

CordViP: Correspondence-based Visuomotor Policy for Dexterous Manipulation in Real-World

Yankai Fu^{1†}, Qiuxuan Feng^{1†}, Ning Chen^{1†}, Zichen Zhou¹, Mengzhen Liu¹,
Mingdong Wu¹, Tianxing Chen², Shanyu Rong¹, Jiaming Liu¹, Hao Dong¹, Shanghang Zhang^{1,3✉}

¹State Key Laboratory of Multimedia Information Processing, School of Computer Science, Peking University,

²The University of Hong Kong, ³Beijing Academy of Artificial Intelligence

[†]Equal Contribution, [✉]Corresponding author

<https://aureleopku.github.io/CordViP>

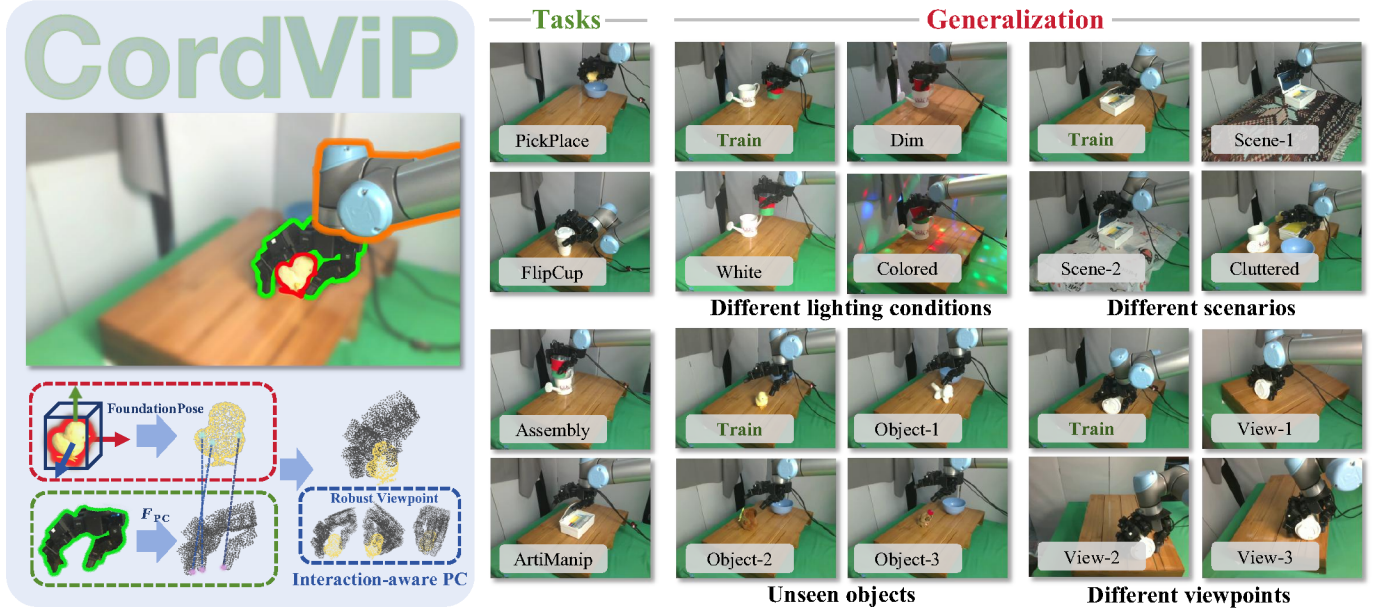


Fig. 1: We propose CordViP, a correspondence-based visuomotor policy for dexterous manipulation in the real world. (a) Left: We present the interaction-aware point clouds, which demonstrate robustness to different viewpoints while establishing correspondences between the object and the hand. (b) Right: Our method achieves promising results across multiple real-world dexterous manipulation tasks, showcasing exceptional generalization capabilities.

Abstract—Achieving human-level dexterity in robots is a key objective in the field of robotic manipulation. Recent advancements in 3D-based imitation learning have shown promising results, providing an effective pathway to achieve this goal. However, obtaining high-quality 3D representations presents two key problems: (1) the quality of point clouds captured by a single-view camera is significantly affected by factors such as camera resolution, positioning, and occlusions caused by the dexterous hand; (2) the global point clouds lack crucial contact information and spatial correspondences, which are necessary for fine-grained dexterous manipulation tasks. To eliminate these limitations, we propose CordViP, a novel framework that constructs and learns correspondences by leveraging the robust 6D pose estimation of objects and robot proprioception. Specifically, we first introduce the interaction-aware point clouds, which establish correspondences between the object and the hand. These point clouds are then used for our pre-training policy, where we also incorporate object-centric contact maps and hand-arm coordination information, effectively capturing both spatial

and temporal dynamics. Our method demonstrates exceptional dexterous manipulation capabilities, achieving state-of-the-art performance in six real-world tasks, surpassing other baselines by a large margin. Experimental results also highlight the superior generalization and robustness of CordViP to different objects, viewpoints, and scenarios. Code and videos are available on <https://aureleopku.github.io/CordViP>.

I. INTRODUCTION

Dexterous manipulation is a fundamental capability in human daily life such as assembling small parts and opening boxes. Achieving human-level dexterity in real-world scenarios is crucial for integrating robots into everyday human activities. Recent advancements in imitation learning have demonstrated significant potential in various robotic manipulation tasks. Some existing methods leverage 2D images as input to directly predict actions [7, 70, 54, 63]. While these vision-based

imitation learning approaches are capable of handling a wide range of tasks, they typically demand extensive demonstrations [19, 7, 13, 57] and, at the same time, fail to capture intricate spatial relationships and 3D structures essential for dexterous manipulation [69, 23, 49, 11], limiting their ability to perform complex, fine-grained tasks.

Recent research has increasingly focused on 3D imitation learning, utilizing point clouds and voxels as representations, to help robots perceive 3D environments and reason about spatial relationships [69, 13, 56, 14, 64, 51]. For example, 3D diffusion-based policies, which aim to enhance robots' ability to perform complex manipulation tasks by providing a more accurate and holistic representation of the environment. These methods for obtaining 3D representations in real-world scenarios often depend on a single-view RGB-D camera to generate point clouds, imposing significant demands on both the camera's quality and its position. However, the multi-fingered nature of dexterous hands often leads to occlusions during object manipulation, which not only loses the spatial information of the object point clouds but also obscures critical correspondence information for precise manipulation. Besides, some works have explored the use of tactile sensors to enhance contact information [17, 66, 18, 62, 2]. For instance, Wu et al. [62] proposed a canonical representation of 3D tactile data, which is concatenated with other visual information. Although tactile sensors hold promise, their high cost and susceptibility to interference from factors such as temperature fluctuations and electromagnetic fields undermine their stability and limit their practical applicability in real-world scenarios. Amidst all the challenges, leveraging efficient and robust 3D imitation learning to understand spatial information in complex, dynamic environments is crucial for dexterous manipulation.

In this paper, we propose a **correspondence-based visuomotor policy (CordViP)**, which focuses on the spatial and temporal consistency between the manipulated object and the dexterous hand, even under significant occlusions. CordViP separately extracts the 3D representations of robotic hand's joint structure and the manipulated object, and utilizes the contact map between them as well as arm-hand cooperative motion information to pretrain the observation encoder. This approach enhances the dexterous hand's ability to understand spatial interactions and better facilitate a range of downstream tasks. Specifically, our framework operates in three phases: (1) We leverage the robust 6D pose estimation of objects combined with the robot's proprioceptive state to construct interaction-aware point clouds, providing ideal 3D observations that facilitates effective learning and inference for the visuomotor policy. (2) We pretrain the encoder network on the play data we collected to predict the contact map between the dexterous hand and the object, and reconstruct the cooperative relationship between the hand and the arm for further understanding of spatial and temporal correspondences. (3) We utilize the pre-trained encoder to extract the semantic features of 3D point clouds and robot state, which encapsulates spatial consistency, contact dynamics and collaborative information, and then are used as conditions for the diffusion policy to predict actions.

To comprehensively evaluate our proposed CordViP, we conduct extensive real-world experiments on six dexterous tasks: PickPlace, FlipCup, Assembly, Artimanip, FlipCap and LongHoriManip. Comparative results demonstrate that our method not only exhibits superior effectiveness, but also achieves remarkable performance with a minor number of expert demonstrations, highlighting its capability to learn efficiently from limited data. Furthermore, we find that CordViP generalizes well to various environmental perturbations, including varying lighting conditions, unseen objects, diverse scenarios, and different camera viewpoints, outperforming other baselines by a large margin.

In summary, our contributions are as follows:

- 1) We develop a pipeline based on robust 6D pose estimation of objects and robot proprioception. This pipeline enables the real-time acquisition of complete 3D representations and semantic flow in real-world environments, and can effectively address the challenges posed by occlusions during dexterous manipulation.
- 2) We propose CordViP, a correspondence-based visuomotor policy that utilizes the contact map and hand-arm coordination information to facilitate the understanding of spatial and temporal consistency.
- 3) We demonstrate the effectiveness and generalization of our method through a range of real-world experiments using a dexterous hand.

II. RELATED WORK

A. Dexterous Manipulation

Dexterous Manipulation is a long-standing research topic in robotics that aims to give robots the ability to perform delicate operations like humans [5, 1, 45, 65]. Traditional methods often rely on trajectory optimization based on dynamic models to solve operational problems [26, 33, 58], but these methods have limitations in complex tasks because they simplify the contact dynamics and are difficult to deal with uncertainties in dynamic environments. In contrast, Reinforcement Learning (RL) does not rely on accurate physical models but learns operational policies through interaction with the environment, which is highly adaptable. RL has achieved remarkable results in many dexterous manipulation tasks, such as object reorientation [3, 43, 39, 20] and sequential manipulation [6, 15]. However, RL methods often suffer from several challenges, such as the need for extensive reward engineering and system design, as well as limited generalization to unseen scenarios. Additionally, while Sim-to-Real is a common technique employed in RL, the gap between simulations and the real world degrades the performance of the policies once the models are transferred into real robots [71]. Imitation learning (IL), as another effective learning method, can quickly learn effective control policies by imitating expert demonstrations [67, 54]. In this work, we propose a correspondence-based visual imitation learning policy that utilizes spatial information between various components, enabling the acquisition of complex skills with a minimal number of expert demonstrations.

