Meta-Reward-Net: Implicitly Differentiable Reward Learning for Preference-based Reinforcement Learning

Anonymous Author(s)
Affiliation
Address
email

Abstract

Setting up a well-designed reward function has been challenging for many reinforcement learning applications. Preference-based reinforcement learning (PbRL) provides a new framework that avoids reward engineering by leveraging human preferences (i.e., preferring apples over oranges) as the reward signal. Therefore, improving the efficacy of data usage for preference data becomes critical. In this work, we propose Meta-Reward-Net (MRN), a data-efficient PbRL framework that incorporates bi-level optimization for both reward and policy learning. The key idea of MRN is to adopt the performance of the Q-function as the learning target. Based on this, MRN learns the Q-function and the policy in the inner level while updating the reward function adaptively according to the performance of the Q-function on the preference data in the outer level. Our experiments on locomotion tasks and robotic manipulation tasks demonstrate that MRN outperforms prior methods in the case of little feedback and significantly improves data efficiency, achieving state-of-the-art in preference-based RL. Ablation studies further demonstrate that MRN learns a more accurate Q-function compared to prior work and shows obvious advantages when only a small amount of feedback is available.

1 Introduction

In recent years, reinforcement learning has achieved great success in solving complex sequential decision-making tasks, such as video or board games [1, 2, 3, 4], autonomous driving [5], quantitative finance [6], automation system [7, 8, 9], etc. For common decision making tasks, the goal of the agent is to maximize the cumulative reward. However, one central challenge to reinforcement learning is how to design reward functions. On the one hand, the quality of the designed reward function largely depends on the problem solver's understanding of task objective, operation logic, and related background knowledge. Even excellent engineers still need plenty of time to try different methods in complex RL tasks. On the other hand, there is a problem that the agent might hack the reward function. In policy learning, the agent utilizes the defect of the reward function to maximize the cumulative reward instead of solving the expected task. Besides, in human-involved scenarios, the objective of the agent is to maximize happiness of humans, making it hard to specify a reward function.

Previous work has provided some ideas to avoid directly constructing reward functions, such as imitation learning [10]. Although imitation learning has an excellent performance in some tasks, its performance is difficult to surpass human level. Preference-based reinforcement learning is a more flexible and convenient alternative method. A human expert can easily give a preference to a trajectory pair of the agent, which implies the desired behaviors, that is to say, the goal of the task. In preference-based RL, the reward function is learned through preferences given by the human teacher on trajectory pairs. The feedback from humans guides the agent to achieve specified goals or learn

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

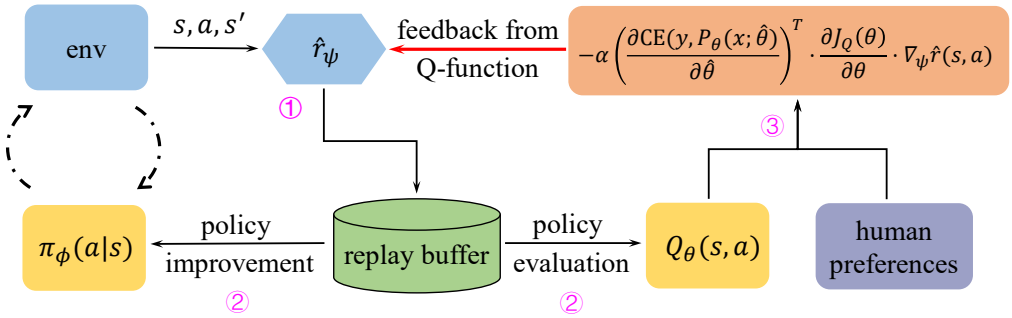

Figure 1: Framework of Meta-Reward-Net. ① Trajectories are sampled by interacting with the environment and reward is labeled by $\hat{r}_\psi$. ② Transitions are sampled from the replay buffer and are relabeled by the up-to-date $\hat{r}_\psi$ for optimizing the policy and the Q-function. ③ The performance of the Q-function on the preference data is evaluated to provide the gradient as a type of feedback for reward learning.

desired behaviors. Recent work [11, 12, 13] shows that providing sufficient feedback achieves better performance under this setting.

However, preference data queried from human experts is expensive. In reality, there is likely to be only a small amount of data available. With the limitation of the amount of feedback, previous methods perform poorly or even do not work. Meanwhile, recent work on semi-supervised learning takes success in computer vision [14], pseudo labels can be adjusted according to the feedback from student networks and further improve student's performance. Inspired by this, we consider utilizing the feedback from the Q-function for reward learning via bi-level optimization. By considering the loss calculated according to the feedback, the optimization of the reward function is aware of the accuracy of the Q-function, which is beneficial for the learning of the Q-function and the policy.

