

DiSPo: Diffusion-SSM based Policy Learning for Coarse-to-Fine Action Discretization

Nayoung Oh, Jaehyeong Jang, Moonkyeong Jung, and Daehyung Park[†]

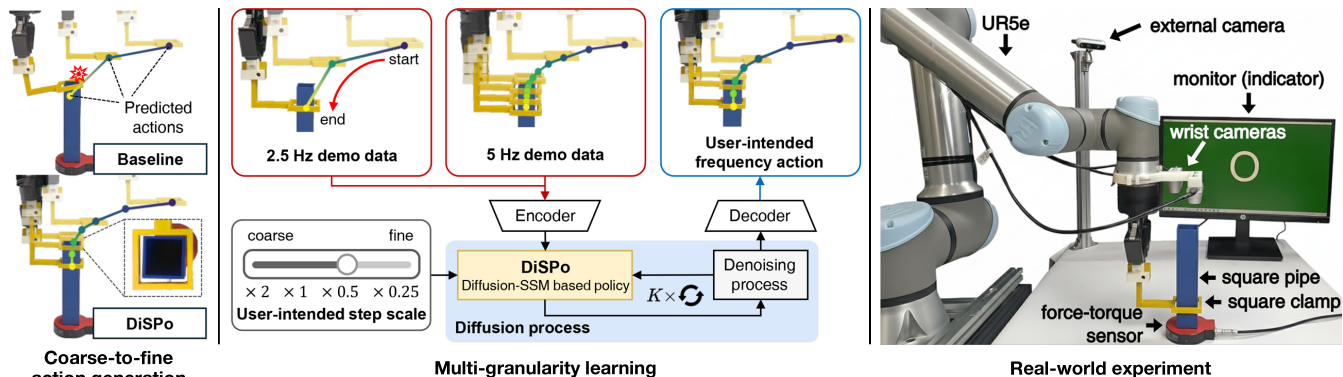


Fig. 1: Overview of DiSPo: a diffusion-SSM based policy for coarse-to-fine imitation learning. Our proposed DiSPo framework generates fine-grained actions from coarse demonstrations, while the baseline method generates rough actions. Leveraging the representation power of diffusion policy and the flexible discretization capabilities of SSMs, DiSPo learns from multi-granularity demonstrations and generates actions at user-intended granularities. A demonstration of a real-world *clamp passing* task, which requires high precision control to avoid collision.

Abstract—We aim to solve the problem of learning user-intended granular skills from multi-granularity demonstrations. Traditional learning-from-demonstration methods typically rely on extensive fine-grained data, interpolation techniques, or dynamics models, which are ineffective at encoding or decoding the diverse granularities inherent in skills. To overcome it, we introduce a novel diffusion-state space model (SSM) based policy (DiSPo) that leverages an SSM, Mamba, to learn from diverse coarse demonstrations and generate multi-scale actions. Our proposed step-scaling mechanism in Mamba is a key innovation, enabling memory-efficient learning, flexible granularity adjustment, and robust representation of multi-granularity data. DiSPo outperforms state-of-the-art baselines on coarse-to-fine benchmarks, achieving up to an 81% improvement in success rates while enhancing inference efficiency by generating inexpensive coarse motions where applicable. We validate DiSPo’s scalability and effectiveness on real-world manipulation scenarios. Code and Videos are available at <https://robo-dispo.github.io>.

I. INTRODUCTION

Researchers have increasingly focused on endowing robots with dexterous, generalizable policies such as human manipulations. These manipulations are often a mixture of coarse-to-fine actions [1], we call *multi-granularity* actions. These involve large positioning movements alongside precise maneuvers critical for tasks such as screwing, welding, and

The authors are with Korea Advanced Institute of Science and Technology, Korea ({nyoh, wogud9019, jmk7791, daehyung}@kaist.ac.kr). [†]D. Park is the corresponding author.

This work was supported by the Institute of Information & communications Technology Planning & Evaluation(IITP) grant funded by the Korea government(MSIT) (No. RS-2024-00336738, RS-2022-II220311, RS-2024-00509279) and the Technology Innovation Program(RS-2025-25453780) funded by the Ministry of Trade Industry & Resources(MOTIR, Korea).

insertion. Learning these locally precise behaviors is crucial to task success [2], [3], [4].

In this context, we aim to solve the problem of generating manipulation skills at multiple levels of granularity through imitation learning (IL), a process we call *multi-granularity learning* as shown in Fig. 1. This requires models to learn from both fine-grained and general coarse demonstrations. Further, the models need to generate precise actions across varying control scales according to user needs, understanding the temporal structure of demonstrations. We term it as *multi-granularity reproduction*.

Traditional IL methods, such as dynamic movement primitives [5], learn complex trajectories by adopting dynamics models [6]. These methods allow for frequency adjustments in output, learning from a specific frequency of input trajectories. In the line of research, state space models (SSMs), such as Mamba [7], offer memory-efficient, powerful encoding. However, their fixed action representations struggle to capture complex or multi-modal behaviors across diverse task conditions or modalities.

Alternatively, neural IL methods, such as behavior transformers [8], [9], [10] and diffusion-based policies [11], [12], [13], [14], are increasingly acquiring attention with expressive power and robustness. These approaches are capable of learning from diverse, high-dimensional multi-modal datasets [15], [16], [17], [18], [19]. However, most approaches learn from a specific frequency of trajectories [20] or an unspecified timescale of state-action pairs [15], without understanding *multi-granularity*. Further, modeling fine-grained skills typically requires high-frequency demonstrations causing storage and computational overhead.