B. Imitation Learning

Imitation learning (IL) allows a robot to directly learn from experts. Behavioral Cloning (BC) is one of the simplest imitation learning algorithms, which treats the problem of learning behavior as a supervised learning task [40]. The modeling methods commonly used in traditional BC, such as MSE, discretization [29], and K-Means [16], have limitations when modeling complex action distributions. They fail to effectively capture the diversity and nuances of human behavior [38]. Over the past few years, diffusion models have emerged as a new modeling approach in BC, becoming powerful tools that enable robots to learn from demonstrations, handle uncertainty, and perform complex multi-step tasks with precision. From the early applications of DDPMs to the recent innovations in BESO [47], OCTO [52], and CrossFormers [10], these models have continually pushed the boundaries of what's possible in robotic behavior generation. While traditional BC policies typically rely on 2D image-based representations [7, 70, 54, 63, 28, 30], recent advancements have extended imitation learning to 3D visual representations [69, 4, 13, 55, 56, 31]. These 3D approaches provide a more comprehensive understanding of spatial relationships and 3D structures, further enhancing robotic behavior learning.

C. Correspondence Learning

Correspondence refers to the relationship or alignment between different entities or components, with the aim of establishing meaningful connections. Correspondence learning has been shown to improve performance in various robotic tasks, including grasping [36, 9], perception [27, 4], pose estimation [21] and garment manipulation [61]. In this paper, correspondence specifically refers to the alignment between hand-object spatial interaction and hand-arm temporal coordination. By incorporating correspondence, we enhance feature extraction capabilities, thereby enabling more accurate and coordinated movements in downstream tasks.

III. METHOD

The overview of our framework is shown in Figure 2, which operates in three phases: (1) Interaction-aware generation of 3D point clouds. We acquire relatively accurate and complete 3D observations during real-world dexterous manipulation tasks even under significant occlusions, as described in III-B. (2) Contact and coordination-enhanced feature extraction. By leveraging large-scale play data and incorporating contact maps and hand-arm coordination, we improve spatial interaction perception and capture cooperative motion features, detailed in III-C. (3) Correspondence-based diffusion policy. The pre-trained encoder is used to extract 3D representations, which guide the training of a visuomotor policy, as outlined in III-D.

A. Problem Formulation

We formulate our problem as learning a visuomotor policy $\pi : \mathcal{O} \mapsto \mathcal{A}$ from expert demonstrations of the form of $\{(o_1, a_1), (o_2, a_2), \dots, (o_n, a_n)\}$, where \mathcal{O} represents the

robot's observations and \mathcal{A} represents the corresponding actions, allowing the robot to generalize beyond the training data distribution. In our approach, each observation o_t is composed of the object's point cloud \mathbf{P}_{obj} , the hand's point cloud \mathbf{P}_{hand} , and the robot's joint states, including a 6-Dof arm and 16-Dof hand configuration. Unlike previous works that rely on global point clouds for 3D feature extraction, our approach prioritizes capturing the individual information of the arm, hand, and object throughout the manipulation process. As a result, CordViP not only effectively addresses occlusion challenges during dexterous manipulation but also significantly improves the model's ability to comprehend spatial interactions and collaborative dynamics. Furthermore, leveraging these observations, we compute contact maps between the robotic hand and the manipulated objects, as well as capture collaborative interaction information between the arm and hand. These elements are critical for modeling spatial and temporal relationships.

B. Interaction-aware Generation of 3D Point Clouds

Motivated by the superior generalization and efficiency of the 3D-based diffusion policy [69, 56, 68, 4], the key intuition behind our solution is to focus on the interactions between the hand and the manipulated object in 3D space. Although intuitively reasonable, achieving this goal is challenging in practice. On the one hand, real-world point cloud data, typically captured using stereo cameras or low-cost RGB-D scanners, suffers from geometric and semantic loss due to factors such as light reflection, material transparency, and limitations of sensor resolution and viewing angle. On the other hand, during dexterous manipulation with multi-fingered hands, occlusions frequently occur, resulting in the loss of critical contact and interaction information, which is vital for precise and effective manipulation. To this end, we propose the interaction-aware generation of 3D point clouds, enabling the reconstruction of crucial spatial information.

Real-to-Sim for Digital Twin Generation. To achieve the goal of obtaining a complete and accurate static point cloud of the manipulated object, we aim to reconstruct the digital twin from a single-view image [34, 35]. Referring to the approach TripoSR [53], we implement a 3D generation technique, which utilizes its strong priors and broad understanding of visual concepts in the 3D world to generate 3D digital assets. To ensure the accuracy of point cloud flow tracking, we maintain consistency between the geometric and material properties of the reconstructed assets and their real-world counterparts. Subsequently, we uniformly sample points on the surface of the generated digital twin to obtain the initial 3D point clouds, providing a robust and accurate initial observation for both 3D spatial perception and pose tracking.

Pose-Driven Point Cloud Tracking. We have successfully obtained the initial point cloud of the object. However, tracking the object's point cloud in complex, real-world environments poses a significant challenge, particularly in scenarios with severe occlusions. To overcome this, we leverage foundation models to ensure precise pose estimation of the manipulated

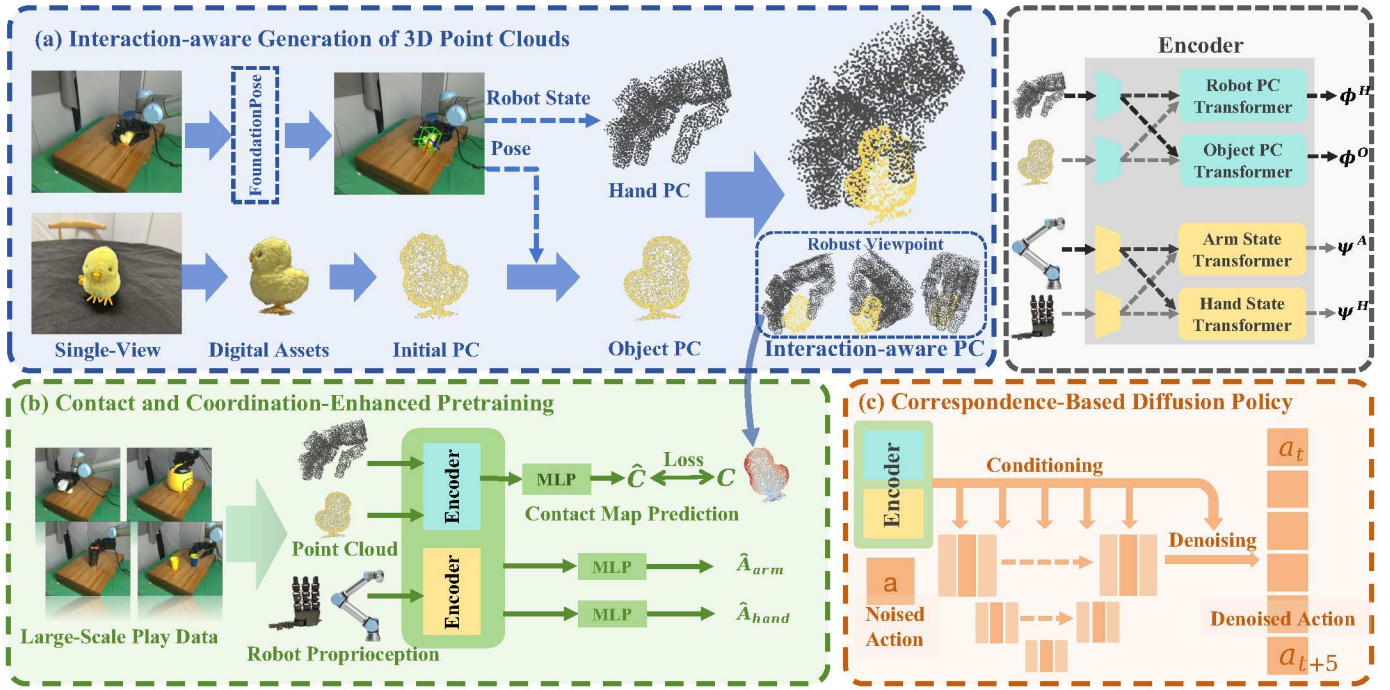


Fig. 2: **Overview Framework** (a) We first employ TripoSR to generate the initial object point cloud and FoundationPose to estimate the 6D pose of the object. In parallel, the hand point cloud is generated based on the robot’s state. They are combined to construct interaction-aware point clouds, which demonstrate robustness to viewpoint variations. (b) During the pre-training phase, the generated point cloud data, combined with the robot’s proprioceptive information, is utilized to enhance spatial understanding and interaction modeling. (c) The pre-trained encoder is subsequently integrated into an imitation learning framework to facilitate downstream tasks in dexterous manipulation.

object, thereby enhancing point cloud tracking. Specifically, we first employ the Segment Anything model [25] to extract the mask of the manipulated object. This is combined with a digital twin generated via real-to-sim techniques, enabling us to use FoundationPose [60] for accurate pose estimation. Using the center of the acquired point cloud as the origin, we then apply the estimated pose to transform the entire point cloud, aligning it with the desired coordinate frame.

Robot Point Cloud Forward Kinematics Model. To obtain the point cloud of the dexterous hand, we develop the robot point cloud forward kinematics model \mathcal{F}_{pc} . We first parse the URDF file to identify individual links of the robot system, and uniformly sample point clouds on the surface for each link. The robot system model is constructed with point clouds denoted as $\{P_{\ell_i}\}_{i=1}^{N_\ell}$, where N_ℓ is the number of links. Following Wei et al. [59], we designed the robot point cloud forward kinematics model that maps any joint configuration to the corresponding pose of the point cloud. In order to focus on the interaction between the hand and the object, we ignore the point cloud of the robotic arm. Therefore, the 3D observation of the hand $\mathbf{P}_{hand} \in \mathbb{R}^{N_P \times 3}$ is defined as:

$$\mathbf{P}_{hand} = \mathcal{F}_{pc}(q, \{P_{\ell_i}\}_{i=1}^{N_\ell}), \quad (1)$$

where N_P represents the downsampled size of the point cloud, which is set to 1024 in practice. Through our practical evaluations, the point cloud forward kinematics model operates




Fig. 3: **Point Clouds Comparison.** We present point clouds of two methods under three different viewpoints. Notably, for better visualization, we have applied color information to the point clouds. However, color information is not used in the policy learning.

at 8Hz, a frequency that is sufficiently high for our experimental setup, where inference is conducted at 5Hz.

