

# Discrete Policy: Learning Disentangled Action Space for Multi-Task Robotic Manipulation

Kun Wu<sup>1,\*</sup>, Yichen Zhu<sup>2,\*</sup>, Jinming Li<sup>3</sup>, Junjie Wen<sup>4</sup>,  
Ning Liu<sup>2</sup>, Zhiyuan Xu<sup>5</sup>, Qinru Qiu<sup>1</sup>, and Jian Tang<sup>5,†</sup>

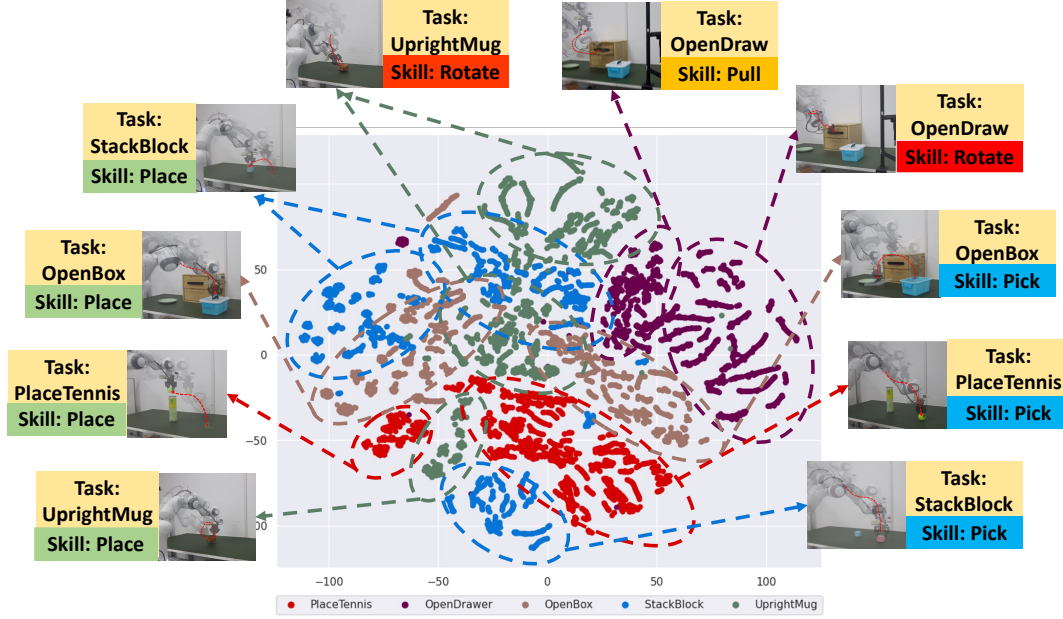


Fig. 1: **Visualization of Discrete Policy.** The t-SNE visualization of feature embeddings from Discrete Policy reveals that skills across different tasks cluster closely together. This pattern suggests that discrete latent spaces are capable of disentangling the complex, multimodal action distributions encountered in multi-task policy learning.

**Abstract**—Learning visuomotor policy for multi-task robotic manipulation has been a long-standing challenge for the robotics community. The difficulty lies in the diversity of action space: typically, a goal can be accomplished in multiple ways, resulting in a multimodal action distribution for a single task. The complexity of action distribution escalates as the number of tasks increases. In this work, we propose Discrete Policy, a robot learning method for training universal agents capable of multi-task manipulation skills. Discrete Policy employs vector quantization to map action sequences into a discrete latent space, facilitating the learning of task-specific codes. These codes are then reconstructed into the action space conditioned on observations and language instruction. We evaluate our method on both simulation and multiple real-world embodiments, including both single-arm and bimanual robot settings. We demonstrate that our proposed Discrete Policy outperforms a well-established Diffusion Policy baseline and

many state-of-the-art approaches, including ACT, Octo, and OpenVLA. For example, in a real-world multi-task training setting with five tasks, Discrete Policy achieves an average success rate that is 26% higher than Diffusion Policy and 15% higher than OpenVLA. As the number of tasks increases to 12, the performance gap between Discrete Policy and Diffusion Policy widens to 32.5%, further showcasing the advantages of our approach. Our work empirically demonstrates that learning multi-task policies within the latent space is a vital step toward achieving general-purpose agents. Our project is at <https://discretepolicy.github.io>.

## I. INTRODUCTION

In the realm of robotics, the ability to train robots for multi-tasking operations presents significant challenges that stem from the inherent complexity of handling diverse tasks simultaneously [1]. Traditional robotic systems often focus on specialized tasks, but the dynamic environments in which modern robots operate demand versatile functionalities that can adapt to various situations.

However, the unique challenges of predicting robotic actions — such as multimodal action distributions, sequential correlations, the need for high precision, and noise in expert demonstration data — make learning a single task through imitation learning a formidable problem. When extending

<sup>1</sup>Syracuse University, NY, USA {kwu102, qiqiu}@syr.edu

<sup>2</sup>Midea Group, AI Research Center, China {zhuyc25, liuning22}@midea.com

<sup>3</sup>Shanghai University, China ljm2022@shu.edu.cn

<sup>4</sup>East China Normal University, China jjwen@cs.ecnu.edu.cn

<sup>5</sup>Beijing Innovation Center of Humanoid Robotics, Beijing, China {Eric.Xu, jian.tang}@x-humanoid.com

\* Co-first author. † Corresponding author: Jian Tang. Work done during Kun Wu’s internship at Midea Group.

this approach to multi-task imitation learning, these challenges intensify. The tasks may involve a mixture of short-, medium-, and long-horizon objectives of varying lengths, further complicating the action space. This complexity results in an entangled, more diverse multimodal action distribution as the characteristics of single tasks merge. This fusion exacerbates the difficulty of accurately learning and executing multiple tasks simultaneously. Naive imitation learning, such as Behavior Cloning [2] and Diffusion Policy [3], directly maps the observation and instruction to the action space. This mapping operates on the task-entangled dimension, which makes it hard for the policy network to distinguish different tasks from them, not to mention learning precise behavior.

In this work, we present Discrete Policy, a framework for multi-task visuomotor policy learning. Discrete Policy consists of two components: a vector-quantized variation autoencoder (VQ-VAE) and a conditional diffusion model. During training, the VQ-VAE is used to project the action sequence into discrete latent space and extract the discrete action modal abstraction  $z$  for the current task. This  $z$  represents the latent feature embedding of the current task. Then, a conditional diffusion model is leveraged to conduct the denoising process in the discrete latent space. Finally, the task-specific latent embedding is mapped to the decoder of VQ-VAE to reconstruct the action space, conditioned on the observation and language instruction.

