

OM²P: Offline Multi-Agent Mean-Flow Policy

AAAI Track

Zhuoran Li

IIIS, Tsinghua University

Beijing, China

lizr20@mails.tsinghua.edu.cn

Xun Wang

IIIS, Tsinghua University

Beijing, China

wang-x24@mails.tsinghua.edu.cn

Hai Zhong

IIIS, Tsinghua University

Beijing, China

zhongh22@mails.tsinghua.edu.cn

Qingxin Xia

ByteDance

Beijing, China

xiaqingxin@bytedance.com

Lihua Zhang

ByteDance

Beijing, China

lizhiyu.0@bytedance.com

Longbo Huang*

longbohuang@tsinghua.edu.cn

IIIS, Tsinghua University

Beijing, China

ABSTRACT

Generative models, especially diffusion and flow-based models, have been promising in offline multi-agent reinforcement learning. However, integrating powerful generative models into this framework poses unique challenges. In particular, diffusion and flow-based policies suffer from low sampling efficiency due to their iterative generation processes, making them impractical in time-sensitive or resource-constrained settings. To tackle these difficulties, we propose Offline Multi-Agent Mean-Flow Policy (OM²P), a novel offline MARL algorithm to achieve efficient *one-step* action generation. To address the misalignment between generative objectives and reward maximization, we introduce a reward-aware optimization scheme that integrates a carefully-designed mean-flow matching loss with Q-function supervision. Additionally, we design a generalized timestep distribution and a derivative-free estimation strategy to reduce memory overhead and improve training stability. Empirical evaluations on Multi-Agent Particle and MuJoCo benchmarks demonstrate that OM²P achieves superior performance, with up to a 3.8× reduction in GPU memory usage and up to a 10.1× speed-up in training time. Our approach represents the first to successfully integrate mean-flow model into offline MARL, paving the way for practical and scalable generative policies in cooperative multi-agent settings.

KEYWORDS

Offline Multi-Agent Reinforcement Learning; Mean-Flow Policy

ACM Reference Format:

Zhuoran Li, Xun Wang, Hai Zhong, Qingxin Xia, Lihua Zhang, and Longbo Huang. 2026. OM²P: Offline Multi-Agent Mean-Flow Policy: AAAI Track. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 10 pages. <https://doi.org/10.65109/670A0504>

*Corresponding author.



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/670A0504>

1 INTRODUCTION

Offline Multi-Agent Reinforcement Learning (Offline MARL) focuses on acquiring coordinated policies from fixed datasets without further environment interaction. This setup is particularly valuable in risk-sensitive or limited-interaction domains such as autonomous driving [51], robotic manipulation [8], and distributed resource allocation [14], where data collection is expensive or unsafe. Recently, generative models—especially diffusion and flow-based models—have shown promise for offline policy learning by modeling expressive, multimodal action distributions. These models, originally developed for image generation [17, 26, 43], have been extended to offline RL [5, 37, 47] and further adapted to multi-agent settings for action generation [25], trajectory planning [54], and data augmentation [11].

However, their reliance on iterative sampling often leads to inefficiencies in training and inference, posing a bottleneck for large-scale or time-sensitive applications. While these generative models offer strong representational capacity, their sampling inefficiency poses a major obstacle in offline MARL. Unlike conventional policy networks, these policies require multiple iterative steps to sample actions, leading to high computational overhead for training and inference. This issue is magnified in multi-agent scenarios due to the need for repeated joint action sampling across agents, which significantly amplifies the computational burden. Improving the sampling efficiency of diffusion or flow policies is therefore critical for their practical deployment in time-sensitive, real-world multi-agent applications [49].

Recent advances in flow-based models, particularly mean-flow model [12], have demonstrated the potential for efficient one-step sampling by replacing instantaneous velocity fields with their averaged counterparts, enabling fast and high-quality generation in image synthesis. Inspired by mean-flow model, we thus explore whether they can serve as policy networks in offline MARL to resolve the efficiency bottleneck by avoiding multi-step inference during training, especially policy distillation [37]. However, directly applying this adaptation poses key challenges: (i) The training objectives of mean-flow model and offline MARL are inherently misaligned: maximizing expected rewards is not equivalent to minimizing negative log-likelihood. (ii) Computing the gradients of the mean-flow velocity, which is required by the partial derivation of the mean-flow objective, introduces significant computational

overhead. (iii) Existing mean-flow training formulations are not tailored for one-step policy learning, especially with uniform timestep sampling, which is critical in RL applications.

Motivated by the limitations of generative models in offline MARL and the burgeoning development of flow-based models, we ask: *Is it possible to design ultra-efficient training and inference flow-based policies in offline MARL?*

To address the efficiency challenges in offline MARL, we introduce *Offline Multi-Agent Mean-Flow Policy* (OM²P), a novel framework that seamlessly incorporates mean-flow model as policy networks for offline MARL. OM²P redefines the training paradigm of mean-flow-based policy networks via generalized timestep sampling and derivative-free target mean-velocity estimation to enable efficient one-step action sampling without relying on distillation and iterative action generation. Our key contributions are summarized as follows:

- We introduce a novel framework of the mean-flow-based model into offline MARL, avoiding policy distillation to achieve one-step action generation. This design significantly improves training and inference efficiency while maintaining strong policy expressiveness.
- We propose a decentralized offline training scheme that leverages a modified mean-flow matching loss in conjunction with Q-function supervision. To further improve efficiency, we adopt a generalized exponential-family distribution for timestep sampling and replace the exact gradient of target mean-velocity with a numerically stable finite-difference approximation, leading to reduced memory overhead and faster optimization.
- We conduct comprehensive experiments on standard offline MARL benchmarks, including Multi-Agent Particle (MPE) and Multi-Agent MuJoCo (MAMuJoCo) environments. OM²P consistently achieves near-optimal performance with up to 3.8× GPU memory reduction and 10.1× speed up in training time, demonstrating both its efficiency and effectiveness.

2 RELATED WORK

2.1 Offline MARL

Offline Multi-Agent Reinforcement Learning (Offline MARL) faces unique challenges such as distributional shift, credit assignment, and multi-agent coordination under fixed datasets [9, 23]. Existing works extend pessimistic value estimation from single-agent RL to the multi-agent setting (e.g., MAICQ [50], MABCQ [19], CFCQL [42], and OMAC [46]), primarily focusing on conservative learning principles. Other methods, for instance, OMAR [36], SIT [44] and InSPO [29] address local optima and data imbalance through actor rectification, reliable credit assignment and sequentially optimizing agent policies in-sample to avoid out-of-distribution joint actions as a potential risk.

Recent methods leverage diffusion models to solve data scarcity and multi-modality in offline MARL, achieving coordinated behavior through trajectory modeling (MADiff [54], DoF [24]), interaction-aware data synthesis (INS [11]), diffusion-based action generation (DOM2 [25]), and decentralized policy learning with score function decomposition (OMSD [40]). Planning-based and world-model approaches (e.g., MOMA-PPO [4], AlberDICE [33]) further improve

coordination under offline constraints, while trajectory generation methods (e.g., MAT [48], MADT [34], and MADTKD [45]) enhance policy learning by modeling the temporal dependencies and structure of agent behaviors.

