
AdaFlow: Imitation Learning with Variance-Adaptive Flow-Based Policies

Xixi Hu, Bo Liu, Xingchao Liu, Qiang Liu
The University of Texas at Austin
{hxixi, bliu, xcliu, lqiang}@cs.utexas.edu

Abstract

Diffusion-based imitation learning improves Behavioral Cloning (BC) on multi-modal decision-making, but comes at the cost of significantly slower inference due to the recursion in the diffusion process. It urges us to design efficient policy generators while keeping the ability to generate diverse actions. To address this challenge, we propose AdaFlow, an imitation learning framework based on flow-based generative modeling. AdaFlow represents the policy with state-conditioned ordinary differential equations (ODEs), which are known as probability flows. We reveal an intriguing connection between the conditional variance of their training loss and the discretization error of the ODEs. With this insight, we propose a variance-adaptive ODE solver that can adjust its step size in the inference stage, making AdaFlow an adaptive decision-maker, offering rapid inference without sacrificing diversity. Interestingly, it automatically reduces to a one-step generator when the action distribution is uni-modal. Our comprehensive empirical evaluation shows that AdaFlow achieves high performance with fast inference speed.

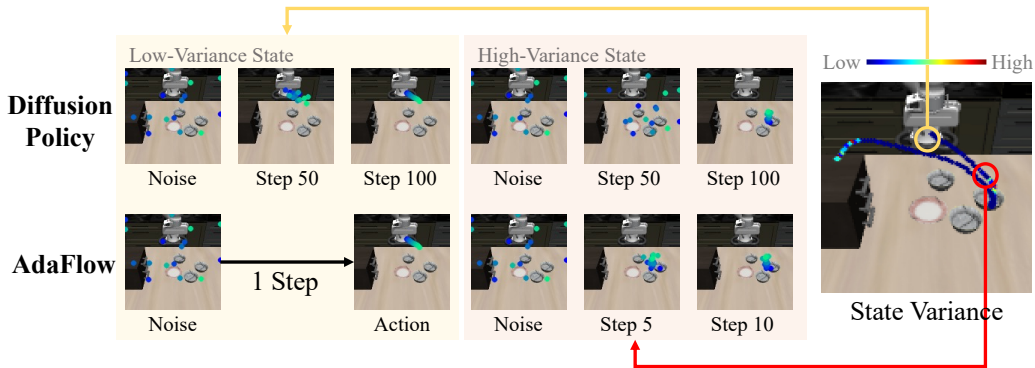


Figure 1: AdaFlow is a fast imitation learning policy. It adaptively adjust the number of simulation steps when generating actions. For low-variance states, it functions as a one-step action generator. For high-variance states, it employs more steps to ensure accurate action generation. This adaptive approach enables AdaFlow to achieve an average generation speed close to one step per task completion.

1 Introduction

Imitation Learning (IL) is a widely adopted method in robot learning [1, 2]. In IL, an agent is given a demonstration dataset from a human expert finishing a certain task, and the goal is for it to complete the task by learning from this dataset. IL is notably effective for learning complex, non-declarative motions, yielding remarkable successes in training real robots [3–6].

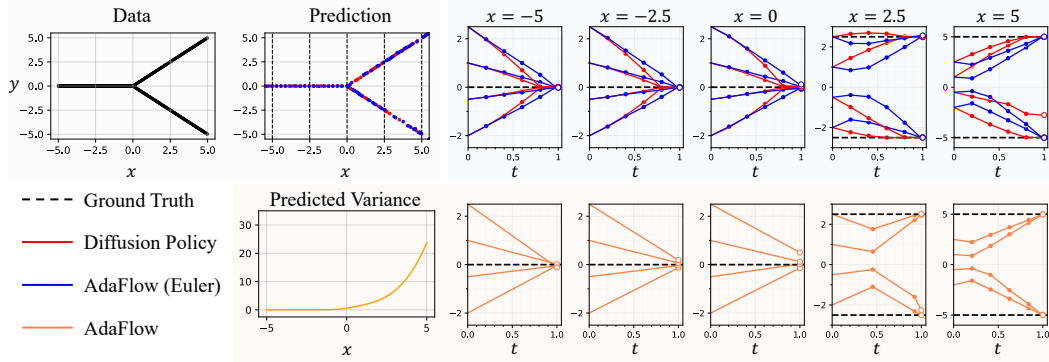


Figure 2: Illustrating the computation adaptivity of AdaFlow (orange) on simple regression task. In the upper portion of the image, we use Diffusion Policy (DDIM) and AdaFlow to predict y given x , with deterministic $y = 0$ when $x \leq 0$, and bimodal $y = \pm x$ when $x > 0$. Both DDIM and AdaFlow fit the demonstration data well. However, the simulated ODE trajectory learned by Diffusion-Policy with DDIM (red) is not straight no matter what x is. By contrast, the simulated ODE trajectory learned by AdaFlow with fixed step (blue) is a straight line when the prediction is deterministic ($x \leq 0$), which means the generation can be exactly done by one-step Euler discretization. At the bottom, we show that AdaFlow can adaptively adjust the number of simulation steps based on the x value according to the estimated variance at x .

The primary approach for IL is Behavioral Cloning (BC) [7–10], where the agent is trained with supervised learning to acquire a deterministic mapping from states to actions. Despite its simplicity, vanilla BC struggles to learn diverse behaviors in states with many possible actions [10, 11, 6, 12]. To improve it, various frameworks have been proposed. For instance, Implicit BC [12] learns an energy-based model for each state and searches the actions that minimize the energy with optimization algorithms. Diffuser [13, 14] and Diffusion Policy [6] adopts diffusion models [15, 16] to generate diverse actions, which has become the default method for training on large-scale robotics data [17–20].

The computational cost of the learned policies at the execution stage is important for an IL framework in a real-world deployment. Unfortunately, none of the previous frameworks enjoys both efficient inference and diversity. Although energy-based models and diffusion models can generate multi-modal action distributions, they require *recursive processes* to generate the actions. These recursive processes usually involve tens (or even hundreds) of queries before reaching their stopping criteria.

In this paper, we propose a new IL framework, named AdaFlow, that learns a dynamic generative policy that can autonomously adjust its computation on the fly, thus cheaply outputs multi-modal action distributions to complete the task. AdaFlow is inspired by recent advancements in flow-based generative modeling [21–24]. We learn probability flows, which are essentially ordinary differential equations (ODEs), to represent the policies. These flows are powerful generative models that precisely capture the complicated distributions, but similar to energy-based models and diffusion models, they still require *multiple recursive iterations* to simulate the ODEs in the inference stage.

AdaFlow differs from existing flow generative models like Rectified Flow [25] and Consistency Models [26], by utilizing the *initially learned ODE* to maintain low training and inference costs, and function as a one-step generator for deterministic target distributions. In contrast, both of these methods require an additional distillation or reflow [25] process to achieve fast inference. To improve the efficiency, we propose an adaptive ODE solver based on the finding that the simulation error of the ODE is closely related to the variance of the training loss at different states. We let the action generation model to output an additional variance scalar alongside the action it produces. During the execution of the policy, we change the step size according to the variance predicted at the current state. Equipping the flow-based policy with the proposed adaptive ODE solver, AdaFlow wisely allocates computational resources, yielding high efficiency without sacrificing the diversity provided by the flow-based generative models. Specifically, in states with deterministic action distributions, AdaFlow generates the action in *one step* – as efficient as naive BC.