To illustrate our method, we present a visualization that employs t-SNE [4] to demonstrate the feature embedding of Discrete Policy on training with five real-world tasks, which is shown in Figure 1. Intriguingly, we find that tasks with similar characteristics are located in proximity to each other. For example, the feature embeddings on the left primarily represent the skills associated with “placing”, while those at the bottom correspond to “picking” skills. These skills, which appear across various tasks, have closely situated feature embeddings, yet they still maintain discernible boundaries between them. Our observations suggest that a discrete latent space may be more effective for disentangling tasks and skills in policy learning. We conducted experiments across simulation and real-world data, including 23 tasks from RL-Bench, 12 tasks from single-arm robots, and 5 tasks from bimanual robots. These task types range from simple object pick-and-place operations to more complex, contact-rich manipulations like placing objects into a drawer. Our results demonstrate that our method, Discrete Policy, significantly outperforms Diffusion Policy [5], a strong baseline, by a considerable margin. Notably, the performance gap between Discrete Policy and Diffusion Policy widens as the number of tasks incorporated into the training increases. Our approach also shows superior performance over multiple state-of-the-art methods including MT-ACT [6], Octo [7], and OpenVLA [8].

## II. RELATED WORK

**Diffusion model for policy learning.** Diffusion models belong to generative models. It progressively transforms random noise into a data sample, which has achieved great

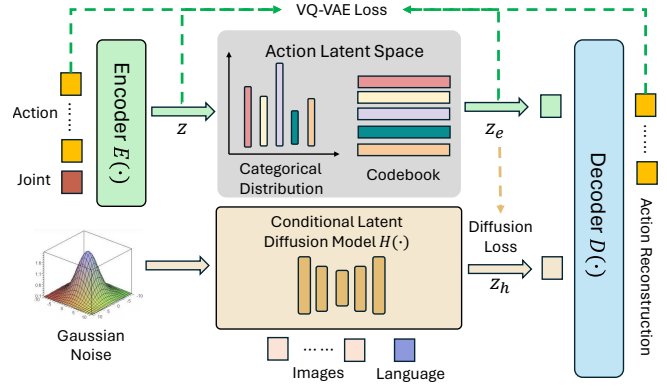


Fig. 2: Overview of Discrete Policy. In the first training stage, as indicated by the green arrow, we train a VQ-VAE that maps actions into discrete latent space with an encoder and then reconstructs the actions based on the latent embeddings using a decoder. In the second training stage, as indicated by the brown arrow, we train a latent diffusion model that predicts task-specific latent embeddings to guide the decoder in predicting accurate actions.

success in generating high-fidelity image and image editing [9]–[12]. Recently, diffusion models have been applied in robotics, combining with reinforcement learning [13]–[19], imitation learning [3], [7], [20]–[34], reward learning [35], [36], grasping [37]–[39], and motion planning [37], [40]–[45]. In this work, we leverage the diffusion model to perform noising-denoising on the discrete latent space, instead of high-dimensional feature space as conventional approaches do.

**Multi-tasking in robotics.** Multi-task robotic manipulation has led to significant progress in the execution of complex tasks and the ability to generalize to new scenarios [25], [27], [46]–[53]. Notable methods often involve the use of extensive interaction data [47], [48], [54] to train multi-task models. For example, RT-1 [46] underscores the benefits of task-agnostic training, demonstrating superior performance in real-world robotic tasks across a variety of datasets. RT-2 [55] trains with mixed robot data and image-text pairs. PerAct [56] encodes language goals and shows its effectiveness in real robot experiments. Octo [7] uses cross-embodiment data for pertaining. This paper proposes a new approach to learning multi-task policy in the discrete latent space.

## III. PRELIMINARIES

**Denoising Diffusion Probabilistic Models (DDPMs)** [57], as a class of expressive generative models, leverage Stochastic Langevin Dynamics [58] to model the output generation process. During the inference phase, starting from a Gaussian noise  $x^K$ , DDPMs perform  $K$  iterations of the denoising process to generate a sequence of intermediate results  $x^{k-1}, \dots, x^0$  and take  $x^0$  as the final denoised output.

To train the DDPMs, a denoising iteration  $k$  is randomly selected for each data sample in the training dataset, and the

corresponding noise  $\epsilon^k$  is sampled from the noise scheduler. Then the noise prediction network  $H(\cdot)$  is trained with the following objective:

$$\mathcal{L}_{ddpm} = \|\epsilon^k - H(x^0 + \epsilon^k, k)\|_2. \quad (1)$$

In Discrete Policy, we employ the Denoising Diffusion Implicit Models (DDIM) [59] as the latent diffusion model, which improves upon DDPMs with faster inference speed.

#### IV. METHODOLOGY

##### A. Overview

In this work, we introduce Discrete Policy, designed to simultaneously learn multiple robotic manipulation tasks. As depicted in the Figure 2, Discrete Policy comprises two main components: 1) **Training stage 1:** We train a Vector Quantized Variational Autoencoder (VQ-VAE) [60] that encodes complex actions into a discrete latent space and subsequently reconstructs these actions using a decoder. 2) **Training stage 2:** We utilize a latent diffusion model to generate task-specific latent embeddings, which guide the decoder in executing the appropriate action modality.

**Formal definition.** At each inference timestamp  $t$ , Discrete Policy take a language instruction  $\mathbf{l}$  and current observations  $\mathbf{o} = \langle \mathbf{c}, \mathbf{m} \rangle$  as input. The  $\mathbf{c} \in \mathbb{R}^{2 \times 3 \times H \times W}$  consists of two RGB images from the external fixed cameras on the left and right, and  $\mathbf{m} \in \mathbb{R}^d$  represents the low-dimensional proprioceptive states. The latent diffusion model, conditioned on  $\mathbf{l}$  and  $\mathbf{o}$ , performs a denoising process to generate the latent action embedding  $z$ . Subsequently, the VQ-VAE decoder uses the embedding  $z$ , along with  $\mathbf{l}$  and  $\mathbf{o}$ , to predict the final action  $\hat{a}$ .