2.2 Diffusion Models and Flow Matching in Offline RL

Diffusion-based methods have been widely adopted in offline RL for modeling action distributions as policies (Diff-QL [47], SfBC [5], CEP [32], EDP [20], and IDQL [16]), behavior cloning (Diff-BC [38], ARQ [13]), trajectory planning (Diffuser [18], Decision Diffuser [1]), and data augmentation (SynthER [31], GTA [22]). The emergence of diffusion policies has also inspired extensions to robotic learning, e.g., Diffusion Policy [6] and DPPO [41]. In parallel, flow-based approaches, e.g., FQL [37], FlowQ [3], QIPO [52], and SSCP [21] offer efficient alternatives to iterative diffusion by enabling action sampling via flow matching or energy-weighted training.

Different from prior diffusion-based and flow-based approaches that rely on multi-step denoising or distillation—which can be computationally prohibitive in multi-agent settings—OM²P introduces a mean-flow policy that enables efficient one-step action generation. By seamlessly integrating a generalized mean-flow model into offline MARL, our method eliminates the need for multi-step action sampling while naturally aligning policy learning with value estimates. In particular, the mean-flow generator is supervised by the target mean velocity derived from the offline Q-function, providing a stable and principled form of value alignment. This unified design not only maintains high efficiency and low memory overhead, but also proves sufficiently expressive across benchmarks, obviating the need for additional energy-based guidance or more complex reward alignment strategies and offering a scalable and efficient solution for high-dimensional multi-agent coordination.

3 BACKGROUND

3.1 Offline MARL

We consider fully cooperative multi-agent tasks modeled as a decentralized partially observable Markov decision process (Dec-POMDP) [35], defined by a tuple $G = \langle \mathcal{I}, \mathcal{S}, \mathcal{O}, \mathcal{A}, \Pi, \mathcal{P}, \mathcal{R}, n, \gamma \rangle$. Here, $\mathcal{I} = \{1, \dots, n\}$ denotes the set of n agents, \mathcal{S} is the global state space, and $\mathcal{O} = (\mathcal{O}_1, \dots, \mathcal{O}_n)$ represents the local observation spaces for each agent. Each agent i selects actions $a_i \in \mathcal{A}_i$, where $\mathcal{A} = (\mathcal{A}_1, \dots, \mathcal{A}_n)$ denotes the joint action space. The environment evolves according to the dynamics function $P(s'|s, a) : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$, where $P(s'|s, a) \in \mathcal{P}$ denotes the probability of transitioning to state s' given the current state s and action a . Each agent receives an observation $o_i \in \mathcal{O}_i$ and a scalar reward $r_i \in \mathbb{R}$ from a shared reward function $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$. Within the framework of cooperative MARL, the agents' collective endeavor is to devise a set of policies $\pi = \{\pi_1, \dots, \pi_n\} \in \Pi$ to maximize the expected cumulative discounted rewards: $\mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t \sum_{i=1}^n r_i^t \right]$, where $\gamma \in [0, 1)$ represents the discount factor. In the offline MARL paradigm, agents are trained exclusively on a pre-collected dataset \mathcal{D} generated by sampling from a potentially unknown behavior policy π_β , thereby decoupling the training process from direct interaction with the environment.

3.2 Flow Matching

To pave the way for introducing our OM²P algorithm, we provide a brief overview of flow matching and mean-flow model, which serve as foundational tools in our design.

Flow matching [2, 26, 28] is a framework of generative model that achieves efficient and high-quality sample generation by learning flow trajectories from a simple base distribution (e.g., Gaussian) to a complex target distribution [7, 27]. Given data $x \sim p(x)$ over \mathbb{R}^d , it constructs a trajectory $x_t = \psi(t, x)$ that transforms noise samples $\psi(0, x) \sim \mathcal{N}(0, I_d)$ into data samples $\psi(1, x) = x \sim p(x)$, where the velocity is defined as $v(t, x) = \frac{d}{dt}\psi(t, x)$. To learn this transformation, the model optimizes the velocity $v_\theta(t, x)$ using the objective: $\min_\theta \mathbb{E}_{x_0 \sim \mathcal{N}(0, I_d), x_1 \sim p(x), t \sim U([0, 1])} \|v_\theta(t, x_t) - v(t, x_t)\|_2^2$. After training, new samples can be generated by numerically integrating the learned velocity field via the ODE: $x_t = x_r + \int_r^t v_\theta(\tau, x_\tau) d\tau$.

The mean-flow framework [12] simplifies the sampling by replacing instantaneous velocities with a mean velocity: $u(x_r, r, t) = \frac{1}{t-r} \int_r^t v(\tau, x_\tau) d\tau$, yielding the closed-form update $x_t = x_r + u(x_r, r, t)(t-r)$. This formulation avoids expensive numerical integration, enabling efficient single- or multi-step sampling. To model the mean velocity, we adopt a neural network $u_\theta(x_r, r, t)$ trained to satisfy the mean-flow identity (equals to Equation (9) in [12]) via:

$$\mathcal{L}(\theta) = \mathbb{E} \left[\|u_\theta(x_r, r, t) - \text{stopgrad}(u_{\text{target}})\|_2^2 \right], \quad (1)$$

where $x_r = (1-r)x_0 + rx_1$, with $x_0 \sim \mathcal{N}(0, I_d)$, $x_1 \sim p(x)$, and $r, t \sim U([0, 1])$. The target velocity is computed as $u_{\text{target}} = v(r, x_r) - (r-t) \frac{d}{dr}u(x_r, r, t)$, where the temporal derivative is approximated by $\frac{d}{dr}u = v(r, x_r) \partial_{x_r}u + \partial_r u$.

3.3 Bridge to Mean-Flow Policy

Flow matching trains a velocity field along interpolation paths between noisy inputs and dataset actions, enabling sample generation via denoising. While effective for static data modeling, directly applying this paradigm to offline MARL is inadequate due to key mismatches in objective, efficiency, and scalability.

Unlike generative models, offline MARL demands policies that maximize cumulative rewards, not just replicate dataset distributions—rendering vanilla flow objectives misaligned [47]. Moreover, multi-step generation undermines training efficiency, especially in multi-agent systems where joint action sampling is costly [54]. Although mean-flow model supports one-step generation, their objectives are not inherently designed for it, and computing target velocity fields remains expensive and reward-agnostic [15]. These issues are magnified in decentralized, multi-agent scenarios [49].

To address these challenges, we propose OM²P—a novel offline MARL framework that repurposes mean-flow model as an efficient, reward-aligned policy network. Through a carefully designed objective and optimization strategy, OM²P enables fast one-step action generation while preserving policy expressiveness and improving training efficiency.

4 THE OM²P ALGORITHM

We present OM²P, a novel algorithm addressing the efficiency and scalability bottlenecks of generative policies in multi-agent settings. By leveraging a decentralized mean-flow architecture (as illustrated

by the representative agent in Figure 1), OM²P enables fast one-step action generation, avoiding the computationally prohibitive iterative sampling and distillation typical of single-agent adaptations. This framework introduces specific modeling and optimization innovations to meet the distinct demands of offline MARL. In the following sections, we begin by detailing the integration of mean-flow model into offline MARL and highlighting the inefficiencies of naively applying standard mean-flow formulations. We then introduce a novel training objective tailored to efficient reward-aware policy learning. Finally, we provide an overview of the full OM²P algorithm and discuss its advantages.

