

Learning Dexterous Manipulation with Quantized Hand State

Ying Feng^{*,1,3}, Hongjie Fang^{*,1,†}, Yinong He^{*,1,4},
 Jingjing Chen¹, Chenxi Wang², Zihao He¹, Ruonan Liu¹, Cewu Lu^{1,2,3,†}

Abstract—Dexterous robotic hands enable robots to perform complex manipulations that require fine-grained control and adaptability. Achieving such manipulation is challenging because the high degrees of freedom tightly couple hand and arm motions, making learning and control difficult. Successful dexterous manipulation relies not only on precise hand motions, but also on accurate spatial positioning of the arm and coordinated arm-hand dynamics. However, most existing visuomotor policies represent arm and hand actions in a single combined space, which often causes high-dimensional hand actions to dominate the coupled action space and compromise arm control. To address this, we propose *DQ-RISE*, which quantizes hand states to simplify hand motion prediction while preserving essential patterns, and applies a continuous relaxation that allows arm actions to diffuse jointly with these compact hand states. This design enables the policy to learn arm-hand coordination from data while preventing hand actions from overwhelming the action space. Experiments show that *DQ-RISE* achieves more balanced and efficient learning, paving the way toward structured and generalizable dexterous manipulation. Project website: <https://rise-policy.github.io/DQ-RISE/>.

I. INTRODUCTION

Dexterous robotic hands have become a central focus in robotics, allowing robots to interact with the physical world in ways that resemble human manipulation [39, 46, 48]. With multiple independently actuated fingers and rich contact surfaces, they can perform complex manipulations beyond parallel grasping [14, 52], including precise in-hand reorientation [1, 5, 6] and adaptive gripping for objects of diverse shapes and sizes [16, 54, 59]. Beyond these fundamental tasks, dexterous hands can use tools, coordinate bimanual actions, and carry out intricate assembly procedures [8, 28, 61], bridging the gap between rigid robotic actuation and human-like flexibility. Together, these capabilities define the scope and promise of robotic dexterity.

However, these abilities come at a cost: the additional degrees of freedom (DoF) make manipulation highly complex and tightly coupled with arm motion. To illustrate, consider the everyday task of opening a jar by hooking onto its lid and rotating it until it comes off, as depicted in Fig. 1. Successful manipulation requires not only *accurate hand motions* to twist the lid, but also two often overlooked aspects regarding the arm: (1) *precise arm localization* —bringing the end-effector exactly to the lid’s position; and (2) *arm-hand coordination* — when the fingers hook the lid, the arm must simultaneously press downward to prevent lifting the jar. Together, these requirements highlight that beyond

Task. Open the jar by hooking onto its lid and rotating it until it comes off.




Fig. 1: Dexterous Manipulation from the Action Prediction Perspective. Beyond hand motion, successful dexterous manipulation also requires precise arm localization and coordinated arm-hand dynamics. (A) Existing visuomotor policies predict arm and hand actions jointly, causing hand actions to dominate the combined action space and arm localization to suffer. (B) Naively separating arm and hand predictions can lead to incoherent coordination. (C) Our approach quantizes hand states to preserve hand motion while jointly diffusing arm actions, enabling precise arm localization and smooth arm-hand coordination.

finger dexterity, the primary goals of dexterous manipulation include spatial accuracy of the arm and cooperative dynamics between arm and hand.

From the perspective of robot action prediction, visuomotor policies should output both arm and hand actions. Most existing policies treat them together in a single combined action space [17, 22, 33, 49, 51], as shown in Fig. 1A. While convenient, this approach often leads to imbalanced learning: the high-DoF hand actions dominate the combined action space and hinder accurate arm control, as confirmed by our experiments in §IV-B. This observation motivates a clearer functional distinction: *robotic arms primarily handle spatial localization, while dexterous hands are responsible for executing fine-grained actions*. Accordingly, visuomotor policies should *focus on spatial localization for the arm while memorizing action patterns for the dexterous hand*. However, naively disentangling arm and hand action generation, as illustrated in Fig. 1B, can break their coordination and limit overall policy performance.

Building on this insight, we propose *DQ-RISE*, an extension of the base visuomotor policy RISE [53] that introduces structured action prediction for dexterous manipulation. We

*Equal Contribution. †Corresponding Authors.

¹Shanghai Jiao Tong University. ²Noematrix.

³Shanghai Innovation Institute. ⁴Carnegie Mellon University.

quantize dexterous hand states into a compact set of task-relevant patterns, reformulating hand action prediction as a classification problem analogous to gripper open/close control. To maintain arm-hand coordination, we further introduce a continuous relaxation process that enables the policy to diffuse arm actions jointly with quantized hand actions [23], thereby preserving fine-grained dexterity while improving arm localization. In addition, we design a hybrid dexterous teleoperation system for demonstration collection, which facilitates intuitive arm-hand control compared to existing systems. Extensive experiments across diverse dexterous tasks demonstrate that *DQ-RISE* achieves more balanced and efficient learning than other action prediction schemes, highlighting a practical pathway toward scalable and generalizable robotic dexterity.

II. RELATED WORKS

A. Dexterous Manipulation

Dexterous manipulation has recently attracted considerable attention in robotics. Owing to the high DoF of dexterous hands, prior works have primarily relied on reinforcement learning to acquire complex manipulation skills [1, 3, 5, 6, 8, 34, 43, 70]. These approaches typically learn state-based policies in simulation with carefully designed reward functions [38]. To bridge the gap between simulation and the real world, where observations differ and privileged state information is unavailable, researchers have applied techniques like sim-to-real transfer [34, 43] and teacher-student distillation [6, 9] to improve the policy’s adaptability.

Another line of work learns dexterous hand behaviors from human demonstrations [31, 42] or teleoperated demonstrations [17, 22, 33, 49, 51] within the imitation learning framework. However, as discussed in §I, existing imitation policies mostly predict arm and hand actions jointly in a naively-combined high-DOF action space, often resulting in unbalanced learning and suboptimal hand-arm coordination. To address this limitation, we propose disentangling the joint prediction problem into a more principled and tractable formulation by leveraging the distinct properties of arm motions and hand actions.

B. Quantization in Robotics

Quantization in robotics can be broadly categorized into two main directions: observation quantization and action discretization. Common tools include VQ-VAE [40, 45, 66] and VQ-GAN [13], which leverage generative models [20, 27] and discrete codebooks to produce discrete representations of high-dimensional data, enabling more efficient learning.