As shown in Figure 3, interaction-aware point clouds significantly enhance the quality of 3D observations compared to RGB-D synthesized point clouds, while demonstrating strong robustness to different viewpoints.

C. Contact and Coordination-Enhanced Feature Extraction

Interaction-aware Generation of 3D Point Clouds provides us with accurate and complete point cloud observation. However, fine-grained dexterous manipulation often requires

more detailed information. pre-training based on contact and coordination, in the context, can encourage the encoder to learn the intrinsic structures, facilitating more effective feature extraction, and promoting the learning of various downstream tasks.

Contact Map Synthesis. A contact map serves as a critical piece of information in dexterous manipulation tasks, which captures the interaction between the hand and the object. Given the complete point clouds, we first calculate object-centric contact map \mathcal{C} as the normalized distances from the object's surface point to the hand surface. Given the point clouds of the object \mathbf{P}_{obj} and the point cloud of the hand \mathbf{P}_{hand} , the aligned distance $\mathcal{D}_{(\mathcal{O}, \mathcal{H})}$ between each point v_o on the object surface and the surface of the dexterous hand is defined as follows:

$$\mathcal{D}_{(\mathcal{O}, \mathcal{H})} = \min_{v_h \in \mathbf{P}_{hand}} e^{\gamma(1 - \langle v_h - v_o, n_o \rangle)} \|v_o - v_h\|_2, \quad (2)$$

where n_o is the surface normal of the object, which is computed using the K-Nearest Neighbors methods [8] by considering the local geometric properties of the point cloud. γ is the scaling factor, which is empirically set to 1.

Based on the aligned distance, we compute the contact map following Jiang et al. [24]:

$$\mathcal{C} = 1 - 2(\text{Sigmod}(\theta \cdot \mathcal{D}_{(\mathcal{O}, \mathcal{H})}) - 0.5) \quad (3)$$

where θ is the scaling factor and each point's contact value $c_i \in \mathcal{C}$ is bounded within the $[0, 1]$ range.

Point Cloud Feature Extraction. We employ two encoders with identical architecture to extract point cloud embeddings, denoted as $f_{\theta_O}(\mathbf{P}_{obj})$ and $f_{\theta_H}(\mathbf{P}_{hand})$, for the object and hand, respectively. In detail, we adopt PointNet [41] as the point cloud encoder, which excels in capturing local structures and integrating global information. To establish correspondences between the hand and the object features, we apply two multi-head cross-attention transformers g_{θ_O} , g_{θ_H} to fuse their respective embeddings, which maps the hand and object features to two sets of aligned representations, denoted as ϕ^R and ϕ^O :

$$\begin{aligned} \phi^H &= g_{\theta_H}(f_{\theta_H}(\mathbf{P}_{hand}), f_{\theta_O}(\mathbf{P}_{obj})) + f_{\theta_H}(\mathbf{P}_{hand}) \\ \phi^O &= g_{\theta_O}(f_{\theta_O}(\mathbf{P}_{obj}), f_{\theta_H}(\mathbf{P}_{hand})) + f_{\theta_O}(\mathbf{P}_{obj}) \end{aligned} \quad (4)$$

To help the point cloud encoder learn the intrinsic features, we designed a contact map prediction task. Since 3D point cloud observations implicitly contain contact information, we utilize a three-layer MLP to predict the object-centric contact map \mathcal{C}^{pred} given the point cloud observations of both the hand and the object. The MSE loss is calculated between \mathcal{C} and \mathcal{C}^{pred} . This pre-training approach enables the encoder to learn the interactions and relationships within the environment.

Hand-Arm Coordination Enhancement. To help the robot system learn the features of hand-arm coordination, we also propose a correspondence-based design for action prediction. The arm and hand states are first projected into vectors of identical dimensionality through a linear layer, after which the same cross-attention transformers are employed to establish

correspondences between the hand and the arm. We predict the action sequence of the robot arm based on the point clouds and the state of the hand. Similarly, we also predict the action sequence of the hand using point clouds and the arm state. We use MSE loss to compute the loss between the reconstructed and original action. By further predicting the action sequence respectively, the encoder is able to learn intrinsic features of motion and capture collaborative dynamics.

Given the aforementioned losses, our overall training objective during the pre-training phase is defined as:

$$\min_{\mathcal{E}} \mathcal{L} = \mathcal{L}_{contact} + \lambda \mathcal{L}_{coordination}, \quad (5)$$

where \mathcal{E} represents the encoder of the observation, and λ is a hyperparameter that controls the relative strengths of the losses.

D. Correspondence-based Diffusion Policy

After obtaining the pre-trained encoder, we utilize an imitation learning framework to learn visuomotor policy for dexterous manipulation tasks. Specifically, we adopt conditional denoising diffusion model [22, 7, 38] as our backbone, which conditions on 3D visual features $\phi^{H,O}$ and robot states features $\psi^{A,H}$. Beginning with a Gaussian noise A^K , the denoising network ϵ_θ performs k iterations to gradually denoise A^K into the noise-free action A^0 :

$$A^{k-1} = \alpha_k(A^k - \gamma_k \epsilon_\theta(\phi^{H,O}, \psi^{A,H}, A^k, k)) + \sigma_k \mathcal{N}(0, I), \quad (6)$$

where $\mathcal{N}(0, I)$ is Gaussian noise, α_k , γ_k and σ_k are functions of k , determined by the noise scheduler. This formulation allows the model to capture the distribution of action without the cost of inferring future states.

We use the DDIM scheduler [50] to accelerate the inference speed in real-world experiments. The training objective is to predict the noise added to the original data:

$$\mathcal{L} = \text{MSE}(\epsilon^k, \epsilon_\theta(\bar{\alpha}_k A^0 + \bar{\beta}_k \epsilon^k, \phi^{H,O}, \psi^{A,H}, k)) \quad (7)$$

Unlike the original diffusion-based policy, we incorporate consistency-related features as a condition for policy learning using the pre-trained encoder and fine-tuning the encoder during downstream tasks training. This encoder implicitly extracts contact and coordination information, thereby enhancing the policy's understanding of spatial relationships.

IV. EXPERIMENTS

We conduct comprehensive real-world experiments to answer the following questions:

- To what extent can our framework promote the learning of the visuomotor policy for dexterous manipulation across diverse real-world scenarios (Section IV-B)?
- How promising is CordViP in terms of sample efficiency and generalization capability (Section IV-C, IV-D)?
- What role does each of the system components play in enhancing its overall performance (Section IV-E, IV-F)?




Fig. 4: **Real robot system.** Our system consists of a Leap Hand and a UR5 Arm, with a fixed Realsense L515 camera employed to capture visual observation. The Realsense D435 camera is only used for data collection during teleoperation and is not involved in the policy learning.

A. Experiment Setup

Robot System Setup. As shown in Figure 4, our system consists of a 6-Dof UR5 robot arm and a 16-Dof Leap Hand [48] with four fingers. A single Intel Realsense L515 RGBD camera is mounted on the side of the robot to capture visual observation.

Tasks. We evaluate our approach on four base dexterous manipulation tasks, along with two advanced contact-rich and fine-grained tasks, as shown in Figure 5. The episode length of each task will be limited to a maximum of 500 steps and each task is evaluated with 20 trials by default. We now briefly describe our tasks:

- 1) **PickPlace.** The Leaphand picks up a toy chicken and places it into a blue bowl.
- 2) **FlipCup.** The Leaphand reaches a cup lying on the table, lifts it up, and then rotates the wrist to position the cup upright on the table.
- 3) **Assembly.** The Leaphand reaches and grasps a cylindrical cup, and insert it into a kettle.
- 4) **ArtiManip.** The Leaphand lifts the lid of a box using its thumb and gently opens it, which involves the manipulation of the articulated objects.
- 5) **FlipCap.** This task requires four-finger coordination: the thumb and middle finger lift and rotate the cap slightly, while the index finger pushes it from the opposite side to complete the flip.
- 6) **LongHoriManip.** This task involves four sequential steps: pull, pick, place, and push, requiring precise and continuous control across a long horizon.

For each task, we incorporate a certain level of randomization to ensure that the policy learns the task-specific features rather than fitting the trajectory. As shown in Figure 6, the objects

are randomly placed within the red rectangular region. More details are provided in Appendix A.

Expert demonstrations. Our training demonstrations are collected by human teleoperation. Since our tasks place significant emphasis on coordination and coherence between the hand and the arm, we employ a vision-based approach to teleoperate the robot. Specifically, we use HaMeR [37] to track human hand pose with a single Realsense D435 camera and use Anyteleop [46] framework to retarget the robot system. The robotic arm is controlled through the 6-Dof end-effector’s pose while the robotic hand is controlled through the retargeted hand joint positions. We collect 50 demonstrations for each task, with all expert demonstrations recorded at a frequency of 5Hz.

Baselines. We first compare our method with six state-of-the-art imitation learning algorithms, *i.e.*, three vision-based policies (BCRNN [32], ACT [70], DP [7]) and three 3D-based policies (BCRNN+3D, ACT+3D, DP3 [69]). BCRNN+3D and ACT+3D are variants of BCRNN and ACT, respectively, in which the image input is substituted with point clouds, encoded using PointNet. We emphasize that the resolution of both the image and depth for all 2D and 3D baseline methods is kept consistent, and the point clouds are synthesized from the RGBD camera.

To further assess the performance under the same privileged setting as CordViP, we also introduce three additional baselines: (1) State-based MLP, a simple behavior cloning policy that takes the robot’s proprioception and the object pose as input; (2) State-based Diffusion Policy, which conditions a diffusion model on the concatenated proprioceptive and object pose features; and (3) G3Flow [4], a recent method that dynamically computes semantic flow based on pose information and integrates it with raw 3D point clouds in the Diffusion Policy framework.

