

CARP: Visuomotor Policy Learning via Coarse-to-Fine Autoregressive Prediction

Zhefei Gong¹, Pengxiang Ding^{12*}, Shangke Lyu¹, Siteng Huang¹², Mingyang Sun¹², Wei Zhao¹,
Zhaoxin Fan³, Donglin Wang^{1†}

¹Westlake University ²Zhejiang University

³Beijing Advanced Innovation Center for Future Blockchain and Privacy Computing

[Project Webpage](#)

Abstract

In robotic visuomotor policy learning, diffusion-based models have achieved significant success in improving the accuracy of action trajectory generation compared to traditional autoregressive models. However, they suffer from inefficiency due to multiple denoising steps and limited flexibility from complex constraints. In this paper, we introduce **Coarse-to-Fine AutoRegressive Policy (CARP)**, a novel paradigm for visuomotor policy learning that redefines the autoregressive action generation process as a coarse-to-fine, next-scale approach. CARP decouples action generation into two stages: first, an action autoencoder learns multi-scale representations of the entire action sequence; then, a GPT-style transformer refines the sequence prediction through a coarse-to-fine autoregressive process. This straightforward and intuitive approach produces highly accurate and smooth actions, matching or even surpassing the performance of diffusion-based policies while maintaining efficiency on par with autoregressive policies. We conduct extensive evaluations across diverse settings, including single-task and multi-task scenarios on state-based and image-based simulation benchmarks, as well as real-world tasks. CARP achieves competitive success rates, with up to a 10% improvement, and delivers **10×** faster inference compared to state-of-the-art policies, establishing a high-performance, efficient, and flexible paradigm for action generation in robotic tasks.

1. Introduction

Policy learning from demonstrations, formulated as the supervised regression task of mapping observations to actions, has proven highly effective across various robotic tasks,

* Project lead.

† Corresponding author.




Figure 1. **Policy Comparison.** The representative performance among Behavior Transformer [52] served as an autoregressive policy, Diffusion Policy [11], and our approach in the state-based Robomimic square task experiment. CARP achieves a superior balance of performance and efficiency.

even in its simplest form. Replacing the policies with generative models, particularly those stemming from the vision community, has opened new avenues for improving performance, enabling robots to achieve the high precision required for complex tasks.

Existing approaches have explored different generative modeling techniques to address challenges in visuomotor policy learning. Autoregressive Modeling (AM) [12, 30, 52, 70] provides a straightforward and efficient out-of-the-box solution, benefitting from its scalability, flexibility, and mature exploration at lower computation requirements. However, AM’s next-token prediction paradigm often fails to




Figure 2. **Structure of Current Policies.** \hat{a} is the predicted action, a^k denotes the refining action at step k , s is the historical condition. a) Autoregressive Policy predicts the action step-by-step in the next-token paradigm. b) Diffusion Policy models the noise process used to refine the action sequence. c) CARP refines action sequence predictions autoregressively from coarse to fine granularity.

capture long-range dependencies, global structure, and temporal coherence [26], leading to poor performance, which is essential for many robotic tasks. Recently, Diffusion Modeling (DM) [11, 51, 61] has emerged as a promising alternative, bridging the precision gap in AM by modeling the gradient of the action score function to learn multimodal distributions. Nevertheless, DM requires multiple steps of sequential denoising, making it computationally prohibitive for robotic tasks, especially for robotic tasks requiring efficient real-time inference in on-board compute-constrained environments. Additionally, DDPM [19]’s rigid generative process lacks flexibility and adaptability for tasks with long-term dependencies, often leading to cumulative errors and reduced robustness over extended time spans.

Both AM and DM have their respective advantages and limitations, which are often orthogonal and difficult to balance in practical applications. In this work, we aim to resolve this trade-off by introducing a novel generative paradigm for robot visuomotor policy learning that predicts entire action sequences from a *coarse-to-fine* granularity in a *next-scale* prediction framework. This approach allows our model to achieve performance levels comparable to DM while achieving AM’s efficiency and flexibility.

Our primary contribution is the introduction of a **Coarse-to-Fine AutoRegressive Policy (CARP)**, a hybrid framework that combines AM’s efficiency with DM’s high performance to meet the demands of real-world robotic manipulation. Specifically, our contributions are as follows:

- **Multi-scale action tokenization:** We propose a multi-scale tokenization method for action sequences that captures the global structure and maintains temporal locality, effectively addressing AM’s myopic limitations.
- **Coarse-to-fine autoregressive prediction:** This mechanism refines action sequences in the latent space using Cross-Entropy loss with relaxed Markovian assumptions during iterations, achieving DM-like performance with high efficiency.
- **Comprehensive sim & real experiments:** Extensive experiments demonstrate CARP’s effectiveness in both simulated and real-world robotic manipulation tasks.

In summary, we present CARP, a novel visuomotor policy framework that synergizes the strengths of AM and DM, offering high performance, efficiency, and flexibility.

2. Background

We start by introducing the background, focusing on three key areas: problem formulation, conventional autoregressive policies, and diffusion-based policies.

2.1. Problem Formulation

Problem formulation will consider a task \mathcal{T} , where there are N expert demonstrations $\{\tau_i\}_{i=1}^N$. Each demonstration τ_i is a sequence of state-action pairs. We formulate robot imitation learning as an action sequence prediction problem [11, 46, 61], training a model to minimize the error in future actions conditioned on historical states. Specifically, imitation learning minimize the behavior cloning loss \mathcal{L}_{bc} formulated as

$$\mathcal{L}_{bc} = \mathbb{E}_{s, a \sim \mathcal{T}} \left[\sum_{t=0}^T \mathcal{L}(\pi_{\theta}(a_H | s_O), a_H) \right], \quad (1)$$

where a represents the action, s denotes the state or observation according to the specific task description, t is the current time step, H is the prediction horizon, and O is the historical horizon. For notational simplicity, we denote the action sequence $a_{t:t+H-1}$ as a_H and the state sequence $s_{t-O+1:t}$ as s_O . Here, \mathcal{L}_{bc} represents a supervised action prediction loss, such as mean squared error or negative log-likelihood, T is the length of the demonstration, and θ represents the learnable parameters of the policy network π_{θ} .

2.2. Autoregressive Policy

Autoregressive policies leverage the efficiency and flexibility of autoregressive models (GPT-style Decoders). Recent advancements, such as action chunking [52, 70], redefine this paradigm as a multi-token, one-pass prediction method (see Appendix B for further details). We refer to this approach as Autoregressive Policy (AP). In this approach, the

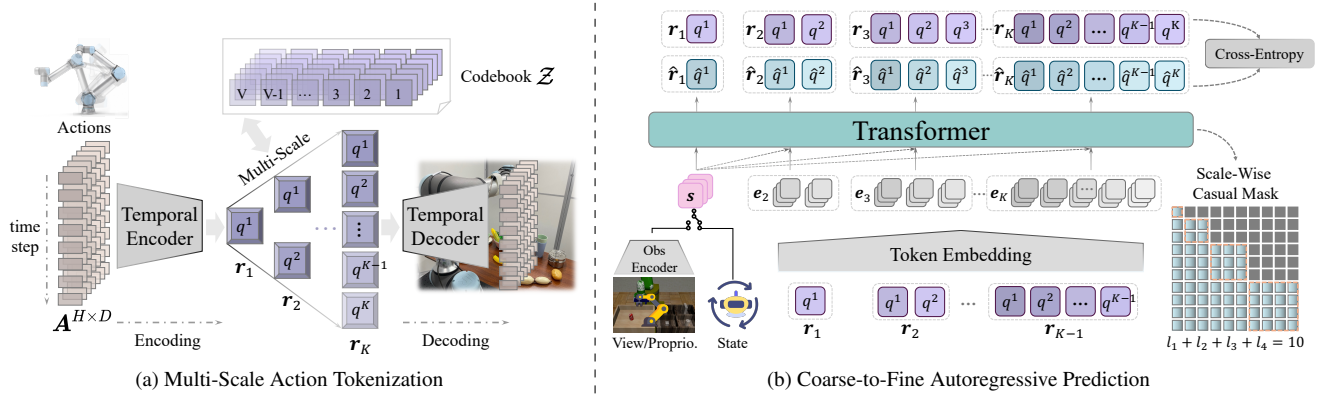


Figure 3. **Overview of the Two Stages of CARP.** a) A multi-scale action autoencoder extracts token maps $\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_K$ to represent the action sequence at different scales, trained using the standard VQVAE loss. b) The autoregressive prediction is reformulated as a *coarse-to-fine, next-scale* paradigm. The sequence is progressively refined from coarse token map \mathbf{r}_1 to finer granularity token map \mathbf{r}_K , where each \mathbf{r}_k contains l_k tokens. An attention mask ensures that each \mathbf{r}_k attends only to the preceding $\mathbf{r}_{1:k-1}$ during training. A standard Cross-Entropy loss is used for training. During inference, the token maps $\mathbf{r}_{1:K}$ are collectively decoded into continuous actions for execution.

