

Offline Meta-Reinforcement Learning with Flow-Based Task Inference and Adaptive Correction of Feature Overgeneralization

Min Wang¹, Xin Li^{1 2*}, Mingzhong Wang³, Hasnaa Bennis¹

¹ Beijing Institute of Technology,

² Key Laboratory of Symbolic Computation and Knowledge Engineering of Ministry of Education, Jilin University,

³ University of the Sunshine Coast

{minwangcs, xinli}@bit.edu.cn, mwang@usc.edu.au, bennishasnaa1920@gmail.com

Abstract

Offline meta-reinforcement learning (OMRL) combines the strengths of learning from diverse datasets in offline RL with the adaptability to new tasks of meta-RL, promising safe and efficient knowledge acquisition by RL agents. However, OMRL still suffers extrapolation errors due to out-of-distribution (OOD) actions, compromised by broad task distributions and Markov Decision Process (MDP) ambiguity in meta-RL setups. Existing research indicates that the generalization of the Q network affects the extrapolation error in offline RL. This paper investigates this relationship by decomposing the Q value into feature and weight components, observing that while decomposition enhances adaptability and convergence in the case of high-quality data, it often leads to policy degeneration or collapse in complex tasks. We observe that decomposed Q values introduce a large estimation bias when the feature encounters OOD samples, a phenomenon we term “feature overgeneralization”. To address this issue, we propose FLORA, which identifies OOD samples by modeling feature distributions and estimating their uncertainties. FLORA integrates a return feedback mechanism to adaptively adjust feature components. Furthermore, to learn precise task representations, FLORA explicitly models the complex task distribution using a chain of invertible transformations. We theoretically and empirically demonstrate that FLORA achieves rapid adaptation and meta-policy improvement compared to baselines across various environments.

Introduction

Offline Meta-Reinforcement Learning (OMRL) focuses on quick adaptation to new tasks through training on a distribution of static and limited offline tasks. Context-based OMRL, which learns to encode task representations from histories, has gained increasing attention for more efficient and stable adaptability. In meta-RL, differences in task distributions between meta-training and meta-testing often lead to Markov Decision Process (MDP) ambiguity, where the policy struggles to distinguish between different tasks, resulting in spurious connections between trajectory data and task representations. In offline settings, extrapolation error frequently arises from the training policy executing out-of-distribution (OOD) actions, and the complexity of a meta-

task dataset composed of multiple tasks further exacerbates this issue. Therefore, addressing the following two challenges is crucial for rapid adaptation in OMRL: 1) accurately inferring the true task distribution; and 2) effectively addressing the extrapolation error problem.

Recent research (Yuan and Lu 2022; Li et al. 2024) has primarily focused on improving the accuracy of task inference through representation learning methods. Nevertheless, these methods ignore the inference and generation of specific probability distributions of task representations, are susceptible to biases introduced by the construction of positive and negative samples, and struggle to capture the latent structure of complex tasks. To better model task distributions and effectively discern different tasks, a strand of research has concentrated on predicting future states and rewards separately using the decoder (Ni et al. 2023) or the context-aware world model (Wang et al. 2024b) to derive disentangled task representations. However, these approaches frequently assume that the distribution of task representations follows a simple prior Gaussian distribution, which restricts the flexibility and expressiveness of task inference models. Moreover, the trade-off between regularization and reconstruction terms presents challenges: a higher proportion of reconstruction loss can lead the decoder to overfit, resulting in inaccuracies in the reconstructed next state and reward, ultimately diminishing the accuracy of task representations.

To address extrapolation error, most OMRL methods leverage existing offline RL algorithms as backbone frameworks, which encourage the training policy π_φ to fully mimic or stay close to the behavior policy π_β . The benefit of offline RL designed to tackle OOD actions originates from their inhibition of overgeneralization of Q networks, rather than only from their ability to prevent overestimation of Q values (Mao et al. 2024). To fully leverage history information and reduce the impact of accumulated errors in temporal difference (TD) updates, Meta-DT (Wang et al. 2024b) strives to leverage the conditional sequence modeling paradigm to effectively improve training efficiency and capture complex temporal relationships. However, the aforementioned methods ignore the disentanglement of dynamics and reward functions in the policy space, only considering the decoupled task representation in feature space. This oversight prevents them from fundamentally addressing the extrapolation error aggravated by different task structures.

*Corresponding author.

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

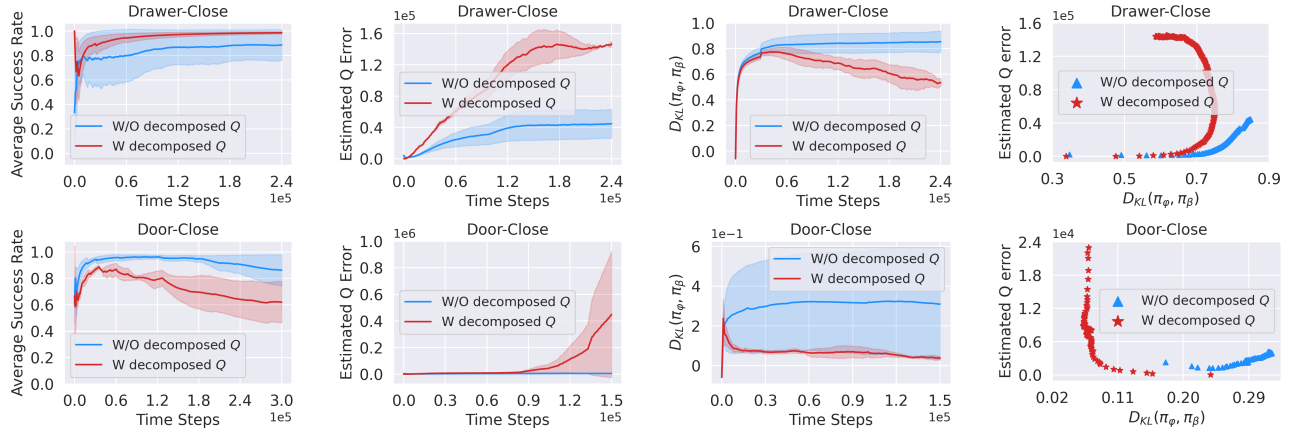


Figure 1: Motivating example: In environments with high-quality datasets or narrow task distributions (e.g. Drawer-Close), the decomposed Q value accelerates the adaptability of the training policy π_φ , and the converged optimal policy outperforms the behavior policy π_β . However, in environments with low-quality datasets or broad task distributions (e.g. Door-Close), the decomposed Q value exacerbates the overestimation issue, ultimately leading to divergence of Q value and failure of π_φ .

