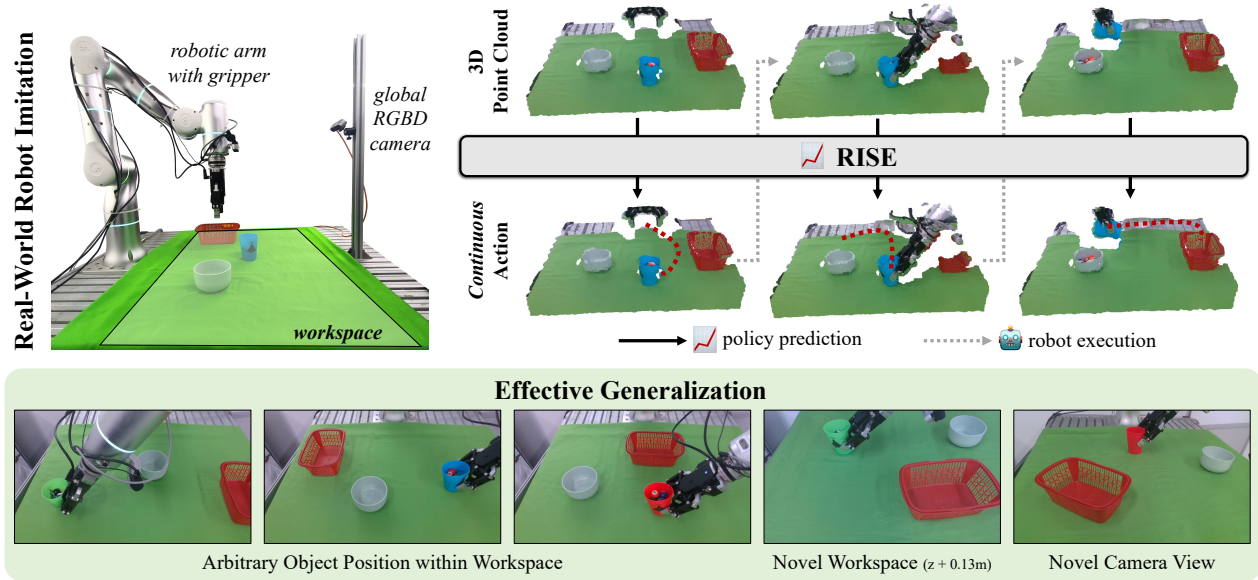


# RISE: 3D Perception Makes Real-World Robot Imitation Simple and Effective

Chenxi Wang<sup>1</sup>, Hongjie Fang<sup>2</sup>, Hao-Shu Fang<sup>2\*</sup>, Cewu Lu<sup>2\*</sup>



**Fig. 1:**  $\checkmark$  RISE focuses on real-world robot imitation settings with a noisy single-view partial point cloud as input, and outputs continuous robot actions. While simple, it shows effective generalization ability across object locations, novel workspaces, and novel camera views.

**Abstract**—Precise robot manipulations require rich spatial information in imitation learning. Image-based policies model object positions from fixed cameras, which are sensitive to camera view changes. Policies utilizing 3D point clouds usually predict keyframes rather than continuous actions, posing difficulty in dynamic and contact-rich scenarios. To utilize 3D perception efficiently, we present RISE, an end-to-end baseline for real-world imitation learning, which predicts continuous actions directly from single-view point clouds. It compresses the point cloud to tokens with a sparse 3D encoder. After adding sparse positional encoding, the tokens are featurized using a transformer. Finally, the features are decoded into robot actions by a diffusion head. Trained with 50 demonstrations for each real-world task, RISE surpasses currently representative 2D and 3D policies by a large margin, showcasing significant advantages in both accuracy and efficiency. Experiments also demonstrate that RISE is more general and robust to environmental change compared with previous baselines. Project website: [rise-policy.github.io](https://rise-policy.github.io).

## I. INTRODUCTION

End-to-end policy learning framework plays an increasingly important role in robotic manipulations. With the advancements in robotics, researchers are recognizing the significance of integrating perception, planning, and execution into a continuous learning process. Recent work has

made significant strides in imitation learning in an end-to-end fashion [1, 4, 9, 10, 58]. These methods allow robots to learn directly from perceptual data, bypassing the need for laborious hand-crafted feature engineering or planning steps, which opens new possibilities for addressing complex manipulation tasks and drives research and progress in the field of manipulation [38].

Spatial information is crucial for precise manipulations. For example, to pour a cup of water, the robot is required to understand the position and orientation of the cup to adjust its movements carefully and prevent water from spilling out. Image-based imitation learning tends to learn implicit spatial representations from fixed camera views [1, 4, 9, 15, 46, 58]. Many of these approaches utilize distinct image encoders for each view and increase the number of cameras to enhance stability and precision, consequently increasing the number of network parameters and computational overhead. In addition, such methods are susceptible to variations in camera poses, introducing a significant challenge to effective model deployment in real-world scenarios.

Spatial information can also be represented using 3D point clouds, which have demonstrated remarkable advantages in robotic tasks like general grasping [8, 45]. Recently, imitation learning based on point clouds is drawing increasing interest in our community [2, 11, 12, 14, 19, 42, 51, 54, 56]. Most of the 3D-based methods learn to predict the next keyframe as opposed to continuous actions, which often struggle with tasks involving frequent contacts and abrupt environmental

<sup>1</sup> Shanghai Noematrix Intelligence Technology Ltd.

<sup>2</sup> Shanghai Jiao Tong University.

\* Hao-Shu Fang and Cewu Lu are the corresponding authors.

Author e-mails: chenxi.wang@noematrix.cn, galaxies@sjtu.edu.cn, fhaoshu@gmail.com, lucewu@sjtu.edu.cn

changes. Meanwhile, addressing the annotation of keyframes at scale for real-world data necessitates additional manual effort. The reason that they struggle to predict continuous actions lies in the difficulty of achieving real-time performance, which relies heavily on efficient 3D feature encoding and remains a challenge for point cloud based methods.