next-token posits the probability of observing the current action \mathbf{a}_t depends solely on its previous states \mathbf{s}_O , which allows for the factorization of the likelihood of sequence with length H as:

$$p(\mathbf{a}_t, \mathbf{a}_{t+1}, \dots, \mathbf{a}_{t+H-1}) = \prod_{k=t}^{t+H-1} p(\mathbf{a}_k | \mathbf{s}_O). \quad (2)$$

However, it introduces several issues that hinder its performance. The linear, step-by-step unidirectional dependency of action prediction may overlook the global structure [26], making it challenging to capture long-range dependencies and holistic coherence or temporal locality [22] in complex scenes or sequences, which restricts their generalizability in tasks requiring bidirectional reasoning, as shown in Fig. 2a. For example, it cannot predict the former actions given the near-terminal state.

2.3. Diffusion Policy

Diffusion-based policies [11] utilize Denoising Diffusion Probabilistic Models [19] to approximate the conditional distribution $p(\mathbf{a}_H | \mathbf{s}_O)$, through modeling the noise during the denoising process from Gaussian noise to noise-free output as

$$\mathbf{a}_H^{k+1} = \alpha(\mathbf{a}_H^k - \gamma \epsilon_\theta(\mathbf{s}_O, \mathbf{a}_H^k, k) + \mathcal{N}), \quad (3)$$

where ϵ_θ is the learnable noise network, k is the current denoising step, \mathcal{N} is the Gaussian noise, α and γ are the hyper-parameters.

Diffusion-based policies show impressive performance in robotic manipulation tasks due to their action generation which gradually refines from random samples which we can abstract as a *coarse-to-fine* process. This kind of

process alike human movement or natural thinking flows in line with human intuition. However, they suffer from poor convergence due to their design. To address this, multiple denoising steps are required, constrained by the Markovian assumption, where \mathbf{a}_H^{k+1} depends only on the previous step \mathbf{a}_H^k (as shown in Fig. 2b), which leads to significant runtime inefficiency. Moreover, they lack the ability to leverage generation context effectively, hindering scalability to complex scenes [16].

3. Method

To address the limitations of existing methods, we propose **Coarse-to-Fine AutoRegressive Policy (CARP)**, a novel visuomotor policy framework that combines the high performance of recent diffusion-based policies with the sample efficiency and flexibility of traditional autoregressive policies.

CARP achieves these advantages through a redesigned autoregressive modeling strategy. Specifically, we shift from conventional *next-token* prediction to a *coarse-to-fine, next-scale* prediction approach, using *multi-scale* action representations. CARP’s training is divided into two stages: *multi-scale* action tokenization and *coarse-to-fine* autoregressive prediction, as shown in Fig. 3.

In this section, we first explain the construction of *multi-scale* action token maps, followed by the *coarse-to-fine* autoregressive prediction approach for action generation. Key implementation details are provided at the end.

3.1. Multi-Scale Action Tokenization

Instead of focusing on individual action steps, we extract representations at multiple scales across the entire action sequence. We propose a novel *multi-scale* action quantiza-

tion autoencoder that encodes a sequence of actions into K discrete token maps, $\mathbf{R} = (\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_K)$, which are used for both training and inference. Our approach builds on the VQVAE architecture [60], incorporating a modified *multi-scale* quantization layer [58] to enable hierarchical encoding.

Encoder and Decoder. As illustrated in Fig. 3a, the actions are first organized into an action sequence $\mathbf{A}^{H \times D}$, where H denotes the prediction horizon and D represents the dimensionality of the action \mathbf{a} . Given that each dimension in the action space is orthogonal to the others, and that there exists a natural temporal dependency within an action sequence, we employ a 1D temporal convolutional network (1D-CNN) along the time dimension, as used in [22], for both the encoder $\mathcal{E}(\cdot)$ and decoder $\mathcal{D}(\cdot)$. Let $\mathbf{F} = \mathcal{E}(\mathbf{A}) \in \mathbb{R}^{L \times C}$, where L denotes the compressed length of the temporal dimension with $L \leq H$, and C represents the dimensionality of the feature map \mathbf{F} .

Quantization. We introduce a quantizer with a learnable codebook $\mathcal{Z} \in \mathbb{R}^{V \times C}$, containing V code vectors, each of dimension C . The quantization process (lines 4-11 in Algorithm 1) generates the action sequence representation iteratively across multiple scales. At scale k , it produces action token map \mathbf{r}_k to represent the sequence, which consists of l_k tokens q . In our implementation, we set $l_k = k$. The function $\text{Lookup}(\mathcal{Z}, v)$ retrieves the v -th code vector from \mathcal{Z} . The quantization function is defined by $\mathbf{r} = \mathcal{Q}(\mathbf{F})$ as following:

$$q = \arg \min_{v \in [V]} \text{dist}(\text{Lookup}(\mathcal{Z}, v), \mathbf{f}) \in [V], \quad (4)$$

where token q represents the nearest vector in \mathcal{Z} for a given feature vector $\mathbf{f} \in \mathbb{R}^{1 \times C}$ in the feature map \mathbf{F} , based on a distance function $\text{dist}(\cdot)$.

Specifically, we adopt a residual-style design [29, 58] for feature maps \mathbf{F} and $\hat{\mathbf{F}}$, as shown in lines 4-11 of Algorithm 1. This design ensures that each finer-scale representation \mathbf{r}_k depends only on its coarser-scale predecessors $(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_{k-1})$, facilitating a *multi-scale* representation. A shared codebook \mathcal{Z} is used across all scales, ensuring that *multi-scale* token maps $\mathbf{R} = (\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_K)$, where K is the number of scales, are drawn from a consistent vocabulary $[V]$. To preserve information during up-sampling, we employ K additional 1D convolutional layers $\{\phi_k\}_{k=1}^K$ [58], as illustrated in lines 9-10 of Algorithm 1.

Loss. The final approximation $\hat{\mathbf{F}}$ of the original feature map \mathbf{F} is combined residually with the *multi-scale* representation derived from the codebook \mathcal{Z} , based on each \mathbf{r} . The reconstructed action sequence is then obtained as $\hat{\mathbf{A}} = \mathcal{D}(\hat{\mathbf{F}})$. To train the quantized autoencoder, a typical VQVAE loss \mathcal{L} is minimized as following

$$\mathcal{L} = \underbrace{\|\mathbf{A} - \hat{\mathbf{A}}\|_2}_{\mathcal{L}_{\text{recon}}} + \underbrace{\|\text{sg}(\mathbf{F}) - \hat{\mathbf{F}}\|_2}_{\mathcal{L}_{\text{quant}}} + \underbrace{\|\mathbf{F} - \text{sg}(\hat{\mathbf{F}})\|_2}_{\mathcal{L}_{\text{commit}}}, \quad (5)$$

where $\mathcal{L}_{\text{recon}}$ minimizes the difference between the original action sequence \mathbf{A} and the reconstruction $\hat{\mathbf{A}}$. $\text{sg}(\cdot)$ denotes stop-gradient. $\mathcal{L}_{\text{quant}}$ aligns the quantized feature map $\hat{\mathbf{F}}$ with the original \mathbf{F} , and $\mathcal{L}_{\text{commit}}$ encourages the encoder to commit to codebook entries, preventing codebook collapse. For $\mathcal{L}_{\text{quant}}$ and $\mathcal{L}_{\text{commit}}$, we calculate every residual calculating moment of each scale \mathbf{r}_k as in Algorithm 1. After training, the autoencoder $\{\mathcal{E}, \mathcal{Q}, \mathcal{D}\}$ tokenizes actions for subsequent *coarse-to-fine* autoregressive modeling.

Discussion. The tokenization strategy described above allows the *multi-scale* tokens \mathbf{R} , extracted from the action sequence $\mathbf{A} \in \mathbb{R}^{H \times D}$ via temporal 1D convolutions, to inherently preserve temporal locality [22]. Additionally, the hierarchical extraction captures the global structure, enabling the model to treat the action sequence as a unified entity.

