

SpatialActor: Exploring Disentangled Spatial Representations for Robust Robotic Manipulation

Hao Shi^{1*}, Bin Xie², Yingfei Liu², Yang Yue¹, Tiancai Wang²,
Haoqiang Fan², Xiangyu Zhang^{3,4}, Gao Huang^{1†}

¹Department of Automation, BNRist, Tsinghua University

²Dexmal

³MEGVII Technology

⁴StepFun

shi-h23@mails.tsinghua.edu.cn, wtc@dexmal.com, gaohuang@tsinghua.edu.cn

Abstract

Robotic manipulation requires precise spatial understanding to interact with objects in the real world. Point-based methods suffer from sparse sampling, leading to the loss of fine-grained semantics. Image-based methods typically feed RGB and depth into 2D backbones pre-trained on 3D auxiliary tasks, but their entangled semantics and geometry are sensitive to inherent depth noise in real-world that disrupts semantic understanding. Moreover, these methods focus on high-level geometry while overlooking low-level spatial cues essential for precise interaction. We propose *SpatialActor*, a disentangled framework for robust robotic manipulation that explicitly decouples semantics and geometry. The Semantic-guided Geometric Module adaptively fuses two complementary geometry from noisy depth and semantic-guided expert priors. Also, a Spatial Transformer leverages low-level spatial cues for accurate 2D-3D mapping and enables interaction among spatial features. We evaluate *SpatialActor* on multiple simulation and real-world scenarios across 50+ tasks. It achieves state-of-the-art performance with 87.4% on RL-Bench and improves by 13.9% to 19.4% under varying noisy conditions, showing strong robustness. Moreover, it significantly enhances few-shot generalization to new tasks and maintains robustness under various spatial perturbations.

Project Page — <https://shihao1895.github.io/SpatialActor>

Code — <https://github.com/shihao1895/SpatialActor>

1 Introduction

Robotic manipulation enables robots to understand scenes and interact with objects to perform precise physical tasks in the real-world environments. Some existing methods (Zeng et al. 2021; Zhao et al. 2023; Brohan et al. 2022; Kim et al. 2024; Chi et al. 2023; Liu et al. 2024a; Shi et al. 2025) rely solely on 2D visual inputs to predict end-effector actions in 3D space, however, they often struggle in scenarios requiring spatial reasoning, occlusion handling, geometric shape comprehension, or fine-grained object interactions due to

their limited understanding of spatial geometry. Given that real-world tasks inherently occur in 3D space, incorporating 3D spatial information is crucial for learning robust and generalizable robotic manipulation policies.

Recent efforts in robotic manipulation have explored various approaches to exploit spatial information. In Fig. 1 (a), point cloud-based approaches (Zhang et al. 2023; Chen et al. 2023; Ze et al. 2024; James et al. 2022) represent 3D geometry explicitly, yet suffer from semantic loss due to sparse sampling and are limited by the high cost of 3D annotations, which constrains pretraining scalability. In contrast, Fig. 1 (b) illustrates image-based methods (Goyal et al. 2023, 2024; Fang et al. 2025; Wang et al. 2024) that utilize multi-view RGB-D to jointly model semantics and geometry in a shared feature space. These methods exploit structured 2D inputs to obtain dense semantics and benefit from strong 2D pretrained priors, enabling competitive performance. However, the entanglement of semantics and geometry makes these methods sensitive to inherent depth noise in the real-world, which degrades semantic and geometric understanding. As shown in Fig. 1 (d), even minor noise can lead to a significant performance drop of 8.9% in RVT2 (Goyal et al. 2024). In reality, depth is often compromised by sensor noise, lighting variations, and surface reflections, which severely limit the practical application of such methods in the real-world. Furthermore, the joint modeling primarily retains high-level geometry while neglecting low-level spatial cues that are critical for precise interaction by providing fine-grained 2D-3D correspondences.

The limitations above call for three critical capabilities in robotic manipulation: 1) fine-grained spatial understanding to enable accurate control; 2) robustness to sensor noise to ensure real-world reliability; and 3) low-level spatial cues to support consistent spatial tokens interaction. This raises a fundamental question: How can we construct a robust spatial representation that fulfills these requirements?

To address this, we propose *SpatialActor*, a novel framework for robust spatial representation in robotic manipulation. Instead of a shared feature space, we decouple semantics and geometry to mitigate cross-modal interference. Furthermore, we decompose geometric information into high-

*Work done during internship at Dexmal.

†Corresponding author: Gao Huang.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

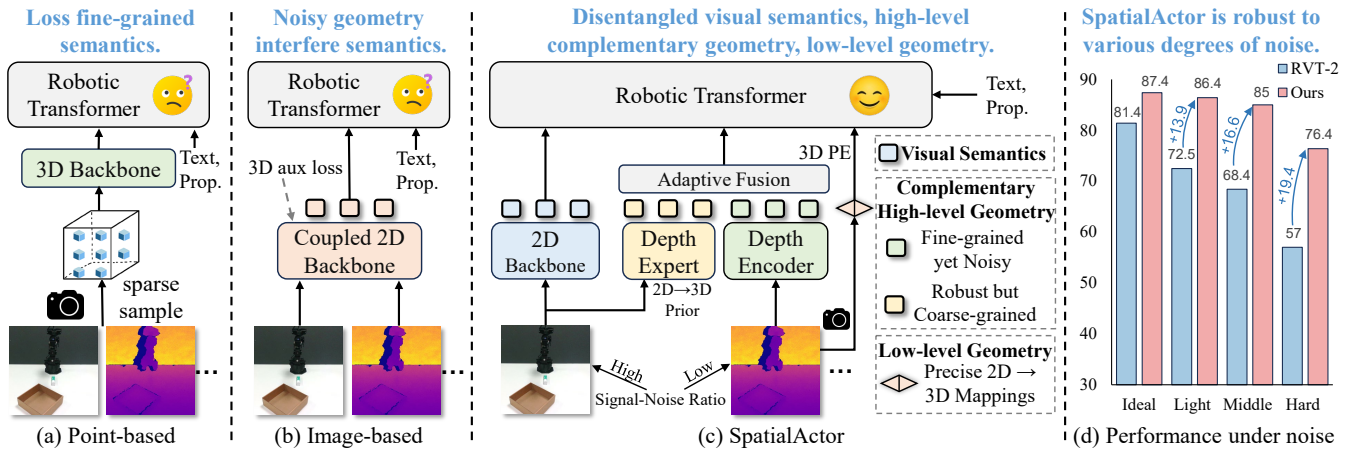


Figure 1: **Methodology comparisons.** (a) Point-based methods suffer from sparse sampling, leading to the loss of fine-grained semantics. (b) Image-based methods typically entangle semantics and geometry, while inherent depth noise in real-world disrupts semantic understanding. (c) SpatialActor disentangle visual semantics, two complementary high-level geometry from noisy depth and expert priors, low-level spatial cues. (d) Performance under various degrees of noise, showing the robustness.

level geometric representations and low-level spatial cues. To construct a robust high-level geometric representation, we propose a Semantic-guided Geometric Module (SGM). Within the SGM, high signal-to-noise semantics from RGB are processed by a large-scale pretrained depth estimation expert (Yang et al. 2024, 2025) to produce a robust but coarse geometric prior. Meanwhile, raw depth inputs retain fine-grained geometric details but are inherently noisy. By adaptively integrating these complementary geometric representations through a gating mechanism, the SGM enhances both robustness and spatial precision, effectively addressing the limitations of individual modalities. For low-level positional cues, we introduce a Spatial Transformer (SPT) that integrates spatial modeling into the transformer layers. By employing spatial position encoding, distinct spatial tokens are endowed with unique spatial indices, facilitating spatial interactions. The model performs view-level interaction to refine token relationships within each view, followed by scene-level interaction that unifies cross-modal cues across the scene, yielding features for the action head.