##### B. Vector Quantized Autoencoder for Multimodal Action

We detail the first training stage for VQ-VAE. Unlike the vanilla Variational Autoencoder (VAE) [61], which uses a continuous latent space and a Gaussian distribution prior, the VQ-VAE’s discrete latent space simplifies the differentiation of skills and action modalities in robotic manipulation tasks. This facilitates more precise action predictions by the decoder.

To unify and compress different action modalities across tasks into a single latent space, the encoder  $E(\cdot)$  processes the actions  $a_{t:t+h}$  over a fixed horizon  $h$  and the corresponding low-dimensional proprioceptive states  $\mathbf{m}_{t:t+h}$ , producing the latent action embedding  $z = E(a_{t:t+h}, \mathbf{m}_{t:t+h})$ . We include the proprioceptive state in the input because it directly correlates with the robotic arm’s actions, where  $a$  represents the 6D pose of the end effector, and  $\mathbf{m}$  denotes the robotic joint positions in our implementation. To ensure the encoder  $E(\cdot)$  focuses on learning from the actions and not the auxiliary inputs, we do not include images  $\mathbf{c}$  or language instructions  $\mathbf{l}$  in the encoder’s input.

To discretize the latent action space in the autoencoder, we maintain a codebook of discrete latent embedding  $e \in \mathbb{R}^{c \times f}$ , where  $c$  represents the number of the latent embedding categories and  $f$  is the embedding dimension. Given the latent action embedding  $z$  from the encoder, the discrete

latent embedding  $z_e$  is selected through a bottleneck layer  $S(\cdot)$  following the nearest neighbor lookup rule as follows:

$$z_e = S(z) = e_j, \text{ where } j = \operatorname{argmin}_i \|z - e_i\|_2, \quad (2)$$

where  $e_i$  is the discrete latent embedding for the  $i$ -th category in the codebook. Then, the decoder  $D(\cdot)$  takes the discrete latent embedding  $z_e$  combined with the language instruction  $\mathbf{l}$  and current observations  $\mathbf{o}$  to predict the action  $\hat{a}$ . Similar to the loss function in VQ-VAE [60], the final objective consists of three terms:

$$\mathcal{L}_{vq} = \|\hat{a} - a\|_1 + \beta \|sg(z) - z_e\|_2^2 + \beta \|z - sg(z_e)\|_2^2, \quad (3)$$

where  $sg(\cdot)$  is the operator for stopping the gradient, the first term is to reconstruct the input, the second term is the quantization loss to train the codebook, and the third term is the commitment loss to encourage the encoder to commit to a code.  $\beta$  is the coefficient to balance the losses, which is set to 1.0 in our implementation

##### C. Latent Diffusion Model

After training the VQ-VAE, a naive sampling method is to use a uniform categorical distribution and select one latent embedding randomly from the codebook  $e$ . However, for the multi-task robotic manipulation problem, this method risks selecting a prior latent embedding  $z_e$  that is inappropriate for the current task and leads to task failure. Therefore, we freeze the trained VQ-VAE and employ a conditional latent diffusion model  $H(\cdot)$  to generate the discrete latent embedding  $z_e$  suitable for the current task. Given the noisy latent embedding  $z^k$  at the denoise iteration  $k$  and the language instruction  $\mathbf{l}$  and current observations  $\mathbf{o}$ , the latent diffusion model  $H(\cdot)$  outputs the corresponding noise for the task-specific latent embedding  $z_h$ . Instead of directly predicting the latent embedding  $z_h$ , the latent diffusion model  $H(\cdot)$  predicts the noise at each iteration and tries to recover the latent embedding  $z_h$  from the Gaussian noise through multi-round denoising processes. Following [21], we use DDIM [62] as the latent diffusion model to improve the inference efficiency and use the same noise scheduler and hyperparameters. The training objective is as follows:

$$\mathcal{L}_{ldm} = \|\epsilon^k - H(z_e^k, \mathbf{l}, \mathbf{o}, k)\|_2, \quad (4)$$

where  $k$  represents the  $k$ -th denoise iteration. During the inference, starting from Gaussian noise, the latent diffusion model  $H(\cdot)$  performs  $K$  times denoise processes and recovers the latent embedding  $z_h$ . Then the latent embedding  $z_h$  is passed to the bottleneck layer  $S(\cdot)$  in VQ-VAE for finding the most similar discrete latent embedding  $z_e = S(z_h)$ , and then the decoder  $D(\cdot)$  predicts the final action  $\hat{a}$  based on the discrete latent embedding  $z_e$ .

##### D. Network Architecture

In our work, we leverage both language instruction and images to predict the robot’s action. In particular, for language instructions, we use a pre-trained DistilBERT [63] to extract the features. For RGB images, we use EfficientNet B3 [64] as



Fig. 3: Real-World Experiment Setup for single-arm Franka robot. We use two external fixed-view Zed cameras. The figure in the upper right corner shows all the objects used in our experiments.




Fig. 4: Real-World Experiment Setup for bimanual UR5 robot. We use four fixed-view RealSense D435i cameras.

the visual backbone and fuse the linguistic and visual features by FiLM layers [65]. We also include proprioceptive states, we use MLP layers to project the input to the features with a fixed dimension of 512. We attached a transformer model to VQ-VAE to unify all input tokens into a single embedding. To be specific, the VQ-VAE is equipped with a transformer encoder with 4 layers, and a transformer decoder with 7 layers. The hidden dimension is set to 512. For the diffusion model which operates on the discrete latent space, we use Unet [66] to keep the input (i.e., noisy latent embedding) and output (i.e., denoised latent embedding) dimensions identical.

## V. EXPERIMENTS

After presenting Discrete Policy, we ask the following key questions about the effectiveness of our algorithm: 1) Can Discrete Policy be effectively deployed to real-world scenarios? 2) Can Discrete Policy be scaled up to multiple complex tasks? 3) Can Discrete Policy effectively distinguish different behavioral modalities across multiple tasks? In order to answer these questions, we built a real-world robotic arm environment, designed a variety of manipulation tasks that contain rich skill requirements and long-horizon tasks, and finally conducted extensive experiments as well as visualization of Discrete Policy.

### A. Experiment Setup

For real-world experiments, we collect the dataset through human demonstration. Given a target task, we randomly place the objects within a specified area and ask the human




Fig. 5: Demonstrations of the 12 tasks in single-arm robot experiments.