Unlike traditional autoregressive policies that predict each action token independently (see Eq. (2) and Fig. 2a), CARP leverages dual capabilities to capture both local temporal dependencies and global structure across the entire action sequence. This approach produces smoother transitions and more stable action sequences. Through *multi-scale* encoding, CARP overcomes the short-sighted limitations of conventional autoregressive models, yielding more robust, coherent, and precise behaviors over extended time horizons.

Algorithm 1 Multi-Scale Action VQVAE

- 1: **Inputs:** Action sequence \mathbf{A}
 - 2: **Hyperparameters:** Number of scales K , length of each scale $(l_k)_{k=1}^K$, length of feature map's temporal dimension L
 - 3: **Initialize:** $\mathbf{F} \leftarrow \mathcal{E}(\mathbf{A})$, $\hat{\mathbf{F}} \leftarrow 0$, $\mathbf{R} \leftarrow \square$
 - 4: **for** $k = 1$ **to** K **do**
 - 5: $\mathbf{r}_k \leftarrow \mathcal{Q}(\text{Interpolate}(\mathbf{F}, l_k))$
 - 6: $\mathbf{R} \leftarrow \mathbf{R} \cup \{\mathbf{r}_k\}$
 - 7: $\mathcal{Z}_k \leftarrow \text{Lookup}(\mathcal{Z}, \mathbf{r}_k)$
 - 8: $\mathbf{Z}_k \leftarrow \text{Interpolate}(\mathcal{Z}_k, L)$
 - 9: $\mathbf{F} \leftarrow \mathbf{F} - \phi_k(\mathbf{Z}_k)$
 - 10: $\hat{\mathbf{F}} \leftarrow \hat{\mathbf{F}} + \phi_k(\mathbf{Z}_k)$
 - 11: **end for**
 - 12: $\hat{\mathbf{A}} \leftarrow \mathcal{D}(\hat{\mathbf{F}})$
 - 13: **Return:** Multi-scale token maps \mathbf{R} , reconstructed action sequence $\hat{\mathbf{A}}$
-

3.2. Coarse-to-Fine Autoregressive Prediction

Using *multi-scale* action sequence representations, we shift from traditional *next-token* prediction to a *next-scale* prediction approach, progressing from coarse to fine granularity.

Prediction. As we receive the *multi-scale* representation tokens $(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_K)$, each \mathbf{r} represents the same action sequence in different scales. The autoregressive likelihood

can be formulated as

$$p(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_K) = \prod_{k=1}^K p(\mathbf{r}_k \mid \mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_{k-1}; \mathbf{s}_O), \quad (6)$$

where each autoregressive unit $\mathbf{r}_k \in [V]^{l_k}$ is the token map at scale k containing l_k tokens q . The prefix sequence $(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_{k-1})$ is served as the condition for \mathbf{r}_k , accompanying with the historical state sequence \mathbf{s}_O . This kind of *next-scale* prediction methodology is what we define as *coarse-to-fine* autoregressive prediction.

Due to the residual-style quantization, during autoregressive prediction, we first embed the previous scale token map \mathbf{r}_{k-1} and then reconstruct the next scale feature map \mathbf{e}_k , as shown in Fig. 3b. This feature map \mathbf{e}_k is then used as input for predicting the next scale token map \mathbf{r}_k . The token maps $\mathbf{r}_{1:K}$ are decoded by the *multi-scale* autoencoder into continuous actions, following the residual-style process.

Loss. During the k -th autoregressive step, all distributions over the l_k tokens in \mathbf{r}_k will be generated in parallel, with the *coarse-to-fine* dependency ensured by a block-wise causal attention mask [58]. To optimize the autoregressive model, we utilize the standard Cross-Entropy loss to capture the difference between the predicted token map $\hat{\mathbf{r}}$ and the token map \mathbf{r} from the ground truth action sequence:

$$\mathcal{L}_{\text{Cross-Entropy}} = \sum_{k=1}^K \sum_{i=1}^{l_k} \left[- \sum_{v=1}^V \mathbf{r}_k^{i,v} \log \hat{\mathbf{r}}_k^{i,v} \right], \quad (7)$$

where l_k is the length of each scale, V is the action vocabulary size, and K is the number of scales.

Discussion. CARP models the entire trajectory holistically, progressively refining actions from high-level intentions to fine-grained details (see Fig. 2c). This approach aligns more closely with natural human behavior, where movement is guided by overarching intentions rather than step-by-step planning.

Instead of the commonly used MSE loss [5, 59, 70], CARP employs Cross-Entropy loss (Eq. (7)) in behavior cloning, as MSE tends to enforce a unimodal distribution that can be detrimental to manipulation tasks [52]. CARP’s iterative refinement process resembles the denoising steps in diffusion models to achieve high accuracy.

Furthermore, our approach models actions directly rather than modeling noise, enabling faster convergence in low-dimensional manifolds [36, 64]. By operating in the latent space with action sequence tokens, CARP mitigates trajectory anomalies that can disrupt prediction, unlike methods that work on raw actions [11]. This latent-space representation allows CARP to focus on the essential components of actions, resulting in smoother and more efficient predictions. In contrast, traditional diffusion models rely on Markovian processes, which compress all information

into progressively noisier inputs from previous levels, often hindering efficient learning and requiring more inference steps [16] (see Eq. (3) and Fig. 2b). CARP relaxes this constraint by allowing each scale \mathbf{r}_k to depend on all prior scales $\mathbf{r}_{1:k-1}$ instead of the only previous scale \mathbf{r}_{k-1} . This structure enables CARP to generate high-quality trajectories with significantly fewer steps.

3.3. Implementation Details

Our primary focus is on the CARP algorithm, and we adopt a simple model architecture design.

Tokenization. Due to the disentangling of the action prediction, imprecise *coarse-to-fine* actions representation can decrease the upper limit of model performance. Considering the discontinuity of the most representation for rotation space like Euler angle or quaternion will increase the unstable training process, we utilize rotation6d [71] to get stable *multi-scale* tokens. And for the distance function $\text{dist}(\cdot)$, we use cosine similarity rather than Euclidean distance which is the cause of unstable training based on our observations. In practice, due to the orthogonality between dimensions, we use a separate VQVAE [58] for each dimension of the action space, to gain stable training. All convolutions we used in CARP are 1D to capture the time-dimension features.

Autoregressive. We adopt the architecture of standard decoder-only transformers akin to GPT-2 [9, 58]. The state \mathbf{s}_O is used to generate the initial coarse-scale token map and then is utilized as adaptive normalization [40] for the subsequent predictions. During training, we observe that incorporating Exponential Moving Average (EMA) [18] enhances both training stability and performance, yielding a 4-5% improvement, consistent with findings in [11]. During the inference, kv-caching can be used and no mask is needed, as shown in Fig. 10.

4. Experiment

In this section, we evaluate CARP on diverse robotics tasks, including state-based and image-based benchmarks in single-task and multi-task settings. CARP’s performance is assessed based on task success rates, inference speed, and model parameter scale. Additionally, we validate CARP’s practical effectiveness by deploying it on real-world tasks using a UR5e robotic arm, comparing its performance against state-of-the-art diffusion-based policies.

Our experiments are structured to address the following key research questions:

- **RQ1:** Can CARP match the accuracy and robustness of current state-of-the-art diffusion-based policies?
- **RQ2:** Does CARP maintain high inference efficiency and achieve fast convergence?
- **RQ3:** Does CARP leverage the flexibility benefits of a GPT-style architecture?




Figure 4. **Single-Task Simulation Setup.** We evaluate three tasks from the Robomimic [38] benchmark—Lift, Can, and Square—ordered by increasing difficulty, along with a Kitchen task [17] on the far left.

4.1. Evaluation on Simulation Benchmark

We first evaluate CARP on a set of simulated tasks commonly used to benchmark diffusion-based models, assessing its ability to match their performance while significantly improving computational efficiency.

Policy	Lift	Can	Square	Params/M	Speed/s
BET [52]	0.96	0.88	0.54	0.27	2.12
DP-C [11]	1.00	0.94	0.94	65.88	35.21
DP-T [11]	1.00	1.00	0.88	8.97	37.83
CARP (Ours)	1.00	1.00	0.98	0.65	3.07

Table 1. **State-Based Simulation Results (State Policy).** We report the average success rate of the top 3 checkpoints, along with model parameter scales and inference time for generating 400 actions. CARP significantly outperforms BET and achieves competitive performance with state-of-the-art diffusion models, while also surpassing DP in terms of model size and inference speed.

