

MP1: MeanFlow Tames Policy Learning in 1-step for Robotic Manipulation

Juyi Sheng^{1,2*}, Ziyi Wang^{1,2*}, Peiming Li^{1,2}, Mengyuan Liu^{1,2†}

¹ State Key Laboratory of General Artificial Intelligence, Peking University, Shenzhen Graduate School

² Guangdong Provincial Key Laboratory of Ultra High Definition Immersive Media Technology, Peking University, Shenzhen Graduate School

Abstract

In robot manipulation, robot learning has become a prevailing approach. However, generative models within this field face a fundamental trade-off between the slow, iterative sampling of diffusion models and the architectural constraints of faster Flow-based methods, which often rely on explicit consistency losses. To address these limitations, we introduce MP1, which pairs 3D point-cloud inputs with the MeanFlow paradigm to generate action trajectories in one network function evaluation (1-NFE). By directly learning the interval-averaged velocity via the “MeanFlow Identity”, our policy avoids any additional consistency constraints. This formulation eliminates numerical ODE-solver errors during inference, yielding more precise trajectories. MP1 further incorporates CFG for improved trajectory controllability while retaining 1-NFE inference without reintroducing structural constraints. Because subtle scene-context variations are critical for robot learning, especially in few-shot learning, we introduce a lightweight *Dispersive Loss* that repels state embeddings during training, boosting generalization without slowing inference. We validate our method on the Adroit and Meta-World benchmarks, as well as in real-world scenarios. Experimental results show MP1 achieves superior average task success rates, outperforming DP3 by 10.2% and FlowPolicy by 7.3%. Its average inference time is only 6.8 ms—19× faster than DP3 and nearly 2× faster than FlowPolicy.

Code: — <https://github.com/LogSSim/MP1.git>

Introduction

Robot learning enables robots to act from visual or text inputs to perform physical tasks like grasping and moving objects. Recent advancements in robot learning have been driven by methods like Transformers (Vaswani et al. 2017), Diffusion Models (Ho, Jain, and Abbeel 2020), and Flow Matching (Lipman et al. 2023). These methods have enhanced the ability of robots to understand and execute complex actions in response to multimodal inputs (Chi et al. 2023, 2024), facilitating more versatile robot learning.

Among them, Diffusion Model-based methods, such as the diffusion policy (Chi et al. 2023, 2024), effectively handle multimodal action distributions by representing robot

*These authors contributed equally.

†Corresponding authors: liumengyuan@pku.edu.cn.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: The proposed method outperforms SOTA methods (DP3 (Ze et al. 2024) and FlowPolicy (Zhang et al. 2024)) on the Adroit and Meta-World tasks, showing superior inference time and success rate, as demonstrated by the MP1 on the comparison plot.

action predictions as probability distributions. Additionally, the DP3 (Ze et al. 2024) introduces 3D point cloud features combined with the diffusion policy, improving the success rate of tasks and reducing the amount of training data. Through a step-by-step denoising process, Diffusion Models can capture multiple possible action choices, thereby increasing the flexibility and accuracy of action generation. However, a notable drawback of Diffusion Models is their relatively long inference time. Since action generation requires multiple time steps to denoise, the inference process can be time-consuming, which may become a bottleneck in applications that demand real-time performance.

In recent years, Flow-based methods have been proposed to overcome this problem, such as AdaFlow (Hu et al. 2024), FlowPolicy (Zhang et al. 2024), which aims to reduce sampling steps and achieve more efficient single-step inference.

These methods greatly accelerate inference by optimizing the generative process, but they typically require additional consistency constraints on the model’s outputs to ensure valid trajectories. In contrast, the recently introduced MeanFlow (Geng et al. 2025) paradigm from image generation avoids any explicit consistency loss by learning a mean velocity field. This innovation simplifies the generation process and achieves genuine one-step sampling, markedly improving real-time performance. For example, MeanFlow can reduce inference latency to around 6.8 ms, far better than the 10–20 steps required by diffusion strategies.

We present the first adaptation of the MeanFlow (Geng et al. 2025) paradigm to robot learning, termed MP1. Conditioned on 3D point-cloud observations, MP1 learns the interval-averaged velocity and bypasses the need to integrate instantaneous velocities. This design eliminates ODE-solver error and yields genuinely one-step (1-NFE) trajectory generation with smooth, dynamically consistent actions. However, a purely regression-based objective fails to impose explicit regularization on the policy’s internal feature space (Wang and He 2025). This representational ambiguity is particularly detrimental in robotics, where subtle variations in scene context are critical, and it undermines generalization in few-shot learning.

To counter this, we add Dispersive Loss, spreading out the latent embeddings of distinct states. Acting as a contrastive-style regularizer without positive pairs, it sharpens state discrimination while the original regression term still aligns each state to its expert trajectory. With only ten expert demonstrations, this combination already closes most performance gaps (Fig. 5), underscoring the method’s few-shot generalization. Because Dispersive Loss is computed once per forward pass and vanishes at inference, MP1 preserves its hallmark 1-NFE speed. We conduct experiments on the Adroit and Meta-World simulation datasets and real-world tasks. MP1 is capable of one-step inference and, compared to state-of-the-art (SOTA) methods, improves the average success rate by 7.3% (Tab. 1) while also achieving superior inference speed (Tab. 2). Our contributions are as follows:

- We introduce MP1, the first MeanFlow-based robot learning framework. Conditioning on 3D point-cloud features, it learns effectively from a handful of demonstrations, yet delivers one-step sampling with SOTA success rates and millisecond-level inference latency.
- We incorporate a lightweight Dispersive Loss that regularizes latent features without affecting runtime, boosting few-shot generalization.
- Extensive experiments on Adroit and Meta-World simulation benchmarks, as well as real-world tasks, show that MP1 surpasses Diffusion- and Flow-based baselines in both success rate and speed.