B. Effectiveness

The results of the effectiveness experiments are given in Table I and Table II. Our proposed CordViP maintains a completion rate of over 85% across four base tasks, significantly outperforming the other baselines, and also demonstrates superior performance on the two advanced tasks. BCRNN and its 3D variant perform poorly in all base tasks. While the dexterous hand can locate and reach the object, a substantial finger jitter is observed during the grasping and contact-rich manipulation phases. The image-based policies, ACT and DP, achieve good performance in the flipcup and pickplace tasks. However, they struggle with the assembly task. This is mainly due to the significant occlusions during manipulation, where these image-based policies fail to effectively leverage spatial information. state-based DP matches and even outperforms vision-based baselines on simpler tasks. However, due to the lack of geometric information, it struggles with more dexterous tasks. Compared to ACT, ACT+3D shows superior results in all tasks, highlighting the crucial importance of geometric structure and spatial information in policy learning. G3Flow, which benefits from both 3D visual and state information, is the strongest baseline overall. However, its performance on




Fig. 5: Visualization of six dexterous manipulation tasks, with the right side showing the end state.




Fig. 6: **Randomization of Object Positions.** The red rectangles mark the range of positions of manipulated objects. For PickPlace and FlipCup, both the toy chicken and the cup are randomly rotated within a certain range.

tasks such as FlipCap is limited due to the impact of 3D noise and the lack of correspondence modeling.

Interestingly, we discover that the recently proposed DP3 underperforms relative to DP in most of our tasks, which aligns with the findings of Wang et al. [56]. Although several studies have demonstrated the effectiveness of DP3 in simulation and relatively ideal real-world scenarios, such as those with multiple

TABLE I: **Main results of four base real-world tasks.** Each experiment is evaluated with 20 trials.

Method	PickPlace	FlipCup	Assembly	ArtiManip	Avg
State-based MLP	0%	0%	0%	0%	0%
BCRNN	0%	0%	0%	5%	1%
BCRNN+3D	0%	5%	0%	10%	4%
ACT	45%	70%	25%	65%	51%
ACT+3D	70%	80%	35%	70%	64%
DP	55%	65%	20%	35%	44%
DP3	30%	20%	0%	40%	23%
State-based DP	40%	35%	40%	40%	39%
G3Flow	65%	65%	80%	85%	74%
CordViP(Ours)	85%	90%	90%	95%	90%

viewpoints, we observe that the quality of the point clouds has a significant impact on DP3’s performance. Further details will be discussed in Section IV-E. In contrast, our method exhibits remarkable robustness to the quality of point clouds. By establishing correspondences between the dexterous hand and objects, CordViP facilitates more effective perception of spatial and interaction information.

Inference Efficiency. CordViP also attains an efficient inference speed, as it eliminates the need for the farthest point sampling of point clouds while utilizing compact 3D representations. We evaluate the inference speed of CordViP compared with DP3 on an Intel i7-14700KF CPU and RTX 4090D GPU, CordViP reaches a maximum of 12.84 FPS, surpassing DP3’s 11.79 FPS. This highlights that our approach not only achieves enhanced performance but also maintains low computational overhead during inference.

TABLE II: **Main results of advanced tasks.** Each experiment is evaluated with 10 trials.

Method	FlipCap	LongHoriManip			
		Pull	Pick	Place	Push
ACT	30%	100%	40%	40%	30%
ACT+3D	40%	100%	40%	20%	20%
DP	30%	70%	30%	10%	10%
State-based DP	20%	80%	30%	10%	0%
G3Flow	60%	100%	80%	50%	30%
CordViP(Ours)	80%	100%	100%	70%	60%




Fig. 7: **Experimental results of efficiency.** We train ACT, DP, DP3, and CordViP on the PickPlace and FlipCup tasks with an increasing number of demonstrations.

C. Efficiency

The number of expert demonstrations plays a crucial role in the performance of imitation learning. To access the learning efficiency, we train ACT, DP, DP3 and our proposed CordViP on dexterous tasks using varying quantities of expert demonstrations. As illustrated in Figure 7, CordViP exhibits superior performance, achieving higher accuracy with fewer demonstrations. Remarkably, even with just 10 demonstrations, CordViP effectively establishes correspondences, extracts spatial and geometric features, and maintains a high success rate. Furthermore, when provided with a sufficient number of demonstrations, CordViP also demonstrates excellent manipulation capabilities.

D. Generalization

Besides the remarkable effectiveness and efficiency, CordViP also showcases excellent generalization capabilities in real-world dexterous manipulation tasks. In this section, we comprehensively investigate its generalization abilities across four aspects, as detailed below.

Generalize to different lighting conditions.

We established three distinct lighting conditions. As shown in Table III, the 3D representation of our proposed CordViP is dynamically maintained through FoundationPose tracking. FoundationPose, which integrates multimodal information (RGB images, depth data, and prior 3D model knowledge of the target object), demonstrated significant robustness to lighting variations in our empirical tests. In contrast, diffusion policies failed to complete tasks under most conditions. These image-based policies typically rely on data augmentation to improve generalization, which may introduce instability during

TABLE III: **Generalization results on different lighting conditions.** We evaluate the policy under three lighting scenarios: dim light(dim), white light(white) and colored lighting(colored).

Method	Dim	PickPlace		Dim	FlipCup	
		White	Colored		White	Colored
DP	0%	25%	0%	0%	50%	0%
ACT+3D	65%	55%	70%	75%	80%	75%
CordViP(Ours)	80%	85%	80%	80%	85%	90%

TABLE IV: **Generalization results in diverse scenarios,** including varying visual appearances and challenging cluttered environments.

Method	Scene-1	PickPlace		Scene-1	FlipCup	
		Scene-2	Cluttered		Scene-2	Cluttered
DP	0%	0%	0%	0%	0%	5%
ACT+3D	50%	70%	5%	80%	70%	20%
CordViP(Ours)	70%	80%	75%	75%	80%	85%

training.

Generalize to different scenarios. CordViP relies on a robust 6D pose estimator to track objects and build interaction-aware 3D point clouds, which enables our method to generalize to different scenarios. The results presented in Table IV show that DP is highly sensitive to visual variations, while ACT+3D performs well in scenarios with different visual appearances due to its reliance on 3D geometric inputs. However, we observe 3D-based policy struggles in more challenging, cluttered scenes, as randomly placed objects can lead to noisy or incomplete point cloud representations. In contrast, our method effectively focuses on the manipulation of subjects and objects, establishing spatial correspondences, and has demonstrated strong generalization to different scenarios.

Generalize to unseen objects. To validate policies' generalization to different objects, we test them on the PickPlace and FlipCup tasks using three unseen objects, each varying in color, shape, and dynamics. As shown in Table V, DP shows poor generalization to unseen objects, while ACT+3D shows a certain level of generalization ability. However, ACT+3D struggles to adjust effectively and promptly when confronted with objects that exhibit significant differences. CordViP demonstrates strong generalization ability when handling unseen objects, making fine-grained adjustments by establishing spatial and temporal correspondences.

Generalize to different viewpoints. Achieving generalization to different viewpoints presents a more significant challenge, as the expert demonstrations are all captured using a

TABLE V: **Generalization results on unseen objects.** For PickPlace and FlipCup tasks, we chose three previously unseen objects, varying in color, shape, and dynamics.





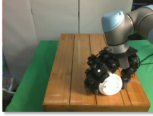

Pick and Place				Flip Cup			
							
origin	object1	object2	object3	origin	object1	object2	object3
Method	PickPlace			FlipCup			
	Object-1	Object-2	Object-3	Object-1	Object-2	Object-3	
DP	0%	5%	0%	15%	35%	0%	
ACT+3D	55%	85%	75%	55%	40%	75%	
CordViP(Ours)	75%	85%	90%	60%	45%	75%	

TABLE VI: **Generalization results to different viewpoints.**

															
Original				View-1				View-2				View-3			
Method	PickPlace			FlipCup				PickPlace				FlipCup			
	View-1	View-2	View-3	View-1	View-2	View-3		View-1	View-2	View-3		View-1	View-2	View-3	
DP	0%	0%	0%	0%	0%	0%		0%	0%	0%		0%	0%	0%	
ACT+3D	20%	45%	60%	0%	40%	30%		20%	45%	60%		0%	40%	30%	
CordViP(Ours)	80%	75%	65%	60%	70%	80%		80%	75%	65%		60%	70%	80%	

fixed camera. We evaluate DP, ACT+3D, and our method with three different camera viewpoints. For 3D-based methods, since the point clouds are represented in the world coordinate system, the camera is recalibrated for each viewpoint. As shown in Table VI, the image-based diffusion policy is highly sensitive to camera viewpoints and completely fails across all three camera views. ACT+3D, which leverages 3D information, demonstrates a certain level of generalization to minor viewpoint changes. However, due to the limitations of the camera viewpoint, the synthesized point clouds can only capture partial spatial information, making it hard to handle significant changes in camera views. By leveraging comprehensive 3D priors of manipulated objects, our method achieves view-agnostic generation of interaction-aware point clouds with full spatial coverage, demonstrating robust performance under significant variations in camera perspectives.

E. Ablations

Effectiveness of components. We conduct a series of ablation experiments on four base real-world tasks to evaluate the effectiveness of different components in our method. As shown in Table VII, our results reveal the critical importance of contact and coordination information for learning dexterous manipulation policy. Specifically, for contact-rich tasks such as Flip-Cup, the absence of pre-training on contact information significantly reduces success rates. This highlights the essential role of contact in understanding physical interactions between the robotic hand and objects. On the other hand, for tasks that require continuous motion of both the arm and hand, such as Assembly and ArtiManip, incorporating coordination information notably enhances the policy’s ability to execute

complex movement patterns. These findings emphasize the necessity of both contact and coordination information to achieve robust manipulation.

We further observe that merging the Interaction-aware point clouds into a unified representation and providing it as the 3D input to DP3 leads to a noticeable improvement in the success rates across all tasks. This also suggests that, compared to ACT+3D, DP3 places higher demands on the quality of the point clouds, and is less robust to noise and occlusion in real-world scenarios. In our settings, we do not carefully position the camera to minimize occlusion of the dexterous hand, which demonstrates the robustness of our method to different viewpoints.

