

DeFlow: Decoupling Behavior-Prior Modeling and Value Maximization for Offline Policy Extraction

Zhancun Mu Chi Zhang

School of Intelligence Science and Technology
Peking University

muzhancun@stu.pku.edu.cn chizhang.cz@pku.edu.cn

muzhancun.github.io/publications/deflow/

Abstract

We present DeFlow, a decoupled offline RL framework for extracting high-value actions from a learned multi-step flow prior. Directly optimizing iterative generative policies typically requires backpropagation through ODE solvers, while shortcut policies trade off expressivity, inference cost, and policy-improvement objectives. DeFlow keeps the flow model as a behavior prior and trains a lightweight action-conditioned residual module for value improvement under an adaptive trust-region penalty. This design bypasses solver differentiation and regularizes the refined policy toward the learned behavior prior. Empirically, DeFlow is competitive with recent generative offline RL baselines, with the clearest gains on selected multimodal manipulation and offline-to-online settings.




Figure 1. DeFlow Framework. (a) **Iterative Flow Prior:** A multi-step flow policy models multimodal behavior. (b) **Shortcut Extraction:** One-step distillation enables direct Q-optimization but may compress the generated action distribution. (c) **DeFlow:** We keep the iterative prior (Blue) and train a lightweight residual module (Orange) to perform local value improvement while penalizing large deviations from the prior.

1 Introduction

Offline reinforcement learning (RL) must improve a policy while avoiding actions that are weakly supported by the static dataset [16]. This becomes harder on large multimodal benchmarks such as OGBench [23], where standard unimodal actors often cannot represent the diverse behavior distribution. Diffusion and flow-matching policies [18, 30] are attractive in this regime because their iterative sampling procedures can model complex action distributions.

However, extracting a high-value policy from these generative models presents a stark *expressivity-optimizability dilemma*. Directly differentiating the Q-objective through iterative generation requires solver backpropagation, which can be costly; alternative objectives such as Q-score matching avoid this full differentiation [27]. Shortcut methods such as Flow Q-Learning (FQL) [24] and SORL [7] reduce this burden through one-step or shortcut parameterizations, sometimes together with inference-time scaling. These methods are strong baselines, but they also motivate a complementary question: can we keep a multi-step behavior prior intact and perform value improvement through a small local extraction module?

In this work, we study this decoupled extraction route. We introduce **DeFlow**, a framework that separates behavior modeling from value-guided policy refinement.

Specifically, DeFlow comprises two distinct components: (1) a **Flow Matching Policy** optimized via supervised behavior modeling; and (2) a lightweight **Action Refinement Module** that maximizes Q-values via deterministic policy gradients (DDPG [17]). Instead of updating the flow model through the ODE solver, the refinement module learns an action-conditioned residual shift on the generated action (Figure 1c).

A critical challenge in this decoupled setting is discouraging large residual moves that exploit unreliable critic values far from the learned prior. We address this with an adaptive Lagrangian penalty on the residual norm. The constraint should be interpreted as a trust-region surrogate around the learned flow prior, not as an exact guarantee of support preservation.

Some recent works [5, 35] also decouple large expressive policies and small RL heads for policy improvement. DeFlow instantiates this pattern for offline policy extraction with a flow-matching prior and an instance-level residual conditioned on the sampled action.

Our contribution is the specific combination of a multi-step behavior prior, proposal-conditioned residual extraction, and adaptive residual regularization for offline RL. Flow matching, residual learning, dual adaptation, and BPTT avoidance are established ingredients rather than separate novelty claims. We study how this combination behaves relative to shortcut extraction, with benchmark comparisons and diagnostics of residual attribution, capacity, and budget sensitivity.

The results support a scoped claim: DeFlow is competitive with recent generative offline RL methods, while its strongest gains appear in selected multimodal manipulation tasks and in offline-to-online settings where the same residual extraction protocol is reused.

2 Preliminaries

We consider the standard Markov Decision Process (MDP) [28] defined by the tuple $(\mathcal{S}, \mathcal{A}, r, p, \gamma)$, where $\mathcal{S}$ and $\mathcal{A}$ denote the state and action spaces, $r(s, a)$ the reward function, $p(s'|s, a)$ the transition dynamics, and $\gamma \in [0, 1)$ the discount factor.

2.1 Offline RL as Constrained Optimization

In offline RL, we are given a static dataset $\mathcal{D} = \{(s, a, r, s')\}$ collected by a behavior policy π_β . The goal is to learn a policy π_θ that maximizes the expected cumulative return. However, due to the distributional shift, simply maximizing values on out-of-distribution (OOD) actions leads to extrapolation errors. Following Levine et al. [16], offline RL is often cast as a constrained optimization problem:

$$\begin{aligned} & \max_{\pi_\theta} \mathbb{E}_{s \sim \mathcal{D}, a \sim \pi_\theta(\cdot|s)} [Q(s, a)] \\ & \text{s.t.} \quad \mathbb{D}(\pi_\theta(\cdot|s) || \pi_\beta(\cdot|s)) \leq \epsilon, \end{aligned} \tag{1}$$

where $Q(s, a)$ is the state-action value function and $\mathbb{D}$ is a divergence measure (e.g., KL divergence or Wasserstein distance) used to discourage unsupported policy shifts.

Critic Learning. To evaluate the policy, we learn a critic Q_{ξ} by minimizing the standard Bellman error [21]. The objective is:

$$\mathcal{L}_Q(\xi) = \mathbb{E}_{(s,a,r,s') \sim \mathcal{D}} \left[(Q_{\xi}(s, a) - r - \gamma \hat{Q}_{\xi}(s', a'))^2 \right]. \quad (2)$$

where $a' \sim \pi_{\theta}(\cdot | s')$ and $\hat{Q}_{\xi}$ utilizes a target network for stability. Critic parameters ξ , prior parameters ψ , and residual parameters ϕ are distinct.

The Actor’s Dilemma. Standard methods typically convert the constrained problem in Eq. (1) into an unconstrained one via a Lagrangian relaxation with a fixed penalty coefficient α :

$$\mathcal{L}_{\pi}(\theta) = \mathbb{E}_{s \sim \mathcal{D}} \left[-\mathbb{E}_{a \sim \pi_{\theta}} [Q_{\xi}(s, a)] + \alpha \mathcal{L}_{\text{BC}}(\pi_{\theta}, \mathcal{D}) \right]. \quad (3)$$

Gaussian policies allow efficient reparameterized actor updates but have limited ability to represent multimodal behavior. Iterative generative policies offer greater expressivity, while direct differentiation through their sampling chains increases the optimization cost. Alternative training objectives can avoid this differentiation. A fixed regularization coefficient also introduces a task-dependent trade-off between value improvement and prior proximity [1, 24].

2.2 Conditional Flow Matching

To address the modeling challenge, we leverage Conditional Flow Matching (CFM) [18], which offers a simulation-free alternative to diffusion training. CFM learns a time-dependent vector field $v_{\psi}(t, x) : [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ that pushes a simple prior distribution p_0 (e.g., $\mathcal{N}(0, I)$) to the data distribution p_1 (e.g., $\pi_{\beta}(a|s)$).

The generative process is defined by an Ordinary Differential Equation (ODE):

$$\frac{d}{dt} \Psi_t(x) = v_{\psi}(t, \Psi_t(x)), \quad \Psi_0(x) = x \sim p_0, \quad (4)$$

where Ψ_t is the flow map. We adopt the straight conditional path $x_t = (1 - t)x_0 + tx_1$ [19]. Noise and data endpoints are sampled independently; this does not assert an optimal coupling between the endpoint distributions. The vector field is trained via the regression objective:

$$\mathcal{L}_{\text{FM}}(\psi) = \mathbb{E}_{t \sim \mathcal{U}[0,1], x_0 \sim p_0, x_1 \sim \mathcal{D}} \|v_{\psi}(t, x_t) - (x_1 - x_0)\|_2^2. \quad (5)$$

Flow as a Policy. In our context, the flow model serves as a state-conditional generator. We denote the generated action as a deterministic function of the initial noise z and state s :

$$a_{\text{flow}} = \text{ODESolve}(v_{\psi}, s, z), \quad z \sim \mathcal{N}(0, I). \quad (6)$$

In our framework, this flow model acts as the learned behavior prior, decoupling distribution modeling from the policy improvement step described next.

3 DeFlow: Disentangling Fidelity and Optimality

DeFlow disentangles behavior modeling from value-guided extraction. Instead of training one generative policy with both behavior-cloning and Q-maximization gradients, we keep the multi-step flow as a proposal distribution and learn a small residual extractor on top of its samples. Algorithm 1 provides an overview of our method.

Algorithm 1 DeFlow

```

1: function  $\mu_\psi(s, z)$ 
2:   for  $m = 0, \dots, M - 1$  do
3:      $z \leftarrow z + v_\psi(m/M, s, z)/M$ 
4:   end for
5:   return  $\text{clip}(z, -1, 1)$ 
6: end function
7: function  $\text{ACT}(s)$ 
8:    $z \sim \mathcal{N}(0, I_d), \quad a_b \leftarrow \mu_\psi(s, z)$ 
9:   return  $\text{clip}(a_b + f_\phi(s, a_b), -1, 1)$ 
10: end function
11: while not converged do
12:   Sample minibatch  $\{(s, a, r, s')\} \sim \mathcal{D}$ 
13:   Update  $Q_\xi$  with Bellman targets using  $a' = \text{ACT}(s')$ 
14:   Update  $v_\psi$  on  $\|v_\psi(t, s, x^t) - (x^1 - x^0)\|_2^2$ 
15:    $a_b \leftarrow \text{sg}[\mu_\psi(s, z)], \quad \Delta \leftarrow f_\phi(s, a_b)$ 
16:   Update  $f_\phi$  on  $-Q_\xi^{\text{norm}}(s, \text{clip}(a_b + \Delta, -1, 1)) + \text{sg}[\alpha]\|\Delta\|_2^2/d$ 
17:   Update  $\eta = \log \alpha$  on  $-\eta \text{sg}(\|\Delta\|_2^2/d - \delta)$ 
18: end while return  $\mu_\psi, f_\phi$ 

```

We formulate the policy as a composite function, where the final action a is derived from a **Base Flow** μ_ψ adjusted by a **Refinement Module** f_ϕ :

$$a = \underbrace{\text{sg}[\mu_\psi(s, z)]}_{\text{Fixed Proposal}} + f_\phi(s, \underbrace{\text{sg}[\mu_\psi(s, z)]}_{\text{Context}}), \quad z \sim \mathcal{N}(0, I) \quad (7)$$

Here, $\text{sg}[\cdot]$ denotes the stop-gradient operator, and the displayed action is before clipping to the action bounds. The flow μ_ψ provides a behavior-prior proposal, while f_ϕ , implemented as an MLP, learns a residual shift toward higher critic values.

The role of Stop-Gradient: The stop-gradient operator cuts the computational graph between the refinement module and the flow model. During the update of f_ϕ , the base flow μ_ψ is treated as a fixed stochastic generator rather than a differentiable layer. This bypasses backpropagation through the ODE solver and keeps the two optimization loops separated.

3.1 Revisiting Existing Coupling Paradigms

To better contextualize DeFlow, we briefly analyze the limitations inherent in two prevalent strategies for integrating flow models with Q-learning.

Paradigm I: End-to-End Policy Steering. Ideally, one might update a flow model π_ψ by backpropagating the Q-gradient through the ODE integration chain:

$$\nabla_\psi \mathcal{L}_{\text{RL}} = \mathbb{E}_{s,z} \left[\frac{\partial Q(s, a)}{\partial a} \cdot \frac{\partial \text{ODESolve}(v_\psi, s, z)}{\partial v_\psi} \right]. \quad (8)$$

While conceptually direct, this “steering” approach requires BPTT or the Adjoint Method across numerous integration steps. This is memory-intensive and can suffer from unstable or weak gradients through the solver.