Empirical results across multiple benchmarks demonstrate that AdaFlow achieve consistently good performance across success rate with high execution efficiency. Specifically, our contributions are:

- We proposed AdaFlow as a generative model-based policy for decision making tasks, capable of generating actions almost as quickly as a single model inference pass.
- We conducted comprehensive experiments across decision making tasks, including navigation and robot manipulation, utilizing benchmarks such as LIBERO [27] and RoboMimic [10]. AdaFlow consistently outperforms existing state-of-the-art models, despite requiring 10x less inference times to generate actions.
- We offer a theoretical analysis of the overall error in action generation by AdaFlow, providing a bound that ensures precision and reliability.

2 Related Work

Diffusion/Flow-based Generative Models and Adaptive Inference Diffusion models [28, 16, 15, 29] succeed in various applications, e.g., image/video generation [30–33], audio generation [34], point cloud generation [35–38], etc.. However, numerical simulation of the diffusion processes typically involve hundreds of steps, resulting in high inference cost. Post-hoc samplers have been proposed to solve this issue [39–44] by transforming the diffusion process into marginal-preserving probability flow ODEs, yet they still use the same number of inference steps for different states. Although adaptive ODE solvers, such as adaptive step-size Runge-Kutta [45], exist, they cannot significantly reduce the number of inference steps. In comparison, the adaptive sampling strategy of AdaFlow is specifically designed based on intrinsic properties of the ODE learned rectified flow, and can achieve one-step simulation for most of the states, making it much faster for decision-making tasks in real-world applications. Recently, new generative models [21, 25, 22, 23, 46, 24, 26, 47] have emerged. These models directly learn probability flow ODEs by constructing linear interpolations between two distributions, or learn to distill a pretrained diffusion model [26, 47] with an additional distillation training phase. Empirically, these methods exhibit more efficient inference due to their preference of straight trajectories. Among them, Rectified flow achieves one step generation with reflow, a process to straighten the ODE. However, it requires a costly synthetic data construction.

By contrast, AdaFlow only leverages the initially learned ODE, but still keeps cheap training and inference costs that are similar to behavior cloning. We achieve this by unveiling a previously overlooked feature of these flow-based generative models: they act as one-step generators for deterministic target distributions, and their variance indicates the straightness of the probability flows for a certain state. Leveraging this feature, we design AdaFlow to automatically change the level of action modalities given on the states.

Diffusion Models for Decision Making For decision making, diffusion models obtain success as in other applications areas [48–51]. In a pioneering work, Janner et al. [13] proposed “Diffuser”, a planning algorithm with diffusion models for offline reinforcement learning. This framework is extended to other tasks in the context of *offline reinforcement learning* [52], where the training dataset includes reward values. For example, Ajay et al. [14] propose to model policies as conditional diffusion models. The application of DDPM [16] and DDIM [43] on visuomotor policy learning for physical robots [6] outperforms counterparts like Behavioral Cloning. Freund et al. [53] exploits two coupled normalizing flows to learn the distribution of expert states, and use that as a reward to train an RL agent for imitation learning. AdaFlow admits a much simpler training and inference pipeline compared with it. Despite the success of adopting generative diffusion models as decision makers in previous works, they also bring redundant computation, limiting their application in real-time, low-latency decision-making scenarios for autonomous robots. AdaFlow propose to leverage rectified flow instead of diffusion models, facilitating adaptive decision making for different states while significantly reducing computational requirements. In this work, similar to Diffusion Policy [6], we focus on offline imitation learning. While AdaFlow could theoretically be adapted for offline reinforcement learning, we leave it for future works.

3 AdaFlow for Imitation Learning

To yield an agent that enjoys both multi-modal decision-making and fast execution, we propose AdaFlow, an imitation learning framework based on flow-based generative policy. The merits of AdaFlow lie in its adaptive ability: it identifies the behavioral complexity at a state before allocating

computation. If the state has a deterministic choice of action, it outputs the required action rapidly; otherwise, it spends more inference time to take full advantage of the flow-based generative policy. This handy adaptivity is made possible by leveraging a combination of elements: 1) a special property of the flow 2) a variance estimation neural network and 3) a variance-adaptive ODE solver. We formally introduce the whole framework in the sequel.

3.1 Flow-Based Generative Policy

Given the expert dataset $D = \{(s^{(i)}, a^{(i)})\}_{i=1}^n$, our goal is to learn a policy π_θ that can generate trajectories following the target distribution π_E . π_θ can be induced from a state-conditioned flow-based model,

$$dz_t = v_\theta(z_t, t | s)dt, \quad z_0 \sim \mathcal{N}(0, I). \quad (1)$$

Here, s is the state and the velocity field is parameterized by a neural network θ that takes the state as an additional input. To capture the expert distribution with the flow-based model, the velocity field can be trained by minimizing a state-conditioned least-squares objective,

$$L(\theta) = \mathbb{E}_{\substack{(s, \mathbf{a}) \sim D \\ \mathbf{x}_0 \sim \mathcal{N}(0, I)}}} \left[\int_0^1 \|\mathbf{a} - \mathbf{x}_0 - v_\theta(\mathbf{x}_t, t | s)\|_2^2 dt \right], \quad (2)$$

where \mathbf{x}_t is the linear interpolation between \mathbf{x}_0 and $\mathbf{x}_1 = \mathbf{a}$:

$$\mathbf{x}_t = t\mathbf{a} + (1-t)\mathbf{x}_0. \quad (3)$$

We should differentiate z_t which is the ODE trajectory in (1) from the linear interpolation \mathbf{x}_t . They do not overlap unless all trajectories of ODE (1) are straight. See Liu et al. [21] for more discussion.

With infinite data sampled from π_E , unlimited model capacity and perfect optimization, it is guaranteed that the policy π_θ generated from the learned flow matches the expert policy π_E [21].

3.2 The Variance-Adaptive Nature of Flow

Typically, to sample from the distribution π_θ at state s , we start with a random sample z_0 from the Gaussian distribution and simulate the ODE (Eq. (1)) with multi-step ODE solvers to get the action. For example, we can exploit N -step Euler discretization,

$$z_{t_{i+1}} = z_{t_i} + \frac{1}{N}v_\theta(z_{t_i}, t_i | s), \quad t_i = \frac{i}{N}, 0 \leq i < N. \quad (4)$$

After running the solver, z_1 is the generated action. This solver requires inference with the network N times for decision making in every state. Moreover, a large N is needed to obtain a smaller numerical error.

However, different states may have different levels of difficulty in deciding the potential actions. For instance, when traveling from a city A to another city B, there could be multiple ways for transportation, corresponding to a multi-modal distribution of actions. After the way of transportation is chosen, the subsequent actions to take will be almost deterministic. This renders using a uniform Euler solver with the same number of inference steps N across all the states a sub-optimal solution. Rather, it is preferred that the agent can vary its decision-making process as the state of the agent changes. The challenge is how to quantitatively estimate the *complexity of a state* and employ the estimation to *adjust the inference of the flow-based policy*.

Variance as a Complexity Indicator We notice an intriguing property of the policy learned by rectified flow, which connects the complexity of a state with the training loss of the flow-based policy: if the distribution of actions is deterministic at a state s (i.e., a Dirac distribution), the trajectory of rectified flow ODE is a straight line, i.e., a *single Euler step* yields an exact estimation of z_1 .