Observation quantization compresses high-dimensional sensory inputs like images into discrete latent codes. Visuomotor policies and vision-language-action models often leverage this for future observation prediction [4, 7, 29, 32, 63, 68], as predicting latent codes of observations is more tractable and useful than forecasting their raw pixel values [32]. This future prediction serves as an auxiliary task that supports action generation and reasoning [29]. Action quantization discretizes continuous control signals into

representative primitives or codebook entries [30, 55, 60]. This simplification improves learning stability and overall performance. However, most existing approaches focus on action chunks that model future robot motion [69].

In contrast, we propose to quantize the dexterous hand state (single-step action) into compact and meaningful discrete codes. Rather than modeling motion trajectories, we directly model hand states, making the quantization more explainable and intuitive. Empirically, we also found that quantizing action chunks often produces less meaningful codes and worse performance compared to state quantization.

C. Dexterous Teleoperation System

Compared with teleoperation systems for grippers [12, 15, 18, 19, 56, 69], dexterous hand teleoperation is substantially more challenging because operators must control high-DoF configurations, coordinate many coupled joints, and handle complex contact dynamics while also managing arm motions. Existing approaches tackle this in different ways: some introduce hardware that directly maps human hand motion to robotic joints [17, 49, 50, 58, 67], while others employ hand pose estimation [10, 25, 41, 44, 62] or motion-capture gloves [51] to capture human keypoints, and then retarget them to the dexterous hand [36, 64]. For joint hand-arm teleoperation, most systems rely on VR/AR joystick-based interfaces for arm control [10, 25, 33, 35, 41, 51], while a few utilize exoskeletons for full-arm mapping [17, 49, 62].

Recently, several studies have simplified human control by defining discrete gestures during data collection [33, 35]. While this is conceptually similar to our insight of quantizing hand states, the key difference is that their approach relies on *manual quantization at collection time*, whereas we apply *automatic quantization during policy training*. Enforcing discrete gestures during collection can increase the cognitive load on operators, bias their control strategies, and limit the advantages of high-DoF dexterous hands over low-DoF end-effectors, often leading to unnatural demonstrations. To this end, we let operators freely control the hand during data collection and perform quantization afterwards for effective and interpretable policy learning.

III. METHOD

In this section, we first tackle data collection by presenting our VR-glove hybrid dexterous teleoperation system (§III-A). We then quantize hand states into discrete latent codes (§III-B). While a straightforward approach is to treat these discrete states directly as a classification problem for policy learning, we find that jointly diffusing arm and hand actions improves action consistency and overall performance. Hence, we re-index the discrete hand states for continuous relaxation (§III-C) and use the resulting relabeled demonstrations to train the base visuomotor policy, ensuring more accurate and robust hand-arm coordination during manipulation (§III-D). An overview of our policy is illustrated in Fig. 3.

A. Hybrid Dexterous Teleoperation System

As shown in Fig. 2, our platform consists of a Flexiv Rizon 4 robotic arm with a 6-DoF OyMotion ROHand. Two




Fig. 2: Robot Platform and Hybrid Dexterous Teleoperation System. Our platform consists of a Flexiv robotic arm equipped with an ROHand. During teleoperation, the arm is controlled via a VR joystick, where the joystick button can be used to pause arm motion and adjust the joystick pose for more intuitive and convenient operation. For hand control, we use a GForce glove to directly operate the ROHand using joint correspondence.

Intel RealSense D415 cameras provide global observations to mitigate occlusions, while a wrist-mounted Intel RealSense D435 camera is used for calibration only.

We design a VR-glove hybrid system for teleoperation with dexterous hands in a single-arm setup. A Meta Quest 3 VR headset tracks its joystick pose to control the arm movements, while the operator uses the other hand with a glove to capture precise hand motions for dexterous hand manipulation, as shown in Fig. 2. For the ROHand, we use an OyMotion GForce glove, which measures the human hand motions corresponding to each dexterous hand joint using pressure tablet sensors, and maps these measurements into joint signals to control the dexterous hand.

Compared to pure vision-based systems [10, 25, 41, 62] and discretized gesture systems [33, 35] for dexterous hand teleoperation, our system leverages motion-capture gloves to provide more precise and intuitive hand control. When combined with arm teleoperation, prior approaches often coupled VR joysticks (or similar localization devices) directly with the gloves [35, 51], which we found to greatly restrict robotic arm rotation during teleoperation. In contrast, our decoupled design allows greater flexibility in arm motion. Furthermore, inspired by [15], we introduce a “pause” mechanism using a joystick button that lets the operator pause arm teleoperation at any time, reposition the joystick, and then resume control, as illustrated in Fig. 2. This further improves the flexibility

of our system, making it particularly suited for tasks that require substantial arm rotations, as demonstrated in §IV-D.

B. Dexterous Hand State Quantization

We obtain N demonstrations during data collection, where each demonstration is a trajectory $\{(o_i, s_i^{(a)}, s_i^{(h)})\}_i$, and $o_i, s_i^{(a)}, s_i^{(h)}$ denote the observation, arm state, and hand state at time step i , respectively. Previous methods [30, 55, 60] typically quantize concatenated arm-hand action chunks $\{(s_{i+k}^{(a)}, s_{i+k}^{(h)})\}_{k=1}^C$ with chunk size C . We argue, however, that this design is inappropriate for two reasons.

First, we should quantize hand actions only, rather than concatenated arm-hand actions (Fig. 4C) As discussed in §I, arm and hand actions serve fundamentally different purposes: the arm primarily manages spatial localization, while the hand governs interaction once the target region is reached. For example, when opening a jar, the arm must first move the hand to the correct position on the lid before the fingers perform the hooking action. Precise arm control is therefore essential for enabling proper finger interaction, whereas small inaccuracies in hand control are often tolerable. This makes hand actions a more natural target for quantization than concatenated arm-hand actions.