Policy	Lift	Can	Square	Params/M	Speed/s
IBC [15]	0.72	0.02	0.00	3.44	32.35
DP-C [11]	1.00	0.97	0.92	255.61	47.37
DP-T [11]	1.00	0.98	0.86	9.01	45.12
CARP (Ours)	1.00	0.98	0.88	7.58	4.83

Table 2. **Image-Based Simulation Results (Visual Policy).** Results show that CARP consistently balances high performance and high efficiency. We highlight our results in light-blue.

Experimental Setup. We use the Robomimic [38] benchmark suite, which is widely adopted for evaluating diffusion-based models [11, 42, 61]. We evaluate CARP on both state-based and image-based datasets collected from expert demonstrations, with each task containing 200 demonstrations. The selected tasks are standard in single-task evaluations, as shown in Fig. 4 and Fig. 12. For evaluating the ability to learn multiple long-horizon tasks, we utilize the Franka Kitchen environment [17] which contains 7 objects for interaction and comes with a human demonstration dataset of 566 demonstrations, each completing 4 tasks in arbitrary order (see Fig. 14 for details).

Baselines. We compare CARP with previous autoregressive policies as well as recent diffusion policies. Behavior Transformer (BET) [52] is an autoregressive policy with action discretization and correction mechanisms, similar to offset-based prediction. Implicit Behavior Cloning (IBC) [15] utilizes energy-based models for supervised robotic behavior learning. Diffusion Policy (DP) [11] combines a denoising process with action prediction, implemented in two variants: CNN-based (DP-C) and Transformer-based (DP-T).

Policy	p1	p2	p3	p4	Params	Speed
BET [52]	0.96	0.84	0.6	0.20	0.30	1.95
DP-C [11]	1.00	1.00	1.00	0.96	66.94	56.14
DP-T [11]	1.00	0.99	0.98	0.96	80.42	56.32
CARP (Ours)	1.00	1.00	0.98	0.98	3.88	2.01

Table 3. **Multi-Stage Task Results (State Policy).** In the Kitchen, p_x represents the frequency of interactions with x or more objects. CARP outperforms BET, especially on challenging metrics like p_4 , and achieves competitive performance compared to DP, with fewer parameters and faster inference speed.

Metrics. For each task, we evaluate the policies by running 50 trials with random initializations to compute the success rate and report the average success rate across the top 3 checkpoints. For the Kitchen task, success rates are reported with increasing difficulty, determined by the number of interactive objects (from p_1 to p_4). Inference speed is measured on an A100 GPU by averaging the time required for 400 action predictions (280 for the Kitchen task) over 5 runs to ensure robustness. We also record the parameter count for each model using the same PyTorch implementation interface.

Implementation Details. For baseline models, we follow the same implementation and training configurations provided by [11]. In state-based experiments, we set the observation horizon $O = 2$ and the prediction horizon $H = 16$ across all models. For image-based experiments, we set $O = 1$ and $H = 16$ for better transferability to real-world scenarios. As per the benchmark, only the first 8 actions in the prediction horizon are executed starting from the current step (see Appendix C). For CARP, we first train an action VQVAE model (see Sec. 3.1) following [29], using $V = 512$, $C = 8$, a batch size of 256, and 300 epochs per task. Given a horizon $H = 16$, we design multi-scale representations with scales of 1, 2, 3 and 4 to capture *coarse-to-fine* information across the action sequence. We then train an autoregressive GPT-2 style, decoder-only transformer (see Sec. 3.2), based on [58], using the same training settings as the benchmark, with a batch size of 256 for state-based experiments (4000 epochs) and a batch size of 64 for image-based experiments (3000 epochs).

Results. As demonstrated in Tables 1, 2, and 3, CARP




Figure 5. **Visualization of the Trajectory and Refining Process.** The left panel shows the final predicted trajectories for each task, with CARP producing smoother and more consistent paths than Diffusion Policy (DP). The right panel visualizes intermediate trajectories during the refinement process for CARP (top-right) and DP (bottom-right). DP displays considerable redundancy, resulting in slower processing and unstable training, as illustrated by 6 selected steps among 100 denoising steps. In contrast, CARP achieves efficient trajectory refinement across all 4 scales, with each step contributing meaningful updates.

consistently achieves comparable performance to state-of-the-art diffusion models across both state-based and image-based tasks, answering **RQ1**. Notably, CARP significantly outperforms diffusion policies in terms of inference speed, being approximately 10 times faster, with only 1-5% of the parameters required by diffusion models, thereby strongly supporting **RQ2** regarding CARP’s efficiency.

Analysis. To further examine CARP’s stability and efficiency, we visualize spatial trajectories along the xyz axes for the Can and Square tasks in Fig. 5. In each task’s left panel, CARP consistently reaches specific regions (light grey area) to perform task-related actions, such as positioning for object grasping or placement. Compared to diffusion-based models, CARP’s trajectories are smoother and more consistent, underscoring its stability and accuracy (supporting **RQ1**). The right panels of Fig. 5 compare CARP’s *coarse-to-fine* action predictions with 6 selected denoising steps of the diffusion model. CARP achieves accurate predictions within just 4 *coarse-to-fine* steps, whereas the diffusion model requires numerous denoising iterations, with many early steps introducing redundant computations. This analysis further supports **RQ2** and highlights CARP’s efficiency and fast convergence advantage over diffusion policies in generating accurate actions with fewer refinement steps, as detailed in Sec. 3.2.

4.2. Evaluation on Multi-Task Benchmark

We evaluate CARP on the MimicGen [39] multi-task simulation benchmark, widely used by the state-of-the-art Sparse Diffusion Policy (SDP) [61], to demonstrate CARP’s flexibility, a result of its GPT-style autoregressive design.




Figure 6. **Multi-Task Simulation Setup.** We evaluate eight tasks from the MimicGen [39] benchmark: Coffee, Hammer Cleanup, Mug Cleanup, Nut Assembly, Square, Stack, Stack Three, and Threading, listed left-to-right and top-to-bottom.

Experimental Setup. MimicGen extends benchmark Robomimic [38] by including 1K–10K human demonstrations per task, with diverse initial state distributions for enhanced generalization in multi-task evaluation. It consists of 12 robosuite [72] tasks powered by MuJoCo and 4 high-precision tasks from Isaac Gym Factory. We select 8 robosuite tasks—Coffee, Hammer, Mug, Nut, Square, Stack, Stack Three, and Threading—for evaluation, each with 1K training trajectories, following the settings in SDP [61]. Visualization is shown in Fig. 6 and Fig. 13.

Baselines. We compare CARP with two baselines: Task-Conditioned Diffusion (TCD) [2, 34]: A basic diffusion-based multi-task policy and Sparse Diffusion Policy (SDP) [61]: A transformer-based diffusion policy that leverages Mixture of Experts (MoE) [53]. Both baselines

Policy	Prms/M	Speed/s	Coffee	Hammer	Mug	Nut	Square	Stack	Stack three	Threading	Avg.
TCD [34]	156.11	107.15	0.77	0.92	0.53	0.44	0.63	0.95	0.62	0.56	0.68
SDP [61]	159.85	112.39	0.82	1.00	0.62	0.54	0.82	0.96	0.80	0.70	0.78
CARP (Ours)	16.08	6.92	0.86	0.98	0.74	0.78	0.90	1.00	0.82	0.70	0.85

Table 4. **Multi-Task Simulation Results (Visual Policy)**. Success rates are averaged across the top three checkpoints for each task, as well as the overall average across all tasks. We also report parameter count and inference time for generating 400 actions. CARP outperforms diffusion-based policies by 9%-25% in average performance, with significantly fewer parameters and over 10× faster inference.

are trained using visual inputs.

Metrics. Success rates are reported for each task as the average of the best three checkpoints. For Nut Assembly, partial success (e.g., placing one block inside the cylinder) is assigned a score of 0.5, while full success is scored as 1. For all other tasks, a score of 1 is awarded only when strict success criteria are fully satisfied. Additionally, we calculate the average success rate across all tasks. To evaluate model efficiency, we report both the parameter scale calculated by the Pytorch interface and the inference time required to predict 400 actions on a single A100 GPU.

Implementation Details. For CARP, we extend the single-task implementation by incorporating a task embedding as an additional condition, alongside the observation sequence s . We also use a moderately deeper decoder-only transformer in GPT-2 style. CARP is trained with a batch size of 512 for 200 epochs on an A100 GPU. Baseline models follow the same training settings as SDP [61]. This minimal modification enables CARP to adapt to multi-task learning seamlessly.