Choice of point cloud encoders. We use PointNet [41], PointNet++ [42], and PointNeXt [44] as point cloud encoders. As shown in Table VIII, PointNet outperforms the other encoders in success rates. We also freeze the point cloud encoder during the training phase of the correspondence-based diffusion policy. The results indicate that fine-tuning further enhances the learning of downstream tasks.

TABLE VIII: **Ablation experiments on point cloud encoders.**

* indicates that the encoder is frozen during the training phase of the correspondence-based diffusion policy.

Encoders	PickPlace	FlipCup	Assembly	ArtiManip	Avg
PointNet	85%	90%	90%	95%	90%
PointNet++	5%	30%	30%	40%	26%
PointNeXt	0%	0%	0%	0%	0%
PointNet*	70%	90%	75%	80%	79%
PointNet++*	5%	15%	5%	30%	14%

F. Transferability

Our proposed CordViP can be easily transferred to different backbones. Following the design of Zhao et al. [70], we choose a Transformer-based CVAE as one of the decision-making backbones for the policy, which outputs a sequence of actions. We replace the 2D image input with our Interaction-aware point clouds and use the pre-trained encoder to establish correspondences between various components. Table IX illustrates that CordViP achieves excellent performance on both diffusion-based and transformer-based backbones, showcasing strong transferability.

TABLE IX: **Transferability to different backbone.**

Method	PickPlace	FlipCup	Assembly	ArtiManip
Diffusion-based	85%	90%	90%	95%
Transformer-based	90%	75%	90%	80%

G. Failure Case Analysis

In this section, we analyze some failure cases of CordViP. We take the flipcup task as an example and visualize the results in Figure 8a. We observe that while CordViP generally achieves good localization and grasping, the hand occasionally hovers in place after wrist rotation. One possible reason is that,

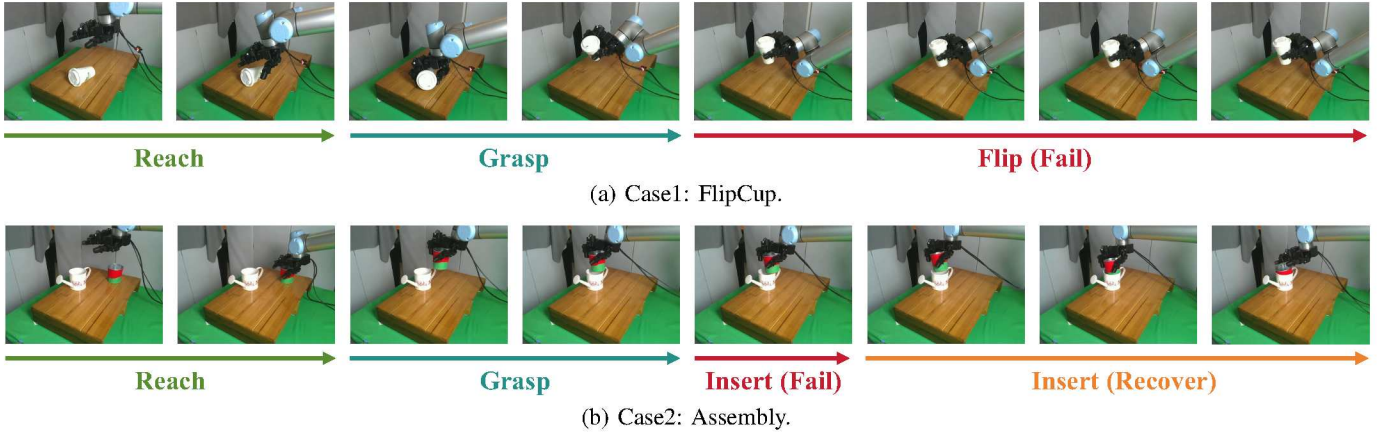


Fig. 8: **Failure case.** (a) Case 1 is a failure case from the Flip Cup task; (b) Case2 is an example from the Assembly task, where after a failure to insert the cup, the policy made adjustments, corrected the orientation of the cup, and successfully completed the task.

TABLE VII: **Ablation experiments on the Effectiveness of different components.** DP3 + Interaction-aware PC refers to using the Interaction-aware point clouds as the visual input for DP3. W/o. contact and coordination pretrain means that the encoder was not pre-trained with contact and coordination data.

Ablation	PickPlace	FlipCup	Assembly	ArtiManip	Avg
DP3 + Interaction-aware PC	65%	75%	60%	85%	71%
CordViP w/o. contact and coordination pretrain	75%	80%	75%	85%	79%
CordViP w/o. contact pretrain	85%	75%	85%	90%	84%
CordViP w/o. coordination pretrain	80%	85%	80%	85%	83%
CordViP(Ours)	85%	90%	90%	95%	90%

during the collection of expert demonstrations, we deliberately reduced wrist speed, resulting in slight jitters that led the model to learn suboptimal solutions. A potential solution could be expanding the policy’s planning horizon by increasing the values of *horizon* and *n_action_steps*.

Nevertheless, we also observe that our model possesses the ability to automatically correct some failure scenarios and recover from them. In the assembly task, inserting the cylindrical cup into the kettle is a delicate operation that relies on strong 3D spatial reasoning capabilities. As shown in Figure 8b, CordViP is able to make timely adjustments, such as when the cup is stuck at the opening of the kettle. It can quickly sense the state of the object and adjust the joint pose to successfully complete the assembly task.

V. CONCLUSIONS AND LIMITATIONS

In this paper, we present CordViP, a novel framework that learns correspondence-based visuomotor policy for dexterous manipulation in the real world. First, we utilize powerful 3D generation methods to obtain 3D models of objects, then use robust 6D object pose estimation based on this modeling information and robot proprioception to obtain interaction-aware point clouds, enhancing the quality of the point clouds and addressing the issues caused by occlusions from the dexterous hand. Second, we use object-centric contact map and coordination information to design a pre-training task that

effectively establishes spatial and temporal correspondences. Finally, the features obtained from the pre-trained encoder are used as conditions to train a visuomotor policy. CordViP significantly outperforms state-of-the-art 2D and 3D baselines on six real-world dexterous manipulation tasks, demonstrating highly competitive performance in both effectiveness and efficiency.

Limitations. Despite the exceptional performance demonstrated by CordViP, there are still certain limitations that could be explored in future work. First, our method struggles to accurately estimate the 6D pose of deformable objects due to the limited expressive capacity of FoundationPose when handling non-rigid geometry. Second, the accuracy of digital twin modeling can significantly impact pose tracking and the quality of initial object point clouds, which are critical for fine-grained dexterous manipulation. Additionally, if the dexterous hand completely occludes the object, FoundationPose may fail to track it. One promising research direction could be the incorporation of visuotactile 6D pose tracking. We leave further exploration of these possibilities for future work.

ACKNOWLEDGEMENT

This work was supported by the National Natural Science Foundation of China (62476011) and the Joint R&D Fund of Beijing Smart-chip Microelectronics Technology Co., Ltd (SGSC0000RGJS2401830, SGSC0000RGJS2401829).

REFERENCES

- [1] Yunfei Bai and C Karen Liu. Dexterous manipulation using both palm and fingers. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1560–1565. IEEE, 2014.
- [2] Raunaq Bhirangi, Abigail DeFranco, Jacob Adkins, Carmel Majidi, Abhinav Gupta, Tess Hellebrekers, and Vikash Kumar. All the feels: A dexterous hand with large-area tactile sensing. *IEEE Robotics and Automation Letters*, 2023.
- [3] Tao Chen, Jie Xu, and Pulkit Agrawal. A system for general in-hand object re-orientation. In *Conference on Robot Learning*, pages 297–307. PMLR, 2022.
- [4] Tianxing Chen, Yao Mu, Zhixuan Liang, Zhanxin Chen, Shijia Peng, Qiangyu Chen, Mingkun Xu, Ruizhen Hu, Hongyuan Zhang, Xuelong Li, et al. G3flow: Generative 3d semantic flow for pose-aware and generalizable object manipulation. *arXiv preprint arXiv:2411.18369*, 2024.
- [5] Yuanpei Chen, Tianhao Wu, Shengjie Wang, Xidong Feng, Jiechuan Jiang, Zongqing Lu, Stephen McAleer, Hao Dong, Song-Chun Zhu, and Yaodong Yang. Towards human-level bimanual dexterous manipulation with reinforcement learning. *Advances in Neural Information Processing Systems*, 35:5150–5163, 2022.
- [6] Yuanpei Chen, Chen Wang, Li Fei-Fei, and C Karen Liu. Sequential dexterity: Chaining dexterous policies for long-horizon manipulation. *arXiv preprint arXiv:2309.00987*, 2023.
- [7] Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research*, page 02783649241273668, 2023.
- [8] T. Cover and P. Hart. Nearest neighbor pattern classification. *IEEE Transactions on Information Theory*, 13(1): 21–27, 1967. doi: 10.1109/TIT.1967.1053964.
- [9] Kairui Ding, Boyuan Chen, Ruihai Wu, Yuyang Li, Zongzheng Zhang, Huan-ang Gao, Siqi Li, Guyue Zhou, Yixin Zhu, Hao Dong, et al. Preafford: Universal affordance-based pre-grasping for diverse objects and environments. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7278–7285. IEEE, 2024.
- [10] Ria Doshi, Homer Walke, Oier Mees, Sudeep Dasari, and Sergey Levine. Scaling cross-embodied learning: One policy for manipulation, navigation, locomotion and aviation. *arXiv preprint arXiv:2408.11812*, 2024.
- [11] Ben Eisner, Harry Zhang, and David Held. Flowbot3d: Learning 3d articulation flow to manipulate articulated objects. *arXiv preprint arXiv:2205.04382*, 2022.
- [12] Ben Eisner, Yi Yang, Todor Davchev, Mel Vecerik, Jonathan Scholz, and David Held. Deep se (3)-equivariant geometric reasoning for precise placement tasks. *arXiv preprint arXiv:2404.13478*, 2024.
- [13] Theophile Gervet, Zhou Xian, Nikolaos Gkanatsios, and Katerina Fragkiadaki. Act3d: Infinite resolution action detection transformer for robotic manipulation. *arXiv preprint arXiv:2306.17817*, 2023.
- [14] Ankit Goyal, Jie Xu, Yijie Guo, Valts Blukis, Yu-Wei Chao, and Dieter Fox. Rvt: Robotic view transformer for 3d object manipulation. In *Conference on Robot Learning*, pages 694–710. PMLR, 2023.
- [15] Abhishek Gupta, Justin Yu, Tony Z Zhao, Vikash Kumar, Aaron Rovinsky, Kelvin Xu, Thomas Devlin, and Sergey Levine. Reset-free reinforcement learning via multi-task learning: Learning dexterous manipulation behaviors without human intervention. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6664–6671. IEEE, 2021.
- [16] William Hebgen Guss, Stephanie Milani, Nicholay Topin, Brandon Houghton, Sharada Mohanty, Andrew Melnik, Augustin Harter, Benoit Buschmaas, Bjarne Jaster, Christoph Berganski, et al. Towards robust and domain agnostic reinforcement learning competitions: Minerl 2020. In *NeurIPS 2020 Competition and Demonstration Track*, pages 233–252. PMLR, 2021.
- [17] Irmak Guzey, Ben Evans, Soumith Chintala, and Lerrel Pinto. Dexterity from touch: Self-supervised pre-training of tactile representations with robotic play. *arXiv preprint arXiv:2303.12076*, 2023.
- [18] Irmak Guzey, Yinlong Dai, Ben Evans, Soumith Chintala, and Lerrel Pinto. See to touch: Learning tactile dexterity through visual incentives. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 13825–13832. IEEE, 2024.
- [19] Siddhant Haldar, Jyothish Pari, Anant Rai, and Lerrel Pinto. Teach a robot to fish: Versatile imitation from one minute of demonstrations. *arXiv preprint arXiv:2303.01497*, 2023.
- [20] Ankur Handa, Arthur Allshire, Viktor Makoviychuk, Aleksei Petrenko, Ritvik Singh, Jingzhou Liu, Denys Makoviichuk, Karl Van Wyk, Alexander Zhurkevich, Balakumar Sundaralingam, et al. Dextreme: Transfer of agile in-hand manipulation from simulation to reality. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5977–5984. IEEE, 2023.
- [21] Rasmus Laurvig Haugaard and Anders Glent Buch. Surfemb: Dense and continuous correspondence distributions for object pose estimation with learnt surface embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6749–6758, 2022.
- [22] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [23] Yueru Jia, Jiaming Liu, Sixiang Chen, Chenyang Gu, Zhilue Wang, Longzan Luo, Lily Lee, Pengwei Wang, Zhongyuan Wang, Renrui Zhang, and Shanghang Zhang. Lift3d foundation policy: Lifting 2d large-scale pretrained models for robust 3d robotic manipulation, 2024. URL <https://arxiv.org/abs/2411.18623>.
- [24] Hanwen Jiang, Shaowei Liu, Jiashun Wang, and Xiaolong

