EFFICIENT FULLY-OFFLINE META-REINFORCEMENT LEARNING VIA DISTANCE METRIC LEARNING AND BEHAVIOR REGULARIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

We study the offline meta-reinforcement learning (OMRL) problem, a paradigm which enables reinforcement learning (RL) algorithms to quickly adapt to unseen tasks without any interactions with the environments, making RL truly practical in many real-world applications. This problem is still not fully understood, for which two major challenges need to be addressed. First, offline RL usually suffers from bootstrapping errors of out-of-distribution state-actions which leads to divergence of value functions. Second, meta-RL requires efficient and robust task inference learned jointly with control policy. In this work, we enforce behavior regularization on learned policy as a general approach to offline RL, combined with a deterministic context encoder for efficient task inference. We propose a novel negative-power distance metric on bounded context embedding space, whose gradients propagation is detached from the Bellman backup. We provide analysis and insight showing that some simple design choices can yield substantial improvements over recent approaches involving meta-RL and distance metric learning. To the best of our knowledge, our method is the first model-free and end-to-end OMRL algorithm, which is computationally efficient and demonstrated to outperform prior algorithms on several meta-RL benchmarks.[1]

1 INTRODUCTION

Applications of reinforcement learning (RL) in real-world problems have been proven successful in many domains such as games (Silver et al., 2017; Vinyals et al., 2019; Ye et al., 2020) and robot control (Johannink et al., 2019). However, the implementations so far usually rely on interactions with either real or simulated environments. In other areas like healthcare (Gottesman et al., 2019), autonomous driving (Shalev-Shwartz et al., 2016) and controlled-environment agriculture (Binas et al., 2019) where RL shows promise conceptually or in theory, exploration in real environments is evidently risky, and building a high-fidelity simulator can be costly. Therefore a key step towards more practical RL algorithms is the ability to learn from fixed data (batch). Such paradigm, termed "offline RL" or "batch RL", would enable better generalization by incorporating diverse prior experience. Moreover, by leveraging and reusing previously collected data, off-policy algorithms such as SAC (Haarnoja et al., 2018) has been shown to achieve far better sample efficiency than on-policy methods. The same applies to offline RL algorithms since they are by nature off-policy.

The aforementioned design principles motivated a surge of recent works on offline/batch RL (Fujimoto et al., 2019; Kumar et al., 2019; Wu et al., 2019; Siegel et al., 2020). These papers propose remedies by regularizing the trained policy to stay close to the training dataset of logged transitions, in order to mitigate the effect of bootstrapping error (Kumar et al., 2019), where errors in evaluation of out-of-distribution state-action pairs are never corrected and hence easily diverge due to inability to collect new data samples for feedback. There exist claims that offline RL can be implemented successfully without explicit correction for distribution mismatch given sufficiently large and diverse training data (Agarwal et al., 2020). However, we find such assumption unrealistic in many

[1]Source code and data will be released at https://github.com/FOCAL-ICLR/FOCAL-ICLR/ and https://anonymous.4open.science/r/e657391f-0364-4670-a366-fb236d83f602/ respectively.

practices as well as in our experiments. In this paper, to tackle the out-of-distribution problem in offline RL in general, we follow the paradigm of behavior regularization in Wu et al. (2019).

For design of practical RL, besides the ability to learn without exploration, it's also ideal to have an algorithm that can generalize to a variety of scenarios. To solve real-world challenges in multi-task setting, such as treating different diseases, driving under various road conditions or growing diverse crops in autonomous greenhouses, a practical/data-efficient agent is expected to quickly transfer and adapt to unseen tasks, especially when the tasks share common structures. Meta-learning methods (Vilalta & Drissi, 2002; Thrun & Pratt, 2012; Finn et al., 2017) address this problem by learning an inductive bias from experience collected across a distribution of tasks, which can be naturally extended to reinforcement learning context. Under the umbrella of this so-called meta-RL, almost all current methods require on-policy data during either both meta-training and testing (Finn et al., 2017; Wang et al., 2016; Duan et al., 2016) or at least online exploration for task inference (Rakelly et al., 2019). An efficient and robust method which incorporates both fully-offline learning and meta-learning in RL, despite few attempts (Li et al., 2019b; Dorfman & Tamar, 2020), has not been fully developed and validated.

In this paper, under the first principle of maximizing practicality of RL algorithm, we propose an efficient method that integrates task inference with RL algorithms in a fully-offline fashion. Our fully-offline context-based actor-critic meta-RL algorithm, or FOCAL, achieves excellent sample efficiency and fast adaptation with only limited logged experience, on a range of deterministic continuous control meta-environments. The primary contribution of this work is designing the first end-to-end and model-free offline meta-RL algorithm which is computationally efficient and effective without any prior knowledge of task identity or reward/transition functions. To achieve efficient task inference, we propose an inverse-power loss for effective learning and clustering of latent variables, in analogy with coulomb potential in electromagnetism, which is also unseen in previous work. Theoretical analyses and insights are provided to shed light on the specific design choices customized for the OMRL problem.

2 Related Work

Meta-RL Our work FOCAL builds upon the meta-learning (Vilalta & Drissi, 2002; Thrun & Pratt, 2012) framework in the context of reinforcement learning. Among all paradigms of meta-RL, this paper is most related to context-based and metric-based approaches. Context-based meta-RL employs models with memory such as recurrent (Duan et al., 2016; Wang et al., 2016; Fakoor et al., 2019), recursive (Mishra et al., 2017) or probabilistic (Rakelly et al., 2019) structures to achieve fast adaptation by aggregating experience into a latent representation on which the policy is conditioned. The design of recent methods usually leverages the temporal or Markovian properties of RL problems.

Metric-based meta-RL focuses on learning effective task representations to facilitate task inference and conditioned control policies, by employing techniques such as distance metric learning (Yang & Jin, 2006). Koch et al. (2015) proposed the first metric-based meta-algorithm for few-shot learning, in which a Siamese network (Chopra et al., 2005) is trained with triplet loss to compare the similarity between a query and supports in the embedding space. Many metric-based meta-RL algorithms extend these works. (Snell et al., 2017; Sung et al., 2018; Li et al., 2019a).

Among all aforementioned meta-learning approaches, this paper is most related to context-based PEARL algorithm (Rakelly et al., 2019) and metric-based prototypical networks (Snell et al., 2017). PEARL achieves SOTA performance for off-policy meta-RL by introducing a probabilistic permutation-invariant context encoder, along with a design which disentangles task inference and control by different sampling strategies. However, it requires exploration during meta-testing. The prototypical networks employ similar design of context encoder as well as an Euclidean distance metric on deterministic embedding space, but tackles meta-learning of classification tasks with squared distance loss as opposed to the inverse-power loss in FOCAL to tackle the more complex OMRL problem.

Offline/Batch RL To address the bootstrapping error (Kumar et al., 2019) problem of offline RL, this paper adopts behavior regularization directly from Wu et al. (2019), which provides a relatively unified framework of several recent offline or off-policy RL methods (Haarnoja et al., 2018; Fuji-

moto et al., 2019; Kumar et al., 2019). It incorporates a divergence function between distributions over actions in the actor-critic objectives. As with SAC (Haarnoja et al., 2018), one limitation of the algorithm is its sensitivity to reward scale and regularization strength. In our experiments, we indeed observed wide spread of optimal hyper-parameters across different meta-RL environments, shown in Table 3.

Offline meta-RL To the best of our knowledge, despite attracting more and more attention, the offline meta-RL problem is still under-developed. We are aware of a few papers that tackle the same problem from different angles (Li et al., 2019b; Dorfman & Tamar, 2020). Li et al. (2019b) focuses on a specific scenario where biased datasets cause the task inference module to spuriously correlate only state-action pairs to the task identity, ignoring the reward information. This problem occurs when datasets of different tasks do not have significant overlap in their state-action visitation frequencies, and is exacerbated by sparse rewards. Even though we observed similar effects during our experiment, in this paper, we do not restrict our algorithm design to such assumptions and instead try to robustify the algorithm in a more general setting, showing experimental results on various training and testing datasets distributions (Fig 7). Another key difference is that our method is end-to-end and model free, whereas Li et al. (2019b) requires training of offline BCQ (Fujimoto et al., 2019) and reward function models for each task, which are computationally demanding. They also adopted the original PEARL strategy of probabilistic context encoder and couples learning of task inference and control by Bellman gradients, for which we chose differently for provably better performance and efficiency.

Dorfman & Tamar (2020) on the other hand, completely circumvents the out-of-distribution problem of offline RL by assuming sufficient data samples. They also formulate the OMRL as a Bayesian RL (Ghavamzadeh et al., 2016) problem and employs a probabilistic approach. Therefore we consider their methodology tangential to ours.