Paradigm II: Single-step Shortcut Distillation. Alternatively, methods like FQL [24] attempt to bypass the ODE complexity by distilling the iterative flow into a single-step policy $\pi_\theta(s, z)$. The actor is typically trained with a composite objective:

$$\mathcal{L}_\pi(\omega) = \underbrace{\mathbb{E}_{s,a^\pi} [-Q_\xi(s, a^\pi)]}_{\text{Q maximization}} + \underbrace{\alpha \mathcal{L}_{\text{Distill}}(\omega)}_{\text{BC regularization}}. \quad (9)$$

While efficient at inference, this paradigm can compress the complex, multimodal distribution represented by the original flow. Furthermore, the scalar α introduces a sensitive hyperparameter trade-off: a large α limits improvement, while a small α allows larger deviations from the behavior prior.

3.2 The DeFlow Approach

DeFlow addresses these challenges by maintaining the *Multi-step Flow* as a behavior prior while employing a *Residual Refinement* for optimization.

The Base Flow μ_ψ is updated solely via the Flow Matching objective $\mathcal{L}_{\text{FM}}(\psi)$, so Q-learning gradients do not directly alter the prior.

Inspired by trust-region methods, we formulate refinement using a fixed, task-level budget δ on the mean squared residual per action coordinate. Let d be the action dimension:

$$\max_{f_\phi} \mathbb{E}_{s,z}[Q(s, a_{\text{flow}} + \Delta a)] \quad \text{s.t.} \quad \mathbb{E}_{s,z}[\|\Delta a\|_2^2/d] \leq \delta, \quad (10)$$

where $\Delta a = f_\phi(s, a_{\text{flow}})$ and $a_{\text{flow}} = \text{sg}[\mu_\psi(s, z)]$. Here s is sampled from replay. The division by d matches the implementation’s coordinate-averaged squared error; δ is a squared-displacement budget, not a per-action radius. We apply action clipping when evaluating the critic and executing the policy. From an optimal-transport view [26, 32], the residual map $T_s(a_b) = a_b + f_\phi(s, a_b)$ is a Monge-style transport from the learned flow prior to the refined policy. Under the coupling induced by sampling $a_b \sim \mu_\psi(\cdot|s)$ and mapping it through T_s , $\mathbb{E}\|\Delta a\|_2^2$ upper-bounds the squared W_2 distance to the learned prior; we state this formally in Section A.

Relation to residual policies and SORL. Classical residual policy learning [29] usually adds a state-dependent correction to a base controller. DeFlow instead uses an instance-level residual $f_\phi(s, a_b)$, allowing different sampled actions from the same state to receive different local corrections. SORL [7] is an important step-size-conditioned shortcut-flow baseline: its actor is trained with flow matching, shortcut-consistency, and Q-maximization losses, and its evaluation can sweep shortcut step counts with optional critic-ranked candidates. We include it in the main benchmark, while DeFlow focuses on local residual extraction around a prior held fixed during each actor update. The prior continues to receive its separate behavior-modeling updates. Action conditioning is not a mode classifier: different proposals can receive different corrections, which may be mode-specific when modes are separated in action space.

To solve Equation (10), we employ the method of Lagrange multipliers. We first apply **Q-Normalization** to the RL objective to mitigate the scale variance of Q-values across different environments. We then introduce a learnable Lagrangian multiplier α (where $\alpha > 0$). The loss function for the refinement module f_ϕ is defined as:

$$\mathcal{L}_{\text{Refine}}(\phi) = \mathbb{E}_{s,z} \left[-\frac{Q(s, \text{clip}(a_{\text{flow}} + \Delta a, -1, 1))}{\text{sg}(\mathbb{E}_{\text{batch}}|Q|)} + \text{sg}[\alpha] \|\Delta a\|_2^2/d \right]. \quad (11)$$

We adapt the positive penalty $\alpha = \exp(\eta)$ using a log-space dual update. In the implementation, the loss minimized with respect to η is:

$$\mathcal{L}_{\text{dual}}(\eta) = \eta \left(\delta - \text{sg} \left[\mathbb{E}_{s,z} \|\Delta a\|_2^2/d \right] \right). \quad (12)$$

When the empirical residual exceeds δ , gradient descent increases η and hence α ; otherwise it decreases the penalty. This resembles SAC’s log-temperature adaptation [11], but controls residual movement rather than entropy. It is a surrogate dual update, not a claim of solving the nonconvex constrained problem exactly. Adaptive α reduces fixed-coefficient tuning; δ still requires selection. Minibatch regularization neither enforces a hard pointwise bound nor guarantees true behavior-support membership.

4 Experiments

We evaluate DeFlow on 73 offline RL tasks and 15 offline-to-online (O2O) tasks. The goal is to test whether local residual extraction around a multi-step flow prior is competitive with recent generative offline RL methods, and where it helps or underperforms.

Table 1. Offline RL results on OGBench and D4RL. We report DeFlow against Gaussian-policy and flow-based/generative policy-extraction baselines. SORL results are included where aligned OGBench numbers are available; missing entries are marked by “-” and are not used when identifying row-wise best scores. Results are averaged over 8 seeds for state tasks and 4 seeds for visual tasks when applicable. Bold entries are at least $0.95\times$ the best reported score in their row. See Table 15 for the full task-level version of the non-SORL baselines. Counts in parentheses indicate the number of tasks.

Task Category	Gaussian Policies			Flow / Generative Extraction					
	BC	IQL	ReBRAC	FAWAC	FBRAC	IFQL	SORL	FQL	DeFlow
antmaze-large (5)	11 $\pm$ 1	53 $\pm$ 3	81 $\pm$ 5	6 $\pm$ 1	60 $\pm$ 6	28 $\pm$ 5	89 $\pm$ 2	79 $\pm$ 3	81 $\pm$ 3
antmaze-giant (5)	0 $\pm$ 0	4 $\pm$ 1	26 $\pm$ 8	0 $\pm$ 0	4 $\pm$ 4	3 $\pm$ 2	9 $\pm$ 6	9 $\pm$ 6	12 $\pm$ 5
humanoidmaze-medium (5)	2 $\pm$ 1	33 $\pm$ 2	22 $\pm$ 8	19 $\pm$ 1	38 $\pm$ 5	60 $\pm$ 14	64 $\pm$ 4	58 $\pm$ 5	48 $\pm$ 4
humanoidmaze-large (5)	1 $\pm$ 0	2 $\pm$ 1	2 $\pm$ 1	0 $\pm$ 0	2 $\pm$ 0	11 $\pm$ 2	5 $\pm$ 2	4 $\pm$ 2	5 $\pm$ 2
antsoccer-arena (5)	1 $\pm$ 0	8 $\pm$ 2	0 $\pm$ 0	12 $\pm$ 0	16 $\pm$ 1	33 $\pm$ 6	69 $\pm$ 2	60 $\pm$ 2	67 $\pm$ 3
cube-single (5)	5 $\pm$ 1	83 $\pm$ 3	91 $\pm$ 2	81 $\pm$ 4	79 $\pm$ 7	79 $\pm$ 2	97 $\pm$ 1	96 $\pm$ 1	96 $\pm$ 2
cube-double (5)	2 $\pm$ 1	7 $\pm$ 1	12 $\pm$ 1	5 $\pm$ 2	15 $\pm$ 3	14 $\pm$ 3	25 $\pm$ 3	29 $\pm$ 2	40 $\pm$ 5
scene (5)	5 $\pm$ 1	28 $\pm$ 1	41 $\pm$ 3	30 $\pm$ 3	45 $\pm$ 5	30 $\pm$ 3	57 $\pm$ 2	56 $\pm$ 2	51 $\pm$ 3
puzzle-3x3 (5)	2 $\pm$ 0	9 $\pm$ 1	21 $\pm$ 1	6 $\pm$ 2	14 $\pm$ 4	19 $\pm$ 1	—	30 $\pm$ 1	43 $\pm$ 4
puzzle-4x4 (5)	0 $\pm$ 0	7 $\pm$ 1	14 $\pm$ 1	1 $\pm$ 0	13 $\pm$ 1	25 $\pm$ 5	—	17 $\pm$ 2	11 $\pm$ 2
D4RL antmaze (6)	17	57	78	44 $\pm$ 3	64 $\pm$ 7	65 $\pm$ 7	—	84 $\pm$ 3	83 $\pm$ 2
D4RL adroit (12)	48	53	59	48 $\pm$ 1	50 $\pm$ 2	52 $\pm$ 1	—	52 $\pm$ 1	52 $\pm$ 1
Visual manipulation (5)	-	42 $\pm$ 4	60 $\pm$ 2	-	22 $\pm$ 2	50 $\pm$ 5	—	65 $\pm$ 2	69 $\pm$ 2

¹ Pixel tasks selectively benchmark methods that are strong on state-based OGBench tasks. ² SORL numbers are taken from the aligned OGBench comparison used in DROL; SORL uses step-size-conditioned shortcut flows and inference-time shortcut-step scaling.

4.1 Experimental Setup

Benchmarks. We use OGBench [23] as the primary benchmark, including 50 state-based reward-conditioned single-task tasks and 5 visual manipulation tasks. We also evaluate on 18 D4RL tasks [9].

Baselines. We compare against Gaussian-policy baselines (BC, IQL [14], ReBRAC [31]) and flow-based extraction baselines (FAWAC, FBRAC, IFQL, FQL [24]). We also include SORL [7] wherever aligned OGBench results are available. SORL is not an apples-to-apples matched-inference baseline because its shortcut-flow implementation conditions on step size and evaluates multiple shortcut step counts, with optional best-of- N critic selection; nevertheless, it is an important recent reference point for efficient generative offline RL.

Offline-to-online Experiments. For O2O experiments, we further compare against Cal-QL [22] and RLPD [3], which are specifically engineered for online adaptation. We adapt DeFlow for this setting by simply augmenting the offline dataset $\mathcal{D}$ with an online replay buffer during training.

Evaluation Protocol. We follow the OGBench protocol [23], reporting performance after a fixed number of gradient steps. Results are averaged over 8 seeds for state-based tasks and 4 seeds for pixel-based tasks. We **bold** scores that are within 95% of the best reported performance in each row. The main results use a fixed four-layer residual; we do not select residual depth using the additional diagnostics. Their four-seed protocols and final-checkpoint evaluation budgets are reported separately in Section C.4.

4.2 Results and Analysis

Offline results. Table 1 shows a mixed but informative comparison. SORL is very strong on navigation and scene-style tasks, including antmaze-large, humanoidmaze-medium, antsoccer, and scene. DeFlow is most competitive on manipulation tasks where local multimodal precision matters: it improves over FQL and the reported SORL numbers on cube-double, and it remains strong on puzzle-3x3 and visual manipulation. On D4RL, DeFlow is close to FQL on AntMaze and Adroit rather than clearly dominating. These results support a scoped claim: residual extraction around a multi-step prior is a useful alternative to shortcut extraction, especially for selected multimodal manipulation regimes.

Table 2. Offline-to-online RL results. The results are averaged over 8 seeds unless otherwise mentioned. Entries show offline $\rightarrow$ online performance; OGBench rows use the default single-task environment.