Second, we should quantize single-step hand actions, i.e., hand states, rather than hand action chunks (Fig. 4D) Hand action chunks capture temporal motion patterns of the dexterous hand, so quantizing them directly encodes these patterns into the codebook [60], leading to rapid codebook expansion compared to state-level quantization. Importantly, when hand action chunks are quantized, they must be classified from a discrete chunk codebook, whereas arm action chunks are typically generated via the diffusion process [2, 11, 19, 53]. This mismatch in action generation methods disentangles the two processes and can severely disrupt arm-hand coordination. For example, when opening a jar, the fingers’ hooking motion should only occur once the arm has correctly positioned the hand on the lid. If the hand motion is triggered too early due to the separation of arm and hand action chunk generation, the fingers may miss the lid or collide with the jar, causing manipulation failures.

To avoid these issues, we adopt single-step hand action quantization, i.e., hand state quantization. Concretely, we extract hand states $s^{(h)}$ from the demonstration dataset \mathcal{D} and train a two-layer residual VQ-VAE [66] to discretize them, as shown in Fig. 3①. Each hand state $s^{(h)}$ is encoded into a latent z_e , quantized by nearest-neighbor lookup in hierarchical codebooks $\{z_q\}$, and then decoded to reconstruct $\hat{s}^{(h)}$. The model is optimized with the standard VQ-VAE loss,

$$\mathcal{L} = \|s^{(h)} - \hat{s}^{(h)}\|_2^2 + \beta \| \text{sg}[z_e] - z_q \|_2^2 + \gamma \|z_e - \text{sg}[z_q]\|_2^2,$$

where $\text{sg}[\cdot]$ denotes the stop-gradient operator, and β, γ are weighting coefficients. The first term enforces reconstruction, while the latter two promote stable codebook usage.

C. Continuous Relaxation of Discretized Hand State

After quantizing hand states, we merge the multi-layer codebooks into a unified codebook with K discrete hand

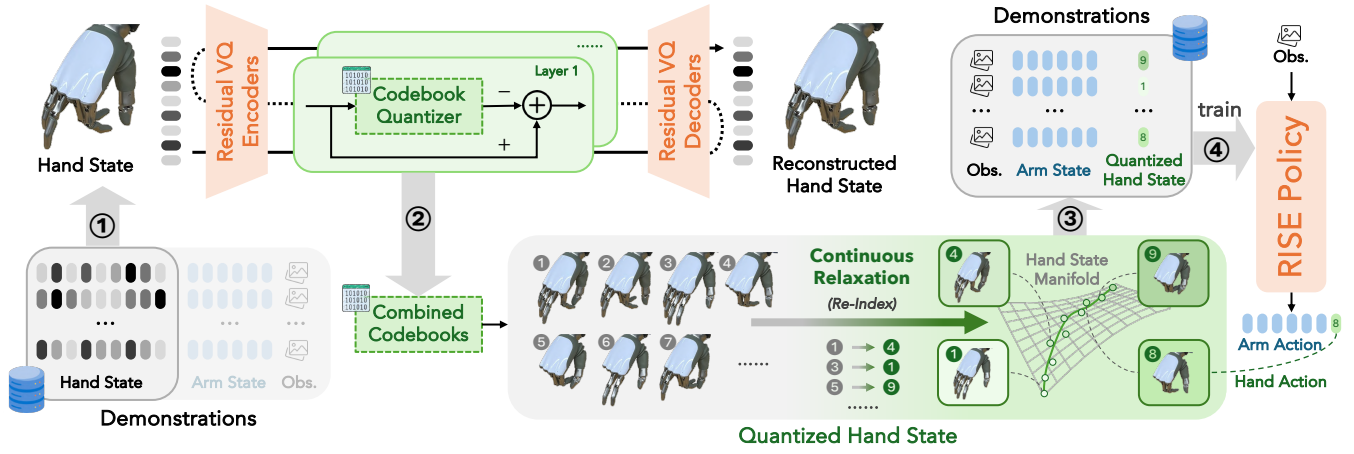


Fig. 3: DQ-RISE Policy Architecture. ① Hand state data from demonstrations are used to train a residual VQ-VAE [66] for hand state quantization (§III-B); ② The trained codebooks yield K quantized hand states, which are re-indexed to maintain consistency between consecutive codes and sequential continuity across all codes (§III-C); ③ The original hand states/actions are replaced by these re-indexed states in the demonstration dataset (§III-D); ④ The visuomotor policy is trained on the transformed dataset, jointly diffusing arm and hand actions; during inference, the predicted continuous hand actions are projected to the nearest quantized actions for execution (§III-D).

state codes. A natural approach is to formulate hand action prediction as a classification problem and use a classification head to predict future hand actions. However, as discussed in §III-B and validated in our experiments, decoupling arm and hand action generation often leads to mismatched actions and degraded performance. Hence, we consider *integrating the discrete hand states into arm action chunk diffusion*. We draw inspiration from the gripper: in practice, we usually only focus on whether it is open or closed, yet the visuomotor policy predicts a continuous value, as the transition from fully closed to fully open is inherently smooth and consistent. By analogy, if a consistent and continuous direction can be identified in the discrete hand state space, these states can likewise be predicted in a continuous manner and seamlessly integrated into the diffusion process.

To this end, we propose a continuous relaxation of the discretized hand states by re-indexing them in a continuous order. Instead of operating in the VQ-VAE latent space, we directly apply principal component analysis (PCA) [24] to the raw 6-DoF hand states. Projecting onto the first principal component, which captures the largest variance, provides a one-dimensional representation that reflects the dominant trend of hand motion. The quantized hand states are then sequentially re-indexed along this axis, as shown in Fig. 3②.

This design ensures that neighboring states in the re-indexed sequence correspond to similar hand configurations in the original hand state space, yielding a coherent ordering of gestures. By contrast, applying PCA on high-dimensional VQ-VAE features does not guarantee that adjacent indices represent semantically consistent hand poses. Our approach therefore provides a more interpretable continuous relaxation of the discretized hand states.