4.1 Mean-Flow Model in Offline MARL

We begin by describing how to seamlessly integrate mean-flow model into the offline multi-agent reinforcement learning (MARL) setting and highlight the inefficiencies that arise from directly applying standard formulations. Specifically, we instantiate the mean-flow model as the policy network for each agent in a decentralized MARL framework (for clarity, we omit the agent index in subsequent notation). That is, the conditional policy $\pi_\theta(a|o)$, which maps local observations to action distributions, is parameterized via a mean-flow model.

Given the observation-action pairs $(o, a) \sim \mathcal{D}$, which are sampled from a fixed dataset induced by a behavior policy π_β , the mean-flow policy constructs an action trajectory $a_t = \psi(t, a|o)$ using the instantaneous velocity $v(a_r, r|o) = \frac{da_r}{dr}$ and the mean velocity $u(a_r, r, t|o) = \frac{1}{t-r} \int_r^t v(a_\tau, \tau|o) d\tau$. This defines a continuous path from the noise $a_0 = \psi(0, a|o) \sim \mathcal{N}(0, I_d)$ to data $a_1 = \psi(1, a|o) \sim \pi_\beta(a|o)$, where the intermediate states satisfy $a_t = a_r + \int_r^t v(a_\tau, \tau|o) d\tau = a_r + (t-r)u(a_r, r, t|o)$.

We use the neural network $u_\theta(a_r, r, t|o)$ to represent the policy $\pi_\theta(a|o)$ for action generation. A natural extension of mean-flow model [12] is to train a velocity field $u_\theta(a_r, r, t|o)$ that maps interpolated the noisy actions to their expected displacement:

$$\mathcal{L}_{bc}(\theta) = \mathbb{E}[\|u_\theta(a_r, r, t|o) - \text{stopgrad}(u_{\text{target}})\|_2^2]. \quad (2)$$

We sample $(o, a) \sim \mathcal{D}$ from the fixed dataset, $(t, r) \sim U([0, 1])$ uniformly, and set $a_r = (1-r)\epsilon + ra$ with $\epsilon \sim \mathcal{N}(0, I)$ to interpolate between the action and Gaussian noise for simplicity as a common practice [12, 27]. The regression target is $u_{\text{target}} = v(a_r, r|o) - (r-t) \frac{d}{dr}u_\theta(a_r, r, t|o)$, where the temporal derivative is given by $\frac{d}{dr}u_\theta = v(a_r, r|o) \partial_{a_r}u_\theta + \partial_r u_\theta$. We apply $\text{stopgrad}(\cdot)$ to prevent double backpropagation and stabilize training. Based on Equation (2), we train the mean velocity field to approximate the behavior policy π_β . At inference, actions are generated via the update rule:

$$a_t = a_r + (t-r)u_\theta(a_r, r, t|o), \quad (3)$$

starting from $a_0 \sim \mathcal{N}(0, I)$ until $t = 1$. For one-step sampling, we compute $\tilde{a} = a_0 + u_\theta(a_0, 0, 1|o)$ using a single network evaluation.

However, directly applying this mechanism to offline MARL faces two limitations: (i) Training Instability and Inefficiency: Standard mean-flow losses require estimating target velocities through both forward and backward passes, involving gradients w.r.t. a_r and t . This leads to unstable supervision, inaccurate gradients, and high memory consumption in multi-agent settings. (ii) Temporal Mismatch: Uniform timestep sampling treats all $t \in [0, 1]$ equally,

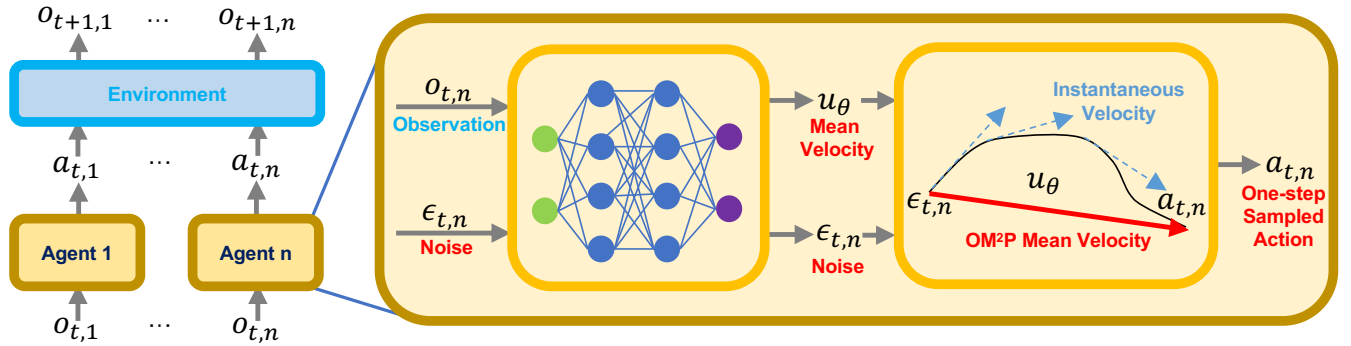


Figure 1: Overview of the decentralized OM²P framework. OM²P illustrates a single representative agent performing scalable, one-step action generation to avoid the computational bottleneck of iterative sampling inherent to multi-agent settings.

while our goal is achieving accurate one-step action generation (i.e., $t = 1$). This mismatch degrades the performance of generating actions. Furthermore, the objective mismatch—where offline MARL aims to maximize the expected episodic return, while mean-flow model focuses on fitting the data distribution—can lead to suboptimal policy learning by failing to emphasize reward-relevant behaviors.

To overcome these challenges, OM²P innovatively introduces *adaptive timestep distribution* and *derivative-free velocity estimation* for more efficient mean-flow generation. Additionally, it leverages Q-value guidance to surpass the behavior policy [37, 47]. Attributed to these elaborate designs, our proposed OM²P consistently outperforms existing baselines.

4.2 Key Components of OM²P

We now outline the essential components of OM²P: an adaptive objective enabled by generalized timestep sampling, derivative-free velocity estimation, and Q-value guidance for policy improvement.

Generalized Timestep Sampling. Instead of [12, 27], we introduce a generalized timestep distribution:

$$p(t; \xi) = \frac{1}{Z(\xi)} \exp(\xi^T h(t)), \quad (4)$$

where ξ is a fixed coefficient vector that parameterizes and thereby determines the functional form of the timestep distribution, $Z(\xi)$ is the normalizing constant, and $h(t) = [\log t, \log(1-t), t, 1]^T$ is a polynomial feature basis. Moreover, we set $r = 1 - t$ to ensure that the distribution of t also governs that of r , enabling coupled control over the interpolation and learning emphasis.

This formulation generalizes hand-crafted schedules (e.g., Beta, exponential) via a reparameterized timestep distribution $p(t; \xi)$ compared with a default uniform distribution in [12, 27], enabling task-adaptive one-step generation by controlling ξ values, with $h(t)$ including polynomial and logarithmic terms for flexible shaping without explicitly computing the normalizing constant $Z(\xi)$. Focusing more on informative timesteps (e.g., near $t=1$) improves gradient quality and accelerates policy learning. As shown in Figure 2 (left), using a generalized distribution ($\xi = [5, 5, 0, 0]$) improves stability over the uniform case ($\xi = [0, 0, 0, 0]$) under pure behavior cloning, demonstrating that more flexible timestep distributions better emphasize informative regions and enhance learning.