Results. As shown in Tab. 4, CARP achieves up to a 25% average improvement in success rates compared to state-of-the-art diffusion-based policies, highlighting its strong performance. With minimal modification, CARP seamlessly transitions from single-task to multi-task learning, further demonstrating its flexibility, a benefit of its GPT-style architecture. Additionally, as shown in Tab. 4, CARP achieves over 10× faster inference speed and uses only 10% of the parameters compared to SDP. Leveraging its GPT-style autoregressive design [7, 47], CARP can seamlessly adapt to multi-task settings without complex design modifications [61], demonstrating the flexibility of this architecture, which strongly supports **RQ3**. These results strongly demonstrate CARP’s flexibility, solidifying its role as a high-performance and high-efficiency approach for visuomotor robotic policies.

4.3. Evaluation on Real-World

In this section, we evaluate our approach, CARP, on real-world tasks under compute-constrained conditions, comparing its performance and efficiency against baseline methods.

Experimental Setup. To validate CARP’s real-world applicability, we design two manipulation tasks:




Figure 7. **Real-World Setup.** The left panel shows the environment used for the experiment and demonstration collection. The right panel shows the trajectory from the Cup and Bowl datasets.

1. *Cup*: The robot must locate a cup on the table, pick it up, move to the right area of the table, and put it down steadily (top-right, Fig. 7).
2. *Bowl*: The robot needs to identify a smaller bowl and a larger pot on the table, pick up the bowl, and place it inside the pot (bottom-right, Fig. 7).

We use a UR5e robotic arm with a Robotiq-2f-85 gripper, equipped with two RGB cameras: one mounted on the wrist and one in a third-person perspective (left panel, Fig. 7). The robot is controlled through 6D end-effector positioning, with inverse kinematics for joint angle calculation. For teleoperation, we collected 50 human demonstration trajectories for each task using a 3D Connexion space mouse.

Baselines. As a baseline, we reproduce the CNN-based image-based diffusion policy [11], adapting the model’s input size to accommodate our observational setup.

Metrics. For each trained policy, we report the average success rate across 20 trials per task, with the initial positions randomized. We also measure inference speed on an NVIDIA GeForce RTX 2060 GPU, reporting action prediction frequency in Hertz.

Implement Details. For both models, the input consists of current visual observations from the wrist and scene cameras (resolution: 120×160), as well as proprioceptive data from the robotic arm. We execute 8 predicted actions out of a horizon of 16 predictions. We train the diffusion policy for 3000 epochs with a batch size of 64. For CARP, we use

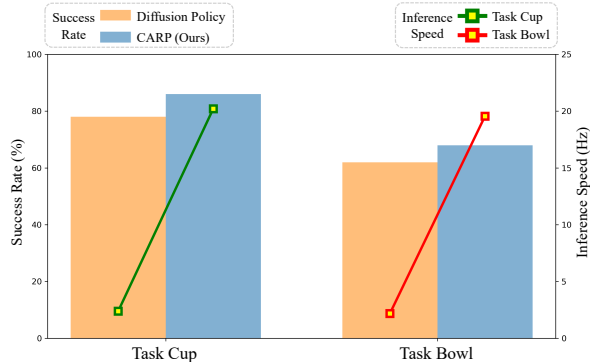


Figure 8. **Real-World Results (Visual Policy)**. We report the average success rate across 20 trials and the inference speed as action prediction frequency. CARP achieves competitive success rates with significantly faster inference compared to DP.

the same visual policy structure as in the simulation tasks, training the action VQVAE for 300 epochs with a batch size of 256, and the decoder-only transformer for 3000 epochs with a batch size of 64.

Results. As shown in Fig. 8 and Fig. 9, CARP achieves comparable or superior performance, with up to a 10% improvement in success rate over the diffusion policy across all real-world tasks, supporting **RQ1**. Additionally, CARP achieves approximately $8\times$ faster inference than the baseline on limited computational resources, demonstrating its suitability for real-time robotic applications, thus supporting **RQ2** (see Fig. 15 for further visualization).




Figure 9. **Visualization of CAPR on Real-World Tasks**. CARP generates smooth and successful trajectories for the Cup and Bowl tasks, with temporal progression from left to right.

5. Related Work

5.1. Visual Generation

Advancements in generative models for visual generation have significantly influenced the robotics community. Autoregressive models generate images in a raster-scan fashion using discrete tokens from image tokenizers [14, 14, 20, 60]. GPT-2-style transformers [4, 29, 48, 63] demonstrate

strong performance by generating tokens sequentially. Recent work has scaled these models to billions of parameters, achieving impressive text-to-image synthesis results [35, 69] and robotic action generation [13, 54]. VAR [58] introduces a new next-scale autoregressive paradigm that shifts image representation from patches to scales. This framework [58], has been applied across tasks [16, 32, 33, 37, 43, 66, 67]. Studies [56, 58] show that autoregressive models can surpass diffusion models in achieving state-of-the-art performance, serving as a key inspiration for our work.

5.2. Visuomotor Policy Learning

Behavior cloning [8] has proven effective, especially in autonomous driving [41] and manipulation [68], offering a simpler alternative to complex reinforcement learning. Explicit policies directly map states or observations to actions [68], enabling efficient inference. However, they struggle with complex tasks. Techniques like action space discretization [65] and Mixture Density Networks (MDNs) [12, 52] have been proposed, though they face issues with exponential action space growth and hyperparameter sensitivity. Implicit policies, often based on Energy-Based Models (EBMs) [15, 23], offer flexibility but are challenging to train due to optimization instabilities. Diffusion models have been applied to robotic policy learning [2, 11, 50], showing their effectiveness for decision-making tasks. However, their multi-step denoising process is computationally expensive. Recent efforts focus on improving generalization [62], adapting to 3D environments [61, 64], and enhancing modularity via Mixture of Experts (MoE) [61]. Consistency models [36, 55] accelerate inference but often compromise action prediction accuracy, along with inflexible model design.

6. Conclusion

In this work, we introduce Coarse-to-Fine Autoregressive Policy (CARP), a novel paradigm for robotic visuomotor policy learning that combines the efficiency of autoregressive modeling (AM) with the high performance of diffusion modeling (DM). CARP incorporates: 1) *multi-scale* action tokenization to capture global structure and temporal locality, addressing AM’s limitations in long-term dependency; 2) *coarse-to-fine* autoregressive prediction that refines actions from high-level intentions to detailed execution, achieving DM-like performance with AM-level efficiency through latent space prediction and relaxed Markovian constraints. The comprehensive evaluations, from simulation to real-world, demonstrate CARP’s effectiveness in balancing high performance, efficiency, and flexibility.

We hope this work will inspire further exploration into next-generation policy learning by leveraging GPT-style autoregressive models, advancing a more unified perspective on current generative modeling techniques (see Appx.A).