Fig. 6: Demonstrations of the 6 tasks in bimanual robot experiments.

demonstrators to finish the task nearly perfectly. We record the RGB stream from multiple camera views and robot states, e.g., joint position, during the whole robot control process. Our model predicts the 6D pose, including position  $(x, y, z)$  and rotation  $(roll, pitch, yaw)$ . For the single-arm Franka robot shown in Figure 3, we follow Droid [21], with two external fixed-view Zed 2 stereo cameras, one on the left and one on the right. For bimanual UR5 robotic arms shown in Figure 4, we used four RealSense D435i cameras in the environment, two of which are hand-eye cameras mounted at the wrists of each arm, and the other two are external cameras mounted above and in front of the operating table.

For the simulation benchmark, we evaluate on BiDexHands [67] and Metaworld [68] Medium level and Hard level, following the settings in MWM [69]. All experiments were trained with 20 demonstrations and evaluated with 3 seeds, and for each seed, the success rate was averaged over five different iterations.

**Task Description for Real-World Experiments.** For the single-arm Franka experiments, we designed two multi-task evaluation protocols named Multi-Task 5 (MT-5) and Multi-Task 12 (MT-12). MT-12 extends MT-5’s task range to 12 tasks, including 3 long-horizon tasks, more varied scenarios, and more complex skill requirements such as flip, press, and pull, which are shown in Figure 5. For the bimanual UR5 robotic experiments, as shown in Figure 6, we designed 6 challenging tasks that require collaboration between two robotic arms. We provide a summary of all the tasks in Table I. For each task, we evaluate 20 times with different initial object states and report the success rates.

### B. Experimental Results

**Results on Single-arm Franka Robot.** Figure 7 shows the comparisons of the success rates of all tasks on MT-5



TABLE I: The tasks summary of our real-world experiments.

#	Task	Horizon Type	# of Demo.	Avg. Traj. Length	Task Instruction
<b>Single-Arm Franka Robot</b>					
T1	PlaceTennis	Short	100	170.7	Place the tennis ball in the tube.
T2	OpenDrawer	Short	100	284.1	Pull open the drawer.
T3	OpenBox	Short	100	231.4	Open the blue plastic box.
T4	StackBlock	Short	100	229.1	Stack the pink block on the blue block.
T5	UprightMug	Short	100	195.4	Turn the tipped mug upright.
T6	CloseDrawer	Short	100	176.1	Push the drawer closed.
T7	PlaceCan	Short	100	237.9	Put the drink can in the cup holder.
T8	ArrangeFlower	Short	100	282.2	Put the bouquet in the vase.
T9	CloseMicrowave	Short	100	164.4	Close the microwave door.
T10	InsertToast	Long	100	475.9	Place the toast in the toaster.
T11	StoreToyCar	Long	100	457.2	Put the toy car in the drawer.
T12	DisposePaper	Long	100	295.2	Put the trash in the bin.
<b>Bimanual UR5 Robot</b>					
T1	TennisBallPack	Long	64	438.5	Unzip the bag, place the tennis ball into the bag.
T2	BreadTransfer	Short	100	250.0	Place bread on the plate.
T3	StackBlock	Short	150	200.0	Stack the orange cube on the purple cube.
T4	BreadDrawer	Long	100	448.9	Place the bread into the drawer.
T5	SweepPaper-1	Long	100	320.0	Sweep trash into the green dustpan.
T6	SweepPaper-2	Long	100	300.0	Sweep trash into the white dustpan.

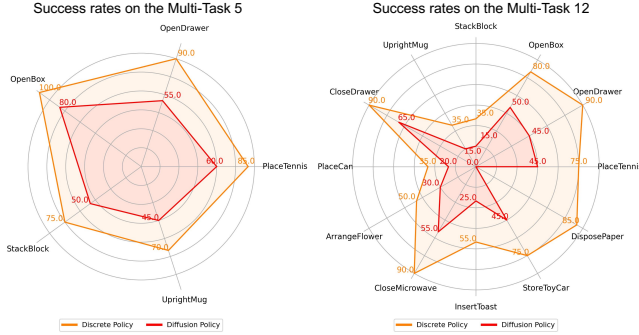


Fig. 7: The figures on the left and right show the success rates on the MT-5 and MT-12, respectively.

and MT-12, respectively. We observe that Discrete Policy achieves average success rates of 84% on MT-5 and 66.3% on MT-12, which significantly outperforms Diffusion Policy by 25% and 32.5%. Another observation is that the gap between average success rates becomes larger as more complex tasks evolve during training. We believe this is due to the increased number of tasks and the increased difficulty of the model distinguishing between individual tasks and learning the optimal policy among them. Notably, both methods suffer from a performance drop as the number of tasks increases. This is due to fixed total training iterations.

In Table II, we compare our method with other state-of-the-art approaches, ranging from pure Transformer architectures trained using behavior cloning to diffusion-based policies. Notably, both OpenVLA and Octo are pre-trained on OpenX datasets, making a direct comparison with our method—trained solely on limited robot data—unfair. Despite this, our method consistently achieves the best performance and is on par with these methods for all tasks. In terms of average success rate, our approach even surpasses OpenVLA by 15% over five tasks.

**Results on Bimanual UR5 robot.** We further conducted experiments using a bimanual UR5 robot. With an increased number of degrees of freedom, the complexity of the tasks rises significantly. In this setting, we trained all methods on six tasks, ranging from extremely long-horizon tasks like placing a tennis ball into a closed bag to more straightforward

TABLE II: Comparisons on **Multi-Task 5 (MT-5) in the single Franka robot**, with success rates reported and the best results highlighted in bold. The symbol \* denotes that methods are pretrained by 970K OpenX [70] robot data.

Method	Task1	Task1	Task3	Task4	Task5	Average
RT-1 [71]	25	30	30	10	15	22
BeT [72]	45	30	65	30	15	37
BESO [29]	40	30	55	25	15	33
MDT [24]	50	35	55	20	25	37
Octo* [73]	40	55	50	30	40	43
OpenVLA* [8]	<b>85</b>	80	90	40	50	69
MT-ACT [6]	80	80	<b>100</b>	55	45	59
Diffusion Policy [5]	60	55	80	50	45	58
Discrete Policy	<b>85</b>	<b>90</b>	<b>100</b>	<b>75</b>	<b>70</b>	<b>84</b>

TABLE III: Comparing Discrete Policy with baseline methods on **six bimanual UR5 robot tasks**.