Task	IQL	ReBRAC	Cal-QL	RLPD	IFQL	FQL	DeFlow
humanoidmaze-medium	21 $\pm$ 13 $\rightarrow$ 16 $\pm$ 8	16 $\pm$ 20 $\rightarrow$ 1 $\pm$ 1	0 $\pm$ 0 $\rightarrow$ 0 $\pm$ 0	0 $\pm$ 0 $\rightarrow$ 8 $\pm$ 10	56 $\pm$ 35 $\rightarrow$ 82 $\pm$ 20	12 $\pm$ 7 $\rightarrow$ 22 $\pm$ 12	13 $\pm$ 5 $\rightarrow$ 65 $\pm$ 12
antsoccer-arena	2 $\pm$ 1 $\rightarrow$ 0 $\pm$ 0	0 $\pm$ 0 $\rightarrow$ 0 $\pm$ 0	0 $\pm$ 0 $\rightarrow$ 0 $\pm$ 0	0 $\pm$ 0 $\rightarrow$ 0 $\pm$ 0	26 $\pm$ 15 $\rightarrow$ 39 $\pm$ 10	28 $\pm$ 8 $\rightarrow$ 86 $\pm$ 5	44 $\pm$ 8 $\rightarrow$ 86 $\pm$ 3
cube-double	0 $\pm$ 1 $\rightarrow$ 0 $\pm$ 0	6 $\pm$ 5 $\rightarrow$ 28 $\pm$ 28	0 $\pm$ 0 $\rightarrow$ 0 $\pm$ 0	0 $\pm$ 0 $\rightarrow$ 0 $\pm$ 0	12 $\pm$ 9 $\rightarrow$ 40 $\pm$ 5	40 $\pm$ 11 $\rightarrow$ 92 $\pm$ 3	48 $\pm$ 13 $\rightarrow$ 93 $\pm$ 4
scene	14 $\pm$ 11 $\rightarrow$ 10 $\pm$ 9	55 $\pm$ 10 $\rightarrow$ 100 $\pm$ 0	1 $\pm$ 2 $\rightarrow$ 50 $\pm$ 53	0 $\pm$ 0 $\rightarrow$ 100 $\pm$ 0	0 $\pm$ 1 $\rightarrow$ 60 $\pm$ 39	82 $\pm$ 11 $\rightarrow$ 100 $\pm$ 1	58 $\pm$ 14 $\rightarrow$ 100 $\pm$ 1
puzzle-4x4	5 $\pm$ 2 $\rightarrow$ 1 $\pm$ 1	8 $\pm$ 4 $\rightarrow$ 14 $\pm$ 35	0 $\pm$ 0 $\rightarrow$ 0 $\pm$ 0	0 $\pm$ 0 $\rightarrow$ 100 $\pm$ 1	23 $\pm$ 6 $\rightarrow$ 19 $\pm$ 33	8 $\pm$ 3 $\rightarrow$ 38 $\pm$ 52	4 $\pm$ 4 $\rightarrow$ 100 $\pm$ 0
antmaze-umaze-v2	77 $\rightarrow$ 96	98 $\rightarrow$ 75	77 $\rightarrow$ 100	0 $\pm$ 0 $\rightarrow$ 98 $\pm$ 3	94 $\pm$ 5 $\rightarrow$ 96 $\pm$ 2	97 $\pm$ 2 $\rightarrow$ 99 $\pm$ 1	96 $\pm$ 3 $\rightarrow$ 99 $\pm$ 2
antmaze-umaze-diverse-v2	60 $\rightarrow$ 64	74 $\rightarrow$ 98	32 $\rightarrow$ 98	0 $\pm$ 0 $\rightarrow$ 94 $\pm$ 5	69 $\pm$ 20 $\rightarrow$ 93 $\pm$ 5	79 $\pm$ 16 $\rightarrow$ 100 $\pm$ 1	87 $\pm$ 6 $\rightarrow$ 99 $\pm$ 1
antmaze-medium-play-v2	72 $\rightarrow$ 90	88 $\rightarrow$ 98	72 $\rightarrow$ 99	0 $\pm$ 0 $\rightarrow$ 98 $\pm$ 2	52 $\pm$ 19 $\rightarrow$ 93 $\pm$ 2	77 $\pm$ 7 $\rightarrow$ 97 $\pm$ 2	76 $\pm$ 5 $\rightarrow$ 97 $\pm$ 2
antmaze-medium-diverse-v2	64 $\rightarrow$ 92	85 $\rightarrow$ 99	62 $\rightarrow$ 98	0 $\pm$ 0 $\rightarrow$ 97 $\pm$ 2	44 $\pm$ 26 $\rightarrow$ 89 $\pm$ 4	55 $\pm$ 19 $\rightarrow$ 97 $\pm$ 3	61 $\pm$ 9 $\rightarrow$ 98 $\pm$ 2
antmaze-large-play-v2	38 $\rightarrow$ 64	68 $\rightarrow$ 32	32 $\rightarrow$ 97	0 $\pm$ 0 $\rightarrow$ 93 $\pm$ 5	64 $\pm$ 14 $\rightarrow$ 80 $\pm$ 5	66 $\pm$ 40 $\rightarrow$ 84 $\pm$ 30	76 $\pm$ 6 $\rightarrow$ 95 $\pm$ 4
antmaze-large-diverse-v2	27 $\rightarrow$ 64	67 $\rightarrow$ 72	44 $\rightarrow$ 92	0 $\pm$ 0 $\rightarrow$ 94 $\pm$ 3	69 $\pm$ 6 $\rightarrow$ 86 $\pm$ 5	75 $\pm$ 24 $\rightarrow$ 94 $\pm$ 3	76 $\pm$ 9 $\rightarrow$ 96 $\pm$ 2
pen-cloned-v1	84 $\rightarrow$ 102	74 $\rightarrow$ 138	-3 $\rightarrow$ -3	3 $\pm$ 2 $\rightarrow$ 120 $\pm$ 10	77 $\pm$ 7 $\rightarrow$ 107 $\pm$ 10	53 $\pm$ 14 $\rightarrow$ 149 $\pm$ 6	64 $\pm$ 13 $\rightarrow$ 142 $\pm$ 7
door-cloned-v1	1 $\rightarrow$ 20	0 $\rightarrow$ 102	-0 $\rightarrow$ -0	0 $\pm$ 0 $\rightarrow$ 102 $\pm$ 7	3 $\pm$ 2 $\rightarrow$ 50 $\pm$ 15	0 $\pm$ 0 $\rightarrow$ 102 $\pm$ 5	0 $\pm$ 0 $\rightarrow$ 99 $\pm$ 2
hammer-cloned-v1	1 $\rightarrow$ 57	7 $\rightarrow$ 125	0 $\rightarrow$ 0	0 $\pm$ 0 $\rightarrow$ 128 $\pm$ 29	4 $\pm$ 2 $\rightarrow$ 60 $\pm$ 14	0 $\pm$ 0 $\rightarrow$ 127 $\pm$ 17	7 $\pm$ 3 $\rightarrow$ 106 $\pm$ 10
relocate-cloned-v1	0 $\rightarrow$ 0	1 $\rightarrow$ 7	-0 $\rightarrow$ -0	0 $\pm$ 0 $\rightarrow$ 2 $\pm$ 2	-0 $\pm$ 0 $\rightarrow$ 5 $\pm$ 3	0 $\pm$ 1 $\rightarrow$ 62 $\pm$ 8	1 $\pm$ 0 $\rightarrow$ 35 $\pm$ 8




Figure 2. Visualization of Action Landscapes. Overlay of Q-value contours and action samples. **Gray:** Base Flow samples (data proxy); **Blue:** FQL (one-step); **Red:** DeFlow (ours). Arrows indicate the refinement trajectory. DeFlow performs local residual updates around base-flow proposals, illustrating smaller local shifts in this representative state.

Mechanism visualization. Figure 2 visualizes critical states in cube-double-play. In these visualized states, FQL samples can fall between nearby modes or toward high-Q regions, while DeFlow starts from multi-step flow samples and applies a regularized residual. This diagnostic is not a formal support guarantee, but it illustrates the intended role of the residual penalty: keep value improvement local to the learned proposal distribution.

Offline-to-online and diagnostics. Table 2 and Figure 3 summarize O2O results. Using the same offline configuration, DeFlow is competitive with FQL on many tasks and is especially strong on cube-double and puzzle-4x4; it is weaker on some cloned Adroit tasks, highlighting dependence on prior coverage and critic quality. Additional diagnostics in Sections C.1 and C.2 report δ sensitivity, an imbalanced human/expert dataset, action-conditioned refinement ablations, and compatibility with floq-style value modeling. Sections C.3 and C.4 add prior-only evaluation and residual-capacity sensitivity; Section D documents the runtime comparison.




Figure 3. Offline-to-online RL results (8 seeds).

4.3 Interpreting the Empirical Pattern

The SORL comparison changes how the results should be read: DeFlow is not uniformly stronger, and the table points to regimes where different extraction mechanisms are useful. SORL and shortcut-style methods can allocate extra evaluation compute across shortcut step counts or critic-ranked candidates, which fits several navigation and scene-style tasks. The strong SORL numbers on antmaze-large, humanoidmaze-medium, antsoccer, and scene are consistent with this view.

DeFlow is intended for a narrower regime: the learned prior already produces useful local modes, but value extraction must make small, sample-specific corrections without collapsing those modes into a single shortcut output. This is a possible explanation for the manipulation results, not an experimentally isolated cause. For example, cube-double and puzzle-3x3 require choosing among nearby multimodal behaviors and then making precise local adjustments. A residual that is conditioned on both the state and the sampled action can move one proposal without forcing all proposals from the same state to share the same correction. This is also why the landscape visualization is included as a diagnostic rather than as a proof: it illustrates the intended local mechanism in a representative state, while the quantitative table determines where that mechanism is actually useful.

The O2O results provide a complementary stress test: we reuse the same residual architecture and budget mechanism while appending online transitions to replay. This transfers well when the offline prior already covers useful behaviors, as in cube-double and puzzle-4x4, but is weaker when prior coverage or critic quality is less reliable. The appendix diagnostics probe sensitivity and component attribution rather than serving as additional headline benchmarks.

Attribution and residual capacity. Disabling the residual at the same trained checkpoints gives 10 ± 2 versus 62 ± 12 with refinement on cube-double task1, and 0.5 ± 0.9 versus 35 ± 11 on puzzle-3x3 task2 (Table 10). These evaluations isolate the residual’s inference-time contribution conditional on the trained prior, not robustness to arbitrary prior quality. The capacity sweep (Tables 11 and 12) also shows task dependence: one layer underperforms two on puzzle-3x3 and several navigation tasks, while greater depth does not consistently help on cube-double. A possible explanation is a balance between insufficient correction capacity and fitting optimistic critic errors. We have not directly measured this causal mechanism. The fixed four-layer main benchmark remains unchanged.

4.4 Runtime and Compute Trade-offs

DeFlow separates behavior-prior sampling from value improvement. The base policy still uses a multi-step flow solver, so DeFlow does not claim the same inference profile as one-step shortcut policies; the practical question is whether the residual module adds much overhead beyond the flow proposal itself. Let M denote the number of flow-solver steps, C_v the cost of one vector-field evaluation, C_f the cost of the residual module, and C_Q the cost of one critic evaluation. A direct BPTT actor update through an iterative generator retains the solver trajectory and backpropagates through M vector-field calls, with memory proportional to the unrolled solver. Shortcut policies reduce this to a one-step or few-step extraction cost, and SORL can spend test-time compute on shortcut step counts or critic-ranked candidates. DeFlow still pays MC_v to sample the prior, but its actor-gradient path only traverses f_ϕ and Q because $a_{\text{flow}} = \text{sg}[\mu_\psi(s, z)]$.

Table 3. Compute trade-offs for generative policy extraction. M is the number of flow-solver steps; C_v , C_f , and C_Q denote vector-field, residual-module, and critic costs.