References

- [1] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, et al. Do as i can, not as i say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022. 13
- [2] Anurag Ajay, Yilun Du, Abhi Gupta, Joshua Tenenbaum, Tommi Jaakkola, and Pulkit Agrawal. Is conditional generative modeling all you need for decision-making? *arXiv preprint arXiv:2211.15657*, 2022. 7, 9
- [3] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35:23716–23736, 2022. 13
- [4] Dosovitskiy Alexey. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 9
- [5] Mariusz Bojarski. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016. 5
- [6] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. *arXiv preprint arXiv:2307.15818*, 2023. 13
- [7] Tom B Brown. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020. 8, 13
- [8] Harish chaandar Ravichandar, Athanasios S. Polydoros, Sonia Chernova, and Aude Billard. Recent advances in robot learning from demonstration. *Annu. Rev. Control. Robotics Auton. Syst.*, 3:297–330, 2020. 9
- [9] Junsong Chen, Jincheng Yu, Chongjian Ge, Lewei Yao, Enze Xie, Yue Wu, Zhongdao Wang, James Kwok, Ping Luo, Huchuan Lu, et al. Pixart-alpha: Fast training of diffusion transformer for photorealistic text-to-image synthesis. *arXiv preprint arXiv:2310.00426*, 2023. 5
- [10] Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Misha Laskin, Pieter Abbeel, Aravind Srivas, and Igor Mordatch. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34:15084–15097, 2021. 13
- [11] Cheng Chi, Siyuan Feng, Yilun Du, Zhenjia Xu, Eric Cousineau, Benjamin Burchfiel, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. In *Proceedings of Robotics: Science and Systems (RSS)*, 2023. 1, 2, 3, 5, 6, 8, 9, 13, 14
- [12] Zichen Jeff Cui, Yibin Wang, Nur Muhammad Mahi Shafiq, and Lerrel Pinto. From play to policy: Conditional behavior generation from uncurated robot data. *arXiv preprint arXiv:2210.10047*, 2022. 1, 9, 13
- [13] Pengxiang Ding, Han Zhao, Wenjie Zhang, Wenxuan Song, Min Zhang, Siteng Huang, Ningxi Yang, and Donglin Wang. Quar-vla: Vision-language-action model for quadruped robots. In *European Conference on Computer Vision*, pages 352–367. Springer, 2025. 9
- [14] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12873–12883, 2021. 9
- [15] Pete Florence, Corey Lynch, Andy Zeng, Oscar A Ramirez, Ayzaan Wahid, Laura Downs, Adrian Wong, Johnny Lee, Igor Mordatch, and Jonathan Tompson. Implicit behavioral cloning. In *Proceedings of the 5th Conference on Robot Learning*, pages 158–168. PMLR, 2022. 6, 9
- [16] Jiatao Gu, Yuyang Wang, Yizhe Zhang, Qihang Zhang, Dinghui Zhang, Navdeep Jaitly, Josh Susskind, and Shuangfei Zhai. Dart: Denoising autoregressive transformer for scalable text-to-image generation. *arXiv preprint arXiv:2410.08159*, 2024. 3, 5, 9
- [17] Abhishek Gupta, Vikash Kumar, Corey Lynch, Sergey Levine, and Karol Hausman. Relay policy learning: Solving long-horizon tasks via imitation and reinforcement learning. *arXiv preprint arXiv:1910.11956*, 2019. 6
- [18] David Haynes, Steven Corns, and Ganesh Kumar Venayagamoorthy. An exponential moving average algorithm. In *2012 IEEE Congress on Evolutionary Computation*, pages 1–8. IEEE, 2012. 5
- [19] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *arXiv preprint arxiv:2006.11239*, 2020. 2, 3, 13
- [20] Mengqi Huang, Zhendong Mao, Zhuowei Chen, and Yongdong Zhang. Towards accurate image coding: Improved autoregressive image generation with dynamic vector quantization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22596–22605, 2023. 9
- [21] Michael Janner, Qiyang Li, and Sergey Levine. Offline reinforcement learning as one big sequence modeling problem. *Advances in neural information processing systems*, 34: 1273–1286, 2021. 13
- [22] Michael Janner, Yilun Du, Joshua Tenenbaum, and Sergey Levine. Planning with diffusion for flexible behavior synthesis. In *International Conference on Machine Learning*, 2022. 3, 4
- [23] Daniel Jarrett, Ioana Bica, and Mihaela van der Schaar. Strictly batch imitation learning by energy-based distribution matching. *Advances in Neural Information Processing Systems*, 33:7354–7365, 2020. 9
- [24] Yunfan Jiang, Agrim Gupta, Zichen Zhang, Guanzhi Wang, Yongqiang Dou, Yanjun Chen, Li Fei-Fei, Anima Anandkumar, Yuke Zhu, and Linxi Fan. Vima: General robot manipulation with multimodal prompts. *arXiv preprint arXiv:2210.03094*, 2(3):6, 2022. 13
- [25] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020. 13
- [26] Maciej Kilian, Varun Jampani, and Luke Zettlemoyer. Computational tradeoffs in image synthesis: Diffusion, masked-token, and next-token prediction. *arXiv preprint arXiv:2405.13218*, 2024. 2, 3

- [27] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, et al. Openvla: An open-source vision-language-action model. *arXiv preprint arXiv:2406.09246*, 2024. 13
- [28] Lucy Lai, Ann Zixiang Huang, and Samuel J Gershman. Action chunking as policy compression. *PsyArXiv*, 2022. 13
- [29] Doyup Lee, Chiheon Kim, Saehoon Kim, Minsu Cho, and Wook-Shin Han. Autoregressive image generation using residual quantization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11523–11532, 2022. 4, 6, 9
- [30] Seungjae Lee, Yibin Wang, Haritheja Etukuru, H Jin Kim, Nur Muhammad Mahi Shafiullah, and Lerrel Pinto. Behavior generation with latent actions. *arXiv preprint arXiv:2403.03181*, 2024. 1, 13
- [31] Xinghang Li, Minghuan Liu, Hanbo Zhang, Cunjun Yu, Jie Xu, Hongtao Wu, Chilam Cheang, Ya Jing, Weinan Zhang, Huaping Liu, et al. Vision-language foundation models as effective robot imitators. *arXiv preprint arXiv:2311.01378*, 2023. 13
- [32] Xiang Li, Hao Chen, Kai Qiu, Jason Kuen, Jiuxiang Gu, Bhiksha Raj, and Zhe Lin. Imagefolder: Autoregressive image generation with folded tokens. *arXiv preprint arXiv:2410.01756*, 2024. 9
- [33] Xiang Li, Kai Qiu, Hao Chen, Jason Kuen, Zhe Lin, Rita Singh, and Bhiksha Raj. Controlvar: Exploring controllable visual autoregressive modeling. *arXiv preprint arXiv:2406.09750*, 2024. 9
- [34] Zhixuan Liang, Yao Mu, Hengbo Ma, Masayoshi Tomizuka, Mingyu Ding, and Ping Luo. Skilldiffuser: Interpretable hierarchical planning via skill abstractions in diffusion-based task execution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16467–16476, 2024. 7, 8
- [35] Yang Liu, Pengxiang Ding, Siteng Huang, Min Zhang, Han Zhao, and Donglin Wang. Pite: Pixel-temporal alignment for large video-language model. In *European Conference on Computer Vision*, pages 160–176. Springer, 2025. 9
- [36] 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. 5, 9
- [37] Xiaoxiao Ma, Mohan Zhou, Tao Liang, Yalong Bai, Tiejun Zhao, Huaian Chen, and Yi Jin. Star: Scale-wise text-to-image generation via auto-regressive representations. *arXiv preprint arXiv:2406.10797*, 2024. 9
- [38] 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. In *arXiv preprint arXiv:2108.03298*, 2021. 6, 7
- [39] Ajay Mandlekar, Soroush Nasiriany, Bowen Wen, Iretiayo Akinola, Yashraj Narang, Linxi Fan, Yuke Zhu, and Dieter Fox. Mimicgen: A data generation system for scalable robot learning using human demonstrations. In *7th Annual Conference on Robot Learning*, 2023. 7
- [40] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2337–2346, 2019. 5
- [41] Dean A Pomerleau. Alvin: An autonomous land vehicle in a neural network. *Advances in neural information processing systems*, 1, 1988. 9
- [42] Aaditya Prasad, Kevin Lin, Jimmy Wu, Linqi Zhou, and Jeannette Bohg. Consistency policy: Accelerated visuomotor policies via consistency distillation. *arXiv preprint arXiv:2405.07503*, 2024. 6
- [43] Kai Qiu, Xiang Li, Hao Chen, Jie Sun, Jinglu Wang, Zhe Lin, Marios Savvides, and Bhiksha Raj. Efficient autoregressive audio modeling via next-scale prediction. *arXiv preprint arXiv:2408.09027*, 2024. 9
- [44] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. 13
- [45] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 13
- [46] Ilija Radosavovic, Baifeng Shi, Letian Fu, Ken Goldberg, Trevor Darrell, and Jitendra Malik. Robot learning with sensorimotor pre-training. *arXiv:2306.10007*, 2023. 2
- [47] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67, 2020. 8
- [48] Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. *Advances in neural information processing systems*, 32, 2019. 9
- [49] Scott Reed, Konrad Zolna, Emilio Parisotto, Sergio Gomez Colmenarejo, Alexander Novikov, Gabriel Barth-Maron, Mai Gimenez, Yury Sulsky, Jackie Kay, Jost Tobias Springenberg, et al. A generalist agent. *arXiv preprint arXiv:2205.06175*, 2022. 13
- [50] Moritz Reuss, Maximilian Li, Xiaogang Jia, and Rudolf Lioutikov. Goal-conditioned imitation learning using score-based diffusion policies. *arXiv preprint arXiv:2304.02532*, 2023. 9
- [51] Moritz Reuss, Ömer Erdiñç Yağmurlu, Fabian Wenzel, and Rudolf Lioutikov. Multimodal diffusion transformer: Learning versatile behavior from multimodal goals. In *First Workshop on Vision-Language Models for Navigation and Manipulation at ICRA 2024*, 2024. 2
- [52] Nur Muhammad Mahi Shafiullah, Zichen Jeff Cui, Ariuntuya Altanzaya, and Lerrel Pinto. Behavior transformers: Cloning k modes with one stone. In *Thirty-Sixth Conference on Neural Information Processing Systems*, 2022. 1, 2, 5, 6, 9, 13
- [53] Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. Outra-