Derivative-Free Velocity Estimation. To mitigate the high memory cost and instability of gradient-based calculation, we introduce a derivative-free velocity estimation approach that circumvents partial derivative computation while maintaining training effectiveness. We avoid costly backpropagation through intermediate r by estimating the time derivative via finite differences:

$$\frac{du_\theta}{dr} \approx \frac{u_\theta(a_{r+\Delta r}, r + \Delta r, t|o) - u_\theta(a_r, r, t|o)}{\Delta r}. \quad (5)$$

Here the calculation of $a_{r+\Delta r}$ and a_r follows the same interpolation path for consistency. By avoiding partial derivatives of the mean velocity and using a forward-only finite difference approximation, our derivative-free method eliminates second-order gradient tracking, substantially reduces memory consumption, and enhances stability and computational efficiency in high-dimensional settings. As illustrated in Figure 2 (right), when $\Delta r \leq 10^{-8}$ as a float32 value, the finite-difference estimator yields stable performance nearly indistinguishable from that using exact gradients. This empirical observation validates the reliability of our derivative-free approach within practical numerical ranges.




Figure 2: Average episodic return measured at different training timesteps in MPE World task via expert dataset. Left: Behavior cloning performance using a Beta distribution ($\xi = [5, 5, 0, 0]$) vs. uniform ($\xi = [0, 0, 0, 0]$). Non-uniform weighting improves stability by emphasizing critical timesteps. Right: Performance under varying finite-difference step sizes Δr for estimating $\frac{du_\theta}{dr}$. Values $\Delta r \leq 10^{-8}$ yield results comparable to exact gradients, validating the reliability of our approximation.

Policy Improvement of One-Step Mean-Flow Policy. While our mean-flow framework effectively models action generation, optimizing only the behavior cloning loss via Equations (2), (4), and

(5) is insufficient for offline MARL. These objectives merely encourage imitation of the dataset behavior, without driving improvement. Even with refined timestep modeling and derivative-free enhancements, the policy remains bound to suboptimal actions unless guidance by reward signals.

To address this limitation, we incorporate a value-based objective that complements the mean-flow loss with Q-value guidance, following the principle that actions should not only match the data distribution but also yield higher expected returns. Specifically, we define the policy loss as:

$$\mathcal{L}(\theta) = \mathcal{L}_{bc}(\theta) + \mathcal{L}_q(\theta) = \mathcal{L}_{bc}(\theta) - \eta \mathbb{E}[Q_\phi(o, \tilde{a})], \quad (6)$$

where $(o, a) \sim \mathcal{D}$ is drawn from the offline dataset, and \tilde{a} is sampled from the mean-flow policy. The Q-guided term promotes actions with higher estimated returns, enabling the policy to depart from the behavior distribution and favor more rewarding choices.

To ensure effective guidance, the critic is trained following [37]:

$$\mathcal{L}_\phi = \mathbb{E} \left[\left(r + \gamma \text{mean}_{j=1,2} Q_{\phi_j'}(o', \tilde{a}') - Q_\phi(o, a) \right)^2 \right], \quad (7)$$

where $(o, a, r, o') \sim \mathcal{D}$, \tilde{a}' is sampled from the mean-flow policy via one-step action sampling and γ is the discounted factor. The integration of reward-based supervision with smooth action modeling facilitates performance beyond behavior cloning.

4.3 OM²P Algorithm

Algorithm 1 Offline Multi-Agent Mean-Flow Policy (OM²P)

- 1: **Input:** Initialize Q-networks $Q_{\phi_j}^1, Q_{\phi_j}^2$, policy network π_j with random parameters $\phi_j^1, \phi_j^2, \theta_j$, target networks with $\bar{\phi}_j^1 \leftarrow \phi_j^1, \bar{\phi}_j^2 \leftarrow \phi_j^2, \bar{\theta}_j \leftarrow \theta_j$ for each agent $j = 1, \dots, N$ and dataset \mathcal{D} . // **Initialization**
 - 2: **for** training step $t = 1$ **to** T **do**
 - 3: **for** agent $j = 1$ **to** n **do**
 - 4: Sample a random minibatch of S samples (o_j, a_j, r_j, o'_j) from dataset \mathcal{D}' . // **Sampling**
 - 5: Update critics ϕ_j^1, ϕ_j^2 to minimize (7). // **Update Critic**
 - 6: Update the actor θ_j to minimize (6). // **Update Mean-Flow Policy**
 - 7: Update target networks: $\bar{\phi}_j^k \leftarrow \rho \phi_j^k + (1 - \rho) \bar{\phi}_j^k, (k = 1, 2), \bar{\theta}_j \leftarrow \rho \theta_j + (1 - \rho) \bar{\theta}_j$.
 - 8: **end for**
 - 9: **end for**
-

Summarizing the above, we present the pseudo-code of OM²P in Algorithm 1. It constitutes a decentralized offline MARL framework, where each agent $i = 1, \dots, N$ maintains its own critic networks $Q_{\phi_1}^i, Q_{\phi_2}^i$ and a policy network π_θ^i . Line 1 initializes all critic and actor networks along with their respective target counterparts. Line 4 performs mini-batch sampling from the offline dataset \mathcal{D} to construct training data. The critic networks are updated in Line 5 by minimizing a Bellman regression loss (Equation (7)), ensuring reliable value estimation. Subsequently, Line 6 updates the actor

networks by minimizing the regularized mean-flow objective (Equation (6)), integrating both behavioral fidelity and reward supervision. Line 7 applies soft update to the target networks, stabilizing the learning process.

OM²P presents a novel integration of mean-flow generative modeling into decentralized offline MARL, offering an efficient solution for one-step action generation without policy distillation. To emphasize its practical advantages, we summarize the framework’s principal strengths below:

Efficient policy learning. In contrast to iterative diffusion- and flow-based methods (e.g., Diff-QL [47], SfBC [5]), OM²P achieves efficient one-step synthesis via a unified objective that strictly couples adaptive mean-flow matching with reward-aware supervision. This tailored integration aligns generative modeling with value maximization, ensuring that policy learning remains grounded in the data distribution while significantly reducing sampling and computational overheads.

Stable training with reduced computation. To overcome the convergence instability of standard flow models in RL, OM²P employs generalized timestep sampling and a derivative-free velocity approximation to eliminate costly second-order gradient computations. Different from FQL [37] that rely on complex rollouts and distillation, our native one-step design enhances training stability and efficiency, offering a specialized solution for the high computational demands of multi-agent environments.

Modular and scalable architecture. OM²P features a modular architecture where each component—optimized for stability or efficiency—avoids additive complexity. Computational costs scale linearly with agents, limited to target estimation and single-pass inference. By bypassing multi-step sampling and distillation pipelines, OM²P ensures scalability, utilizing value guidance within the mean-flow framework to maintain robust performance and mitigate risks in high-dimensional offline MARL.

5 EXPERIMENTS

We assess the effectiveness of our proposed method across diverse multi-agent environments and datasets. Our evaluation emphasizes these key aspects: (1) **Performance** — assessing how OM²P compares to existing state-of-the-art offline MARL algorithms in standard benchmarks; (2) **Efficiency** — examining whether our algorithm can maintain competitive performance while reducing memory consumption during training, and whether it can leverage efficient one-step sampling to generate high-quality policies during both training and inference.