Family	Update path	Eval. budget	Proposal use
Iterative BPTT	solver + critic	MC_v	update generator
Shortcut / FQL	shortcut + critic	1–few steps	compress actor
SORL-style	shortcut + critic	k steps; optional ranking	scale inference
DeFlow	residual + critic	$MC_v + C_f$	local residual

Table 4 reports wall-clock measurements on the same NVIDIA 3090 setup used for our main experiments. The measured overhead is modest in this implementation: DeFlow is slightly slower than FQL in offline training and inference because it preserves multi-step proposal sampling, while the residual MLP adds little extra cost relative to the flow and critic updates. These timings are implementation-specific and do not establish a speed advantage over SORL or a scaling law for larger models. DeFlow can also refine several proposals and rank them by the critic, at additional sampling and critic cost; that variant is not evaluated here. Applying the same separation to a large VLA-style prior is a motivation for future work, not an empirical result of this study.

Table 4. Training and inference time comparison in milliseconds per step.

Method	Offline train	Online train	Inference
FQL	2.2 $\pm$ 1.0	5.1 $\pm$ 0.5	4.2 $\pm$ 0.7
DeFlow	2.4 $\pm$ 1.0	5.3 $\pm$ 0.4	4.5 $\pm$ 0.5
DeFlow (frozen prior)	–	5.2 $\pm$ 0.5	–

5 Related Work

Offline RL under behavior constraints. Offline RL often improves value under constraints that keep the policy close to the static dataset [16]. Conservative and behavior-regularized methods use value penalties or actor regularization [10, 15]; IQL and SQL avoid explicit out-of-distribution maximization [14, 36]. OGBench [23] exposes a harder version of the same problem: the policy class must represent multimodal behavior, but the extraction step must still avoid critic exploitation. DeFlow follows this constrained-improvement view with a residual budget around samples from a learned flow prior.

Diffusion and flow policies for offline RL. Diffusion policies and flow-matching policies are attractive for offline RL because they can model complex conditional action distributions. Value guidance and weighted policy objectives offer several ways to incorporate critic information into expressive policies [13, 20, 25, 37]. Other methods backpropagate value gradients through the iterative denoising or sampling process [4, 34], which provides a more direct policy-improvement signal but introduces substantial memory and stability costs. Q-score matching [27] instead connects the policy score to the action gradient of the Q-function and differentiates through the denoising model rather than the full sampling chain. Thus, avoiding BPTT is an established design property; DeFlow’s distinction is to apply value gradients to a proposal-conditioned action residual while training the prior only through behavior modeling. Flow matching offers straighter generation paths and has recently been used for policy and value modeling in FQL, floq, SAC Flow, MeanFlow-QL, SORL, and Value Flows [1, 6, 7, 18, 24, 33, 38]. DeFlow is closest to this flow-based line, but keeps the iterative flow model as a behavior prior and applies critic gradients only to a small residual extractor.

Shortcut extraction and SORL. FQL and related shortcut methods reduce the optimization burden by distilling or parameterizing a one-step policy, making direct Q optimization cheaper than differentiating through a full solver [8, 24]. SORL extends this efficient extraction direction with step-size-conditioned shortcut models, shortcut-consistency training, and inference-time evaluation over shortcut step counts, with optional critic best-of- N selection; its strong results on several OGBench navigation and scene tasks make it an important comparison point [7]. DeFlow explores a different trade-off: retain a multi-step

learned prior and regularize a local residual instead of using a shortcut policy. The empirical pattern in Table 1 is consistent with different extraction mechanisms being useful in different regimes; the unmatched inference budgets do not isolate the cause of this pattern.

Residual policy learning and refinement. Residual learning adds learned corrections on top of existing models, from residual networks [12] to residual policy learning [29]. Recent work trains small refinement modules on larger policies, often with online interaction [2, 5, 35]. DeFlow uses the same broad decomposition in offline policy extraction, with an action-conditioned residual so different samples from the same state can receive different local corrections.

6 Discussions and Conclusion

We presented DeFlow, a BPTT-free framework for extracting policies from multi-step flow priors through action-conditioned residual refinement. The empirical picture is deliberately scoped: DeFlow is competitive with recent flow-based offline RL methods, improves selected multimodal manipulation and O2O tasks, and underperforms strong shortcut or inference-scaling methods in some navigation regimes. This suggests that offline generative policy extraction should be evaluated along several axes at once: behavior-prior fidelity, local value improvement, inference-time compute, and robustness to critic error.

Limitations. DeFlow depends on coverage of useful proposals: a local residual has no reliable mechanism for recovering distant behaviors absent from the data or missed by the prior. Proximity to an imperfect prior does not certify true data support. Both the budget δ and residual capacity affect performance, and adaptive α does not remove these choices or prevent critic exploitation. Our attribution tests do not systematically vary prior quality, and the critic-optimism explanation remains a hypothesis. Multi-step sampling retains a cost proportional to the solver budget, and the SORL comparison is not compute-matched.

Acknowledgments and Disclosures

Zhancun Mu thanks his former advisor, Yitao Liang, for his support. This work was supported by the Shanghai Artificial Intelligence Innovation Center under Grant No. 2025SHZDZX026D22. The authors declare no competing interests relevant to this work.

References

- [1] Bhavya Agrawalla, Michal Nauman, Khush Agrawal, and Aviral Kumar. floq: Training critics via flow-matching for scaling compute in value-based rl. *arXiv preprint arXiv:2509.06863*, 2025.
- [2] Lars Ankile, Anthony Simeonov, Idan Shenfeld, Marcel Torne, and Pulkit Agrawal. From imitation to refinement – residual rl for precise assembly. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 01–08. IEEE, 2025.
- [3] Philip J Ball, Laura Smith, Ilya Kostrikov, and Sergey Levine. Efficient online reinforcement learning with offline data. In *International Conference on Machine Learning (ICML)*, 2023.
- [4] Zihan Ding and Chi Jin. Consistency models as a rich and efficient policy class for reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2024.
- [5] Perry Dong, Qiyang Li, Dorsa Sadigh, and Chelsea Finn. Expo: Stable reinforcement learning with expressive policies, 2025. URL <https://arxiv.org/abs/2507.07986>.
- [6] Perry Dong, Chongyi Zheng, Chelsea Finn, Dorsa Sadigh, and Benjamin Eysenbach. Value flows, 2025. URL <https://arxiv.org/abs/2510.07650>.
- [7] Nicolas Espinosa-Dice, Yiyi Zhang, Yiding Chen, Bradley Guo, Owen Oertell, Gokul Swamy, Kianté Brantley, and Wen Sun. Scaling offline rl via efficient and expressive shortcut models. *arXiv preprint arXiv:2505.22866*, 2025.
- [8] Kevin Frans, Danijar Hafner, Sergey Levine, and Pieter Abbeel. One step diffusion via shortcut models. In *International Conference on Learning Representations (ICLR)*, 2025.
- [9] Justin Fu, Aviral Kumar, Ofir Nachum, G. Tucker, and Sergey Levine. D4rl: Datasets for deep data-driven reinforcement learning. *ArXiv*, abs/2004.07219, 2020.

- [10] Scott Fujimoto and Shixiang Shane Gu. A minimalist approach to offline reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 2021.
- [11] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *arXiv*, 2018. URL <https://arxiv.org/pdf/1801.01290.pdf>.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [13] Bingyi Kang, Xiao Ma, Chao Du, Tianyu Pang, and Shuicheng Yan. Efficient diffusion policies for offline reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 2023.
- [14] Ilya Kostrikov, Ashvin Nair, and Sergey Levine. Offline reinforcement learning with implicit q-learning. In *International Conference on Learning Representations (ICLR)*, 2022.
- [15] Aviral Kumar, Aurick Zhou, G. Tucker, and Sergey Levine. Conservative q-learning for offline reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 2020.
- [16] Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint*, 2020.
- [17] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Manfred Otto Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2016.
- [18] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. In *International Conference on Learning Representations (ICLR)*, 2023.
- [19] Yaron Lipman, Marton Havasi, Peter Holderrieth, Neta Shaul, Matt Le, Brian Karrer, Ricky T. Q. Chen, David Lopez-Paz, Heli Ben-Hamu, and Itai Gat. Flow matching guide and code. *ArXiv*, abs/2412.06264, 2024.
- [20] Cheng Lu, Huayu Chen, Jianfei Chen, Hang Su, Chongxuan Li, and Jun Zhu. Contrastive energy prediction for exact energy-guided diffusion sampling in offline reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2023.
- [21] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin A. Riedmiller. Playing atari with deep reinforcement learning. *ArXiv*, abs/1312.5602, 2013.
- [22] Mitsuhiro Nakamoto, Yuexiang Zhai, Anika Singh, Max Sobol Mark, Yi Ma, Chelsea Finn, Aviral Kumar, and Sergey Levine. Cal-ql: Calibrated offline rl pre-training for efficient online fine-tuning. In *Neural Information Processing Systems (NeurIPS)*, 2023.
- [23] Seohong Park, Kevin Frans, Benjamin Eysenbach, and Sergey Levine. Ogbench: Benchmarking offline goal-conditioned rl. In *International Conference on Learning Representations (ICLR)*, 2025.
- [24] Seohong Park, Qiyang Li, and Sergey Levine. Flow q-learning. *ArXiv*, abs/2502.02538, 2025.
- [25] Xue Bin Peng, Aviral Kumar, Grace Zhang, and Sergey Levine. Advantage-weighted regression: Simple and scalable off-policy reinforcement learning. *ArXiv*, abs/1910.00177, 2019.
- [26] Gabriel Peyré and Marco Cuturi. Computational optimal transport: With applications to data science. *Foundations and Trends in Machine Learning*, 11(5–6):355–607, 2019.
- [27] Michael Psenka, Alejandro Escontrela, Pieter Abbeel, and Yi Ma. Learning a diffusion model policy from rewards via q-score matching. In *Proceedings of the 41st International Conference on Machine Learning*, 2024.
- [28] Martin L. Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- [29] Tom Silver, Kelsey Allen, Josh Tenenbaum, and Leslie Kaelbling. Residual policy learning, 2019. URL <https://arxiv.org/abs/1812.06298>.
- [30] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning (ICML)*, 2015.
- [31] Denis Tarasov, Vladislav Kurenkov, Alexander Nikulin, and Sergey Kolesnikov. Revisiting the minimalist approach to offline reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 2023.
- [32] Cédric Villani et al. *Optimal transport: old and new*. Springer, 2009.
- [33] Zeyuan Wang, Da Li, Yulin Chen, Ye Shi, Liang Bai, Tianyuan Yu, and Yanwei Fu. One-step generative policies with q-learning: A reformulation of meanflow. *arXiv preprint arXiv:2511.13035*, 2025.
- [34] Zhendong Wang, Jonathan J Hunt, and Mingyuan Zhou. Diffusion policies as an expressive policy class for offline reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2023.
- [35] Wenli Xiao, Haotian Lin, Andy Peng, Haoru Xue, Tairan He, Yuqi Xie, Fengyuan Hu, Jimmy Wu, Zhengyi Luo, Linxi "Jim" Fan, Guanya Shi, and Yuke Zhu. Self-improving vision-language-action models with data generation via residual rl, 2025. URL <https://arxiv.org/abs/2511.00091>.
- [36] Haoran Xu, Li Jiang, Jianxiong Li, Zhuoran Yang, Zhaoran Wang, Victor Chan, and Xianyuan Zhan. Offline rl with no ood actions: In-sample learning via implicit value regularization. In *International Conference on Learning*

Representations (ICLR), 2023.