Related Work

2D Input Robot Learning

Most methods that utilize 2D visual input predict robot actions based on images. For example, BC-Z (Jang et al. 2022) aligned 2D visual-language features, allowing the robot to

generalize to new target tasks. ALOHA (Zhao et al. 2023) employed an ACT network to model the relationship between 2D vision and robotic actions. Similarly, other approaches, such as DP (Chi et al. 2023, 2024), HPT (Wang et al. 2024b), also leveraged 2D inputs for robot learning policies. However, 2D inputs often lack depth information, which limits the accuracy in completing tasks.

3D Input Robot Learning

To overcome the limitations of 2D inputs, 3D inputs have gained prominence. Approaches such as PerACT (Shridhar, Manuelli, and Fox 2022) and ACT3D (Gervet et al. 2023) utilized voxel data. However, due to the high computational cost associated with voxel data, point clouds have become the dominant form of 3D input. Methods like RVT (Goyal et al. 2023), RVT2 (Goyal et al. 2024), GPA-RAM (Sheng et al. 2025) and DP3 (Ze et al. 2024) leveraged point clouds, and their introduction has been shown to enhance task success rates.

Diffusion-Based Robot Learning

Diffusion models have recently made significant progress in image and action generation. This model simulates the diffusion process of data by gradually adding noise and then denoising to generate data. In robot learning, diffusion models are used to generate continuous control actions and strategies to tackle complex tasks. Initially, diffusion policy (Chi et al. 2023, 2024) addressed the issues of multimodal action generation and action consistency. Later, DP3 (Ze et al. 2024) introduced point cloud information, greatly improving task success rates. HDP (Ma et al. 2024b) established spatial relationships between task points in point clouds using diffusion. RDT (Liu et al. 2024) built a large-scale robot dual-arm model based on diffusion. Additionally, models like EquiDiff (Wang et al. 2024a) and Instant Policy (Vosylius and Johns 2025) are also based on diffusion. However, diffusion still faces challenges related to inference time.

Flow-Based Robot Learning

Flow matching is a method for training continuous normalized flows by learning vector fields associated with probabilistic paths. Recent research has shown that Flow matching performs well in image generation (Ma et al. 2024a; Esser et al. 2024; Polyak, Zohar, and Andrew Brown 2025), and Consistency Models (Song et al. 2023; Song and Dhariwal 2023) further improve the sampling efficiency of Flow methods, enabling one-step sampling. Consequently, AdaFlow (Hu et al. 2024) was introduced as a robot learning framework using flow-based generative modeling to represent the policy via state-conditioned ordinary differential equations. Building on this approach, FlowPolicy (Zhang et al. 2024) adapted it to robot learning, achieving one-step sampling by enforcing input consistency constraints.

3D point clouds are particularly well-suited for robot learning (Gervet et al. 2023; Ze et al. 2024), which is why this paper utilizes them as input. Methods based on diffusion and flow focus on understanding the environment and generating actions. While diffusion-based methods effectively address these challenges, they introduce significant delays in

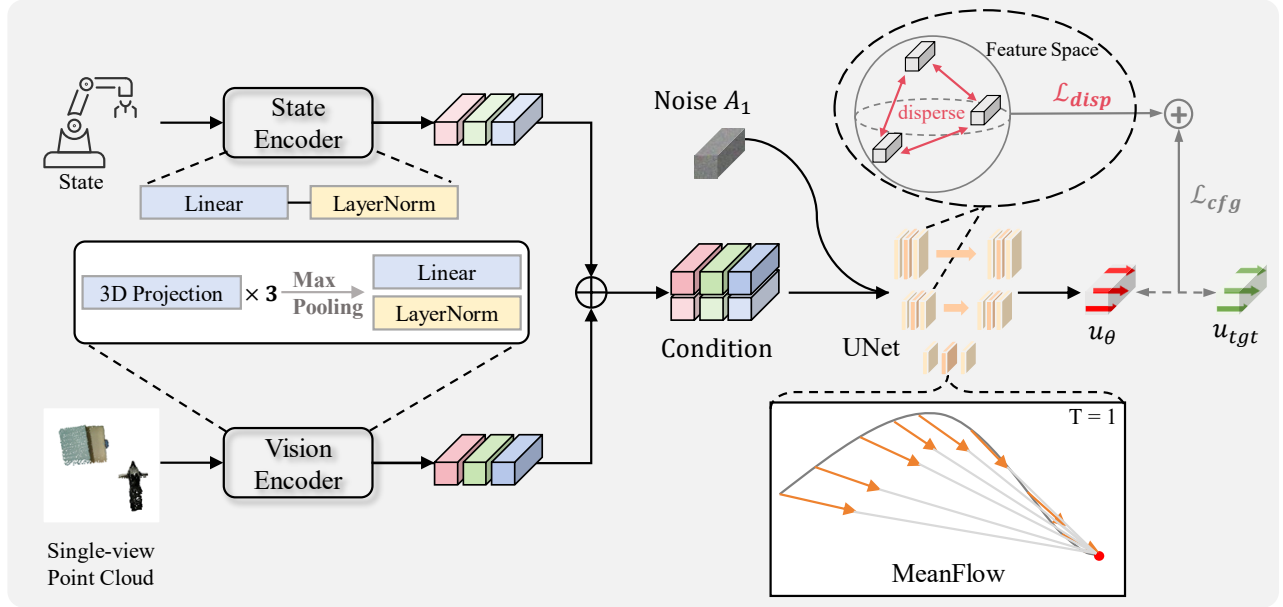


Figure 2: Overview of MP1. The MP1 takes the historical observation point cloud and the robot’s state as inputs. These inputs are processed through a visual encoder and a state encoder, respectively, and then serve as conditional inputs to the UNet-integrated MeanFlow. After passing through the MeanFlow, the model computes regression loss (\mathcal{L}_{cfg}) between the mean velocity generated from the initial noise and the target velocity. This \mathcal{L}_{cfg} is combined with a Dispersive Loss (\mathcal{L}_{disp}) imposed on the UNet’s hidden states to jointly optimize the network parameters.