Inspired by Successor Features (SFs) (Barreto et al. 2017), prior work in online meta-RL (Wang et al. 2024a) has proposed a decomposition of the Q value to share common task structures and enable rapid adaptation to new tasks. This is expressed as: $Q = \psi^T W$, where context-aware SFs ψ are related to the transition dynamics, and reward weights W correspond to the reward function. This decomposition decouples task structures, facilitating the reuse of shared task features and assisting in exploring the relationship between representation and overestimation in offline settings. We found that in offline meta-RL, although decomposed Q values promote rapid policy adaptation, they present new challenges. In Fig. 1, we observe that environments with expert-level data or lower task variability, such as Drawer-Close, benefit from this decomposition, which enhances the adaptability of policy to different tasks and allows for rapid convergence. During early training, as the KL divergence between π_φ and π_β increases, the estimation error gradually amplifies. Interestingly, for the decomposed Q , forcing π_φ to align intimately with π_β further increases the estimated error in the later phase. This reflects that π_φ eventually surpasses π_β and converges to optimal, and the estimation error correspondingly stabilizes. In contrast, environments like Door-Close, which contain more suboptimal trajectories or have a broader task distribution, make it easier for agents to encounter OOD actions and lead to feature overgeneralization. Importantly, overestimation caused by an overgeneralized ψ in the decomposed Q value can be particularly problematic. As training progresses, the estimation error increases exponentially even if the KL divergence is gradually small, resulting in divergent Q values and causing π_φ to collapse.

The key to overcoming this dilemma is the identification and processing of OOD samples. In this paper, we propose **FL**ow-based task infe**R**rence and **A**daptive correction of feature overgeneralization caused by OOD actions (FLORA) in offline meta-RL scenarios. In detail, FLORA starts by employing a chain of invertible

linear-time transformations for task inference, flexibly approximating more complex task distributions and deriving compact task representations for policy learning. Subsequently, FLORA models the distribution of context-aware SFs and estimates *epistemic uncertainty* through double-feature learning to detect OOD actions. Ultimately, FLORA incorporates a reward feedback mechanism: it reduces uncertainty to discourage the agent from trying low-reward OOD actions and increases uncertainty to encourage high-reward actions, mitigating overgeneralization issues and simultaneously ensuring meta-policy improvement. Our main contributions are summarized as follows:

1. We formalize a feature overgeneralization issue that leads to degradative policies in offline meta-RL. To achieve a balance between mitigating overgeneralization and maintaining policy improvement in decoupled policy space, we propose FLORA, which frames adaptive adjustment of feature learning as the choice of conservative level of estimated uncertainty. By leveraging flow-based task inference, FLORA learns complex task distributions more flexibly and efficiently, enhancing the task inference.
2. We theoretically prove that FLORA, by leveraging a decoupled representation and policy space, achieves a superior policy compared to standard Q -value estimation.
3. We extensively evaluate FLORA across diverse environments of MuJoCo and Meta-World, demonstrating that it significantly outperforms baselines in adaptation efficiency and achieves superior asymptotic performance.

Related Work

Offline Meta-Reinforcement Learning

OMRL aims to enable rapid adaptation to new tasks by training with static and limited offline meta-tasks. It can be categorized as gradient-based or context-based. Although gradient-based methods (Mitchell et al. 2021) can adapt to a

broad range of tasks, they still suffer from the inherent sample inefficiency of nested policy gradient updates.

Context-based OMRL, in contrast, alleviates the sample complexity issue by amortizing task adaptation into simple task inference via a context encoder. Initial attempts, such as FOCAL (Li, Yang, and Luo 2021), employ a distance metric to cluster samples from the same task while separating those from different tasks. Subsequent work has extended these ideas by leveraging generative adversarial networks (e.g., ER-TRL (Nakhaeinezhadfar, Scannell, and Pajarinen 2025)), mutual information (Gao et al. 2024), and contrastive objectives (Yuan and Lu 2022) to enhance task representations in feature space. To obtain disentangled task representations, recent research (Li et al. 2024; Zhou et al. 2024) explores the integration of variational task inference.

Despite these advancements, they still suffer from extrapolation errors when querying OOD actions in TD updates. Most OMRL methods address this issue by leveraging behavioral policy regularization or behavior cloning as the base algorithm to keep training samples in distribution, such as IDAQ (Wang et al. 2023). Another line of research (Wang et al. 2024b) has introduced sequence modeling to tackle the accumulation of extrapolation errors. Nevertheless, these approaches overlook the impact of decoupling dynamics and rewards in the policy space on distributional shifts.

Successor Features

Pioneering works (Dayan 1993) introduced the idea of successor representation, which proposes that the value function can be decomposed into reward weights and state representations that characterize the state transition dynamics. Subsequently, Barreto et al. (2017, 2018) formally defined the concept as Successor Features (SFs), supporting the separation of environmental dynamics from the reward structure for more efficient learning. Considering the task structure in meta-RL, Han and Wu (2022) first proposes applying SFs to decouple transition dynamics and rewards in the representation space, thereby enhancing the accurate identification of task representations. Wang et al. (2024a) extends this disentanglement into the policy space and incorporates task uncertainty into the task inference module, which strengthens the learning of meaningful task representations. Unlike these works, which focus solely on the online meta-RL setting, this paper leverages the decomposition in a more challenging offline setting where we address the extrapolation error caused by feature overgeneralization through flexible value estimation via uncertainty correction.