In this work, we focus on the efficiency of feedback in the learning of preference-based RL. We use bi-level optimization method for reward learning which contains two loops. In the inner loop, we update the Q-function and the policy via the reward function, while optimizing the reward function according to the performance of the Q-function on the preference data in the outer loop. Our experiments demonstrate that our method considerably improves feedback efficiency. Besides, the results of further evaluation show that MRN exceeds other methods by a large margin when a small amount of feedback is available and learns a more accurate Q-function.

In summary, the main contributions of our work are three-fold. Firstly, we propose a new preference-based RL algorithm Meta-Reward-Net, which utilizes bi-level optimization methods in reward learning. Secondly, we show that MRN substantially improves the feedback efficiency and outperforms preference-based RL baselines on a variety of robotic manipulation tasks from Meta-world [15] and locomotion tasks from DeepMind Control Suite (DMControl) [16, 17]. Lastly, we demonstrate that benefiting from bi-level optimization, the advantage becomes obvious compared to PEBBLE when only a small amount of feedback is provided. We also show that the feedback from the Q-function to the reward function is beneficial for agent learning, leading to a more accurate Q-function and a better policy.

2 Related Work

Reinforcement Learning. Reinforcement learning has gradually become an effective and powerful method to solve complex sequential decision-making problems. In recent years, much prior work has proven the capacity of reinforcement learning. Reinforcement learning applications include video games [1, 2, 3], robot control [18], manipulation [8, 9], board games [4], autonomous driving [5] and so on. In addition, there are many applications of reinforcement learning in the fields of computer vision [19], natural language processing [20], and recommendation system [21, 22]. In the framework of reinforcement learning, an agent obtains data from interaction with the environment to optimize policy to maximize the expected return. In this process, the reward function plays a crucial role. Our

method differs in that we do not assume that there is a reward function from engineering. Instead, we utilize the preferences of humans to guide the agent to learn desired behaviors.

Preference-based Reinforcement Learning. Prior work has successfully trained the agent to complete specific tasks or achieve goals through the feedback from the teacher. [23] provides a general learning framework of preference-based reinforcement learning. [24] utilizes two kinds of feedback, initializes with imitation learning policy, and further improves the performance of the policy with feedback from humans. In this work, we mainly focus on feedback efficiency in the learning of preference-based RL. Much previous work [25, 26, 27, 28, 29] considered learning reward function from the most informative data to be consistent with the human preferences. Recently, several feedback-efficient preference-based RL algorithms have been proposed. PEBBLE [11] combines unsupervised pre-training and the technique of relabeling experience to improve feedback efficiency. SURF [12] learns the reward function by semi-supervised learning and data augmentation. RUNE [13] facilitates exploration via reward uncertainty to reduce the amount of feedback. However, these methods only focus on guiding reward learning through supervised loss between human preferences and estimated preferences. We have a different approach in that in addition to utilizing the supervised loss, we also consider the performance of the Q-function on a labeled dataset as feedback to assist reward learning, thus beneficial for agent learning. By introducing extra feedback, MRN gets more task-related information from human preferences to improve feedback efficiency remarkably.

Bi-level Optimization. In computer vision, there are several bi-level optimization algorithms that achieve great success. Meta-Weight-Net [30] provides weights for samples from an unbalanced dataset in a bi-level manner, while Meta Label Correction [31] view this problem as a label correction problem. Meta Pseudo Labels [14] combines Pseudo Labels methods and bi-level optimization to generate high quality pseudo labels. In reinforcement learning, LIIR [32] learns an intrinsic reward function for each agent to achieve better cooperation in multi-agent reinforcement learning. CAIL [33] proposes to reweight demonstrations of different optimality in imitation learning. Our method uses a similar bi-level method for reward learning. To the best of our knowledge, we are the first to introduce bi-level optimization into preference-based RL.

3 Preliminaries

Preference-based Reinforcement Learning. In standard reinforcement learning framework, a finite Markov decision process (MDP) can be presented as a tuple of $\langle \mathcal{S}, \mathcal{A}, R, P, \gamma \rangle$, which consists of state space $\mathcal{S}$, action space $\mathcal{A}$, transition function, reward function, and discount factor. $P(s'|s,a)$ represents stochastic dynamics of the environment, which is the probability of selecting action a to transit to s' in a given state s. $R(s,a)$ represents the reward obtained by selecting an action a in a given state s. The policy $\pi(a|s)$ is a mapping from state space to action space. The objective of the agent is to collect trajectories from interaction with the environment to maximize the expected return.