Method / Task	T1	T2	T3	T4	T5	T6	Avg.
BeT [72]	30	10	0	30	40	20	21.7
MT-ACT [6]	<b>70</b>	40	10	70	80	60	55.0
Diffusion Policy [5]	30	35	0	45	65	50	37.5
Discrete Policy	<b>70</b>	<b>55</b>	<b>30</b>	<b>85</b>	<b>85</b>	<b>75</b>	<b>65.8</b>

tasks such as transferring a piece of bread. The experimental results are illustrated in Table III. We compare our approach to BeT [72], MT-ACT [6], and Diffusion Policy [5]. Our approach achieves an average success rate exceeding 65%, while Diffusion Policy achieves only 37.5%. Our method also outperforms both MT-ACT and BeT. These empirical results demonstrate the effectiveness of our method in multi-task scenarios, even in the challenging bimanual manipulation setting.

**Results on Simulation.** Finally, we evaluated our approach on simulation benchmarks, selecting several complex tasks, including 6 tasks from Bi-DexHands [67], 11 tasks from Metaworld Medium [68], and 6 tasks from Metaworld Hard [68], [69]. The experiments are shown in Table IV. Compared to state-of-the-art methods, our approach demonstrates superior performance across all settings, further validating its effectiveness.

### C. Visualization on Action Space

To illustrate that Discrete Policy can effectively disentangle action space between diverse tasks, we visualized the features in the latent space using the t-SNE [4] in Figures 1 and 8. We compare our approach to the Diffusion Policy. For MT-5, we observe that Discrete Policy clearly delineates different features while aggregating similar ones, indicating that Discrete Policy effectively distinguishes between multi-modal actions and skills across multiple tasks. In contrast, the features of Diffusion Policy are fragmented and overlap a lot. In the MT-12 scenario, where the number of tasks is greater, delineating the distribution of actions becomes more challenging. Nonetheless, Discrete Policy successfully distinguishes the action distributions in most tasks. Conversely, the complex action distributions confound the Diffusion Policy.

TABLE IV: Experimental results of Bi-Dexhands [67] and Metaworld [68], a simulation benchmark. The numbers in parentheses indicate the number of tasks for the simulation benchmark.

Method	Bi-Dex-Hands(6)	Metaworld-Medium(11)	Metaworld-Hard(6)
VINN [74]	20.6	5.2	2.7
BeT [72]	24.5	9.1	0.9
MT-ACT [6]	47.6	15.4	4.8
Diffusion Policy [5]	35.7	16.2	3.1
Discrete Policy	<b>54.3</b>	<b>19.6</b>	<b>7.9</b>

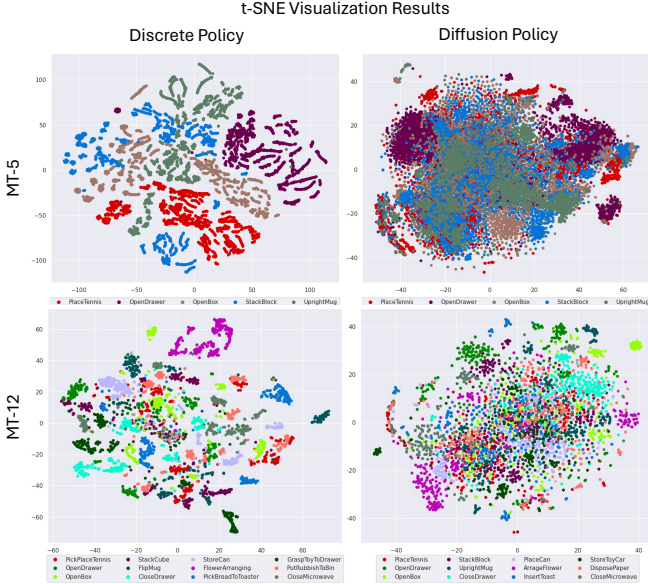


Fig. 8: The t-SNE visualization results of the latent features of Discrete Policy and Diffusion Policy on MT-5 and MT-12, respectively.

#### D. Ablation Study

We conducted extensive ablation studies to evaluate the effects of various hyperparameters, including action chunk size, codebook categories, latent embedding dimension, and the number of latent embeddings. The results of these experiments are summarized in Table V.

For action chunk size, we observed that success rates increase as the chunk size grows, particularly when increasing from 16 to 32. We also experimented with different numbers of codebook categories to assess whether increasing them from 256 to 1024 would impact model performance. Our results show that increasing the number of codebook categories consistently improves the average success rates across all tasks. This improvement is likely due to the larger capacity of the codebook, which allows it to capture a broader range of behavioral patterns.

Latent embeddings play a crucial role in our method. We evaluate both the number and dimensionality of these embeddings. Empirically, we found that increasing the dimensionality of the latent embeddings enhances performance, leading

TABLE V: Ablation study on Multi-Task 5 (MT-5) in the single Franka robot environment in terms of the success rates with the best results in bold.

Factor	Value	Task1	Task2	Task3	Task4	Task5
Chunk Size	16	70	<b>85</b>	<b>100</b>	60	50
	32	<b>85</b>	<b>85</b>	<b>100</b>	<b>75</b>	<b>60</b>
	64	60	80	90	65	50
Codebook Category	256	60	<b>90</b>	90	65	55
	512	<b>80</b>	<b>90</b>	<b>100</b>	<b>70</b>	<b>60</b>
	1024	<b>80</b>	<b>90</b>	<b>100</b>	<b>70</b>	<b>60</b>
Latent Embed. Dim.	32	50	75	85	50	55
	64	60	75	<b>95</b>	<b>70</b>	65
	128	<b>80</b>	<b>90</b>	<b>95</b>	<b>70</b>	<b>70</b>
Latent Embed. Num.	1	60	85	95	65	55
	8	<b>80</b>	<b>90</b>	<b>100</b>	<b>75</b>	<b>75</b>
	16	<b>80</b>	<b>90</b>	<b>100</b>	70	70




Fig. 9: Skill Composition of Discrete Policy.

to higher success rates. However, increasing the number of latent embeddings does not necessarily yield better results. Our observations indicate that setting the number of latent embeddings to eight provides optimal performance.

#### E. Skill Composition

Skill composition allows the model to map language queries to specific action spaces. By decomposing the action sequence into discrete latent embeddings, our Discrete Policy aims to align language inputs with this latent space. Intuitively, skills represented in the latent space can be aligned and composed to form new skills [75], [76]. To evaluate whether our method is capable of this, we combined two instructions from the training data and assessed whether the model could successfully complete the newly composed task.