- [37] Shiyuan Zhang, Weitong Zhang, and Quanquan Gu. Energy-weighted flow matching for offline reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2025.
- [38] Yixian Zhang, Shu'ang Yu, Tonghe Zhang, Mo Guang, Haojia Hui, Kaiwen Long, Yu Wang, Chao Yu, and Wenbo Ding. Sac flow: Sample-efficient reinforcement learning of flow-based policies via velocity-reparameterized sequential modeling, 2025. URL <https://arxiv.org/abs/2509.25756>.

A Residual Trust Region as an OT Surrogate

We use the standard Wasserstein formulation of optimal transport, where a Monge map deterministically pushes one distribution to another and W_2^2 is the minimum squared Euclidean cost over all couplings [26, 32].

Proposition A.1. *Fix a state s and let $\nu_s = \mu_\psi(\cdot|s)$ be the learned flow prior. Assume that ν_s and the residual have finite second moments. Define the refined policy by the deterministic map $T_s(a_b) = a_b + f_\phi(s, a_b)$ with $a_b \sim \nu_s$, i.e., $\pi_\phi(\cdot|s) = (T_s)_\# \nu_s$. Then*

$$W_2^2(\pi_\phi(\cdot|s), \nu_s) \leq \mathbb{E}_{a_b \sim \nu_s} [\|f_\phi(s, a_b)\|_2^2]. \quad (13)$$

Proof. The joint distribution induced by $(a_b, T_s(a_b))$ is a valid coupling between ν_s and $\pi_\phi(\cdot|s)$. The squared transport cost under this coupling is $\mathbb{E}_{a_b \sim \nu_s} \|a_b - T_s(a_b)\|_2^2 = \mathbb{E}_{a_b \sim \nu_s} \|f_\phi(s, a_b)\|_2^2$. Since W_2^2 is the minimum over all couplings, the inequality follows. $\square$

This proposition only relates the refined policy to the learned flow prior. It does not imply exact preservation of the true data support, which depends on the quality and coverage of μ_ψ . With the implementation’s coordinate-averaged penalty, the upper bound is $d \mathbb{E} \|f_\phi(s, a_b)\|_2^2 / d$. Averaging over replay states gives $\mathbb{E}_s W_2^2 \leq d\delta$ if the population constraint is satisfied. This is an average-state statement, not a per-state guarantee; a minibatch estimate alone is not a certified population bound. For bounded actions, the base proposal is clipped before refinement. Projection of the refined action onto the same box cannot increase its distance to that base proposal, so the same bound holds after action clipping.

B Implementation Details

We implement DeFlow based on the FQL [24] and floq [1] codebases. The flow model architecture follows the same design as FQL, using a 4-layer MLP with 512 hidden units. The refinement module uses the same architecture for clear comparison. We use Q-Normalization as in FQL to stabilize Q-value scales across different tasks. For other hyperparameters except the target divergence δ , we follow the FQL/floq’s default settings. The experiments use public OGBench and D4RL benchmarks and the FQL/floq implementation bases under their respective terms and licenses.

B.1 Hyperparameter Settings

For the target divergence δ , we set it based on the estimated Intrinsic Action Variance (IAV) of the dataset. Specifically, we compute the average variance of actions among k -nearest neighbors ($k = 5$) for each state in the following tasks and normalize it by the action dimension:

	cube-double	antmaze-large	puzzle-3x3	antsoccer
IAV	0.013	0.018	0.007	0.026

Leveraging prior knowledge of task complexity [23], we use the Intrinsic Action Variance (IAV) as a rough guide for choosing the constraint δ :

- **Fine Manipulation** (e.g., cube-double-play): High precision requires a tighter constraint, $\delta \approx 0.1 \times \text{IAV}$.
- **Navigation/Locomotion** (e.g., antsoccer): Higher tolerance allows a looser constraint, $\delta \approx 1.0 \times \text{IAV}$.
- **D4RL Benchmarks**: A generic setting suffices based on data quality (e.g., $\delta = 10^{-3}$).

IAV is a rough scale heuristic, not a validated predictor of an optimal budget. The actual settings include exceptions: humanoidmaze uses 0.001, puzzles use 0.0005, D4RL AntMaze uses 0.015, and pen uses 0.01. The complete task table, rather than the generic heuristic above, specifies the reported configuration. We maintain the same δ for online fine-tuning without further adjustment. This strict protocol may limit

asymptotic performance on specific outlier tasks, but it tests whether the residual constraint can transfer without task-specific retuning. Detailed hyperparameter settings for all experiments are provided in Table 14.

Selection protocol and reporting scope. Table 5 separates inherited settings, task-dependent choices, and post-submission diagnostics. The main benchmark used a four-layer residual throughout; additional depth comparisons do not replace those scores. We retain the same task budget during O2O adaptation. The original records do not provide a complete count of candidate trials or environment evaluations used during hyperparameter development. Thus, the reported configuration and diagnostic budgets are reproducible, but they should not be interpreted as a complete audit of historical selection cost.

Table 5. Hyperparameter choices and diagnostic scope. The sensitivity grids are diagnostic settings, not a claimed search history for selecting the main results.

Quantity	Main benchmark protocol	Additional analysis / budget
Residual budget δ	Task-dependent IAV-guided choice; retained for O2O	$\{0.0005, 0.001, 0.005, 0.01, 0.02\}$ diagnostic in Table 6
Penalty α	Adapted in log space; not a fixed-coefficient grid	Original nominal fixed settings $\{0.03, 0.1, 0.3, 1, 3\}$ in Table 9
Residual architecture	Four hidden layers, width 512; no per-task depth tuning	One/two layers, width 512; four seeds per setting; Section C.4
Prior and critic	FQL-based configuration; prior learns by flow matching	Held at the same architecture within each capacity comparison
Inference budget	Multi-step proposal followed by one residual	Ten Euler steps in capacity runs; no candidate ranking
Evaluation selection	Fixed-step main evaluation	Final scheduled checkpoint for capacity; 20 or 50 episodes per seed

Implementation conventions. The residual penalty is the minibatch mean of squared coordinates, $\hat{g} = (Bd)^{-1} \sum_{i=1}^B \|\Delta a_i\|_2^2$. The actor treats α as constant, and the dual update treats $\hat{g}$ as constant. In the audited adaptive implementation, $\eta = \log \alpha$ starts at 1 and uses the shared Adam learning rate 3×10^{-4} . Q-normalization divides the actor objective by the stopped minibatch mean absolute Q-value. Clipping is applied to the base proposal and the refined action. These conventions matter when translating a reported δ into an action-space displacement.

C Additional Experimental Results

We report the full offline RL results across all 73 tasks in Table 15. The results are averaged over 8 seeds (4 seeds for pixel-based tasks) and we bold scores that are within 95% of the best performance as in FQL [24].

C.1 Additional Diagnostics

This section collects diagnostics that are referenced in the main text but are not used as primary benchmark claims. They are intended to probe failure modes and component attribution.

Table 6. Sensitivity to the residual budget δ on a representative manipulation diagnostic. Larger budgets permit larger residual moves and can quickly degrade performance.

δ	0.0005	0.001	0.005	0.01	0.02
Score	55 ± 8	50 ± 7	27 ± 4	8 ± 5	4 ± 2

Table 7. Diagnostic checks for prior coverage and action-conditioned refinement. The mixed dataset uses 50% human and 50% expert data; the refinement ablation removes action conditioning from $f_\phi(s, a_b)$.

Diagnostic	Setting	Baseline	DeFlow / Variant
Mixed human/expert data	score	FQL: 21 ± 3	DeFlow: 37 ± 5
Refinement module	full action-conditioned $f_\phi(s, a_b)$	-	40 ± 5
Refinement module	remove action conditioning	-	5 ± 3

Table 8. Success rates on manipulation tasks when combining DeFlow with floq’s value function.

Task	FQL	floq	DeFlow	DeFlow+floq
cube-double-play	25 ± 6	47 ± 15	40 ± 5	58 ± 5
puzzle-3x3-play	29 ± 5	37 ± 7	43 ± 4	50 ± 3

C.2 Ablations on α Adaptation

We retain the original comparison of adaptive penalties with fixed-coefficient settings on cube-double-play and puzzle-3x3-play, with the implementation qualification below.

Table 9. Original fixed-penalty sweep versus adaptation, separated by training phase. The fixed columns retain the nominal sweep setting c ; see the implementation qualification below.

Task	Phase	$c = 0.03$	0.1	0.3	1	3	Adaptive
cube-double	Offline	0 ± 0	0 ± 0	0 ± 0	44 ± 2	7 ± 2	48 ± 13
	Online	0 ± 0	0 ± 0	0 ± 0	10 ± 3	89 ± 5	93 ± 4
puzzle-3x3	Offline	0 ± 0	0 ± 0	7 ± 7	0 ± 0	35 ± 8	23 ± 8
	Online	100 ± 1	100 ± 0	100 ± 0	0 ± 0	0 ± 0	100 ± 0

The original sweep illustrates task- and phase-dependent sensitivity: small nominal settings perform poorly offline on cube-double, whereas a setting that works offline on puzzle-3x3 can fail online. Adaptation is competitive in these examples but does not win every offline comparison. These returns do not by themselves establish that failures are caused by OOD actions or that the adaptive constraint is always satisfied.

Implementation qualification. The current fixed-penalty code path multiplies the configured coefficient twice, giving an effective coefficient c^2 . The historical code version for this original sweep has not been recovered, so we retain its nominal labels rather than assert a calibrated effective- α axis. This table is descriptive evidence about the reported settings, not a matched-effective-penalty comparison. The new capacity diagnostics all use the adaptive branch and are unaffected by that fixed-branch issue.

C.3 Prior-Only and Residual Attribution

We evaluate the same four trained checkpoints per task with and without the residual at inference, using 50 episodes per checkpoint and mode. No component is retrained. Scores in Table 10 are success percentages, reported as mean and standard deviation across checkpoints. This tests whether the residual changes performance for the particular learned priors; it does not isolate the effects of prior architecture, prior training quality, or dataset coverage.

Table 10. Inference-time residual attribution at the same trained checkpoints. These are separate evaluations, not replacements for the main benchmark scores.

Task	Base flow only	Base flow + residual
cube-double task1	10 ± 2	62 ± 12
puzzle-3x3 task2	0.5 ± 0.9	35 ± 11

The improvement shows that these scores do not arise from the flow prior alone. It is complementary to the action-input ablation in Table 7: disabling the residual tests its inference-time contribution, whereas removing action conditioning changes the correction function being learned. Neither establishes a guarantee of mode preservation. The residual is conditioned on the proposal’s position in action space, without explicit mode labels.

C.4 Residual-Capacity Sensitivity

We vary only the residual hidden-layer configuration within each new comparison, holding the flow prior and critic at four hidden layers of width 512. Each configuration uses four seeds, a fixed task budget, ten flow steps, batch size 256, and learning rate 3×10^{-4} . We report the last scheduled evaluation: one million offline updates for OGBench and 500,000 for pen-cloned. Cube-double uses seeds 2026–2029 and 20 evaluation episodes per seed; the other capacity runs use seeds 1996–1999 and 50 episodes. Means and population standard deviations are computed across the four final seed scores and rounded to whole points.

Table 11. Residual-depth diagnostics. All hidden layers have width 512. Cube’s four-layer entry is a matched four-seed rerun (+); puzzle and scene four-layer entries are the original eight-seed benchmark references (*), not matched reruns.

Task	One layer	Two layers	Four layers
cube-double task1	81 ± 9	74 ± 23	$69 \pm 10 +$
puzzle-3x3 task2	6 ± 3	55 ± 7	$38 \pm 11 *$
scene task2	52 ± 12	70 ± 13	$59 \pm 17 *$

The original cube-double task1 main-table result is 65 ± 6 across eight seeds; it is unchanged. The four-seed cube rerun above is distinct. Across these tasks, the results do not support a monotonic benefit from increasing residual depth. Comparisons to the original puzzle and scene references also differ in seeds and evaluation provenance, so they should not be treated as paired estimates of a depth effect.