inference time due to the need for multiple sampling steps. On the other hand, Flowpolicy achieves one-step sampling by applying consistency constraints to the input, but it relies on certain assumptions. The MP1 method proposed in this paper is derived from the definitions of average and instantaneous velocities. This approach eliminates the need for assumption-based constraints, enabling efficient one-step inference and the generation of high-quality robot actions.

Method

Our goal is to develop an effective and highly generalized robotic learning policy. Efficiency, particularly single-step (1-NFE) inference, is crucial for real-time applications, while generalization is essential for leveraging limited demonstrations and adapting to subtle environmental variations. To address these challenges, we propose the MP1 (Fig. 2). Unlike Diffusion-based methods, our approach does not require multi-step denoising; distinct from existing Flow-based approaches, the MP1 does not rely on ODE solvers, consistency constraints, and the integration of interval-wise instantaneous velocities. Furthermore, by encouraging the latent embeddings of different input states to disperse, we improve the model’s generalization abilities and task success rate, all without sacrificing inference speed.

MP1: One-Step Trajectory Generation

In the context of robot learning, the policy’s task is to map a sequence of observations, including 3D point clouds \mathbf{P} and robotic states \mathbf{S} , to a future action trajectory \mathbf{A} . We first process these inputs using encoders to obtain conditional fea-

tures. The point cloud $\mathbf{P} \in \mathbb{R}^{n_o \times np \times 3}$ is passed through 3D Projection to produce a visual feature vector \mathbf{f}_v , while the robot state $\mathbf{S} \in \mathbb{R}^{n_o \times s_a}$ is encoded into a state feature \mathbf{f}_s . These are combined into a single conditional vector $\mathbf{c} = (\mathbf{f}_v, \mathbf{f}_s)$ that guides the action generation process.

To achieve single-step generation, we model the policy as a conditional MeanFlow. Unlike standard Flow Matching (FM), which learns an instantaneous velocity field $v(z_t, t)$ and requires solving an ODE for sampling, MeanFlow learns the average velocity field $u(z_t, r, t)$ over an interval $[r, t]$. The average velocity is defined as the total displacement divided by the time duration:

$$u(z_t, r, t) \triangleq \frac{1}{t-r} \int_r^t v(z_\tau, \tau) d\tau \quad (1)$$

This formulation is key to bypassing iterative integration. Training a network u_θ to model this field directly from its integral definition is intractable. Instead, MeanFlow leverages the “MeanFlow Identity”, a local differential equation derived by differentiating the integral definition to t :

$$u(z_t, r, t) = v(z_t, t) - (t-r) \frac{d}{dt} u(z_t, r, t) \quad (2)$$

where the total derivative $\frac{d}{dt} u(z_t, r, t)$ expands to $v(z_t, t) \partial_z u + \partial_t u$. This allows us to train u_θ with a simple regression objective that enforces this identity:

$$\mathcal{L}(\theta) = \mathbb{E}_{t,r,x,\epsilon} \|u_\theta(z_t, r, t) - sg(u_{tgt})\|_2^2 \quad (3)$$

The target u_{tgt} is constructed using the known instantaneous velocity v_t of the probability path and the network’s own

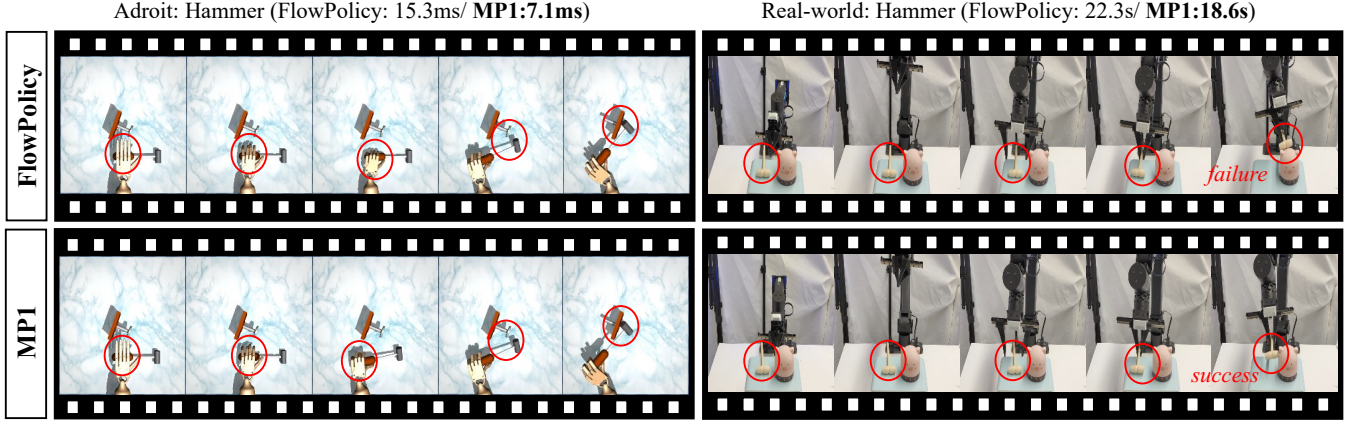


Figure 3: Qualitative comparison of the proposed MP1 and the previous SOTA method (FlowPolicy (Zhang et al. 2024)) on Adroit Hammer and real-world Hammer tasks. Our method is faster, with 7.1ms in the simulated hammer and 18.6s in the real-world scenario. Moreover, our method successfully completes the real-world hammer task, whereas FlowPolicy fails.