In the general preference-based RL from [23], there is no reward function from reward engineering and a reward function estimator $\widehat{r}_\psi$ should be learned to be consistent with preferences from the human expert. Specifically, a segment σ is a sequence of states and actions which is $(s_{t+1}, a_{t+1}, \cdots, s_{t+k}, a_{t+k})$. Human expert provides a preference y on given two segments (σ^0, σ^1) and y is the distribution over $\{0,1\}$, $y \in \{(1,0),(0,1),(0.5,0.5)\}$. Following the Bradley-Terry model [34], a preference predictor constructed by the reward function estimate $\widehat{r}_\psi$ is formulated as:

$$P_\psi[\sigma^0 \succ \sigma^1] = \frac{\exp \sum_t \widehat{r}_\psi(s_t^0, a_t^0)}{\exp \sum_t \widehat{r}_\psi(s_t^0, a_t^0) + \exp \sum_t \widehat{r}_\psi(s_t^1, a_t^1)}, \tag{1}$$

where $\sigma^0 \succ \sigma^1$ denotes σ^0 is more consistent with the expectations of human experts compared with σ^1. The reward function learning can be solved by minimizing the cross-entropy loss between predictions from preference predictors and human preferences.

$$\mathcal{L}_{\text{supervised}}(\psi) = -\mathop{\mathbb{E}}_{(\sigma^0,\sigma^1,y)\sim\mathcal{D}} \Big[y(0) \log P_\psi[\sigma^0 \succ \sigma^1] + y(1) \log P_\psi[\sigma^1 \succ \sigma^0] \Big]. \tag{2}$$

We refer to this objective as supervised loss in the following sections. By optimizing the reward function using this loss, segments that are more in line with human preferences obtain a higher cumulative reward.

Soft Actor-Critic. SAC [35] is an off-policy algorithm based on the maximum entropy RL, which encourages the agent to explore the environment by acting as randomly as possible. SAC consists of

soft Q-function $Q_\theta(s,a)$ with parameters θ and policy $\pi_\phi(a|s)$ with parameters ϕ. Q-function with parameters θ is defined as the expectation of return:

$$Q_\theta(s_t,a_t) = \mathbb{E}\left[\sum_{t'=t}^{T}\gamma^{t'-t}r_{t'}\,\middle|\, S_t = s_t, A_t = a_t\right], \tag{3}$$

where $\gamma \in [0,1]$ is a discount factor.

The parameters θ of soft Q-function are trained by minimizing the soft Bellman residual:

$$J_Q(\theta) = \mathbb{E}_{\tau_t\sim\mathcal{B}}\Big[\big(Q_\theta(s_t,a_t) - r_t - \gamma\bar{V}(s_{t+1})\big)^2\Big], \tag{4}$$

where $\bar{V}(s_t) = \mathbb{E}_{a_t\sim\pi_\phi}\big[Q_{\bar{\theta}}(s_t,a_t) - \alpha\log\pi_\phi(a_t|s_t)\big]$, $\tau_t = (s_t, a_t, s_{t+1}, r_t)$ is the transition at time step t, α is a learnable temperature parameter that controls the item of entropy, $\bar{\theta}$ are parameters of the target soft Q-function, and $\mathcal{B}$ is replay buffer. After the updating of the Q-function, policy π_ϕ is updated by minimizing the loss:

$$J_\pi(\phi) = \mathbb{E}_{s_t\sim\mathcal{B},a_t\sim\pi_\phi}\Big[\alpha\log\pi_\phi(a_t|s_t) - Q_\theta(s_t,a_t)\Big]. \tag{5}$$

By performing policy evaluation and policy improvement alternately, SAC trains an agent with excellent and stable performance. In this work, we consider using SAC as our backbone reinforcement learning algorithm.

4 Meta-Reward-Net

In this section, we formally present Meta-Reward-Net, which includes two key components, optimizing the reward function based on the performance of the Q-function in outer loop and learning the agent in the inner loop. In the following, we first provide a new perspective that the Q-function can be used to compute preference labels, then define the objective of MRN and formulate a bi-level optimization problem.

4.1 The Objective

The probability that segment σ^0 is preferred is proportional to the exponential return of it. Motivated by this, we use $Q_\theta(s_0^0, a_0^0)$ and $Q_\theta(s_0^1, a_0^1)$ to respectively measure of the return of σ^0 and σ^1 since the Q-value equal to the expectation of segment return. Therefore, the probability that σ^0 is preferred to σ^1 is computed through the Q-function:

$$P_\theta[\sigma^0 \succ \sigma^1] = \frac{\exp Q_\theta(s_0^0, a_0^0)}{\exp Q_\theta(s_0^0, a_0^0) + \exp Q_\theta(s_0^1, a_0^1)}. \tag{6}$$

Given human preference labels y, the performance of the Q-function is measured by the cross-entropy loss between preference predictions computed by (6) and ground-truth labels:

$$\mathcal{L}_{\text{feedback}}(\theta(\psi)) = -\mathop{\mathbb{E}}_{(\sigma^0,\sigma^1,y)\sim\mathcal{D}}\Big[y(0)\log P_{\theta(\psi)}[\sigma^0 \succ \sigma^1] + y(1)\log P_{\theta(\psi)}[\sigma^1 \succ \sigma^0]\Big], \tag{7}$$

where $\theta(\psi)$ denotes the updating of θ depends on the reward provided by $\widehat{r}_\psi$.