Problem Formulation

OMRL generally assumes tasks are drawn from a distribution $\mathcal{M}_i = \langle \mathcal{S}, \mathcal{A}, \mathcal{P}_i, \mathcal{R}_i, \gamma \rangle \sim P(\mathcal{M})$, where \mathcal{S} and \mathcal{A} denote the state and action spaces, \mathcal{P} and \mathcal{R} represent the transition dynamics and reward functions, respectively. γ is the discount factor. These tasks, or MDPs, share identical state and action spaces but differ in their transition dynamics and rewards. For each training task \mathcal{M}_i , an offline dataset \mathcal{D}_i is collected using a behavior policy π_{β_i} . During meta-training, the task representation \mathbf{z} is inferred by a context encoder E based on history trajectories τ (referred to

as *context*) sampled from the offline datasets \mathcal{D} . Prohibited from interacting with the environment, the contextual policy $\pi(a|s, \mathbf{z})$ is trained by employing \mathcal{D} to perform downstream tasks. During meta-testing, given a test task $\mathcal{M}_{i'} \sim P(\mathcal{M})$, the agent accesses a limited dataset $\mathcal{D}_{i'}$ to perform task inference and then interacts with the environment to evaluate the expected return of π . The overall objective of OMRL is to maximize the expected discounted cumulative rewards:

$$\begin{aligned} \max_{\pi} \mathcal{R} &= \mathbb{E}_{\mathcal{M} \sim P(\mathcal{M})} [\sum_{t=0}^{\infty} \gamma^t r(s, a)], \\ a &\sim \pi(s, \mathbf{z}), \quad \mathbf{z} \sim E(\tau), \quad s, \tau \sim \mathcal{D}. \end{aligned} \quad (1)$$

Methodology

We first elucidate how to decompose Q value, conditioned on task representations, into features of the transition dynamics and reward weights, enabling shared structures and rapid adaptation in the offline meta-RL, while also introducing the issue of feature overgeneralization. To mitigate extrapolation errors from this overgeneralization, we elaborate on our method for detecting and correcting OOD samples through uncertainty estimation and return feedback, which together enable adaptive feature learning. Finally, we describe how to implement flow-based task inference, which is crucial for deriving appropriate task representations.

Overgeneralization Problem

Given the offline meta-task dataset \mathcal{D} , which contains trajectories $\tau^{\mathcal{M}_i} = \{(s_t, a_t, r_t, s_{t+1})\}_{t=1}^H$ with horizon H from various training tasks drawn from an identical task distribution, that is, $\mathcal{M}_{\text{train}} \sim P(\mathcal{M})$, we leverage a context encoder E parameterized by ω to derive task representations \mathbf{z} as:

$$\mathbf{z} = E_{\omega}(\tau_t^{\mathcal{M}}), \tau_t^{\mathcal{M}} = \{s_{t+c}, a_{t+c}, r_{t+c}, s_{t+c+1}\}_{c=0}^h, \quad (2)$$

where $\tau_t^{\mathcal{M}}$ is a trajectory segment of horizon h sampled starting from a randomly chosen t in the offline dataset \mathcal{D} .

Following Barreto et al. (2017), we assume the reward function can be linearly decomposed as follows:

$$r(s, a, s', \mathbf{z}) = \phi(s, a, s', \mathbf{z})^{\top} W(\mathbf{z}), \quad (3)$$

where $\phi(s, a, s', \mathbf{z}) \in \mathbb{R}^d$ are context-aware features reflecting variations in state transition dynamics within tasks, and $W(\mathbf{z}) \in \mathbb{R}^d$ are context-aware reward weights that mirror changes in the reward function across tasks. The Q value under policy π is rewritten as:

$$\begin{aligned} Q^{\pi}(s, a, \mathbf{z}) &= \mathbb{E}_{\pi} [r_{t+1} + \gamma r_{t+2} + \dots | s_t = s, a_t = a] \\ &= \mathbb{E}_{\pi} [\phi_{t+1}^{\top} W(\mathbf{z}) + \gamma \phi_{t+2}^{\top} W(\mathbf{z}) + \dots | s, a] \\ &= \mathbb{E}_{\pi} [\sum_{n=t}^{\infty} \gamma^{n-t} \phi_{n+1} | s, a]^{\top} W(\mathbf{z}) \\ &= \psi^{\pi}(s, a, \mathbf{z})^{\top} W(\mathbf{z}), \end{aligned} \quad (4)$$

where $\psi^{\pi}(s, a, \mathbf{z})$ represents context-aware SFs, equivalent to the expected discounted cumulative ϕ . This decoupling of dynamics and rewards within the policy space facilitates the investigation of the impact of features on extrapolation errors and promotes the sharing of similar task structures.

In our offline setting with a decomposed Q value, context-aware SFs ψ and reward weights W are updated separately.

The W network and context-aware features ϕ are jointly trained by minimizing the prediction error of rewards:

$$\mathcal{L}_r = \mathbb{E}_{\substack{(s,a,s',r) \sim \mathcal{D} \\ \mathbf{z} \sim E}} \left[\frac{1}{2} (r - \phi_\xi(s, a, s', \mathbf{z})^\top W_\mu(\mathbf{z}))^2 \right]. \quad (5)$$

And ψ_θ is estimated by minimizing the following TD loss:

$$\begin{aligned} \mathcal{L}_\psi(\theta) = & \mathbb{E}_{(s,a,s') \sim \mathcal{D}, \mathbf{z} \sim E, a' \sim \pi} \left[\psi_\theta(s, a, \mathbf{z}) \right. \\ & \left. - \left(\phi_\xi(s, a, s', \mathbf{z}) + \gamma \psi_{\hat{\theta}}(s', a', \mathbf{z}) \right) \right]^2, \quad (6) \end{aligned}$$

where $\psi_{\hat{\theta}}$ is a target context-aware SFs network with soft parameter updates, and ϕ serves as an immediate reward.

However, in offline scenarios with decomposed Q values, as illustrated in Eq. 6, the extrapolation error is primarily caused by $\psi_{\hat{\theta}}(s', a', \mathbf{z})$ when the next action a' taken by the policy π may not be contained in the dataset \mathcal{D} . This can lead to overgeneralization in ψ_θ and ultimately result in a biased Q value. We term this problem as *feature overgeneralization* and observe that overgeneralized ψ generally triggers a severe estimation bias in Q value, as in Door-Close in Fig. 1.