Proposition 3.1. *Let v^* be the optimum of Eq. (2). If $\text{var}_{\pi_E}(\mathbf{a} | s) = 0$ where $\mathbf{a} \sim \pi_E(\cdot | s)$, then the learned ODE conditioned on s is*

$$dz_t = v^*(z_t, t | s)dt = (\mathbf{a} - z_0)dt, \quad \forall t \in [0, 1], \quad (5)$$

whose trajectories are straight lines pointing to z_1 and hence can be calculated exactly with one step of Euler step:

$$z_1 = z_0 + v^*(z_0, 0 | s).$$

Algorithm 1 AdaFlow: Execution

- 1: **Input:** Current state \mathbf{s} , minimal step size ϵ_{\min} , error threshold η , pre-trained networks v_θ and σ_ϕ .
 - 2: Initialize $\mathbf{z}_0 \sim \mathcal{N}(0, I)$, $t = 0$.
 - 3: **while** $t < 1$ **do**
 - 4: Compute step size
$$\epsilon_t = \text{Clip} \left(\frac{\eta}{\sigma_\phi(\mathbf{z}_t, t \mid \mathbf{s})}, [\epsilon_{\min}, 1 - t] \right).$$
 - 5: Update $t \leftarrow t + \epsilon_t, \mathbf{z}_t \leftarrow \mathbf{z}_t + \epsilon_t v_\theta(\mathbf{z}_t, t \mid \mathbf{s})$.
 - 6: **end while**
 - 7: Execute action $a = \mathbf{z}_1$.
-

Note that the straight trajectories of (5) satisfies $\mathbf{z}_t = t\mathbf{a} + (1 - t)\mathbf{z}_0$, which makes it coincides with the linear interpolation \mathbf{x}_t . As show in [21], this happens only when all the linear trajectories do not intersect on time $t \in [0, 1)$.

More generally, we can expect that the straightness of the ODE trajectories depends on how deterministic the expert policy π_E is. Moreover, the straightness can be quantified by a conditional variance metric defined as follows:

$$\begin{aligned} \sigma^2(x, t \mid \mathbf{s}) &= \text{var}(\mathbf{a} - \mathbf{x}_0 \mid \mathbf{x}_t = x, \mathbf{s}) \\ &= \mathbb{E} \left[\|\mathbf{a} - \mathbf{x}_0 - v^*(\mathbf{x}_t, t \mid \mathbf{s})\|^2 \mid \mathbf{x}_t = x, \mathbf{s} \right]. \end{aligned} \tag{6}$$

Proposition 3.2. *Under the same condition as Proposition 3.1, we have $\sigma^2(\mathbf{z}_t, t \mid \mathbf{s}) = 0$ from (5).*

The proof of the above propositions is in Appendix A.1. To summarize, the variance of the state-conditioned loss function at (\mathbf{z}_t, t) can be an indicator of the multi-modality of actions. When the variance is zero, the flow-based policy can generate the expected action with only one query of the velocity field, saving a huge amount of computation. In Section 3.3, we will show the variance can be used to bound the discretization error, thereby enabling the design of an adaptive ODE solver.

Variance Estimation Network In practice, the conditional variance $\sigma^2(x, t \mid \mathbf{s})$ can be empirically approximated by a neural network $\sigma_\phi^2(x, t \mid \mathbf{s})$ with parameter ϕ . Once the neural velocity v_θ is learned, we can estimate σ_ϕ by minimizing the following Gaussian negative log-likelihood loss:

$$\min_{\phi} \mathbb{E} \left[\int_0^1 \frac{\|\mathbf{a} - \mathbf{x}_0 - v_\theta(\mathbf{x}_t, t \mid \mathbf{s})\|^2}{2\sigma_\phi^2(\mathbf{x}_t, t \mid \mathbf{s})} + \log \sigma_\phi^2(\mathbf{x}_t, t \mid \mathbf{s}) dt \right]. \tag{7}$$

We adopt a two-stage training strategy by first training the velocity network v_θ then the variance estimation network σ_ϕ . In practice, the second stage just involves fine-tuning a few linear layers on top of the trained velocity network. Alternatively, we can optimize both the variance estimation and action generation simultaneously, which can extend training time. Our experiments show that joint training and two-stage training yield comparable performance.

3.3 Variance-Adaptive Flow-Based Policy

Because the variance indicates the straightness of the ODE trajectory, it allows us to develop an adaptive approach to set the step size to yield better estimation with lower error during inference.

To derive our method, let us consider to advance the ODE with step size ϵ_t at \mathbf{z}_t :

$$\mathbf{z}_{t+\epsilon_t} = \mathbf{z}_t + \epsilon_t v^*(\mathbf{z}_t, t \mid \mathbf{s}). \tag{8}$$

The problem is how to set the step size ϵ_t properly. If ϵ_t is too large, the discretized solution will significantly diverge from the continuous solution; if ϵ_t is too small, it will take excessively many steps to compute.

We propose an adaptive ODE solver based on the principle of matching the discretized marginal distribution p_t of \mathbf{z}_t from (8), and the ideal marginal distribution p_t^* when following the exact ODE

(1). This is made possible with a key insight below showing that the discretization error can be bounded by the conditional variance $\sigma^2(\mathbf{z}_t, t \mid \mathbf{s})$.

Proposition 3.3. *Let p_t^* be the marginal distribution of the exact ODE $d\mathbf{z}_t = v^*(\mathbf{z}_t, t \mid \mathbf{s})dt$. Assume $\mathbf{z}_t \sim p_t = p_t^*$ and $p_{t+\epsilon_t}$ the distribution of $\mathbf{z}_{t+\epsilon_t}$ following (8). Then we have*

$$W_2(p_{t+\epsilon_t}^*, p_{t+\epsilon_t})^2 \leq \epsilon_t^2 \mathbb{E}_{\mathbf{z}_t \sim p_t} [\sigma^2(\mathbf{z}_t, t \mid \mathbf{s})],$$

where W_2 denotes the 2-Wasserstein distance.

We provide the proof in Appendix A.2. Hence, given a threshold η , to ensure that an error of $W_2(p_{t+\epsilon_t}^*, p_{t+\epsilon_t})^2 \leq \eta^2$, we can bound the step size by $\epsilon_t \leq \eta/\sigma(\mathbf{z}_t, t \mid \mathbf{s})$. Because ϵ_t at time t should not be large than $1 - t$, we suggest the following rule for setting the step size ϵ_t at \mathbf{z}_t at time t :

$$\epsilon_t = \text{Clip} \left(\frac{\eta}{\sigma(\mathbf{z}_t, t \mid \mathbf{s})}, [\epsilon_{\min}, 1 - t] \right), \quad (9)$$

where we impose an additional lower bound ϵ_{\min} to avoid ϵ_t to be unnecessarily small. Besides, the proposed adaptive strategy guarantees to instantly arrive at the terminal point when $\sigma^2(\mathbf{z}_t, t \mid \mathbf{s}) = 0$ as $\epsilon_t = 1 - t$. Moreover, it aligns with Section. 3.2 since for states with deterministic actions, it sets $\epsilon_0 = 1$ to generate the action in one step. We incorporate the above insights to the execution in Algorithm 1.

Global Error Analysis Proposition 3.3 provides the local error at each Euler step. In the following, we provide an analysis of the overall error for generating \mathbf{z}_1 when we simulate ODE while following the adaptive rule (9). To simplify the notation, we drop the dependency on the state \mathbf{s} , and write $v_t^*(\cdot) = v^*(\cdot, t \mid \mathbf{s})$.