5.1 Experimental Setup

Environments. We perform comprehensive evaluations of our method on two commonly used benchmarks for multi-agent reinforcement learning: the Multi-Agent Particle Environment (MPE) [30] and the Multi-Agent MuJoCo (MAMuJoCo) suite [39]. For our experiments, we choose the Predator-Prey, World, and Cooperative Navigation scenarios from MPE, and a two-agent HalfCheetah setup from MAMuJoCo.

Datasets. Following the dataset generation protocol proposed in [10], we construct four types of datasets to reflect a range of policy qualities. These include medium-replay, medium, medium-expert,

Predator–Prey	OMAR	MA-SfBC	MA-FQL	OM ² P
Medium-Replay	86.8±43.7	26.1±10.0	203.8±43.3	203.1±59.3
Medium	116.9±45.2	127.0±50.9	255.1±65.0	257.3±52.0
Medium-Expert	128.3±35.2	152.3±41.2	233.0±42.5	230.2±71.1
Expert	202.8±27.1	256.0±26.9	289.5±28.6	330.3±40.9
Avg.	133.7±37.8	140.4±32.3	245.4±44.9	255.2±55.9
World	OMAR	MA-SfBC	MA-FQL	OM ² P
Medium-Replay	21.1±15.6	9.1±5.9	107.0±30.0	106.0±10.2
Medium	45.6±16.0	54.2±22.7	136.1±18.5	142.0±18.6
Medium-Expert	71.5±28.2	60.6±22.9	120.7±16.9	126.9±30.2
Expert	84.8±21.0	97.3±19.1	139.4±13.9	175.0±21.5
Avg.	55.8±20.2	55.3±17.7	128.1±19.8	137.5±20.1
Cooperative Navigation	OMAR	MA-SfBC	MA-FQL	OM ² P
Medium-Replay	260.7±37.7	196.1±11.1	375.6±41.7	380.3±61.8
Medium	348.7±51.7	276.3±8.8	469.6±42.1	441.2±21.3
Medium-Expert	450.3±39.0	299.8±16.8	525.0±25.9	510.0±14.5
Expert	564.6±8.6	553.0±41.1	582.6±26.0	614.0±4.8
Avg.	406.1±34.2	331.3±19.4	488.2±33.9	486.4±25.6
HalfCheetah-v2	OMAR	MA-SfBC	MA-FQL	OM ² P
Medium-Replay	1674.8±201.5	-128.3±71.3	2553.6±337.4	2525.9±441.3
Medium	2797.0±445.7	1386.8±248.8	2172.3±242.3	2673.9±225.1
Medium-Expert	2900.2±403.2	1392.3±190.3	2655.2±368.6	2917.0±301.2
Expert	2963.8±410.5	2386.6±440.3	3869.8±329.3	4151.6±275.0
Avg.	2583.9±365.2	1259.3±237.7	2812.7±319.4	3067.1±310.7

Table 1: Comparative performance of OM²P against OMAR, MA-SfBC and MA-FQL under different tasks and datasets. Values within 1% of the best performance in each row are highlighted in bold.

and expert datasets. Each type captures different levels of behavior performance and data coverage. Full descriptions of the dataset construction process and characteristics are listed below:

- **Medium-replay:** includes all transitions stored in the replay buffer during training until the policy reaches medium-level performance.
- **Medium:** consists of 1 million transitions generated by unrolling a policy whose performance corresponds to the medium level.
- **Expert:** consists of 1 million transitions collected by executing a well-trained expert policy.
- **Medium-expert:** formed by mixing medium and expert data in fixed proportions (90% medium and 10% expert for MPE; 99.9% medium and 0.1% expert for MAMuJoCo).

Network Structures. In OM²P, we employ a multi-layer perceptron (MLP) to parameterize the mean-velocity function, which serves as the core of our policy network. The MLP takes the observation, noise and timestep concatenation as inputs and outputs the estimated mean-velocity vector used for one-step action generation. This design offers a lightweight and efficient architecture, enabling faster inference and lower memory consumption. All critic networks, including those in baseline methods such as OMAR and

MA-FQL [36, 37], similarly adopt MLPs that map concatenated state-action pairs to scalar Q-values. Unless otherwise noted, all networks use GeLU activation and layer normalization after activation, four hidden layers with 512 neurons, and a consistent feature dimension across agents.

Hyperparameters. Across all tasks in MPE and MAMuJoCo, we use a unified learning rate of 3×10^{-4} for both the Q-function and the mean-flow policy network. Due to variations in dataset quality and underlying data distributions, we adopt different values for the Q-weight and the timestep sampling distribution. The specific configurations used for each task are provided in Table 2. Other fixed hyperparameters are shown in Table 3.

Baselines. We benchmark our proposed OM²P algorithm against several representative offline multi-agent reinforcement learning methods, including OMAR [36], a multi-agent adaptation of the diffusion-based SfBC policy [5] (denoted MA-SfBC), and a multi-agent adaptation of flow-based FQL [37] (denoted MA-FQL). These methods are implemented based on decentralized training with independent actor-critic structures. Besides, we compare our algorithm with MADIFF[54], CFCQL[42] and DoF[24] as a centralized

Table 2: The η value and the ξ value in OM²P.

Predator Prey	η	ξ
Medium-Replay	100.0	[0, 0, 0, 0]
Medium	10.0	[0, 0, 0.2, 0]
Medium-Expert	100.0	[2, 0, 0, 0]
Expert	1.0	[0, 0, 0, 0]
World	η	ξ
Medium-Replay	100.0	[0, 0, 0.2, 0]
Medium	1000.0	[0, 0, 0.2, 0]
Medium-Expert	10000.0	[0, 0, 0, 0]
Expert	1.0	[0, 0, 0, 0]
Cooperative Navigation	η	ξ
Medium-Replay	1.0	[0, 0, 0, 0]
Medium	1.0	[0, 0, 0, 0]
Medium-Expert	0.1	[0, 0, 0.05, 0]
Expert	1.0	[0, 0, 0.1, 0]
HalfCheetah-v2	η	ξ
Medium-Replay	1.0	[0, 0, 0, 0]
Medium	0.125	[0, 0, 0, 0]
Medium-Expert	0.02	[0, 0, 0, 0]
Expert	0.008	[0, 0, 0, 0]

training decentralized execution algorithm in MPE tasks (Due to the dataset inconsistency, the results are shown in the page¹).

Devices and Experimental Details. We run all algorithms on a single NVIDIA RTX A6000 GPU. To ensure statistical reliability, we conduct each algorithm using five different random seeds and report the mean and standard deviation of the evaluation returns. Each model is trained for 1 million timesteps, with evaluations performed every 100 timesteps. During each evaluation, the policy is tested over 10 episodes, and the final performance is reported as the mean and standard deviation across the five seeds at the same timestep.

Table 3: Hyperparameters in OM²P.

Hyperparameter	Value
Batch Size	256
Discount Factor γ	0.99
Target Network Update Rate ρ	0.005
Temporal Difference Δr	10^{-12}

5.2 Empirical Results

Standard Tasks and Datasets. OM²P consistently achieves superior performance across both MPE and MA-MuJoCo benchmarks, as shown in Table 1. It outperforms OMAR and MA-SfBC by substantial margins and matches or surpasses MA-FQL in most scenarios. Thanks to its one-step policy design, OM²P demonstrates stable learning across datasets of varying quality, while significantly improving training and inference efficiency. On the more challenging

multi-agent HalfCheetah-v2 domain, OM²P maintains strong scalability and robustness, achieving nearly the highest average return under all dataset regimes and excelling particularly on expert data. These results collectively highlight OM²P’s effectiveness as a scalable, efficient, and high-performing solution for offline multi-agent reinforcement learning.