The core of MRN is utilizing the feedback from the Q-function in reward learning, which means that the optimization of the reward function takes the performance of current Q-function into consideration. Based on this idea, we formulate the objective using (7) as a measurement of the Q-function. The objective of MRN is to minimize the loss of Q_θ on a labeled preference dataset and the Q-function is trained by minimizing the Bellman residual. The overall objective is formulated as:

$$\begin{aligned} \min_{\psi,\theta} \quad & \mathcal{L}_{\text{feedback}}(\theta(\psi)), \\ \text{s.t.} \quad & \theta(\psi) = \arg\min_{\theta} J_Q(\psi,\theta). \end{aligned} \tag{8}$$

By formulating MRN as a bi-level optimization algorithm, this allows the reward function to provide rewards that are beneficial for improving the performance of the Q-function, which further leads to a better policy π_ϕ.

4.2 Bi-level Optimization

The objective in (8) is solved by the following bi-level optimization algorithm: θ is optimized by the reward estimation from $\widehat{r}_\psi$ in inner loop while ψ is updated according to the performance of the Q-function on the labeled data in outer loop.

Pseudo Updating: Building Connection between θ and ψ. To utilize the feedback information from the Q-function and improve the performance of Q-function, we formulate the outer loop updating as optimizing the loss of θ on the labeled preference data with respect to ψ. However, we can not directly optimize this objective since the objective is independent of ψ. So the first step is to build a connection between θ and ψ. Sample a mini-batch state-action pairs from the replay buffer and use them to query current reward function with parameters $\psi^{(k)}$ to obtain reward estimation $\widehat{r}_\psi(s_t, a_t)$, where k denotes current iteration step. Then (4) becomes:

$$J_Q(\theta) = \mathbb{E}_{\tau_t \sim \mathcal{B}}\Big[\big(Q_\theta(s_t, a_t) - \widehat{r}_\psi(s_t, a_t) - \gamma \bar{V}(s_{t+1})\big)^2 \Big]. \tag{9}$$

At each bi-level optimization step, we first pseudo update the parameters of the Q-function. Pseudo updating means that we do not directly perform the updating on the Q-function, but update the parameters of a copy of current Q-function by minimizing (9) with learning rate α:

$$\hat{\theta}^{(k)} = \theta^{(k)} - \alpha \, \nabla_\theta J_Q(\psi, \theta)|_{\theta^{(k)}} \, , \tag{10}$$

where $\hat{\theta}^{(k)}$ denotes the updated copy of $\theta^{(k)}$. By performing (10), the connection between $\hat{\theta}^{(k)}$ and $\psi^{(k)}$ is built.

Outer Loop: Optimizing ψ to Improve the Performance of Q_θ on Labeled Data. After building connections through pseudo updating, the copy of Q-function with parameters $\hat{\theta}^{(k)}$ is tested on labeled preference data. The predicted preference label $P_\theta(x)$ is computed by the Q-function using (6), where x denotes a segment pair (σ^0, σ^1). We use implicit differentiation in our method. The objective of outer loop is formulated in (7), and the implicit derivative of the outer loss with respect to ψ is calculated using the chain rule:

$$\begin{aligned} g^{(k)}_{\text{feedback}} &= \nabla_\theta \mathcal{L}_{\text{feedback}}|_{\hat{\theta}^{(k)}} \; \nabla_\psi \hat{\theta}^{(k)} \Big|_{\psi^{(k)}} \\ &= h \cdot \nabla_\psi \widehat{r}_\psi(s_t, a_t)|_{\psi^{(k)}} \, , \end{aligned} \tag{11}$$

where $h = -\alpha \cdot \left(\nabla_\theta \mathcal{L}_{\text{feedback}} \left(\hat{\theta}^{(k)} \right) \right)^\top \cdot \nabla_\theta J_Q \left(\theta^{(k)} \right)$ and full derivation can be found in Appendix A. (11) formulates the gradient of feedback from the Q-function and this can be done easily using automatic differentiation in Pytorch [36]. Since pseudo updating is performed and the connection between θ and ψ is built, ψ is updated to improve the performance of the Q-function by minimizing the cross-entropy between preference labels from Q_θ and ground-truth labels:

$$\psi^{(k+1)} = \psi^{(k)} - \beta \; g^{(k)}_{\text{feedback}} \Big|_{\psi^{(k)}} \, , \tag{12}$$

where β is the learning rate of the outer loop.