estimate of the total derivative, with a stop-gradient $sg(\cdot)$ to ensure stability:

$$u_{tgt} = v_t - (t - r)(v_t \partial_z u_\theta + \partial_t u_\theta) \quad (4)$$

For our MP1, we adapt this framework to generate action trajectories \mathbf{A} . The network $u_\theta(\mathbf{A}_t, r, t|\mathbf{c})$ is trained to predict the average velocity that transforms a noise vector $\mathbf{A}_1 \sim \mathcal{N}(0, I)$ into an expert action trajectory \mathbf{A}_0 . To improve control, we integrate Classifier-Free Guidance (CFG) by training the network to model a CFG-aware average velocity, u_θ^{cfg} . This is achieved with a modified regression loss based on a guided instantaneous velocity field:

$$\mathcal{L}_{cfg}(\theta) = \mathbb{E} \|u_\theta^{cfg}(\mathbf{A}_t, r, t|\mathbf{c}) - sg(u_{tgt})\|_2^2 \quad (5)$$

where the target u_{tgt} is now computed using a guided velocity $\tilde{v}_t \triangleq \omega v_t(\mathbf{A}_t|\mathbf{A}_0, \mathbf{c}) + (1 - \omega)u_\theta^{cfg}(\mathbf{A}_t, t, t|\theta)$, blending the conditional and unconditional predictions.

Enhancing Representational generalization with Dispersive Loss

While the MeanFlow objective \mathcal{L}_{cfg} excels at learning the temporal dynamics required to produce accurate output trajectories, it provides only an indirect signal for learning the complex mapping from high-dimensional conditional inputs \mathbf{c} to actions. This can lead to a form of “feature collapse”, where the policy network maps distinct environmental states that demand fundamentally different actions to nearly identical points in its latent space. Such ambiguity is particularly detrimental in robot learning, where subtle differences in object pose or scene configuration are critical for success, especially in few-shot learning regimes.

To directly address this, we incorporate Dispersive Loss, a principled and self-contained regularizer that operates on the model’s internal representations. The core idea is to encourage the latent embeddings of different input samples within a training batch to disperse, thereby enforcing a more discriminative feature space. Dispersive Loss functions as a “contrastive loss without positive pairs”; the repulsive force is

supplied by the loss itself, while the alignment of a state to its correct action is handled implicitly by the primary \mathcal{L}_{cfg} regression objective. Let $\mathbf{z}_\mathbf{A} = f_{policy}(\mathbf{A}_t, t, \mathbf{c})$ be the intermediate representation from a chosen layer within our policy network. The Dispersive Loss is defined as:

$$\mathcal{L}_{Disp}(\theta) = \log \mathbb{E}_{i,j \in \mathcal{B}} \left[\exp \left(-\frac{\|\mathbf{z}_{\mathbf{A},i} - \mathbf{z}_{\mathbf{A},j}\|_2^2}{\tau} \right) \right], \quad (6)$$

where \mathcal{B} is a mini-batch of training samples, $\mathbf{z}_{\mathbf{A},i}$ and $\mathbf{z}_{\mathbf{A},j}$ are the intermediate representations for two samples in the batch, and τ is a temperature hyperparameter. Specifically, we adopt the InfoNCE-based variant of Dispersive Loss (using ℓ_2 distance, with temperature $\tau = 1$) and apply it to the output features of each down-sampling block in the U-Net backbone of u_θ .

Training Objective and Inference

Our final training objective synergistically combines the trajectory generation and representation regularization goals:

$$\mathcal{L}_{total}(\theta) = \mathcal{L}_{cfg}(\theta) + \lambda \mathcal{L}_{Disp}(\theta) \quad (7)$$

Here, \mathcal{L}_{cfg} (Eq. 5) ensures the policy generates dynamically correct action trajectories, while \mathcal{L}_{Disp} (Eq. 6) promotes a well-structured and discriminative latent space to improve generalization. When λ is set to 0.5, it balances the contribution of the two loss terms.

Crucially, \mathcal{L}_{Disp} , as a training-time regularizer, adds no computational overhead during inference, preserving the policy’s efficiency. Given the current conditional features \mathbf{c} and a random noise sample $\mathbf{A}_1 \sim \mathcal{N}(0, I)$, the MP1 generates the entire K-step action \mathbf{A}_0 in a single forward pass:

$$\mathbf{A}_0 = \mathbf{A}_1 - u_\theta^{cfg}(\mathbf{A}_1, 0, 1|\mathbf{c}) \quad (8)$$

where $r = 0, t = 1$. This preserves the 1-NFE capability that is vital for real-time robotic control, while benefiting from the enhanced generalization conferred by a more structured and discriminative representational foundation.