- Wang. Hand-object contact consistency reasoning for human grasps generation. In *Proceedings of the International Conference on Computer Vision*, 2021.
- [25] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. *arXiv:2304.02643*, 2023.
- [26] Vikash Kumar, Yuval Tassa, Tom Erez, and Emanuel Todorov. Real-time behaviour synthesis for dynamic hand-manipulation. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6808–6815. IEEE, 2014.
- [27] Zihang Lai, Senthil Purushwalkam, and Abhinav Gupta. The functional correspondence problem. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15772–15781, 2021.
- [28] Zhixuan Liang, Yao Mu, Yixiao Wang, Tianxing Chen, Wenqi Shao, Wei Zhan, Masayoshi Tomizuka, Ping Luo, and Mingyu Ding. Dexhanddiff: Interaction-aware diffusion planning for adaptive dexterous manipulation, 2024. URL <https://arxiv.org/abs/2411.18562>.
- [29] Chu-Cheng Lin, Aaron Jaech, Xin Li, Matthew R Gormley, and Jason Eisner. Limitations of autoregressive models and their alternatives. *arXiv preprint arXiv:2010.11939*, 2020.
- [30] Yushan Liu, Shilong Mu, Xintao Chao, Zizhen Li, Yao Mu, Tianxing Chen, Shoujie Li, Chuqiao Lyu, Xiao-ping Zhang, and Wenbo Ding. Avr: Active vision-driven robotic precision manipulation with viewpoint and focal length optimization. *arXiv preprint arXiv:2503.01439*, 2025.
- [31] Guanxing Lu, Zifeng Gao, Tianxing Chen, Wenxun Dai, Ziwei Wang, and Yansong Tang. Manicm: Real-time 3d diffusion policy via consistency model for robotic manipulation. *arXiv preprint arXiv:2406.01586*, 2024.
- [32] Ajay Mandlekar, Danfei Xu, Josiah Wong, Soroush Nasiriany, Chen Wang, Rohun Kulkarni, Li Fei-Fei, Silvio Savarese, Yuke Zhu, and Roberto Martín-Martín. What matters in learning from offline human demonstrations for robot manipulation. *arXiv preprint arXiv:2108.03298*, 2021.
- [33] Igor Mordatch, Zoran Popović, and Emanuel Todorov. Contact-invariant optimization for hand manipulation. In *Proceedings of the ACM SIGGRAPH/Eurographics symposium on computer animation*, pages 137–144, 2012.
- [34] Yao Mu, Tianxing Chen, Shijia Peng, Zanxin Chen, Zeyu Gao, Yude Zou, Lunkai Lin, Zhiqiang Xie, and Ping Luo. Robotwin: Dual-arm robot benchmark with generative digital twins (early version). *arXiv preprint arXiv:2409.02920*, 2024.
- [35] Yao Mu, Tianxing Chen, Zanxin Chen, Shijia Peng, Zhiqian Lan, Zeyu Gao, Zhixuan Liang, Qiaojun Yu, Yude Zou, Mingkun Xu, Lunkai Lin, Zhiqiang Xie, Mingyu Ding, and Ping Luo. Robotwin: Dual-arm robot benchmark with generative digital twins, 2025. URL <https://arxiv.org/abs/2504.13059>.
- [36] Timothy Patten, Kiru Park, and Markus Vincze. Dgcm-net: dense geometrical correspondence matching network for incremental experience-based robotic grasping. *Frontiers in Robotics and AI*, 7:120, 2020.
- [37] Georgios Pavlakos, Dandan Shan, Ilija Radosavovic, Angjoo Kanazawa, David Fouhey, and Jitendra Malik. Reconstructing hands in 3D with transformers. In *CVPR*, 2024.
- [38] Tim Pearce, Tabish Rashid, Anssi Kanervisto, Dave Bignell, Mingfei Sun, Raluca Georgescu, Sergio Valcarcel Macua, Shan Zheng Tan, Ida Momennejad, Katja Hofmann, et al. Imitating human behaviour with diffusion models. *arXiv preprint arXiv:2301.10677*, 2023.
- [39] Johannes Pitz, Lennart Röstel, Leon Sievers, and Berthold Bäuml. Dextrous tactile in-hand manipulation using a modular reinforcement learning architecture. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1852–1858. IEEE, 2023.
- [40] Dean A Pomerleau. Alvin: An autonomous land vehicle in a neural network. *Advances in neural information processing systems*, 1, 1988.
- [41] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *arXiv preprint arXiv:1612.00593*, 2016.
- [42] Charles R Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *arXiv preprint arXiv:1706.02413*, 2017.
- [43] Haozhi Qi, Ashish Kumar, Roberto Calandra, Yi Ma, and Jitendra Malik. In-hand object rotation via rapid motor adaptation. In *Conference on Robot Learning*, pages 1722–1732. PMLR, 2023.
- [44] Guocheng Qian, Yuchen Li, Houwen Peng, Jinjie Mai, Hasan Hammoud, Mohamed Elhoseiny, and Bernard Ghanem. Pointnext: Revisiting pointnet++ with improved training and scaling strategies. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [45] Yuzhe Qin, Yueh-Hua Wu, Shaowei Liu, Hanwen Jiang, Ruihan Yang, Yang Fu, and Xiaolong Wang. Dexmv: Imitation learning for dexterous manipulation from human videos, 2021.
- [46] Yuzhe Qin, Wei Yang, Binghao Huang, Karl Van Wyk, Hao Su, Xiaolong Wang, Yu-Wei Chao, and Dieter Fox. Anyteleop: A general vision-based dexterous robot arm-hand teleoperation system. *arXiv preprint arXiv:2307.04577*, 2023.
- [47] Moritz Reuss, Maximilian Li, Xiaogang Jia, and Rudolf Lioutikov. Goal-conditioned imitation learning using score-based diffusion policies. *arXiv preprint arXiv:2304.02532*, 2023.
- [48] Kenneth Shaw, Ananye Agarwal, and Deepak Pathak. Leap hand: Low-cost, efficient, and anthropomorphic hand for robot learning. *arXiv preprint arXiv:2309.06440*, 2023.
- [49] Mohit Shridhar, Lucas Manuelli, and Dieter Fox.