- geously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*, 2017. 7
- [54] Wenxuan Song, Han Zhao, Pengxiang Ding, Can Cui, Shangke Lyu, Yaning Fan, and Donglin Wang. Germ: A generalist robotic model with mixture-of-experts for quadruped robot. *arXiv preprint arXiv:2403.13358*, 2024. 9
- [55] Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. *arXiv preprint arXiv:2303.01469*, 2023. 9
- [56] Peize Sun, Yi Jiang, Shoufa Chen, Shilong Zhang, Bingyue Peng, Ping Luo, and Zehuan Yuan. Autoregressive model beats diffusion: Llama for scalable image generation. *arXiv preprint arXiv:2406.06525*, 2024. 9
- [57] Garrett Thomas, Ching-An Cheng, Ricky Loynd, Felipe Vieira Frujeri, Vibhav Vineet, Mihai Jalobeanu, and Andrey Kolobov. Plex: Making the most of the available data for robotic manipulation pretraining. In *Conference on Robot Learning*, pages 2624–2641. PMLR, 2023. 13
- [58] Keyu Tian, Yi Jiang, Zehuan Yuan, Bingyue Peng, and Liwei Wang. Visual autoregressive modeling: Scalable image generation via next-scale prediction. *arXiv preprint arXiv:2404.02905*, 2024. 4, 5, 6, 9
- [59] Faraz Torabi, Garrett Warnell, and Peter Stone. Behavioral cloning from observation. *arXiv preprint arXiv:1805.01954*, 2018. 5
- [60] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 4, 9
- [61] Yixiao Wang, Yifei Zhang, Mingxiao Huo, Ran Tian, Xiang Zhang, Yichen Xie, Chenfeng Xu, Pengliang Ji, Wei Zhan, Mingyu Ding, et al. Sparse diffusion policy: A sparse, reusable, and flexible policy for robot learning. *arXiv preprint arXiv:2407.01531*, 2024. 2, 6, 7, 8, 9
- [62] Jingyun Yang, Zi-ang Cao, Congyue Deng, Rika Antonova, Shuran Song, and Jeannette Bohg. Equibot: Sim (3)-equivariant diffusion policy for generalizable and data efficient learning. *arXiv preprint arXiv:2407.01479*, 2024. 9
- [63] Jiahui Yu, Xin Li, Jing Yu Koh, Han Zhang, Ruoming Pang, James Qin, Alexander Ku, Yuanzhong Xu, Jason Baldridge, and Yonghui Wu. Vector-quantized image modeling with improved vqgan. *arXiv preprint arXiv:2110.04627*, 2021. 9
- [64] 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. 5, 9
- [65] Andy Zeng, Pete Florence, Jonathan Tompson, Stefan Welker, Jonathan Chien, Maria Attarian, Travis Armstrong, Ivan Krasin, Dan Duong, Vikas Sindhwani, et al. Transporter networks: Rearranging the visual world for robotic manipulation. In *Conference on Robot Learning*, pages 726–747. PMLR, 2021. 9
- [66] Jinzhi Zhang, Feng Xiong, and Mu Xu. G3pt: Unleash the power of autoregressive modeling in 3d generation via cross-scale querying transformer. *arXiv preprint arXiv:2409.06322*, 2024. 9
- [67] Qian Zhang, Xiangzi Dai, Ninghua Yang, Xiang An, Ziyong Feng, and Xingyu Ren. Var-clip: Text-to-image generator with visual auto-regressive modeling. *arXiv preprint arXiv:2408.01181*, 2024. 9
- [68] Tianhao Zhang, Zoe McCarthy, Owen Jow, Dennis Lee, Xi Chen, Ken Goldberg, and Pieter Abbeel. Deep imitation learning for complex manipulation tasks from virtual reality teleoperation. In *2018 IEEE international conference on robotics and automation (ICRA)*, pages 5628–5635. IEEE, 2018. 9
- [69] Han Zhao, Min Zhang, Wei Zhao, Pengxiang Ding, Siteng Huang, and Donglin Wang. Cobra: Extending mamba to multi-modal large language model for efficient inference. *arXiv preprint arXiv:2403.14520*, 2024. 9
- [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. 1, 2, 5, 13
- [71] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5745–5753, 2019. 5
- [72] Yuke Zhu, Josiah Wong, Ajay Mandlekar, Roberto Martín-Martín, Abhishek Joshi, Soroush Nasiriany, Yifeng Zhu, and Kevin Lin. robosuite: A modular simulation framework and benchmark for robot learning. In *arXiv preprint arXiv:2009.12293*, 2020. 7

Supplementary Material

A. Limitations and Future Work

In this work, we propose CARP, a next-generation paradigm for robotic visuomotor policy learning, which effectively balances the long-standing trade-off between high performance and high efficiency seen in previous autoregressive modeling (AM) and diffusion modeling (DM) approaches. Despite these advancements, there remain several limitations and opportunities for improvement in future research.

First, the model structure of CARP can be further optimized for simplicity. Currently, CARP employs a two-stage design, where the first stage utilizes separate *multi-scale* action VQVAE modules for each action dimension to address their orthogonality. Future work could focus on developing a unified one-stage method that integrates *multi-scale* tokenization with the *coarse-to-fine* prediction process, resulting in a more efficient and streamlined framework without compromising performance.

Second, the multimodal capabilities of CARP remain underdeveloped. To mitigate the unimodality issue inherent in traditional autoregressive policies trained with MSE loss, CARP adopts a cross-entropy loss, which retains the potential for multimodality. Compared to the Diffusion Policy’s multimodal ability [11] likely coming from DDPM’s integration over a Stochastic Differential Equation [19], CARP’s approach is relatively straightforward. Further enhancements are needed to fully realize and leverage its multimodal potential.

Third, CARP’s adoption of the GPT-style paradigm opens up promising yet unexplored possibilities. Beyond the flexibility already demonstrated, the contextual understanding capabilities inherent in GPT-style architectures [44] suggest that CARP could be extended to support multimodal inputs [3, 45] like tactile and auditory information and address robotic tasks requiring long-term dependency reasoning [1]. Additionally, its potential for generalization in few-shot and zero-shot learning scenarios [7] represents a particularly valuable direction for future exploration.

Finally, but not exhaustively, the scaling properties of CARP represent a promising frontier. The scaling laws established in existing GPT-style models [25] could be seamlessly applied to CARP, suggesting that increasing network capacity and leveraging larger pre-training datasets could lead to substantial performance gains. Furthermore, recent advances in Vision-Language-Action (VLA) [6, 27, 31] models present an opportunity to integrate CARP into such frameworks, which could further validate its scalability and capabilities.




Figure 10. **Coarse-to-Fine Autoregressive Inference.** During inference, we leverage kv-caching to enable *coarse-to-fine* prediction without the need for causal masks. The full set of token maps, $\mathbf{r}_{1:K}$, is collectively decoded by the action VQVAE into executable actions for the robotic arm.




Figure 11. **Conventional Autoregressive Policy.** In reinforcement learning, conventional autoregressive policies generate action tokens sequentially, where each token is predicted based on the previously generated tokens. This differs from the action chunking prediction shown in Fig. 2a.

B. Definition of Autoregressive Policy

Autoregressive policies naturally capitalize on the efficiency and flexibility of autoregressive models. Initial works from a reinforcement learning perspective applied models like Transformers to predict the next action using states or rewards as inputs [10, 21], as shown in Fig. 11 and formalized as:

$$p(\mathbf{a}_t, \mathbf{a}_{t+1}, \dots, \mathbf{a}_{t+H-1}) = \prod_{k=t}^{t+H-1} p(\mathbf{a}_k | \mathbf{a}_{k-H:k-1}, \mathbf{s}_{k-H:k}), \quad (8)$$

where $\mathbf{s}_{k-H:k-1}$ represents the states or observations corresponding to the previous actions $\mathbf{a}_{k-H:k-1}$, and \mathbf{s}_k is the current state or observation. Following this paradigm, several subsequent works [24, 49, 57] employ autoregressive models to predict one action at a time during inference.