Training Efficiency. To justify the effectiveness of OM²P algorithm, we evaluate the training efficiency of OM²P in terms of GPU memory usage and training time over 10,000 steps, evaluating 10 episodes per 100 steps under the MPE World task using the expert dataset, which is shown in Table 4. Notably, OM²P maximizes efficiency via derivative-free, native one-step generation, bypassing the high memory costs and distillation overheads typical of baselines (MA-FQL). Specifically, OM²P reduces GPU memory usage by over 37% (from 1036MB to 650MB) and cuts training time by more than 90% (from 5674s to 564s for 10.1× improvement) relative to the diffusion-based MA-SfBC. It also achieves over 28% lower memory usage (906MB to 650MB) and approximately 50% faster evaluation (169s to 85s) than the flow-based MA-FQL. Furthermore, when replacing our derivative-free approximation with full gradient computation, memory usage jumps to over 2.4GB—demonstrating the importance of our efficient training design to achieve 3.8× GPU memory reduction. The results demonstrate that OM²P achieves consistently low memory overhead and fast runtime, highlighting its effectiveness in resource-constrained environments.

Table 4: Training and inference efficiency comparison via the GPU memory usage, training time (training 10,000 steps) and evaluating time (summation of evaluating 10 episodes per 100 steps during the process of training 10000 steps) across different methods and ablations with exact gradient value under the MPE World task using the expert dataset.

Algorithm	GPU Use	Training time	Evaluation time
MA-SfBC	1036MB	5674s	975s
MA-FQL	906MB	585s	169s
OM ² P w/ grad.	2442MB	633s	81s
OM²P (ours)	650MB	564s	85s

Scalability. To evaluate the scalability of our framework, we conduct our experiments on the Cooperative Navigation task for large number of agents (The baseline task has three agents). The datasets are collected using the policies trained by HATD3[53] algorithm for tasks with more agents. The dataset has 40000 trajectories and each trajectory has 25 timesteps. Table 5 below shows that OM²P algorithm outperforms other baseline algorithms, which means that our algorithm holds strong scalability among more agents.

Table 5: Performance comparison with more agents in Cooperative Navigation tasks with 4 and 5 agents.

Algorithm	4 Agents	5 Agents
OMAR	447.1±68.1	379.8±66.9
MA-SfBC	450.6±43.1	367.0±47.8
MA-FQL	586.1±43.4	615.7±46.9
OM²P (Ours)	587.2±26.9	621.8±70.1

¹<https://arxiv.org/abs/2508.06269>

Overall, the results demonstrate the advantages of OM²P across benchmarks. It consistently matches or surpasses diffusion- and flow-based baselines while maintaining stable performance across datasets of varying quality. Its one-step mean-flow design yields substantial efficiency gains, reducing GPU memory and cutting training time up to tenfold compared to prior generative approaches. OM²P also scales effectively to larger agent populations, outperforming strong baselines in cooperative tasks with four and five agents. These findings highlight OM²P’s combination of high performance, efficiency, and scalability.

5.3 Ablation Study

We perform ablation studies to investigate these key aspects: the sensitivity of critical hyperparameters, the isolated contributions of core components within OM²P, the robustness across dataset sizes, and the scalability to large numbers of agents.

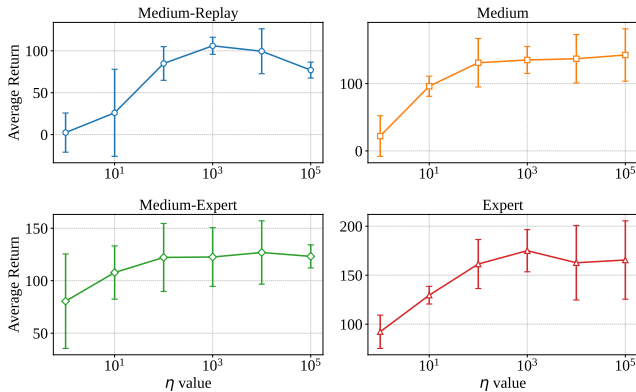


Figure 3: Performance impact of the coefficient η in MPE World.

Hyperparameters. We examine the role of the mixing coefficient η , which balances the behavior cloning (BC) loss and the Q-learning objective. Figure 3 presents the performance of our method across four datasets in the MPE World task under varying η values. The results indicate that η is a crucial hyperparameter that must be tuned in accordance with the quality of the offline dataset to achieve optimal results. Notably, small η values lead to suboptimal learning due to insufficient Q-value guidance, while moderate to large values (e.g., 10^3) consistently deliver stronger performance, particularly on high-quality datasets. This underscores the importance of properly calibrating η to strike a balance between policy imitation and reward-driven optimization.

Ablation on different components. To assess the contribution of each key component in OM²P, we conduct modular ablations by systematically removing or modifying core elements of the framework. Specifically, we evaluate three simplified variants: using only the Q-learning objective (removing the BC term), using only behavior cloning (removing the value-guided objective), and replacing the general timestep distribution with a uniform distribution. As shown in Figure 4, each ablation leads to a noticeable drop in performance compared to the full OM²P. These results collectively demonstrate that the synergy among all components—Q supervision, BC guidance, and timestep reweighting—is critical to realizing the full potential of OM²P.




Figure 4: Effect of removing key components from OM²P on the MPE World task. Each ablation variant shows reduced average returns compared to the full model, confirming the importance of each module.

Robustness across dataset scales. To further evaluate the robustness of our method, we conduct ablation studies on different sizes of offline datasets under the *World* task expert dataset. As shown in Table 6, OM²P consistently maintains high performance even when the dataset size is significantly reduced, demonstrating strong robustness and data efficiency in offline MARL.

Table 6: Performance of OM²P under different dataset sizes on the *World* task expert dataset.

Number of Trajectories	4e4	3.5e4	3e4
Performance of OM ² P	175.0±21.5	160.2±5.2	154.3±10.4

The ablation results highlight the effectiveness and robustness of OM²P. Proper tuning of the mixing coefficient η is essential to balance policy imitation and reward guidance, with moderate to large values yielding consistently strong performance. Modular ablations confirm that each core component—Q-function supervision, behavior cloning, and generalized timestep weighting—contributes substantially to overall performance, and removing any element degrades returns. Furthermore, OM²P maintains high performance even under reduced dataset sizes, demonstrating strong data efficiency and robustness. Collectively, these findings validate the design choices of OM²P and underscore its ability to deliver stable, efficient, and scalable offline multi-agent reinforcement learning.

6 CONCLUSION

In this paper, we propose OM²P, a novel offline multi-agent RL algorithm that integrates mean-flow model into one-step policy generation. OM²P addresses core limitations of existing generative methods by aligning training with reward-based objectives, introducing adaptive timestep sampling, and employing a memory-efficient derivative-free velocity estimation. Extensive experiments on MPE and MAMuJoCo demonstrate that OM²P achieves state-of-the-art performance with significantly improved efficiency. OM²P efficiently integrates mean-flow model into offline MARL, enabling scalable, efficient, and high-quality policy learning in complex multi-agent settings.