To comprehensively evaluate our method, SpatialActor, we conduct experiments on 50+ robotic manipulation tasks in both simulation and real-world. In RL Bench (18 tasks with 249 variations), SpatialActor achieves an 87.4% average success rate, surpassing state-of-the-art methods by approximately 6.0%, with a notable 53.3% improvement in high-precision spatial tasks like Insert Peg. Our method also shows strong robustness, maintaining higher success rates under noise conditions with improvements of 13.9%, 16.9%, and 19.4% at light, medium, and heavy noise levels, respectively. On ColosseumBench, which evaluates 20 tasks under spatial perturbations, SpatialActor consistently outperforms baselines, showcasing superior spatial generalization. Additionally, in a few-shot setting, adapting a multi-task pretrained model to 19 novel tasks with only 10 demonstrations per task, SpatialActor achieves 79.2% success compared to 46.9% for RVT-2. Real-world experiments further validate

these results, as SpatialActor outperforms RVT-2 across 8 tasks and 15 variations, demonstrating its strong robustness and generalization across diverse scenarios.

2 Related Works

2.1 Representation Learning for Manipulation

Early methods relied on proprioceptive sensing (Deng et al. 2020; Andrychowicz et al. 2020), which limited their generalization. With the rise of large-scale visual pretraining, many 2D-based approaches (Nair et al. 2022; Chi et al. 2023; Zhao et al. 2023; Yue et al. 2025; Zeng et al. 2024; Zhong et al. 2025; Xie et al. 2025) leverage strong visual priors to extract semantics. However, they often lack 3D spatial understanding, limiting their effectiveness in precise manipulation. Point cloud-based methods (Fang et al. 2023; Chen et al. 2023; Jia et al. 2024; Ze et al. 2024; Zhang et al. 2023; Sun et al. 2025) capture explicit 3D structures, offering geometry but are hampered by sparsity. Voxel-based representations (Shridhar, Manuelli, and Fox 2023; James et al. 2022) reduce sparsity by discretizing space for structured reasoning, yet they incur high computational costs. Multi-view RGB-D approaches (Goyal et al. 2023, 2024; Zhang et al. 2024; Fang et al. 2025; Wang et al. 2024; Seo et al. 2023) integrate dense 2D semantics with geometry via early fusion or auxiliary supervision, yet such shared feature spaces remain vulnerable to sensor noise and often lack precise spatial corresponding for fine-grained interaction. To address these limitations, we decouple semantics and geometry, and construct geometric representations by fusing complementary high-level expert priors and raw depth together with low-level spatial cues for precise manipulation.

2.2 Vision Foundation Models for Robotics

Vision foundation models have significantly enhanced robotic perception by incorporating semantic and geometric priors. Visual and multimodal models (Radford et al. 2021;

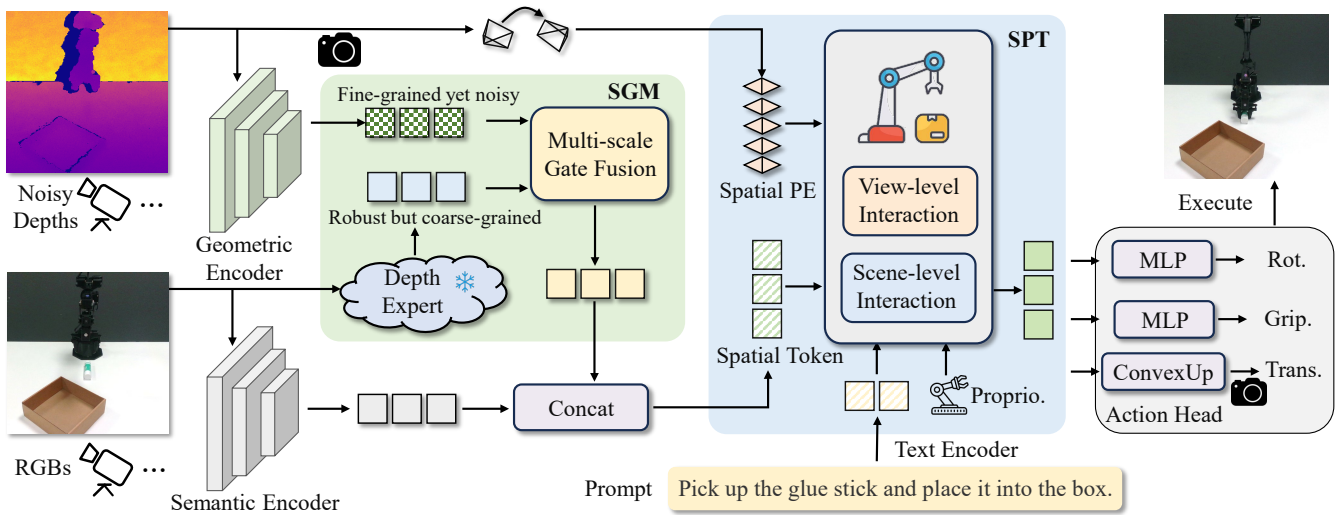


Figure 2: **Overall framework of SpatialActor.** The architecture employs separate vision and depth encoders. Semantic-guided Geometric Module (SGM) adaptively fuses robust yet coarse geometric priors from a pretrained depth expert with noisy depth features via gated fusion to yield high-level geometric representations. In the Spatial Transformer (SPT), low-level spatial cues are encoded as positional embeddings to drive spatial interactions. Finally, view-level interactions refine intra-view features, while scene-level interactions consolidate cross-modal information across views to support the subsequent action head.

Li et al. 2022; Feng et al. 2023; Liu et al. 2024b; Wang et al. 2025b; Wu et al. 2025) leverage diverse datasets to learn strong semantic priors that improve visual understanding, which benefits downstream robotic tasks. However, they focus on the 2D domain and lack spatial understanding capabilities. 3D vision models (Zhu et al. 2024; Zheng et al. 2024; Qian et al. 2022; Zheng et al. 2025; Kang et al. 2024; Zhang et al. 2025) integrate semantic information with explicit spatial structures to facilitate effective geometric perception. However, the acquisition and annotation of 3D data are inherently expensive and labor-intensive, which restricts scalability and limits their application in real-world scenarios. Depth estimation experts (Yang et al. 2024, 2025; Bhat et al. 2023; Wang et al. 2025a) leverage large-scale pretraining on diverse datasets to translate semantics in images into corresponding geometric structures, robustly inferring geometric information even under challenging conditions such as sensor noise and occlusions. In this paper, we leverage the strong semantic alignment of vision models together with robust geometric priors from depth estimation experts.

3 Method

3.1 Overall Framework

Fig. 2 illustrates the overall framework of our approach. The inputs to the robot’s control system are given by

$$X = \{I^v, D^v\}_{v=1}^V, P, L, \quad (1)$$

where $I^v \in \mathbb{R}^{H \times W \times 3}$ and $D^v \in \mathbb{R}^{H \times W}$ denote the RGB image and depth map for view v (with V views in total), $P \in \mathbb{R}^{d_p}$ represents the robot’s proprioceptive state (d_p indicating its dimension), and L denotes the language prompt.