Adaptive Correction of Overgeneralization (ACO)

To accurately identify OOD samples that trigger overgeneralization during the TD update of context-aware SFs ψ_θ , we model the target value $\psi_{\hat{\theta}}$ in Eq. 6 as a Gaussian distribution with mean $\bar{\psi}$ and standard deviation σ_ψ :

$$\Psi(s', a', \mathbf{z}) \triangleq \bar{\psi}(s', a', \mathbf{z}) + \sigma_\psi(s', a', \mathbf{z}). \quad (7)$$

This design adeptly handles the inherent variability and uncertainty of the environment, accommodating more challenging offline scenarios. Moreover, it preserves the multimodality of Q value distributions, promoting stable learning.

We approximate Ψ using a set of D feature atoms $\psi^{(d)}$ ($d = 1, 2, \dots, D$), representing discrete feature values of the ψ_θ network. Specifically, we employ double ψ -learning, with corresponding double feature atoms $\psi^{(d)}$ parameterized by $\hat{\theta}$, to minimize estimation bias. The mean and standard deviation of the feature atoms are calculated as:

$$\bar{\psi}^{(d)} = \frac{1}{2} \left(\psi_{\hat{\theta}_1}^{(d)} + \psi_{\hat{\theta}_2}^{(d)} \right), \quad \sigma^{(d)} = \sqrt{\sum_{e=1}^2 \left(\psi_{\hat{\theta}_e}^{(d)} - \bar{\psi}^{(d)} \right)^2}. \quad (8)$$

The standard deviation $\sigma^{(d)}$ quantifies the *epistemic uncertainty* for detecting OOD samples. By combining $\bar{\psi}^{(d)}$ and $\sigma^{(d)}$, we define the belief distribution $\tilde{\Psi}$:

$$\tilde{\Psi}(s', a', \mathbf{z}) = \bar{\psi}^{(d)}(s', a', \mathbf{z}) + \alpha \sigma^{(d)}(s', a', \mathbf{z}), \quad (9)$$

where α is a dynamically adjusted parameter.

We consider multi-armed bandit strategies (Moskowitz et al. 2021) and associate each bandit arm with a specific α value. Additionally, we introduce feedback on policy performance to update these α values. With U bandit arms, where $\alpha \in \{\alpha_u\}_{u=1}^U$, the adaptive adjustment of the weight for each α_u in the next episode, $w_{g+1}(\alpha_u)$, is formulated as:

$$\begin{aligned} w_{g+1}(\alpha_u) = & w_g(\alpha_u) + \lambda \frac{\mathcal{R}_g - \mathcal{R}_{g-1}}{p_g(\alpha)}, \quad (10) \\ \alpha \sim & p_g(\alpha) \propto \exp(w_g(\alpha_u)), \quad \lambda > 0, \end{aligned}$$

where $p_g(\alpha)$ denotes the exponentially weighted distribution of α , \mathcal{R}_g represents the return from the current episode g , and λ is the update step size. Eq. 10 indicates that changes in the return from one episode to the next, signifying either improvements or degradation in policy performance, trigger updates to the weight of α_u . Conversely, if the policy performance remains unchanged, the weight remains constant.

In offline scenarios, we set $\alpha \leq 0$ to suppress feature overgeneralization and stabilize training. Through iterative training of $w(\alpha_u)$, the probability of $\alpha < 0$ increases when the agent encounters low-return OOD actions, leading to conservative estimates for ψ . Conversely, when no OOD actions are encountered or when the returns from OOD actions are relatively high, the probability of $\alpha = 0$ increases, utilizing the estimated mean as the predicted output. This strategy not only prevents overestimation but also encourages the agent to explore high-return actions outside the dataset, thereby improving the meta-policy and reducing extrapolation error.

We employ belief context-aware SFs $\tilde{\psi}^{(d)}$, sampled from the $\tilde{\Psi}$ distribution as the TD target. To robustly address OOD actions, we leverage the Huber loss (Dabney et al. 2018) for TD errors as the optimization objective for ψ_θ :

$$\mathcal{J}_\psi(\theta_e) = \sum_{d=1}^D \mathcal{L}_{\text{Huber}} \left(\psi_{\theta_e}^{(d)} - \left(\phi_\xi^{(d)} + \gamma \tilde{\psi}^{(d)} \right) \right). \quad (11)$$

The corrected Q value is estimated by calculating the minimal product of ψ and W :

$$\tilde{Q}(s, a, \bar{\mathbf{z}}) = \min_{e=1,2} \sum_{d=1}^D \psi_{\theta_e}^{(d)}(s, a, \bar{\mathbf{z}})^\top W_\mu^{(d)}(\bar{\mathbf{z}}), \quad (12)$$

where $\bar{\mathbf{z}}$ prevents the backpropagation of gradients. From the perspective of distributional RL, these feature atoms construct a support set $\{\psi^{(d)}, 1 \leq d \leq D\}$, with the corresponding probabilities given by $W^{(d)}$. We can then approximate the full distribution of \tilde{Q} values, effectively capturing uncertainties and singular values of training samples.

The loss function for the meta-policy π_φ is structured as

$$\mathcal{L}_\pi = -\mathbb{E}_{s \sim \mathcal{D}} \left[\mathbb{E}_{\tilde{a} \sim \pi_\varphi} \left[\tilde{Q}(s, \tilde{a}, \bar{\mathbf{z}}) \right] - \rho_1 D_{\text{KL}}(\pi_\varphi, \pi_\beta) \right], \quad (13)$$

where D_{KL} denotes the KL divergence between the behavior policy $\pi_\beta(\cdot | s, \mathbf{z})$ and the learned $\pi_\varphi(\cdot | s, \mathbf{z})$, with ρ_1 being a trainable parameter. The policy superiority of FLORA is guaranteed in Theorem 1, with the proof in Appendix A.