D. Visuomotor Policy Learning

After re-indexing the discretized hand states, we relabel each hand action $a^{(h)}$ in the dataset with its corresponding quantized, ordered index $z^{(h)}$. The demonstration trajectory

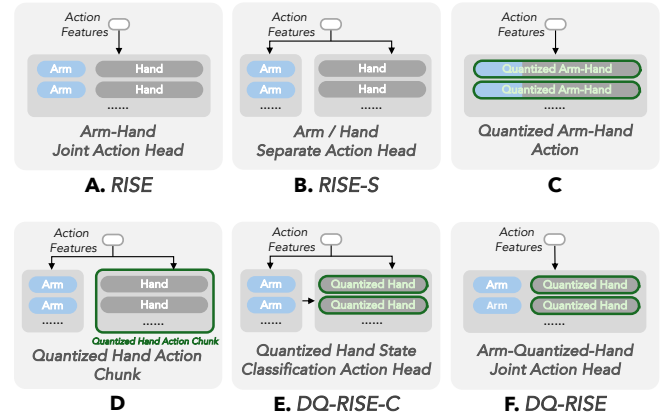


Fig. 4: Different Action Prediction Frameworks. We select RISE, RISE-S, DQ-RISE-C as baselines and compare with our DQ-RISE.

can thus be represented as $\{(o_i, (s_i^{(a)}, z_i^{(h)}))\}_i$, as shown in Fig. 3③. We then train a base visuomotor policy by using the observation o_i as input and a chunk of concatenated future arm and re-indexed hand actions $\{(s_{i+k}^{(a)}, z_{i+k}^{(h)})\}_{k=1}^C$ as output, as shown in Fig. 3④. During inference, the predicted continuous hand action $\hat{z}^{(h)}$ is mapped to its nearest quantized code $\text{idx} = \lceil \hat{z}^{(h)} \rceil$, from which the corresponding hand state $s_{\text{idx}}^{(h)}$ is retrieved for execution.

We adopt RISE [53] as the base visuomotor policy for its strong spatial generalization. Point clouds from the two cameras are first combined using the calibrated extrinsics, and cropped to the workspace region, then fed into the RISE policy to predict future arm-hand action chunks.

IV. EXPERIMENTS

In the experiments section, we aim to address the following research questions: (Q1) Can our **DQ-RISE** policy handle diverse dexterous manipulation tasks? (Q2) Which action prediction scheme yields the best performance in visuomotor policies for dexterous manipulation? (Q3) Is the continuous relaxation process of the **DQ-RISE** policy essen-

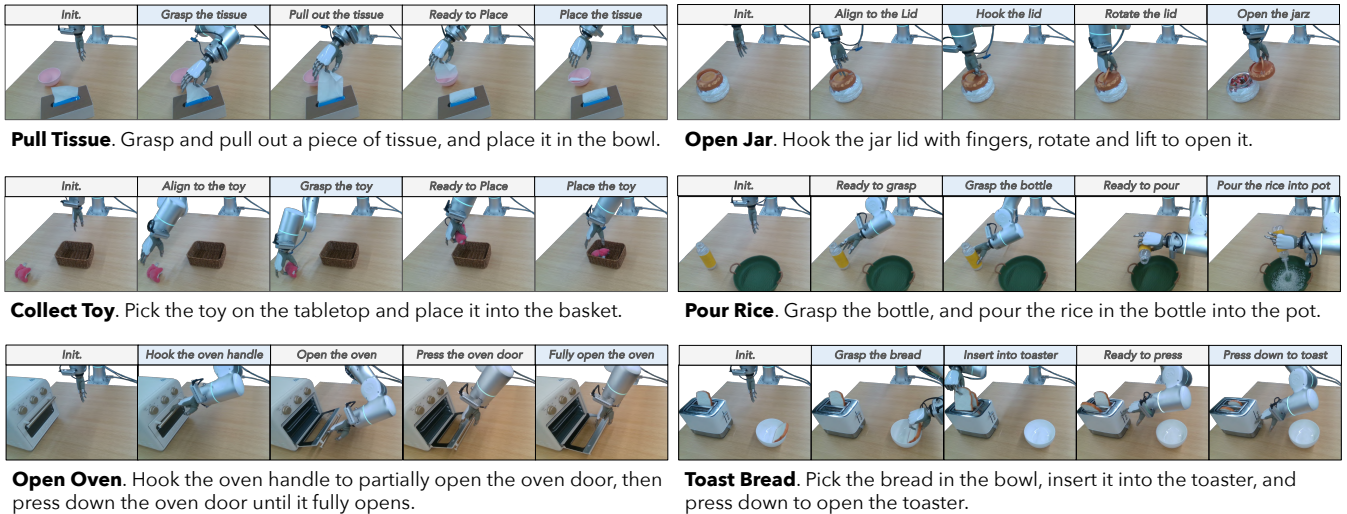


Fig. 5: Task Descriptions. We evaluate six tasks covering pick-and-place (*Pull Tissue*, *Collect Toy*), articulated object manipulation (*Open Jar*, *Open Oven*), tasks requiring large rotations (*Open Jar*, *Pour Rice*), and a long-horizon task (*Toast Bread*). Each task is illustrated with several phases, with the stages used for success rate evaluations highlighted in blue.

Policy	<i>Pull Tissue</i>		<i>Open Jar</i>		<i>Collect Toy</i>		<i>Pour Rice</i>		<i>Open Oven</i>		<i>Toast Bread</i>			Avg.
	<i>Grasp</i>	<i>Place</i>	<i>Hook</i>	<i>Open</i>	<i>Grasp</i>	<i>Place</i>	<i>Grasp</i>	<i>Pour</i>	<i>Hook</i>	<i>Press</i>	<i>Grasp</i>	<i>Insert</i>	<i>Press</i>	
<i>RISE</i> [53]	75%	45%	80%	55%	60%	60%	90%	80%	100%	90%	80%	20%	0%	55.00%
<i>RISE-S</i>	75%	55%	60%	45%	75%	70%	95%	85%	95%	95%	75%	25%	20%	61.67%
<i>DQ-RISE-C</i>	15%	10%	0%	0%	0%	0%	0%	0%	20%	5%	0%	0%	0%	2.50%
<i>DQ-RISE (ours)</i>	95%	85%	95%	90%	95%	80%	100%	100%	100%	100%	100%	65%	60%	85.83%

TABLE I: Evaluation Results. We report the success rates of every policy in certain task phases (Fig. 5). *DQ-RISE* outperforms other action prediction variants and can complete various dexterous manipulation tasks effectively, even on the challenging *Toast Bread* task.

tial for effective learning? (Q4) How does manual hand-state quantization during data collection differ from our automatic quantization in policy training? (Q5) Does our VR-glove hybrid teleoperation system provide a more intuitive single-arm teleoperation interface for dexterous hands?

A. Setup

Tasks. We design 6 tasks to evaluate the policies on dexterous manipulation in the real world, including pick-and-place operations, articulated object manipulation, tasks with significant rotations, and long-horizon tasks. Please refer to Fig. 5 for detailed descriptions of each task.