In this work, we propose an end-to-end imitation baseline, **RISE**, a method leveraging 3D perception to make real-world robot imitation simple and effective. RISE takes point clouds as input directly, and outputs continuous action trajectories for the immediate future. It employs a shallow 3D encoder built with sparse convolutions [5], which effectively utilizes the advantages of conventional convolution architectures and avoids redundant computations on empty 3D spaces. Such design makes RISE able to obtain precise spatial information with a single camera, reducing the additional cost of hardware accumulation. The encoded point features are then mapped to the action space by a transformer. Considering the unordered characteristic of 3D points, we use sparse positional encoding, a function of coordinates, to help the transformer master the relative relationships among different point tokens in 3D space. Although the point tokens are not distributed in continuous positions like language or image tokens, we can still observe stable results in scenes with variable object locations. With sparse positional encoding, the point features can be easily modeled by transformers and naturally embedded into multimodal inputs. To show the benefits of 3D perception, we only consider RISE with point cloud input in this work. Finally, the action features are decoded into continuous trajectories by a diffusion head [4].

We test RISE in 6 real-world tasks, including *pick-and-place*, *6-DoF pouring*, *push-to-goal* (with tool), and *long-horizon* tasks. To verify the generalization ability of the policy with object locations, all the objects are randomly arranged throughout the entire workspace. Trained on 50 demonstrations for each task, RISE significantly outperforms other representative methods and keeps stable when the number of objects increases. We also find that RISE is more robust to environmental disturbance, such as changing camera views and increasing table height. Such generalization enhances the error tolerance of real-world deployment.

## II. RELATED WORK

### A. Imitation Learning for Robotics

Imitation learning is a machine learning paradigm where a robot learns to operate by observing and mimicking expert demonstrations. Behavior cloning (BC) [34], as the most direct form of imitation learning, aims to identify a mapping from observations to corresponding robot actions with the supervision of the given demonstrations. Despite its simplicity, BC has shown promising potential in learning robotic manipulations [1, 4, 21, 27, 42, 46, 58].

**2D Imitation Learning.** 2D image data is commonly used in imitation learning. One intuitive approach is to utilize pre-trained representation models for images [24, 25, 26, 30, 37] to convert them into 1D representations and map these transformed observations to the action space either through

a BC policy [55] or non-parametric nearest neighbour [33]. Unfortunately, current pre-trained representation models are not general enough to handle diverse experimental environments, encountering trouble achieving satisfactory results in real-world settings. Thus, many researchers learn such mapping in an end-to-end manner [1, 4, 21, 27, 31, 39, 46, 58, 62] and have demonstrated impressive performance across many tasks. Specifically, ACT [58] adopts a CVAE scheme [43] with transformer backbones [47] and ResNet image encoders [17] to model the variability of human data, while Diffusion Policy [4] directly utilizes diffusion process [18] to express multimodal action distributions generatively. Nonetheless, these policies are sensitive to camera positions and often fail to capture 3D spatial information about the objects in the environments.

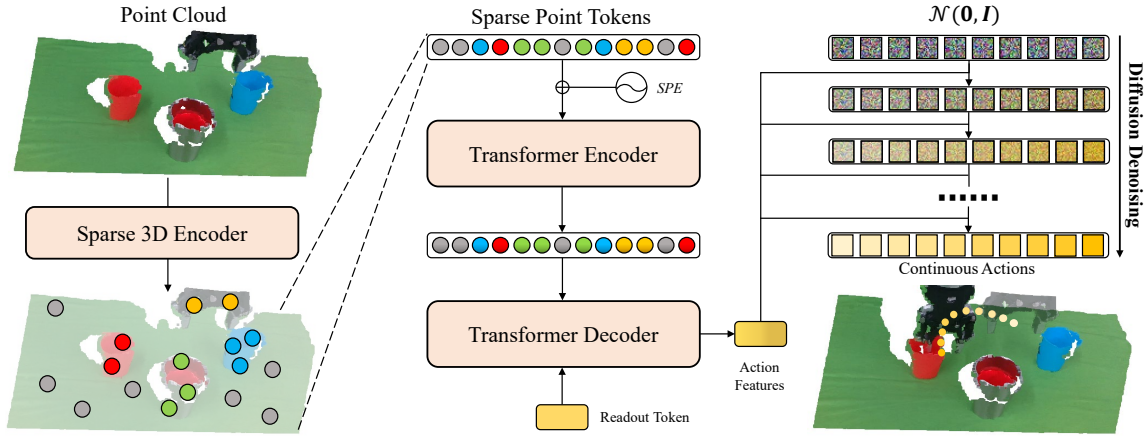
**3D Imitation Learning.** The formulation of incorporating 3D information in the imitation learning framework is under active exploration. The most straightforward method is to apply projections to transform the 3D point cloud to several 2D image views and transfer the task to multi-view image-based policy learning [12, 14], which requires the virtual viewpoints to be carefully designed to ensure performance. Moreover, due to sparse and noisily sensed point clouds, [12] fails to grasp slim objects like marker pens in real-world experiments. [19, 42, 54] process point clouds to dense voxel grids and apply 3D convolutions. Since high-resolution 3D feature maps require expensive computes, these methods have to trade off performance against cost. [61] proposed an object-centric representation for learning which requires an additional segmentation process. [11, 51] featurize point clouds by projecting multi-view image features to 3D world to avoid dense convolutions. However, such feature fusion techniques struggle to capture the consistent 3D representation from different views accurately. Recently, a concurrent work DP3 [53] also leverages 3D perception in robotic manipulation policies, but our real-world evaluations in §IV-F demonstrate that it cannot handle demonstrations with various representations limited by its network capacity. Dex-Cap [48] also proposed a PointNet [35] with diffusion head architecture for dexterous manipulation.

As mentioned before, most of the current 3D robotic imitation learning methods predict keyframes instead of continuous action, which makes it hard to annotate and limits their capacity. Besides, many of these methods only show results in simulation environments like RL Bench [20] and CALVIN [28]. In this work, we aim to evaluate our method in a more challenging setting: continuous action control with a noisy single-view partial point cloud in the real world.