Methods	Publication	NFE	Meta-World							Average
			Hammer	Adroit Door	Pen	Easy (21)	Medium (4)	Hard (4)	Very Hard (5)	
DP	RSS'23	10	16±10	34±11	13±2	50.7±6.1	11.0±2.5	5.25±2.5	22.0±5.0	35.2±5.3
Adaflow	NeuRIPS'24	-	45±11	27±6	18±6	49.4±6.8	12.0±5.0	5.75±4.0	24.0±4.8	35.6±6.1
CP	arxiv'24	1	45±4	31±10	13±6	69.3±4.2	21.2±6.0	17.5±3.9	30.0±4.9	50.1±4.7
DP3	RSS'24	10	100±0	56±5	46±10	87.3±2.2	44.5±8.7	32.7±7.7	39.4±9.0	68.7±4.7
Simple DP3	RSS'24	10	98±2	40±17	36±4	86.8±2.3	42.0±6.5	38.7±7.5	35.0±11.6	67.4±5.0
FlowPolicy	AAAI'25	1	98±1	61±2	54±4	84.8±2.2	58.2±7.9	40.2±4.5	52.2±5.0	71.6±3.5
MP1	-	1	100±0	69±2	58±5	88.2±1.1	68.0±3.1	58.1±5.0	67.2±2.7	78.9±2.1

Table 1: Performance of different methods on 37 Tasks. We evaluate the performance of our method on 3 Adroit and 34 Meta-World tasks with three random seeds, comparing it to SOTA methods based on Diffusion and Flow. Our method with NFE=1 outperforms the best Diffusion-based method (DP3 (Ze et al. 2024)) by 10.2% in success rate. Compared to FlowPolicy (Zhang et al. 2024), which uses 1-NFE sampling, our method achieves a 7.3% higher average success rate across all tasks.

Methods	Publication	NFE	Adroit /ms			Meta-World /ms				Average /ms
			Hammer	Door	Pen	Easy (21)	Medium (4)	Hard (4)	Very Hard (5)	
DP3	RSS'24	10	129.5±13.9	141.3±14.8	145.1±12.3	129.3±10.7	134.7±11.5	131.9±12.4	138.4±10.8	132.2±11.2
Simple DP3	RSS'24	10	103.1±11.4	111.3±10.2	128.2±13.1	91.9±8.6	98.3±9.1	101.3±9.7	103.8±10.2	97.0±9.2
FlowPolicy	AAAI'25	1	15.3±1.1	13.2±4.0	12.0±2.8	12.0±1.4	12.2±1.5	13.5±1.4	14.5±1.6	12.6±1.5
MP1	-	1	7.1±0.2	7.2±0.1	7.4±0.3	6.7±0.0	6.7±0.1	6.7±0.1	6.8±0.1	6.8±0.1

Table 2: Comparison of inference times for different methods evaluated on the Meta-World and Adroit benchmark. Due to its multi-step denoising process, Diffusion-based approaches run slower than Flow-based ones. MP1 achieves SOTA inference speed across all sub-tasks, with an average latency of just 6.8 ms-nearly 2× faster than the best FlowPolicy (which relies on consistency constraints for its 1-NFE sampling) and 14× faster than Simple DP3 (Ze et al. 2024).

Simulation Experiments

To demonstrate the efficacy of the MP1, we conduct fair comparative experiments on multiple datasets (e.g., Adroit (Rajeswaran et al. 2017) and Meta-World (Yu et al. 2020)), comparing our proposed method against existing approaches in terms of success rate and inference time.

Simulation benchmark

In simulation, we evaluate the proposed method on three tasks from the Adroit benchmark. To further assess the algorithm’s generality and generalization across scenarios, we also select thirty-four tasks from the Meta-World benchmark, comprising twenty-one Easy tasks, four Medium tasks, four Hard tasks, and five Very Hard tasks.

Baselines

In our comparison with existing SOTA methods, we include DP (Chi et al. 2023, 2024), AdaFlow (Hu et al. 2024), and CP (Prasad et al. 2024) with 2D inputs, as well as DP3 (Ze et al. 2024), Simple DP3 (Ze et al. 2024), and FlowPolicy (Zhang et al. 2024) with 3D inputs. DP, DP3, and Simple DP3 all employ 10-NFE, whereas CP and FlowPolicy use 1-NFE, and AdaFlow operates with a variable NFE.

Implementation Details

We generate ten expert demonstrations for training on Adroit and Meta-World. For the point-cloud data, farthest point sampling (FPS) sampling reduces the number of points to either 512 or 1024; for the image data, resolution is down-sampled to 84 × 84 pixels. The proposed MP1 and SOTA methods are each trained and tested using three random

seeds (0, 1, and 2). For each random seed, the model is trained for 3,000 epochs on Adroit and 1,000 epochs on Meta-World, with performance evaluated every 200 epochs. From the evaluation results for each seed, the five highest success rates are selected and averaged. Finally, the overall success rate and standard deviation for the task are computed across all three seeds. All training and testing are performed on an NVIDIA RTX 4090 GPU, with a batch size of 128, optimization uses the AdamW optimizer with a learning rate of 0.0001 (Adroit and Meta-World apply the same learning rate), an observation window of 2 steps, a history length of 4 states, and a prediction horizon of 4 steps, MeanFlow ratio is 0.5 and the CFG scale is 2.0. We conduct inference speed tests for three seeds on the same GPU. First, we calculate the average speed for the three stable-state values. Then, we compute the average and standard deviation of these speeds as the experimental results.

Results of Simulation

Tab. 1 demonstrates that the MP1 achieves SOTA performance across all sub-tasks. The overall average success rate reaches 78.9%±2.1%, which significantly outperforms the previous best method, FlowPolicy, at 71.6%±3.5%. Compared to existing approaches, our method yields a 10.2% improvement over DP3 and a 7.3% improvement over FlowPolicy in terms of average success rate, indicating that mean-velocity flows are better suited for the robot learning. In the Meta-World tasks, it can be observed that the MP1 achieves a higher success rate on “Very Hard” tasks compared to its performance on “Hard” tasks. The proposed method demonstrates notable improvements across 13 tasks of Meta-World (Medium, Hard, and Very Hard), with success rates in-

Task / Method	Adroit		Meta-World							
	Door	Pen	Reach	Coffee Pull	Hand Insert	Pick Place	Push	Disassemble	Stick Pull	Pick Place Wall
MP1	58±5	69±2	24.7±3.3	92.3±3.7	10.0±2.9	50.7±9.1	74.0±6.7	74.0±1.4	74.0±1.4	64.3±1.2
-Loss _{dis}	55±6	68±4	19.7±1.2	90.7±2.1	9.3±1.7	48.7±8.2	50.7±7.6	72.7±0.5	72.0±5.0	60.3±2.4

Table 3: Ablation Study on *Dispersive Loss* for Adroit and Meta-World Tasks. -Loss_{dis} signifies that the Dispersive Loss term has been omitted.