Baselines. We compare our *DQ-RISE* policy against three baselines for integrating dexterous hand action prediction: (1) the base visuomotor policy (*RISE*, Fig. 4A), which predicts concatenated arm-hand action chunks; (2) the base visuomotor policy with separate diffusions (*RISE-S*, Fig. 4B), which uses two diffusion heads for arm and hand action generation; and (3) the base visuomotor policy with quantized hand action classification (*DQ-RISE-C*, Fig. 4E), which diffuses arm actions first and then classifies quantized hand actions using action features and the predicted arm action. We omit other baselines (Fig. 4C and Fig. 4D) according to the previous discussions in §III-B.

Implementations. We train a two-layer residual VQ-VAE [66] with a codebook size of 4 for each layer, resulting in $K = 16$ quantized hand states per task. We set the both the commitment cost and codebook usage weight $\beta = \gamma = 1.67$. The residual VQ-VAE is optimized using Adam [26] with a

learning rate of 3×10^{-4} , a batch size of 256 for 1500 epochs. Other policy hyperparameters follow *RISE* [53].

Protocols. We use our VR-glove hybrid dexterous teleoperation system (§III-A) to collect 50 teleoperated demonstrations per task, and use these demonstrations to train the policies. We deploy the policies on a workstation with an NVIDIA RTX 3090 GPU. Following [12, 19, 57], object positions are randomized within the workspace before each task. Each policy is evaluated over 20 trials per task, and we report the overall success rate as the primary task completion metric, along with success rates broken down by task phases.

B. Results

***DQ-RISE* policy is able to effectively handle a wide range of dexterous manipulation tasks (Q1).** It achieves the highest success rates across all six evaluated tasks, with an average success rate of 85.83%. Beyond basic pick-and-place operations and articulated object manipulation such as opening an oven or a jar, *DQ-RISE* can also perform more complex long-horizon tasks like *Toast Bread*. Many of these tasks require the hand to adaptively switch between distinct poses at different stages. For example, in the *Toast Bread* task, the robot first uses its thumb and index finger to grasp the bread, and after placing it into the toaster, employs its index and middle fingers to press the button. Importantly, our quantized hand states prove sufficient for such tasks, reinforcing our design motivation: *the arm primarily handles localization, while the hand only needs to memorize certain action patterns*. Our quantization dramatically reduces the

size of the hand state space, enabling the policy to concentrate on the more challenging problem of precise arm localization during manipulation.

Our action prediction scheme with joint arm and quantized hand action prediction, achieves the best performance among all alternatives (Q2). This result further demonstrates that quantizing hand states facilitates effective arm action learning and arm-hand coordination, while still preserving fine-grained hand actions to a large extent. In contrast, the vanilla approach of directly concatenating arm and hand action spaces (*RISE*) struggles with fine-grained localization. For example, in the *Pull Tissue* and *Collect Toys* tasks, accurate localization is essential not only for grasping the tissue or toy but also for placing them precisely into the bowl or basket. These difficulties verify our hypothesis from §1 that the high-DoF hand state tends to dominate the action space, making learning less effective. Naively separating the arm and hand action predictions (*RISE-S*) alleviates this issue and improves performance on most tasks, but fails on the *Open Jar* task, where tight arm-hand coordination is crucial to hook and rotate the lid, as illustrated in Fig. 1.

Combining classification with diffusion-based action generation on the same conditioning feature can introduce inconsistent gradient flows during training, ultimately degrading policy performance (Q2). In our experiments, *DQ-RISE-C* rarely succeeds in completing tasks. We observe that while the policy often predicts arm actions quite well, failures in hand state classification — such as prematurely changing the hand pose — distort subsequent observations, push the arm trajectory out of distribution, and ultimately cause rollouts to collapse. One possible explanation is that, since quantized hand actions are classified conditioned on the predicted arm action, distribution shift in arm predictions during inference could harm classification. However, an ablation study on whether to include this arm-conditioning route (Fig. 6A) shows a negligible effect, ruling out this factor. We thus attribute the failure primarily to inconsistent gradient flows between the arm diffusion head and the hand classification head, which hinder effective joint optimization and result in suboptimal learning. Similar interference between classification and regression objectives has been documented in multi-task learning [47, 65], and comparable observations have been made in other visuomotor policies [21]. These findings further validate our design choice to apply continuous relaxation to quantized hand states, avoiding a classification objective and ensuring consistent gradient propagation during training.

C. Ablations

The continuous relaxation process is essential for effective policy learning (Q3). We select the *Open Jar* task as an example to ablate the function of the continuous relaxation process. As shown in Fig. 6B, removing re-indexing leads to a substantial drop in policy performance, whereas our policy achieves a much higher success rate. Without continuous ordering, neighboring code indices may correspond to very different hand states, making policy

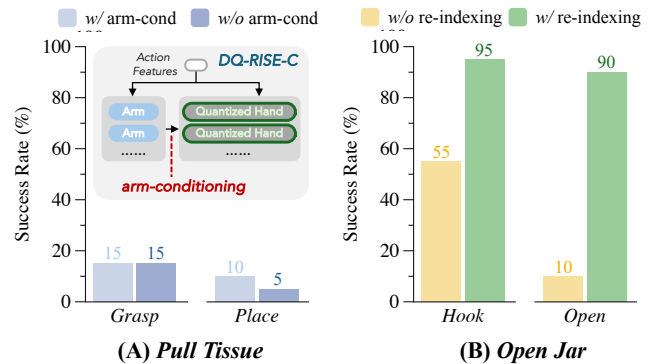


Fig. 6: (A) Ablation of Arm Conditioning in *DQ-RISE-C*. The policies perform similarly regardless of whether arm conditioning is applied during hand code classification. **(B) Ablation of Continuous Relaxation in *DQ-RISE*.** Perform continuous relaxation (*i.e.*, re-indexing) enhances the interpretability of the predicted hand codes, and significantly boosts performance.

learning difficult and unstable. Moreover, non-continuous codes reduce the policy’s tolerance to prediction errors: even a small mistake could map to a drastically different hand configuration. In contrast, continuous code indices ensure that nearby predictions correspond to similar hand states, improving robustness and allowing the policy to coordinate arm and hand actions more reliably. This confirms that continuous relaxation is a key component in *DQ-RISE*.