### B. 3D Perception

3D perception has received considerable attention from researchers in the computer vision and robotics communities. It can be roughly divided into the following three categories:

**Projection-based.** This approach initially projects the 3D point cloud onto multiple images on different planes and then employs traditional multi-view image perception techniques. It is widely applied in shape recognition [16],



**Fig. 2:** Overview of RISE architecture. The input of RISE is a noisy point cloud captured from the real world. A 3D encoder built with sparse convolution is employed to compress the point cloud into tokens. The tokens are fed into the transformer encoder after adding sparse positional encoding. A readout token is used to query the action features from the transformer decoder. Conditioned on the action features, the Gaussian samples are denoised into continuous actions iteratively using a diffusion head.

object detection [3, 23] and robotic manipulations [12, 14] due to its simplicity. However, the projections can lead to the geometric information loss of the 3D data, and the sensitivity to the choice of projection planes may result in inferior performance [57].

**Point-based.** Early researchers directly utilized 3D convolutional neural networks (CNNs) to process 3D point cloud data based on dense volumetric representations [6, 50, 60]. Still, the sparsity of 3D data makes the vanilla approaches inefficient and memory-intensive. To solve this problem, researchers have explored using octrees for memory footprint reduction [40, 49], utilizing sparse convolutions to minimize unnecessary computations in inactive regions to improve efficiency and effectiveness [5, 13], and aggregating features across point sets directly using different network architectures [32, 35, 36, 57].

**NeRF-based.** Neural radiance fields (NeRFs) [29] have demonstrated impressive performance on high-fidelity 3D scene synthesis and scene representation extractions. In recent years, some studies [7, 41, 52, 54] have employed features extracted from pre-trained 2D foundational models as additional supervisory signals in NeRF training for scene feature extraction and distillation. Nevertheless, NeRF training requires image data from multiple views, which poses obstacles for scaling up in real-world environments. Additionally, it does not align with our single-view setting.

### III. METHOD

Given a point cloud  $\mathcal{O}^t = \{P_i^t = (x_i^t, y_i^t, z_i^t, r_i^t, g_i^t, b_i^t)\}$  as the observation at time  $t$ , RISE aims to predict the next  $n$ -step robot actions  $\mathcal{A}^t = \{A_{t+1}, A_{t+2}, \dots, A_{t+n}\}$ , where  $A_i$  contains the translation, rotation and width of the gripper. Due to the large domain gap between point clouds and robot actions, it is challenging to learn the approximation  $f: \mathcal{O}^t \rightarrow \mathcal{A}^t$  directly. To model the process, RISE is decomposed into three functions: a sparse 3D encoder  $h_E: \mathcal{O}^t \rightarrow \mathcal{F}_P^t$ , a transformer  $h_T: \mathcal{F}_P^t \rightarrow \mathcal{F}_A^t$  and an action decoder  $h_D: \mathcal{F}_A^t \rightarrow \mathcal{A}^t$ , where  $\mathcal{F}_P^t$  and  $\mathcal{F}_A^t$  denote the features of point clouds and actions respectively.

#### A. Modeling Point Clouds using Sparse 3D Encoder

The most significant difference between point cloud data and images is that point clouds are sparse and unorganized, which makes CNNs unsuitable to be applied to the points. For inputs at different scales, the computation efficiency and flexibility of a model should be taken into consideration. We employ a 3D encoder built on sparse convolution [5]. It keeps most of the standard convolution, while only computes outputs on predefined coordinates. Such an operator saves computation and inherits the core advantage of conventional convolution.

The sparse 3D encoder  $h_E$  adopts a shallow ResNet architecture [17]. It is composed of one initial convolution layer, four residual blocks, and one final convolution layer, with five  $2 \times$  sparse pooling layers between every two components. The number of layers can be freely increased, while the evaluation results demonstrate that a shallow encoder is sufficient for our experiments.

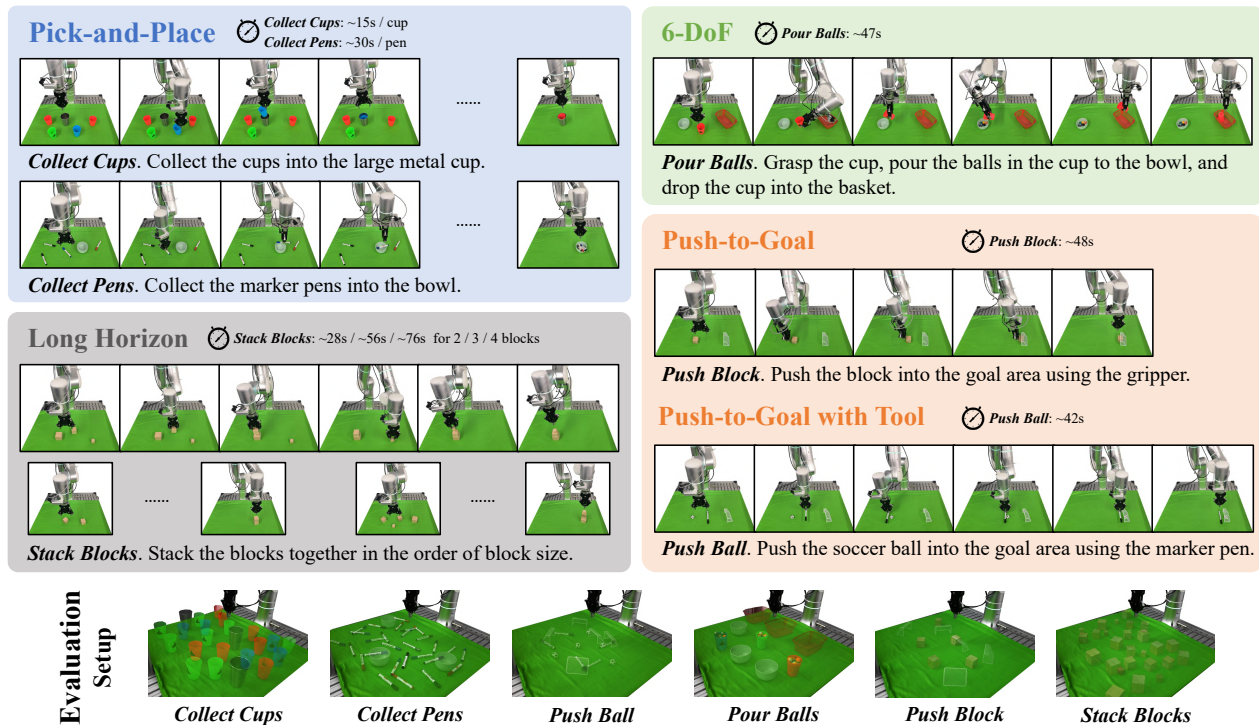
By  $h_E$ , the voxelized point cloud  $\mathcal{O}^t$  is encoded to sparse point features  $\mathcal{F}_P^t$  in an efficient way, avoiding redundant computing on huge empty space.  $\mathcal{F}_P^t$  is then fed into the transformer  $h_T$  as sparse tokens. For  $\mathcal{O}^t$  cropped in  $1 \times 1 \times 1\text{m}^3$  space,  $\mathcal{F}_P^t$  contains only 60 ~ 80 tokens. Although the token number is less than the one in ACT [58] (300 per image), experiments in §IV-F show that point cloud based ACT still outperforms the original implementation.