More recently, the concept of action chunking [28], derived from neuroscience, has demonstrated notable benefits for imitation learning [12, 30, 52, 70]. In action chunking, individual actions are grouped and executed as cohesive units, leading to improved efficiency in storage and

execution, as depicted in Fig. 2a. This paradigm extends the capabilities of GPT-style decoders by modifying them to generate chunks of actions in one forward pass, replacing the traditional single-step autoregressive operation with a multi-token, pseudo-autoregressive process. The action generation process for this paradigm is described as:

$$p(\mathbf{a}_t, \mathbf{a}_{t+1}, \dots, \mathbf{a}_{t+H-1}) = \prod_{k=t}^{t+H-1} p(\mathbf{a}_k | \mathbf{s}_O), \quad (9)$$

where O is the historical horizon. The model predicts the entire action sequence in one forward pass without strictly adhering to step-by-step autoregressive operations.

Given the significant performance improvements enabled by action chunking, we adopt this multi-token, one-forward-pass framework throughout the article and experiments when referring to autoregressive policies (AP).

C. Implementation Consistent with Diffusion Policy

We adopt similar experimental settings with 1 or 2 observations, a prediction horizon of 16, and an executable action length of 8, following the standard setup used in Diffusion Policy [11].

It is important to note that our classical formulation introduces a minor discrepancy in the horizon definition compared to the implementation of Diffusion Policy [11]. Specifically, in our formulation, the horizon H encompasses the past observed steps, meaning that the index of the current next predicted action is O , rather than 0. In contrast, as outlined in formulation Eq. (1), the horizon H does not include past observations, with the first prediction step corresponding to the next time step. While the rationale behind this design remains unclear due to limitations in the author’s understanding, we retain the horizon definition introduced by Diffusion Policy [11] to ensure consistency in our experimental comparisons.




Figure 12. Visualization of Tasks in Single-Task Experiment.

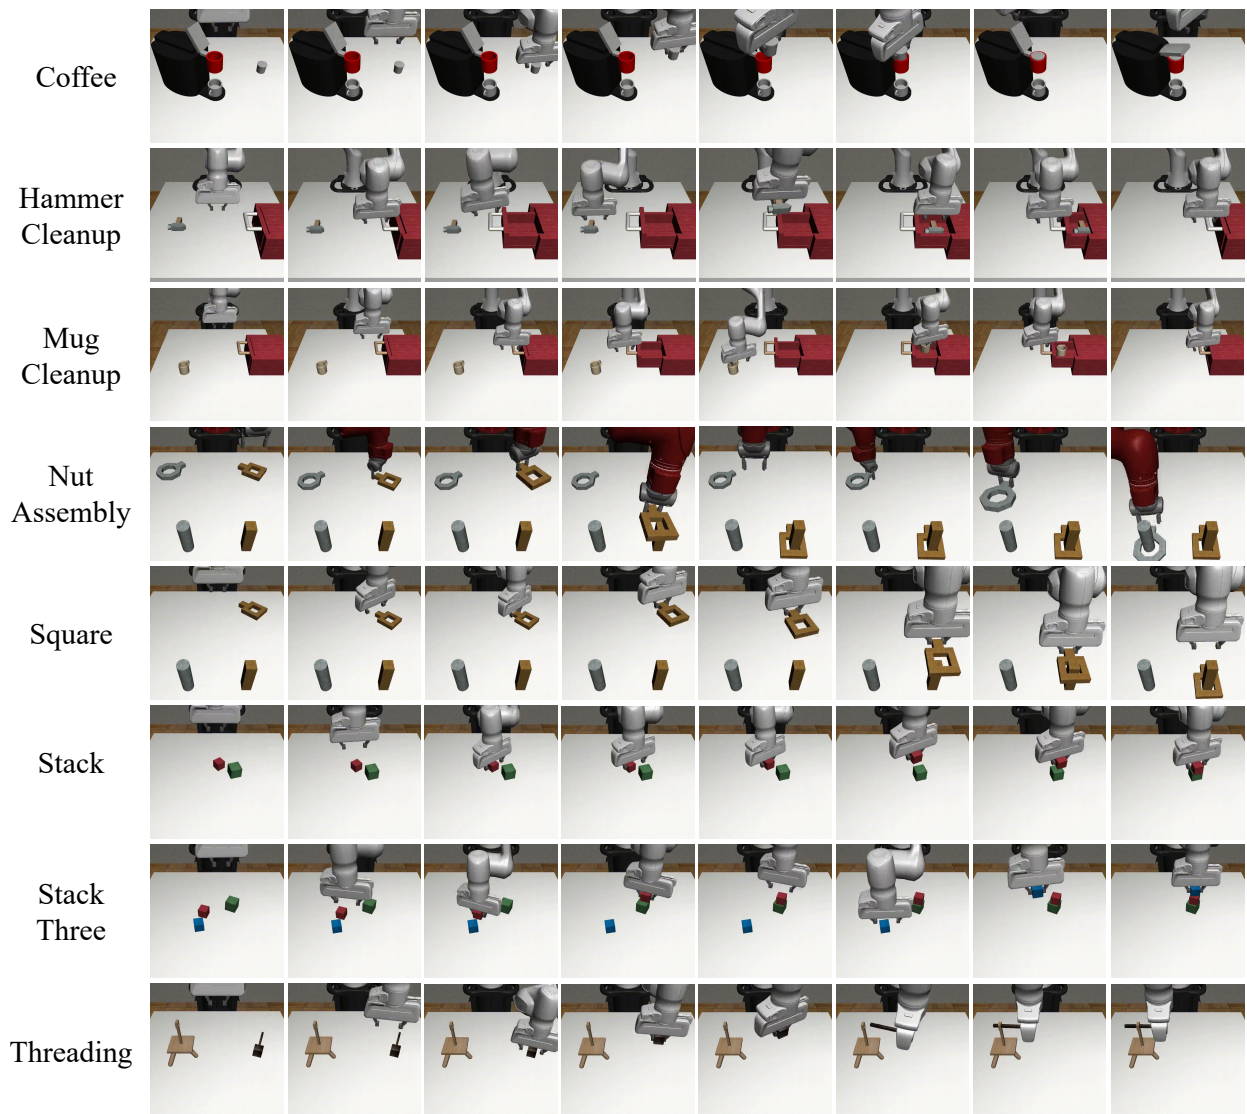


Figure 13. Visualization of Tasks in Multi-Task Experiment.



Figure 14. Visualization of All Interaction Tasks in Kitchen Experiment.

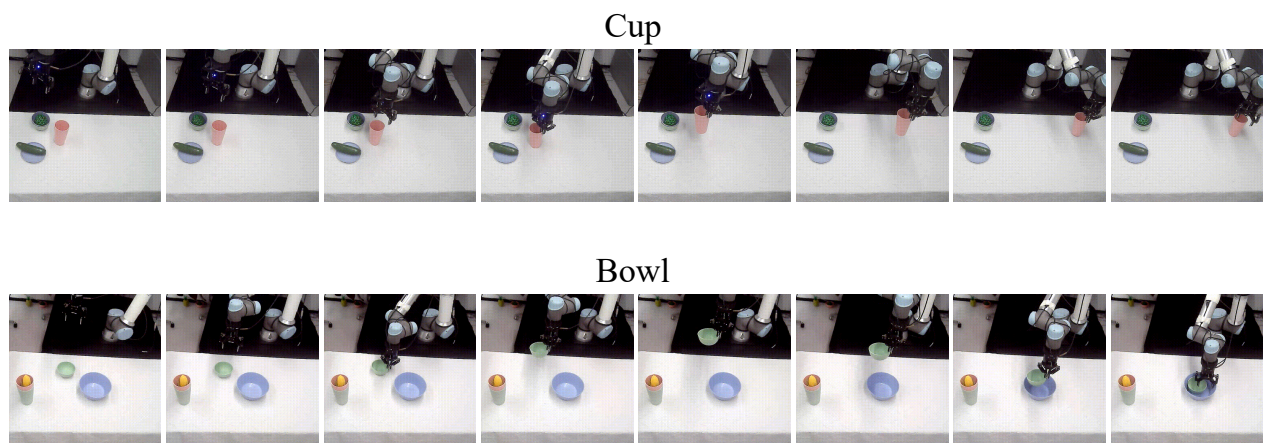


Figure 15. Visualization of Tasks in Real-World Setup.