As shown in Figure 9, we tasked the robot with two new objectives: 1) placing a tennis ball into a cup holder, and 2) placing a tennis ball into a drawer. Our Discrete Policy successfully interpreted these new instructions and completed the tasks, demonstrating the effectiveness of using discrete latent embeddings in action generation. This further supports the motivation behind our work, where a disentangled action space can be effectively composed into meaningful sequences.

## VI. CONCLUSION

This paper explores innovative strategies for multi-task learning in robotic systems. Our method, called Discrete Policy, learns action patterns in the latent space, enabling better disentanglement of feature representations across different skills. Through extensive simulations and real-world

experiments, our approach demonstrates superior performance in multi-task settings compared to various state-of-the-art methods. Overall, we believe that the Discrete Policy approach offers a compelling and practical new perspective on learning multi-task policies for embodied control.

## REFERENCES

- [1] K. Bousmalis, G. Vezzani, D. Rao, C. Devin, A. X. Lee, M. Bauza, T. Davchev, Y. Zhou, A. Gupta, A. Raju, *et al.*, “Robocat: A self-improving foundation agent for robotic manipulation,” *arXiv preprint arXiv:2306.11706*, 2023.
- [2] P. Florence, C. Lynch, A. Zeng, O. A. Ramirez, A. Wahid, L. Downs, A. Wong, J. Lee, I. Mordatch, and J. Tompson, “Implicit behavioral cloning,” in *Conference on Robot Learning*. PMLR, 2022, pp. 158–168.
- [3] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song, “Diffusion policy: Visuomotor policy learning via action diffusion,” *RSS*, 2023.
- [4] L. van der Maaten and G. Hinton, “Visualizing data using t-sne,” *Journal of Machine Learning Research*, vol. 9, no. 86, pp. 2579–2605, 2008. [Online]. Available: <http://jmlr.org/papers/v9/vandermaten08a.html>
- [5] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song, “Diffusion policy: Visuomotor policy learning via action diffusion,” *arXiv preprint arXiv:2303.04137*, 2023.
- [6] H. Bharadhwaj, J. Vakil, M. Sharma, A. Gupta, S. Tulsiani, and V. Kumar, “Roboagent: Generalization and efficiency in robot manipulation via semantic augmentations and action chunking,” *arXiv preprint arXiv:2309.01918*, 2023.
- [7] O. M. Team, D. Ghosh, H. Walke, K. Pertsch, K. Black, O. Mees, S. Dasari, J. Hejna, T. Kreiman, C. Xu, *et al.*, “Octo: An open-source generalist robot policy,” *arXiv preprint arXiv:2405.12213*, 2024.
- [8] M. Kim, K. Pertsch, S. Karamcheti, T. Xiao, A. Balakrishna, S. Nair, R. Rafailov, E. Foster, G. Lam, P. Sanketi, Q. Vuong, T. Kollar, B. Burchfiel, R. Tedrake, D. Sadigh, S. Levine, P. Liang, and C. Finn, “Openvla: An open-source vision-language-action model,” *arXiv preprint arXiv:2406.09246*, 2024.
- [9] J. Ho, A. Jain, and P. Abbeel, “Denoising diffusion probabilistic models,” *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020.
- [10] J. Song, C. Meng, and S. Ermon, “Denoising diffusion implicit models,” *arXiv preprint arXiv:2010.02502*, 2020.
- [11] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, “High-resolution image synthesis with latent diffusion models,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 10 684–10 695.
- [12] C. Saharia, W. Chan, S. Saxena, L. Li, J. Whang, E. L. Denton, K. Ghasemipour, R. Gontijo Lopes, B. Karagol Ayan, T. Salimans, *et al.*, “Photorealistic text-to-image diffusion models with deep language understanding,” *Advances in neural information processing systems*, vol. 35, pp. 36 479–36 494, 2022.
- [13] L. Yang, Z. Huang, F. Lei, Y. Zhong, Y. Yang, C. Fang, S. Wen, B. Zhou, and Z. Lin, “Policy representation via diffusion probability model for reinforcement learning,” *arXiv preprint arXiv:2305.13122*, 2023.
- [14] B. Mazouze, W. Talbott, M. A. Bautista, D. Hjelm, A. Toshev, and J. Susskind, “Value function estimation using conditional diffusion models for control,” *arXiv preprint arXiv:2306.07290*, 2023.
- [15] J. Brehmer, J. Bose, P. De Haan, and T. S. Cohen, “Edgi: Equivariant diffusion for planning with embodied agents,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [16] S. Venkatraman, S. Khaitan, R. T. Akella, J. Dolan, J. Schneider, and G. Berseth, “Reasoning with latent diffusion in offline reinforcement learning,” *arXiv preprint arXiv:2309.06599*, 2023.