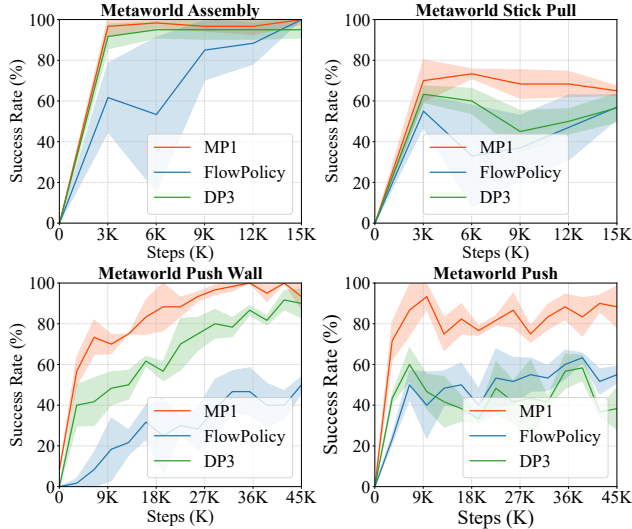


Figure 4: Success rate curves of different methods on multiple Meta-World tasks. We compare the performance of MP1, FlowPolicy, and DP3 on four tasks. The x-axis represents training steps, and the y-axis shows the success rate. Shaded areas represent the standard deviation across different random seeds. The proposed method achieves higher success rates with smaller variance.

creased by 9.8%, 17.9%, and 15.0% respectively over the FlowPolicy. On the 21 “Easy” tasks in Meta-World, the proposed approach achieves a success rate of 88.2%, representing a 3.4% improvement over the FlowPolicy. Moreover, the proposed method maintains a consistently low standard deviation on certain subtasks (for example, the standard deviation for Adroit Hammer, Door, and the Easy tasks in Meta-World ranges from only 0 to 2.0%), and the average fluctuation in success rate is merely 2.1%, which is lower than that of other methods. These results further indicate that the proposed approach exhibits greater stability and reliability under different random seeds.

Fig. 4 shows the training success rate curves of MP1, FlowPolicy, and DP3 on four Meta-World tasks: Assembly, Stick Pull, Push Wall, and Push. As the number of training steps increases, all methods demonstrate improved success rates; however, MP1 achieves faster convergence and higher final success rates across all tasks. Furthermore, the shaded area representing the standard deviation across different random seeds is generally narrower for MP1, indicating better stability and generalization.

Task / Flow Ratio	Adroit	Meta-World				Avg.
	Pen	Dial Turn	Coffee Pull	Assembly	Disassemble	
0 ($r = t$)	53±5	81±1	62±4	97±2	70±3	72.6±3.0
0.25	49±6	89±2	92±0	99±1	71±4	80.0±2.6
0.50	58±5	90±2	92±4	98±1	74±1	82.4±2.6
0.75	54±4	90±2	90±5	98±1	71±3	80.6±3.0
1.0	0±0	0±0	12±5	0±0	0±0	2.4±1.0

Table 4: Performance of different flow ratios (when $r \neq t$) in Adroit Pen and Meta-World tasks.

Fig. 1 and Tab. 2 summarize the inference speeds of various methods on the Adroit and Meta-world benchmarks. The results show that Diffusion-based approaches are much slower than Flow-based methods. The proposed method achieves SOTA performance in all subtasks, with approximately a 2× speedup over FlowPolicy and a 14× speedup over Simple DP3. On an NVIDIA 4090, our method attains an average inference time of 6.8 ms. Furthermore, MP1 is more stable than prior methods, unaffected by GPU load fluctuations, with a standard deviation of ±0.1 ms.

Overall, Flow-based methods can achieve 1-NFE, enabling much faster inference than Diffusion-based approaches (Fig. 1). Our proposed MP1 achieves SOTA inference speed (6.8 ms) and performance, with a 78.9% success rate across 37 tasks.

Ablation Study

We primarily conduct ablation experiments on the dispersive loss, flow ratio, and number of demonstrations in MP1.

Dispersive Loss

Tab. 3 compares the standard MP1 with a variant in which the Dispersive Loss is removed. Introducing the *Dispersive Loss* (see the “MP1” row) improves success rates on all ten tasks across the Adroit and Meta-World benchmarks, with average gains of roughly 4-5 percentage points. These results show that Dispersive Loss enhances policy performance in diverse manipulation scenarios.