- Perceiver-actor: A multi-task transformer for robotic manipulation. In *Conference on Robot Learning*, pages 785–799. PMLR, 2023.
- [50] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.
- [51] Yup Su, Xinyu Zhan, Hongjie Fang, Yong-Lu Li, Cewu Lu, and Lixin Yang. Motion before action: Diffusing object motion as manipulation condition. *arXiv preprint arXiv:2411.09658*, 2024.
- [52] Octo Model Team, Dibya Ghosh, Homer Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Tobias Kreiman, Charles Xu, et al. Octo: An open-source generalist robot policy. *arXiv preprint arXiv:2405.12213*, 2024.
- [53] Dmitry Tochilkin, David Pankratz, Zexiang Liu, Zixuan Huang, , Adam Letts, Yangguang Li, Ding Liang, Christian Laforte, Varun Jampani, and Yan-Pei Cao. Tripso: Fast 3d object reconstruction from a single image. *arXiv preprint arXiv:2403.02151*, 2024.
- [54] Chen Wang, Linxi Fan, Jiankai Sun, Ruohan Zhang, Li Fei-Fei, Danfei Xu, Yuke Zhu, and Anima Anandkumar. Mimicplay: Long-horizon imitation learning by watching human play. *arXiv preprint arXiv:2302.12422*, 2023.
- [55] Chen Wang, Haochen Shi, Weizhuo Wang, Ruohan Zhang, Li Fei-Fei, and C. Karen Liu. Dexcap: Scalable and portable mocap data collection system for dexterous manipulation. *arXiv preprint arXiv:2403.07788*, 2024.
- [56] Chenxi Wang, Hongjie Fang, Hao-Shu Fang, and Cewu Lu. Rise: 3d perception makes real-world robot imitation simple and effective. *arXiv preprint arXiv:2404.12281*, 2024.
- [57] Dian Wang, Stephen Hart, David Surovik, Tarik Kelestemur, Haojie Huang, Haibo Zhao, Mark Yeatman, Jiuguang Wang, Robin Walters, and Robert Platt. Equivariant diffusion policy. In *8th Annual Conference on Robot Learning*, 2024. URL <https://openreview.net/forum?id=wD2kUVLT1g>.
- [58] Ruicheng Wang, Jialiang Zhang, Jiayi Chen, Yinzen Xu, Puhao Li, Tengyu Liu, and He Wang. Dexgraspnet: A large-scale robotic dexterous grasp dataset for general objects based on simulation. *arXiv preprint arXiv:2210.02697*, 2022.
- [59] Zhenyu Wei, Zhixuan Xu, Jingxiang Guo, Yiwen Hou, Chongkai Gao, Zhehao Cai, Jiayu Luo, and Lin Shao. D(r,o) grasp: A unified representation of robot and object interaction for cross-embodiment dexterous grasping. *arXiv preprint arXiv:2410.01702*, 2024.
- [60] Bowen Wen, Wei Yang, Jan Kautz, and Stan Birchfield. FoundationPose: Unified 6d pose estimation and tracking of novel objects. In *CVPR*, 2024.
- [61] Ruihai Wu, Haoran Lu, Yiyan Wang, Yubo Wang, and Hao Dong. Unigarmentmanip: A unified framework for category-level garment manipulation via dense visual correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2024.
- [62] Tianhao Wu, Jinzhou Li, Jiyao Zhang, Mingdong Wu, and Hao Dong. Canonical representation and force-based pretraining of 3d tactile for dexterous visuo-tactile policy learning. *arXiv preprint arXiv:2409.17549*, 2024.
- [63] Shangning Xia, Hongjie Fang, Cewu Lu, and Hao-Shu Fang. Cage: Causal attention enables data-efficient generalizable robotic manipulation. *arXiv preprint arXiv:2410.14974*, 2024.
- [64] Zhou Xian, Nikolaos Gkanatsios, Theophile Gervet, Tsung-Wei Ke, and Katerina Fragkiadaki. Chaineddiffuser: Unifying trajectory diffusion and keypose prediction for robotic manipulation. In *7th Annual Conference on Robot Learning*, 2023.
- [65] Zhao-Heng Yin, Binghao Huang, Yuzhe Qin, Qifeng Chen, and Xiaolong Wang. Rotating without seeing: Towards in-hand dexterity through touch. *Robotics: Science and Systems*, 2023.
- [66] Zhao-Heng Yin, Binghao Huang, Yuzhe Qin, Qifeng Chen, and Xiaolong Wang. Rotating without seeing: Towards in-hand dexterity through touch. *arXiv preprint arXiv:2303.10880*, 2023.
- [67] Maryam Zare, Parham M. Kebria, Abbas Khosravi, and Saeid Nahavandi. A survey of imitation learning: Algorithms, recent developments, and challenges. *IEEE Transactions on Cybernetics*, 54(12):7173–7186, 2024. doi: 10.1109/TCYB.2024.3395626.
- [68] Yanjie Ze, Zixuan Chen, Wenhao Wang, Tianyi Chen, Xialin He, Ying Yuan, Xue Bin Peng, and Jiajun Wu. Generalizable humanoid manipulation with improved 3d diffusion policies. *arXiv preprint arXiv:2410.10803*, 2024.
- [69] Yanjie Ze, Gu Zhang, Kangning Zhang, Chenyuan Hu, Muhan Wang, and Huazhe Xu. 3d diffusion policy: Generalizable visuomotor policy learning via simple 3d representations. In *ICRA 2024 Workshop on 3D Visual Representations for Robot Manipulation*, 2024.
- [70] Tony Z Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual manipulation with low-cost hardware. *arXiv preprint arXiv:2304.13705*, 2023.
- [71] Wenshuai Zhao, Jorge Peña Queralta, and Tomi Westerlund. Sim-to-real transfer in deep reinforcement learning for robotics: a survey. In *2020 IEEE symposium series on computational intelligence (SSCI)*, pages 737–744. IEEE, 2020.

A. Real-World Task Description

PickPlace: This task requires coordinated motion of the robot’s hand and arm, involving all four fingers. The Leaphand first locates the position of the toy chicken based on visual input. The palm then approaches the chicken, with all four fingers gradually wrapping around it. Once the chicken is grasped, the wrist is raised, and the hand is moved towards a blue bowl. Upon reaching a position directly above the bowl, the fingers are released, allowing the chicken to be placed inside. This task presents challenges in accurately locating and grasping the toy chicken, coordinating finger movements for a stable grip, and precisely manipulating the hand to place the object into the bowl while avoiding obstacles and ensuring a gentle release. Success is achieved if the toy chicken is placed into the bowl.

FlipCup: This task requires the robot to flip a cup from a lying position on the table to an upright standing position. The Leaphand approaches the cup and places its hand on the top of the cup, lifts it, and then flips it to an upright position. The hand must apply controlled force to rotate the cup while maintaining stability to prevent it from tipping over. The challenge lies in the fact that the cup may undergo changes in its orientation, requiring the hand to dynamically adjust its position and force to stabilize the cup’s motion and achieve the desired outcome. Success is achieved when the cup stands upright on the table.

Assembly: This task requires the robot to assemble a cylindrical cup onto a kettle. The Leaphand first approaches the cup and grasps it, positioning it accurately to align with the kettle’s opening. Using coordinated finger and wrist movements, the hand carefully attaches the cup to the kettle, ensuring a secure fit. The challenges lie in precise alignment and making fine adjustments based on feedback while handling high occlusion and ambiguity. Success is achieved when the cup is securely assembled onto the kettle.

ArtiManip: This task requires the robot to open a box using its thumb and index fingers. The robot needs to first reach the box, grasp the box’s lid, and then carefully adjust its fingers to open the box without pushing it. The hand must coordinate the motion of the thumb with fine adjustments to apply the right amount of force, ensuring the lid is opened smoothly and safely. This task presents a challenge in handling articulated objects with multiple moving parts while maintaining delicate control over the manipulation process. Success is achieved when the lid is fully open.

FlipCap: This task requires the robot hand to flip over a small cap using coordinated motion of multiple fingers. The robot must first reach the cap and stabilize it with the thumb and middle finger, gently lifting and rotating it to a tilted position. Simultaneously, the index finger is required to push from the opposite side with precise force to complete the flipping motion. This task demands high-level coordination between fingers, particularly in managing the contact points during rotation. Accurate finger placement and dynamic re-adjustments are crucial to avoid slipping or excessive force.

Success is defined as fully flipping the cap without displacing it.

LongHoriManip: This task requires the robot hand to complete a long-horizon manipulation sequence involving four consecutive subtasks: pull, pick, place, and push. The dexterous hand must locate and grasp the drawer handle, pulling it open with sufficient force and precision. Once the drawer is fully extended, the hand proceeds to pick up a small toy chicken from the table using a coordinated grasp. It then carefully places the toy into the open drawer, ensuring stability and proper placement. Finally, the hand pushes the drawer closed, completing the sequence. This task poses a significant challenge due to its extended temporal dependencies and the need for continuous, adaptive control across multiple object interactions. Each subtask is individually evaluated for success to better understand performance at different stages. The overall task is considered successful when the toy is placed inside and the drawer is fully closed.

We list the parameters of expert demonstrations for different tasks in Table X. For all demonstrations of a given task, we maintain a consistent number of steps. “Demo” refers to the number of demonstrations collected for each task, “Episode Length” denotes the duration of each episode in a task, “Teleop. Times” indicates the teleoperation time required to collect a single demonstration, and “Max Steps” represents the maximum execution time for a task during evaluation.

TABLE X: **Parameters of expert demonstrations for real-world tasks.** “Demo” refers to the number of demonstrations, “Episode Length” denotes the duration of each episode in a task, “Teleop. Times” indicates the teleoperation time per demonstration, and “Max Steps” represents the maximum execution time for a task during evaluation.

Task Name	Demos	Episode Length	Teleop. Times(s)	Max steps
<i>PickPlace</i>	50	150	30	400
<i>Flipcup</i>	50	150	30	300
<i>Assembly</i>	50	175	35	500
<i>ArtiManip</i>	50	190	38	600
<i>Flipcap</i>	50	150	30	400
<i>LongHoriManip</i>	50	250	50	800

B. Implementation Details

Network Architecture. For point cloud encoding, we first use PointNet[41] to process point cloud data without RGB information, outputting a set of point feature vectors at the dimension of 1024. The PointNet consists of three fully connected layers, each followed by LayerNorm for normalization and ReLU activation.

For the cross-attention transformer, we adopted the architecture design from Eisner et al. [12], using a multi-head attention block of 4 heads. The state features of the robotic arm and the dexterous hand are each passed through a linear layer, mapped to 16 dimensions. The features are then processed through the same Transformer architecture for cross-attention, enabling feature fusion. The fused features are subsequently combined with the original features using a residual connection.



Fig. 10: 3D digital asset generation from a single view.

Demonstrations Process. We utilize the RealSense L515 camera to capture RGB-D images with a resolution of 480×640 . The depth data are aligned with the RGB data to ensure accurate spatial correspondence. All data collection is managed through ROS and data recording begins once both the camera feed and robot teleoperation inputs are received. For our method, we use only RGB and depth data to track the object’s pose. In contrast, for other baselines, we synthesize the point cloud from RGBD data, and both the pose and the point clouds are transformed into the world coordinate system. We crop point clouds with the range of $x \in [-0.4m, 0.1m]$, $y \in [-0.7, -0.4]$, $z \in [0.1, 0.51]$, which has been verified to be suitable for observation, as shown in Figure 9.




Fig. 9: **Visualization of Point Cloud Processing.** The point cloud is synthesized from RGBD data. The point cloud is then cropped and processed using farthest point sampling (FPS) to generate 1024 points.