Flow-based Task Inference (FTI)

In context-based meta-RL, task representations \mathbf{z} are typically derived by a context encoder that processes recent trajectories: $\mathbf{z} \sim E_\omega(\mathbf{z}|\tau), \tau \sim \mathcal{D}$. As meta-task distributions are often inherently multimodal, a property further amplified in offline settings due to the heterogeneity of π_β for collecting datasets \mathcal{D} , a unimodal Gaussian prior is insufficient to capture this complex distribution $E_\omega(\mathbf{z}|\tau)$. We perform flow-based task inference to explicitly model this distribution by transforming a simple Gaussian distribution through a series of invertible mappings, thereby enhancing

the expressiveness of task representations and better matching the true posterior. We utilize an invertible transformation (Rezende and Mohamed 2015) defined as follows:

$$f_\eta(\mathbf{z}) = \mathbf{z} + \mathbf{u}l(\mathbf{w}^\top \mathbf{z} + b), \quad (14)$$

where \mathbf{w} is the weight vector, b is the bias, \mathbf{u} is a scaling vector, and $l(\cdot)$ is a smooth function. This transformation is part of a series of *planar flows* F_η , which transforms \mathbf{z} drawn from an initial Gaussian-distributed prior $p(\mathbf{z})$ into \mathbf{z}_K :

$$\mathbf{z}_K = F_\eta^K(\mathbf{z}) = f_{\eta_K} \circ \dots \circ f_{\eta_2} \circ f_{\eta_1}(\mathbf{z}), \mathbf{z} \sim p(\mathbf{z}). \quad (15)$$

The final posterior distribution $E_\omega^K(\mathbf{z}|\tau)$, obtained by successively propagating the initial distribution $E_\omega^0(\mathbf{z}|\tau)$ through a chain of K transformations f_{η_k} , is given by:

$$\log E_\omega^K(\mathbf{z}|\tau) = \log E_\omega^0(\mathbf{z}|\tau) - \sum_{k=1}^K \log \left| \det \frac{\partial f_{\eta_k}}{\partial \mathbf{z}_{k-1}} \right|. \quad (16)$$

To extract compact task representations, E_ω leverages variational inference to model the contextual state-action value function $Q(s, a, \mathbf{z})$, with the evidence lower bound $\mathcal{J}_{\text{ELBO}}$ for training E_ω defined as:

$$\mathbb{E}_{\tau \sim \mathcal{D}, \mathbf{z} \sim E_\omega} [\log p(Q) - \rho_2 D_{\text{KL}}(E_\omega^K(\mathbf{z}|\tau) \| p(\mathbf{z}_K))], \quad (17)$$

where ρ_2 is a hyperparameter set to 0.1 in all experiments. More details on $\mathcal{J}_{\text{ELBO}}$ are in Appendix A. We incorporate the training loss of reward \mathcal{L}_r and context-aware successor features \mathcal{J}_ψ to collectively update the context encoder E_ω , facilitating an accurate capture of task changes (Wang et al. 2024a). The overall loss function for training E_ω is:

$$\mathcal{L}_E = \rho_3(\mathcal{L}_r + \mathcal{J}_\psi) - \mathcal{J}_{\text{ELBO}}, \quad (18)$$

where ρ_3 is a hyperparameter set to 0.01 in all experiments. The pseudocode of FLORA is illustrated in Alg. 1.

Theorem 1. (Policy Superiority Guarantee) Consider a meta-task \mathcal{M}_i with an optimal policy π^* whose action-value is $Q_i^{\pi^*}$. Let $Q_i^{\pi_j^*}$ be the action-value of an optimal policy of \mathcal{M}_j when performed on \mathcal{M}_i . Given the set $\{\hat{Q}_i^{\pi_1^*}, \hat{Q}_i^{\pi_2^*}, \dots, \hat{Q}_i^{\pi_L^*}\}$ such that $|\hat{Q}_i^{\pi_j^*} - \hat{Q}_i^{\pi_j^*}| \leq \epsilon$ for all $s \in \mathcal{S}, a \in \mathcal{A}, \mathbf{z} \in \mathcal{Z}$. Define the upper bounds of the distance to the optimal policy for FLORA and the standard Q value as follows: $|Q_i^{\pi^*} - Q_i^{\text{flora}}| \leq \delta_{\text{flora}}$ and $|Q_i^{\pi^*} - Q_i^{\pi_{s-q}}| \leq \delta_{s-q}$, respectively. Then we have $\delta_{\text{flora}} \leq \delta_{s-q}$.

Experiment

Experimental Settings

Setups We evaluated FLORA on two benchmarks: (1) **Meta-World** (Yu et al. 2019) comprises diverse robotic manipulation tasks and is generally considered more challenging due to its broader distribution. We follow the ML1 evaluation protocol, which evaluates the few-shot adaptability with random initial objects and goal positions. (2) **MuJoCo** (Todorov, Erez, and Tassa 2012) focuses solely on parametric diversity. We evaluated two distinct scenarios: a) tasks with changes in reward functions (e.g., goal position for Point-Robot), and b) tasks with changes in transition dynamics (e.g., wind conditions for Point-Robot-Wind).

Offline Data Collection For each environment in Meta-World, we sampled 40 training tasks and 10 testing tasks from the meta-task distribution. For MuJoCo, we sampled 8 training tasks and 2 testing tasks. Following IDAQ, we utilized the Soft Actor-Critic (SAC) algorithm (Haarnoja et al. 2018) to train policies for each task and considered the SAC policy at different training phases as behavior policies. These policies were then rolled out to generate 50 trajectories per benchmark, creating comprehensive offline datasets. To ensure consistency and fairness, all methods were trained on the same dataset in each environment. Moreover, all baseline methods utilize BRAC as the backbone algorithm to address the overestimation of Q value in offline meta-RL. More experimental details are provided in Appendix C.