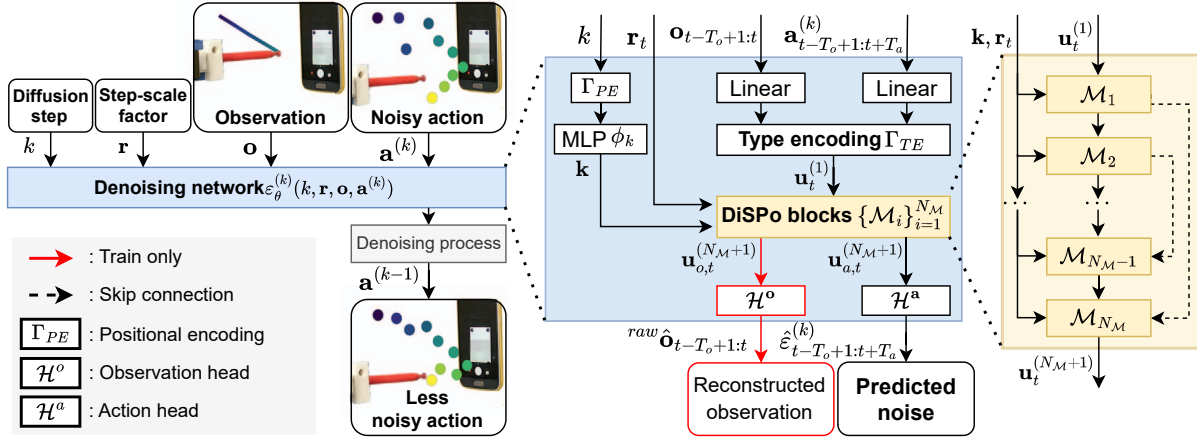


Fig. 2: Illustration of the DiSPo architecture. DiSPo takes diffusion step k , step-scale factors \mathbf{r}_t , encoded observations $\mathbf{o}_{t-T_o+1:t}$, and noisy actions $\mathbf{a}_{t-T_o+1:t+T_a}^{(k)}$. The model identifies the noise $\hat{\varepsilon}_{t-T_o+1:t+T_a}^{(k)}$ within the input noisy actions through stacked DiSPo blocks and utilizes the identified noise to generate the less noisy action $\mathbf{a}_{t-T_o+1:t+T_a}^{(k-1)}$.

We propose a novel coarse-to-fine imitation learning algorithm, diffusion-SSM based policy (DiSPo), combining the representation power of diffusion models with the flexible discretization power of SSMs. We particularly adopt a state-of-the-art SSM, Mamba, to enable DiSPo to learn and reproduce trajectories at *multi-granularity* through data-efficient training strategies as shown in Fig. 1. We show that DiSPo is capable of producing varying scales of behavior, not only learning from multiple rates of coarse demonstrations but also modulating the discretization level of trajectories through a granularity predictor online. To the best of our knowledge, this is the first attempt to modulate Mamba’s discrete model for fine-grained manipulations. We introduce novel coarse-to-fine IL benchmarks evaluating our method against state-of-the-art visuomotor policy learning methods. The evaluation shows the modulation of step size in DiSPo generates finer movements with expert-like behaviors.

II. PRELIMINARIES

An SSM describes a dynamic system that accepts an input sequence $\mathbf{u}(t) \in \mathbb{R}$, produces an output sequence $\mathbf{y}(t) \in \mathbb{R}$, and updates a set of internal states $\mathbf{h}(t) \in \mathbb{R}^N$, where N denotes the dimensions of the state. The system consists of first-order differential equations, known as state and output equations: $\dot{\mathbf{h}}(t) = \mathbf{A}\mathbf{h}(t) + \mathbf{B}\mathbf{u}(t)$, $\mathbf{y}(t) = \mathbf{C}\mathbf{h}(t)$, where $\mathbf{A} \in \mathbb{R}^{N \times N}$, $\mathbf{B} \in \mathbb{R}^{N \times 1}$, and $\mathbf{C} \in \mathbb{R}^{1 \times N}$ are the state, input, and output parameters, respectively.

For discrete computations, the SSM transforms the continuous-time system into a discrete-time system, defined over a discrete input sequence $\mathbf{u}_{1:L} \in \mathbb{R}^{L \times 1}$ and output sequence $\mathbf{y}_{1:L} \in \mathbb{R}^{L \times 1}$ where L denotes the sequence length. Given a step size $\Delta \in \mathbb{R}$, the discrete-time system at time step t is $\mathbf{h}_t = \bar{\mathbf{A}}\mathbf{h}_{t-1} + \bar{\mathbf{B}}\mathbf{u}_t$, $\mathbf{y}_t = \mathbf{C}\mathbf{h}_t$, where the discrete parameters are

$$\bar{\mathbf{A}} = \exp(\Delta\mathbf{A}), \quad \bar{\mathbf{B}} = (\Delta\mathbf{A})^{-1}(\exp(\Delta\mathbf{A}) - \mathbf{I}) \cdot \Delta\mathbf{B}, \quad (1)$$

following the zero-order hold (ZOH) discretization rule. In contrast to S4 [21] with fixed step sizes, Mamba makes

parameters $(\mathbf{B}, \mathbf{C}, \Delta)$ as a function of the input \mathbf{u}_t ,

$$\mathbf{B}_t = f_B(\mathbf{u}_t), \quad \mathbf{C}_t = f_C(\mathbf{u}_t), \quad \Delta_t = \text{SoftPlus}(f_\Delta(\mathbf{u}_t)), \quad (2)$$

where f_B , f_C , and f_Δ are trainable layers, and SoftPlus is an activation function. In this work, we extend the SSM to process D -dimensional input and output sequences, $\mathbf{u}_{1:L} \in \mathbb{R}^{L \times D}$ and $\mathbf{y}_{1:L} \in \mathbb{R}^{L \times D}$. We handle each of the D channels independently, but share the matrix \mathbf{A} across all channels [21], [7].

III. DIFFUSION-SSM BASED POLICY MODEL

Our proposed model, DiSPo, generates multi-granularity actions using a Mamba-based diffusion process. The model progressively denoises an input sequence by predicting the parameter of a discrete SSM (e.g., \mathbf{A} and \mathbf{B}), conditioned on a user-defined step-scale factor \mathbf{r} . Fig. 2 illustrates the architecture, incorporating a denoising diffusion probabilistic model (DDPM) [22] with N_M stacked DiSPo blocks $\{\mathcal{M}_i\}_{i=1}^{N_M}$, where each DiSPo block is a variant of the Mamba block. Inspired by the decoder-only Mamba (D-Ma) [23], we design the architecture to learn denoising networks $\varepsilon_\theta^{(k)}$, parameterized by θ . These networks generate a less noisy action sequence $\mathbf{a}^{(k-1)}$, at the k -th denoising step ($k \in [1, \dots, K]$), conditioned on a history of observations \mathbf{o} , noisy actions $\mathbf{a}^{(k)}$, and step-scale factors \mathbf{r} :

$$\mathbf{a}^{(k-1)} = \alpha \left(\mathbf{a}^{(k)} - \gamma \varepsilon_\theta^{(k)}(k, \mathbf{r}, \mathbf{o}, \mathbf{a}^{(k)}) + \mathcal{N}(0, \sigma^2 \mathbf{I}) \right), \quad (3)$$

where α , γ , and σ are the noise schedule parameters following the DDPM formulation [22]. For notational simplicity, we omit the time index t . Starting from an initial Gaussian noise sample, $\mathbf{a}^{(K)}$, DiSPo recursively applies the denoising process to generate an imitated action sequence.