We collect both the robot’s state and actions using joint angles in radians, including the 6-DOF joints of the robotic arm and the 16-DOF joints of the Leaphand.

Real-to-Sim for Digital Twin Generation. We use TripoSR [53] to generate digital twins from a single-view image, which enables the creation of high-quality 3D assets. The visual results are shown in the figure 10.

6D Pose Estimation. We utilize FoundationPose [60] to perform robust 6D pose estimation for various objects across tasks. For the PickPlace task, we estimate the 6D pose of both the chicken and the bowl to capture the spatial relationships between the objects. For the FlipCup task, we focus on accurately estimating the 6D pose of the cup. For the Assembly task, we separately estimate the 6D pose of the cylindrical cup

and the kettle, enabling precise interactions with each object during assembly. For the Artimaip task, we first decompose it into two distinct parts: the box body and the lid, and then perform estimation for each part using FoundationPose. For the FlipCap task, we estimate the 6D pose of the cap. For the LongHoriManip task, we separately estimate the 6D poses of the toy chicken and the drawer handle.

Normalizations. The range of training data has a significant impact on the training stability of CordViP. We linearly scale the minimum and maximum values of each action and observation dimension to the range of $[-1, 1]$. This step is necessary for DDIM [50] and DDPM [22], as they clip the predicted results to the range of $[-1, 1]$ for training stability.

Hyperparameters. We list the training hyperparameters used in CordViP in Table XI.

TABLE XI: Training hyperparameters in CordViP.

Hyperparameters	Value
Robot point cloud size	1024*3
Object point cloud size	1024*3
Contact map scaling factor γ	1
Contact map scaling factor θ	10
Contact map size	1024*1
Loss weight λ	1
<i>horizon</i>	12
<i>n_obs_steps</i>	4
<i>n_action_steps</i>	6
Optimizer	AdamW

Baselines settings. We train the Diffusion Policy for 600 epochs with *horizon*=12, *n_obs_steps*=4, and *n_action_steps*=8. The Diffusion Policy baseline utilizes ResNet18 as the visual encoder and employs CNN-based backbones. The 3D Diffusion Policy is trained for 8000 epochs with *horizon*=12, *n_obs_steps*=4, *n_action_steps*=8. It uses DP3 Encoder as the point cloud encoder. For ACT, we train the model for 5000 epochs with *action_chunk*=30. In ACT3D, train for 1600 epochs with an *action_chunk* size of 30. Point clouds are used to replace the original image inputs. The point clouds are encoded using PointNet, and the extracted point cloud features are given learnable positional encodings, similar to other input information such as joint states and latent inputs. For BCRNN, we train the model for 1500 epochs with *horizon*=10, *n_obs_steps*=1, *n_action_steps*=1. The BCRNN3D is trained for 3000 epochs with *horizon*=10, *n_obs_steps*=1, *n_action_steps*=1, where the observations are replaced from images to point clouds. It uses PointNet as the point cloud encoder.

C. Failure Analysis of Baselines

3D Diffusion Policy. DP3 [69] appears to struggle in learning meaningful actions from our demonstration data. We have already analyzed in the paper that potential reasons include factors such as the camera’s viewpoint and the quality of the point cloud. Wang et al. [56] also points out that the type of motion patterns can affect the quality of demonstration learning. Instead of nature actions, axis-wise actions were used

in DP3’s demonstration data. This is because the robotic arm is controlled by the keyboard, which inherently limits the motion representation to axis-wise actions.




Fig. 11: **Comparison of Motion Patterns.** DP3 Uses Axis-Wise Actions.

RISE. RISE [56] is a recently proposed end-to-end baseline for real-world imitation learning, which predicts continuous actions directly from single-view point clouds. It takes voxelized point clouds as input to the policy and assumes that the end-effector pose of the robotic arm is implicitly encoded within the point cloud. However, this approach proves unsuitable for our dexterous hand scenario, where the hand has a high degree of freedom and often experiences occlusions. We evaluated RISE in our settings and observed that the robotic arm exhibited excessively abrupt movements.

D. More Visualization Results

We present additional point cloud visualization results in Figure 12, demonstrating that interaction-aware point clouds can effectively enhance the quality of 3D observations.

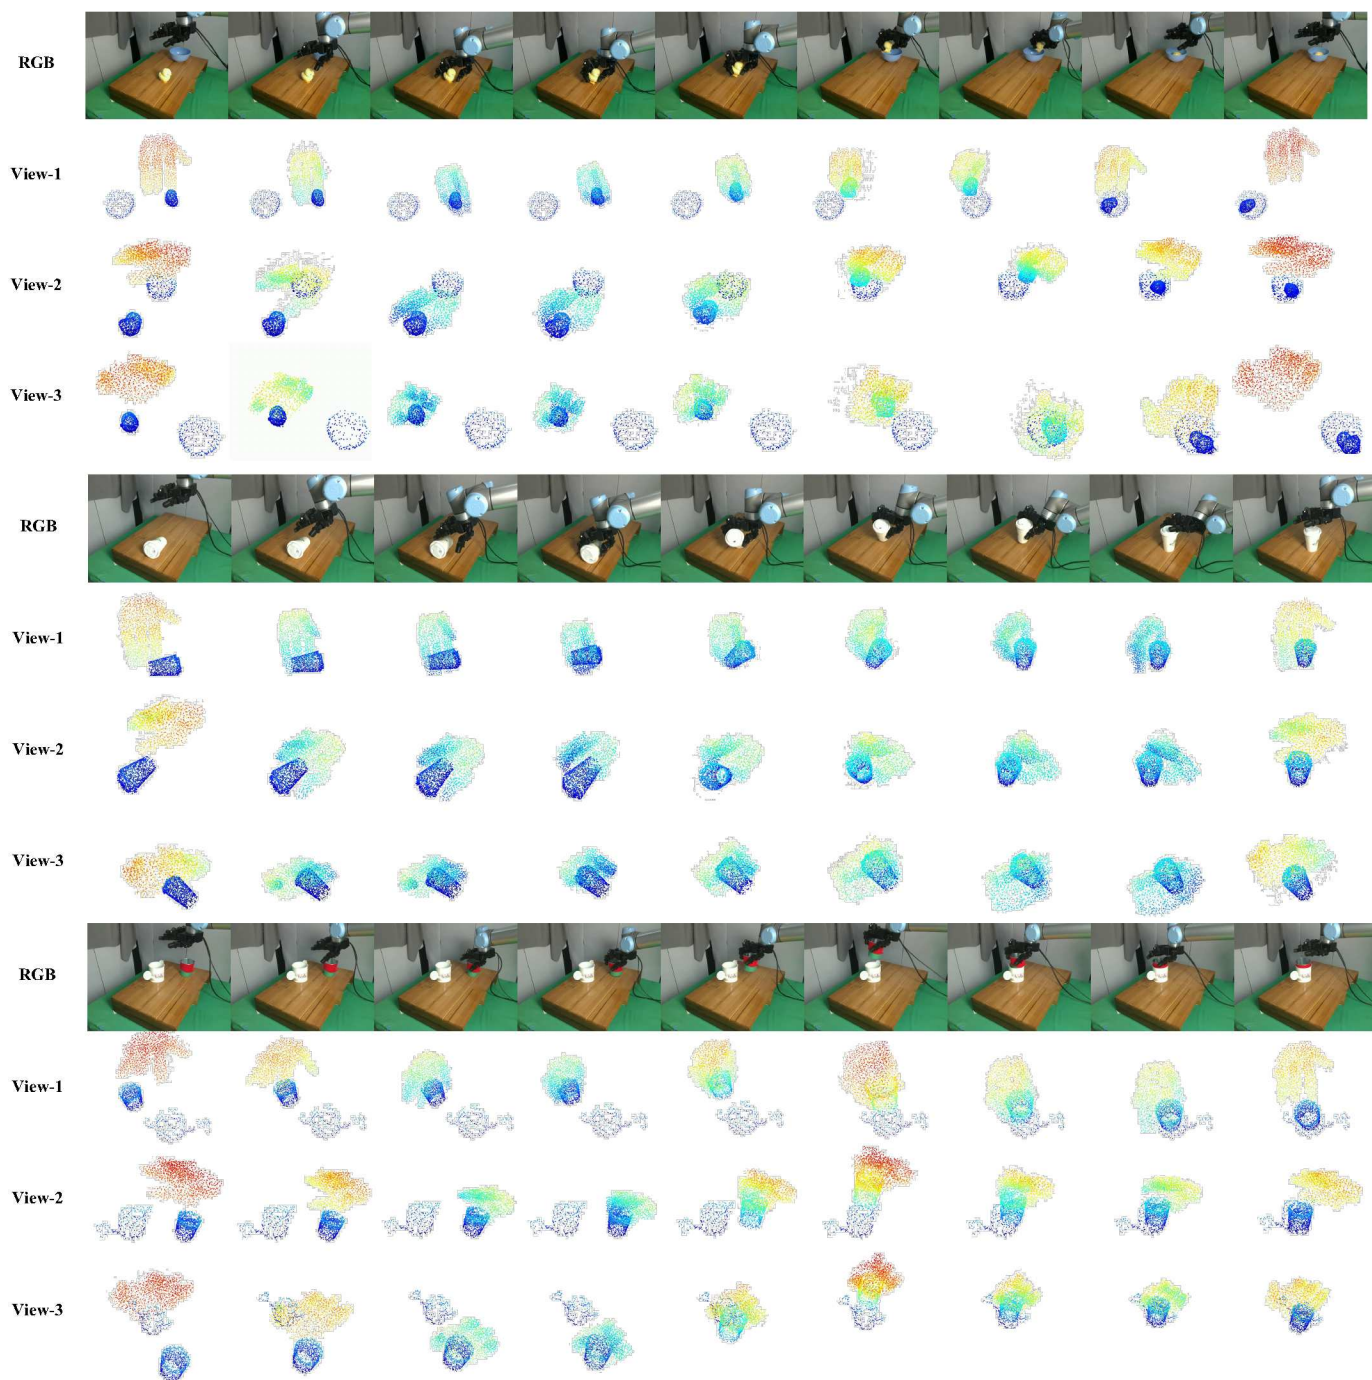


Fig. 12: Visualization of Point Clouds During the Task Process.