7 ACKNOWLEDGEMENT

This work is supported by the National Natural Science Foundation of China Grant 52494974.

REFERENCES

- [1] Anurag Ajay, Yilun Du, Abhi Gupta, Joshua B Tenenbaum, Tommi S Jaakkola, and Pulkit Agrawal. 2023. Is Conditional Generative Modeling all you need for Decision Making?. In *The Eleventh International Conference on Learning Representations*.
- [2] Michael Samuel Albergo and Eric Vanden-Eijnden. 2023. Building Normalizing Flows with Stochastic Interpolants. In *The Eleventh International Conference on Learning Representations*.
- [3] Marvin Alles, Nutan Chen, Patrick van der Smagt, and Botond Cseke. 2025. FlowQ: Energy-Guided Flow Policies for Offline Reinforcement Learning. arXiv:2505.14139 [cs.LG] <https://arxiv.org/abs/2505.14139>
- [4] Paul Barde, Jakob Foerster, Derek Nowrouzezahrai, and Amy Zhang. 2024. A Model-Based Solution to the Offline Multi-Agent Reinforcement Learning Coordination Problem. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems (Auckland, New Zealand) (AAMAS '24)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 141–150.
- [5] Huayu Chen, Cheng Lu, Chengyang Ying, Hang Su, and Jun Zhu. 2023. Offline Reinforcement Learning via High-Fidelity Generative Behavior Modeling. In *The Eleventh International Conference on Learning Representations*.
- [6] Cheng Chi, Siyuan Feng, Yilun Du, Zhenjia Xu, Eric Cousineau, Benjamin Burchfiel, and Shuran Song. 2023. Diffusion Policy: Visuomotor Policy Learning via Action Diffusion. In *Proceedings of Robotics: Science and Systems (RSS)*.
- [7] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. 2024. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*.
- [8] Yuming Feng, Chuye Hong, Yaru Niu, Shiqi Liu, Yuxiang Yang, Wenhao Yu, Tingnan Zhang, Jie Tan, and Ding Zhao. 2025. Learning Multi-Agent Locomotion for Long-Horizon Quadrupedal Pushing. arXiv:2411.07104 [cs.RO] <https://arxiv.org/abs/2411.07104>
- [9] Claude Formanek, Asad Jeewa, Jonathan Shock, and Arnu Pretorius. 2023. Off-the-Grid MARL: Datasets and Baselines for Offline Multi-Agent Reinforcement Learning. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*. 2442–2444.
- [10] Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. 2020. D4rl: Datasets for deep data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219* (2020).
- [11] Yuqian Fu, Yuanheng Zhu, Jian Zhao, Jiajun Chai, and Dongbin Zhao. 2025. INS: Interaction-aware Synthesis to Enhance Offline Multi-agent Reinforcement Learning. In *The Thirteenth International Conference on Learning Representations*. <https://openreview.net/forum?id=kx2DLIPr40>
- [12] Zhengyang Geng, Mingyang Deng, Xingjian Bai, J Zico Kolter, and Kaiming He. 2025. Mean Flows for One-step Generative Modeling. *arXiv preprint arXiv:2505.13447* (2025).
- [13] Wonjoon Goo and Scott Niekum. 2022. Know Your Boundaries: The Necessity of Explicit Behavioral Cloning in Offline RL. arXiv:2206.00695 [cs.LG] <https://arxiv.org/abs/2206.00695>
- [14] Mohamad A. Hady, Siyi Hu, Mahardhika Pratama, Jimmy Cao, and Ryszard Kowalczyk. 2025. Multi-Agent Reinforcement Learning for Resources Allocation Optimization: A Survey. arXiv:2504.21048 [cs.MA] <https://arxiv.org/abs/2504.21048>
- [15] Tiankai Hang, Shuyang Gu, Xin Geng, and Baining Guo. 2024. Improved Noise Schedule for Diffusion Training. arXiv:2407.03297 [cs.CV] <https://arxiv.org/abs/2407.03297>
- [16] Philippe Hansen-Estruch, Ilya Kostrikov, Michael Janner, Jakub Grudzien Kuba, and Sergey Levine. 2023. Idql: Implicit q-learning as an actor-critic method with diffusion policies. *arXiv preprint arXiv:2304.10573* (2023).
- [17] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* 33 (2020), 6840–6851.
- [18] Michael Janner, Yilun Du, Joshua Tenenbaum, and Sergey Levine. 2022. Planning with Diffusion for Flexible Behavior Synthesis. In *International Conference on Machine Learning*. PMLR, 9902–9915.
- [19] Jiechuan Jiang and Zongqing Lu. 2023. Offline Decentralized Multi-Agent Reinforcement Learning. In *ECAL*. 1148–1155.
- [20] Bingyi Kang, Xiao Ma, Chao Du, Tianyu Pang, and Shuicheng Yan. 2023. Efficient diffusion policies for offline reinforcement learning. *Advances in Neural Information Processing Systems* 36 (2023), 67195–67212.
- [21] Prajwal Kairala and Cody Fleming. 2025. Flow-Based Single-Step Completion for Efficient and Expressive Policy Learning. arXiv:2506.21427 [cs.LG] <https://arxiv.org/abs/2506.21427>
- [22] Jaewoo Lee, Sujin Yun, Taeyoung Yun, and Jinkyoo Park. 2024. Gta: Generative trajectory augmentation with guidance for offline reinforcement learning. *Advances in Neural Information Processing Systems* 37 (2024), 56766–56801.
- [23] Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. 2020. Offline Reinforcement Learning: Tutorial, Review, and Perspectives on Open Problems. arXiv:2005.01643 [cs.LG] <https://arxiv.org/abs/2005.01643>
- [24] Chao Li, Ziwei Deng, Chenxing Lin, Wenqi Chen, Yongquan Fu, Weiquan Liu, Chenglu Wen, Cheng Wang, and Siqi Shen. 2025. DoF: A Diffusion Factorization Framework for Offline Multi-Agent Decision Making. In *International Conference on Learning Representations*.
- [25] Zhuoran Li, Ling Pan, and Longbo Huang. 2023. Beyond Conservatism: Diffusion Policies in Offline Multi-agent Reinforcement Learning. arXiv:2307.01472 [cs.AI] <https://arxiv.org/abs/2307.01472>
- [26] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. 2023. Flow Matching for Generative Modeling. In *The Eleventh International Conference on Learning Representations*.
- [27] Yaron Lipman, Marton Havasi, Peter Holderrieth, Neta Shaul, Matt Le, Brian Karrer, Ricky TQ Chen, David Lopez-Paz, Heli Ben-Hamu, and Itai Gat. 2024. Flow matching guide and code. *arXiv preprint arXiv:2412.06264* (2024).
- [28] Xingchao Liu, Chengyue Gong, et al. 2023. Flow Straight and Fast: Learning to Generate and Transfer Data with Rectified Flow. In *The Eleventh International Conference on Learning Representations*.
- [29] Zongkai Liu, Qian Lin, Chao Yu, Xiawei Wu, Yile Liang, Donghui Li, and Xuetao Ding. 2025. Offline Multi-Agent Reinforcement Learning via In-Sample Sequential Policy Optimization. *Proceedings of the AAAI Conference on Artificial Intelligence* 39, 18 (Apr. 2025), 19068–19076. <https://doi.org/10.1609/aaai.v39i18.34099>
- [30] Ryan Lowe, Yi I Wu, Aviv Tamar, Jean Harb, OpenAI Pieter Abbeel, and Igor Mordatch. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in neural information processing systems* 30 (2017).
- [31] Cong Lu, Philip Ball, Yee Whye Teh, and Jack Parker-Holder. 2023. Synthetic experience replay. *Advances in Neural Information Processing Systems* 36 (2023), 46323–46344.
- [32] Cheng Lu, Huayu Chen, Jianfei Chen, Hang Su, Chongxuan Li, and Jun Zhu. 2023. Contrastive Energy Prediction for Exact Energy-Guided Diffusion Sampling in Offline Reinforcement Learning. *arXiv preprint arXiv:2304.12824* (2023).
- [33] Daiki E. Matsunaga, Jongmin Lee, Jaeseok Yoon, Stefanos Leonardos, Pieter Abbeel, and Kee-Eung Kim. 2023. AlberDICE: addressing out-of-distribution joint actions in offline multi-agent RL via alternating stationary distribution correction estimation. In *Proceedings of the 37th International Conference on Neural Information Processing Systems (New Orleans, LA, USA) (NIPS '23)*. Curran Associates Inc., Red Hook, NY, USA, Article 3175, 31 pages.
- [34] Linghui Meng, Muning Wen, Chenyang Le, Xiyun Li, Dengpeng Xing, Weinan Zhang, Ying Wen, Haifeng Zhang, Jun Wang, Yaodong Yang, et al. 2023. Offline pre-trained multi-agent decision transformer. *Machine Intelligence Research* 20, 2 (2023), 233–248.
- [35] Frans A Oliehoek and Christopher Amato. 2016. *A concise introduction to decentralized POMDPs*. Springer.
- [36] Ling Pan, Longbo Huang, Tengyu Ma, and Huazhe Xu. 2022. Plan better amid conservatism: Offline multi-agent reinforcement learning with actor rectification. In *International Conference on Machine Learning*. PMLR, 17221–17237.
- [37] Seohong Park, Qiyang Li, and Sergey Levine. 2025. Flow Q-Learning. arXiv:2502.02538 [cs.LG] <https://arxiv.org/abs/2502.02538>
- [38] Tim Pearce, Tabish Rashid, Anssi Kanervisto, Dave Bignell, Mingfei Sun, Raluca Georgescu, Sergio Valcarcel Macua, Shan Zheng Tan, Ida Momennjad, Katja Hofmann, et al. 2024. Imitating Human Behaviour with Diffusion Models. In *The Eleventh International Conference on Learning Representations*.
- [39] Bei Peng, Tabish Rashid, Christian Schroeder de Witt, Pierre-Alexandre Kamieny, Philip Torr, Wendelin Böhrer, and Shimon Whiteson. 2021. Facmac: Factored multi-agent centralised policy gradients. *Advances in Neural Information Processing Systems* 34 (2021), 12208–12221.
- [40] Dan Qiao, Wenhao Li, Shanchao Yang, Hongyuan Zha, and Baoxiang Wang. 2025. Offline Multi-agent Reinforcement Learning via Score Decomposition. arXiv:2505.05968 [cs.LG] <https://arxiv.org/abs/2505.05968>
- [41] Allen Z Ren, Justin Lidard, Lars Lien Ankile, Anthony Simeonov, Pulkit Agrawal, Anirudha Majumdar, Benjamin Burchfiel, Hongkai Dai, and Max Simchowitz. 2024. Diffusion Policy Policy Optimization. In *CoRL 2024 Workshop on Mastering Robot Manipulation in a World of Abundant Data*.
- [42] Jianzhun Shao, Yun Qu, Chen Chen, Hongchang Zhang, and Xiangyang Ji. 2023. Counterfactual Conservative Q Learning for Offline Multi-agent Reinforcement Learning. In *Thirty-seventh Conference on Neural Information Processing Systems*. <https://openreview.net/forum?id=62zmO4mv8X>
- [43] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. 2021. Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations*.
- [44] Qi Tian, Kun Kuang, Furui Liu, and Baoxiang Wang. 2023. Learning from Good Trajectories in Offline Multi-Agent Reinforcement Learning. *Proceedings of the AAAI Conference on Artificial Intelligence* 37, 10 (Jun. 2023), 11672–11680. <https://doi.org/10.1609/aaai.v37i10.26379>
- [45] Wei-Cheng Tseng, Tsun-Hsuan Johnson Wang, Yen-Chen Lin, and Phillip Isola. 2022. Offline Multi-Agent Reinforcement Learning with Knowledge Distillation. *Advances in Neural Information Processing Systems* 35 (2022), 226–237.
- [46] Xiangsen Wang and Xianyan Zhan. 2023. Offline Multi-Agent Reinforcement Learning with Coupled Value Factorization. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*. 2781–2783.