Assumption 3.4. *Assume $\|v_t^*\|_{Lip} \leq L$ for $t \in [0, 1]$, and the solutions of $d\mathbf{z}_t = v_t(\mathbf{z}_t)dt$ has bounded second curvature $\|\ddot{\mathbf{z}}_t\| \leq M$ for $t \in [0, 1]$.*

This is a standard assumption in numerical analysis, under which Euler’s method with a constant step size of ϵ_{\min} admits a global error of order $O(\epsilon_{\min})$.

Proposition 3.5. *Under Assumption 3.4, assume we follow Euler step (8) with step size ϵ_t in (9), starting from $\mathbf{z}_0 = \mathbf{x}_0 \sim p_0^*$. Let p_t be the distribution of \mathbf{z}_t we obtained in this way, and p_t^* that of \mathbf{x}_t in (3). Note that p_1^* is the true data distribution. Set $\eta = M_\eta \epsilon_{\min}^2 / 2$ for some $M_\eta > 0$, and $\epsilon_{\min} = 1/N_{\max}$.*

Let N_{ada} be the number of steps we arrive at \mathbf{z}_1 following the adaptive schedule. We have

$$W_2(p_1^*, p_1) \leq C \times \frac{N_{\text{ada}}}{N_{\max}} \times \epsilon_{\min},$$

where C is a constant depending on M , M_η and L .

The idea is that the error is proportional to $\frac{N_{\text{ada}}}{N_{\max}}$, suggesting that the algorithm claims an improved error bound in the good case when it takes a smaller number of steps than the standard Euler method with constant step size ϵ_{\min} . We provide the proof in Appendix A.3.

Discussion of AdaFlow and Rectified Flow. Rectified Flow operates in two stages: the first is learning an ordinary differential equation (ODE), and the second involves a technique called "reflow" used to straighten the learned trajectory. Theoretically, reflow allows for one-step action generation. However, using reflow introduces two major drawbacks: 1) It significantly prolongs training time, particularly because generating the required pseudo noise-data pairs through ODE simulation is computationally expensive; 2) It leads to poorer generation quality due to straightened ODE. In contrast, our method utilizes only the original ODE, eliminating the need for an additional reflow or distillation process, and consistently achieves more accurate action generation.

4 Experiments

We conducted comprehensive experiments on four sets of tasks: **1)** a simple 1D toy example to demonstrate the computational adaptivity of AdaFlow; **2)** a 2D navigation problem; and two robot

manipulation task suites on **3) RoboMimic** [14] following past works [6] and **4) LIBERO** [3], provide diverse and realistic scenarios for evaluation.

Our results show that AdaFlow improves the success rate of completing both navigation and manipulation tasks, outperforming state-of-the-art methods such as BC and its variants, as well as Diffusion Policy, across a range of tasks. Additionally, AdaFlow drastically reduces the inference cost. Further experiments demonstrate that AdaFlow is robust to changes in hyperparameters and can adaptively adjust its inference speed according to different states, ensuring efficient and reliable performance.

	BC	Diffusion Policy	Rectified Flow	AdaFlow
Behavior Diversity	✗	✓	✓	✓
Fast Action Generation	✓	✗	✓	✓
No Distillation / Reflow	✓	✓	✗	✓

Table 1: Comparison between BC, Diffusion Policy, Rectified Flow and AdaFlow.

4.1 Regression

We start with a 1D regression task designed to demonstrate the adaptivity nature of AdaFlow. The goal is to learn a mapping from x to y where

$$y = \begin{cases} 0 & \text{for } x \leq 0 \\ \pm x & \text{for } x > 0. \end{cases} \quad (10)$$

Note that $y | x$ is deterministic when $x \leq 0$ and stochastic otherwise. The training and testing data are uniformly sampled from the ground-truth function with $x \in [-5, 5]$. Details about the setup and the hyperparameters are provided in Appendix.

AdaFlow can achieve 1-step generation for deterministic states. Figure 2 (top-right) shows the generation trajectories of Diffusion Policy and AdaFlow with 5 step. Notably, when $x \leq 0$, AdaFlow generates *straight* trajectories and is therefore able to predict y with a single step, aligning our analysis in Proposition 3.1 and 3.2. In contrast, Diffusion Policy generates curved trajectories when step = 5, and hence cannot predict y accurately with a single step. The bottom of Figure 2 shows the estimated variance by AdaFlow across $x \in [-5, 5]$, which accurately aligns with the expected variance. In addition, as x increases, AdaFlow adaptively increases the required number of simulation steps.

4.2 Navigating a 2D Maze

We create two sets of maze navigation tasks to validate AdaFlow’s performance of modeling multi-modal behavior. In particular, we create two *single-task* environments where the agent starts and ends at a fixed point and two *multi-task* environments where the agent can start and end at different points. All four environments are simulated in D4RL Maze2D [54] using MuJoCo. The environments and demonstrations are visualized in Figure 7.

Method	NFE↓	Maze 1	Maze 2	Maze 3	Maze 4
<i>Needs reflow</i>					
Rectified Flow	1	0.82	1.00	1.00	0.80
BC	1	1.00	1.00	0.92	0.76
BC-GMM	1	0.84	1.00	0.88	0.72
Diffusion Policy	1	0.00	0.32	0.16	0.08
Diffusion Policy	5	0.58	1.00	0.84	0.76
Diffusion Policy	20	0.62	0.98	0.84	0.82
AdaFlow	1.56	0.98	1.00	0.96	0.86

Table 2: Performance on maze navigation tasks. The table showcases the success rate for each model across different maze complexities. The highest success rate for each task are highlighted in **bold**. NFE denotes Number of Function Evaluations.

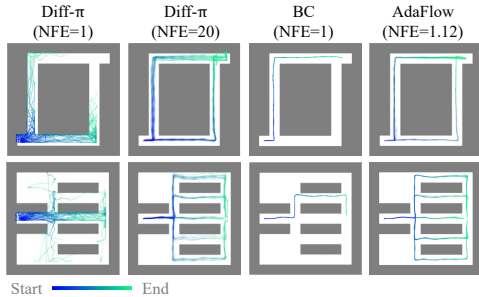


Figure 3: **Generated trajectories.** We visualize the trajectories generated by different policies, with the agent’s starting point fixed.

Method	NFE↓	Lift		Can		Square		Transport		ToolHang	Push-T
		ph	mh	ph	mh	ph	mh	ph	mh	ph	ph
Rectified Flow (<i>Needs reflow</i>)	1	1.00	1.00	0.94	1.00	0.94	0.92	0.90	0.76	0.88	0.92
LSTM-GMM	1	1.00	1.00	1.00	1.00	0.95	0.86	0.76	0.62	0.67	0.69
IBC	1	0.79	0.15	0.00	0.01	0.00	0.00	0.00	0.00	0.00	0.75
BET	1	1.00	1.00	1.00	1.00	0.76	0.68	0.38	0.21	0.58	-
Diffusion Policy	1	0.04	0.04	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.04
Diffusion Policy	2	0.64	0.98	0.52	0.66	0.56	0.12	0.84	0.68	0.68	0.34
Diffusion Policy	100	1.00	1.00	1.00	1.00	1.00	0.97	0.90	0.72	0.90	0.91
AdaFlow	1.17	1.00	1.00	1.00	0.96	0.98	0.96	0.92	0.80	0.88	0.96

Table 3: Success rate on RoboMimic Benchmark. The highest success rate for each task are highlighted in **bold**.




Figure 4: **LIBERO tasks**. We visualize the demonstrated trajectories of the robot’s end effector.