Table 12. Broader one- versus two-layer sensitivity diagnostic. OGBench entries are success percentages; pen-cloned uses D4RL normalized return. No best-depth selection is applied to the main benchmark.

Task	One layer	Two layers
antmaze-large task2	48 ± 14	73 ± 6
humanoidmaze-medium task2	43 ± 18	73 ± 10
antsoccer-arena task2	28 ± 12	76 ± 3
cube-single task2	98 ± 1	97 ± 3
puzzle-4x4 task1	22 ± 2	24 ± 7
pen-cloned	63 ± 20	64 ± 8

Interpretation and limits. Two layers improve over one on several navigation tasks and puzzle-3x3, whereas cube-single is nearly saturated and the pen-cloned and puzzle-4x4 differences are small relative to their variability. A plausible explanation is a trade-off between representational underfitting and optimization against an imperfect offline critic [16]. Insufficient capacity may miss useful proposal-dependent corrections. Greater capacity may fit sharper directions of optimistic Q-error under a DDPG-style actor objective; a smaller network could sometimes provide useful implicit regularization. These experiments do not measure Q-calibration, residual smoothness, or optimization conditioning directly, so they do not establish that explanation or show that smaller networks are intrinsically smoother. The same δ sets the same target penalty budget, but actual residual norms can differ because the constraint is only encouraged during training.

Table 13. Task-specific settings for the new capacity diagnostics. Target aggregation refers to the Bellman target, not candidate ranking.

Task	δ	Discount	Target aggregation
cube-double task1	0.001	0.99	mean
puzzle-3x3 task2	0.0005	0.99	mean
scene task2	0.001	0.99	mean
antmaze-large task2	0.01	0.99	min
humanoidmaze-medium task2	0.001	0.995	mean
antsoccer-arena task2	0.01	0.995	mean
cube-single task2	0.001	0.99	mean
puzzle-4x4 task1	0.0005	0.99	mean
pen-cloned	0.01	0.99	min

D Runtime and Compute Accounting

Table 4 reports the original NVIDIA RTX 3090 per-step timings. FQL and DeFlow take 2.2 ± 1.0 and 2.4 ± 1.0 ms per offline training step, respectively, and 4.2 ± 0.7 and 4.5 ± 0.5 ms per inference step. DeFlow’s online training time is 5.3 ± 0.4 ms, versus 5.2 ± 0.5 ms when freezing the prior. These measurements characterize this implementation, not a controlled decomposition of residual-head overhead or an equal-compute comparison with SORL.

The additional capacity runs use ten Euler flow steps, batch size 256, and one million training updates for each OGBench run (500,000 for pen-cloned). Evaluation counts, seeds, and task-specific settings are listed in Section C.4 and Table 13. These update counts exclude evaluation and compilation time. The original timing record does not specify enough detail about warm-up, synchronization, timing repeats, or evaluation overhead to infer total experimental GPU-hours from per-step timings. We therefore do not report an extrapolated total-compute estimate.

The architecture avoids an actor gradient through the flow solver but still pays for its forward evaluations. A larger pretrained action prior could make this separation attractive, including VLA-style priors, but neither VLA adaptation nor large-model runtime scaling is tested here.

Table 14. Task-specific hyperparameters for offline RL. We take hyperparameters from FQL [24] for comparison.

Task	IQL α	ReBRAC (α_1, α_2)	IDQL N	SRPO β	CAC η	FAWAC α	FBRAC α	IFQL N	FQL α	DeFlow δ
antmaze-large-navigate-singletask-task1-v0	10	(0.003, 0.01)	32	0.3	1	3	3	32	10	0.01
antmaze-large-navigate-singletask-task2-v0	10	(0.003, 0.01)	32	0.3	1	3	3	32	10	0.01
antmaze-large-navigate-singletask-task3-v0	10	(0.003, 0.01)	32	0.3	1	3	3	32	10	0.01
antmaze-large-navigate-singletask-task4-v0	10	(0.003, 0.01)	32	0.3	1	3	3	32	10	0.01
antmaze-large-navigate-singletask-task5-v0	10	(0.003, 0.01)	32	0.3	1	3	3	32	10	0.01
antmaze-giant-navigate-singletask-task1-v0	10	(0.003, 0.01)	32	0.3	1	3	10	32	10	0.01
antmaze-giant-navigate-singletask-task2-v0	10	(0.003, 0.01)	32	0.3	1	3	10	32	10	0.01
antmaze-giant-navigate-singletask-task3-v0	10	(0.003, 0.01)	32	0.3	1	3	10	32	10	0.01
antmaze-giant-navigate-singletask-task4-v0	10	(0.003, 0.01)	32	0.3	1	3	10	32	10	0.01
antmaze-giant-navigate-singletask-task5-v0	10	(0.003, 0.01)	32	0.3	1	3	10	32	10	0.01
humanoidmaze-medium-navigate-singletask-task1-v0	10	(0.01, 0.01)	32	0.3	0.03	3	30	32	30	0.001
humanoidmaze-medium-navigate-singletask-task2-v0	10	(0.01, 0.01)	32	0.3	0.03	3	30	32	30	0.001
humanoidmaze-medium-navigate-singletask-task3-v0	10	(0.01, 0.01)	32	0.3	0.03	3	30	32	30	0.001
humanoidmaze-medium-navigate-singletask-task4-v0	10	(0.01, 0.01)	32	0.3	0.03	3	30	32	30	0.001
humanoidmaze-medium-navigate-singletask-task5-v0	10	(0.01, 0.01)	32	0.3	0.03	3	30	32	30	0.001
humanoidmaze-large-navigate-singletask-task1-v0	10	(0.01, 0.01)	32	0.3	1	3	30	32	30	0.001
humanoidmaze-large-navigate-singletask-task2-v0	10	(0.01, 0.01)	32	0.3	1	3	30	32	30	0.001
humanoidmaze-large-navigate-singletask-task3-v0	10	(0.01, 0.01)	32	0.3	1	3	30	32	30	0.001
humanoidmaze-large-navigate-singletask-task4-v0	10	(0.01, 0.01)	32	0.3	1	3	30	32	30	0.001
humanoidmaze-large-navigate-singletask-task5-v0	10	(0.01, 0.01)	32	0.3	1	3	30	32	30	0.001
antsoccer-arena-navigate-singletask-task1-v0	1	(0.01, 0.01)	32	0.03	1	10	30	64	10	0.01
antsoccer-arena-navigate-singletask-task2-v0	1	(0.01, 0.01)	32	0.03	1	10	30	64	10	0.01
antsoccer-arena-navigate-singletask-task3-v0	1	(0.01, 0.01)	32	0.03	1	10	30	64	10	0.01
antsoccer-arena-navigate-singletask-task4-v0	1	(0.01, 0.01)	32	0.03	1	10	30	64	10	0.01
antsoccer-arena-navigate-singletask-task5-v0	1	(0.01, 0.01)	32	0.03	1	10	30	64	10	0.01
cube-single-play-singletask-task1-v0	1	(1, 0)	32	0.03	0.003	1	100	32	300	0.001
cube-single-play-singletask-task2-v0	1	(1, 0)	32	0.03	0.003	1	100	32	300	0.001
cube-single-play-singletask-task3-v0	1	(1, 0)	32	0.03	0.003	1	100	32	300	0.001
cube-single-play-singletask-task4-v0	1	(1, 0)	32	0.03	0.003	1	100	32	300	0.001
cube-single-play-singletask-task5-v0	1	(1, 0)	32	0.03	0.003	1	100	32	300	0.001

Table 14. Task-specific hyperparameters for offline RL (continued).

Task	IQL α	ReBRAC (α_1, α_2)	IDQL N	SRPO β	CAC η	FAWAC α	FBRAC α	IFQL N	FQL α	DeFlow δ
cube-double-play-singletask-task1-v0	0.3	(0.1, 0)	32	0.1	0.3	0.3	100	32	300	0.001
cube-double-play-singletask-task2-v0	0.3	(0.1, 0)	32	0.1	0.3	0.3	100	32	300	0.001
cube-double-play-singletask-task3-v0	0.3	(0.1, 0)	32	0.1	0.3	0.3	100	32	300	0.001
cube-double-play-singletask-task4-v0	0.3	(0.1, 0)	32	0.1	0.3	0.3	100	32	300	0.001
cube-double-play-singletask-task5-v0	0.3	(0.1, 0)	32	0.1	0.3	0.3	100	32	300	0.001
scene-play-singletask-task1-v0	10	(0.1, 0.01)	32	0.1	0.3	0.3	100	32	300	0.001
scene-play-singletask-task2-v0	10	(0.1, 0.01)	32	0.1	0.3	0.3	100	32	300	0.001
scene-play-singletask-task3-v0	10	(0.1, 0.01)	32	0.1	0.3	0.3	100	32	300	0.001
scene-play-singletask-task4-v0	10	(0.1, 0.01)	32	0.1	0.3	0.3	100	32	300	0.001
scene-play-singletask-task5-v0	10	(0.1, 0.01)	32	0.1	0.3	0.3	100	32	300	0.001
puzzle-3x3-play-singletask-task1-v0	10	(0.3, 0.01)	32	0.1	0.01	0.3	100	32	1000	0.0005
puzzle-3x3-play-singletask-task2-v0	10	(0.3, 0.01)	32	0.1	0.01	0.3	100	32	1000	0.0005
puzzle-3x3-play-singletask-task3-v0	10	(0.3, 0.01)	32	0.1	0.01	0.3	100	32	1000	0.0005
puzzle-3x3-play-singletask-task4-v0	10	(0.3, 0.01)	32	0.1	0.01	0.3	100	32	1000	0.0005
puzzle-3x3-play-singletask-task5-v0	10	(0.3, 0.01)	32	0.1	0.01	0.3	100	32	1000	0.0005
puzzle-4x4-play-singletask-task1-v0	3	(0.3, 0.01)	32	0.1	0.01	0.3	300	32	1000	0.0005
puzzle-4x4-play-singletask-task2-v0	3	(0.3, 0.01)	32	0.1	0.01	0.3	300	32	1000	0.0005
puzzle-4x4-play-singletask-task3-v0	3	(0.3, 0.01)	32	0.1	0.01	0.3	300	32	1000	0.0005
puzzle-4x4-play-singletask-task4-v0	3	(0.3, 0.01)	32	0.1	0.01	0.3	300	32	1000	0.0005
puzzle-4x4-play-singletask-task5-v0	3	(0.3, 0.01)	32	0.1	0.01	0.3	300	32	1000	0.0005
antmaze-umaze-v2	-	-	-	-	0.01	3	10	32	10	0.015
antmaze-umaze-diverse-v2	-	-	-	-	0.01	3	10	32	10	0.015
antmaze-medium-play-v2	-	-	-	-	0.01	3	10	32	10	0.015
antmaze-medium-diverse-v2	-	-	-	-	0.01	3	10	32	10	0.015
antmaze-large-play-v2	-	-	-	-	4.5	3	1	32	3	0.015
antmaze-large-diverse-v2	-	-	-	-	3.5	3	1	32	3	0.015
pen-human-v1	-	-	32	0.03	0.003	0.03	30000	32	10000	0.01
pen-cloned-v1	-	-	32	0.1	0.003	0.3	10000	32	10000	0.01
pen-expert-v1	-	-	32	0.1	0.03	0.1	30000	32	3000	0.01
door-human-v1	-	-	32	0.01	0.03	1	30000	32	30000	0.001
door-cloned-v1	-	-	32	0.03	0.03	1	10000	128	30000	0.001
door-expert-v1	-	-	32	0.01	0.03	3	30000	32	30000	0.001
hammer-human-v1	-	-	128	0.1	0.03	3	30000	32	30000	0.001
hammer-cloned-v1	-	-	32	0.1	0.003	0.03	10000	32	10000	0.001
hammer-expert-v1	-	-	32	0.03	0.03	3	30000	32	30000	0.001
relocate-human-v1	-	-	32	0.03	0.01	0.3	30000	128	10000	0.001
relocate-cloned-v1	-	-	64	0.03	0.01	0.1	3000	32	30000	0.001
relocate-expert-v1	-	-	32	0.01	0.003	1	30000	32	30000	0.001
visual-cube-single-play-singletask-task1-v0	1	(1, 0)	-	-	-	-	100	32	300	0.001
visual-cube-double-play-singletask-task1-v0	0.3	(0.1, 0)	-	-	-	-	100	32	100	0.001
visual-scene-play-singletask-task1-v0	10	(0.1, 0.01)	-	-	-	-	100	32	100	0.001
visual-puzzle-3x3-play-singletask-task1-v0	10	(0.3, 0.01)	-	-	-	-	100	32	300	0.0005
visual-puzzle-4x4-play-singletask-task1-v0	3	(0.3, 0.01)	-	-	-	-	300	32	300	0.0005

Table 15. Full offline RL results. We present the full results on the 73 OGBench and D4RL tasks. (*) indicates the default task in each environment. The results are averaged over 8 seeds (4 seeds for pixel-based tasks) unless otherwise mentioned.