Inner Loop: Optimizing θ and ϕ. In the inner loop, the objective is the same as (4) and (5) in the training of SAC. To calculate the new reward $\widehat{r}_\psi(s_t, a_t)$, use the same batch of state-action pairs in pseudo updating to query the updated reward function. However, we do not need the connection in the inner loop since the connection is used for outer level optimization. With newly obtained reward estimation, we update Q-function Q_θ by minimizing (4) with learning rate α:

$$\theta^{(k+1)} = \theta^{(k)} - \alpha \, \nabla_\theta J_Q(\theta)|_{\theta^{(k)}} \, , \tag{13}$$

and update policy π_ϕ by minimizing (5) with learning rate α:

$$\phi^{(k+1)} = \phi^{(k)} - \alpha \, \nabla_\theta J_\pi(\phi)|_{\phi^{(k)}} \, . \tag{14}$$

Auxiliary Loss. In addition to optimizing the loss in the outer loop, the reward function is augmented with a supervised loss generally used in preference-based RL, which is formulated in (2). Our intuition is that the optimization of the reward function needs both supervised learning and the feedback from the Q-function, and neither of them can be removed. On the one hand, supervised loss

is necessary for it ensures the reward estimation is consistent with human preferences. On the other hand, the feedback loss at the outer level is beneficial because it improves the performance of the Q-function, leading to a more accurate Q-function and finally a better policy.

The full procedure of our method is detailed in Algorithm 1. Before reward learning, we first initialize the policy and replay buffer with unsupervised exploration, which is proposed in PEBBLE [11] and can be found in Appendix B. The reward function is updated using supervised loss per K iterations, while the bi-level optimization is performed per N iterations. We use off-policy RL algorithm SAC to collect transitions and save them in the replay buffer. The Q-function and the policy is optimized in each training step.

Algorithm 1 Meta-Reward-Net

Input: supervised reward learning frequency K, bi-level updating frequency N
Input: number of human's preference labels per session M
Initialize θ and ψ
Initialize a preference dataset $\mathcal{D} \leftarrow \emptyset$
Initialize $\mathcal{B}$ and ϕ with unsupervised exploration
for each iteration **do**
  Take action $a_t \sim \pi_\phi(a_t|s_t)$ and obtain $s_{t+1} \sim p(s_{t+1} \mid s_t, a_t)$
  Store transition $\{(s_t, a_t, s_{t+1}, \widehat{r}_\psi(s_t, a_t))\}$ in $\mathcal{B}$
  Sample minibatch $\{(\tau_j)\}_{j=1}^{B} \sim \mathcal{B}$
  if iteration $\%\ K == 0$ **then**
    Query a human teacher for M preference labels and store them in $\mathcal{D}$
    Sample preference data in $\mathcal{D}$
    Optimize (2) with respect to ψ
    Use updated $\widehat{r}_\psi$ to relabel the replay buffer $\mathcal{B}$
  end if
  if iteration $\%\ N == 0$ **then**
    Sample preference data in $\mathcal{D}$
    Pseudo update θ using (10)
    Update ψ using (12)
    Use updated $\widehat{r}_\psi$ to relabel the replay buffer $\mathcal{B}$
  end if
  Update θ and ϕ using (13) and (14), respectively
end for
Output: policy π_ϕ

5 Experiments

In this section, our method is evaluated on a variety of robotic manipulation tasks from Meta-world [15] and locomotion tasks from DeepMind Control Suite (DMControl) [16, 17]. The tasks used in our experiments are shown in Appendix C, which are the same as the tasks used in SURF.

5.1 Setup

Baselines. Reward-based SAC and three state-of-the-art preference-based RL algorithms are used for comparison:

- SAC [35]: SAC is considered as the ground-truth algorithm since the agent is provided with real reward function, which is not the case in preference-based RL. SAC is evaluated in our experiments because it is the backbone RL algorithm of PEBBLE.
- Preference PPO [23]: the method is a re-implementation using PPO [37]. It uses an ensemble of reward functions and disagreement sampling for querying.
- PEBBLE [11]: the method is a preference-based RL method with unsupervised exploration and reward relabeling.
- SURF [12]: the method combines temporal data augmentation and pseudo labels in semi-supervised learning, which is the state-of-the-art algorithm in preference-based RL.

- Meta-Reward-Net (MRN): the proposed method, which utilizes the feedback from the Q-function in reward learning through bi-level optimization.

Implementation Details. For all methods, we use unsupervised exploration proposed in PEBBLE [11]. For the sampling of queries, disagreement-based sampling is used for all preference-based RL methods, following the setting in [23]. An ensemble of three reward functions is used and the reward output is computed by averaging output of three reward functions. To systematically evaluate the performance and speed up the training process, following the setting in PEBBLE [11] and SURF [12], we consider a script teacher that can always provide the ground-truth preference label of a segment pair. Concretely, this is implemented by comparing the ground-truth return of each segment, but the reward is not accessible to the agent under the setting of preference-based RL.