#### B. Transformer with Sparse Point Tokens

We adopt transformer [47] to implement the mapping from point features  $\mathcal{F}_P^t$  to action features  $\mathcal{F}_A^t$ . While the positional encoding for image tokens is dense and natural, sparse point tokens could not be processed in the same manner. We instead introduce sparse positional encoding for point tokens.

Let  $(x, y, z)$  be the coordinate of the point token  $P$  with  $d$ -dimension, the position of  $P$  is defined as

$$\begin{aligned} pos_k &= \frac{k}{v} + c, \quad k \in \{x, y, z\}, \\ pos &= [pos_x, pos_y, pos_z], \end{aligned} \quad (1)$$



**Fig. 3:** Definition of the tasks in the experiments. During evaluation, each task is randomly initialized within the robot workspace. For each task, only 3 to 5 setups from the evaluations are depicted in the figure for clarity.

where  $c$  and  $v$  are fixed offsets, and  $[\cdot]$  stands for vector concatenation. The encoding dimension along each axis is  $d_x = d_y = \lfloor d/3 \rfloor$ ,  $d_z = d - d_x - d_y$ . The position encoding of  $P$  is computed by  $SPE = [SPE^x, SPE^y, SPE^z]$  where

$$\begin{cases} SPE_{(pos,2i)}^k = \sin \frac{pos_k}{10000^{2i/d_k}} \\ SPE_{(pos,2i+1)}^k = \cos \frac{pos_k}{10000^{2i/d_k}} \end{cases}, \quad k \in \{x, y, z\} \quad (2)$$

With the help of sparse positional encoding, we effectively capture intricate 3D spatial relationships among unordered points, which enables seamless embedding of the 3D features into conventional transformers.

The transformer  $h_T$  utilizes an encoder-decoder architecture, taking point features  $\mathcal{F}_P^t$  as input tokens without other proprioceptive signals. In the transformer decoding step, we use one readout token to query action features  $\mathcal{F}_A^t$ .

### C. Diffusion as Action Decoder

The action decoder  $h_D$  is implemented as a denoising process by diffusion [4, 18, 22]. Conditioning on  $\mathcal{F}_A^t$ ,  $h_D$  denoises the Gaussian noises  $\mathcal{N}(0, \sigma^2 I)$  to actions  $\mathcal{A}^t$  iteratively. The denoising process of step  $k$  is

$$\mathcal{A}_k^t = \alpha(\mathcal{A}_k^t - \gamma \epsilon_\theta(\mathcal{O}^t, \mathcal{A}_k^t, k) + \mathcal{N}(0, \sigma^2 I)), \quad (3)$$

where  $\epsilon_\theta$  is a network predicting noises with parameters  $\theta$ ,  $\alpha$ ,  $\gamma$  and  $\sigma$  are hyperparameters related to  $k$  in noise schedule. The objective function is the simplified objective in [18]. We use the DDIM scheduler [44] to accelerate the inference speed in real-world experiments.

The regression head is also frequently employed due to its simplicity [11, 14, 21, 58], whereas the diffusion head

excels in handling scenes with multiple targets. Moreover, diffusion produces diverse trajectories to the same target, as opposed to averaging learned trajectories [4].

For all tasks in our experiments, RISE adopts a unified action representation in the camera coordinate system, which is composed of translations, rotations, and gripper widths. We opt for absolute position for translation and 6D representation [59] for rotation with continuity considerations.

## IV. EXPERIMENTS

### A. Setup

In the experiments, we use a Flexiv Rizon robotic arm<sup>1</sup> equipped with a Dahuan AG-95 gripper<sup>2</sup> for interacting with objects. Two Intel RealSense D435 RGBD cameras<sup>3</sup> are installed for scene perception. One global camera is positioned in front of the robot, while the other in-hand camera is mounted on the end-effector of the arm. For 3D perception, only the global camera is used to generate a noisy single-view partial point cloud; while for image-based policies, both cameras are used for a better understanding of spatial geometries. All devices are linked to a workstation with an Intel Core i9-10900K CPU and an NVIDIA RTX 3090 GPU for both data collection and evaluation.

We carefully designed 6 tasks from 4 types for the experiments as illustrated in Fig. 3: the *pick-and-place* tasks (*Collect Cups* and *Collect Pens*), the *6-DoF* tasks (*Pour Balls*), the *push-to-goal* tasks (*Push Block* and *Push Ball*)

<sup>1</sup><https://www.flexiv.com/product/rizon>

<sup>2</sup><https://en.dh-robotics.com/product/ag>

<sup>3</sup><https://www.intelrealsense.com/depth-camera-d435>

and the *long-horizon* tasks (**Stack Blocks**). The data collection setup and process are the same as [9], *i.e.*, end-effector tele-operations using a haptic device. Unless specifically stated, we gathered 50 expert demonstrations for each task as the training data for the policy, and each policy was tested for 20 consecutive trials to evaluate its performance. During evaluations, objects in the task are randomly initialized within the robot workspace of approximately 50cm  $\times$  70cm.