The continuous relaxation process makes quantized hand state codes interpretable and allows their indices to be predicted continuously (Q3). As illustrated in Fig. 7, neighboring code indices correspond to smoothly varying hand poses, and interpolation between indices yields meaningful intermediate hand configurations. This continuous, structured representation enables the policy to predict arm and hand actions jointly, without needing separate treatments for arm prediction and hand prediction.

D. User Study

We conduct a user study to evaluate our VR-glove hybrid dexterous teleoperation system against three alternatives: (1) *a coupled arm-hand control system*, where the VR joystick is attached to the user’s wrist and the glove controls the robotic hand; (2) *a discretized gesture system*, in which users control the hand via discrete gestures (*e.g.*, keyboard inputs) while using the same VR device for arm motion; and (3) *a variant without a pausing mechanism*, which does not allow pausing the arm during teleoperation. Six participants with

Teleoperation System	Open Jar		Avg. Rank ↓
	Success Rate ↑	Time (s) ↓	
Coupled arm-hand control	5 / 6	25.17	3.83
Ours w/ discretized gesture	6 / 6	20.50	2.83
Ours w/o pausing	6 / 6	16.67	2.25
Ours	6 / 6	13.83	1.08

TABLE II: User Study Results. Our arm-hand decoupled teleoperation system with a pausing mechanism is both intuitive and convenient to control the arm and the dexterous hand, improving success rates and reducing completion time during teleoperation.

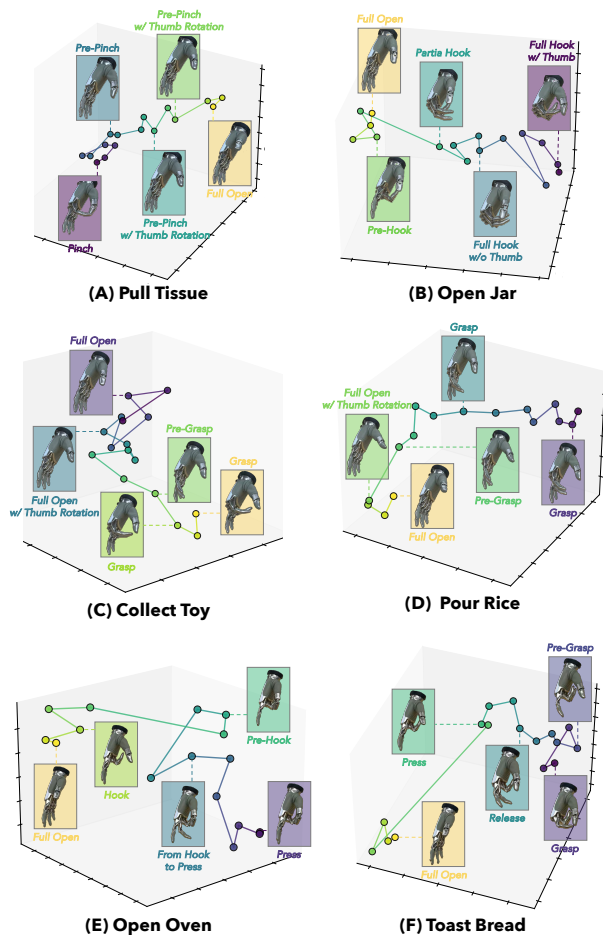


Fig. 7: Quantized Hand State after Re-Indexing. Hand states are projected into 3D points via UMAP [37], with selected points annotated by their corresponding hand poses for reference. Re-indexing in the continuous relaxation process makes code transitions continuous and interpretable in the hand states, supporting further joint arm action and quantized hand action diffusing.

varying teleoperation experience has 3 minutes to familiarize themselves with each system before attempting the *Open Jar* task in randomized order. Afterward, they are asked to rank the systems by intuitiveness and convenience.

Our VR-glove hybrid teleoperation system provides a more intuitive single-arm interface for dexterous hands (Q5). As shown in Tab. II, it achieves the highest success rate (6/6), shortest completion time (13.83s), and best average rank (1.08). The coupled arm-hand system performs the worst due to the severe rotations required for the *Open Jar* task, making the teleoperation control difficult. The discretized gesture and no-pausing variants improve over it, but still fall short in intuitiveness and convenience, respectively. These results demonstrate that integrating continuous glove-based hand control with VR-based arm control and a pausing mechanism enhances efficiency, reduces cognitive load, and improves coordination, confirming our system’s superiority for single-arm dexterous manipulation.

V. CONCLUSION

In this work, we explore visuomotor policy learning on dexterous manipulations from the action prediction per-

spective. Our study highlights that successful manipulation depends not only on fine-grained hand motions but also on precise arm localization and coherent arm-hand coordination. Existing approaches, which either combine or fully separate arm and hand action spaces, often suffer from imbalanced learning or degraded coordination. To address these challenges, we propose *DQ-RISE*, which quantizes dexterous hand states into a compact set of task-relevant patterns and jointly diffuses them with arm actions, allowing the policy to focus on accurate arm localization while retaining dexterous hand capabilities. We also introduce a hybrid dexterous teleoperation system to support intuitive and convenient demonstration collection. Experiments across diverse dexterous tasks show that *DQ-RISE* consistently achieves the best overall performance, validating the effectiveness of our structured action prediction framework.

Looking ahead, we see two promising directions for future research. First, extending *DQ-RISE* to multi-task learning presents challenges, as hand state quantization may introduce diverse hand states as codes when task distributions vary widely, potentially requiring adaptive or hierarchical quantization schemes. Second, the continuous relaxation that underpins joint arm-hand diffusion could become more difficult to stabilize in multi-task or long-horizon settings, motivating the exploration of improved relaxation techniques or hybrid discrete-continuous formulations. Addressing these challenges would further advance the scalability and generalization of visuomotor policies, paving the way for robust and versatile dexterous manipulation in real-world environments.

ACKNOWLEDGEMENT

This work was supported in part by the National Natural Science Foundation of China (No. 62595774), the Shanghai Committee of Science and Technology, China (Grant No. 24511103200), Shanghai Artificial Intelligence Laboratory, XPLOER PRIZE grants.