3 PRELIMINARIES

3.1 NOTATIONS AND PROBLEM STATEMENT

We consider fully-observed Markov Decision Process (MDP) (Puterman, 2014) in deterministic environments such as MuJoCo (Todorov et al., 2012). An MDP can be modeled as $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, R, \rho_0, \gamma)$ with state space $\mathcal{S}$, action space $\mathcal{A}$, transition function $P(s'|s,a)$, bounded reward function $R(s,a)$, initial state distribution $\rho_0(s)$ and discount factor $\gamma \in (0,1)$. The goal is to find a policy $\pi(a|s)$ to maximize the cumulative discounted reward starting from any state. We introduce the notion of multi-step state marginal of policy π as $\mu_\pi^t(s)$, which denotes the distribution over state space after rolling out π for t steps starting from state s. The notation $R_\pi(s)$ denotes the expected reward at state s when following policy π: $R_\pi(s) = \mathbb{E}_{a\sim\pi}[R(s,a)]$. The state-value function (a.k.a. value function) and action-value function (a.k.a Q-function) are therefore

$$V_\pi(s) = \sum_{t=0}^{\infty} \gamma^t \mathbb{E}_{s_t \sim \mu_\pi^t(s)}[R(s_t)] \tag{1}$$

$$Q_\pi(s,a) = R(s,a) + \gamma \mathbb{E}_{s' \sim P(s'|s,a)}[V_\pi(s')] \tag{2}$$

Q-learning algorithms are implemented by iterating the Bellman optimality operator $\mathcal{B}$, defined as:

$$(\mathcal{B}\hat{Q})(s,a) := R(s,a) + \gamma \mathbb{E}_{P(s'|s,a)}[\max_{a'} \hat{Q}(s',a')] \tag{3}$$

when the state space is large/continuous, $\hat{Q}$ is used as a hypothesis from the set of function approximators (e.g. neural networks).

In the offline context of this work, given a distribution of tasks $p(\mathcal{T})$ where every task is an MDP, we study off-policy meta-learning from collections of static datasets of transitions $\mathcal{D}_i = \{(s_{i,t}, a_{i,t}, s'_{i,t}, r_{i,t}) | t = 1, ..., N\}$ generated by a set of behavior policies $\{\beta_i(a|s)\}$ associated with each task index i. A key underlying assumption of meta-learning is that the tasks share some common structures. Accordingly, by definition of MDP, in this paper we restrict our attention on tasks with shared state and action space, but differ in transition and reward functions.

We define the meta-optimization objective as

$$\mathcal{L}(\theta) = \mathbb{E}_{\mathcal{T}_i \sim p(\mathcal{T})} \left[\mathcal{L}_{\mathcal{T}_i}(\theta) \right] \tag{4}$$

where $\mathcal{L}_{\mathcal{T}_i}(\theta)$ is the objective evaluated on transition samples drawn from task $\mathcal{T}_i$. A common choice of $p(\mathcal{T})$ is the uniform distribution on the set of given tasks $\{\mathcal{T}_i | i = 1, ..., n\}$. In this case, the meta-training procedure turns into minimizing the average losses across all training tasks

$$\hat{\theta}_{\text{meta}} = \arg\min_{\theta} \frac{1}{n} \sum_{k=1}^{n} \mathbb{E}\left[\mathcal{L}_k(\theta)\right] \tag{5}$$

3.2 Behavior Regularized Actor Critic

Similar to SAC, to constrain the out-of-distribution error in offline RL, for each individual task $\mathcal{T}_i$, behavior regularization (Wu et al., 2019) introduces a divergence measure between the learner π_θ and the behavior policy π_b in both value and target Q-functions. For simplicity, we ignore task index in this section:

$$V_\pi^D(s) = \sum_{t=0}^{\infty} \gamma^t \mathbb{E}_{s_t \sim \mu_\pi^t(s)} \left[R_\pi(s) - \alpha D(\pi_\theta(\cdot|s_t), \pi_b(\cdot|s_t))\right] \tag{6}$$

$$\bar{Q}_\psi^D(s, a) = \bar{Q}_\psi(s, a) - \gamma \alpha \hat{D}(\pi_\theta(\cdot|s), \pi_b(\cdot|s)) \tag{7}$$

where $\bar{Q}$ denotes a target Q-function without gradients and $\hat{D}$ denotes a sample-based estimate of the divergence function D. In actor-critic framework, the loss functions of Q-value and policy learning are given by, respectively,

$$\mathcal{L}_{critic} = \mathbb{E}_{\substack{(s,a,r,s') \sim \mathcal{D} \\ a' \sim \pi_\theta(\cdot|s')}} \left[\left(r + \gamma \bar{Q}_\psi^D(s', a') - Q_\psi(s, a)\right)^2\right] \tag{8}$$

$$\mathcal{L}_{actor} = -\mathbb{E}_{(s,a,r,s') \sim \mathcal{D}} \left[\mathbb{E}_{a'' \sim \pi_\theta(\cdot|s)}[Q_\psi(s, a'')] - \alpha \hat{D}\right] \tag{9}$$

3.3 Context-Based meta-RL

Context-based meta-RL algorithms aggregate context information, typically in form of task-specific transitions, into a latent space $\mathcal{Z}$. Adaptation at test time can be viewed as a special case of RL on partially-observed MDP (Kaelbling et al., 1998) in which a latent representation z as the unobserved part of the state needs to be inferred. Once given complete information of z and s combined as the full state, the reinforcement learning of policy $\pi_\theta(s, z)$ becomes RL on regular MDP, and properties of regular RL such as the existence of optimal policy and value functions hold naturally. We therefore formulate the context-based meta-RL problem as solving a *task-augmented MDP* (TA-MDP). The formal definitions are provided in Appendix B.

4 Method

Based on our formulation of context-based meta-RL problem, FOCAL first learns an effective representation of meta-training tasks on latent space $\mathcal{Z}$, then solves the offline RL problem on TA-MDP with behavior regularized actor critic method. We illustrate our training procedure in Figure 1 and describe the detailed algorithm in Appendix A. Our method consists of three key design choices: deterministic context encoder, distance metric learning on latent space as well as decoupled training of task inference and control.

4.1 Deterministic Context Encoder

Similar to Rakelly et al. (2019), we introduce an inference network $q_\phi(z|c)$, parameterized by ϕ, to encode task identity from context c. In terms of the context encoder design, recent meta-RL methods either employ recurrent neural networks (Duan et al., 2016; Wang et al., 2016) to capture the temporal property of MDP, or use probabilistic models (Rakelly et al., 2019) for reasoning of uncertainties. These design choices are proven effective in on-policy or partially-offline off-policy algorithms. However, since our approach aims to address the fully-offline meta-RL problem, we argue that a deterministic context encoder works better in this scenario, given a few assumptions. First, we consider only *deterministic MDP* in this paper, where the transition function is a Dirac

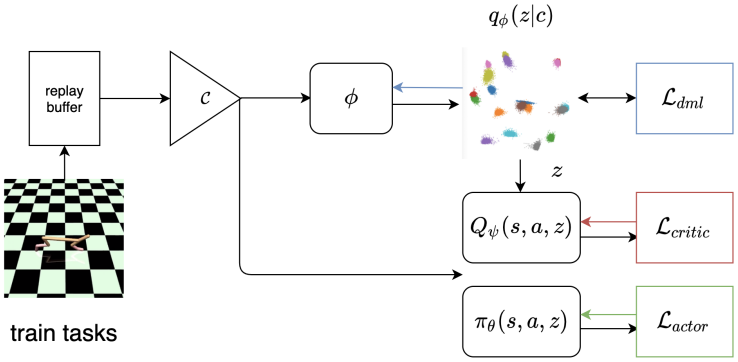

Figure 1: **Meta-training procedure.** The inference network q_ϕ uses context data to determine the latent context variable z, which conditions the actor and critic, and is optimized by the DML objective. The actor and critic objectives are optimized in parallel without any mixing of gradients.

delta distribution. We assume that all meta-learning tasks in this paper are deterministic MDPs, which is satisfied by common RL benchmarks such as MuJoCo. The formal definitions are detailed in Appendix B. Second, we assume all tasks share the same state and action space, while each is characterized by a unique set of transition and reward functions. Mathematically, this means there exists an injective function $f : \mathcal{T} \to \mathcal{P} \times \mathcal{R}$, where $\mathcal{P}$ and $\mathcal{R}$ are functional spaces of transition probability $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \to \{0, 1\}$ and bounded reward $R : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$ respectively. A stronger condition of this injective property is that for every task $\mathcal{T} \sim p(\mathcal{T})$ and any state-action pair (s, a), the corresponding transition and reward are unique across all tasks, which brings the *Task-Transition Correspondence* assumption:

Assumption 1 (Task-Transition Correspondence) *We consider meta-RL with a task distribution $p(\mathcal{T})$ to satisfy task-transition correspondence if and only if* $\forall \mathcal{T}_1, \mathcal{T}_2 \sim p(\mathcal{T})$, $(s, a) \in \mathcal{S} \times \mathcal{A}$:

$$P_1(\cdot|s, a) = P_2(\cdot|s, a), R_1(s, a) = R_2(s, a) \iff \mathcal{T}_1 = \mathcal{T}_2 \tag{10}$$

Under the deterministic MDP assumption, the transition probability function $P(\cdot|s, a)$ is associated with the transition map $t : \mathcal{S} \times \mathcal{A} \to \mathcal{S}$ (Definition 3). The task-transition correspondence suggests that, given the action-state pair (s, a) and task $\mathcal{T}$, there exists a unique transition-reward pair (s', r). Based on these assumptions, one can define a task-specific map $f_\mathcal{T} : \mathcal{S} \times \mathcal{A} \to \mathcal{S} \times \mathbb{R}$ on the set of transitions $\mathcal{D}_\mathcal{T}$:

$$f_\mathcal{T}(s_t, a_t) = (s'_t, r_t), \quad \forall \mathcal{T} \sim p(\mathcal{T}), (s_t, a_t, s'_t, r_t) \in \mathcal{D}_\mathcal{T} \tag{11}$$

Recall that all tasks defined in this paper share the same state-action space, hence $\{f_\mathcal{T}|\mathcal{T} \sim p(\mathcal{T})\}$ forms a function family defined on the transition space $\mathcal{S} \times \mathcal{A} \times \mathcal{S} \times \mathbb{R}$, which is also by definition the context space $\mathcal{C}$. This lends a new interpretation that as a task inference module, the context encoder $q_\phi(z|c)$ enforces an embedding of the task-specific map $f_\mathcal{T}$ on the latent space $\mathcal{Z}$, i.e. $q_\phi : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \times \mathbb{R} \to \mathcal{Z}$. Following Assumption 1, every transition $\{s_i, a_i, s'_i, r_i\}$ corresponds to a unique task $\mathcal{T}_i$, which means in principle, task identity can be inferred from any single transition tuple. This implies the context encoder should be permutation-invariant and deterministic, since the embedding of context does not depend on the order of the transitions nor involve any uncertainty. This observation is crucial since it provides theoretical basis for few-shot learning (Snell et al., 2017; Sung et al., 2018) in our settings. In particular, when learning in a fully-offline fashion, any meta-RL algorithm at test-time cannot perform adaptation by exploration. The theoretical guarantee that a few randomly-chosen transitions can enable effective task inference ensures that FOCAL is feasible and efficient.

4.2 Distance Metric Learning (DML) of Latent Variables

In light of our analysis on the context encoder design, the goal of task inference is to learn a robust and effective representation of context for better discrimination of task identities. Unlike PEARL, which requires Bellman gradients to train the inference network, our main insight is to disentangle

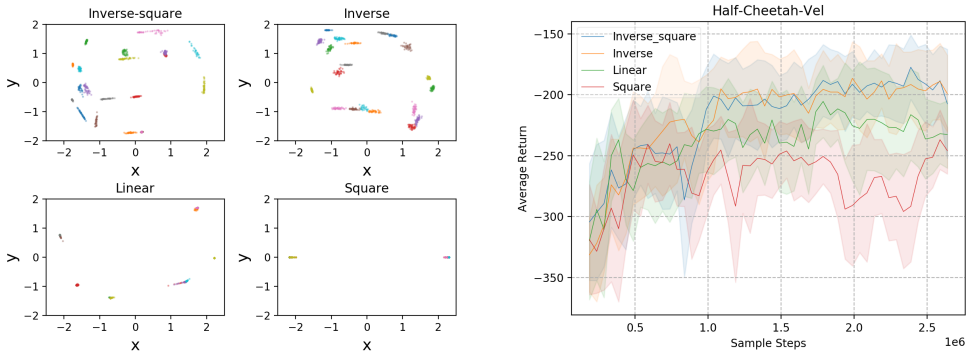

(a) context embedding

(b) test-task performance

Figure 2: (a) A 2D projection of the 5D embedding vectors of samples drawn from 20 randomized tasks on Half-Cheetah-Vel by principal component analysis (PCA). Trials cover four DML losses of different powers. Data points are color-coded according to task identity. (b) FOCAL trained with inverse-power losses also outperform the linear and square distance losses.

the learning of context encoder from the learning of control policy. As explained in previous reasoning about the deterministic encoder, the latent variable is a representation of the task properties involving only transition probability and reward, which in principle should be completely captured by the transition datasets. Given continuous neural networks as function approximators, the learned value functions conditioned on latent variable z cannot account for the difference between tasks if the corresponding embedding vectors are too close (Appendix C). Therefore for implementation, we formulate the latent variable learning problem as obtaining the map $q_\phi : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \times \mathbb{R} \to \mathcal{Z}$ of transition data $\mathcal{D}_i = \{(s_{i,t}, a_{i,t}, s'_{i,t}, r_{i,t}) | t = 1, ..., N\}$ that clusters similar data points (same task) while pushes away dissimilar data points (different tasks) on the embedding space $\mathcal{Z}$, which is essentially distance metric learning (DML) (Sohn, 2016). A common loss function in DML is contrasitive loss (Chopra et al., 2005; Hadsell et al., 2006). Given input data $x_i, x_j \in \mathcal{X}$ and label $y \in \{1, ..., L\}$, it is written as

$$\mathcal{L}^m_{cont}(x_i, x_j; q) = \mathbf{1}\{y_i = y_j\}||q_i - q_j||_2^2 + \mathbf{1}\{y_i \neq y_j\}\max(0, m - ||q_i - q_j||_2)^2 \tag{12}$$

where m is a constant parameter, q_i is the embedding vector of x_i. For data point of different tasks/labels, contrastive loss rewards the distance between their embedding vectors by L^2 norm, which is weak when the distance is small, as in the case when z is normalized and q_ϕ is randomly initialized. Empirically, we observe that objectives with positive powers of distance lead to degenerate representation of tasks, forming clusters that contain embedding vectors of multiple tasks (Figure 2(a)). Theoretically, this is due to the fact that an accumulative L^2 loss of distance between data points is proportional to the dataset variance, which may lead to degenerate distribution such as Bernoulli distribution. This is proven in Appendix B. To build robust and efficient task inference module, it's crucial to ensure *every* task embedding cluster to be separated from each other. We therefore introduce a negative-power variant of contrastive loss as follows:

$$\mathcal{L}^m_{dml}(x_i, x_j; q) = \mathbf{1}\{y_i = y_j\}||q_i - q_j||_2^2 + \mathbf{1}\{y_i \neq y_j\}\beta \cdot \frac{1}{||q_i - q_j||_2^n + \epsilon} \tag{13}$$

where $\epsilon > 0$ is a small coefficient added to avoid division by zero, the power n can be any non-negative number. We experimented with 1 and 2 in this paper and compare them with the classical L^1, L^2 metrics in 5.2.2.

5 Experiments

In our experiments, we assess the performance of FOCAL by comparing it with several baseline algorithms on meta-RL benchmarks. Specific design choices are examined through three ablations. All return curves are averaged over 3 random seeds. Supplementary experiments are provided in Appendix E.

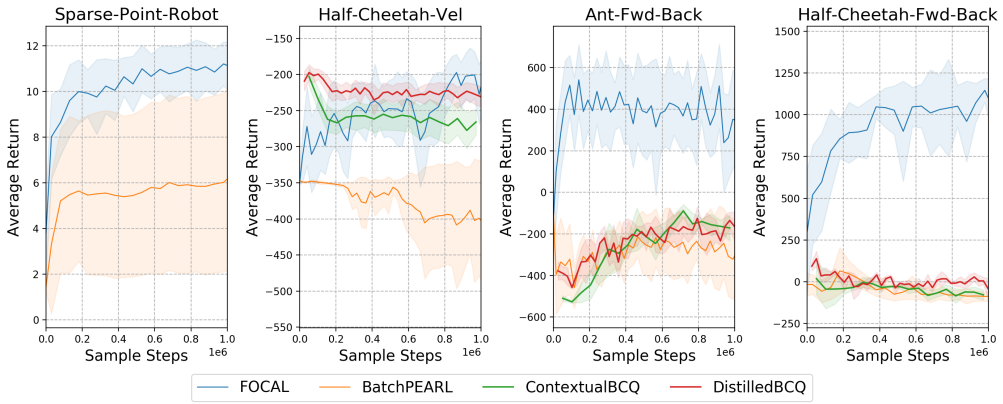

Figure 3: Average episodic testing return of FOCAL vs. other baselines on four meta-environments.