For each view v , the RGB images and noisy depth maps are processed separately. The images I^v and the language

instruction L are fed into a vision-language model (e.g., CLIP (Radford et al. 2021)) to extract semantic features F_{sem}^v and text features F_{text} . Meanwhile, raw depth maps D^v are processed by a depth encoder to yield fine-grained but noisy geometric features F_{geo}^v . Subsequently, F_{geo}^v is enhanced via a Semantic-guided Geometric Module (SGM). In SGM, large-scale pre-trained depth estimation expert is employed to obtain robust yet coarse geometric priors \hat{F}_{geo}^v . A multi-scale gated fusion module then adaptively fuses F_{geo}^v with \hat{F}_{geo}^v to produce refined geometric features $F_{\text{fuse-geo}}^v$, preserving details while reducing noise, which are concatenated with F_{sem}^v to form the final spatial representation H^v .

We further introduce a Spatial Transformer (SPT). Within the SPT, intrinsic and extrinsic parameters, along with depth values, are used to construct a spatial encoding that captures the low-level spatial cues between spatial tokens. The SPT first applies view-level interaction to consolidate intra-view context, followed by scene-level cross-modal interaction to aggregate cross-modal cues into a unified scene representation. Finally, an action head predicts the robot’s 3D end-effector pose and gripper state.

3.2 Semantic-guided Geometric Module

Real-world depth measurements are often noisy due to sensor limitations and environmental interference, whereas RGB images provide high signal-to-noise semantic cues. Large-scale pretrained depth estimation models (e.g., Depth Anything (Yang et al. 2024, 2025)) learn a smooth semantic-to-geometric mapping, offering robust and generalizable geometric priors. In contrast, raw depth features retain fine-grained, pixel-level details but are highly sensitive to noise.

To leverage these complementary strengths, we extract robust yet coarse-grained geometric priors from RGB inputs

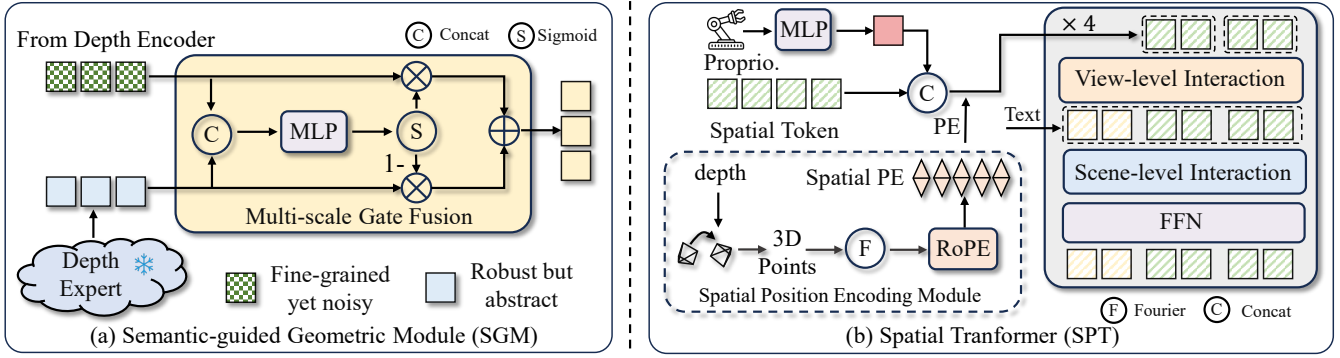


Figure 3: **Semantic-guided Geometric Module and Spatial Transformer.** (a) SGM adaptively combines two complementary geometric representations via a gating mechanism. (b) SPT converts 3D points into spatial positional embeddings using RoPE to establish 2D–3D correspondences, followed by view-level and scene-level interactions for spatial token refinement.

via a frozen large-scale pre-trained depth estimation expert:

$$\hat{F}_{\text{geo}}^v = \mathcal{E}_{\text{expert}}(I^v) \in \mathbb{R}^{H \times W \times C}, \quad (2)$$

and extract fine-grained but noisy geometry from raw depth using a depth encoder (e.g. ResNet-50 (He et al. 2016)):

$$F_{\text{geo}}^v = \mathcal{E}_{\text{raw}}(D^v) \in \mathbb{R}^{H \times W \times C}. \quad (3)$$

As shown in Fig. 3 (a), a multi-scale gating mechanism then adaptively fuses these features to yield an optimized geometric representation that preserves fine details while reducing noise and aligning with the semantic cues.

$$G^v = \sigma(\text{MLP}(\text{Concat}(\hat{F}_{\text{geo}}^v, F_{\text{geo}}^v))), \quad (4)$$

$$F_{\text{fuse-geo}}^v = G^v \odot F_{\text{geo}}^v + (1 - G^v) \odot \hat{F}_{\text{geo}}^v, \quad (5)$$

where σ denotes sigmoid activation and \odot element-wise multiplication. The gate G^v learns to retain reliable depth details while suppressing noise.

3.3 Spatial Transformer

For each view v , we denote the spatial features as $H^v \in \mathbb{R}^{N_v \times D}$. The proprioceptive input P is projected via an MLP and fused with H^v by element-wise addition:

$$\tilde{H}^v = H^v + \text{MLP}(P). \quad (6)$$

Given a pixel (x', y') with depth $d = D^v(x', y')$, its 3D coordinate $[x, y, z]^\top$ in the robot-centric coordinate system is computed via perspective projection:

$$[x, y, z, 1]^\top = E^v (d \cdot (K^v)^{-1} [x', y', 1]^\top \| 1), \quad (7)$$

where $K^v \in \mathbb{R}^{3 \times 3}$ and $E^v \in \mathbb{R}^{4 \times 4}$ denote the intrinsic and extrinsic matrices, and $\|$ denotes vector concatenation.

To encode spatial cues, we apply rotary positional encoding to \tilde{H}^v , where each axis is assigned $D/3$ dimensions. We define a set of frequencies:

$$\omega_k = \lambda^{-2k/d}, \quad k = 0, 1, \dots, \frac{d}{2} - 1, \quad d = D/3, \quad (8)$$

with $\lambda = 10000$ to control the frequency bandwidth. In the spirit of Fourier feature mappings, we compute axis-wise sinusoidal embeddings as:

$$\cos_{\text{pos}} = [\cos(\omega_k u)]_{u \in \{x, y, z\}, k=0, \dots, d/2-1}, \quad (9)$$

$$\sin_{\text{pos}} = [\sin(\omega_k u)]_{u \in \{x, y, z\}, k=0, \dots, d/2-1}. \quad (10)$$

The final position-encoded features are given by:

$$T^v = \tilde{H}^v \odot \cos_{\text{pos}} + \text{rot}(\tilde{H}^v) \odot \sin_{\text{pos}}, \quad (11)$$

where \odot denotes element-wise multiplication, and $\text{rot}(\cdot)$ rotates each (f_{2i}, f_{2i+1}) feature pair as $(-f_{2i+1}, f_{2i})$.

At the view level, self-attention followed by a feed-forward network (FFN) refines each view's token representation. At the scene level, tokens from all views are concatenated with language features F_{text} . Another round of self-attention and an FFN then fuses cross-view and language context, producing the final refined tokens. The tokens are fed into a lightweight decoder (ConvexUp) to generate a per-view 2D heatmap. The target 2D position is obtained via argmax and lifted into 3D using the camera model. The action head then uses an MLP on local features around this position to regress the rotation $\theta = (\theta_x, \theta_y, \theta_z)$ and gripper state g . Together with the 3D translation (x, y, z) , these form the final action $A = (x, y, z, \theta_x, \theta_y, \theta_z, g)$.