AdaFlow achieves high diversity and success with low NFE; Diffusion Policy and BC lag in comparison. We compare AdaFlow against baseline methods in Table 2. We additionally visualize the rollout trajectories from each learned policy in Figure 3 as a qualitative comparison of the learned behavior across different methods. From the results, we see that AdaFlow with an average Number of Function Evaluation (NFE) of 1.56 NFE can achieve highly diverse behavior and high success rate in the meantime. By contrast, Diffusion Policy only demonstrates diverse behavior when NFE is larger than 5 and falls behind in success rate even with 20 NFE compared to AdaFlow. BC, on the other hand, has high success rate while performing relatively poorly in terms of behavior diversity.

4.3 Robot Manipulation Tasks

Experiment Setup To further validate how AdaFlow performs on practical robotics tasks, we compare AdaFlow against baselines on a Push-T task [6], the RoboMimic [10] benchmark (Lift, Can, Square, Transport, ToolHang) and the LIBERO [27] benchmark. For the Push-T task and the tasks in RoboMimic, we follow the exact experimental setup described in Diffusion Policy [6]. Following the Diffusion Policy, we add three additional baseline methods: 1) LSTM-GMM, BC with the LSTM model and a Gaussian mixture head, 2) IBC, the implicit behavioral cloning [12], an energy-based model for generative decision-making, and 3) BET [11]. For the LIBERO tasks, we pick a subset of six Kitchen tasks and follow the setup described in the LIBERO paper (Check Figure 4 for the description of the six tasks).

Method	NFE↓	Task 1	Task 2	Task 3	Task 4	Task 5	Task 6	Average
Rectified Flow (<i>Needs reflow</i>)	1	0.90	0.82	0.98	0.82	0.82	0.96	0.88
Diffusion Policy	1	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Diffusion Policy	2	0.00	0.58	0.36	0.66	0.36	0.32	0.38
Diffusion Policy	20	0.94	0.84	0.98	0.78	0.82	0.92	0.88
AdaFlow	1.27	0.98	0.80	0.98	0.82	0.90	0.96	0.91

Table 4: Success Rate on LIBERO Benchmark. The highest success rate for each task are highlighted in **bold**.

AdaFlow consistently outperforms competitors in varied robot manipulation tasks with high efficiency. The results of the Push-T task and the RoboMimic benchmark are summarized in Table 3. From the table, we observe that AdaFlow consistently achieves comparable or higher success rates across different challenging manipulation tasks, compared against all baselines, with only an average

NFE of 1.17. Note that Diffusion Policy, while showing high success rates using NFE = 100, falls behind when NFE = 1. Results for the six LIBERO tasks are presented in Table 4. Aligning with findings from our previous experiments, AdaFlow once again outperforms BC and Diffusion Policy in terms of success rate with an average NFE of 1.27. We additionally visualize the variance predicted by AdaFlow in Figure 5. It can be seen that the model identifies the high variance when the robot’s end-effector is close to the object or target area, matching the variance from the demonstration data.

4.4 Ablation Study

We valid how AdaFlow performs against baselines regarding the training and inference efficiency. In addition, we examine how critical the variance estimation network is.




Figure 5: **Predicted variance.** We visualize the variance predicted by AdaFlow. The variance is computed on states from the expert’s demonstration and averaged over all simulation steps (e.g., t from 0 to 1). Then we normalize the variance to $[0, 1]$ by the largest variance found at all states.




Figure 6: Ablation studies on AdaFlow.

Higher Training and Inference Efficiency. Figure 6 (top) examines changes in success rate relative to the NFE. AdaFlow maintains a high success rate with a very low NFE, whereas the Diffusion Policy generally requires more than three NFE to perform well. Although BC performs well with one NFE, it demonstrates very limited behavioral diversity and struggles to model multi-modal behavior. Figure 6 (bottom) illustrates training efficiency by displaying the success rate over epochs. It shows that AdaFlow has a better area-under-curve than Diffusion Policy, indicating faster learning. As expected, due to its simplicity, Behavioral Cloning (BC) achieves the best learning efficiency.

Robustness to η . In Figure 6, the NFEs in AdaFlow are calculated at various η values. It shows that AdaFlow is robust to changes in η .

On the Importance of Variance Estimation. In Table 5, we provide the performance of AdaFlow with and without the variance estimation network on the four mazes from Section 4.2. From the results, it is clear that the variance estimation network not only makes inference faster, but can also lead to better performance.

	Maze1	Maze1	Maze3	Maze4
w/o Variance Estimation	0.78	1.00	0.92	0.80
AdaFlow (Ours)	0.98	1.00	0.96	0.86

Table 5: Ablation study on the use of estimated variance to determine inference steps. Euler sampler is used for AdaFlow without variance estimation.

5 Conclusion

We present AdaFlow, a novel imitation learning algorithm adept at efficiently generating diverse and adaptive policies, addressing the trade-off between computational efficiency and behavioral diversity inherent in current models. Through extensive experimentation across various settings, AdaFlow demonstrated superior performance across multiple dimensions including success rate, behavioral diversity, and training/execution efficiency. This work lays a robust foundation for future research on adaptive imitation learning methods in real-world scenarios.

References

- [1] Schaal, S. Is imitation learning the route to humanoid robots? *Trends in cognitive sciences*, 3(6):233–242, 1999.
- [2] Osa, T., J. Pajarinen, G. Neumann, et al. An algorithmic perspective on imitation learning. *Foundations and Trends® in Robotics*, 7(1-2):1–179, 2018.
- [3] Liu, B., X. Xiao, P. Stone. A lifelong learning approach to mobile robot navigation. *IEEE Robotics and Automation Letters*, 6(2):1090–1096, 2021.
- [4] Zhu, Y., A. Joshi, P. Stone, et al. Viola: Imitation learning for vision-based manipulation with object proposal priors. *6th Annual Conference on Robot Learning (CoRL)*, 2022.
- [5] Brohan, A., N. Brown, J. Carbajal, et al. Rt-1: Robotics transformer for real-world control at scale. *arXiv preprint arXiv:2212.06817*, 2022.
- [6] Chi, C., S. Feng, Y. Du, et al. Diffusion policy: Visuomotor policy learning via action diffusion. *arXiv preprint arXiv:2303.04137*, 2023.
- [7] Pomerleau, D. A. Alvin: An autonomous land vehicle in a neural network. *Advances in neural information processing systems*, 1, 1988.
- [8] Ross, S., G. Gordon, D. Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635. JMLR Workshop and Conference Proceedings, 2011.
- [9] Torabi, F., G. Warnell, P. Stone. Behavioral cloning from observation. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization, 2018.
- [10] Mandlekar, A., D. Xu, J. Wong, et al. What matters in learning from offline human demonstrations for robot manipulation. *arXiv preprint arXiv:2108.03298*, 2021.
- [11] Shafiullah, N. M., Z. Cui, A. A. Altanzaya, et al. Behavior transformers: Cloning k modes with one stone. *Advances in neural information processing systems*, 35:22955–22968, 2022.
- [12] Florence, P., C. Lynch, A. Zeng, et al. Implicit behavioral cloning. In *Conference on Robot Learning*, pages 158–168. PMLR, 2022.
- [13] Janner, M., Y. Du, J. Tenenbaum, et al. Planning with diffusion for flexible behavior synthesis. In *International Conference on Machine Learning*, pages 9902–9915. PMLR, 2022.
- [14] Ajay, A., Y. Du, A. Gupta, et al. Is conditional generative modeling all you need for decision making? In *The Eleventh International Conference on Learning Representations*. 2022.
- [15] Song, Y., J. Sohl-Dickstein, D. P. Kingma, et al. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*. 2020.
- [16] Ho, J., A. Jain, P. Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