5.1 Sample Efficiency and Performance

We evaluate FOCAL on four continuous control meta-environments of robotic locomotion, three of which are simulated via the MuJoCo simulator (Todorov et al., 2012), with exception of a 2D navigation problem with sparse reward called Sparse-Point-Robot. All environments require adaptation across reward functions. Additionally, we experimented on a unique meta-environment Walker-2D-Params where tasks differ by transition functions and carry out ablations in 5.2.3. These meta-RL benchmarks were previously introduced by Finn et al. (2017) and are described in Appendix D. We train stochastic SAC models for every single task and roll out policies saved at every checkpoint to collect trajectories. The offline training datasets can be collected as a selection of the saved trajectories, which facilitates tuning of the performance level and state-action distributions of the datasets for each task.

For OMRL, there are two natural baselines. The first is by modifying PEARL to train and test from solely logged data without exploration. The second one is Contextual BCQ introduced by Li et al. (2019b). By incorporating a latent space z in the state space for task inference, it can be interpreted as the task-augmented variant of offline BCQ algorithm (Fujimoto et al., 2019). Like PEARL, the task inference module is trained using Bellman gradients. Lastly, we include comparison with the method proposed by Li et al. (2019a) which we term "Distilled BCQ". Although as discussed earlier, Distilled BCQ is a model-based, two-stage method as opposed to our model-free and end-to-end approach, we consider it by far the most competitive and related OMRL algorithm to FOCAL, due to the lack of other OMRL methods.

As shown in Figure 3, we observe that FOCAL outperforms prior offline meta-RL methods across all domains in terms of both sample efficiency and performance. Most notably, when compared to Batch PEARL, a vanilla version of FOCAL without behavior regularization or distance metric learning, the gap in performance is evident due to divergence of value functions as well as sub-optimal representation of the context encoder (Appendix E).

We started experiments with expert-level datasets. However, for some tasks such as Ant-Fwd-Back, we observed that a diverse training sets (generated by expert, medium, random policies) result in a better meta policy (Table 2). *We conjecture that mixed datasets, despite sub-optimal actions, provides a broader support for state-action distributions, making it easier for the context encoder to learn the correct correlation between task identity and transition tuples (i.e., transition/reward functions). While using expert trajectories, there might be little overlap between state-action distributions across tasks (Fig 7), which may cause the agent to learn spurious correlation. This is the exact problem Li et al. (2019b) aims to address.* Moreover, similar to other offline methods such as classical supervised learning, the performance of learned policy is sensitive to data distribution according on our observation (Table 5).

To incorporate a more competitive baseline, we also compare our design to a variant of FOCAL where the context encoder is trained by Bellman gradients, jointly with the control policy, to demonstrate the effectiveness of our proposed decoupled training. We discuss the ablation in 5.2.3.

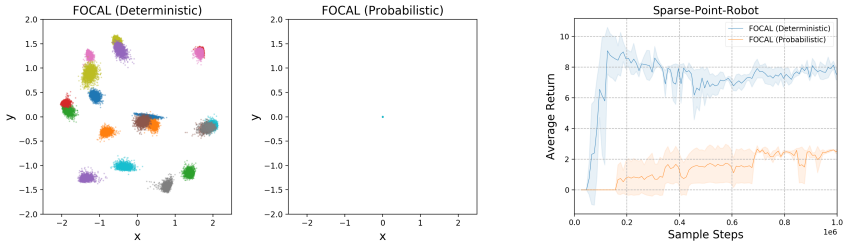

(a) context embedding (b) test-task performance

Figure 4: (a) 2D PCA projection of the embedding vectors drawn from 20 randomized tasks on Walker-2D-Params. Data points are color-coded according to task identity. (b) FOCAL trained with deterministic encoder also outperforms the probabilistic variant on Sparse-Point-Robot.

5.2 Ablations

Based on our previous analysis, we examine and validate three key design choices of FOCAL by the following ablations.

5.2.1 Deterministic vs. Probabilistic Context Encoder

Despite abundance successes of probabilistic/variational inference models in previous work (Kingma & Welling, 2013; Alemi et al., 2016; Rakelly et al., 2019), we observe experimentally that a deterministic context encoder performs significantly better on both simple (Sparse-Point-Robot) and complex (Walker-2D-Params) tasks in fully offline settings, shown in Figure 4. As reflected in our design principles, this is due to

1. Offline meta-RL does not require exploration. Also when the *Task-Transition Correspondence* is satisfied, there is not need for reasoning about uncertainty during adaption.
2. The deterministic context encoder in FOCAL is trained with carefully designed metric-based learning objective, detached from the Bellman update, which provides better efficiency and stability for meta-learning.

In principle, however, the deterministic context encoder which is built upon the *Task-Transition Correspondence* assumption, is less effective in presence of sparse reward or stochastic environments. In this case, a non-negligible portion of samples violate Assumption 1 and provide confounding information for task inference, effective reducing the size of the dataset. Due to these limitations, the fully-offline algorithm may be prone to over-fitting. Indeed, we observe mild performance instability and deterioration similar to overfitting on Half-Cheetah-Vel (Fig 3) and Sparse-Point-Robot (Fig 4(b)). *Nevertheless, we consider Assumption 1 a sufficient condition for the feasibility of our deterministic context encoder design, but not necessary for FOCAL to work well. In fact, even on environments where Assumption 1 is not satisfied (Sparse-Point-Robot), FOCAL still obtains excellent performance overall (Figure 3) given no significant distribution shift between training and testing dataset (Table 5). This is corroborated in our full-range experiments in Fig 6.*

5.2.2 Power Law of Distance Metric Loss

To show the effectiveness of our proposed negative-power distance metrics for OMRL problem, we tested context embedding loss with different powers of distance, from L^{-2} to L^2. A 2D PCA projection view of the high-dimensional embedding space in Figure 2 demonstrates that, distance metric loss with negative powers are more effective in separating embedding vectors of different tasks, whereas positive powers exhibit degenerate behaviors, leading to less robust and effective conditioned policies. By a physical analogy, the inverse-power losses provide "repulsive forces" that drive apart all data points, regardless of the initial distribution. In electromagnetism, consider the latent space as a 3D metal cube and the embedding vectors as positions of positive charges. By Gauss's law, at equilibrium state, all "charges" are distributed on the surface of the cube with densities positively related to the local curvature of the surface. Indeed, we observe from the "Inverse-square" and "Inverse" trials that almost all vectors are located near the edge of the latent space, with higher concentration around the vertices, which have the highest local curvatures (Fig 8(a)).

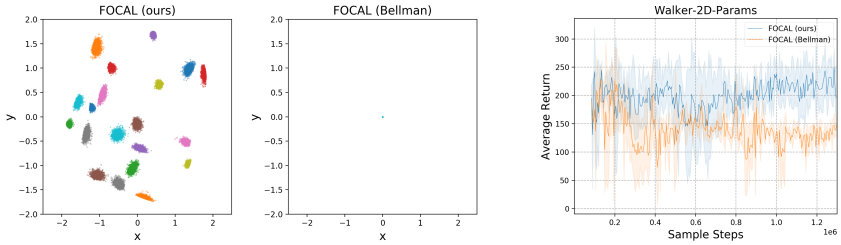

(a) context embedding (b) test-task performance

Figure 5: (a) 2D PCA projection of the embedding vectors drawn from 20 randomized tasks on Walker-2D-Params. Data points are color-coded according to task identity. (b) FOCAL with proposed training strategy also outperforms the counterpart trained via Bellman gradients on Walker-2D-Params.

Loss	RMS	ESR
Inverse-square	1.282	0.430
Inverse	1.217	0.420
Linear	1.385	0.409
Square	1.415	0.253

Table 1: Embedding Statistics on Half-Cheetah-Vel

To evaluate the effectiveness of different powers of DML loss, we define a metric called *effective separation rate* (ESR) which computes the percentage of embedding vector pairs of different tasks whose distance on latent space $\mathcal{Z}$ is larger than the expectation of randomly distributed vector pairs, i.e., $\sqrt{2l/3}$ on $(-1, 1)^l$. Table 1 demonstrates that DML losses of negative power are more effective in maintaining distance between embeddings of different tasks, while no significant distinction is shown in terms of RMS distance, which is aligned with our insight that RMS or effectively classical L^2 objective, can be optimized by degenerate distributions (Lemma 1). This is the core challenge addressed by our proposed inverse-power loss.

