a language-specified instruction to open the designated drawer, pick up the target object, place it inside the drawer, and then close the drawer. The key challenges involve handling the random initialization of both the cabinet's position and the objects' locations.

Stack: In the Stack task, four cups of different colors and sizes are provided. The robot must stack the cups in a specific order based on their colors. The main challenges lie in accurately recalling the stacking sequence and executing precise placement, as the cups are closely sized and must fit together properly.

Pour: In the Pour task, three distinct cups and three different containers are placed in randomized initial positions. The robot must generalize to novel object configurations while maintaining the precision necessary to pour the contents from the cups into the containers without spilling. The primary challenge lies in adapting to varying spatial arrangements while executing controlled and accurate pouring motions.

Sweep: Unlike standard Pick-and-Place tasks, this task requires the robot to acquire a novel sweeping skill. In the Sweep task, the positions of the broom, dustpan, and trash vary across trials, and even the number of trash items changes. The key challenge is manipulating deformable trash materials that differ from those encountered during training, requiring the policy to exhibit strong generalization and adaptability.

Fold: The Fold task requires precise manipulation skills. The goal is to neatly fold a towel that is randomly oriented at the start of each trial. The primary challenge lies in accurately handling the soft, deformable material to achieve a clean and consistent fold despite varying initial conditions.

Arrange: In the Arrange task, the setup includes a mixing machine and a container. The robot must follow a specific sequence: first, open the mixing machine; next, place the container on the designated pad; and finally, close the machine. This task primarily emphasizes long-horizon planning, requiring the robot to execute a multi-step procedure in the correct order.

C Hyper Parameters

Hyperparameter	CALVIN ABC	RoboCasa	Real World	
Number of x-Blocks	10	8	6	
Attention Heads	8	8	8	
Action Chunk Size	10	10	10	
History Length	1	1	1	
Embedding Dimension	2048	768	2048	
Image Encoder	FiLM-ResNet50	ConvNextV2	FiLM-ResNet50	
Goal Lang Encoder	CLIP ViT-B/32	CLIP ViT-B/32	CLIP ViT-B/32	
Attention Dropout	0.3	0.3	0.3	
Residual Dropout	0.1	0.1	0.1	
MLP Dropout	0.1	0.1	0.1	
Optimizer	AdamW	AdamW	AdamW	
Betas	[0.9, 0.95]	[0.9, 0.95]	[0.9, 0.95]	
Learning Rate	1e-4	1e-4	1e-4	
Transformer Weight Decay	0.05	0.05	0.05	
Other weight decay	0.05	0.05	0.05	
Batch Size	128	128	128	
Train Steps in Thousands	25	15	30	
$\sigma_{ m max}$	80	80	80	
$\sigma_{ m min}$	0.001	0.001	0.001	
σ_t	0.5	0.5	0.5	
EMA	True	True	True	
Time steps	Exponential	Exponential	Exponential	
Sampler	DDIM	DDIM	DDIM	
Trainable Parameters (Millions)	147	111	96	

Table 6: Summary of all the Hyperparameters for our experiments.

We export all the hyper parameters across three benchmarks for reproduction.

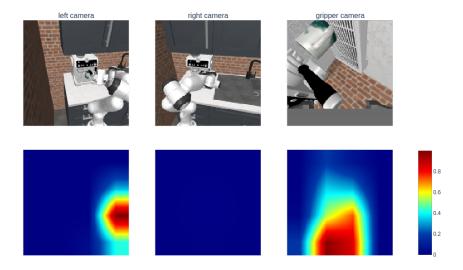
D Activation Map Analysis

To gain qualitative insights into what regions the visual encoders attend to during action inference, we visualize activation maps using Grad-CAM++ [54]. Unlike classification tasks, our diffusion-based policy does not predict discrete categories, therefore, we apply Grad-CAM++ using the diffusion loss as the target signal, following the approach of highlighting input regions that most influence the denoised trajectory prediction. We generate the heatmaps using the Grad-CAM++ implementation³, and compute activations for each camera view across three RoboCasa tasks: OpenDrawer, Turn On Sink Faucet, and Coffee Serve Mug. In all figures, we use a ConvNeXtv2 encoder and extract Grad-CAM++ heatmaps from the final convolutional block before normalization. Each visualization consists of six images arranged in two rows. The top row shows the raw visual input (RGB or XYZ visualized in color), and the bottom row displays the corresponding Grad-CAM++ heatmaps for each of the three camera views: static left, static right, and wrist-mounted. These maps highlight spatial regions with the greatest impact on predicted actions.

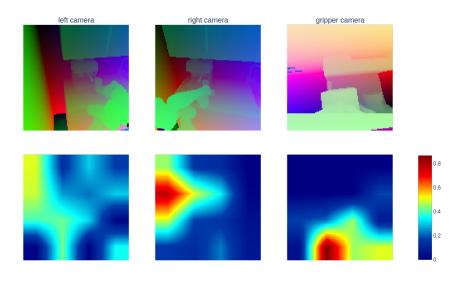
Overall, the attention patterns are consistent with task-relevant visual cues. For example, activations commonly focus on the robot gripper, the manipulated object, or the goal location, depending on the modality and perspective.