A distinct feature of DiSPo is the integration of factors \mathbf{r} into Mamba, inspired by manual adjustment of rates in time-invariant SSMs [21], [24]. The adjustable factor \mathbf{r} represents the discretization scale (i.e., temporal granularity) of the sequences in the discrete SSM. This allows DiSPo to learn

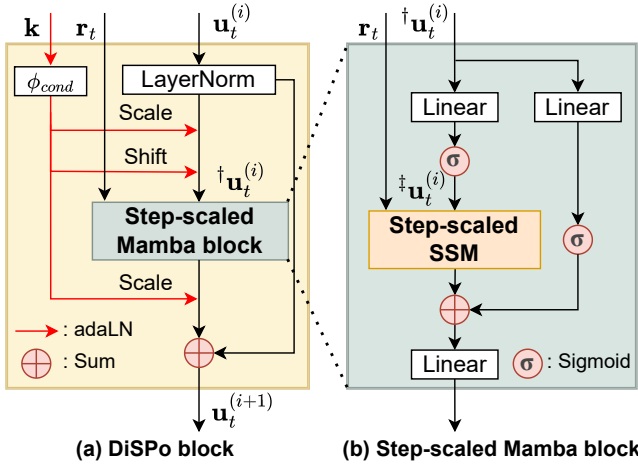


Fig. 3: (a) A DiSPo block \mathcal{M}_i refines noise-related features in the type encoded sequence $\mathbf{u}_t^{(i)}$ using adaLN conditioned on the diffusion step embedding \mathbf{k} . (b) A step-scaled Mamba block takes \mathbf{r}_t and $\dagger\mathbf{u}_t^{(i)}$.

from multiple rates of demonstrations and to adjust step sizes for discrete SSM parameters. We describe the details below.

A. Mamba-based denoising process

Consider an input sequence $\mathbf{u}_t^{(1)} \in \mathbb{R}^{L \times D}$ in the k -th diffusion step and the time step t , where L and D are the length and dimension of the input sequence, respectively. Note that, to simplify the notation, we omit k and retain i for the variables defined in the k -th step below (e.g. $\mathbf{u}_t^{(k,1)} = \mathbf{u}_t^{(1)}$). The Mamba-based denoising network predicts the action noise $\hat{\varepsilon}^{(k)}$ by updating the sequences $\mathbf{u}_t^{(i)}$ with noise-relevant features through the $\{\mathcal{M}_i\}_{i=1}^{N_{\mathcal{M}}}$ blocks. Then it transforms the action component of the last updated sequence $\mathbf{u}_{a,t}^{(N_{\mathcal{M}}+1)}$ into the action noise through an output action head \mathcal{H}^a ,

$$\mathbf{u}_t^{(i+1)} = \mathcal{M}_i(\mathbf{k}, \mathbf{r}_t, \mathbf{u}_t^{(i)}) \quad \text{and} \quad \hat{\varepsilon}^{(k)} = \mathcal{H}^a(\mathbf{u}_{a,t}^{(N_{\mathcal{M}}+1)}), \quad (4)$$

where $i \in [1, \dots, N_{\mathcal{M}}]$, $\mathbf{k} \in \mathbb{R}^D$ is an embedding for the diffusion step k , and $\mathbf{r}_t \in \mathbb{R}^L$. Each \mathcal{M}_i block processes input sequences with the same size, $\mathbf{u}_t^{(i)} \in \mathbb{R}^{L \times D}$. The denoising process consists of three parts: input encoding, diffusion process, and noise prediction.

Input encoding: The first DiSPo block takes the input sequence $\mathbf{u}_t^{(1)}$, a diffusion step embedding \mathbf{k} , and step-scale factors \mathbf{r}_t at each k -th step. The input sequence $\mathbf{u}_t^{(1)}$ consists of observation and noisy-action embeddings over lengths T_o and $T_o + T_a$, respectively. We represent it as

$$\mathbf{u}_t^{(1)} = [\Gamma_{\text{TE}}(\mathbf{f}_{o,t-T_o+1}), \dots, \Gamma_{\text{TE}}(\mathbf{f}_{o,t}), \Gamma_{\text{TE}}(\mathbf{f}_{a,t-T_o+1}), \dots, \Gamma_{\text{TE}}(\mathbf{f}_{a,t+T_a})], \quad (5)$$

where $\mathbf{f}_{o,t}$ and $\mathbf{f}_{a,t}$ are observation and action features, respectively. $\Gamma_{\text{TE}} : \mathbb{R}^D \rightarrow \mathbb{R}^D$ represent *type encoding*, which injects a learnable vector to the input ($\in \mathbb{R}^D$). Note that $L = 2T_o + T_a$. An observation feature $\mathbf{f}_{o,t} \in \mathbb{R}^D$ is an embedding vector of the observation $\mathbf{o}_t \in \mathbb{R}^{D_o}$ preprocessed

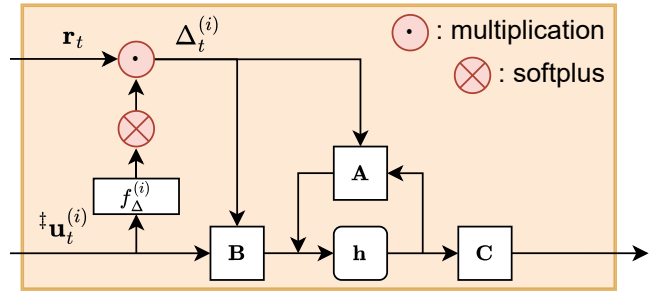


Fig. 4: A step-scaled SSM takes input sequence $\dagger\mathbf{u}_t^{(i)}$ and \mathbf{r}_t to scale $\Delta_t^{(i)}$, and discretizes the learned SSM parameters using the step sizes.

from raw sensory observations $\text{raw}\mathbf{o}_t$ at a timestep t . The embedding process is a linear projection by $\mathbf{f}_{o,t} = \mathbf{w}_o\mathbf{o}_t + \mathbf{b}_o$ with a weight matrix $\mathbf{w}_o \in \mathbb{R}^{D \times D_o}$ and a bias $\mathbf{b}_o \in \mathbb{R}^D$. In this work, we use \mathbf{o}_t as a concatenated vector of an image encoding from ResNet18 [25] with attentional pooling [26] and a proprioception vector (e.g., end-effector positions) normalized in the range of $[-1, 1]$.

An action feature $\mathbf{f}_{a,t} \in \mathbb{R}^D$ is an embedding vector of the action $\mathbf{a}_t \in \mathbb{R}^{D_a}$, obtained either by normalizing the raw action command $\text{raw}\mathbf{a}_t$ with noise during training or by denoising the noisy action from the previous diffusion step during inference. The embedding process is a linear projection by $\mathbf{f}_{a,t} = \mathbf{w}_a\mathbf{a}_t + \mathbf{b}_a$ with a weight matrix $\mathbf{w}_a \in \mathbb{R}^{D \times D_a}$ and a bias $\mathbf{b}_a \in \mathbb{R}^D$. In this work, we use a pose vector as a command, normalizing in the range of $[-1, 1]$.