5.2.3 Context Encoder Training Strategies

The last key design choice of FOCAL is to train the context encoder via distance-based metrics, fully decoupled from control (Bellman updates). To illustrate the necessity of such design, we compare FOCAL to a more competitive baseline than Batch PEARL by applying the same behavior regularization on the latter. Without DML loss, we train the encoder entirely by Bellman gradients. Since on simple environments such as Sparse-Point-Robot, FOCAL can learn expert-level policy without behavior regularization ($\alpha = 0$ in Table 3), this ablation is effectively addressed in Fig 3. We therefore performed an additional ablation on the unique meta-environment Walker-2D-Params, where tasks differ by transition functions. Fig 5 shows that our proposed strategy achieves effective clustering of task context and therefore better control policy, whereas training via Bellman gradients cannot. We conjecture that on complex tasks where behavior regularization is necessary to ensure convergence, the value functions can often grow to an enormous scale (10^7 for instance, see Fig 9) with large variance, causing the Bellman gradients to dominate over the contribution of task inference module. Evetually, context embedding collapses and fails to learn anything useful. Further analyses are provided in Appendix E.2

6 Conclusion

In this paper, we propose a novel fully-offline meta-Rl algorithm, FOCAL, in pursuit of more practical RL. Our method involves independent distance metric learning of a deterministic context encoder for efficient task inference, combined with an actor-critic apparatus with regularization to learn effective policy from static data. By re-formulating the meta-RL tasks as *task-augmented MDPs* with *Task-Transition Correspondence* assumption, we shed light on the effectiveness of our design choices in both theory and experiments. Our approach achieves superior results compared to existing OMRL algorithms on a diverse set of continous control meta-RL domains. Despite the success, we also discuss the breaking down of our fundamental assumptions in extreme scenarios such as sparse reward or stochastic environments, which opens up avenues for future work of more general and sophisticated OMRL algorithms.

REFERENCES

Rishabh Agarwal, Dale Schuurmans, and Mohammad Norouzi. An optimistic perspective on offline reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2020.

Alexander A Alemi, Ian Fischer, Joshua V Dillon, and Kevin Murphy. Deep variational information bottleneck. *arXiv preprint arXiv:1612.00410*, 2016.

Jonathan Binas, Leonie Luginbuehl, and Yoshua Bengio. Reinforcement learning for sustainable agriculture. In *ICML 2019 Workshop Climate Change: How Can AI Help*, 2019.

Sumit Chopra, Raia Hadsell, and Yann LeCun. Learning a similarity metric discriminatively, with application to face verification. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 1, pp. 539–546. IEEE, 2005.

Ron Dorfman and Aviv Tamar. Offline meta reinforcement learning. *arXiv preprint arXiv:2008.02598*, 2020.

Yan Duan, John Schulman, Xi Chen, Peter L Bartlett, Ilya Sutskever, and Pieter Abbeel. Rl˜2: Fast reinforcement learning via slow reinforcement learning. *arXiv preprint arXiv:1611.02779*, 2016.

Rasool Fakoor, Pratik Chaudhari, Stefano Soatto, and Alexander J Smola. Meta-q-learning. *arXiv preprint arXiv:1910.00125*, 2019.

Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *arXiv preprint arXiv:1703.03400*, 2017.

Scott Fujimoto, David Meger, and Doina Precup. Off-policy deep reinforcement learning without exploration. In *International Conference on Machine Learning*, pp. 2052–2062, 2019.

Mohammad Ghavamzadeh, Shie Mannor, Joelle Pineau, and Aviv Tamar. Bayesian reinforcement learning: A survey. *arXiv preprint arXiv:1609.04436*, 2016.

Omer Gottesman, Fredrik Johansson, Matthieu Komorowski, Aldo Faisal, David Sontag, Finale Doshi-Velez, and Leo Anthony Celi. Guidelines for reinforcement learning in healthcare. *Nat Med*, 25(1):16–18, 2019.

Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*, 2018.

Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pp. 1735–1742. IEEE, 2006.

Tobias Johannink, Shikhar Bahl, Ashvin Nair, Jianlan Luo, Avinash Kumar, Matthias Loskyll, Juan Aparicio Ojea, Eugen Solowjow, and Sergey Levine. Residual reinforcement learning for robot control. In *2019 International Conference on Robotics and Automation (ICRA)*, pp. 6023–6029. IEEE, 2019.

Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2):99–134, 1998.

Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

Gregory Koch, Richard Zemel, and Ruslan Salakhutdinov. Siamese neural networks for one-shot image recognition. In *ICML deep learning workshop*, volume 2. Lille, 2015.

Aviral Kumar, Justin Fu, Matthew Soh, George Tucker, and Sergey Levine. Stabilizing off-policy q-learning via bootstrapping error reduction. In *Advances in Neural Information Processing Systems*, pp. 11784–11794, 2019.

Hongyang Li, David Eigen, Samuel Dodge, Matthew Zeiler, and Xiaogang Wang. Finding task-relevant features for few-shot learning by category traversal. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–10, 2019a.

Jiachen Li, Quan Vuong, Shuang Liu, Minghua Liu, Kamil Ciosek, Henrik Iskov Christensen, and Hao Su. Multi-task batch reinforcement learning with metric learning. *arXiv*, pp. arXiv–1909, 2019b.

Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. A simple neural attentive meta-learner. *arXiv preprint arXiv:1707.03141*, 2017.

Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.

Kate Rakelly, Aurick Zhou, Chelsea Finn, Sergey Levine, and Deirdre Quillen. Efficient off-policy meta-reinforcement learning via probabilistic context variables. In *International conference on machine learning*, pp. 5331–5340, 2019.

Shai Shalev-Shwartz, Shaked Shammah, and Amnon Shashua. Safe, multi-agent, reinforcement learning for autonomous driving. *arXiv preprint arXiv:1610.03295*, 2016.

Noah Y Siegel, Jost Tobias Springenberg, Felix Berkenkamp, Abbas Abdolmaleki, Michael Neunert, Thomas Lampe, Roland Hafner, and Martin Riedmiller. Keep doing what worked: Behavioral modelling priors for offline reinforcement learning. *arXiv preprint arXiv:2002.08396*, 2020.

David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.

Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in neural information processing systems*, pp. 4077–4087, 2017.

Kihyuk Sohn. Improved deep metric learning with multi-class n-pair loss objective. In *Advances in neural information processing systems*, pp. 1857–1865, 2016.

Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1199–1208, 2018.

Sebastian Thrun and Lorien Pratt. *Learning to learn*. Springer Science & Business Media, 2012.

Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–5033. IEEE, 2012.

Ricardo Vilalta and Youssef Drissi. A perspective view and survey of meta-learning. *Artificial intelligence review*, 18(2):77–95, 2002.

Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.

Jane X Wang, Zeb Kurth-Nelson, Dhruva Tirumala, Hubert Soyer, Joel Z Leibo, Remi Munos, Charles Blundell, Dharshan Kumaran, and Matt Botvinick. Learning to reinforcement learn. *arXiv preprint arXiv:1611.05763*, 2016.

Yifan Wu, George Tucker, and Ofir Nachum. Behavior regularized offline reinforcement learning. *arXiv preprint arXiv:1911.11361*, 2019.

Liu Yang and Rong Jin. Distance metric learning: A comprehensive survey. *Michigan State Universiy*, 2(2):4, 2006.

Deheng Ye, Zhao Liu, Mingfei Sun, Bei Shi, Peilin Zhao, Hao Wu, Hongsheng Yu, Shaojie Yang, Xipeng Wu, Qingwei Guo, et al. Mastering complex control in moba games with deep reinforcement learning. In *AAAI*, pp. 6672–6679, 2020.