The action supervision includes three parts: a cross-entropy loss on per-view 2D heatmaps for translation, cross-entropy losses on discretized Euler angles for rotation, and a binary classification loss for the gripper state.

4 Experiments

To comprehensively evaluate the effectiveness of SpatialActor, we conduct experiments in both simulation and real-world settings. Specifically, we aim to answer the following key questions: (1) How does SpatialActor compare to state-of-the-art robotic manipulation policies? (2) How robust is SpatialActor under noisy conditions? (3) How well does SpatialActor generalize to few-shot settings? (4) How does SpatialActor perform under spatial perturbations? (5) What is the impact of different components of SpatialActor? (6) How does SpatialActor perform in real-robot setups?

Models	Avg. Success \uparrow	Avg. Rank \downarrow	Close Jar	Drag Stick	Insert Peg	Meat off Grill	Open Drawer	Place Cups	Place Wine	Push Buttons
C2F-ARM-BC (James et al. 2022)	20.1	9.5	24.0	24.0	4.0	20.0	20.0	0.0	8.0	72.0
HiveFormer (Guhur et al. 2023)	45.3	7.8	52.0	76.0	0.0	100.0	52.0	0.0	80.0	84.0
PolarNet (Chen et al. 2023)	46.4	7.3	36.0	92.0	4.0	100.0	84.0	0.0	40.0	96.0
PerAct (Shridhar et al. 2023)	49.4	7.1	55.2 \pm 4.7	89.6 \pm 4.1	5.6 \pm 4.1	70.4 \pm 2.0	88.0 \pm 5.7	2.4 \pm 3.2	44.8 \pm 7.8	92.8 \pm 3.0
RVT (Goyal et al. 2023)	62.9	5.3	52.0 \pm 2.5	99.2 \pm 1.6	11.2 \pm 3.0	88.0 \pm 2.5	71.2 \pm 6.9	4.0 \pm 2.5	91.0 \pm 5.2	100.0 \pm 0.0
Act3D (Gervet et al. 2023)	65.0	5.3	92.0	92.0	27.0	94.0	93.0	3.0	80.0	99.0
SAM-E (Zhang et al. 2024)	70.6	2.9	82.4 \pm 3.6	100.0 \pm 0.0	18.4 \pm 4.6	95.2 \pm 3.3	95.2 \pm 5.2	0.0 \pm 0.0	94.4 \pm 4.6	100.0 \pm 0.0
3D Diffuser Actor (Ke et al. 2024)	81.3	2.8	96.0 \pm 2.5	100.0 \pm 0.0	65.6 \pm 4.1	96.8 \pm 1.6	89.6 \pm 4.1	24.0 \pm 7.6	93.6 \pm 4.8	98.4 \pm 2.0
RVT-2 (Goyal et al. 2024)	81.4	2.8	100.0 \pm 0.0	99.0 \pm 1.7	40.0 \pm 0.0	99.0 \pm 1.7	74.0 \pm 11.8	38.0 \pm 4.5	95.0 \pm 3.3	100.0 \pm 0.0
SpatialActor (Ours)	87.4 \pm 0.8	2.3	94.0 \pm 4.2	100.0 \pm 0.0	93.3 \pm 4.8	98.7 \pm 2.1	82.0 \pm 3.3	56.7 \pm 8.5	94.7 \pm 4.8	100.0 \pm 0.0

Models	Put in Cupboard	Put in Drawer	Put in Safe	Screw Bulb	Slide Block	Sort Shape	Stack Blocks	Stack Cups	Sweep Dustpan	Turn Tap
C2F-ARM-BC (James et al. 2022)	0.0	4.0	12.0	8.0	16.0	8.0	0.0	0.0	0.0	68.0
HiveFormer (Guhur et al. 2023)	32.0	68.0	76.0	8.0	64.0	8.0	8.0	0.0	28.0	80.0
PolarNet (Chen et al. 2023)	12.0	32.0	84.0	44.0	56.0	12.0	4.0	8.0	52.0	80.0
PerAct (Shridhar et al. 2023)	28.0 \pm 4.4	51.2 \pm 4.7	84.0 \pm 3.6	17.6 \pm 2.0	74.0 \pm 13.0	16.8 \pm 4.7	26.4 \pm 3.2	2.4 \pm 2.0	52.0 \pm 0.0	88.0 \pm 4.4
RVT (Goyal et al. 2023)	49.6 \pm 3.2	88.0 \pm 5.7	91.2 \pm 3.0	48.0 \pm 5.7	81.6 \pm 5.4	36.0 \pm 2.5	28.8 \pm 3.9	26.4 \pm 8.2	72.0 \pm 0.0	93.6 \pm 4.1
Act3D (Gervet et al. 2023)	51.0	90.0	95.0	47.0	93.0	8.0	12.0	9.0	92.0	94.0
SAM-E (Zhang et al. 2024)	64.0 \pm 2.8	92.0 \pm 5.7	95.2 \pm 3.3	78.4 \pm 3.6	95.2 \pm 1.8	34.4 \pm 6.1	26.4 \pm 4.6	0.0 \pm 0.0	100.0 \pm 0.0	100.0 \pm 0.0
3D Diffuser Actor (Ke et al. 2024)	85.6 \pm 4.1	96.0 \pm 3.6	97.6 \pm 2.0	82.4 \pm 2.0	97.6 \pm 3.2	44.0 \pm 4.4	68.3 \pm 3.3	47.2 \pm 8.5	84.0 \pm 4.4	99.2 \pm 1.6
RVT-2 (Goyal et al. 2024)	66.0 \pm 4.5	96.0 \pm 0.0	96.0 \pm 2.8	88.0 \pm 4.9	92.0 \pm 2.8	35.0 \pm 7.1	80.0 \pm 2.8	69.0 \pm 5.9	100.0 \pm 0.0	99.0 \pm 1.7
SpatialActor (Ours)	72.0 \pm 3.6	98.7 \pm 3.3	96.7 \pm 3.9	88.7 \pm 3.9	91.3 \pm 6.9	73.3 \pm 6.5	56 \pm 7.6	81.3 \pm 4.1	100.0 \pm 0.0	95.3 \pm 3.0

Table 1: **Performance on RL Bench.** We report success rates on 18 RL Bench tasks with 249 variations. SpatialActor achieves the highest overall performance, surpassing the previous state-of-the-art RVT-2 by 6.0%. Notably, on tasks requiring high spatial precision, such as Insert Peg and Sort Shape, SpatialActor outperforms RVT-2 by 53.3% and 38.3%, respectively.

4.1 Comparison with State-of-the-Art Policies

Simulation Environment and Datasets. We evaluate SpatialActor on RL Bench (James et al. 2020), a mainstream multi-task 3D manipulation benchmark built on CoppeliaSim (Rohmer, Singh, and Freese 2013). The simulation environment features a Franka robotic arm with a parallel gripper operating in a tabletop scenario. Observations come from four fixed RGB-D cameras (front, left/right shoulder, wrist) at 128×128 resolution. The action space consists of 3D translation, rotation of the end-effector, and binary gripper control. An OMPL-based motion planner (Sucan, Moll, and Kavraki 2012) is utilized to compute feasible trajectories. Following PerAct (Shridhar, Manuelli, and Fox 2023), we use 18 tasks with 249 variations covering diverse manipulation skills, each with 100 expert demonstrations for training and 25 unseen episodes for evaluation.