### B. Pick-and-Place Tasks

As the cornerstone of robotic manipulations, *pick-and-place* tasks emphasize precise manipulations of objects and efficient generalizations of the robot policy within the workspace. As shown in Fig. 3, the **Collect Cups** task is designed to evaluate the policy performance in predicting the translation part of the action, whereas the **Collect Pens** task is utilized to further assess the ability of the policy in predicting the planner rotation part of the action.

We employ two representative image-based policies as our baselines: ACT [58] and Diffusion Policy [4]. We also evaluate a keyframe-based 3D policy Act3D [11], the current state-of-the-art policy on RL Bench [20] in the experiments. For Act3D, we utilize a simple action planner to execute the predicted keyposes, preventing collisions with other objects in the workspace.

For each of the five scenarios where there are 1, 2, 3, 4, and 5 objects in the workspace, we collected 10 demonstrations, resulting in a total of 50 demonstrations used for the policy training. During evaluations, we conducted 10 trials for each of the five scenarios, tallying the number of cups placed into the large metal cup or the number of pens placed into the bowl, and calculating the completion rate. For each object, we imposed a runtime limit of 20 keyframes for keyframe-based policies and 300 steps for continuous-control policies.




Fig. 4: Experimental results of the *pick-and-place* tasks.

The evaluation results are depicted in Fig. 4. In the **Collect Cups** task, RISE achieves a completion rate of over 90% when the number of cups is less than 3. Even in complex environments with 4 or 5 cups, RISE maintains a completion rate of over 65%, showcasing capabilities surpassing all baselines. Similar trends are observed in the **Collect Pens** task: RISE consistently outperforms all baselines, demonstrating its ability to not only predict the translation part but also accurately forecast planner rotation. We also discover that Act3D performs comparably to image-based baselines. Moreover, given that Act3D requires specially designed motion planners for more complicated actions and

cannot provide immediate responses to sudden changes in the environment, we therefore only employ ACT and Diffusion Policy as baselines in our subsequent experiments.

### C. 6-DoF Tasks

The **6-DoF Pour Balls** task is designed to test the ability of robot policies to forecast actions with complex spatial rotations, instead of the simple planner rotation in the *pick-and-place* tasks. As shown in Fig. 3, the robotic arm needs to undergo complex spatial rotations to complete the task, at times approaching its kinematic limits. During evaluations, 10 balls are initialized in the cup. Besides the action success rates of the policy, the number of balls poured into the cup is also recorded to calculate the completion rates. The runtime limit for this task is set as 1200 steps.

The experimental results are shown in Tab. I. From the action success rates, it is evident that RISE can learn actions with complex spatial rotations more effectively compared to image-based policies. Additionally, its execution of the pouring action is more precise regarding pouring positions, resulting in higher task completion rates. This also highlights the effectiveness of 3D perception, which can capture more accurate spatial relationships between objects.


Method	Success Rate (%)			Completion Rate (%)	
	Grasp	Pour	Place	Overall	If Poured
ACT [58]	30	30	0	13.0	43.3
Diffusion Policy [4]	55	55	35	30.5	55.5
RISE (ours)	<b>80</b>	<b>80</b>	<b>70</b>	<b>49.0</b>	<b>61.3</b>

TABLE I: Experimental results of the *6-DoF* task **Pour Balls**. “If Poured” denotes the completion rate in the context of having completed the pouring action.

### D. Push-to-Goal Tasks

Robot policies should generate immediate feedback to environmental dynamics, enabling adaptation to object movements within the environment to accomplish tasks. To this end, we designed two *push-to-goal* tasks **Push Block** and **Push Ball**, as shown in Fig. 3. Furthermore, the **Push Ball** task requires the robot to use a tool (a marker pen) to complete the task, thereby testing the ability of the policy to utilize tools. In the evaluation process, we compute the distance  $d$  from the object center to the goal area. If the object center is within the goal area ( $d = 0$ ), it is considered a success, as shown in Tab. II (left). We then calculate the task success rate and the average distance as metrics. The runtime limit for this task is set as 1200 steps.

The evaluation results are shown in Tab. II (right). In the **Push Block** task, RISE slightly surpasses Diffusion Policy in terms of success rate, while pushing the block closer to the goal area. However, in the **Push Ball** task, RISE outperforms Diffusion Policy by a significant margin, demonstrating its effective 3D perception of object position changes and ability to rapidly adjust policy action outputs accordingly. We also observe that ACT struggles in both tasks and frequently causes the robot to make hard contact with the objects and triggers the mechanical emergency stop, which might be attributed to its imprecision in scene perception.



Task	Method	Success Rate (%)	Average Distance $d$ (cm) ↓
<i>Push Block</i>	Diffusion Policy [4]	50	5.67
	RISE (ours)	<b>55</b>	<b>3.51</b>
<i>Push Ball</i>	Diffusion Policy [4]	30	6.05
	RISE (ours)	<b>60</b>	<b>4.89</b>

TABLE II: Evaluation metrics illustrations (left) and experimental results (right) of the *push-to-goal* tasks *Push Block* and *Push Ball*.

### E. Long-Horizon Tasks

Long-horizon tasks are crucial in robotic manipulations as they highlight the impact of error accumulation in actions over extended horizons, providing insights into the robustness and adaptability of the policy. Therefore, we designed the *long-horizon* task *Stack Blocks* to assess the policy’s ability in this regard, given that blocks are more likely to topple as the stack grows. We emphasize that this task is more challenging than previous *pick-and-place* tasks because (a) some blocks are only slightly smaller than gripper width, requiring precise controls in grasping; (b) the policy needs to recognize the sizes of the blocks and select a suitable stacking order to ensure the stability of the resulting stack; and (c) as the stack grows, the policy needs to dynamically adjust the placement height of the blocks to ensure that they do not collide with existing blocks and can be smoothly placed on top of them.

For three cases where there are 2, 3, and 4 blocks in the workspace, we collected 10, 20, and 20 demonstrations respectively, totaling 50 demonstrations for the policy training. During evaluations, 10 trials are conducted for each case, and the average number of successfully stacked blocks is reported to measure the policy performance. We set the runtime limit for the task as 600, 1200, and 1800 steps for each case respectively.