Lastly, as a part of input conditions, we embed the diffusion step k into a D -dimensional vector $\mathbf{k} = \phi_k(\Gamma_{\text{PE}}(k))$ by sinusoidal *positional encoding* $\Gamma_{\text{PE}} : \mathbb{R} \rightarrow \mathbb{R}^D$ followed by a multi-layer perceptron $\phi_k : \mathbb{R}^D \rightarrow \mathbb{R}^D$. We provide further details on the step-scale factors \mathbf{r}_t in Sec. III-B.

Diffusion process: At each diffusion step, we update the sequence $\mathbf{u}_t^{(i)}$ with noise-relevant features through stacked DiSPo blocks $\{\mathcal{M}_i\}_{i=1}^{N_{\mathcal{M}}}$ with skip connections. Fig. 3 shows a DiSPo block that is a step-scaled Mamba block with adaptive layer normalization (adaLN) [27], performing dimension-wise scaling and shifting $\mathbf{u}_t^{(i)}$ into $\dagger\mathbf{u}_t^{(i)}$ conditioned on the diffusion step embedding \mathbf{k} . Taking $\dagger\mathbf{u}_t^{(i)}$, the step-scaled Mamba block adjusts the parameters of discrete-time SSM, according to user needs, i.e., a vector of step-scale factors $\mathbf{r}_t \in \mathbb{R}_{>0}^L$, and then updates the input sequence via the internal step-scaled SSM. In contrast to conventional Mamba blocks, we exclude convolutional layers that limit handling diverse granularity of input sequences due to fixed-size receptive fields.

Fig. 4 shows the proposed step-scaled SSM for *multi granularity*. Our SSM predicts the appropriate step size $\Delta_t^{(i)} \in \mathbb{R}_{>0}^{L \times D}$ with respect to the input sequence $\dagger\mathbf{u}_t^{(i)}$, a non-linear projection of $\dagger\mathbf{u}_t^{(i)}$, and the user-intended scales \mathbf{r}_t ,

$$\Delta_t^{(i)} = \mathbf{r}_t \cdot \text{SoftPlus} \left(f_{\Delta}^{(i)} \left(\dagger\mathbf{u}_t^{(i)} \right) \right), \quad (6)$$

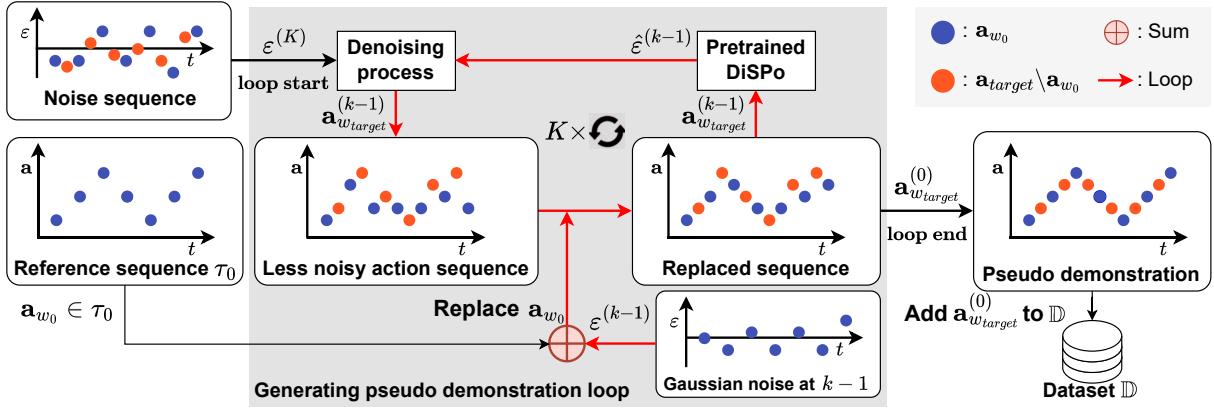


Fig. 5: Generating a pseudo demonstration for fine-tuning. Starting from Gaussian noise $\varepsilon^{(K)}$ and a reference sequence τ_0 , the model iteratively denoises and replaces w_0 frequency actions in the less noisy action sequence with noise added $\mathbf{a}_{w_0} \in \tau_0$. We repeat this process until the model generates a noise-less action sequence at target frequency $\mathbf{a}_{w_{\text{target}}}^{(0)}$, which we refer to as a pseudo demonstration.

where $f_{\Delta}^{(i)}$ is a block-wise trainable linear layer used in Eq. (2). We use $\Delta_t^{(i)}$ to calculate $\bar{\mathbf{A}}$ and $\bar{\mathbf{B}}$ following Eq. (1). **Noise prediction:** After $N_{\mathcal{M}}$ times of feature updates, the action head \mathcal{H}^a predicts the action noise $\hat{\varepsilon}_{t-T_o+1:t+T_a}^{(k)}$ with respect to $\mathbf{u}_{a,t}^{(N_{\mathcal{M}}+1)}$ that corresponds to the noisy action input $\mathbf{a}_{t-T_o+1:t+T_a}^{(k)}$ for the k -th denoising process. We then use the predicted noise to find the denoised action input $\mathbf{a}_{t-T_o+1:t+T_a}^{(k-1)}$ for the next diffusion step $k-1$, following Eq. (3) in inference.

In addition, during training, we enable our model to reconstruct the given raw observation $\text{raw } \mathbf{o}_t$ decoding the updated sequence $\mathbf{u}_{o,t}^{(N_{\mathcal{M}}+1)}$ through an observation head \mathcal{H}^o . The reconstruction helps the model to keep capturing fine details in observations across layers. Here, the decoder consists of a linear layer for low-dimensional observation and a ResNet18 decoder for image-based observation.

B. Multi-granularity reproduction

To control the granularity of generated actions, DiSPo takes a vector of step-scale factors, $\mathbf{r}_t = [r_{t-T_o+1}^o, \dots, r_t^o, r_{t-T_o+1}^a, \dots, r_t^a, \dots, r_{t+T_a}^a]$, where r_t^o and r_t^a represent the step size scales, we call factors, of the observation \mathbf{o}_t and the action \mathbf{a}_t relative to those of a reference sequence. We apply identical scales to the observation and past action sequences, such that $\mathbf{r}_{t-T_o+1:t-1}^o = \mathbf{r}_{t-T_o+1:t-1}^a$.

To define the reference step size, we use a mode selection approach that chooses the most frequently observed step size in demonstrations. DiSPo then allows for manual selection of the desired step-scale factors. For example, we set $\mathbf{r}_t = \mathbf{1}_{t-T_o+1:t} + \mathbf{1}_{t-T_o+1:t-1} + 0.5_{t:t+T_a}$ when we want to achieve twice finer actions and $\mathbf{r}_t = \mathbf{1}_{t-T_o+1:t} + \mathbf{1}_{t-T_o+1:t-1} + 2_{t:t+T_a}$ for twice coarser actions. In addition, DiSPo includes a step-scale factor predictor ϕ_r , implemented as an MLP, which predicts a factor \mathbf{r}_t to accomplish the task given the observation sequence $\mathbf{o}_{t-T_o+1:t}$.