Implementation Details. SpatialActor is trained for approximately 40k iterations using a cosine learning rate schedule with an initial 2k-iteration warm-up. Training is performed using 8 GPUs with a total batch size of 192 (24 per GPU) and an initial learning rate of 2.4×10^{-3} . Data augmentation includes random spatial translations of up to 12.5 cm along the x , y , and z axes, as well as rotations of up to 45° around the z axis. We follow RVT (Goyal et al. 2023, 2024), incorporating its virtual view design and two-stage process. Furthermore, we employ CLIP (Radford et al. 2021) as our vision-language encoder to provide aligned cross-modal representations, and Depth Anything v2 (Yang et al. 2025) as our geometry expert.

Performance on RL Bench 18 Tasks. Tab. 1 summarizes the performance of various methods on 18 RL Bench tasks with 249 variations. SpatialActor achieves an average success rate of **87.4%**, surpassing the previous state-of-the-art by 6.0%. Notably, SpatialActor shows substantial improvements on tasks requiring high spatial precision, such as Insert Peg and Sort Shape. It achieves success rates of 93.3% and 73.3% on these tasks, outperforming RVT-2 by **53.3%** and **38.3%**, respectively. These results highlight SpatialActor’s superior spatial handling capability.

4.2 Robustness under Noisy Conditions

Experimental Setup. Depth measurements are inherently affected by sensor noise, lighting variations, and surface reflections. To simulate these challenges, we inject controlled Gaussian noise into reconstructed point clouds. Specifically, we design three noise levels: *Light* corrupts 20% of the points with a Gaussian standard deviation of 0.05, *Middle* corrupts 50% of the points with a standard deviation of 0.1, and *Heavy* corrupts 80% of the points with a standard deviation of 0.1. This setup allows us to evaluate the robustness of our approach under progressively severe noisy conditions.

Performance Evaluation. Tab. 2 shows that under *Light*, *Middle*, and *Heavy* noise, SpatialActor improves average success rates over RVT-2 by 13.9%, 16.9%, and 19.4%, respectively. Notably, in tasks requiring high spatial precision, these gains are even more pronounced. For instance, on Insert Peg task, SpatialActor outperforms RVT-2 by 88.0%, 78.6%, and 61.3% under the respective noise levels.

Models	Noise type	Avg. Success \uparrow	Close Jar	Drag Stick	Insert Peg	Meat off Grill	Open Drawer	Place Cups	Place Wine	Push Buttons
RVT-2	<i>Light</i>	72.5 \pm 0.5	92.0 \pm 4.0	100.0 \pm 0.0	6.7 \pm 4.6	100.0 \pm 0.0	82.7 \pm 10.1	25.3 \pm 6.1	96.0 \pm 4.0	74.7 \pm 8.3
SpatialActor (Ours)		86.4 \pm 0.4	97.3 \pm 2.3	98.7 \pm 2.3	94.7 \pm 6.1	96.0 \pm 0.0	73.3 \pm 10.1	54.7 \pm 8.3	92.0 \pm 4.0	98.7 \pm 2.3
RVT-2	<i>Middle</i>	68.4 \pm 0.9	85.3 \pm 2.3	100.0 \pm 0.0	2.7 \pm 2.3	94.7 \pm 2.3	82.7 \pm 11.5	20.0 \pm 0.0	89.3 \pm 4.6	73.3 \pm 4.6
SpatialActor (Ours)		85.3 \pm 0.9	100.0 \pm 0.0	98.7 \pm 2.3	81.3 \pm 6.1	96.0 \pm 4.0	78.7 \pm 8.3	45.3 \pm 10.1	89.3 \pm 4.6	97.3 \pm 4.6
RVT-2	<i>Heavy</i>	57.0 \pm 0.9	49.3 \pm 6.1	94.7 \pm 4.6	0.0 \pm 0.0	97.3 \pm 2.3	86.7 \pm 2.3	8.0 \pm 4.0	86.7 \pm 2.3	64.0 \pm 4.0
SpatialActor (Ours)		76.4 \pm 0.5	82.7 \pm 2.3	98.7 \pm 2.3	61.3 \pm 6.1	100.0 \pm 0.0	80.0 \pm 4.0	21.3 \pm 4.6	92.0 \pm 0.0	92.0 \pm 4.6

Models	Put in Cupboard	Put in Drawer	Put in Safe	Screw Bulb	Slide Block	Sort Shape	Stack Blocks	Stack Cups	Sweep to Dustpan	Turn Tap
RVT-2	57.3 \pm 2.3	100.0 \pm 0.0	92.0 \pm 4.0	81.3 \pm 6.1	62.7 \pm 23.1	46.7 \pm 6.1	53.3 \pm 2.3	45.3 \pm 2.3	96.0 \pm 6.9	93.3 \pm 4.6
SpatialActor (Ours)	81.3 \pm 2.3	98.7 \pm 2.3	98.7 \pm 2.3	88.0 \pm 4.0	72.0 \pm 4.0	76.0 \pm 6.9	62.7 \pm 2.3	82.7 \pm 2.3	97.3 \pm 4.6	93.3 \pm 2.3
RVT-2	50.7 \pm 12.2	98.7 \pm 2.3	98.7 \pm 2.3	76.0 \pm 4.0	57.3 \pm 2.3	38.7 \pm 10.1	45.3 \pm 12.2	25.3 \pm 6.1	96.0 \pm 4.0	96.0 \pm 6.9
SpatialActor (Ours)	74.7 \pm 6.1	100.0 \pm 0.0	94.7 \pm 2.3	88.0 \pm 4.0	81.3 \pm 15.1	76.0 \pm 4.0	58.7 \pm 6.1	77.3 \pm 8.3	100.0 \pm 0.0	97.3 \pm 2.3
RVT-2	20.0 \pm 6.9	97.3 \pm 2.3	93.3 \pm 2.3	58.7 \pm 2.3	57.3 \pm 8.3	13.3 \pm 6.1	13.3 \pm 6.1	1.3 \pm 2.3	92.0 \pm 0.0	92.0 \pm 4.0
SpatialActor (Ours)	64.0 \pm 4.0	100.0 \pm 0.0	100.0 \pm 0.0	78.7 \pm 8.3	58.7 \pm 2.3	52.0 \pm 4.0	42.7 \pm 6.1	70.7 \pm 6.1	82.7 \pm 2.3	97.3 \pm 4.6

Table 2: **Performance Under Various Noise Levels.** We report success rates under three noise conditions: *Light* noise corrupts 20% of the points in the reconstructed point cloud with random Gaussian noise (std = 0.05), *Middle* noise corrupts 50% with noise of std = 0.1, and *Heavy* noise corrupts 80% with noise of std = 0.1. Under these conditions, SpatialActor improves average success rates by approximately 13.9%, 16.9%, and 19.4% over RVT-2 at the Light, Middle, and Heavy noise levels, respectively.

Models	Avg. Success \uparrow	Close Laptop	Put Rubbish in Bin	Beat Buzz	Close Microwave	Put Shoes in Box	Get Ice	Change Clock	Close Box	Reach Target
RVT-2	46.9 \pm 1.5	76.0 \pm 6.1	10.3 \pm 5.1	47.4 \pm 8.5	61.7 \pm 9.8	7.4 \pm 4.3	93.7 \pm 3.9	72.6 \pm 2.8	49.1 \pm 8.6	12.0 \pm 5.7
SpatialActor	79.2 \pm 2.7	90.0 \pm 7.5	100.0 \pm 0.0	92.0 \pm 2.5	95.3 \pm 11.4	25.3 \pm 13.8	96.0 \pm 2.5	83.3 \pm 7.3	95.3 \pm 4.7	86.0 \pm 2.2

Models	Close Door	Remove Cups	Close Drawer	Spatula Scoop	Close Fridge	Put Knife on Board	Screw Nail	Close Grill	Plate in Rack	Meat on Grill
RVT-2	4.0 \pm 3.3	33.7 \pm 13.8	96.0 \pm 0.0	70.9 \pm 6.8	81.7 \pm 8.6	14.3 \pm 7.3	38.9 \pm 15.1	66.3 \pm 8.9	24.6 \pm 7.1	30.0 \pm 8.5
SpatialActor	36.0 \pm 14.1	66.0 \pm 8.3	96.7 \pm 3.9	84.7 \pm 8.2	95.3 \pm 5.3	66.0 \pm 2.2	62.7 \pm 6.0	96.0 \pm 0.0	48.0 \pm 8.0	90.0 \pm 2.8

Table 3: **Few-Shot Generalization.** We adapt pre-trained model to 19 new tasks using only 10 demonstrations per task (1/10th of original data). We reports success rates, showing that SpatialActor, significantly outperforms RVT-2 in the few-shot setting.