- [47] Zhendong Wang, Jonathan J Hunt, and Mingyuan Zhou. 2023. Diffusion Policies as an Expressive Policy Class for Offline Reinforcement Learning. In *The Eleventh International Conference on Learning Representations*.
- [48] Muning Wen, Jakub Kuba, Runji Lin, Weinan Zhang, Ying Wen, Jun Wang, and Yaodong Yang. 2022. Multi-agent reinforcement learning is a sequence modeling problem. *Advances in Neural Information Processing Systems* 35 (2022), 16509–16521.
- [49] Yuqing Xie, Chao Yu, Hongzhi Zang, Feng Gao, Wenhao Tang, Jingyi Huang, Jiayu Chen, Botian Xu, Yi Wu, and Yu Wang. 2025. Multi-UAV Formation Control with Static and Dynamic Obstacle Avoidance via Reinforcement Learning. arXiv:2410.18495 [cs.RO] <https://arxiv.org/abs/2410.18495>
- [50] Yiqin Yang, Xiaoteng Ma, Chenghao Li, Zewu Zheng, Qiyuan Zhang, Gao Huang, Jun Yang, and Qianchuan Zhao. 2021. Believe what you see: Implicit constraint approach for offline multi-agent reinforcement learning. *Advances in Neural Information Processing Systems* 34 (2021), 10299–10312.
- [51] Ruiqi Zhang, Jing Hou, Florian Walter, Shangding Gu, Jiayi Guan, Florian Röhrbein, Yali Du, Panpan Cai, Guang Chen, and Alois Knoll. 2024. Multi-Agent Reinforcement Learning for Autonomous Driving: A Survey. arXiv:2408.09675 [cs.AI] <https://arxiv.org/abs/2408.09675>
- [52] Shiyuan Zhang, Weitong Zhang, and Quanquan Gu. 2025. Energy-Weighted Flow Matching for Offline Reinforcement Learning. In *The Thirteenth International Conference on Learning Representations*.
- [53] Yifan Zhong, Jakub Grudzien Kuba, Xidong Feng, Siyi Hu, Jiaming Ji, and Yaodong Yang. 2024. Heterogeneous-Agent Reinforcement Learning. *Journal of Machine Learning Research* 25, 32 (2024), 1–67. <http://jmlr.org/papers/v25/23-0488.html>
- [54] Zhengbang Zhu, Minghuan Liu, Liyuan Mao, Bingyi Kang, Minkai Xu, Yong Yu, Stefano Ermon, and Weinan Zhang. 2024. Madiff: Offline multi-agent learning with diffusion models. *Advances in Neural Information Processing Systems* 37 (2024), 4177–4206.