Appendices

A PSEUDO-CODE

Algorithm 1: FOCAL Meta-training

- Pre-collected batch $\mathcal{D}_i = \{(s_j, a_j, s'_j, r_j)\}_{j:1...N}$ of a set of training tasks $\{\mathcal{T}_i\}_{i=1...n}$ drawn from $p(\mathcal{T})$
- Learning rates $\alpha_1, \alpha_2, \alpha_3$

Initialize context replay buffer $\mathcal{C}_i$ for each task $\mathcal{T}_i$
Initialize inference network $q_\phi(z|c)$, learning policy $\pi_\theta(a|s, z)$ and Q-network $Q_\psi(s, z, a)$ with parameters ϕ, θ and ψ
while not done **do**
  for each $\mathcal{T}_i$ **do**
    for t = 0, $T - 1$ **do**
      Sample mini-batches of B transitions $\{(s_{i,t}, a_{i,t}, s'_{i,t}, r_{i,t})\}_{t:1...B} \sim \mathcal{D}_i$ and update $\mathcal{C}_i$
    end
  end
  Sample meta-batches of M tasks $\sim p(\mathcal{T})$
  for step in training steps **do**
    for each $\mathcal{T}_i$ **do**
      Sample mini-batches c_i and $b_i \sim \mathcal{C}_i$ for context encoder and policy training (b_i, c_i **are identical by default**)
      for each $\mathcal{T}_j$ **do**
        Sample mini-batches c_j from $\mathcal{C}_j$
        $\mathcal{L}^{ij}_{dml} = \mathcal{L}_{dml}(c_i, c_j; q)$
      end
      $\mathcal{L}^{i}_{actor} = \mathcal{L}_{actor}(b_i, q(c_i))$
      $\mathcal{L}^{i}_{critic} = \mathcal{L}_{critic}(b_i, q(c_i))$
    end
    $\phi \leftarrow \phi - \alpha_1 \nabla_\phi \sum_{ij} \mathcal{L}^{ij}_{dml}$
    $\theta \leftarrow \theta - \alpha_2 \nabla_\theta \sum_{i} \mathcal{L}^{i}_{actor}$
    $\psi \leftarrow \psi - \alpha_3 \nabla_\psi \sum_{i} \mathcal{L}^{i}_{critic}$
  end
end

Algorithm 2: FOCAL Meta-testing

- Pre-collected batch $\mathcal{D}_{i'} = \{(s_{j'}, a_{j'}, s'_{j'}, r_{j'})\}_{j':1...M}$ of a set of testing tasks $\{\mathcal{T}_{i'}\}_{i'=1...m}$ drawn from $p(\mathcal{T})$

Initialize context replay buffer $\mathcal{C}_{i'}$ for each task $\mathcal{T}_i$
for each $\mathcal{T}_{i'}$ **do**
  for t = 0, $T - 1$ **do**
    Sample mini-batches of B transitions $c_{i'} = \{(s_{i',t}, a_{i',t}, s'_{i',t}, r_{i',t})\}_{t:1...B} \sim \mathcal{D}_{i'}$ and update $\mathcal{C}_{i'}$
    Compute $z_{i'} = q_\phi(c_{i'})$
    Roll out policy $\pi_\theta(a|s, z_{i'})$ for evaluation
  end
end

B DEFINITIONS AND PROOFS

Lemma 1 *The contrastive loss of a given dataset $\mathcal{X} = \{x_i | i = 1, ..., N\}$ is proportional to the variance of the random variable $X \sim \mathcal{X}$ Consider the contrastive loss $\sum_{i \neq j}(x_i - x_j)^2$, which consists of $N(N-1)$ pairs of different samples (x_i, x_j) drawn from $\mathcal{X}$. It can be written as*

$$\sum_{i \neq j}(x_i - x_j)^2 = 2\left((N-1)\sum_i x_i^2 - \sum_{i \neq j} x_i x_j\right) \tag{14}$$

The variance of $X \sim \mathcal{X}$ is expressed as

$$\begin{aligned} Var(X) &= \overline{(X - \overline{X})^2} && (15)\\ &= \overline{X^2} - (\overline{X})^2 && (16)\\ &= \frac{1}{N}\sum_i x_i^2 - \frac{1}{N^2}(\sum_i x_i)^2 && (17)\\ &= \frac{1}{N^2}\left((N-1)\sum_i x_i^2 - \sum_{i \neq j} x_i x_j\right) && (18) \end{aligned}$$

where $\overline{X}$ denotes the expectation of X. By substituting Eqn. 18 into 14, we have

$$\sum_{i \neq j}(x_i - x_j)^2 = 2N^2(Var(X)) \tag{19}$$

Definition 1 (Task-Augmented MDP) *A task-augmented Markov Decision Process (TA-MDP) can be modeled as $\mathcal{M} = (\mathcal{S}, \mathcal{Z}, \mathcal{A}, P, R, \rho_0, \gamma)$ where*

- *$\mathcal{S}$: state space*
- *$\mathcal{Z}$: contextual latent space*
- *$\mathcal{A}$: action space*
- *P: transition function $P(s', z'|s, z, a)$*
- *R: reward function $R(s, z, a)$*
- *$\rho_0(s, z)$: joint initial state and task distribution*
- *$\gamma \in (0, 1)$: discount factor*

Definition 2 *The Bellman optimality operator $\mathcal{B}_z$ on TA-MDP is defined as*

$$(\mathcal{B}_z \hat{Q})(s, z, a) := R(s, z, a) + \gamma \mathbb{E}_{P(s', z'|s, z, a)}[\max_{a'} \hat{Q}(s', z', a')] \tag{20}$$

Definition 3 (Deterministic MDP) *For a deterministic MDP, a transition map $t : \mathcal{S} \times \mathcal{A} \to \mathcal{S}$ exists such that:*

$$P(s'|s, a) = \delta(s' - t(s, a)) \tag{21}$$

where $\delta(x - y)$ is the Dirac delta function that is zero almost everywhere except $x = y$.

C IMPORTANCE OF DISTANCE METRIC LEARNING FOR META-RL ON TASK-AUGMENTED MDPS

We provide an informal argument that enforcing distance metric learning (DML) is crucial for meta RL on task-augmented MDPs (TA-MDPs). Consider a classical *continuous* neural network with $L \in \mathbb{N}$ layers, $n_l \in \mathbb{N}$ many nodes at the l-th hidden layer for $l = 1, ..., L$, input dimension n_0, output dimension n_{L+1} and nonlinear continuous activation function $\sigma : \mathbb{R} \to \mathbb{R}$ expressed as

$$N_\theta(\boldsymbol{x}) := A_{L+1} \circ \sigma_L \circ A_L \circ \cdots \circ \sigma_1 \circ A_1(\boldsymbol{x}) \tag{22}$$

where $A_l : \mathbb{R}^{n_{l-1}} \to \mathbb{R}^{n_l}$ is an affine linear map defined by $A_l(\boldsymbol{x}) = \boldsymbol{W}_l x + \boldsymbol{b}_l$ for $n_l \times n_{l-1}$ dimensional weight matrix $\boldsymbol{W}_l$ and n_l dimensional bias vector $\boldsymbol{b}_l$ and $\sigma_l : \mathbb{R}^{n_l} \to \mathbb{R}^{n_l}$ is an element-wise nonlinear continuous activation map defined by $\sigma_l(\boldsymbol{z}) := (\sigma(z_1), ..., \sigma(z_{n_l}))$. θ denotes the network parameters. Since every affine and activation map is continuous, their composition N_σ is also continuous, which means by definition of continuity:

$$\forall \epsilon > 0, \quad \exists \eta > 0 \quad \text{s.t.} \tag{23}$$

$$|x_1 - x_2| < \eta \Rightarrow |N_\theta(x_1) - N_\theta(x_2)| < \epsilon \tag{24}$$

where $|\cdot|$ in principle denotes any valid metric defined on Euclidean space $\mathbb{R}^n$. A classical example is the Euclidean distance.

Now consider N_θ as the value function on TA-MDP with deterministic embedding, approximated by a neural network parameterized by θ:

$$\hat{Q}_\theta(s, a, z) \approx Q_\theta(s, a, z) = R_z(s, a) + \gamma \mathbb{E}_{s' \sim P_z(s'|s,a)}[V_\theta(s')] \tag{25}$$

The continuity of neural network implies that for a pair of sufficiently close embedding vectors (z_i, z_j), there exists sufficiently small $\eta > 0$ and $\epsilon > 0$ that

$$z_1, z_2 \in \mathcal{Z}, |z_1 - z_2| < \eta \Rightarrow |\hat{Q}_\theta(s, a, z_1) - \hat{Q}_\theta(s, a, z_2)| < \epsilon \tag{26}$$

Eqn 26 implies that for a pair of different tasks $(\mathcal{T}_i, \mathcal{T}_j) \sim p(\mathcal{T})$, if their embedding vectors are sufficiently close in the latent space $\mathcal{Z}$, the mapped values of meta-learned functions approximated by continuous neural networks are suffciently close too. Since by Eqn 25, due to different transition functions $P_{z_i}(s'|s, a)$, $P_{z_j}(s'|s, a)$ and reward functions $R_{z_i}(s'|s, a)$, $R_{z_j}(s'|s, a)$ of $(\mathcal{T}_i, \mathcal{T}_j)$, the distance between the **true values** of two Q-functions $|Q_\theta(s, a, z_1) - Q_\theta(s, a, z_2)|$ is not guaranteed to be small. This suggests that a meta-RL algorithm with inefficient representation of context embedding $z = q_\phi(c)$, which fails in maintaining effective distance between two distinct tasks $\mathcal{T}_i$, $\mathcal{T}_j$, is unlikely to accurately learn the value functions (or any policy-related functions) for both tasks simultaneously.