MeanFlow Ratio

In (2), the average speed is reflected when $r \neq t$; when $r = t$ (ratio=0), it degenerates to the standard flow matching method. We extract five tasks from Adroit and Meta-World and conduct success rate tests under different flow ratios (0, 0.25, 0.50, 0.75, 1.0). As shown in Tab. 4, performance declines when it reverts to Flow Matching (ratio = 0), compared to the case where $r \neq t$. The results from Tab. 4 indi-

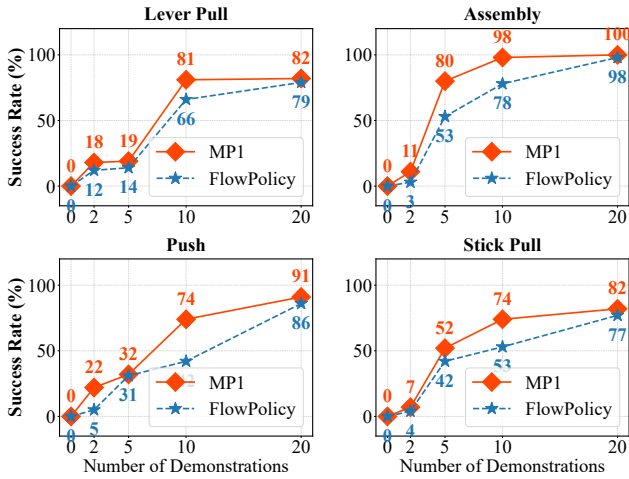


Figure 5: The effect of the number of demonstrations on different methods. As the number increases, the success rate gradually improves.

Task / Method	Hammer	Drawer Close	Heat Water	Stack Block	Spoon
MP1	90% / 18.6s	100% / 8.8s	90% / 23.4s	80% / 27.2s	90% / 22.6s
FlowPolicy	70% / 22.3s	90% / 15.7s	60% / 31.1s	50% / 29.6s	80% / 26.7s
DP3	70% / 31.1s	80% / 20.2s	70% / 38.8s	60% / 35.1s	70% / 28.3s

Table 5: Success rates (%) and per-task completion times (s) of different methods in real-world experiments.

cate that the flow ratio ($r \neq t$) has minimal impact on task success rates for most tasks, except for extreme cases (1.0 flow ratio) where the success rate drops significantly.

Number of Demonstrations

In Fig. 5, we evaluate the impact of different numbers of demonstrations (0, 2, 5, 10, 20) on task performance. As the number of demonstrations increases, the success rate improves significantly across various tasks. Our method, the MP1, consistently outperforms the FlowPolicy, especially with fewer demonstrations. These results show that increasing the number of demonstrations enhances model performance for most tasks, with the proposed method excelling in few-shot learning scenarios.

Real-world Experiments

Experimental Setup

We conduct experiments on the ARX R5 dual-arm robot, using the same camera configuration as DP3, the RealSense L515. The experimental environment is shown in Fig. 6.

Real-world Tasks

We evaluate our approach on five real-world tasks: 1) Hammer – Grasp a hammer and strike a pig. 2) Drawer Close – Closing a drawer. 3) Heat Water – Position a kettle in a suitable location. 4) Stack Block - Stack a block. 5) Spoon - Put the spoon in the bowl.




Figure 6: Real-world setup.

Real-world Experimental Results

In Fig. 3, we present the performance of MP1 and Flowpolicy on the hammer task in the simulation environment, as well as the experimental results for the hammer task in the real-world environment. In the real environment, Flowpolicy failed to complete the task, while MP1 successfully completed it with a success rate of 90% and a completion time of 18.6 seconds. MP1’s faster inference speed enables the robot to perform the task more smoothly. More examples can be found in the Project Page: <https://mp1-2254.github.io/>.

Tab. 5 reports the performance of different methods in real-world robotic experiments, measured by success rate (%) and average task completion time (s). Each task was trained and evaluated using 20 human demonstrations under a unified setting of $horizon = 16$ and observation stride $n_{obs} = 2$. MP1 achieves SOTA performance in both success rate and task completion time across all five tasks, with an average success rate of 90% and an average completion time of 20.1s. In comparison, FlowPolicy has an average success rate of 70% (25.1s), and DP3 has 70% (30.7s). This demonstrates the effectiveness of MP1 in real-world scenarios.

Conclusion

In this paper, we address the limitations of existing Diffusion-based and Flow-based approaches by introducing MeanFlow into robot learning. The proposed MP1 eliminates the reliance on multi-step denoising, interval-wise instantaneous velocity integration, and consistency constraints. MP1 enables one-step action generation for robot manipulation by modeling the velocity field, making it suitable for real-time control. Furthermore, the introduction of Dispersive Loss helps disperse the latent embeddings of distinct states, improving generalization under low-data conditions. Experimental results demonstrate that MP1 delivers significant performance gains across various robotic tasks, confirming its effectiveness and wide applicability.

Acknowledgments

This work was supported by National Natural Science Foundation of China (No. 62473007), Guangdong Provincial

Key Laboratory of Ultra High Definition Immersive Media Technology (No. 2024B1212010006), Shenzhen Innovation in Science and Technology Foundation for The Excellent Youth Scholars (No. RCYX20231211090248064).