For the implementation of SAC, Preference PPO and PEBBLE, we use the publicly released repository of B-Pref [38] in our experiments.[1] In their implementation, Preference PPO is re-implemented using on-policy algorithm PPO. SURF is also implemented using their released code.[2] For SAC, PEBBLE and SURF, the hyperparameters and network architectures we use are the same as them (e.g., number of network layers, learning rate, frequency of supervised reward learning). For the amount of human's preference feedback, we set 100 for Walker, Cheetah, Button Press and Window Open, 700 for Quadruped, 1000 for Door Open and Drawer Open, 4000 for Sweep Into, and 10000 for Hammer.

Our method is implemented by using PEBBLE as the backbone. We use bi-level updating frequency $N = 5000$ for Cheetah, Hammer, Door Open, Button Press, Drawer Open and Window Open, $N = 1000$ for Walker, $N = 3000$ for Quadruped, and $N = 10000$ for Sweep Into.

For each task, we run all algorithms independently for ten times and report the average with a standard deviation. Tasks of Meta-world are measured on success rate while the tasks of DMControl are measured on ground-truth episode return. The experiments are run on a single machine with one NVIDIA RTX 2080 Ti GPU. Details on hyperparameters, network architectures can be found in Appendix C.

5.2 Results

Meta-world Tasks. Examples of the six continuous control tasks from Meta-world are shown in Appendix C. These tasks are selected for our experiments, including robotic manipulation skills of various difficulty. The details of the tasks can be found in Appendix C.

Figure 2 shows the training curves of MRN and the baselines on the Meta-world tasks. In this figure, SAC achieves the best performance in each task by using the ground-truth reward function. Since little feedback is provided, we observe that there is a gap between all preference-based RL methods and the best performance, but MRN still exceeds the preference-based RL baselines by a large margin. These results demonstrate that MRN considerably improves the performance when only a small amount of feedback is available. We also notice that Preference PPO does not work in most tasks.

DMControl Tasks. For DMControl, three locomotion tasks in Appendix C are used for evaluation, including Walker-walk, Cheetah-run and Quadruped-walk. These tasks encourage the agent to move forward by providing the agent with the reward that is positively correlated with agent's velocity. However, the agent is not accessible to the ground-truth reward function, and all preference-based RL methods are provided with human preference labels. Similar to tasks of Meta-world, we only provide the baselines and MRN with few but the same amount of preference labels.

Figure 3 shows the results of five methods on DMControl tasks. SURF achieves the same return as PEBBLE, while MRN shows obvious advantages compared with baseline methods. The results demonstrate that MRN performs well with a small amount of feedback and improves feedback efficiency.

We remark that MRN can be regarded as introducing feedback from the Q-function through bi-level optimization based on PEBBLE. Comparing the results of MRN and PEBBLE in Figure 2 and Figure 3, we find that using bi-level optimization can significantly improve performance when only a small amount of feedback is provided.