4.3 Few-Shot Generalization

We evaluate the few-shot generalization ability of SpatialActor by adapting the multi-task pre-trained model to 19 novel tasks using only 10 demonstrations per task, just one-tenth of the data used during multi-task training. In this few-shot adaptation scenario, the model is initialized with its pre-trained weights and then fine-tuned on the limited data. As shown in Tab. 3, our experiments demonstrate that SpatialActor effectively transfers previously learned skills to new tasks with minimal adaptation data. Overall, SpatialActor achieves an average success rate of 79.2%, compared to 46.9% for RVT-2, yielding an improvement of approximately 32.3%. This significant boost underscores the superior few-shot generalization capability of our approach.

4.4 Spatial Perturbations on ColosseumBench

Setup. We evaluate the robustness on the Colosseum benchmark (Pumacay et al. 2024), which assesses manipulation policies under environmental variations. We report results on 20 tasks under the baseline setting (no perturbation) and three spatial perturbations: manipulation object size (MO-Size), which scales the object being manipulated; receiver object size (RO-Size), which scales an indirectly used object such as a container; and camera pose perturbation (Camera-Perturb), which randomly shifts the camera’s position and orientation to vary the observation viewpoint.

Performance Evaluation. The results in Tab. 4 indicate that under the no-perturbation condition, our method achieves a task-average success rate of 57.4%. When spatial perturbations are introduced, SpatialActor attains 59.2% under MO-Size variations, 62.0% under RO-Size changes, and 54.2% with camera pose perturbations. These results consistently outperform competing methods, demonstrating strong robustness and generalization under spatial variations.

4.5 Ablation Study

Our ablation study on 18 tasks (Tab. 5) shows that decoupling semantics and geometry improves performance in both noise-free and heavy-noise settings, increasing success rates from 81.4% to 85.1% and from 57.0% to 68.7%, respectively. Introducing the Semantic-guided Geometry Module (SGM) further boosts performance, especially under heavy noise, where performance rises to 73.9%. Finally, the Spatial Transformer (SPT), which provides precise low-level spatial cues, brings the success rates to 87.4% and 76.4% in noise-free and noisy conditions, respectively.

4.6 Real-World Evaluation

Setup. In real-world experiments, we use a WidowX single-arm robot equipped with an Intel RealSense D435i RGB-D camera. The camera is statically mounted to capture a front view of the workspace. We perform both intrinsic

Method	No-Vars \uparrow	MO-Size \uparrow	RO-Size \uparrow	Cam Pose \uparrow
R3M (Nair et al. 2022)	2.9	1.8	0.0	0.8
MVP (Radosavovic et al. 2022)	3.4	4.4	0.5	2.6
VoxPoser (Huang et al. 2023)	5.4	3.3	6.5	6.2
PerAct (Shridhar et al. 2023)	34.5	35.6	29.3	36.3
RVT (Goyal et al. 2023)	43.6	35.3	40.5	42.2
SpatialActor (Ours)	57.4\pm3.0	59.2\pm2.4	62.0\pm3.2	54.2\pm1.8

Table 4: **Performance Under Spatial Perturbations.** We report average success rates on 20 ColosseumBench tasks under 4 conditions: No-Vars, manipulation object size, receiver object size, and camera pose.

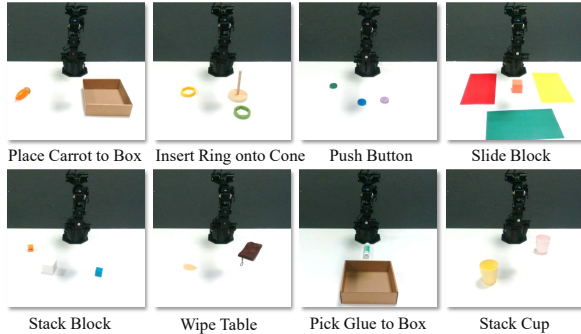


Figure 4: **Real-world tasks.** We employed 8 distinct tasks with a total of 15 variants in real-world experiments.

and extrinsic calibration between the camera and the robot to accurately transform the observed point clouds into the robot’s base coordinate system. The system is integrated using a ROS package. Images are originally captured at a resolution of 1280×720 and are downsampled to 128×128 .

Dataset Collection. We conduct experiments on a series of real-world tasks (Fig. 4), including (1) Pick Glue to Box, (2) Stack Cup, (3) Push Button, (4) Slide Block, (5) Place Carrot to Box, (6) Stack Block, (7) Insert Ring onto Cone, and (8) Wipe Table. For each task, we collect 25 demonstrations that capture diverse spatial configurations and object variations. Some tasks are instantiated with multiple variations, for example, the Slide Block task includes yellow, green, and red variants, resulting in a total of 15 variations across the 8 tasks. The trajectories are recorded at 30 fps, and key-frames are extracted to construct the training set.

Evaluation. We evaluate SpatialActor against RVT-2 on various real-world tasks. Single-variant tasks are tested 20 times, and multi-variant tasks 10 times per variant. As shown in Tab. 6, SpatialActor consistently outperforms RVT-2, with an average improvement of around 20% across tasks, demonstrating effectiveness in real-world scenarios.

To evaluate robustness to distribution shifts, we test SpatialActor under variations in manipulated object, receiver object, lighting, and background (Fig. 5). SpatialActor maintains consistently high performance across these diverse and challenging conditions, clearly demonstrating strong robustness and generalization in complex real-world scenarios.

Decouple	SGM	SPT	Avg. success on 18 tasks \uparrow	
			No noise	Heavy noise
			81.4	57.0
\checkmark			85.1	68.7
\checkmark	\checkmark		86.4	73.9
\checkmark	\checkmark	\checkmark	87.4	76.4

Table 5: **Ablation Study.** We analyze the contribution of each module to overall performance and their effect on robustness under heavy noisy conditions.

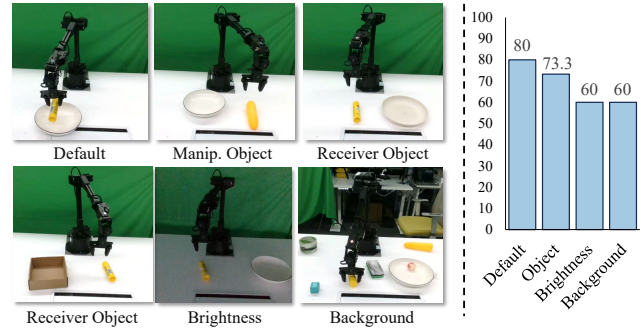


Figure 5: **Real-world Generalization Evaluation.** We assess SpatialActor under variations in manipulated object, receiver object, brightness, and background. Performance remains robust across challenging settings.