Method	# Blocks		
	2	3	4
ACT [58]	0.6 / 1	0.5 / 2	0.3 / 3
Diffusion Policy [4]	0.7 / 1	0.5 / 2	0.5 / 3
RISE (ours)	<b>0.8 / 1</b>	<b>1.5 / 2</b>	<b>0.9 / 3</b>

TABLE III: Experimental results (average stacked blocks) of the *long-horizon* task *Stack Blocks*.


The experimental results are shown in Tab. III. In the simple scenario with only two blocks, all policies yielded similar results; however, as the number of blocks increased, RISE progressively surpassed the baselines significantly, showcasing its ability to adapt well to *long-horizon* tasks and control accumulated errors. Moreover, we observe that the baselines exhibit a higher frequency of the aforementioned issues (b) and (c) compared to RISE, implying that powered by 3D perception, RISE has a deeper understanding of the scene, and its predictions of actions are also more precise.

### F. Effectiveness of 3D Perception

In this section, we explore how 3D perception enhances the performance of robot manipulation policies on the *Collect Cups* task with 5 cups. We replace the image encoder of the image-based policies ACT and Diffusion Policy with the

Method	3D	# Cameras	Completion Rate (%)
ACT [58]		2	12
	✓	1	32 ↑ <sup>20</sup>
Diffusion Policy [4]		2	24
	✓	1	36 ↑ <sup>12</sup>
DP3* [53]	✓	1	-
Act3D [11]	✓	1	28
RISE (ours)	✓	1	<b>66</b>

TABLE IV: Effectiveness test of 3D perception on the *Collect Cups* task with 5 cups (10 trials). 2D version of policies take images from both global and in-hand cameras as input. \* DP3 fails to learn in our setting, see text for a more detailed analysis.



Method	Completion Rate (%)	
	Axis-wise	Natural
DP3, hor. 4	0	0
DP3, hor. 8	0	0
DP3, hor. 16	20	0
DP3, hor. 24	40	0
RISE	<b>80</b>	<b>100</b>

TABLE V: Analysis of the failures of DP3 in our experiments. (left) Illustrations of the axis-wise and natural action. (right) Results of the *Collect Cups* task with 1 cup when using demonstrations with different action representations for training (10 trials).

sparse 3D encoder used in RISE. The experiment results are shown in Tab. IV. We observe a significant improvement in the performance of ACT and Diffusion Policy after applying 3D perception even with fewer camera views, surpassing the 3D policy Act3D, which reflects the effectiveness of our 3D perception module in manipulation policies.

We also evaluate the recently proposed DP3 [53] in this experimental setting. However, DP3 appears to struggle in learning meaningful actions from our demonstration data. After communicating with the authors, one potential reason is that they are using RealSense L515 in their original experiments, and we adopted RealSense D435 in our experiments. The point cloud from D435 is noisier, making it more challenging for networks to learn. By using sparse convolution, RISE is more robust to the noise in the point cloud. Besides, after delving into their real robot experiments, we found that instead of natural actions, axis-wise actions are used in their demonstration data, as illustrated in Tab. V (left). Hence, we collect 50 demonstrations on the *Collect Cups* task with 1 cup using axis-wise and natural action representations respectively. These demonstrations are then used for DP3 policy training. After carefully tuning hyper-parameters such as horizons and color utilizations, we report the evaluation results in Tab. V, with the best completion rate of 40%. We suspect that the limited network capacity of the 3D




TABLE VI: Generalization test setup and experimental results of the *Collect Pens* task with 1 pen (10 trials).

encoder of DP3 prevents it from modeling the diverse state-action pairs present in the real-world human teleoperated demonstrations, leading it to only handle a smaller set of state-action pairs under the axis-wise action representations. On the contrary, RISE can handle real-world demonstrations with various action representations and maintain satisfactory performances. Lastly, compared to the evaluation setup in the DP3 paper, we allow objects to be placed anywhere in the entire workspace. This results in a greater variation of object locations, making the task more challenging.

### G. Generalization Test

We evaluate the generalization ability of different methods on the *Collect Pens* task with 1 pen under different levels of environmental disturbances as follows.

- L1.** Generalize to objects with similar shapes but different colors. In this task, we replace the original green bowl with a pink one, which is denoted as **Bowl** in Tab. VI.
- L2.** Generalize to different light conditions in the environment, denoted as **Light** in Tab. VI.
- L3.** Generalize to new workspace configurations in the environment. In this task, we elevate the workspace by 13cm to form a new workspace configuration, which is denoted as **Height** in Tab. VI.
- L4.** Generalize to new camera viewpoints, denoted as **CamView** in Tab. VI.

The graphical illustration of the environmental disturbances and the evaluation results are shown in Tab. VI. We can observe that the image-based policies can achieve a decent L1-level and some L2-level generalizations, but they cannot reach the L3-level and L4-level generalizations involving spatial transformations. Act3D, as a 3D policy, demonstrates good generalization up to L3-level disturbances; however, it nearly completely fails in the L4-level generalization test. RISE exhibits strong generalization abilities across all levels of testing, even in the most challenging L4-level tests involving changes in the camera view.

## V. CONCLUSION

In this paper, we present RISE, an efficient end-to-end policy utilizing 3D perception for real-world robot manipulation. RISE compresses point clouds with a sparse 3D encoder, followed by sparse positional encoding and a transformer to obtain action features. The features are decoded into continuous actions by a diffusion head. RISE significantly outperforms currently representative 2D and 3D policies in multiple tasks, demonstrating great advantages in both accuracy and efficiency. Our ablations verify the effectiveness of

Method	Completion Rate (%)					
	Original	Disturbance				CamView
		Bowl	Light	Height	Height	
ACT [58]	80	70 $\downarrow 10$	40 $\downarrow 40$	0 $\downarrow 80$	0 $\downarrow 80$	0 $\downarrow 80$
Diffusion Policy [4]	70	50 $\downarrow 20$	30 $\downarrow 40$	0 $\downarrow 70$	0 $\downarrow 70$	0 $\downarrow 70$
Act3D [11]	70	40 $\downarrow 30$	60 $\downarrow 10$	50 $\downarrow 20$	10 $\downarrow 60$	10 $\downarrow 60$
RISE (ours)	<b>90</b>	<b>80</b> $\downarrow 10$	<b>80</b> $\downarrow 10$	<b>80</b> $\downarrow 10$	<b>50</b> $\downarrow 40$	<b>50</b> $\downarrow 40$

3D perception and the generalization of RISE under different levels of environmental disturbances. We hope our baseline inspires the integration of 3D perception into real-world policy learning.