IV. MULTI-GRANULARITY POLICY LEARNING

We introduce a *multi-granularity learning* scheme to improve the prediction performance of high-frequency actions that are not available in the demonstration dataset \mathbb{D} . Our scheme consists of two steps: 1) pretraining with sample-rate augmentation and 2) fine-tuning with pseudo actions.

In pretraining, to handle various granularities, we first augment the dataset \mathbb{D} with random step-scale factors. We randomly draw a reference sequence $\tau_0 = [(\text{raw } \mathbf{o}_1, \text{raw } \mathbf{a}_1), \dots, (\text{raw } \mathbf{o}_T, \text{raw } \mathbf{a}_T)] \in \mathbb{D}$ with length T and sample frequency ω_0 . By selecting a frequency $\omega_j \leq \omega_0$, we resample a sequence τ_j with step-scale factors $\mathbf{r}_j = \frac{\mathbf{1}_L}{(\omega_j/\omega_0)}$ from τ_0 . Repetition of these enhancements creates the N_{ω} number of random frequency sequences: $\tau = \{\tau_1, \dots, \tau_{N_{\omega}}\}$. We then introduce a total loss $\mathcal{L} = \mathcal{L}_{MSE}^{\varepsilon} + \lambda \cdot \mathcal{L}_{MSE}^o$, where $\mathcal{L}_{MSE}^{\varepsilon}$, \mathcal{L}_{MSE}^o , and λ are a noise prediction error loss, an observation reconstruction loss, and a weighting coefficient ($\in \mathbb{R}_{>0}$), respectively. In detail, $\mathcal{L}_{MSE}^{\varepsilon}$ uses the mean squared error (MSE) to minimize a variational bound on the KL divergence between the true denoising process and that modeled by DiSPo:

$$\mathcal{L}_{MSE}^{\varepsilon} = \text{MSE}(\varepsilon^{(k)}, \varepsilon_{\theta}(k, \mathbf{r}_t, \mathbf{o}_t, \mathbf{a}_t^{(0)} + \varepsilon^{(k)})). \quad (7)$$

where $k \in [0, \dots, K]$. Likewise, \mathcal{L}_{MSE}^o is the MSE between an input observation $\text{raw } \mathbf{o}_t$ and its reconstruction from the observation head \mathcal{H}^o .

In fine-tuning, we co-finetune DiSPo on original and pseudo demonstration dataset to produce high-frequency actions not available in the dataset \mathbb{D} . Fig. 5 shows the process of generating fine-grained pseudo demonstrations. We randomly draw a reference sequence τ_0 with its frequency w_0 and generate a fine-grained sequence using the pretrained DiSPo by selecting a target frequency $\omega_{\text{target}} > \omega_0$ with $\mathbf{r} = \frac{\mathbf{1}_L}{(\omega_{\text{target}}/\omega_0)}$. In practice, starting from Gaussian noise, we perform the diffusion process K times to generate a noise-less action sequence $\mathbf{a}_{1:T'}^{(0)}$, where $T' = T \cdot \omega_{\text{target}}/\omega_0$.

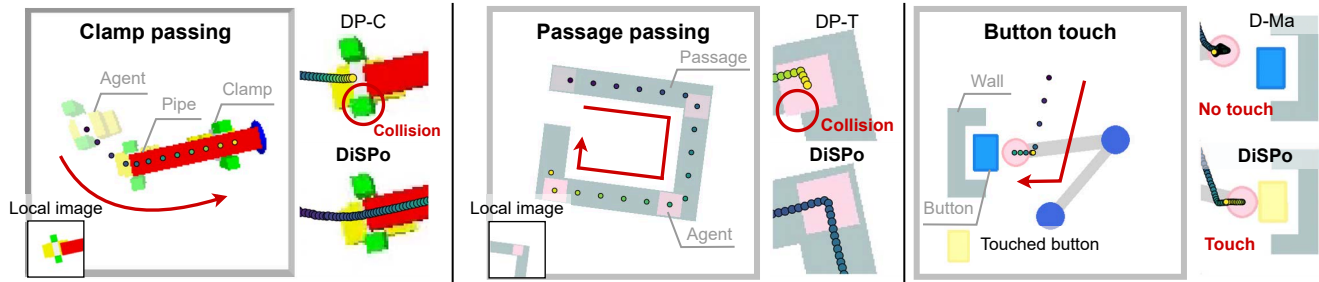


Fig. 6: Illustrations of three simulation benchmarks, *clamp passing*, *passage passing*, and *button touch*. Dots denote either demonstrations at 2.5 Hz or predicted actions from DiSPo and baselines.

However, the pretrained DiSPo is not sufficient to accurately produce high-frequency actions yet.

To figure it out, we decompose the predicted high-frequency actions $\mathbf{a}_{1:T'}$ into a subset with ω_0 frequency of actions $\mathbf{a}_{w_0}^{(k)}$ and its complement, $\mathbf{a}_{1:T'} \setminus \mathbf{a}_{w_0}^{(k)}$. At each k -th denoising process, we replace $\mathbf{a}_{w_0}^{(k)}$ with the demonstration actions in τ_0 with noise, $\mathbf{a}_{1:T'} + \epsilon^{(k)}$. This replacement helps in producing fine-grained pseudo actions that remain close to the demonstrations. In addition, as DiSPo predicts an action chunk, it produces multiple actions at a timestep. We aggregate these repeated predictions by weighted averaging, following the temporal ensemble in ACT [28], to obtain the final fine-grained action sequence. In contrast, the generation of fine-grained observations remains challenging. Thus, we retain the original frequency ω_0 of the observations by setting $r_t^o = 1$ and $r_t^a = \omega_0/\omega_{\text{target}}$ in fine-tuning. We call each outcome sequence a pseudo demonstration.

We fine-tune DiSPo using both pseudo demonstrations and original demonstrations. In practice, we repeat the generation of pseudo demonstrations and fine-tuning, gradually increasing the target frequency ω_{target} . Note that we fine-tune the model with the loss \mathcal{L} corresponding to $\mathbf{a}_{w_0}^{(k)}$ only since the predicted actions $\mathbf{a}_{1:T'} \setminus \mathbf{a}_{w_0}^{(k)}$ are not reliable as original demonstrations. However, sequential prediction with finer step-scale factors helps fine-tuning it as SSM internal state propagates through a sequence.

V. EXPERIMENTAL SETUP

We conduct quantitative and qualitative evaluations using three simulated benchmarks and two real-world manipulation tasks. The benchmarks statistically assess the ability to generate fine-grained actions from coarse demonstrations. Below, we describe each benchmark in detail.