D EXPERIMENTAL DETAILS

D.1 DETAILS OF THE MAIN EXPERIMENTAL RESULT (FIGURE 3)

The main experimental result in the paper is the comparative study of performance of FOCAL with three baseline OMRL algorithms (Batch PEARL, Contextual BCQ and Distilled BCQ), shown in Figure 3. Here we plot the same data as in Figure 3 for the full number of time steps taken in our experiments, in Figure 6. Due to limited computational budget, we only performed experiments of baseline algorithms Contextual BCQ and Distilled BCQ on three meta-environments (Half-Cheetah-Vel, Half-Cheetah-Fwd-Back and Ant-Fwd-Back) for 10^6 steps.

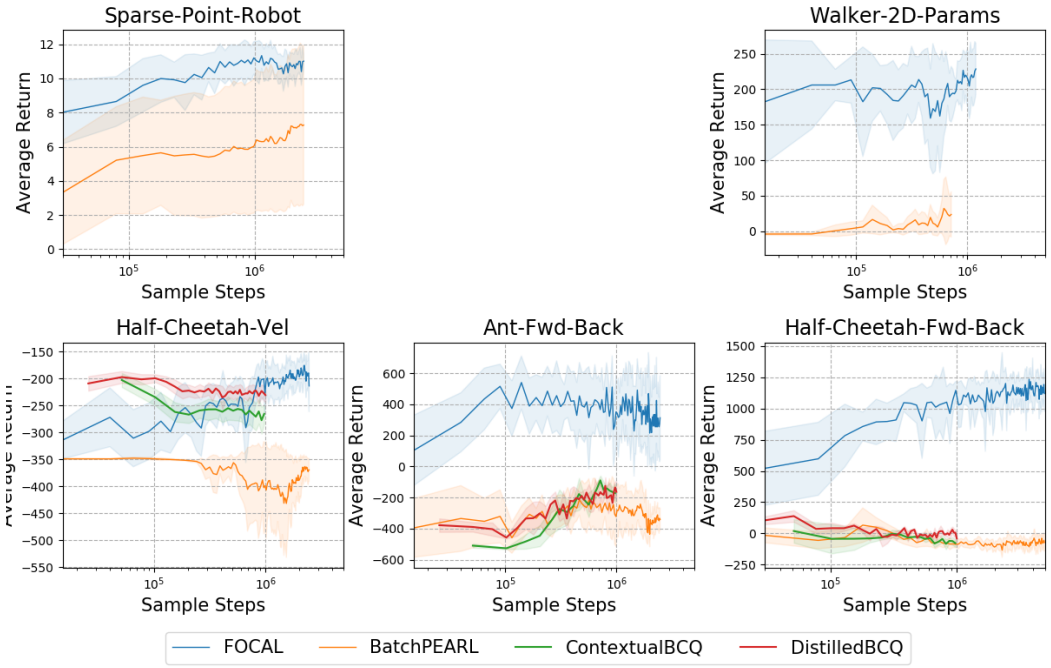

Figure 6: Average episodic testing return of FOCAL vs. other baselines on five meta-environments.

In addition, we provide experimental details on how we produce the result in this section. The performance levels of the training/testing data for the experiments are given in Table 2, which are selected for the best test-time performance over four levels: expert, medium, random, mixed (containing all logged trajectories of trained SAC models from beginning (random quality) to end (expert quality). For mixed data, the diversity of samples is optimal but the average performance level is lower than expert:

Meta Env	Training Data	Testing Data
Sparse-Point-Robot	expert	expert
Half-Cheetah-Vel	expert	expert
Ant-Fwd-Back	mixed	mixed
Half-Cheetah-Fwd-Back	mixed	mixed
Walker-2D-Params	mixed	mixed

Table 2: Quality of data used for best test-time performance. We maintain the same quality of data for training and testing due to algorithm's sensitivity to distribution shift. From our experiments, we observe that for some envs/tasks, datasets with the best performance generate the best testing result, whereas for some envs/tasks, the diversity of data matters the most.

The DML loss used in experiments in Figure 3 is inverse-squared, which gives the best performance among the four power laws we experimented with in Figure 2.

D.2 Description of the Meta Environments

- Sparse-Point-Robot: A 2D navigation problem introduced in PEARL (Rakelly et al., 2019). Starting from the origin, each task is to guide the agent to a specific goal located on the unit circle centered at the origin. Non-sparse reward is defined as the negative distance from the current location to the goal. In sparse-reward scenario, reward is truncated to 0 when the agent is outside a neighborhood of the goal controlled by the goal radius. While inside the neighborhood, agent is rewarded by 1 - distance, a positive value.
- Half-Cheetah-Fwd-Back: Control a Cheetah robot to move forward or backward. **Reward function** is dependent on the walking direction.
- Half-Cheetah-Vel: Control a Cheetah robot to achieve a target velocity running forward. **Reward function** is dependent on the target velocity.
- Ant-Fwd-Back: Control an Ant robot to move forward or backward. **Reward function** is dependent on the walking direction.
- Walker-2D-Params: Agent is initialized with some system dynamics parameters randomized and must move forward, it is a unique environment compared to others since tasks differ in transition function. **Transitions function** is dependent on randomized task-specific parameters such as mass, inertia and friction coefficients.

D.3 Hyperparameter Settings

The details of important hyperparameters used to produce the experimental results in the paper are presented in Table 3 and 4.

Hyperparameters	Sparse-Point-Robot	Half-Cheetah-Vel	Ant-Fwd-Back	Half-Cheetah-Fwd-Back	Walker-2D-Params
reward scale	100	5	5	5	5
DML loss weight(β)	1	1	1	1	1
behavior regularization strength(α)	0	500	1e6	500	500
buffer size (per task)	1e4	1e4	1e4	1e4	1e4
batch size	256	256	256	256	256
meta batch size	16	16	4	4	16
g_lr(f-divergence discriminator)	1e-4	1e-4	1e-4	1e-4	1e-4
dml_lr(α_1)	1e-3	1e-3	1e-3	1e-3	1e-3
actor_lr(α_2)	1e-3	1e-3	1e-3	1e-3	1e-3
critic_lr(α_3)	1e-3	1e-3	1e-3	1e-3	1e-3
discount factor	0.9	0.99	0.99	0.99	0.99
# training tasks	80	80	2	2	40
# testing tasks	20	20	2	2	10
goal radius	0.2	N/A	N/A	N/A	N/A
latent space dimension	5	5	20	5	5
network width (context encoder)	200	200	200	200	200
network depth (context encoder)	3	3	3	3	3
network width (others)	300	300	300	300	300
network depth (others)	3	3	3	3	3
maximum episode length	20	200	200	200	200

Table 3: Hyperparameters used to produce Figure 3. Meta batch size refers to the number of tasks used for computing the DML loss $\mathcal{L}_{dml}^{ij}$ at a time. Larger meta batch size leads to faster convergence but requires greater computational power.

Hyperparameters	Half-Cheetah-Vel
reward scale	5
behavior regularization strength(α)	500
buffer size (per task)	1e4
batch size	256
meta batch size	16
g_lr(f-divergence discriminator)	1e-4
dml_lr(α_1)	1e-3
actor_lr(α_2)	1e-3
critic_lr(α_3)	1e-3
discount factor	0.99
# training tasks	80
# testing tasks	20
latent space dimension	5
network width (context encoder)	200
network depth (context encoder)	3
network width (others)	300
network depth (others)	3
maximum episode length	200

(a) Compared to Half-Cheetah-Vel experiment in Table 3, the meta batch and latent space dimension were reduced to speed up computation.

Trials	β	ϵ
Inverse-Square	1	0.1
Inverse	2	0.1
Linear	8	0.1
Square	16	0.1

(b) The DML loss weight β and coefficient ϵ (defined in Eqn 13) used in experiments of Figure 2 to match the scale of objective functions of different power laws. The weights are chosen such that all terms are equal when the average distance of x_i and x_j per dimension is 0.5, a reasonable value given $x \in (-1, 1)^l$.