## VI. ACKNOWLEDGEMENT

This work was supported by the National Key Research and Development Project of China (No. 2022ZD0160102), the National Key Research and Development Project of China (No. 2021ZD0110704), Shanghai Artificial Intelligence Laboratory, XPLOER PRIZE grants.

## REFERENCES

- [1] Anthony Brohan et al. “RT-1: Robotics Transformer for Real-World Control at Scale”. In: *Robotics: Science and Systems*. 2023.
- [2] Shizhe Chen et al. “PolarNet: 3D Point Clouds for Language-Guided Robotic Manipulation”. In: *Conference on Robot Learning*. PMLR. 2023, pp. 1761–1781.
- [3] Xiaozhi Chen et al. “Multi-View 3D Object Detection Network for Autonomous Driving”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2017, pp. 1907–1915.
- [4] Cheng Chi et al. “Diffusion Policy: Visuomotor Policy Learning via Action Diffusion”. In: *Robotics: Science and Systems*. 2023.
- [5] Christopher Choy, JunYoung Gwak, and Silvio Savarese. “4D Spatio-Temporal ConvNets: Minkowski Convolutional Neural Networks”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019, pp. 3075–3084.
- [6] Angela Dai et al. “ScanNet: Richly-Annotated 3D Reconstructions of Indoor Scenes”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2017, pp. 5828–5839.
- [7] Danny Driess et al. “Reinforcement Learning with Neural Radiance Fields”. In: *Advances in Neural Information Processing Systems* 35 (2022), pp. 16931–16945.
- [8] Hao-Shu Fang et al. “AnyGrasp: Robust and Efficient Grasp Perception in Spatial and Temporal Domains”. In: *IEEE Transactions on Robotics* (2023).
- [9] Hao-Shu Fang et al. “RH20T: A Comprehensive Robotic Dataset for Learning Diverse Skills in One-Shot”. In: *IEEE International Conference on Robotics and Automation*. 2024.
- [10] Hongjie Fang et al. “AirExo: Low-Cost Exoskeletons for Learning Whole-Arm Manipulation in the Wild”. In: *IEEE International Conference on Robotics and Automation*. 2024.
- [11] Théophile Gervet et al. “Act3D: 3D Feature Field Transformers for Multi-Task Robotic Manipulation”. In: *Conference on Robot Learning*. PMLR. 2023, pp. 3949–3965.
- [12] Ankit Goyal et al. “RVT: Robotic View Transformer for 3D Object Manipulation”. In: *Conference on Robot Learning*. PMLR. 2023, pp. 694–710.
- [13] Benjamin Graham, Martin Engelcke, and Laurens Van Der Maaten. “3D Semantic Segmentation with Submanifold Sparse Convolutional Networks”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2018, pp. 9224–9232.
- [14] Pierre-Louis Guhur et al. “Instruction-Driven History-Aware Policies for Robotic Manipulations”. In: *Conference on Robot Learning*. PMLR. 2022, pp. 175–187.
- [15] Huy Ha, Pete Florence, and Shuran Song. “Scaling Up and Distilling Down: Language-Guided Robot Skill Acquisition”. In: *Conference on Robot Learning*. PMLR. 2023, pp. 3766–3777.