- [17] Sridhar, A., D. Shah, C. Glossop, et al. Nomad: Goal masked diffusion policies for navigation and exploration. *arXiv preprint arXiv:2310.07896*, 2023.
- [18] Team, O. M., D. Ghosh, H. Walke, et al. Octo: An open-source generalist robot policy, 2023.
- [19] Shah, D., A. Sridhar, N. Dashora, et al. Vint: A foundation model for visual navigation. *arXiv preprint arXiv:2306.14846*, 2023.
- [20] Hansen-Estruch, P., I. Kostrikov, M. Janner, et al. Idql: Implicit q-learning as an actor-critic method with diffusion policies. *arXiv preprint arXiv:2304.10573*, 2023.
- [21] Liu, X., C. Gong, Q. Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *The Eleventh International Conference on Learning Representations*. 2022.
- [22] Lipman, Y., R. T. Chen, H. Ben-Hamu, et al. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*. 2022.
- [23] Albergo, M. S., N. M. Boffi, E. Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *arXiv preprint arXiv:2303.08797*, 2023.
- [24] Heitz, E., L. Belcour, T. Chambon. Iterative α -(de)blending: A minimalist deterministic diffusion model. In *ACM SIGGRAPH 2023 Conference Proceedings, SIGGRAPH '23*. Association for Computing Machinery, New York, NY, USA, 2023.
- [25] Liu, Q. Rectified flow: A marginal preserving approach to optimal transport. *arXiv preprint arXiv:2209.14577*, 2022.
- [26] Song, Y., P. Dhariwal, M. Chen, et al. Consistency models. 2023.
- [27] Liu, B., Y. Zhu, C. Gao, et al. Libero: Benchmarking knowledge transfer for lifelong robot learning. *arXiv preprint arXiv:2306.03310*, 2023.
- [28] Sohl-Dickstein, J., E. Weiss, N. Maheswaranathan, et al. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015.
- [29] Song, Y., S. Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.
- [30] Ho, J., T. Salimans, A. A. Gritsenko, et al. Video diffusion models. In *Advances in Neural Information Processing Systems*.
- [31] Zhang, S., X. Yang, Y. Feng, et al. Hive: Harnessing human feedback for instructional visual editing. *arXiv preprint arXiv:2303.09618*, 2023.
- [32] Wu, L., C. Gong, X. Liu, et al. Diffusion-based molecule generation with informative prior bridges. *Advances in Neural Information Processing Systems*, 35:36533–36545, 2022.
- [33] Saharia, C., W. Chan, S. Saxena, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022.
- [34] Kong, Z., W. Ping, J. Huang, et al. Diffwave: A versatile diffusion model for audio synthesis. In *International Conference on Learning Representations*. 2020.
- [35] Luo, S., W. Hu. Diffusion probabilistic models for 3d point cloud generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2837–2845. 2021.
- [36] Liu, X., L. Wu, M. Ye, et al. Learning diffusion bridges on constrained domains. In *The Eleventh International Conference on Learning Representations*. 2023.
- [37] Luo, S., W. Hu. Score-based point cloud denoising. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4583–4592. 2021.

- [38] Wu, L., D. Wang, C. Gong, et al. Fast point cloud generation with straight flows. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9445–9454. 2023.
- [39] Karras, T., M. Aittala, T. Aila, et al. Elucidating the design space of diffusion-based generative models. In *Advances in Neural Information Processing Systems*. 2022.
- [40] Lu, C., Y. Zhou, F. Bao, et al. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems*, 35:5775–5787, 2022.
- [41] Liu, L., Y. Ren, Z. Lin, et al. Pseudo numerical methods for diffusion models on manifolds. In *International Conference on Learning Representations*. 2021.
- [42] Lu, C., Y. Zhou, F. Bao, et al. Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. *arXiv preprint arXiv:2211.01095*, 2022.
- [43] Song, J., C. Meng, S. Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*. 2020.
- [44] Bao, F., C. Li, J. Zhu, et al. Analytic-dpm: an analytic estimate of the optimal reverse variance in diffusion probabilistic models. In *International Conference on Learning Representations*. 2021.
- [45] Press, W. H., S. A. Teukolsky. Adaptive stepsize runge-kutta integration. *Computers in Physics*, 6(2):188–191, 1992.
- [46] Albergo, M. S., E. Vanden-Eijnden. Building normalizing flows with stochastic interpolants. In *The Eleventh International Conference on Learning Representations*. 2022.
- [47] Salimans, T., J. Ho. Progressive distillation for fast sampling of diffusion models. In *International Conference on Learning Representations*. 2022.
- [48] Kapelyukh, I., V. Vosylius, E. Johns. Dall-e-bot: Introducing web-scale diffusion models to robotics. *IEEE Robotics and Automation Letters*, 2023.
- [49] Yang, S., O. Nachum, Y. Du, et al. Foundation models for decision making: Problems, methods, and opportunities. *arXiv preprint arXiv:2303.04129*, 2023.
- [50] Pearce, T., T. Rashid, A. Kanervisto, et al. Imitating human behaviour with diffusion models. In *Deep Reinforcement Learning Workshop NeurIPS 2022*. 2022.
- [51] Chang, W.-D., J. C. G. Higuera, S. Fujimoto, et al. Il-flow: Imitation learning from observation using normalizing flows. *arXiv preprint arXiv:2205.09251*, 2022.
- [52] Wang, Z., H. Zheng, P. He, et al. Diffusion-gan: Training gans with diffusion. In *The Eleventh International Conference on Learning Representations*. 2022.
- [53] Freund, G. J., E. Sarafian, S. Kraus. A coupled flow approach to imitation learning. In *International Conference on Machine Learning*, pages 10357–10372. PMLR, 2023.
- [54] Fu, J., A. Kumar, O. Nachum, et al. D4rl: Datasets for deep data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219*, 2020.

A Appendix

A.1 Proof of Proposition 3.1 and Proposition 3.2.

PROOF 1. $\text{var}_{\pi_E}(\mathbf{a}|s) = 0$ means that the action $\mathbf{a} = a$ equals a deterministic value a given s . With $\mathbf{x}_t = ta + (1-t)\mathbf{x}_0$, note that

$$\mathbf{a} - \mathbf{x}_0 = \frac{1}{1-t}(a - \mathbf{x}_t).$$

Therefore, $\mathbf{a} - \mathbf{x}_0$ is deterministically decided by \mathbf{x}_t and s . This yields

$$v^*(x, t | s) = \mathbb{E}[a - \mathbf{x}_0 | \mathbf{x}_t = x] = \frac{1}{1-t}(a - x).$$

Therefore, we have

$$dz_t = v^*(z_t, t | s) = \frac{1}{1-t}(a - z_t)dt.$$

Solving ODE this yields

$$z_t = ta + (1-t)z_0 = (1-t)v^*(z_0, 0 | s).$$

Differentiating it also yields

$$z_t = (a - z_0)dt.$$

We also have $\sigma^2(x, t | s)$ again because $a - \mathbf{x}_0$ is deterministic given \mathbf{x}_t and s :

$$\sigma^2(x, t | s) = \text{var}(\mathbf{a} - \mathbf{x}_0 | \mathbf{x}_t = x, s) = 0.$$