References

- Chi, C.; Feng, S.; Du, Y.; Xu, Z.; Cousineau, E.; Burchfiel, B.; and Song, S. 2023. Diffusion Policy: Visuomotor Policy Learning via Action Diffusion. In *Proceedings of Robotics: Science and Systems (RSS)*.
- Chi, C.; Xu, Z.; Feng, S.; Cousineau, E.; Du, Y.; Burchfiel, B.; Tedrake, R.; and Song, S. 2024. Diffusion Policy: Visuomotor Policy Learning via Action Diffusion. *The International Journal of Robotics Research*.
- Esser, P.; Kulal, S.; Blattmann, A.; Entezari, R.; Müller, J.; Saini, H.; Levi, Y.; Lorenz, D.; Sauer, A.; Boesel, F.; Podell, D.; Dockhorn, T.; English, Z.; Lacey, K.; Goodwin, A.; Marek, Y.; and Rombach, R. 2024. Scaling Rectified Flow Transformers for High-Resolution Image Synthesis. arXiv:2403.03206.
- Geng, Z.; Deng, M.; Bai, X.; Kolter, J. Z.; and He, K. 2025. Mean flows for one-step generative modeling. *arXiv preprint arXiv:2505.13447*.
- Gervet, T.; Xian, Z.; Gkanatsios, N.; and Fragkiadaki, K. 2023. Act3D: 3D Feature Field Transformers for Multi-Task Robotic Manipulation. arXiv:2306.17817.
- Goyal, A.; Blukis, V.; Xu, J.; Guo, Y.; Chao, Y.-W.; and Fox, D. 2024. RVT-2: Learning Precise Manipulation from Few Demonstrations. arXiv:2406.08545.
- Goyal, A.; Xu, J.; Guo, Y.; Blukis, V.; Chao, Y.-W.; and Fox, D. 2023. RVT: Robotic View Transformer for 3D Object Manipulation. arXiv:2306.14896.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising Diffusion Probabilistic Models. arXiv:2006.11239.
- Hu, X.; Liu, B.; Liu, X.; and Liu, Q. 2024. AdaFlow: Imitation Learning with Variance-Adaptive Flow-Based Policies. arXiv:2402.04292.
- Jang, E.; Irpan, A.; Khansari, M.; Kappler, D.; Ebert, F.; Lynch, C.; Levine, S.; and Finn, C. 2022. BC-Z: Zero-Shot Task Generalization with Robotic Imitation Learning. arXiv:2202.02005.
- Lipman, Y.; Chen, R. T. Q.; Ben-Hamu, H.; Nickel, M.; and Le, M. 2023. Flow Matching for Generative Modeling. arXiv:2210.02747.
- Liu, S.; Wu, L.; Li, B.; Tan, H.; Chen, H.; Wang, Z.; Xu, K.; Su, H.; and Zhu, J. 2024. RDT-1B: a Diffusion Foundation Model for Bimanual Manipulation. *arXiv preprint arXiv:2410.07864*.
- Ma, N.; Goldstein, M.; Albergo, M. S.; Boffi, N. M.; Vanden-Eijnden, E.; and Xie, S. 2024a. Sit: Exploring flow and diffusion-based generative models with scalable interpolant transformers. In *European Conference on Computer Vision*, 23–40. Springer.
- Ma, X.; Patidar, S.; Haughton, I.; and James, S. 2024b. Hierarchical Diffusion Policy for Kinematics-Aware Multi-Task Robotic Manipulation. *CVPR*.
- Polyak, A.; Zohar, A.; and Andrew Brown, e. a. 2025. Movie Gen: A Cast of Media Foundation Models. arXiv:2410.13720.
- Prasad, A.; Lin, K.; Wu, J.; Zhou, L.; and Bohg, J. 2024. Consistency Policy: Accelerated Visuomotor Policies via Consistency Distillation. arXiv:2405.07503.
- Rajeswaran, A.; Kumar, V.; Gupta, A.; Vezzani, G.; Schulman, J.; Todorov, E.; and Levine, S. 2017. Learning complex dexterous manipulation with deep reinforcement learning and demonstrations. *arXiv preprint arXiv:1709.10087*.
- Sheng, J.; Liu, Y.; Xu, S.; Yang, Z.; and Liu, M. 2025. GPA-RAM: Grasp-Pretraining Augmented Robotic Attention Mamba for Spatial Task Learning. arXiv:2504.19683.
- Shridhar, M.; Manuelli, L.; and Fox, D. 2022. Perceiver-Actor: A Multi-Task Transformer for Robotic Manipulation. In *Proceedings of the 6th Conference on Robot Learning (CoRL)*.
- Song, Y.; and Dhariwal, P. 2023. Improved Techniques for Training Consistency Models. arXiv:2310.14189.
- Song, Y.; Dhariwal, P.; Chen, M.; and Sutskever, I. 2023. Consistency Models. arXiv:2303.01469.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L. u.; and Polosukhin, I. 2017. Attention is All you Need. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Vosylius, V.; and Johns, E. 2025. Instant Policy: In-Context Imitation Learning via Graph Diffusion. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Wang, D.; Hart, S.; Surovik, D.; Kelestemur, T.; Huang, H.; Zhao, H.; Yeatman, M.; Wang, J.; Walters, R.; and Platt, R. 2024a. Equivariant Diffusion Policy. In *8th Annual Conference on Robot Learning*.
- Wang, L.; Chen, X.; Zhao, J.; and He, K. 2024b. Scaling Proprioceptive-Visual Learning with Heterogeneous Pre-trained Transformers. arXiv:2409.20537.
- Wang, R.; and He, K. 2025. Diffuse and Disperse: Image Generation with Representation Regularization. *arXiv preprint arXiv:2506.09027*.
- Yu, T.; Quillen, D.; He, Z.; Julian, R.; Hausman, K.; Finn, C.; and Levine, S. 2020. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on robot learning*, 1094–1100. PMLR.
- Ze, Y.; Zhang, G.; Zhang, K.; Hu, C.; Wang, M.; and Xu, H. 2024. 3D Diffusion Policy: Generalizable Visuomotor Policy Learning via Simple 3D Representations. In *Proceedings of Robotics: Science and Systems (RSS)*.
- Zhang, Q.; Liu, Z.; Fan, H.; Liu, G.; Zeng, B.; and Liu, S. 2024. FlowPolicy: Enabling Fast and Robust 3D Flow-based Policy via Consistency Flow Matching for Robot Manipulation. arXiv:2412.04987.
- Zhao, T. Z.; Kumar, V.; Levine, S.; and Finn, C. 2023. Learning fine-grained bimanual manipulation with low-cost hardware. *arXiv preprint arXiv:2304.13705*.