Clamp passing: A gripper agent (yellow) manipulates a clamp (green) to precisely approach and pass through a 2D pipe (red) without collision, as shown in Fig. 6 (Left). The raw observation ${}^{\text{raw}}\mathbf{o}_t$ comprises the agent’s pose ($\in \mathbb{R}^3$) and two RGB images ($\in \mathbb{Z}^{96 \times 96 \times 3}$), one focusing on the agent (Fig. 6 left local image) and the other capturing the entire scene. The raw action ${}^{\text{raw}}\mathbf{a}_t$ is the agent’s target pose ($\in \mathbb{R}^3$). We randomize the initial agent pose and vary the pipes’ geometric properties (length and thickness) and spatial pose (position and orientation) using Pybullet [29].




Fig. 7: A capture of the real-world *button touch* environment.

Passage passing: A rectangular agent (pink) precisely maneuvers through a narrow 2D passage (gray) navigating corners without collision, as shown in Fig. 6 (Middle). The observation ${}^{\text{raw}}\mathbf{o}_t$ includes the agent’s pose ($\in \mathbb{R}^2$) and two RGB images as in the *clamp passing* benchmark. The action is the 2D target position aligning the agent with the passage boundary. We randomize the passage’s shape, width, and orientation using Pymunk [30] and Pygame [31].

Button touch: A two-link planar arm precisely touches a button (blue) without causing a collision between the button and wall, as shown in Fig. 6 (Right). The observation ${}^{\text{raw}}\mathbf{o}_t$ consists of the end-effector position ($\in \mathbb{R}^2$) and an RGB image. The action ${}^{\text{raw}}\mathbf{a}_t$ is the desired end-effector position ($\in \mathbb{R}^2$). We randomize the initial arm configuration and button placement, using Pymunk and Pygame.

For evaluation, we record 90 demonstrations for each benchmark at frequencies ranging from 2.5 Hz to 20 Hz using the toppra path planning library [32]. We then train our method and four baselines on these demonstrations, selecting the best checkpoints based on performance across 50 randomly sampled validation environments.

We finally evaluate the approaches on 100 unseen test environments. The four baselines are as follows:

- DiffusionPolicy-C (DP-C) and DiffusionPolicy-T (DP-T) [12]: CNN- and Transformer-based diffusion policies, respectively.
- D-Ma [23]: A decoder-only variant of the Mamba-based diffusion model, MaLL.
- VQ-BeT [9]: A vector-quantized behavior Transformer (BeT) tokenizing continuous actions.




Fig. 8: Comparison of task success rates [%] across four frequencies of demonstrations per simulated benchmark. We train each method with a source frequency (x -axis) of demonstrations and test a 20 Hz target frequency of actions in new environments. Note that Tracker is a complexity indicator.

As baselines require fixed step sizes, we linearly interpolate their action sequences. In contrast, DiSPo generates fine-grained actions on demand using user-intended or predicted step-scale factors from the learned predictor ϕ_r . For comparison, we compute step-scale factors based on the demonstration and required frequency. We also use relative poses as desired actions when advantageous for baselines; baselines adopt relative poses as action representations for the *clamp passing* and *passage passing* tasks except DP-C, known to perform better with absolute positions [12]. We also report the performance of Tracker, which is not a baseline but a trajectory follower of downsampled ground-truth demonstrations, as a task complexity indicator.

Finally, we demonstrate our method and a baseline, D-Ma, on real-world *clamp passing* as shown in Fig. 1 and *button touch* as shown in Fig. 7 tasks using a UR5e manipulator. Unlike the simulated benchmarks, we extend the action space to 3D translation and horizontal rotation ($\in \mathbb{R}^4$) for *clamp passing* and 3D translation for *button touch*. Each task uses three RealSense cameras: two for local views and one for a fixed global view. We collect 95 human demonstrations at 2.5 Hz for *clamp passing* and a mixture of 2.5 Hz and 5 Hz for *button touch*, and train both methods on these coarse demonstrations. We compare two methods in 10 random environmental setup for each task. For real-time control, we use the denoising diffusion implicit model (DDIM) [33].

VI. EVALUATION

We first evaluate coarse-to-fine IL performance across three benchmarks using demonstrations at various frequencies. As shown in Fig. 8, DiSPo consistently achieves the highest success rates of over 81% across all frequencies, whereas baseline performances significantly drop given 2.5 Hz and 5 Hz demonstrations. For example, baseline methods usually fail at the corner of *passage passing* where DiSPo generates sharp motion as shown in Fig. 6. Occasionally, DiSPo without fine-tuning underperforms compared to baselines at tasks with fine-grained demonstrations, since the tasks are still solvable with low-frequencies demonstrations as the Tracker shows 100% performances. Nevertheless, our fine-tuning method improves performance by up to 19%,

with an average gain of 6%, without additional data collection. In contrast, the Tracker and baseline performances drop to near zero at 2.5 Hz, failing to reproduce abrupt corner maneuvering. These results highlight DiSPo’s data efficiency and its ability to accurately learn feature spaces from coarse datasets.

TABLE I: Task success rates (%) under a mixed-frequency (2.5 and 5 Hz) training dataset in the *button touch* task.

Method	DP-C	DP-T	D-Ma	VQ-BeT	DiSPo
Success Rate	81	81	83	54	93




Fig. 9: Comparison of DiSPo using fixed versus data-driven step-scaling factors from the predictor ϕ_r . We normalize the number of action steps taken in successful cases by the maximum step limit allowed in the task.

We also evaluate *multi-granularity learning* by training methods with demonstrations at mixed frequencies, 2.5 Hz and 5 Hz. As shown in Table I, DiSPo achieves the highest success rate of 93% on the *button touch* task, outperforming all baselines. DiSPo distinguishes sample-wise frequency differences, enabling effective *multi-granularity learning* without performance degradation. In contrast, baselines naively model heterogeneous frequency of state-action pairs, producing actions at an inappropriate speed that cause repetitive small back-and-forth motions near the button. Using DiSPo trained on 5 Hz demonstrations, we further evaluate the *multi-granularity reproduction* capability of DiSPo applying adaptive step scaling guided by the proposed predictor ϕ_r . As shown in Fig. 9, adopted scaling




Fig. 10: Real-world experiments result, Fig. 11: Representative samples showing the UR5e manipulator performing *clamp passing* and *button touch* in a real-world environment.

reduces the number of required steps by 32% with only a minor drop in task success rate on the *button touch* task, still significantly outperforming all baselines. These results demonstrate that DiSPo effectively modulates action discretization levels online, producing coarse motions in less critical regions (e.g., free space) to reduce inference overhead while maintaining fine-grained control in critical areas.