[1] `https://github.com/rll-research/BPref`

[2] `https://github.com/alinlab/SURF`

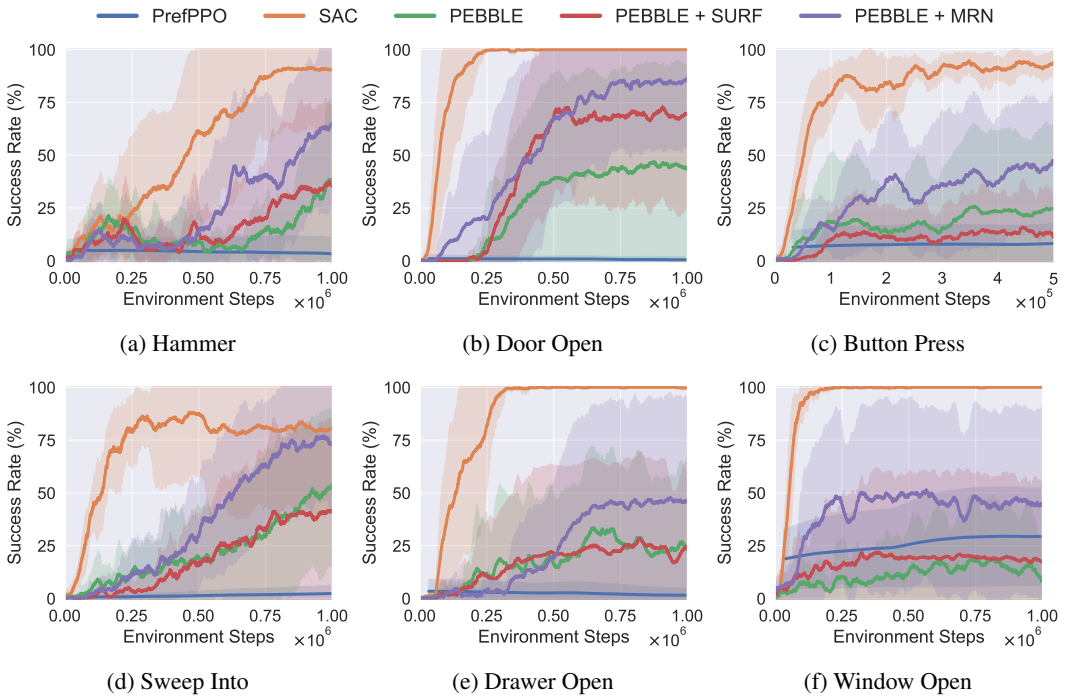

(a) Hammer (b) Door Open (c) Button Press

(d) Sweep Into (e) Drawer Open (f) Window Open

Figure 2: Training curves on six continuous control tasks from Meta-world. The solid line and shaded regions respectively denote mean and standard deviation of success rate, across ten runs.

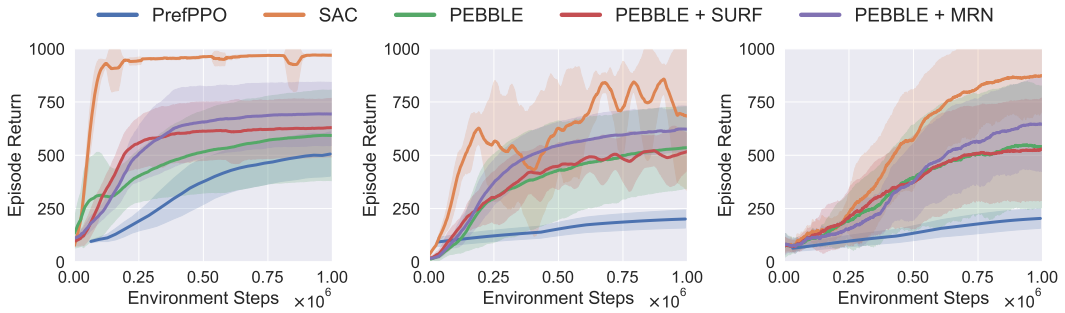

(a) Walker (b) Cheetah (c) Quadruped

Figure 3: Training curves on three continuous control tasks from DMControl. The solid line and shaded regions respectively denote mean and standard deviation of success rate, across ten runs.

5.3 Ablation Study

Number of Human's Feedback. Figure 2 and 3 show that our method outperforms Preference PPO, PEBBLE and SURF under relatively small amounts of human's feedback. To further analyze the performance of different algorithms with different amounts of feedback, Extensive experiments are conducted on Walker and Door Open with varying amounts of feedback: $\{100, 400, 1000, 2000\}$. The results in Figure 4 suggest that our method performs better when the number of human's feedback is small. As the number increases, the performance of three algorithms becomes closer. Intuitively, this phenomenon could be considered that as feedback information becomes more sufficient, the performance gap caused by feedback efficiency gradually disappears.

Accuracy of Q-function. To compare the quality of Q-function trained by MRN and baselines, we the mean squared error between ground-truth Q-values and the output of Q-function. Each method is evaluated on the same ten trajectories by calculating the MSE of the output of the Q-function and the ground-truth Q-values. We report mean and standard deviation across ten runs in Table 1. SAC is presented as the upper bound of the quality of Q-function and MRN achieves the lowest