Table 4: Hyperparameters used to produce Figure 2

E Additional Experiments

E.1 Sensitivity to Distribution Shift

Since in OMRL, all datasets are static and fixed, many challenges from classical supervised learning such as as over-fitting apply. By developing FOCAL, we are also interested in its sensitivity to distribution shift for better understanding of OMRL algorithms. Since for each task $\mathcal{T}_i$, our data-generating behavior policies $\beta_i(a|s)$ are trained from random to expert level, we select three performance levels (expert, medium, random) of datasets to study how combinations of training/testing sets with different quality/distribution affect performance. An illustration of the three quality levels on Sparse-Point-Robot is shown in Fig 7.

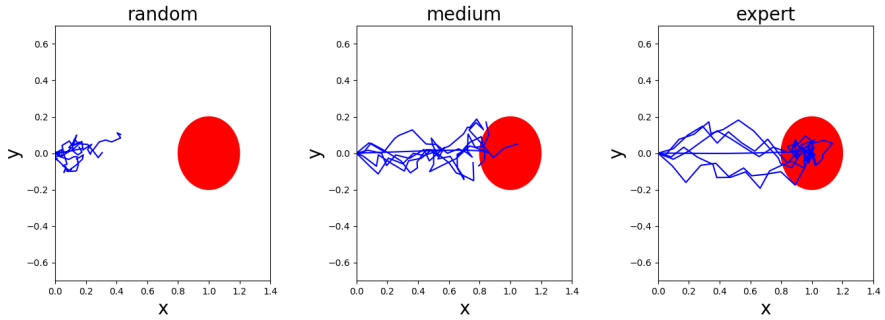

Figure 7: Distribution of rollout trajectories of trained SAC policies of three performance levels: random, medium and expert. Since reward is sparse, only states that lie in the red circle are given non-zero rewards, making meta-learning more challenging and sensitive to data distributions.

Table 5 shows the average return at test-time for various training and testing distributions. Sensitivity to distribution shift is confirmed since training/testing on the similar distribution of data result in relatively higher performance. In particular, this is significant in sparse reward scenario when Assumption 1 is no longer satisfied. With severe over-fitting and task mis-identification , performance of meta-RL policy is inevitably compromised by distribution mismatch between training/testing datasets.

Training Set	Testing Set	Avg. Testing Return
expert	expert	8.16 ± 1.31
medium	medium	8.44 ± 2.59
random	random	2.34 ± 0.76
medium	expert	8.25 ± 2.91
expert	medium	7.12 ± 1.63
medium	random	6.76 ± 1.90
expert	random	4.43 ± 2.04

Table 5: Average testing return of FOCAL on Sparse-Point-Robot tasks with different qualities/distributions of training/testing sets. Training/testing on both expert and medium level datasets give the best results.

E.2 Context Encoder Training Strategies: 5.2.3 Continued

As discussed several times in the paper, one of the key design choices of FOCAL is to decouple learning of task inference and control in terms of gradient computation, rather than sampling strategy as in PEARL. We present our experimental evidence on Sparse-Point-Robot and Half-Cheetah-Vel in Fig 8. Since it was found empirically that Sparse-Point-Robot is the only environment among the four that does not require behavior regularization to work, our control group of Batch PEARL can be seen as a variant of FOCAL which learns from Bellman update (same as PEARL) instead of DML loss.

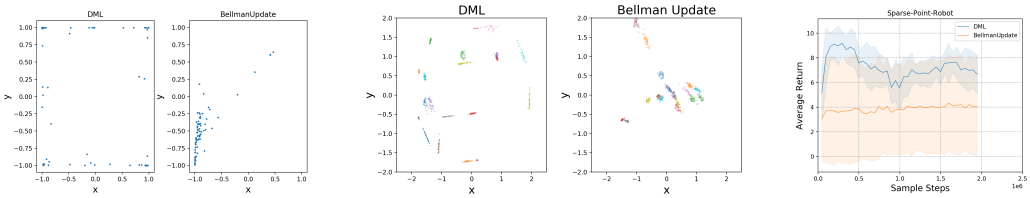

(a) context embedding (b) context embedding (PCA) (c) test-task performance

Figure 8: FOCAL with context encoder q_ϕ trained by DML loss compared to encoder trained by Bellman update (the critic objective defined in Eqn. 15 in the paper). (a) On Half-Cheetah-Vel, viewed from a 2D projection of the latent space, the encoder trained by DML loss separates the embedding vectors more effectively. (b) On Half-Cheetah-Vel, training with DML loss demonstrate faster convergence, better asymptotic performance and less variance. (c) Return curve on Sparse-Point-Robot.

E.3 Divergence of Q-functions in Offline Setting

The necessity of applying behavior regularization on environment like Ant-Fwd-Back to prevent divergence of value functions is demonstrated in Figure 9.

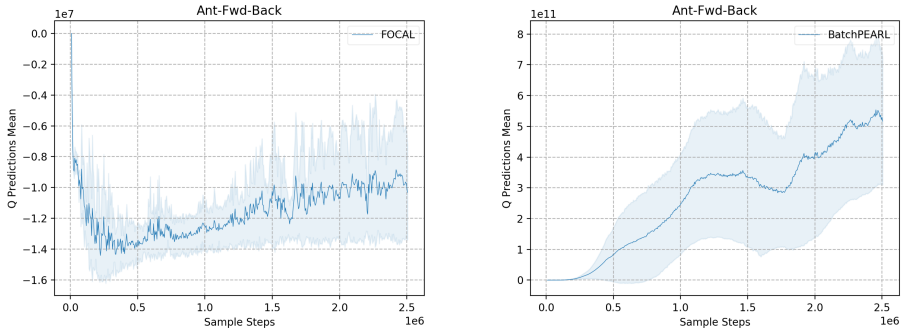

(a) FOCAL

(b) Batch PEARL

Figure 9: On complex environment such as Ant-Fwd-Back, the Q-function learned by Batch PEARL diverges ($> 10^{11}$) whereas the Q-function of FOCAL, despite its large order of magnitude, converges eventually given proper regularization ($\alpha = 10^6$)

F IMPLEMENTATION

We build our algorithm on top of PEARL and BRAC, both are derivatives of the SAC algorithm. SAC is an off-policy actor-critic method with a maximum entropy RL objective which encourages exploration and learning a stochastic policy. Although exploration is not needed in fully-offline scenarios, we found empirically that a maximum entropy augmentation is still beneficial for OMRL, which is likely due to the fact that different actions result in same next state and reward in environments such as Ant.

All function approximators in FOCAL are implemented as neural networks with MLP structures. For normalization, the last activation layer of context encoder and policy networks are invertible squashing operators (tanh), making $\mathcal{Z}$ a bounded Euclidean space $(-1, 1)^l$, which is reflected in Figure 1(a).

Since the training of context encoder and policy are completely decoupled, the whole FOCAL pipeline involves three main objectives. The DML loss for training the inference network $q_\phi(z|c)$ is given by Eqn 13, for mini-batches of transitions drawn from training datasets: $x_i \sim \mathcal{D}_i$, $x_j \sim \mathcal{D}_j$. The embedding vector q_i, q_j are computed as the average embedding over x_i and x_j. The actor and critic losses are the task-augmented version of Eqn 8 and 9:

$$\mathcal{L}_{\text{critic}} = \mathbb{E}_{\substack{(s,a,r,s')\sim\mathcal{D} \\ a'\sim\pi_\theta(\cdot|s')}} \left[\left(r + \gamma \bar{Q}^D_\psi(s', \bar{z}, a') - Q_\psi(s, \bar{z}, a) \right)^2 \right] \tag{27}$$

$$\mathcal{L}_{\text{actor}} = -\mathbb{E}_{(s,a,r,s')\sim\mathcal{D}} \left[\mathbb{E}_{a''\sim\pi_\theta(\cdot|s)} [Q_\psi(s, \bar{z}, a'')] - \alpha \hat{D} \right] \tag{28}$$

where $\bar{Q}$ is a target network and $\bar{z}$ indicates that gradients are not being computed through it. As discussed in (Kumar et al., 2019; Wu, Tucker, and Nachum, 2019), the divergence function $\hat{D}$ can take form of Kernel MMD (Gretton et al., 2012), Wasserstein Divergence (Arjovsky, Chintala, and Bottou, 2017) or f-divergences (Nowozin et al., 2016) such as KL divergence. In this paper, we use the dual form (Nowozin, Cseke, and Tomioka, 2016) of KL divergence, which learns a discriminator g with minimax optimization to circumvent the need of a cloned policy for density estimation.

In principle, as a core design choice of PEARL, the context used to infer $q_\phi(z|c)$ can be distinct from the data used to construct the critic loss. In OMRL however, we found this treatment unnecessary since there is no exploration. Therefore training of DML and actor-critic objectives share the same sampled data, which form an end-to-end algorithm described in Algorithm 1 and 2.