A.2 Proof of Proposition 3.3

PROOF 2. Following the property of rectified flow, the distribution of $\mathbf{x}_1 = ta + (1-t)\mathbf{x}_0$ coincides with p_t for all $t \in [0, 1]$. Hence, we can assume that $z_t = \mathbf{x}_t \sim p_t^*$. In this case, we have $z_{t+\epsilon_t} = \mathbf{x}_t + \epsilon_t v^*(z_t, t | s)$ and $\mathbf{x}_{t+\epsilon_t} = \mathbf{x}_t + \epsilon_t(\mathbf{a} - \mathbf{x}_0)$. We have

$$\begin{aligned} & W_2(p_{t+\epsilon_t}^*, p_{t+\epsilon_t})^2 \\ & \leq \mathbb{E} \left[\|\mathbf{z}_{t+\epsilon_t} - \mathbf{x}_{t+\epsilon_t}\|_2^2 \right] \\ & = \mathbb{E} \left[\mathbb{E} \left[\|\mathbf{z}_{t+\epsilon_t} - \mathbf{x}_{t+\epsilon_t}\|_2^2 \mid \mathbf{x}_t \right] \right] \\ & = \mathbb{E} \left[\mathbb{E} \left[\|\epsilon_t v^*(z_t, t | s) - \epsilon_t(\mathbf{a} - \mathbf{x}_0)\|_2^2 \mid \mathbf{x}_t \right] \right] \\ & = \epsilon_t^2 \mathbb{E}_{z_t \sim p_t} [\sigma^2(z_t, t | s)]. \end{aligned}$$

A.3 Proof of Proposition 3.5

PROOF 3. Assume the adaptive algorithm visits the time grid of $0 = t_0, t_1, \dots, t_N = 1$.

Define $z_t^{t_i}$ be the result when we implement the adaptive discretization algorithm upto t_i and then switch to follow the exact ODE afterward, that is, we have $dz_t^{t_i} = v_t(z_t^{t_i})dt$ for $t \geq t_i$. In this way, we have $z_t^1 = z_t$, and $z_t^0 = z_t^*$, where z_t^* is the trajectory of the exact ODE $dz_t^* = v_t^*(z_t^*)dt$.

From Lemma A.1, we have

$$\left\| z_1^{t_{i-1}} - z_1^{t_i} \right\| \leq \exp(L(1-t_i)) \left\| z_{t_i}^{t_i} - z_{t_i}^{t_{i-1}} \right\|.$$

Let $p_t^{t_i}$ be the distribution of $z_t^{t_i}$. Then we have $p_t^1 = p_t$ and $p_t^0 = p_t^*$. Then

$$\begin{aligned} W_2(p_1^{t_{i-1}}, p_1^{t_i}) & \leq \mathbb{E} \left[\left\| z_1^{t_{i-1}} - z_1^{t_i} \right\|^2 \right]^{1/2} \\ & = \exp(L(1-t_i)) \mathbb{E} \left[\left\| z_{t_i}^{t_{i-1}} - z_{t_i}^{t_i} \right\|^2 \right]^{1/2} \\ & = \exp(L(1-t_i)) \max(\eta, \epsilon_{\min}^2 M/2) \\ & = C \epsilon_{\min}^2 \exp(-Lt_i), \end{aligned}$$

where $C = \frac{1}{2} \max(M, M_\eta) \exp(L(1 - t_i))$. Here we use the bound in the proof of Proposition 3.1 and Lemma A.1. Hence,

$$\begin{aligned} W_2(p_1^*, p_1) &= \sum_{i=1}^{N_{\text{ada}}} W_2(p_1^{t_i-1}, p_1^{t_i}) \\ &\leq \sum_{i=1}^{N_{\text{ada}}} C \epsilon_{\min}^2 \exp(-Lt_i) \\ &\leq C \times \frac{N_{\text{ada}}}{N_{\text{max}}} \times \epsilon_{\min}, \end{aligned}$$

where $C = \exp(L) \max(M, M_\eta)$.

Lemma A.1. Let $\|v_t\|_{Lip} \leq L$ for $t \in [0, 1]$. Assume x_t and y_t solve $dx_t = v_t(x_t)dt$ and $dy_t = v_t(y_t)dt$ starting from x_0, y_0 , respectively. We have

$$\|x_t - y_t\| \leq \exp(Lt) \|x_0 - y_0\|, \quad \forall t \in [0, 1]. \quad (11)$$

PROOF 4.

$$\begin{aligned} \frac{d}{dt} \|x_t - y_t\|^2 &= 2(x_t - y_t)^\top (v_t(x_t) - v_t(y_t)) \\ &\leq 2L \|x_t - y_t\|^2, \end{aligned}$$

where we used $\|v_t(x_t) - v_t(y_t)\| \leq L \|x_t - y_t\|$. Using Gronwall's inequality yields the result.

Lemma A.2. Under Assumption 3.4, we have

$$\|x_{t+\epsilon} - (x_t + \epsilon v_t(x_t))\| \leq \frac{\epsilon^2 M}{2},$$

for $0 \leq t \leq \epsilon + t \leq 1$.

PROOF 5. Direct application of Taylor approximation.

A.4 Visualization of Tasks

We provide a visualization of the 2D Maze Figure 7.




Figure 7: Trajectories of 100 demonstrations for each maze.

A.5 Planner for Maze2D task

We generate the demonstration data in Maze toy using planner similar to [54]. The planner devises a path in a maze environment by calculating waypoints between the start and target points. It begins by transforming the given continuous-state space into a discretized grid representation. Employing Q-learning, it evaluates the optimal actions and subsequently computes the waypoints by performing a rollout in the grid, introducing random perturbations to the waypoints for diversity. The controller connects these waypoints in an ordered manner to form a feasible path. In runtime, it dynamically adjusts the control action based on the proximity to the next waypoint and switches waypoints when close enough, ensuring the trajectory remains adaptive and efficient.




Figure 8: Predicted variance by AdaFlow on the Maze task.

A.6 Comparative Analysis of Separate and Joint Training

In this section, we provide a comparison between the two training strategies employed in our proposed solution: separate training and joint training. Our primary objective is to investigate whether there is a substantial difference in performance and efficiency between these two training approaches.

Experiment Setup. To conduct this comparative analysis, we designed experiments using our proposed framework with both training strategies. Specifically, we consider two approaches: separate and joint training. In **Separate Training** setting, we train the variance prediction network and the policy function separately, as described in our main paper. In **Joint Training** setting, we train both the variance prediction network and the policy function simultaneously in an end-to-end manner. The goal is to assess the impact of these training strategies on the overall performance.

Results and Discussion. As shown in Table 6, the performance were consistent between the two training approaches, indicating the effectiveness of our two-stage framework in balancing policy accuracy and uncertainty estimation. Separate training exhibited faster computational speed, making it the preferred choice once the policy function was robustly trained. Joint training required more computational resources and time.

	Maze 1	Maze 2	Maze 3	Maze 4
AdaFlow (Separate)	0.98	1.00	0.96	0.86
AdaFlow (Joint)	1.00	1.00	0.96	0.88

Table 6: Performance comparison of separate training and joint training of AdaFlow in Maze tasks.

A.7 Visualization of Exact Variance.

In the main paper, we showed the variance predictions made by AdaFlow across different states within a robot’s state space. Here, we explain how we compute the *exact* variance for different states, to provide a ground truth of variance for reference. To achieve this, we first train a 1-Rectified Flow model for the task, then we can compute the exact variance by sampling:

$$\frac{1}{N_t} \frac{1}{N_z} \sum_t \sum_{z_0} \mathbb{E}[\|y - z_0 - v(z_t, t; x)\|^2], \text{ where } z_t = ty + (1 - t)z_0, (x, y) \sim p^*. \quad (12)$$

For each states, we randomly sample 10 time steps ($N_t = 10$) and 10 noises ($N_z = 10$).