Task	#Var.	RVT-2	SpatialActor
(1) Pick Glue to Box	1	50%	85%
(2) Stack Cup	2	30%	30%
(3) Push Button	3	67%	90%
(4) Slide Block	3	60%	67%
(5) Place Carrot to Box	1	30%	65%
(6) Stack Block	2	40%	35%
(7) Insert Ring Onto Cone	2	20%	50%
(8) Wipe Table	1	50%	80%
All tasks	15	43%	63%

Table 6: **Real-World Results.** We report success rates for each task and overall performance across 8 tasks with 15 variations. SpatialActor, consistently outperforms RVT-2, indicating superior robustness in real-world scenarios.

5 Conclusion

In this work, we present SpatialActor, a framework for robust spatial representation in robotic manipulation. SpatialActor disentangles semantic and geometric information, with the geometric branch comprising a high-level module (SGM), which fuses semantic-guided geometric priors with depth features, and a low-level module (SPT), which captures fine-grained spatial cues through position-aware interactions. Experiments across 50+ simulated and real-world tasks show that SpatialActor achieves higher success rates and strong robustness, underscoring the value of disentangled spatial representations for reliable manipulation.

Acknowledgments

This work was supported by the National Science and Technology Major Project of China under Grant No. 2023ZD0121300, the Scientific Research Innovation Capability Support Project for Young Faculty under Grant No. ZYGXQNJSKYCXNLZCXM-I20, and the National Natural Science Foundation of China under Grant No. U24B20173.

References

- Andrychowicz, O. M.; Baker, B.; Chociej, M.; Jozefowicz, R.; McGrew, B.; Pachocki, J.; Petron, A.; Plappert, M.; Powell, G.; Ray, A.; et al. 2020. Learning dexterous in-hand manipulation. *The International Journal of Robotics Research*, 39(1): 3–20.
- Bhat, S. F.; Birkl, R.; Wofk, D.; Wonka, P.; and Müller, M. 2023. Zoedepth: Zero-shot transfer by combining relative and metric depth. *arXiv preprint arXiv:2302.12288*.
- Brohan, A.; Brown, N.; Carbajal, J.; Chebotar, Y.; Dabis, J.; Finn, C.; Gopalakrishnan, K.; Hausman, K.; Herzog, A.; Hsu, J.; et al. 2022. Rt-1: Robotics transformer for real-world control at scale. *arXiv preprint arXiv:2212.06817*.
- Chen, S.; Garcia, R.; Schmid, C.; and Laptev, I. 2023. Polarnet: 3d point clouds for language-guided robotic manipulation. *arXiv preprint arXiv:2309.15596*.
- Chi, C.; Xu, Z.; Feng, S.; Cousineau, E.; Du, Y.; Burchfiel, B.; Tedrake, R.; and Song, S. 2023. Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research*, 02783649241273668.
- Deng, X.; Xiang, Y.; Mousavian, A.; Eppner, C.; Bretl, T.; and Fox, D. 2020. Self-supervised 6d object pose estimation for robot manipulation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 3665–3671. IEEE.
- Fang, H.; Grotz, M.; Pumacay, W.; Wang, Y. R.; Fox, D.; Krishna, R.; and Duan, J. 2025. SAM2Act: Integrating Visual Foundation Model with A Memory Architecture for Robotic Manipulation. *arXiv preprint arXiv:2501.18564*.
- Fang, H.-S.; Wang, C.; Fang, H.; Gou, M.; Liu, J.; Yan, H.; Liu, W.; Xie, Y.; and Lu, C. 2023. Anygrasp: Robust and efficient grasp perception in spatial and temporal domains. *IEEE Transactions on Robotics*, 39(5): 3929–3945.
- Feng, T.; Shi, H.; Liu, X.; Feng, W.; Wan, L.; Zhou, Y.; and Lin, D. 2023. Open compound domain adaptation with object style compensation for semantic segmentation. *Advances in Neural Information Processing Systems*, 36: 63136–63149.
- Gervet, T.; Xian, Z.; Gkanatsios, N.; and Fragkiadaki, K. 2023. Act3d: 3d feature field transformers for multi-task robotic manipulation. *arXiv preprint arXiv:2306.17817*.
- Goyal, A.; Blukis, V.; Xu, J.; Guo, Y.; Chao, Y.-W.; and Fox, D. 2024. Rvt-2: Learning precise manipulation from few demonstrations. *arXiv preprint arXiv:2406.08545*.
- Goyal, A.; Xu, J.; Guo, Y.; Blukis, V.; Chao, Y.-W.; and Fox, D. 2023. Rvt: Robotic view transformer for 3d object manipulation. In *Conference on Robot Learning*, 694–710. PMLR.
- Guhur, P.-L.; Chen, S.; Pinel, R. G.; Tapaswi, M.; Laptev, I.; and Schmid, C. 2023. Instruction-driven history-aware policies for robotic manipulations. In *Conference on Robot Learning*, 175–187. PMLR.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Huang, W.; Wang, C.; Zhang, R.; Li, Y.; Wu, J.; and Fei-Fei, L. 2023. VoxPoser: Composable 3D Value Maps for Robotic Manipulation with Language Models. *arXiv preprint arXiv:2307.05973*.
- James, S.; Ma, Z.; Arrojo, D. R.; and Davison, A. J. 2020. Rlbench: The robot learning benchmark & learning environment. *IEEE Robotics and Automation Letters*, 5(2): 3019–3026.
- James, S.; Wada, K.; Laidlow, T.; and Davison, A. J. 2022. Coarse-to-fine q-attention: Efficient learning for visual robotic manipulation via discretisation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13739–13748.
- Jia, Y.; Liu, J.; Chen, S.; Gu, C.; Wang, Z.; Luo, L.; Lee, L.; Wang, P.; Wang, Z.; Zhang, R.; et al. 2024. Lift3d foundation policy: Lifting 2d large-scale pretrained models for robust 3d robotic manipulation. *arXiv preprint arXiv:2411.18623*.
- Kang, B.; Yue, Y.; Lu, R.; Lin, Z.; Zhao, Y.; Wang, K.; Huang, G.; and Feng, J. 2024. How far is video generation from world model: A physical law perspective. *arXiv preprint arXiv:2411.02385*.
- Ke, T.-W.; Gkanatsios, N.; and Fragkiadaki, K. 2024. 3d diffuser actor: Policy diffusion with 3d scene representations. *arXiv preprint arXiv:2402.10885*.
- Kim, M. J.; Pertsch, K.; Karamcheti, S.; Xiao, T.; Balakrishna, A.; Nair, S.; Rafailov, R.; Foster, E.; Lam, G.; Sanketi, P.; et al. 2024. Openvla: An open-source vision-language-action model. *arXiv preprint arXiv:2406.09246*.
- Li, J.; Li, D.; Xiong, C.; and Hoi, S. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International conference on machine learning*, 12888–12900. PMLR.
- Liu, S.; Wu, L.; Li, B.; Tan, H.; Chen, H.; Wang, Z.; Xu, K.; Su, H.; and Zhu, J. 2024a. Rdt-1b: a diffusion foundation model for bimanual manipulation. *arXiv preprint arXiv:2410.07864*.
- Liu, S.; Zeng, Z.; Ren, T.; Li, F.; Zhang, H.; Yang, J.; Jiang, Q.; Li, C.; Yang, J.; Su, H.; et al. 2024b. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *European Conference on Computer Vision*, 38–55. Springer.
- Nair, S.; Rajeswaran, A.; Kumar, V.; Finn, C.; and Gupta, A. 2022. R3m: A universal visual representation for robot manipulation. *arXiv preprint arXiv:2203.12601*.
- Pumacay, W.; Singh, I.; Duan, J.; Krishna, R.; Thomason, J.; and Fox, D. 2024. The colosseum: A benchmark for evaluating generalization for robotic manipulation. *arXiv preprint arXiv:2402.08191*.