Algorithm 1: FLORA algorithm

Meta-training

Input: Offline training datasets $\mathcal{D}^{\text{Train}} = \{\tau^{\mathcal{M}_i}\}_{i=1}^I$ of tasks $\{\mathcal{M}_{\text{Train}}^i\}_{i=1}^I \sim P(\mathcal{M})$, context encoder $E_\omega(\mathbf{z}|\tau)$, planar flows F_η , meta-policy $\pi_\varphi(a|s, \mathbf{z})$, context-aware features $\phi_\xi(s, a, s', \mathbf{z})$ and reward weights $W_\mu(\mathbf{z})$, context-aware SFs $\psi_{\theta_e}(s, a, \mathbf{z})$ and $e \in \{1, 2\}$

- 1: Initialize the distribution $p_0(\alpha) \leftarrow \mathcal{U}([-1, 0]^U)$
- 2: **while** not done **do**
- 3: **for** each episode $g = 0, \dots, G - 1$ **do**
- 4: Sample adaptive parameter $\alpha \sim p_g(\alpha)$
- 5: **for** each training step **do**
- 6: **for** each training task $\mathcal{M}_i^{\text{train}}$ **do**
- 7: Sample $\tau^{\mathcal{M}_i} = \{(s_t, a_t, r_t, s_{t+1})\}_{t=1}^H \sim \mathcal{D}^{\text{Train}}$ for training E_ω and π_φ
- 8: Infer task representations $\mathbf{z} \sim E_\omega(\mathbf{z}|\tau^{\mathcal{M}_i})$
- 9: Transform \mathbf{z} via planar flows $\mathbf{z}_K = f_\eta(\mathbf{z})$
- 10: **end for**
- 11: Train $E_\omega: \omega \leftarrow \omega - \lambda_\omega \hat{\nabla}_\omega \mathcal{L}_E$ (Eq. 18)
- 12: Train $W_\mu: \mu \leftarrow \mu - \lambda_\mu \hat{\nabla}_\mu \mathcal{L}_r$ (Eq. 5)
- 13: Train $\psi_{\theta_e}: \theta_e \leftarrow \theta_e - \lambda_\psi \hat{\nabla}_{\theta_e} \mathcal{J}_\psi(\theta_e)$ (Eq. 11)
- 14: Train $\pi_\varphi: \varphi \leftarrow \varphi - \lambda_\varphi \hat{\nabla}_\varphi \mathcal{L}_\pi$ (Eq. 13)
- 15: Train target $\psi_{\hat{\theta}_e}: \hat{\theta}_e \leftarrow \zeta \theta_e + (1 - \zeta) \hat{\theta}_e$
- 16: Train $\phi_\xi: \xi \leftarrow \xi - \lambda_\xi \hat{\nabla}_\xi \mathcal{L}_r$ (Eq. 5)
- 17: Train $F_\eta: \eta \leftarrow \eta - \lambda_\eta \hat{\nabla}_\eta \mathcal{L}_E$ (Eq. 18)
- 18: **end for**
- 19: Update weight of α distribution w_{g+1} using Eq. 10
- 20: **end for**
- 21: **end while**

Meta-testing

Input: Offline testing datasets $\mathcal{D}^{\text{Test}} = \{\tau^{\mathcal{M}_{i'}}\}_{i'=1}^{I'}$ of testing tasks $\{\mathcal{M}_{\text{Test}}^{i'}\}_{i'=1}^{I'} \sim P(\mathcal{M})$, learned meta-policy π_φ and context encoder E_ω

- 1: **for** each testing task $\mathcal{M}_{i'}^{\text{test}}$ **do**
- 2: **for** $t=0, \dots, T-1$ **do**
- 3: Sample trajectories $\tau^{\mathcal{M}_{i'}} \sim \mathcal{D}^{\text{Test}}$
- 4: Infer task representations $\mathbf{z} \sim E_\omega(\mathbf{z}|\tau^{\mathcal{M}_{i'}})$
- 5: Roll out policy $\pi_\varphi(a|s, \mathbf{z})$ for evaluation
- 6: **end for**
- 7: **end for**

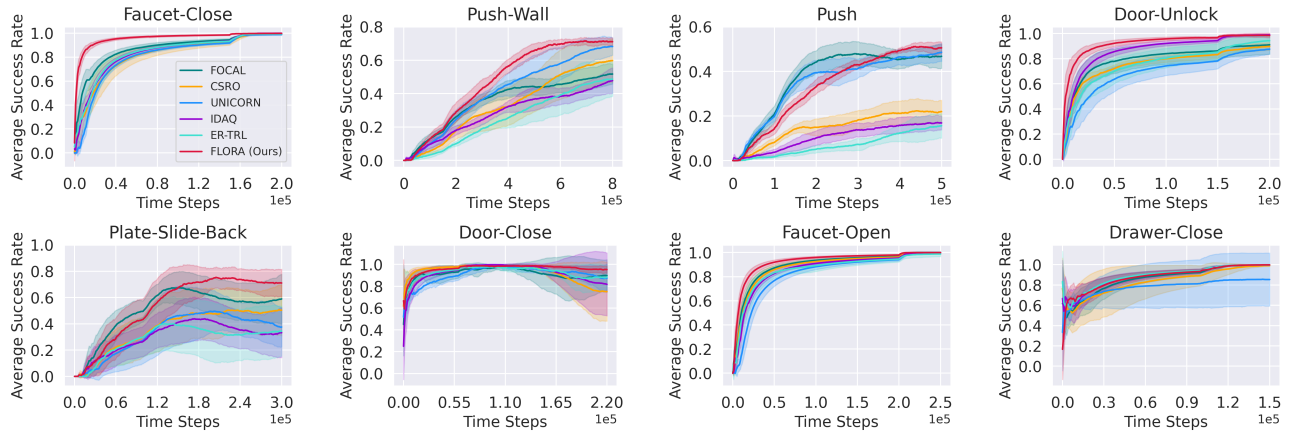


Figure 2: Testing average performance on 8 ML1 environments over 6 random seeds.

	Faucet-Close	Push-Wall	Push	Door-Unlock	Plate-Slide-Back	Door-Close	Faucet-Open	Drawer-Close
FOCAL	99.67±0.75	52.73± 13.55	45.63±13.29	89.96±12.61	61.75±37.69	91.56±15.74	99.96±0.09	99.80±0.39
IDAQ	99.94±0.12	50.90±18.10	18.00± 10.33	98.15± 3.01	36.67±38.97	80.22±36.88	99.83±0.36	99.90±0.22
CSRO	99.89±0.25	58.27±20.54	21.03±10.55	87.61±10.30	51.83±36.24	74.22±39.50	99.92±0.18	99.73±0.60
UNICORN	99.56±0.90	69.47±13.82	49.57±11.45	84.44±13.17	37.33±36.12	83.06±28.34	99.99±0.02	85.77±31.74
ER-TRL	99.61±0.78	50.80±19.40	16.80±11.15	90.74±11.45	38.92±36.13	92.33±14.19	99.51±1.10	99.77±0.52
FLORA	100.00±0.00	71.03±17.39	50.07±13.11	98.27±3.44	73.25±31.87	95.00±9.89	100.00±0.00	99.93±0.15

Table 1: Converged average test success rate ± standard error (%) across 6 random seeds on Meta-World.