REFERENCES

- [1] Marcin Andrychowicz et al. “Learning Dexterous In-Hand Manipulation”. In: *Int. J. Robotics Res.* 39.1 (2020).
- [2] Kevin Black et al. “ π_0 : A Vision-Language-Action Flow Model for General Robot Control”. In: *RSS*. 2025.
- [3] Vittorio Caggiano, Sudeep Dasari, and Vikash Kumar. “MyoDex: A Generalizable Prior for Dexterous Manipulation”. In: *ICML*. 2023, pp. 3327–3346.
- [4] Chi-Lam Cheang et al. “GR-2: A Generative Video-Language-Action Model with Web-Scale Knowledge for Robot Manipulation”. In: *arXiv preprint arXiv:2410.06158* (2024).
- [5] Tao Chen, Jie Xu, and Pulkit Agrawal. “A System for General In-Hand Object Re-Orientation”. In: *CoRL*. 2021, pp. 297–307.
- [6] Tao Chen et al. “Visual Dexterity: In-hand Reorientation of Novel and Complex Object Shapes”. In: *Sci. Robotics* 8.84 (2023).
- [7] Yi Chen et al. “Moto: Latent Motion Token as the Bridging Language for Robot Manipulation”. In: *ICCV*. 2025.
- [8] Yuanpei Chen et al. “Bi-DexHands: Towards Human-Level Bimanual Dexterous Manipulation”. In: *IEEE Trans. Pattern Anal. Mach. Intell.* 46.5 (2024), pp. 2804–2818.
- [9] Zeyuan Chen et al. “ClutterDexGrasp: A Sim-to-Real System for General Dexterous Grasping in Cluttered Scenes”. In: *CoRL*. 2025.
- [10] Xuxin Cheng et al. “Open-TeleVision: Teleoperation with Immersive Active Visual Feedback”. In: *CoRL*. 2024, pp. 2729–2749.
- [11] Cheng Chi et al. “Diffusion Policy: Visuomotor Policy Learning via Action Diffusion”. In: *RSS*. 2023.