- [17] K. Lee, S. Kim, and J. Choi, “Refining diffusion planner for reliable behavior synthesis by automatic detection of infeasible plans,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [18] S. Zhou, Y. Du, S. Zhang, M. Xu, Y. Shen, W. Xiao, D.-Y. Yeung, and C. Gan, “Adaptive online replanning with diffusion models,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [19] H. Chen, C. Lu, C. Ying, H. Su, and J. Zhu, “Offline reinforcement learning via high-fidelity generative behavior modeling,” *arXiv preprint arXiv:2209.14548*, 2022.
- [20] A. Prasad, K. Lin, J. Wu, L. Zhou, and J. Bohg, “Consistency policy: Accelerated visuomotor policies via consistency distillation,” *arXiv preprint arXiv:2405.07503*, 2024.
- [21] A. Khazatsky, K. Pertsch, S. Nair, A. Balakrishna, S. Dasari, S. Karamcheti, S. Nasiriany, M. K. Srirama, L. Y. Chen, K. Ellis, *et al.*, “Droid: A large-scale in-the-wild robot manipulation dataset,” *arXiv preprint arXiv:2403.12945*, 2024.
- [22] V. Vosylius, Y. Seo, J. Uruç, and S. James, “Render and diffuse: Aligning image and action spaces for diffusion-based behaviour cloning,” *arXiv preprint arXiv:2405.18196*, 2024.
- [23] Y. Zhu, Z. Ou, X. Mou, and J. Tang, “Retrieval-augmented embodied agents,” 2024.
- [24] M. Reuss, Ö. E. Yağmurlu, F. Wenzel, and R. Lioutikov, “Multimodal diffusion transformer: Learning versatile behavior from multimodal goals,” in *Robotics: Science and Systems*, 2024.
- [25] Y. Ze, G. Zhang, K. Zhang, C. Hu, M. Wang, and H. Xu, “3d diffusion policy: Generalizable visuomotor policy learning via simple 3d representations,” in *Proceedings of Robotics: Science and Systems (RSS)*, 2024.
- [26] Z. Xian, N. Gkanatsios, T. Gervet, and K. Fragkiadaki, “Unifying diffusion models with action detection transformers for multi-task robotic manipulation,” in *7th Annual Conference on Robot Learning*, 2023.
- [27] Y. Ze, G. Yan, Y.-H. Wu, A. Macaluso, Y. Ge, J. Ye, N. Hansen, L. E. Li, and X. Wang, “Gnfactor: Multi-task real robot learning with generalizable neural feature fields,” in *Conference on Robot Learning*. PMLR, 2023, pp. 284–301.
- [28] H. Ha, P. Florence, and S. Song, “Scaling up and distilling down: Language-guided robot skill acquisition,” in *Conference on Robot Learning*. PMLR, 2023, pp. 3766–3777.
- [29] M. Reuss, M. Li, X. Jia, and R. Lioutikov, “Goal conditioned imitation learning using score-based diffusion policies,” in *Robotics: Science and Systems*, 2023.
- [30] T. Pearce, T. Rashid, A. Kanervisto, D. Bignell, M. Sun, R. Georgescu, S. V. Macua, S. Z. Tan, I. Momennejad, K. Hofmann, *et al.*, “Imitating human behaviour with diffusion models,” *arXiv preprint arXiv:2301.10677*, 2023.
- [31] Y. Wang, Y. Zhang, M. Huo, R. Tian, X. Zhang, Y. Xie, C. Xu, P. Ji, W. Zhan, M. Ding, *et al.*, “Sparse diffusion policy: A sparse, reusable, and flexible policy for robot learning,” *arXiv preprint arXiv:2407.01531*, 2024.
- [32] X. Jia, Q. Wang, A. Donat, B. Xing, G. Li, H. Zhou, O. Celik, D. Blessing, R. Lioutikov, and G. Neumann, “Mail: Improving imitation learning with mamba,” *arXiv preprint arXiv:2406.08234*, 2024.
- [33] M. Zhu, Y. Zhu, J. Li, J. Wen, Z. Xu, N. Liu, R. Cheng, C. Shen, Y. Peng, F. Feng, *et al.*, “Scaling diffusion policy in transformer to 1 billion parameters for robotic manipulation,” *arXiv preprint arXiv:2409.14411*, 2024.
- [34] J. Wen, Y. Zhu, J. Li, M. Zhu, K. Wu, Z. Xu, R. Cheng, C. Shen, Y. Peng, F. Feng, *et al.*, “Tinyvla: Towards fast, data-efficient vision-language-action models for robotic manipulation,” *arXiv preprint arXiv:2409.12514*, 2024.
- [35] M. Psenka, A. Escontrela, P. Abbeel, and Y. Ma, “Learning a diffusion model policy from rewards via q-score matching,” *arXiv preprint arXiv:2312.11752*, 2023.
- [36] T. Huang, G. Jiang, Y. Ze, and H. Xu, “Diffusion reward: Learning rewards via conditional video diffusion,” *arXiv preprint arXiv:2312.14134*, 2023.
- [37] J. Urain, N. Funk, J. Peters, and G. Chalvatzaki, “Se (3)-diffusionfields: Learning smooth cost functions for joint grasp and motion optimization through diffusion,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 5923–5930.
- [38] A. Simeonov, A. Goyal, L. Manuelli, L. Yen-Chen, A. Sarmiento, A. Rodriguez, P. Agrawal, and D. Fox, “Shelving, stacking, hanging: Relational pose diffusion for multi-modal rearrangement,” *arXiv preprint arXiv:2307.04751*, 2023.
- [39] T. Wu, M. Wu, J. Zhang, Y. Gan, and H. Dong, “Learning score-based grasping primitive for human-assisting dexterous grasping,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [40] Z. Zhong, D. Rempe, Y. Chen, B. Ivanovic, Y. Cao, D. Xu, M. Pavone, and B. Ray, “Language-guided traffic simulation via scene-level diffusion,” in *Conference on Robot Learning*. PMLR, 2023, pp. 144–177.
- [41] W. Liu, T. Hermans, S. Chernova, and C. Paxton, “Strucdiffusion:

- Object-centric diffusion for semantic rearrangement of novel objects,” in *Workshop on Language and Robotics at CoRL 2022*, 2022.
- [42] K. Saha, V. Mandadi, J. Reddy, A. Srikanth, A. Agarwal, B. Sen, A. Singh, and M. Krishna, “Edmp: Ensemble-of-costs-guided diffusion for motion planning,” *arXiv preprint arXiv:2309.11414*, 2023.
- [43] J. Chang, H. Ryu, J. Kim, S. Yoo, J. Seo, N. Prakash, J. Choi, and R. Horowitz, “Denoising heat-inspired diffusion with insulators for collision free motion planning,” *arXiv preprint arXiv:2310.12609*, 2023.
- [44] M. Zhu, Y. Zhu, J. Li, J. Wen, Z. Xu, *et al.*, “Language-conditioned robotic manipulation with fast and slow thinking,” 2024.
- [45] J. Wen, Y. Zhu, M. Zhu, J. Li, Z. Xu, *et al.*, “Object-centric instruction augmentation for robotic manipulation,” 2024.
- [46] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu, *et al.*, “Rt-1: Robotics transformer for real-world control at scale,” *arXiv preprint arXiv:2212.06817*, 2022.
- [47] C. Chen, H.-Y. Li, X. Zhang, X. Liu, and U.-X. Tan, “Towards robotic picking of targets with background distractors using deep reinforcement learning,” in *2019 WRC Symposium on Advanced Robotics and Automation (WRC SARA)*. IEEE, 2019, pp. 166–171.
- [48] M. Shridhar, L. Manuelli, and D. Fox, “Cliport: What and where pathways for robotic manipulation,” in *Conference on robot learning*. PMLR, 2022, pp. 894–906.
- [49] Y. Ze, Y. Liu, R. Shi, J. Qin, Z. Yuan, J. Wang, and H. X. H-index, “Visual reinforcement learning with hand-informed representations for dexterous manipulation,” 2023.
- [50] T. Z. Zhao, J. Tompson, D. Driess, P. Florence, S. K. S. Ghasemipour, C. Finn, and A. Wahid, “Aloha unleashed: A simple recipe for robot dexterity,” in *8th Annual Conference on Robot Learning*.
- [51] J. Aldaco, T. Armstrong, R. Baruch, J. Bingham, S. Chan, K. Draper, D. Dwibedi, C. Finn, P. Florence, S. Goodrich, *et al.*, “Aloha 2: An enhanced low-cost hardware for bimanual teleoperation,” *arXiv preprint arXiv:2405.02292*, 2024.
- [52] X. Li, C. Mata, J. Park, K. Kahatapitiya, Y. S. Jang, J. Shang, K. Ranasinghe, R. Burgert, M. Cai, Y. J. Lee, *et al.*, “Llara: Supercharging robot learning data for vision-language policy,” *arXiv preprint arXiv:2406.20095*, 2024.
- [53] D. Niu, Y. Sharma, G. Biamby, J. Quenum, Y. Bai, B. Shi, T. Darrell, and R. Herzig, “Llarva: Vision-action instruction tuning enhances robot learning,” *arXiv preprint arXiv:2406.11815*, 2024.
- [54] E. Jang, A. Irpan, M. Khansari, D. Kappler, F. Ebert, C. Lynch, S. Levine, and C. Finn, “Bc-z: Zero-shot task generalization with robotic imitation learning,” in *Conference on Robot Learning*. PMLR, 2022, pp. 991–1002.
- [55] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, X. Chen, K. Choro-manski, T. Ding, D. Driess, A. Dubey, C. Finn, *et al.*, “Rt-2: Vision-language-action models transfer web knowledge to robotic control,” *arXiv preprint arXiv:2307.15818*, 2023.
- [56] A. Kumar, A. Singh, F. Ebert, Y. Yang, C. Finn, and S. Levine, “Pre-training for robots: Offline rl enables learning new tasks from a handful of trials,” *arXiv preprint arXiv:2210.05178*, 2022.
- [57] J. Ho, A. Jain, and P. Abbeel, “Denoising diffusion probabilistic models,” *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020.
- [58] M. Welling and Y. W. Teh, “Bayesian learning via stochastic gradient langevin dynamics,” in *Proceedings of the 28th international conference on machine learning (ICML-11)*. Citeseer, 2011, pp. 681–688.
- [59] J. Song, C. Meng, and S. Ermon, “Denoising diffusion implicit models,” *arXiv preprint arXiv:2010.02502*, 2020.
- [60] A. Van Den Oord, O. Vinyals, *et al.*, “Neural discrete representation learning,” *Advances in neural information processing systems*, vol. 30, 2017.
- [61] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” *arXiv preprint arXiv:1312.6114*, 2013.
- [62] J. Song, C. Meng, and S. Ermon, “Denoising diffusion implicit models,” in *International Conference on Learning Representations*, 2021. [Online]. Available: <https://openreview.net/forum?id=StlgiaRCHLP>
- [63] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, “Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter,” *arXiv preprint arXiv:1910.01108*, 2019.
- [64] M. Tan and Q. Le, “Efficientnet: Rethinking model scaling for convolutional neural networks,” in *International conference on machine learning*. PMLR, 2019, pp. 6105–6114.
- [65] E. Perez, F. Strub, H. De Vries, V. Dumoulin, and A. Courville, “Film: Visual reasoning with a general conditioning layer,” in *Proceedings of the AAAI conference on artificial intelligence*, vol. 32, no. 1, 2018.
- [66] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5–9, 2015, proceedings, part III 18*. Springer, 2015, pp. 234–241.
- [67] Y. Chen, T. Wu, S. Wang, X. Feng, J. Jiang, Z. Lu, S. McAleer, H. Dong, S.-C. Zhu, and Y. Yang, “Towards human-level bimanual dexterous manipulation with reinforcement learning,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 5150–5163, 2022.
- [68] T. Yu, D. Quillen, Z. He, R. Julian, K. Hausman, C. Finn, and S. Levine, “Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning,” in *Conference on robot learning*. PMLR, 2020, pp. 1094–1100.
- [69] Y. Seo, D. Hafner, H. Liu, F. Liu, S. James, K. Lee, and P. Abbeel, “Masked world models for visual control,” in *Conference on Robot Learning*. PMLR, 2023, pp. 1332–1344.
- [70] A. Padalkar, A. Pooley, A. Jain, A. Bewley, A. Herzog, A. Irpan, A. Khazatsky, A. Rai, A. Singh, A. Brohan, *et al.*, “Open x-embodiment: Robotic learning datasets and rt-x models,” *arXiv preprint arXiv:2310.08864*, 2023.
- [71] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu, *et al.*, “Rt-1: Robotics transformer for real-world control at scale,” *arXiv preprint arXiv:2212.06817*, 2022.
- [72] N. M. Shafiuallah, Z. Cui, A. A. Altanzaya, and L. Pinto, “Behavior transformers: Cloning k k modes with one stone,” *Advances in neural information processing systems*, vol. 35, pp. 22955–22968, 2022.
- [73] Octo Model Team, D. Ghosh, H. Walke, K. Pertsch, K. Black, O. Mees, S. Dasari, J. Hejna, C. Xu, J. Luo, T. Kreiman, Y. Tan, P. Sanketi, Q. Vuong, T. Xiao, D. Sadigh, C. Finn, and S. Levine, “Octo: An open-source generalist robot policy,” in *Proceedings of Robotics: Science and Systems*, Delft, Netherlands, 2024.
- [74] S. Parisi, A. Rajeswaran, S. Purushwalkam, and A. Gupta, “The unsurprising effectiveness of pre-trained vision models for control,” in *International Conference on Machine Learning*. PMLR, 2022, pp. 17359–17371.
- [75] M. L. De Guevara, J. Echevarria, Y. Li, Y. Hold-Geoffroy, C. Smith, and D. Ito, “Cross-modal latent space alignment for image to avatar translation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 520–529.
- [76] Z. Löhner and M. Moeller, “On the direct alignment of latent spaces,” in *Proceedings of UniReps: the First Workshop on Unifying Representations in Neural Models*, ser. Proceedings of Machine Learning Research, M. Fumero, E. Rodolá, C. Domine, F. Locatello, K. Dziugaite, and C. Mathilde, Eds., vol. 243. PMLR, 15 Dec 2024, pp. 158–169. [Online]. Available: <https://proceedings.mlr.press/v243/lohner24a.html>