- [16] Abdullah Hamdi, Silvio Giancola, and Bernard Ghanem. "MVTN: Multi-View Transformation Network for 3D Shape Recognition". In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2021, pp. 1–11.
- [17] Kaiming He et al. "Deep Residual Learning for Image Recognition". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2016, pp. 770–778.
- [18] Jonathan Ho, Ajay Jain, and Pieter Abbeel. "Denosing Diffusion Probabilistic Models". In: *Advances in Neural Information Processing Systems* 33 (2020), pp. 6840–6851.
- [19] Stephen James et al. "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*. 2022, pp. 13739–13748.
- [20] Stephen James et al. "RLBench: The Robot Learning Benchmark & Learning Environment". In: *IEEE Robotics and Automation Letters* 5.2 (2020), pp. 3019–3026.
- [21] Eric Jang et al. "BC-Z: Zero-Shot Task Generalization with Robotic Imitation Learning". In: *Conference on Robot Learning*. PMLR, 2021, pp. 991–1002.
- [22] Michael Janner et al. "Planning with Diffusion for Flexible Behavior Synthesis". In: *International Conference on Machine Learning*. PMLR, 2022, pp. 9902–9915.
- [23] Bo Li, Tianlei Zhang, and Tian Xia. "Vehicle Detection from 3D Lidar Using Fully Convolutional Network". In: *Robotics: Science and Systems*. 2016.
- [24] Yecheng Jason Ma et al. "LIV: Language-Image Representations and Rewards for Robotic Control". In: *International Conference on Machine Learning*. PMLR, 2023, pp. 23301–23320.
- [25] Yecheng Jason Ma et al. "VIP: Towards Universal Visual Reward and Representation via Value-Implicit Pre-Training". In: *International Conference on Learning Representations*. 2023.
- [26] Arjun Majumdar et al. "Where are We in the Search for an Artificial Visual Cortex for Embodied Intelligence?" In: *ICRA 2023 Workshop on Pretraining for Robotics*. 2023.
- [27] Ajay Mandlekar et al. "What Matters in Learning from Offline Human Demonstrations for Robot Manipulation". In: *Conference on Robot Learning*. PMLR, 2021, pp. 1678–1690.
- [28] Oier Mees et al. "CALVIN: A Benchmark for Language-Conditioned Policy Learning for Long-Horizon Robot Manipulation Tasks". In: *IEEE Robotics and Automation Letters* 7.3 (2022), pp. 7327–7334.
- [29] Ben Mildenhall et al. "NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis". In: *Communications of the ACM* 65.1 (2021), pp. 99–106.
- [30] Suraj Nair et al. "R3M: A Universal Visual Representation for Robot Manipulation". In: *Conference on Robot Learning*. PMLR, 2022, pp. 892–909.
- [31] Abhishek Padalkar et al. "Open X-Embodiment: Robotic Learning Datasets and RT-X Models". In: *IEEE International Conference on Robotics and Automation*. 2024.
- [32] Xuran Pan et al. "3D Object Detection with PointFormer". In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2021, pp. 7463–7472.
- [33] Jyothish Pari et al. "The Surprising Effectiveness of Representation Learning for Visual Imitation". In: *Robotics: Science and Systems*. 2022.
- [34] Dean A Pomerleau. "ALVINN: An Autonomous Land Vehicle in a Neural Network". In: *Advances in Neural Information Processing Systems*. Vol. 1. 1988.
- [35] Charles R Qi et al. "PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2017, pp. 652–660.
- [36] Guocheng Qian et al. "PointNeXt: Revisiting PointNet++ with Improved Training and Scaling Strategies". In: *Advances in Neural Information Processing Systems* 35 (2022), pp. 23192–23204.
- [37] Ilija Radosavovic et al. "Real-World Robot Learning with Masked Visual Pre-Training". In: *Conference on Robot Learning*. PMLR, 2022, pp. 416–426.
- [38] Rouhollah Rahmatizadeh et al. "Vision-Based Multi-Task Manipulation for Inexpensive Robots using End-to-End Learning from Demonstration". In: *IEEE International Conference on Robotics and Automation*. IEEE, 2018, pp. 3758–3765.
- [39] Scott E. Reed et al. "A Generalist Agent". In: *Transactions on Machine Learning Research* (2022).
- [40] Gernot Riegler, Ali Osman Ulusoy, and Andreas Geiger. "OctNet: Learning Deep 3D Representations at High Resolutions". In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2017, pp. 3577–3586.
- [41] William Shen et al. "Distilled Feature Fields Enable Few-Shot Language-Guided Manipulation". In: *Conference on Robot Learning*. PMLR, 2023, pp. 405–424.
- [42] Mohit Shridhar, Lucas Manuelli, and Dieter Fox. "Perceiver-Actor: A Multi-Task Transformer for Robotic Manipulation". In: *Conference on Robot Learning*. PMLR, 2022, pp. 785–799.
- [43] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. "Learning Structured Output Representation using Deep Conditional Generative Models". In: *Advances in Neural Information Processing Systems* 28 (2015).
- [44] Jiaming Song, Chenlin Meng, and Stefano Ermon. "Denosing Diffusion Implicit Models". In: *The International Conference on Learning Representations*. 2021.
- [45] Martin Sundermeyer et al. "Contact-Graspnet: Efficient 6-DoF Grasp Generation in Cluttered Scenes". In: *IEEE International Conference on Robotics and Automation*. IEEE, 2021, pp. 13438–13444.
- [46] Octo Model Team et al. *Octo: An Open-Source Generalist Robot Policy*. 2023.
- [47] Ashish Vaswani et al. "Attention is All You Need". In: *Advances in Neural Information Processing Systems* 30 (2017).
- [48] Chen Wang et al. "DexCap: Scalable and Portable Mocap Data Collection System for Dexterous Manipulation". In: *Proceedings of Robotics: Science and Systems (RSS)* (2024).
- [49] Peng-Shuai Wang et al. "O-CNN: Octree-Based Convolutional Neural Networks for 3D Shape Analysis". In: *ACM Transactions On Graphics* 36.4 (2017), pp. 1–11.
- [50] Zhirong Wu et al. "3D ShapeNets: A Deep Representation for Volumetric Shapes". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2015, pp. 1912–1920.
- [51] Zhou Xian et al. "ChainedDiffuser: Unifying Trajectory Diffusion and Keypose Prediction for Robotic Manipulation". In: *Conference on Robot Learning*. PMLR, 2023, pp. 2323–2339.
- [52] Jianglong Ye, Naiyan Wang, and Xiao-long Wang. "FeatureNeRF: Learning Generalizable NeRFs by Distilling Foundation Models". In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2023, pp. 8962–8973.
- [53] Yanjie Ze et al. "3D Diffusion Policy: Generalizable Visuomotor Policy Learning via Simple 3D Representations". In: *Proceedings of Robotics: Science and Systems (RSS)* (2024).
- [54] Yanjie Ze et al. "GNFactor: Multi-Task Real Robot Learning with Generalizable Neural Feature Fields". In: *Conference on Robot Learning*. PMLR, 2023, pp. 284–301.
- [55] Tianhao Zhang et al. "Deep Imitation Learning for Complex Manipulation Tasks from Virtual Reality Teleoperation". In: *IEEE International Conference on Robotics and Automation*. IEEE, 2018, pp. 5628–5635.
- [56] Tong Zhang et al. "A Universal Semantic-Geometric Representation for Robotic Manipulation". In: *CoRL*. Vol. 229. Proceedings of Machine Learning Research. PMLR, 2023, pp. 3342–3363.
- [57] Hengshuang Zhao et al. "Point Transformer". In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2021, pp. 16259–16268.
- [58] Tony Z Zhao et al. "Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware". In: *Robotics: Science and Systems*. 2023.
- [59] Yi Zhou et al. "On the Continuity of Rotation Representations in Neural Networks". In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019, pp. 5745–5753.
- [60] Yin Zhou and Oncel Tuzel. "VoxelNet: End-to-End Learning for Point Cloud Based 3D Object Detection". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2018, pp. 4490–4499.
- [61] Yifeng Zhu et al. "Learning generalizable manipulation policies with object-centric 3d representations". In: *7th Annual Conference on Robot Learning*. 2023.
- [62] Brianna Zitkovich et al. "RT-2: Vision-Language-Action Models Transfer Web Knowledge to Robotic Control". In: *Conference on Robot Learning*. 2023, pp. 2165–2183.