- [12] Cheng Chi et al. “Universal Manipulation Interface: In-The-Wild Robot Teaching Without In-The-Wild Robots”. In: *RSS*. 2024.
- [13] Patrick Esser, Robin Rombach, and Björn Ommer. “Taming Transformers for High-Resolution Image Synthesis”. In: *CVPR*. 2021, pp. 12873–12883.
- [14] Hao-Shu Fang et al. “AnyGrasp: Robust and Efficient Grasp Perception in Spatial and Temporal Domains”. In: *IEEE Trans. Robotics* 39.5 (2023), pp. 3929–3945.
- [15] Hao-Shu Fang et al. “RH20T: A Comprehensive Robotic Dataset for Learning Diverse Skills in One-Shot”. In: *ICRA*. 2024, pp. 653–660.
- [16] Hao-Shu Fang et al. “AnyDexGrasp: General Dexterous Grasping for Different Hands with Human-level Learning Efficiency”. In: *arXiv preprint arXiv:2502.16420* (2025).
- [17] Hao-Shu Fang et al. “DEXOP: A Device for Robotic Transfer of Dexterous Human Manipulation”. In: *arXiv preprint arXiv:2509.04441* (2025).
- [18] Hongjie Fang et al. “AirExo: Low-Cost Exoskeletons for Learning Whole-Arm Manipulation in the Wild”. In: *ICRA*. 2024, pp. 15031–15038.
- [19] Hongjie Fang et al. “AirExo-2: Scaling up Generalizable Robotic Imitation Learning with Low-Cost Exoskeletons”. In: *CoRL*. 2025.
- [20] Ian J. Goodfellow et al. “Generative Adversarial Nets”. In: *NeurIPS*. 2014, pp. 2672–2680.
- [21] Zihao He et al. “FoAR: Force-Aware Reactive Policy for Contact-Rich Robotic Manipulation”. In: *IEEE Robotics Autom. Lett.* 10.6 (2025), pp. 5625–5632.
- [22] Liang Heng et al. “ViTacFormer: Learning Cross-Modal Representation for Visuo-Tactile Dexterous Manipulation”. In: *arXiv preprint arXiv:2506.15953* (2025).
- [23] Jonathan Ho, Ajay Jain, and Pieter Abbeel. “Denoising Diffusion Probabilistic Models”. In: *NeurIPS*. 2020, pp. 6840–6851.
- [24] Harold Hotelling. “Analysis of A Complex of Statistical Variables into Principal Components.” In: *Journal of Educational Psychology* 24.6 (1933), p. 417.
- [25] Tatsuya Kamijo, Cristian C Beltran-Hernandez, and Masashi Hamaya. “Learning Variable Compliance Control from a Few Demonstrations for Bimanual Robot with Haptic Feedback Teleoperation System”. In: *IROS*. 2024, pp. 12663–12670.
- [26] Diederik P. Kingma and Jimmy Ba. “Adam: A Method for Stochastic Optimization”. In: *ICLR*. 2015.
- [27] Diederik P. Kingma and Max Welling. “Auto-Encoding Variational Bayes”. In: *ICLR*. 2014.
- [28] Dong-Hyuk Lee et al. “Peg-in-Hole Assembly With Dual-Arm Robot and Dexterous Robot Hands”. In: *IEEE Robotics Autom. Lett.* 7.4 (2022), pp. 8566–8573.
- [29] Jason Lee et al. “MolmoAct: Action Reasoning Models that Can Reason in Space”. In: *arXiv preprint arXiv:2508.07917* (2025).
- [30] Seungjae Lee et al. “Behavior Generation with Latent Actions”. In: *ICML*. OpenReview.net, 2024.
- [31] Kailin Li et al. “ManipTrans: Efficient Dexterous Bimanual Manipulation Transfer via Residual Learning”. In: *CVPR*. 2025, pp. 6991–7003.
- [32] Shuang Li et al. “Unified Video Action Model”. In: *RSS*. 2025.
- [33] Toru Lin et al. “Learning Visuotactile Skills With Two Multifingered Hands”. In: *ICRA*. 2025, pp. 5637–5643.
- [34] Toru Lin et al. “Sim-to-real reinforcement learning for vision-based dexterous manipulation on humanoids”. In: *CoRL*. 2025.
- [35] Yuhao Lin et al. “TypeTele: Releasing Dexterity in Teleoperation by Dexterous Manipulation Types”. In: *CoRL*. 2025.
- [36] Zhao Mandi et al. “DexMachina: Functional Retargeting for Bimanual Dexterous Manipulation”. In: *arXiv preprint arXiv:2505.24853* (2025).
- [37] Leland McInnes et al. “UMAP: Uniform Manifold Approximation and Projection”. In: *Journal of Open Source Software* 3.29 (2018), p. 861.
- [38] Anusha Nagabandi et al. “Deep Dynamics Models for Learning Dexterous Manipulation”. In: *CoRL*. 2019, pp. 1101–1112.
- [39] Allison M. Okamura, Niels Smaby, and Mark R. Cutkosky. “An Overview of Dexterous Manipulation”. In: *ICRA*. 2000, pp. 255–262.
- [40] Aäron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. “Neural Discrete Representation Learning”. In: *NeurIPS*. 2017, pp. 6306–6315.
- [41] Younghyo Park et al. “DART: Dexterous Augmented Reality Teleoperation Platform for Large-Scale Robot Data Collection in Simulation”. In: 2025, pp. 13883–13889.
- [42] Yuzhe Qin et al. “DexMV: Imitation Learning for Dexterous Manipulation from Human Videos”. In: *ECCV*. 2022, pp. 570–587.
- [43] Yuzhe Qin et al. “DexPoint: Generalizable Point Cloud Reinforcement Learning for Sim-to-Real Dexterous Manipulation”. In: *CoRL*. 2022, pp. 594–605.
- [44] Yuzhe Qin et al. “AnyTeleop: A General Vision-Based Dexterous Robot Arm-Hand Teleoperation System”. In: *RSS*. 2023.
- [45] Ali Razavi, Aäron van den Oord, and Oriol Vinyals. “Generating Diverse High-Fidelity Images with VQ-VAE-2”. In: *NeurIPS*. 2019, pp. 14837–14847.
- [46] Branden Romero et al. “EyeSight Hand: Design of a Fully-Actuated Dexterous Robot Hand with Integrated Vision-Based Tactile Sensors and Compliant Actuation”. In: *IROS*. 2024, pp. 1853–1860.
- [47] Ozan Sener and Vladlen Koltun. “Multi-Task Learning as Multi-Objective Optimization”. In: *NeurIPS*. 2018, pp. 525–536.
- [48] Kenneth Shaw, Ananye Agarwal, and Deepak Pathak. “LEAP Hand: Low-Cost, Efficient, and Anthropomorphic Hand for Robot Learning”. In: *RSS*. 2023.
- [49] Kenneth Shaw et al. “Bimanual Dexterity for Complex Tasks”. In: *CoRL*. 2024, pp. 5166–5183.
- [50] Zilin Si et al. “Tilde: Teleoperation for Dexterous In-Hand Manipulation Learning with a DeltaHand”. In: *RSS*. 2024.
- [51] Chen Wang et al. “DexCap: Scalable and Portable Mocap Data Collection System for Dexterous Manipulation”. In: *RSS*. 2024.
- [52] Chenxi Wang et al. “Graspness Discovery in Clutters for Fast and Accurate Grasp Detection”. In: *ICCV*. 2021, pp. 15944–15953.
- [53] Chenxi Wang et al. “RISE: 3D Perception Makes Real-World Robot Imitation Simple and Effective”. In: *IROS*. 2024, pp. 2870–2877.
- [54] Ruicheng Wang et al. “DexGraspNet: A Large-Scale Robotic Dexterous Grasp Dataset for General Objects Based on Simulation”. In: *ICRA*. 2023, pp. 11359–11366.
- [55] Yating Wang et al. “VQ-VLA: Improving Vision-Language-Action Models via Scaling Vector-Quantized Action Tokenizers”. In: *ICCV*. 2025.
- [56] Philipp Wu et al. “GELLO: A General, Low-Cost, and Intuitive Teleoperation Framework for Robot Manipulators”. In: *IROS*. 2024, pp. 12156–12163.
- [57] Shangning Xia et al. “CAGE: Causal Attention Enables Data-Efficient Generalizable Robotic Manipulation”. In: *ICRA*. 2025, pp. 13242–13249.
- [58] Mengda Xu et al. “DexUMI: Using Human Hand as the Universal Manipulation Interface for Dexterous Manipulation”. In: *CoRL*. 2025.
- [59] Yinzhen Xu et al. “UniDexGrasp: Universal Robotic Dexterous Grasping via Learning Diverse Proposal Generation and Goal-Conditioned Policy”. In: *CVPR*. 2023, pp. 4737–4746.
- [60] Chenyu Yang, Davide Liconti, and Robert K Katzschmann. “VQ-ACE: Efficient Policy Search for Dexterous Robotic Manipulation via Action Chunking Embedding”. In: *arXiv preprint arXiv:2411.03556* (2024).
- [61] Fan Yang et al. “Task-Oriented Tool Manipulation With Robotic Dexterous Hands: A Knowledge Graph Approach From Fingers to Functionality”. In: *IEEE Trans. Cybern.* 55.1 (2025), pp. 395–408.
- [62] Shiqi Yang et al. “ACE: A Cross-platform and visual-Exoskeletons System for Low-Cost Dexterous Teleoperation”. In: *CoRL*. 2024, pp. 4895–4911.
- [63] Seonghyeon Ye et al. “Latent Action Pretraining from Videos”. In: *ICLR*. 2025.
- [64] Zhao-Heng Yin et al. “Geometric Retargeting: A Principled, Ultrafast Neural Hand Retargeting Algorithm”. In: *IROS*. 2025.
- [65] Tianhe Yu et al. “Gradient Surgery for Multi-Task Learning”. In: *NeurIPS*. 2020.
- [66] Neil Zeghidour et al. “SoundStream: An End-to-End Neural Audio Codec”. In: *IEEE ACM Trans. Audio Speech Lang. Process.* 30 (2022), pp. 495–507.
- [67] Han Zhang et al. “DOGlove: Dexterous Manipulation with a Low-Cost Open-Source Haptic Force Feedback Glove”. In: *RSS*. 2025.
- [68] Jianke Zhang et al. “UP-VLA: A Unified Understanding and Prediction Model for Embodied Agent”. In: *ICML*. 2025.
- [69] Tony Z Zhao et al. “Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware”. In: *RSS*. 2023.
- [70] Henry Zhu et al. “Dexterous Manipulation with Deep Reinforcement Learning: Efficient, General, and Low-Cost”. In: *ICRA*. 2019, pp. 3651–3657.