Figure 3: Testing average performance of FLORA and baselines on MuJoCo over 6 random seeds.

Baselines 1) **FOCAL** devised an inverse power distance to cluster similar tasks; 2) **IDAQ** tackled distribution shift via return-based uncertainty; 3) **CSRO** used mutual information to reduce behavior policy bias; 4) **UNICORN** leveraged a decoder to reconstruct task representations; 5) **Meta-DT** utilized sequence modeling and a context-aware world model to achieve efficient generalization; 6) **ER-TRL** employed a GAN to maximize behavior policy entropy.

Performance Comparison

Meta-World Results. Fig. 2 and Table 1 summarize the testing performance and convergent results, respectively. FLORA consistently achieves the highest final performance in all environments and demonstrates superior adaptation efficiency. Notably, in high-quality datasets, such as Door-Unlock, IDAQ shows relatively quick convergence but struggles in more complex Push, where offline datasets contain a higher proportion of suboptimal trajectories. IDAQ

misclassifies more suboptimal data as in-distribution context, leading to significant performance degradation. In contrast, FLORA maintains rapid adaptation and achieves exceptional final performance with relatively minimal variance regardless of dataset quality. CSRO fails to adapt to new tasks in most environments, possibly because the minimization of mutual information results in overly sparse and independent task representations. In challenging environments such as Push, the relationships of different tasks are not effectively captured by CSRO, disrupting the shared structure across tasks and making it difficult for the meta-policy to generalize to new tasks. UNICORN, which decomposes task structures in the representation space, achieves faster policy convergence and higher final performance on Push and Push-Wall. However, UNICORN’s decoder may introduce noise when reconstructing task structures to interfere with the encoder’s task inference. Moreover, the trade-off between reconstruction loss and the regularization term can lead to intrinsic information loss, making UNICORN fall behind other baselines. ER-TRL faces a similar dilemma: although entropy maximization promotes action diversity, the inherent training instability of the generator may lead to mode collapse. Surprisingly, FOCAL exceeds most baselines in adaptability and exhibits impressive final performance in Push and Door-Close. However, the similarity between tasks may be difficult to define accurately across the broader distributed Meta-World, making FOCAL’s distance metric unable to effectively distinguish diverse tasks. Conversely, FLORA consistently outperforms other baselines regarding data efficiency and final testing success rate.

In the Door-Close and Plate-Slide-Back, all baselines exhibit varying performance degradation. They fail to detect and address extrapolation errors promptly, causing the errors to accumulate over time and ultimately resulting in policy collapse. In comparison, FLORA accurately approximates complex task distributions through a chain of invertible transformations and effectively mitigates these errors with adaptive feature learning, consistently achieving rapid adaptation and superior asymptotic performance.

MuJoCo Results. Fig. 3 demonstrates that FLORA has exceptional adaptation efficiency and converged performance with minimal variance. By decoupling transition dynamics and reward functions in both the representation and policy spaces, FLORA accurately captures and shares similar structures of different tasks, thus facilitating fast and stable adaptation to unseen tasks. Specifically, in Point-Robot-Wind, FLORA mitigates the overestimation of the decomposed features that represent dynamics through adaptive correction, effectively reducing extrapolation errors. In contrast, other baselines ignore the impact of meta-task structures on extrapolation errors, resulting in severe fluctuations and increased variance during policy learning.




Figure 4: Ablation study.

Ablation Study. When meta-policies are trained with suboptimal offline datasets in a broader distribution of tasks, context-aware SFs can easily overgeneralize in the presence of OOD actions. This leads to an overestimation of Q values and a subsequent decline in policy performance, as observed in the performance of the model without FTI and ACO in Fig. 4. Specifically, the ACO module plays a critical role in the prompt detection of OOD samples by estimating the uncertainty of training samples. It conservatively estimates the rewards of low-reward OOD samples through a reward feedback mechanism, effectively suppressing feature overgeneralization while preserving the performance gains from policy space decoupling for both in-distribution samples and high-return OOD samples. FTI models more accurate task representation distributions in challenging tasks, thereby enhancing the compactness of task representations. As shown in Fig. 4, ACO effectively mitigates policy decline and prevents policy collapse. Meanwhile, it accelerates policy convergence and improves overall performance.

Rapid adaptation in sparse tasks. We also validate the adaptation efficiency of FLORA on sparse point-robot, where the agent is trained to navigate to training goals (light blue circles in Fig. 5) and tested on a distinct unseen goal (dark blue). A reward is given only when the agent is within a certain radius of the goal. As in Fig. 5, FLORA navigates to the area around the test goal more quickly. This indicates




Figure 5: Visualization of sparse 2D navigation.

that the decomposed representation and policy space facilitate learning common structures in the sparse setting.

Evaluation on Non-Stationary Tasks. Following the setup of (Bing et al. 2023), we evaluated FLORA and the baselines on more challenging tasks, where the agent learns to adapt to time-evolving task structures (with reward functions that vary over time). In our setup, the task period is sampled from a Gaussian distribution $N(250, 10)$ for each 1000 steps per epoch. As illustrated in Fig. 6 (left), our FLORA consistently achieves the highest success rate with minimal variance and maintains excellent adaptability.




Figure 6: Evaluation on non-stationary tasks (left) and with different Q value estimation in Point-Robot-Wind (right).