A.8 Visualization of Predicted Variance on Maze task.

We present the predicted variance by AdaFlow in Figure 8.

A.9 Additional Experimental Details.

Model Architectures. For the 1D toy example, we used a MLP constructed with 5 fully connected layers and SiLU activation functions. We integrated temporal information by extending the time input into a 100-dimensional time-encoding vector through the cosine transformation of a random vector, $\text{cost} * z_T$, where z_T is sampled from a Gaussian distribution. This time feature is then concatenated

with the noise and condition (x) inputs to for time-aware predictions. The network comprises 4 hidden layers, each with 100 neurons, and the output layer predict a single y value. The dataset consists of 10000 single-dimensional samples uniformly distributed in the range $[-5, 5]$.

For navigation and robot manipulation tasks, we adopted the model architecture from Diffusion Policy [6]. For navigation task, we use the same architecture as used in Push-T task. In the RoboMimic and LIBERO experiments, we used the Diffusion Policy-C architecture. To ensure a fair comparison across different methods, we maintained a consistent architecture for all methods in our experiments, except where specifically noted. Detailed parameters are available in Table 7.

Hyperparameter	1D Toy			Maze			RoboMimic & LIBERO		
	RF & AdaFlow	BC	Diffusion Policy	RF & AdaFlow	BC	Diffusion Policy	RF & AdaFlow	BC	Diffusion Policy
Learning rate	1e-2	1e-2	1e-2	1e-4	1e-4	1e-4	1e-4	1e-4	1e-4
Optimizer	Adam	Adam	Adam	AdamW	AdamW	AdamW	AdamW	AdamW	AdamW
β_1	0.9	0.9	0.9	0.9	0.9	0.9	0.95	0.95	0.95
β_2	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999
Weight decay	0	0	0	1e-6	1e-6	1e-6	1e-6	1e-6	1e-6
Batch size	1000	1000	1000	256	256	256	64	64	64
Epochs	200	200	400	200	200	200	500(L) / 3000(RM)	500(L) / 3000(RM)	500(L) / 3000(RM)
Learning rate scheduler	cosine	cosine	cosine	cosine	cosine	cosine	cosine	cosine	cosine
EMA decay rate	-	-	-	0.9999	0.9999	0.9999	0.9999	0.9999	0.9999
Number of training time steps	-	-	100	-	-	20	-	-	100
Number of Inference time steps	100 (RF)	-	100(DDPM)	-	-	20(DDPM)	-	-	100(DDPM)
η	0.1	-	-	1.5	-	-	1.0	-	-
ϵ_{\min}	5	-	-	5	-	-	10	-	-
Action prediction horizon	-	-	-	16	16	16	16	16	16
Number of observation input	-	-	-	2	2	2	2	2	2
Action execution horizon	-	-	-	8	8	8	8	8	8
Observation input size	1	1	1	4 (Single-task) / 6(Multi-task)			76 × 76	76 × 76	76 × 76

Table 7: Hyperparameters used for training AdaFlow and baseline models.

Implementation of Baselines. In our studies, **BC** was implemented as a baseline, applying behavior cloning in its most straightforward form and using a Mean Squared Error loss function between the predicted and ground truth actions. The implementations for DDPM and DDIM remained consistent with those outlined in [6]. Across all experiments, consistency was maintained regarding architecture, input, and output, with all methods adhering to a similar experimental pipeline. We just use a 4 layer MLP with SiLU activation for the variance prediction, with hidden dimension of 512, which is a very small network whose computational overhead can be neglected compared to the full model.

Implementation of Vairance Prediction Network. In the 1D toy experiment, we designed the variance prediction network as a 4-layer MLP, mirroring the main model’s architecture for simplicity. In theory, the variance estimation network takes the same input as rectified flow model, so its input can be just the intermediate features extracted by the main model. Hence in the navigation and manipulation experiments, the inputs of variance prediction networks are the bottle-neck features extracted by the U-Net model.

Training on RoboMimic. Training Diffusion Models on RoboMimic is very resource-intensive. Training and evaluating a Transport task requires over a month of GPU hours. More complex tasks, such as ToolHang, can demand up to three times longer ¹Given the challenges in replicating the results from [6], we opted to start with their open-sourced pretrained model. We then fine-tuned the baselines and our method for 500 epochs and subsequently compared the performance of different models.

A.10 Comparison with standard Rectified Flow.

For the purpose of policy learning, we can consider standard Rectified Flow as a subset of our method, which can be recovered with specific choices of η and ϵ_{\min} . In this section, we compare our approach with the standard Rectified Flow, particularly focusing on the generation within a single step. Standard Rectified Flow requires a reflow or distillation stage to straighten the ODE process. During this reflow stage, the model simulates data using the initial 1-Rectified Flow. These data are then used in distillation training, resulting in what is termed a 2-Rectified Flow. Theoretically, a 2-Rectified Flow is capable of producing a straight generation trajectory, which enables one-step generation. In contrast, the 1-Rectified Flow tends to be less straight, necessitating multiple steps for sample generation.

In Table 8, we compare the performance of 1-Rectified Flow, 2-Rectified Flow, and our method in the maze task. Furthermore, Figure 9 illustrates the trajectories produced by both standard Rectified Flow

¹See this link

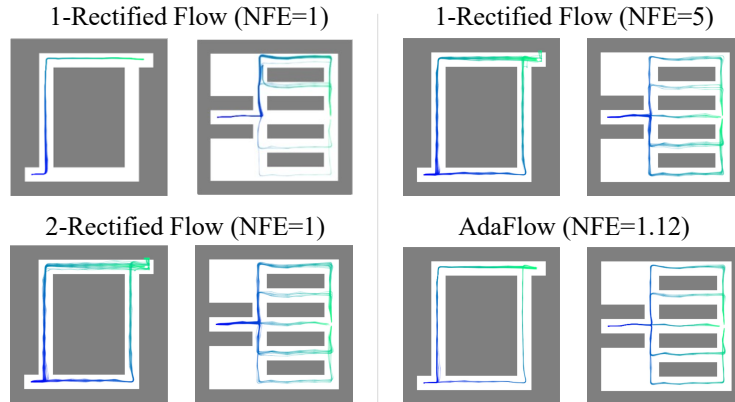


Figure 9: **Generated trajectories.** We visualize the trajectories generated by standard Rectified Flow and AdaFlow, with the agent’s starting point remaining fixed. 0

and our method. It’s evident that the standard 1-Rectified Flow struggles to generate a diverse range of actions in a single step. In contrast, our method is able to produce diverse behaviors in nearly one step. This efficiency is attributed to our method’s ability to estimate the variance across different states, identifying those that require multi-step generation.

	NFE↓	Maze 1	Maze 2	Maze 3	Maze 4
1-RF	1	1.00	1.00	0.98	0.80
1-RF	5	0.82	1.00	0.94	0.80
2-RF (reflow)	1	0.82	1.00	1.00	0.80
AdaFlow ($\eta = 1.5$)	1.56	0.98	1.00	0.96	0.86
AdaFlow ($\eta = 2.5$)	1.12	1.00	1.00	0.94	0.78

Table 8: Performance on maze navigation tasks. The table showcases the success rate (**SR**) for each model across different maze complexities. The highest success rate for each task are highlighted in **bold**. Note that 2-RF needs an expensive distillation training stage.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics [https://neurips.cc/public/EthicsGuidelines?](https://neurips.cc/public/EthicsGuidelines)

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to

generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.

- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.