Task	Gaussian Policies			Diffusion Policies			Flow Policies				
	BC	IQL	ReBRAC	IDQL	SRPO	CAC	FAWAC	FBRAC	IFQL	FQL	DeFlow
antmaze-large-navigate-singletask-task1-v0 (*)	0 $\pm$ 0	48 $\pm$ 9	91 $\pm$ 10	0 $\pm$ 0	0 $\pm$ 0	42 $\pm$ 7	1 $\pm$ 1	70 $\pm$ 20	24 $\pm$ 17	80 $\pm$ 8	88 $\pm$ 6
antmaze-large-navigate-singletask-task2-v0	6 $\pm$ 3	42 $\pm$ 6	88 $\pm$ 4	14 $\pm$ 8	4 $\pm$ 4	1 $\pm$ 1	0 $\pm$ 1	35 $\pm$ 12	8 $\pm$ 3	57 $\pm$ 10	71 $\pm$ 10
antmaze-large-navigate-singletask-task3-v0	29 $\pm$ 5	72 $\pm$ 7	51 $\pm$ 18	26 $\pm$ 8	3 $\pm$ 2	49 $\pm$ 10	12 $\pm$ 4	83 $\pm$ 15	52 $\pm$ 17	93 $\pm$ 3	88 $\pm$ 5
antmaze-large-navigate-singletask-task4-v0	8 $\pm$ 3	51 $\pm$ 9	84 $\pm$ 7	62 $\pm$ 25	45 $\pm$ 19	17 $\pm$ 6	10 $\pm$ 3	37 $\pm$ 18	18 $\pm$ 8	80 $\pm$ 4	77 $\pm$ 5
antmaze-large-navigate-singletask-task5-v0	10 $\pm$ 3	54 $\pm$ 22	90 $\pm$ 2	2 $\pm$ 2	1 $\pm$ 1	55 $\pm$ 6	9 $\pm$ 5	76 $\pm$ 8	38 $\pm$ 18	83 $\pm$ 4	84 $\pm$ 3
antmaze-giant-navigate-singletask-task1-v0 (*)	0 $\pm$ 0	0 $\pm$ 0	27 $\pm$ 22	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 1	0 $\pm$ 0	4 $\pm$ 5	4 $\pm$ 4
antmaze-giant-navigate-singletask-task2-v0	0 $\pm$ 0	1 $\pm$ 1	16 $\pm$ 17	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	4 $\pm$ 7	0 $\pm$ 0	9 $\pm$ 7	36 $\pm$ 13
antmaze-giant-navigate-singletask-task3-v0	0 $\pm$ 0	0 $\pm$ 0	34 $\pm$ 22	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 1	2 $\pm$ 3
antmaze-giant-navigate-singletask-task4-v0	0 $\pm$ 0	0 $\pm$ 0	5 $\pm$ 12	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	9 $\pm$ 4	0 $\pm$ 0	14 $\pm$ 23	5 $\pm$ 8
antmaze-giant-navigate-singletask-task5-v0	1 $\pm$ 1	19 $\pm$ 7	49 $\pm$ 22	0 $\pm$ 1	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	6 $\pm$ 10	13 $\pm$ 9	16 $\pm$ 28	16 $\pm$ 19
humanoidmaze-medium-navigate-singletask-task1-v0 (*)	1 $\pm$ 0	32 $\pm$ 7	16 $\pm$ 9	1 $\pm$ 1	0 $\pm$ 0	38 $\pm$ 19	6 $\pm$ 2	25 $\pm$ 8	69 $\pm$ 19	19 $\pm$ 12	21 $\pm$ 11
humanoidmaze-medium-navigate-singletask-task2-v0	1 $\pm$ 0	41 $\pm$ 9	18 $\pm$ 16	1 $\pm$ 1	1 $\pm$ 1	47 $\pm$ 35	40 $\pm$ 2	76 $\pm$ 10	85 $\pm$ 11	94 $\pm$ 3	69 $\pm$ 5
humanoidmaze-medium-navigate-singletask-task3-v0	6 $\pm$ 2	25 $\pm$ 5	36 $\pm$ 13	0 $\pm$ 1	2 $\pm$ 1	83 $\pm$ 18	19 $\pm$ 2	27 $\pm$ 11	49 $\pm$ 49	74 $\pm$ 18	63 $\pm$ 14
humanoidmaze-medium-navigate-singletask-task4-v0	0 $\pm$ 0	0 $\pm$ 1	15 $\pm$ 16	1 $\pm$ 1	1 $\pm$ 1	5 $\pm$ 4	1 $\pm$ 1	1 $\pm$ 2	1 $\pm$ 1	3 $\pm$ 4	6 $\pm$ 7
humanoidmaze-medium-navigate-singletask-task5-v0	2 $\pm$ 1	66 $\pm$ 4	24 $\pm$ 20	1 $\pm$ 1	3 $\pm$ 3	91 $\pm$ 5	31 $\pm$ 7	63 $\pm$ 9	98 $\pm$ 2	97 $\pm$ 2	81 $\pm$ 7
humanoidmaze-large-navigate-singletask-task1-v0 (*)	0 $\pm$ 0	3 $\pm$ 1	2 $\pm$ 1	0 $\pm$ 0	0 $\pm$ 0	1 $\pm$ 1	0 $\pm$ 0	0 $\pm$ 1	6 $\pm$ 2	7 $\pm$ 6	3 $\pm$ 4
humanoidmaze-large-navigate-singletask-task2-v0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 4
humanoidmaze-large-navigate-singletask-task3-v0	1 $\pm$ 1	7 $\pm$ 3	8 $\pm$ 4	3 $\pm$ 1	1 $\pm$ 1	2 $\pm$ 3	1 $\pm$ 1	10 $\pm$ 2	48 $\pm$ 10	11 $\pm$ 7	10 $\pm$ 6
humanoidmaze-large-navigate-singletask-task4-v0	1 $\pm$ 0	1 $\pm$ 0	1 $\pm$ 1	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 1	0 $\pm$ 0	0 $\pm$ 0	1 $\pm$ 1	2 $\pm$ 3	8 $\pm$ 3
humanoidmaze-large-navigate-singletask-task5-v0	0 $\pm$ 1	1 $\pm$ 1	2 $\pm$ 2	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	1 $\pm$ 1	0 $\pm$ 0	1 $\pm$ 3	1 $\pm$ 1
antsoccer-arena-navigate-singletask-task1-v0	2 $\pm$ 1	14 $\pm$ 5	0 $\pm$ 0	44 $\pm$ 12	2 $\pm$ 1	1 $\pm$ 3	22 $\pm$ 2	17 $\pm$ 3	61 $\pm$ 25	77 $\pm$ 4	84 $\pm$ 7
antsoccer-arena-navigate-singletask-task2-v0	2 $\pm$ 2	17 $\pm$ 7	0 $\pm$ 1	15 $\pm$ 12	3 $\pm$ 1	0 $\pm$ 0	8 $\pm$ 1	8 $\pm$ 2	75 $\pm$ 3	88 $\pm$ 3	89 $\pm$ 4
antsoccer-arena-navigate-singletask-task3-v0	0 $\pm$ 0	6 $\pm$ 4	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	8 $\pm$ 19	11 $\pm$ 5	16 $\pm$ 3	14 $\pm$ 22	61 $\pm$ 6	56 $\pm$ 7
antsoccer-arena-navigate-singletask-task4-v0 (*)	1 $\pm$ 0	3 $\pm$ 2	0 $\pm$ 0	0 $\pm$ 1	0 $\pm$ 0	0 $\pm$ 0	12 $\pm$ 3	24 $\pm$ 4	16 $\pm$ 9	39 $\pm$ 6	43 $\pm$ 5
antsoccer-arena-navigate-singletask-task5-v0	0 $\pm$ 0	2 $\pm$ 2	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	9 $\pm$ 2	15 $\pm$ 4	0 $\pm$ 1	36 $\pm$ 9	39 $\pm$ 9
cube-single-play-singletask-task1-v0	10 $\pm$ 5	88 $\pm$ 3	89 $\pm$ 5	95 $\pm$ 2	89 $\pm$ 7	77 $\pm$ 28	81 $\pm$ 9	73 $\pm$ 33	79 $\pm$ 4	97 $\pm$ 2	96 $\pm$ 3
cube-single-play-singletask-task2-v0 (*)	3 $\pm$ 1	85 $\pm$ 8	92 $\pm$ 4	96 $\pm$ 2	82 $\pm$ 16	80 $\pm$ 30	81 $\pm$ 9	83 $\pm$ 13	73 $\pm$ 3	97 $\pm$ 2	99 $\pm$ 1
cube-single-play-singletask-task3-v0	9 $\pm$ 3	91 $\pm$ 5	93 $\pm$ 3	99 $\pm$ 1	96 $\pm$ 2	98 $\pm$ 1	87 $\pm$ 4	82 $\pm$ 12	88 $\pm$ 4	98 $\pm$ 2	100 $\pm$ 1
cube-single-play-singletask-task4-v0	2 $\pm$ 1	73 $\pm$ 6	92 $\pm$ 3	93 $\pm$ 4	70 $\pm$ 18	91 $\pm$ 2	79 $\pm$ 6	79 $\pm$ 20	79 $\pm$ 6	94 $\pm$ 3	95 $\pm$ 4
cube-single-play-singletask-task5-v0	3 $\pm$ 3	78 $\pm$ 9	87 $\pm$ 8	90 $\pm$ 6	61 $\pm$ 12	80 $\pm$ 20	78 $\pm$ 10	76 $\pm$ 33	77 $\pm$ 7	93 $\pm$ 3	92 $\pm$ 5

Table 15. Full offline RL results (continued).