- Qian, G.; Li, Y.; Peng, H.; Mai, J.; Hammoud, H.; Elhoseiny, M.; and Ghanem, B. 2022. Pointnext: Revisiting pointnet++ with improved training and scaling strategies. *Advances in neural information processing systems*, 35: 23192–23204.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.
- Radosavovic, I.; Xiao, T.; James, S.; Abbeel, P.; Malik, J.; and Darrell, T. 2022. Real-World Robot Learning with Masked Visual Pre-training. *CoRL*.
- Rohmer, E.; Singh, S. P.; and Freese, M. 2013. V-REP: A versatile and scalable robot simulation framework. In *2013 IEEE/RSJ international conference on intelligent robots and systems*, 1321–1326. IEEE.
- Seo, Y.; Kim, J.; James, S.; Lee, K.; Shin, J.; and Abbeel, P. 2023. Multi-view masked world models for visual robotic manipulation. In *International Conference on Machine Learning*, 30613–30632. PMLR.
- Shi, H.; Xie, B.; Liu, Y.; Sun, L.; Liu, F.; Wang, T.; Zhou, E.; Fan, H.; Zhang, X.; and Huang, G. 2025. Memoryvla: Perceptual-cognitive memory in vision-language-action models for robotic manipulation. *arXiv preprint arXiv:2508.19236*.
- Shridhar, M.; Manuelli, L.; and Fox, D. 2023. Perceiver-actor: A multi-task transformer for robotic manipulation. In *Conference on Robot Learning*, 785–799. PMLR.
- Sucan, I. A.; Moll, M.; and Kavraki, L. E. 2012. The open motion planning library. *IEEE Robotics & Automation Magazine*, 19(4): 72–82.
- Sun, L.; Xie, B.; Liu, Y.; Shi, H.; Wang, T.; and Cao, J. 2025. Geovla: Empowering 3d representations in vision-language-action models. *arXiv preprint arXiv:2508.09071*.
- Wang, J.; Chen, M.; Karaev, N.; Vedaldi, A.; Rupprecht, C.; and Novotny, D. 2025a. Vggt: Visual geometry grounded transformer. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 5294–5306.
- Wang, W.; Lei, Y.; Jin, S.; Hager, G. D.; and Zhang, L. 2024. Vihe: Virtual in-hand eye transformer for 3d robotic manipulation. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 403–410. IEEE.
- Wang, Y.; Yue, Y.; Yue, Y.; Wang, H.; Jiang, H.; Han, Y.; Ni, Z.; Pu, Y.; Shi, M.; Lu, R.; et al. 2025b. Emulating human-like adaptive vision for efficient and flexible machine visual perception. *Nature Machine Intelligence*, 1–19.
- Wu, D.; Fu, Y.; Huang, S.; Liu, Y.; Jia, F.; Liu, N.; Dai, F.; Wang, T.; Anwer, R. M.; Khan, F. S.; et al. 2025. RAG-Net: Large-scale Reasoning-based Affordance Segmentation Benchmark towards General Grasping. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 11980–11990.
- Xie, B.; Zhou, E.; Jia, F.; Shi, H.; Fan, H.; Zhang, H.; Li, H.; Sun, J.; Bin, J.; Huang, J.; et al. 2025. Dexbotic: Open-Source Vision-Language-Action Toolbox. *arXiv preprint arXiv:2510.23511*.
- Yang, L.; Kang, B.; Huang, Z.; Xu, X.; Feng, J.; and Zhao, H. 2024. Depth anything: Unleashing the power of large-scale unlabeled data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10371–10381.
- Yang, L.; Kang, B.; Huang, Z.; Zhao, Z.; Xu, X.; Feng, J.; and Zhao, H. 2025. Depth anything v2. *Advances in Neural Information Processing Systems*, 37: 21875–21911.
- Yue, Y.; Wang, Y.; Kang, B.; Han, Y.; Wang, S.; Song, S.; Feng, J.; and Huang, G. 2025. DeeR-VLA: Dynamic Inference of Multimodal Large Language Models for Efficient Robot Execution. *Advances in Neural Information Processing Systems*, 37: 56619–56643.
- Ze, Y.; Zhang, G.; Zhang, K.; Hu, C.; Wang, M.; and Xu, H. 2024. 3d diffusion policy: Generalizable visuomotor policy learning via simple 3d representations. *arXiv preprint arXiv:2403.03954*.
- Zeng, A.; Florence, P.; Tompson, J.; Welker, S.; Chien, J.; Attarian, M.; Armstrong, T.; Krasin, I.; Duong, D.; Sindhwani, V.; et al. 2021. Transporter networks: Rearranging the visual world for robotic manipulation. In *Conference on Robot Learning*, 726–747. PMLR.
- Zeng, J.; Bu, Q.; Wang, B.; Xia, W.; Chen, L.; Dong, H.; Song, H.; Wang, D.; Hu, D.; Luo, P.; et al. 2024. Learning manipulation by predicting interaction. *arXiv preprint arXiv:2406.00439*.
- Zhang, J.; Bai, C.; He, H.; Xia, W.; Wang, Z.; Zhao, B.; Li, X.; and Li, X. 2024. SAM-E: leveraging visual foundation model with sequence imitation for embodied manipulation. *arXiv preprint arXiv:2405.19586*.
- Zhang, T.; Hu, Y.; Cui, H.; Zhao, H.; and Gao, Y. 2023. A universal semantic-geometric representation for robotic manipulation. *arXiv preprint arXiv:2306.10474*.
- Zhang, Y.; Wu, D.; Shi, H.; Liu, Y.; Wang, T.; Fan, H.; and Dong, X. 2025. Grounding Beyond Detection: Enhancing Contextual Understanding in Embodied 3D Grounding. *arXiv preprint arXiv:2506.05199*.
- Zhao, T. Z.; Kumar, V.; Levine, S.; and Finn, C. 2023. Learning fine-grained bimanual manipulation with low-cost hardware. *arXiv preprint arXiv:2304.13705*.
- Zheng, H.; Shi, H.; Chng, Y. X.; Huang, R.; Ni, Z.; Tan, T.; Peng, Q.; Weng, Y.; Shi, Z.; and Huang, G. 2024. DenseG: Alleviating Vision-Language Feature Sparsity in Multi-View 3D Visual Grounding. *Autonomous Grand Challenge CVPR 2024 Workshop*.
- Zheng, H.; Shi, H.; Peng, Q.; Chng, Y. X.; Huang, R.; Weng, Y.; Shi, Z.; and Huang, G. 2025. Densegrounding: Improving dense language-vision semantics for ego-centric 3d visual grounding. *arXiv preprint arXiv:2505.04965*.
- Zhong, Y.; Bai, F.; Cai, S.; Huang, X.; Chen, Z.; Zhang, X.; Wang, Y.; Guo, S.; Guan, T.; Lui, K. N.; et al. 2025. A Survey on Vision-Language-Action Models: An Action Tokenization Perspective. *arXiv preprint arXiv:2507.01925*.
- Zhu, C.; Wang, T.; Zhang, W.; Pang, J.; and Liu, X. 2024. Llava-3d: A simple yet effective pathway to empowering llms with 3d-awareness. *arXiv preprint arXiv:2409.18125*.