Figures 8, 9 and 10 show results for the RoboCasa tasks Coffee Serve Mug, Open Drawer and Turn On Sink Faucet, respectively.

³https://github.com/jacobgil/pytorch-grad-cam

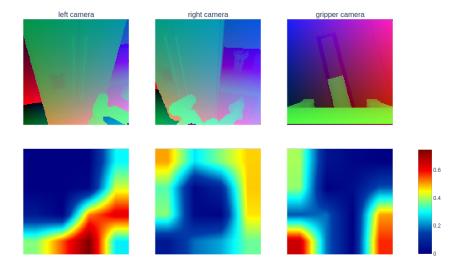


(a) RGB-only ConvNeXtv2 encoder. Top: raw 128x128 RGB input frames provided to the agent. Bottom: Grad-CAM++ heatmaps from the final convolutional layer.

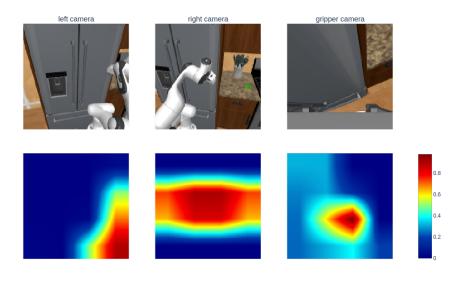


(b) PMP-xyz ConvNeXtv2 encoder. Top: $128x128\ XYZ$ input visualized as RGB. Bottom: Grad-CAM++ heatmaps from the final convolutional layer.

Figure 8: Raw RGB, XYZ input frames and Grad-CAM++ activations on the Coffee Serve Mug RoboCasa task for RGB-only and Point-map-only ConvNeXtv2 visual encoders.

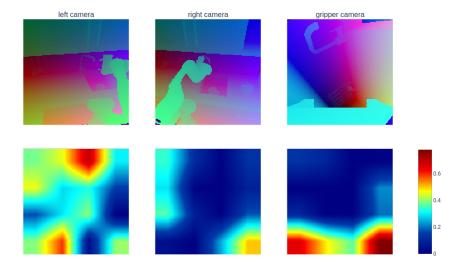


(a) RGB-only ConvNeXtv2 encoder. Top: raw 128x128 RGB input frames provided to the agent. Bottom: Grad-CAM++ heatmaps from the final convolutional layer.

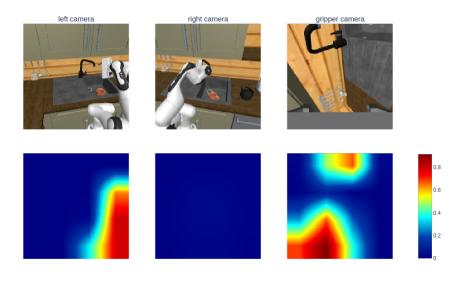


(b) PMP-xyz ConvNeXtv2 encoder. Top: XYZ input visualized as RGB. Bottom: Grad-CAM++ heatmaps from the final convolutional layer.

Figure 9: Raw RGB, XYZ input frames and Grad-CAM++ activations on the Open Drawer RoboCasa task for RGB-only and Point-map-only ConvNeXtv2 visual encoders.



(a) RGB-only ConvNeXtv2 encoder. Top: raw 128x128 RGB input frames provided to the agent. Bottom: Grad-CAM++ heatmaps from the final convolutional layer.



(b) PMP-xyz ConvNeXtv2 encoder. Top: XYZ input visualized as RGB. Bottom: Grad-CAM++ heatmaps from the final convolutional layer.

Figure 10: Raw RGB, XYZ input frames and Grad-CAM++ activations on the Turn On Sink Faucet RoboCasa task for RGB-only and Point-map-only ConvNeXtv2 visual encoders.

E Compute Resources

For the CALVIN experiments, PMP-Cat employs Film-ResNet50 as encoders for both RGB images and point maps, with 10 x-Blocks as backbones (512 latent dimensions), totaling 147M trainable parameters. Training utilizes 4 Nvidia RTX 6000 Ada GPUs with 128 samples per GPU (512 total batch size). Each epoch completes in approximately 13 minutes, allowing full training (25 epochs) in under 6 hours, excluding evaluation time.

For the RoboCasa experiments, PMP-Cat employs ConvNeXtv2 as encoders with 8 x-Blocks using 512 latent dimensions. Training utilizes 1 NVIDIA A100-SXM4-40GB with a 128 batch size.

For the real-robot experiments, PMP-Cat employs Film-ResNet50 as encoders for both images and point maps, with 6 x-Blocks using 256 latent dimensions. Training utilizes 1 Nvidia RTX 6000 Ada GPUs with 128 batch size.