Task	Gaussian Policies			Diffusion Policies			Flow Policies				
	BC	IQL	ReBRAC	IDQL	SRPO	CAC	FAWAC	FBRAC	IFQL	FQL	Deflow
cube-double-play-singletask-task1-v0	8 $\pm$ 3	27 $\pm$ 5	45 $\pm$ 6	39 $\pm$ 19	7 $\pm$ 6	21 $\pm$ 8	21 $\pm$ 7	47 $\pm$ 11	35 $\pm$ 9	61 $\pm$ 9	65 $\pm$ 6
cube-double-play-singletask-task2-v0 (*)	0 $\pm$ 0	1 $\pm$ 1	7 $\pm$ 3	16 $\pm$ 10	0 $\pm$ 0	2 $\pm$ 2	2 $\pm$ 1	22 $\pm$ 12	9 $\pm$ 5	36 $\pm$ 6	49 $\pm$ 13
cube-double-play-singletask-task3-v0	0 $\pm$ 0	0 $\pm$ 0	4 $\pm$ 1	17 $\pm$ 8	0 $\pm$ 1	3 $\pm$ 1	1 $\pm$ 1	4 $\pm$ 2	8 $\pm$ 5	22 $\pm$ 5	34 $\pm$ 15
cube-double-play-singletask-task4-v0	0 $\pm$ 0	0 $\pm$ 0	1 $\pm$ 1	0 $\pm$ 1	0 $\pm$ 0	0 $\pm$ 1	0 $\pm$ 0	0 $\pm$ 1	1 $\pm$ 1	5 $\pm$ 2	7 $\pm$ 4
cube-double-play-singletask-task5-v0	0 $\pm$ 0	4 $\pm$ 3	4 $\pm$ 2	1 $\pm$ 1	0 $\pm$ 0	3 $\pm$ 2	2 $\pm$ 1	2 $\pm$ 2	17 $\pm$ 6	19 $\pm$ 10	46 $\pm$ 11
scene-play-singletask-task1-v0	19 $\pm$ 6	94 $\pm$ 3	95 $\pm$ 2	100 $\pm$ 0	94 $\pm$ 4	100 $\pm$ 1	87 $\pm$ 8	96 $\pm$ 8	98 $\pm$ 3	100 $\pm$ 0	98 $\pm$ 2
scene-play-singletask-task2-v0 (*)	1 $\pm$ 1	12 $\pm$ 3	50 $\pm$ 13	33 $\pm$ 14	2 $\pm$ 2	50 $\pm$ 40	18 $\pm$ 8	46 $\pm$ 10	0 $\pm$ 0	76 $\pm$ 9	59 $\pm$ 17
scene-play-singletask-task3-v0	1 $\pm$ 1	32 $\pm$ 7	55 $\pm$ 16	94 $\pm$ 4	4 $\pm$ 4	49 $\pm$ 16	38 $\pm$ 9	78 $\pm$ 14	54 $\pm$ 19	98 $\pm$ 1	85 $\pm$ 5
scene-play-singletask-task4-v0	2 $\pm$ 2	0 $\pm$ 1	3 $\pm$ 3	4 $\pm$ 3	0 $\pm$ 0	0 $\pm$ 0	6 $\pm$ 1	4 $\pm$ 4	0 $\pm$ 0	5 $\pm$ 1	13 $\pm$ 8
scene-play-singletask-task5-v0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0
puzzle-3x3-play-singletask-task1-v0	5 $\pm$ 2	33 $\pm$ 6	97 $\pm$ 4	52 $\pm$ 12	89 $\pm$ 5	97 $\pm$ 2	25 $\pm$ 9	63 $\pm$ 19	94 $\pm$ 3	90 $\pm$ 4	89 $\pm$ 7
puzzle-3x3-play-singletask-task2-v0	1 $\pm$ 1	4 $\pm$ 3	1 $\pm$ 1	0 $\pm$ 1	0 $\pm$ 1	0 $\pm$ 0	4 $\pm$ 2	2 $\pm$ 2	1 $\pm$ 2	16 $\pm$ 5	38 $\pm$ 11
puzzle-3x3-play-singletask-task3-v0	1 $\pm$ 1	3 $\pm$ 2	3 $\pm$ 1	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	1 $\pm$ 0	1 $\pm$ 1	0 $\pm$ 0	10 $\pm$ 3	35 $\pm$ 11
puzzle-3x3-play-singletask-task4-v0 (*)	1 $\pm$ 1	2 $\pm$ 1	2 $\pm$ 1	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	1 $\pm$ 1	2 $\pm$ 2	0 $\pm$ 0	16 $\pm$ 5	23 $\pm$ 8
puzzle-3x3-play-singletask-task5-v0	1 $\pm$ 0	3 $\pm$ 2	5 $\pm$ 3	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	1 $\pm$ 1	2 $\pm$ 2	0 $\pm$ 0	16 $\pm$ 3	31 $\pm$ 8
puzzle-4x4-play-singletask-task1-v0	1 $\pm$ 1	12 $\pm$ 2	26 $\pm$ 4	48 $\pm$ 5	24 $\pm$ 9	44 $\pm$ 10	1 $\pm$ 2	32 $\pm$ 9	49 $\pm$ 9	34 $\pm$ 8	23 $\pm$ 3
puzzle-4x4-play-singletask-task2-v0	0 $\pm$ 0	7 $\pm$ 4	12 $\pm$ 4	14 $\pm$ 5	0 $\pm$ 1	0 $\pm$ 0	0 $\pm$ 1	5 $\pm$ 3	4 $\pm$ 4	16 $\pm$ 5	9 $\pm$ 2
puzzle-4x4-play-singletask-task3-v0	0 $\pm$ 0	9 $\pm$ 3	15 $\pm$ 3	34 $\pm$ 5	21 $\pm$ 10	29 $\pm$ 12	1 $\pm$ 1	20 $\pm$ 10	50 $\pm$ 14	18 $\pm$ 5	9 $\pm$ 4
puzzle-4x4-play-singletask-task4-v0 (*)	0 $\pm$ 0	5 $\pm$ 2	10 $\pm$ 3	26 $\pm$ 6	7 $\pm$ 4	1 $\pm$ 1	0 $\pm$ 0	5 $\pm$ 1	21 $\pm$ 11	11 $\pm$ 3	7 $\pm$ 4
puzzle-4x4-play-singletask-task5-v0	0 $\pm$ 0	4 $\pm$ 1	7 $\pm$ 3	24 $\pm$ 11	1 $\pm$ 1	0 $\pm$ 0	0 $\pm$ 1	4 $\pm$ 3	2 $\pm$ 2	7 $\pm$ 3	7 $\pm$ 4
antmaze-umaze-v2	55	77	98	94	97	66 $\pm$ 5	90 $\pm$ 6	94 $\pm$ 3	92 $\pm$ 6	96 $\pm$ 2	96 $\pm$ 3
antmaze-umaze-diverse-v2	47	54	84	80	82	66 $\pm$ 11	55 $\pm$ 7	82 $\pm$ 9	62 $\pm$ 12	89 $\pm$ 5	87 $\pm$ 8
antmaze-medium-play-v2	0	66	90	84	81	49 $\pm$ 24	52 $\pm$ 12	77 $\pm$ 7	56 $\pm$ 15	78 $\pm$ 7	80 $\pm$ 7
antmaze-medium-diverse-v2	1	74	84	85	75	0 $\pm$ 1	44 $\pm$ 15	77 $\pm$ 6	60 $\pm$ 25	71 $\pm$ 13	76 $\pm$ 19
antmaze-large-play-v2	0	42	52	64	54	0 $\pm$ 0	10 $\pm$ 6	32 $\pm$ 21	55 $\pm$ 9	84 $\pm$ 7	78 $\pm$ 6
antmaze-large-diverse-v2	0	30	64	68	54	0 $\pm$ 0	16 $\pm$ 10	20 $\pm$ 17	64 $\pm$ 8	83 $\pm$ 4	78 $\pm$ 6
pen-human-v1	71	78	103	76 $\pm$ 10	69 $\pm$ 7	64 $\pm$ 8	67 $\pm$ 5	77 $\pm$ 7	71 $\pm$ 12	53 $\pm$ 6	61 $\pm$ 9
pen-cloned-v1	52	83	103	64 $\pm$ 7	61 $\pm$ 7	56 $\pm$ 10	62 $\pm$ 10	67 $\pm$ 9	80 $\pm$ 11	74 $\pm$ 11	75 $\pm$ 4
pen-expert-v1	110	128	152	140 $\pm$ 6	134 $\pm$ 4	103 $\pm$ 9	118 $\pm$ 6	119 $\pm$ 7	139 $\pm$ 5	142 $\pm$ 6	138 $\pm$ 5
door-human-v1	2	3	-0	6 $\pm$ 2	3 $\pm$ 3	5 $\pm$ 2	2 $\pm$ 1	4 $\pm$ 2	7 $\pm$ 2	0 $\pm$ 0	1 $\pm$ 0
door-cloned-v1	-0	3	0	0 $\pm$ 0	0 $\pm$ 0	1 $\pm$ 0	0 $\pm$ 1	0 $\pm$ 0	2 $\pm$ 2	2 $\pm$ 1	2 $\pm$ 1
door-expert-v1	105	107	106	105 $\pm$ 1	105 $\pm$ 0	98 $\pm$ 3	103 $\pm$ 1	104 $\pm$ 1	104 $\pm$ 2	104 $\pm$ 1	104 $\pm$ 1
hammer-human-v1	3	2	0	2 $\pm$ 1	1 $\pm$ 1	2 $\pm$ 0	2 $\pm$ 1	2 $\pm$ 1	3 $\pm$ 1	1 $\pm$ 1	1 $\pm$ 0
hammer-cloned-v1	1	2	5	2 $\pm$ 1	2 $\pm$ 1	1 $\pm$ 1	1 $\pm$ 0	2 $\pm$ 1	2 $\pm$ 1	11 $\pm$ 9	7 $\pm$ 4
hammer-expert-v1	127	129	134	125 $\pm$ 4	127 $\pm$ 0	92 $\pm$ 11	118 $\pm$ 3	119 $\pm$ 9	117 $\pm$ 9	125 $\pm$ 3	127 $\pm$ 1
relocate-human-v1	0	0	0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0	0 $\pm$ 0
relocate-cloned-v1	-0	0	2	-0 $\pm$ 0	-0 $\pm$ 0	-0 $\pm$ 0	-0 $\pm$ 0	1 $\pm$ 1	-0 $\pm$ 0	-0 $\pm$ 0	1 $\pm$ 8
relocate-expert-v1	108	106	108	107 $\pm$ 1	106 $\pm$ 2	93 $\pm$ 6	105 $\pm$ 3	105 $\pm$ 2	104 $\pm$ 3	107 $\pm$ 1	109 $\pm$ 2
visual-cube-single-play-singletask-task1-v0 ¹	-	70 $\pm$ 12	83 $\pm$ 6	-	-	-	-	55 $\pm$ 8	49 $\pm$ 7	81 $\pm$ 12	86 $\pm$ 9
visual-cube-double-play-singletask-task1-v0 ¹	-	34 $\pm$ 23	4 $\pm$ 4	-	-	-	-	6 $\pm$ 2	8 $\pm$ 6	21 $\pm$ 11	53 $\pm$ 7
visual-scene-play-singletask-task1-v0 ¹	-	97 $\pm$ 2	98 $\pm$ 4	-	-	-	-	46 $\pm$ 4	86 $\pm$ 10	98 $\pm$ 3	97 $\pm$ 2
visual-puzzle-3x3-play-singletask-task1-v0 ¹	-	7 $\pm$ 15	88 $\pm$ 4	-	-	-	-	7 $\pm$ 2	100 $\pm$ 0	94 $\pm$ 1	97 $\pm$ 2
visual-puzzle-4x4-play-singletask-task1-v0 ¹	-	0 $\pm$ 0	26 $\pm$ 6	-	-	-	-	0 $\pm$ 0	8 $\pm$ 15	33 $\pm$ 6	14 $\pm$ 2

¹ Due to the high computational cost of pixel-based tasks, we selectively benchmark 5 methods that achieve strong performance on state-based OGBench tasks.