Finally, we evaluate DiSPo and D-Ma on a UR5e manipulator in two real-world tasks: *clamp passing* and *button touch*. Fig. 10 shows DiSPo achieves higher success rates than D-Ma in both setups. As shown in Fig. 11, DiSPo successfully inserts a square ring clamp with radial clearance 2.5 mm from random initial positions and precisely touches the shutter button by generating fine-grained, collision-free actions. While D-Ma captures rough motions well, it often causes pipe scratching or stops near the button.

VII. CONCLUSION

We proposed DiSPo, a diffusion-SSM based policy that learns and reproduces *multi-granularity* motions by adaptively modulating the action step size with a step-scaling factor. By integrating pseudo-demonstration generation and step-scale prediction, we show that DiSPo reduces storage and computational overhead. On new coarse-to-fine benchmarks, our experimental results demonstrate that DiSPo yields smoother and more accurate motions compared to state-of-the-art methods. Finally, we validate the applicability and superiority of DiSPo through real-world experiments.

VIII. RELATED WORK

Learning from demonstration (LfD). LfD is a field of methodologies that learn state-action mapping (i.e., policy) from a set of example behaviors [34]. With the advancements of neural network technology, transformer [35] architectures have enabled the encoding of multiple skills with multi-modal inputs [16], [17], [18]. To facilitate multi-modal continuous action predictions, researchers have developed categorization-based behavior transformer [8], [9], [10] and diffusion-based behavior policy [11], [12], [13], [14]. Diffusion policy [12], a representative diffusion model, employs a convolution neural network or transformer backbones for predicting action sequences.

State space models. Traditional SSMs are to represent a dynamics system with state variables. Adopting the model in deep learning, Gu et al. introduce deep SSMs memorizing continuous-time signals [36] followed by S4 architecture [21] showing fast inference and linear scalability to long sequences. As a recent variant, Mamba introduces input-dependent adaptability and has gained widespread adoption across diverse research areas by leveraging linear computational complexity and long-range dependency modeling [7], [37], [38]. In robotics, Liu et al. integrate the Mamba language model with a vision encoder to construct a reasoning and manipulation model [39].

Diffusion Models with SSMs. Diffusion models typically depend on Transformer backbones for iterative denoising [27]. Recent works integrate diffusion models with SSMs to improve efficiency in multiple domains, including time series data [40], [41], vision [42], [43], and motion generation [44]. Despite advances in other domains, SSM-driven diffusion remains underexplored in robotics, with prior work focusing mainly on accuracy and efficiency [23], [45]. In this work, we instead focus on the discretization feature of SSMs to generate fine-grained motion without additional training.

Coarse-to-fine strategy in robot learning. Most coarse-to-fine manipulation strategies produce actions by refining observations or altering action representations. For observations, researchers either iteratively zoom in on regions of interest to obtain precise data, such as point clouds [46], [47], or integrate coarse global observations with fine-grained local details [48], [49]. For actions, they design hierarchical control approaches, alternating between coarse-action generation (e.g., waypoints) and fine-action policies [1], [50], [51]. However, these works primarily aim to improve efficiency when learning from fine demonstrations. Instead, we present a unified policy that handles both coarse and fine actions while enabling learning from coarse demonstrations.

REFERENCES

- [1] E. Johns, “Coarse-to-fine imitation learning: Robot manipulation from a single demonstration,” in *Proc. Int’l Conf. on Robotics and Automation*, 2021.
- [2] W. Lian, T. Kelch, D. Holz, A. Norton, and S. Schaal, “Benchmarking off-the-shelf solutions to robotic assembly tasks,” in *Proc. RSJ Int’l Conf. on Intelligent Robots and Systems*, 2021.