Comparison of Different Q Value Estimations. We investigated the influence of Q value learning across 6 random seeds in Point-Robot-Wind, replacing our backbone (decomposed Q network) with a vanilla (standard and undecomposed) and a distributional Q (Dist. Q) network to measure uncertainty, respectively. As in Fig. 6 (right), Dist. Q alleviates the accumulation of errors by fitting the distribution of Q values more effectively than vanilla Q . Furthermore, Dec. Q enhances task distinction and uncertainty identification, achieving the highest return with minimal variance.

Conclusion

OMRL primarily confronts two major challenges: accurately inferring task distributions and alleviating extrapolation errors. This paper proposes FLORA to effectively address these issues by: 1) flexibly and efficiently approximating complex task distributions through flow-based task inference to derive accurate task representations, and 2) precisely identifying OOD samples via uncertainty estimation of decoupled features and mitigating feature overgeneralization through adaptive correction. The experimental results on challenging Meta-World and MuJoCo demonstrate that FLORA outperforms other baselines in accurately discerning diverse tasks and rapidly adapting to new tasks.

Acknowledgments

This work was partially supported by the NSFC under Grants 92270125 and 62276024; by the Fundamental Research Funds for the Central Universities, JLU, under Grant 93K172025K01; and by the Fundamental Research Funds for the Central Universities under Grant 2025CX01010.

References

- Barreto, A.; Borsa, D.; Quan, J.; Schaul, T.; Silver, D.; Hessel, M.; Mankowitz, D.; Zidek, A.; and Munos, R. 2018. Transfer in deep reinforcement learning using successor features and generalised policy improvement. In *International Conference on Machine Learning*, 501–510. PMLR.
- Barreto, A.; Dabney, W.; Munos, R.; Hunt, J. J.; Schaul, T.; van Hasselt, H. P.; and Silver, D. 2017. Successor features for transfer in reinforcement learning. *Advances in neural information processing systems*, 30.
- Bing, Z.; Lerch, D.; Huang, K.; and Knoll, A. 2023. Meta-Reinforcement Learning in Non-Stationary and Dynamic Environments. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 45(03): 3476–3491.
- Dabney, W.; Rowland, M.; Bellemare, M.; and Munos, R. 2018. Distributional reinforcement learning with quantile regression. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Dayan, P. 1993. Improving generalization for temporal difference learning: The successor representation. *Neural computation*, 5(4): 613–624.
- Gao, Y.; Zhang, R.; Guo, J.; Wu, F.; Yi, Q.; Peng, S.; Lan, S.; Chen, R.; Du, Z.; Hu, X.; et al. 2024. Context shift reduction for offline meta-reinforcement learning. *Advances in Neural Information Processing Systems*, 36.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, 1861–1870. PMLR.
- Han, X.; and Wu, F. 2022. Meta Reinforcement Learning with Successor Feature Based Context. *arXiv preprint arXiv:2207.14723*.
- Li, L.; Yang, R.; and Luo, D. 2021. FOCAL: Efficient Fully-Offline Meta-Reinforcement Learning via Distance Metric Learning and Behavior Regularization. In *International Conference on Learning Representations*.
- Li, L.; Zhang, H.; Zhang, X.; Zhu, S.; Zhao, J.; and Heng, P.-A. 2024. Towards an Information Theoretic Framework of Context-Based Offline Meta-Reinforcement Learning. *arXiv preprint arXiv:2402.02429*.
- Mao, Y.; Wang, Q.; Qu, Y.; Jiang, Y.; and Ji, X. 2024. Doubly mild generalization for offline reinforcement learning. *arXiv preprint arXiv:2411.07934*.
- Mitchell, E.; Rafailov, R.; Peng, X. B.; Levine, S.; and Finn, C. 2021. Offline meta-reinforcement learning with advantage weighting. In *International Conference on Machine Learning*, 7780–7791. PMLR.
- Moskovitz, T.; Parker-Holder, J.; Pacchiano, A.; Arbel, M.; and Jordan, M. 2021. Tactical optimism and pessimism for deep reinforcement learning. *Advances in Neural Information Processing Systems*, 34: 12849–12863.
- Nakhaeinezhadfar, M.; Scannell, A.; and Pajarinen, J. 2025. Entropy Regularized Task Representation Learning for Offline Meta-Reinforcement Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 19616–19623.
- Ni, F.; Hao, J.; Mu, Y.; Yuan, Y.; Zheng, Y.; Wang, B.; and Liang, Z. 2023. Metadiffuser: Diffusion model as conditional planner for offline meta-rl. In *International Conference on Machine Learning*, 26087–26105. PMLR.
- Rezende, D.; and Mohamed, S. 2015. Variational inference with normalizing flows. In *International conference on machine learning*, 1530–1538. PMLR.
- Todorov, E.; Erez, T.; and Tassa, Y. 2012. MuJoCo: A physics engine for model-based control. *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 5026–5033.
- Wang, J.; Zhang, J.; Jiang, H.; Zhang, J.; Wang, L.; and Zhang, C. 2023. Offline meta reinforcement learning with in-distribution online adaptation. In *International Conference on Machine Learning*, 36626–36669. PMLR.
- Wang, M.; Li, X.; Zhang, L.; and Wang, M. 2024a. MetaCARD: Meta-Reinforcement Learning with Task Uncertainty Feedback via Decoupled Context-Aware Reward and Dynamics Components. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 15555–15562.
- Wang, Z.; Zhang, L.; Wu, W.; Zhu, Y.; Zhao, D.; and Chen, C. 2024b. Meta-DT: Offline Meta-RL as Conditional Sequence Modeling with World Model Disentanglement. *arXiv preprint arXiv:2410.11448*.
- Yu, T.; Quillen, D.; He, Z.; Julian, R. C.; Hausman, K.; Finn, C.; and Levine, S. 2019. Meta-World: A Benchmark and Evaluation for Multi-Task and Meta Reinforcement Learning. *ArXiv*, abs/1910.10897.
- Yuan, H.; and Lu, Z. 2022. Robust task representations for offline meta-reinforcement learning via contrastive learning. In *International Conference on Machine Learning*, 25747–25759. PMLR.
- Zhou, R.; Gao, C.-X.; Zhang, Z.; and Yu, Y. 2024. Generalizable Task Representation Learning for Offline Meta-Reinforcement Learning with Data Limitations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 17132–17140.