- [3] X. Li, M. Baum, and O. Brock, "Augmentation enables one-shot generalization in learning from demonstration for contact-rich manipulation," in *Proc. RSJ Int'l Conf. on Intelligent Robots and Systems*, 2023.
- [4] G. Papagiannis and E. Johns, "MILES: Making imitation learning easy with self-supervision," in *Proc. Conf. on robot learning*, 2025.
- [5] A. J. Ijspeert, J. Nakanishi, and S. Schaal, "Movement imitation with nonlinear dynamical systems in humanoid robots," in *Proc. Int'l Conf. on Robotics and Automation*, 2002.
- [6] C. Finn and S. Levine, "Deep visual foresight for planning robot motion," in *Proc. Int'l Conf. on Robotics and Automation*, 2017.
- [7] A. Gu and T. Dao, "Mamba: Linear-time sequence modeling with selective state spaces," *arXiv preprint arXiv:2312.00752*, 2023.
- [8] N. M. Shafiqullah, Z. Cui, A. A. Altanzaya, and L. Pinto, "Behavior transformers: Cloning k modes with one stone," in *Conf. on Neural Information Processing Systems*, 2022.
- [9] S. Lee, Y. Wang, H. Etukuru, H. J. Kim, N. M. M. Shafiqullah, and L. Pinto, "Behavior generation with latent actions," in *Proc. Int'l Conf. on Machine Learning*, 2024.
- [10] Z. Gong, P. Ding, S. Lyu, S. Huang, M. Sun, W. Zhao, Z. Fan, and D. Wang, "CARP: Visuomotor policy learning via coarse-to-fine autoregressive prediction," in *Proc. Int'l Conf. on Computer Vision*, 2025.
- [11] M. Janner, Y. Du, J. Tenenbaum, and S. Levine, "Planning with diffusion for flexible behavior synthesis," in *Proc. Int'l Conf. on Machine Learning*, 2022.
- [12] C. Chi, Z. Xu, S. Feng, E. Cousineau, Y. Du, B. Burchfiel, R. Tedrake, and S. Song, "Diffusion policy: Visuomotor policy learning via action diffusion," *Int'l J. of Robotics Research*, 2023.
- [13] U. A. Mishra, S. Xue, Y. Chen, and D. Xu, "Generative skill chaining: Long-horizon skill planning with diffusion models," in *Proc. Conf. on robot learning*, 2023.
- [14] D. Ghosh, H. Walke, K. Pertsch, K. Black, O. Mees, S. Dasari, J. Hejna, T. Kreiman, and C. e. a. Xu, "Octo: An open-source generalist robot policy," in *Proc. Robotics: Science and Systems*, 2024.
- [15] A. O'Neill, A. Rehman, A. Maddukuri, A. Gupta, A. Padalkar, A. Lee, A. Pooley, A. Gupta, A. Mandlekar, and A. e. a. Jain, "Open X-Embodiment: Robotic learning datasets and rt-x models," in *Proc. Int'l Conf. on Robotics and Automation*, 2024.
- [16] A. Brohan, N. Brown, and J. C. et al., "RT-1: Robotics transformer for real-world control at scale," in *Proc. Robotics: Science and Systems*, 2023.
- [17] B. Zitkovich, T. Yu, S. Xu, and P. e. a. Xu, "RT-2: Vision-language-action models transfer web knowledge to robotic control," in *Proc. Conf. on robot learning*, 2023.
- [18] M. Shridhar, L. Manuelli, and D. Fox, "Perceiver-Actor: A multi-task transformer for robotic manipulation," in *Proc. Conf. on robot learning*, 2023.
- [19] C. Wang, L. Fan, J. Sun, R. Zhang, L. Fei-Fei, D. Xu, Y. Zhu, and A. Anandkumar, "MimicPlay: Long-horizon imitation learning by watching human play," in *Proc. Conf. on robot learning*, 2023.
- [20] A. Khazatsky, K. Pertsch, S. Nair, A. Balakrishna, S. Dasari, S. Karamcheti, S. Nasiriany, M. K. Srirama, L. Y. Chen, and K. e. a. Ellis, "DROID: A large-scale in-the-wild robot manipulation dataset," in *Proc. Robotics: Science and Systems*, 2024.
- [21] A. Gu, K. Goel, and C. Re, "Efficiently modeling long sequences with structured state spaces," in *Proc. Int'l Conf. on Learning Representations*, 2022.
- [22] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in *Conf. on Neural Information Processing Systems*, 2020.
- [23] X. Jia, Q. Wang, A. Donat, B. Xing, G. Li, H. Zhou, O. Celik, D. Blessing, R. Lioutikov, and G. Neumann, "MaLL: Improving imitation learning with selective state space models," in *Proc. Conf. on robot learning*, 2025.
- [24] A. Mondal, S. Alletto, and D. Tome, "HumMUSS: Human motion understanding using state space models," in *Proc. Conf. on Computer Vision and Pattern Recognition*, 2024.
- [25] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. Conf. on Computer Vision and Pattern Recognition*, 2016.
- [26] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, and J. e. a. Clark, "Learning transferable visual models from natural language supervision," in *Proc. Int'l Conf. on Machine Learning*, 2021.
- [27] W. Peebles and S. Xie, "Scalable diffusion models with transformers," in *Proc. Int'l. Conf. on Computer Vision*, 2023.
- [28] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, "Learning fine-grained bimanual manipulation with low-cost hardware," in *Proc. Robotics: Science and Systems*, 2023.
- [29] E. Coumans and Y. Bai, "PyBullet, a python module for physics simulation for games, robotics and machine learning." <http://pybullet.org>, 2016.
- [30] V. Blomqvist, "Pymunk: A easy-to-use pythonic rigid body 2d physics library." <https://www.pymunk.org>, 2007.
- [31] "pygame." <https://www.pygame.org/>, 2007.
- [32] H. Pham and Q.-C. Pham, "A new approach to time-optimal path parameterization based on reachability analysis," *Trans. on Robotics*, 2018.
- [33] J. Song, C. Meng, and S. Ermon, "Denoising diffusion implicit models," in *Proc. Int'l. Conf. on Learning Representations*, 2021.
- [34] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robotics and Autonomous Systems*, 2009.
- [35] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," *Conf. on Neural Information Processing Systems*, 2017.
- [36] A. Gu, T. Dao, S. Ermon, A. Rudra, and C. Ré, "Hippo: Recurrent memory with optimal polynomial projections," *Conf. on Neural Information Processing Systems*, 2020.
- [37] T. Dao and A. Gu, "Transformers are SSMs: Generalized models and efficient algorithms through structured state space duality," in *Proc. Int'l Conf. on Machine Learning*, 2024.
- [38] H. Qu, L. Ning, R. An, W. Fan, T. Derr, H. Liu, X. Xu, and Q. Li, "A survey of mamba," *arXiv preprint arXiv:2408.01129*, 2024.
- [39] J. Liu, M. Liu, Z. Wang, P. An, X. Li, K. Zhou, S. Yang, R. Zhang, Y. Guo, and S. Zhang, "RoboMamba: Efficient vision-language-action model for robotic reasoning and manipulation," *Conf. on Neural Information Processing Systems*, 2025.
- [40] K. Goel, A. Gu, C. Donahue, and C. Ré, "It's raw! audio generation with state-space models," in *Proc. Int'l Conf. on Machine Learning*, 2022.
- [41] J. L. Alcaraz and N. Strodthoff, "Diffusion-based time series imputation and forecasting with structured state space models," *Trans. on Machine Learning Research*, 2022.
- [42] J. N. Yan, J. Gu, and A. M. Rush, "Diffusion models without attention," in *Proc. Conf. on Computer Vision and Pattern Recognition*, 2024.
- [43] V. T. Hu, S. A. Baumann, M. Gui, O. Grebenkova, P. Ma, J. Schusterbauer, and B. Ommer, "ZigMa: A dit-style zigzag mamba diffusion model," in *Proc. European Conf. on Computer Vision*, 2024.
- [44] Z. Zhang, A. Liu, I. Reid, R. Hartley, B. Zhuang, and H. Tang, "Motion mamba: Efficient and long sequence motion generation with hierarchical and bidirectional selective ssm," in *Proc. European Conf. on Computer Vision*, 2024.
- [45] J. Cao, Q. Zhang, J. Sun, J. Wang, H. Cheng, Y. Li, J. Ma, Y. Shao, W. Zhao, and G. e. a. Han, "Mamba policy: Towards efficient 3d diffusion policy with hybrid selective state models," in *Proc. RSJ Int'l Conf. on Intelligent Robots and Systems*, 2025.
- [46] S. James, K. Wada, T. Laidlow, and A. J. Davison, "Coarse-to-fine q-attention: Efficient learning for visual robotic manipulation via discretisation," in *Proc. Conf. on Computer Vision and Pattern Recognition*, 2022.
- [47] T. Gervet, Z. Xian, N. Gkanatsios, and K. Fragkiadaki, "Act3d: 3d feature field transformers for multi-task robotic manipulation," in *Proc. Conf. on robot learning*, 2023.
- [48] S. Ling, Y. Wang, R. Wu, S. Wu, Y. Zhuang, T. Xu, Y. Li, C. Liu, and H. Dong, "Articulated object manipulation with coarse-to-fine affordance for mitigating the effect of point cloud noise," in *Proc. Int'l Conf. on Robotics and Automation*, 2024.
- [49] P. Sundaresan, H. Hu, Q. Vuong, J. Bohg, and D. Sadigh, "What's the move? hybrid imitation learning via salient points," in *Proc. Int'l. Conf. on Learning Representations*, 2025.
- [50] E. Valassakis, N. Di Palo, and E. Johns, "Coarse-to-fine for sim-to-real: Sub-millimetre precision across wide task spaces," in *Proc. RSJ Int'l Conf. on Intelligent Robots and Systems*, 2021.
- [51] S. Belkhale, Y. Cui, and D. Sadigh, "Hydra: Hybrid robot actions for imitation learning," in *Proc. Conf. on robot learning*, 2023.