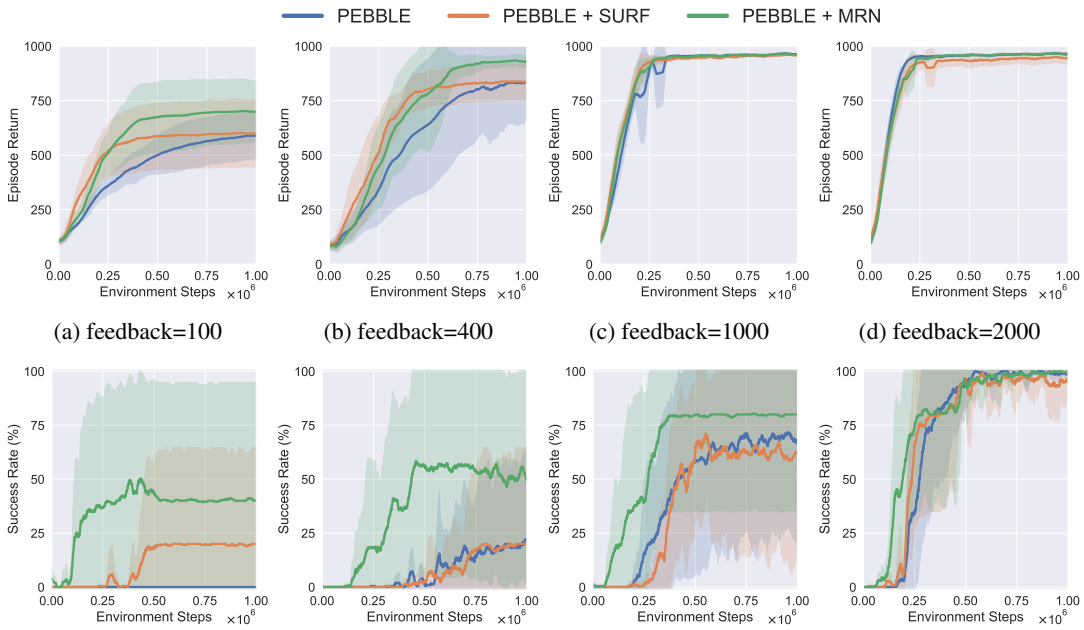

(a) feedback=100 (b) feedback=400 (c) feedback=1000 (d) feedback=2000

(e) feedback=100 (f) feedback=400 (g) feedback=1000 (h) feedback=2000

Figure 4: Training curves on Walker (first row) and Door Open (second row), measured by the ground-truth reward and success rate, respectively. The solid line and shaded regions respectively denote mean and standard deviation, across five runs.

MSE among three methods. In MRN, reward learning additionally takes the feedback from the Q-function into consideration. So the reward provided by the reward function is not only consistent with the ground-truth preference label but also suitable for the performance of the current Q-function. Therefore, the Q-function learned by MRN is more accurate compared to the baselines.

Table 1: Mean squared error of learned Q-function across ten runs.

Task/Method	PEBBLE	PEBBLE+SURF	PEBBLE+MRN (Ours)	SAC
Walker	0.12 ± 0.04	0.11 ± 0.04	$\mathbf{0.10} \pm \mathbf{0.02}$	0.07 ± 0.02
Door Open	0.42 ± 0.91	0.31 ± 0.58	$\mathbf{0.12} \pm \mathbf{0.12}$	0.03 ± 0.02

6 Conclusion

In this work, we propose MRN, a novel feedback-efficient preference-based RL method. By incorporating bi-level optimization for reward and policy learning, we demonstrate MRN outperforms prior methods when a small amount of feedback is available and considerably improves the feedback efficiency on a variety of robotic tasks. In particular, our method exceeds the baselines by a large margin when the amount of feedback is small. From empirical results and analysis, we conclude that the efficiency of feedback in our method mainly benefits from: firstly, our method learns a more accurate Q-function. Secondly, our method learns the Q-function and policy in the inner loop and optimizes reward function according to the performance of Q-function on the preference data in the outer loop. By this way, MRN successfully establishes an efficient mode of information transmission, which can extract more information. To our knowledge, we do not see any potential negative social influences of our work. We hope our method can provide inspiration for future work and encourage preference-based reinforcement learning to be better extended to practical applications.

References

[1] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin A. Riedmiller. Playing atari with deep reinforcement learning. *CoRR*, abs/1312.5602, 2013.

[2] Oriol Vinyals, Igor Babuschkin, Wojciech M. Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H. Choi, Richard Powell, Timo Ewalds, Petko Georgiev, Junhyuk Oh, Dan Horgan, Manuel Kroiss, Ivo Danihelka, Aja Huang, Laurent Sifre, Trevor Cai, John P. Agapiou, Max Jaderberg, Alexander Sasha Vezhnevets, Rémi Leblond, Tobias Pohlen, Valentin Dalibard, David Budden, Yury Sulsky, James Molloy, Tom Le Paine, Çaglar Gülçehre, Ziyu Wang, Tobias Pfaff, Yuhuai Wu, Roman Ring, Dani Yogatama, Dario Wünsch, Katrina McKinney, Oliver Smith, Tom Schaul, Timothy P. Lillicrap, Koray Kavukcuoglu, Demis Hassabis, Chris Apps, and David Silver. Grandmaster level in starcraft II using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.

[3] Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemyslaw Debiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Christopher Hesse, Rafal Józefowicz, Scott Gray, Catherine Olsson, Jakub Pachocki, Michael Petrov, Henrique Pondé de Oliveira Pinto, Jonathan Raiman, Tim Salimans, Jeremy Schlatter, Jonas Schneider, Szymon Sidor, Ilya Sutskever, Jie Tang, Filip Wolski, and Susan Zhang. Dota 2 with large scale deep reinforcement learning. *CoRR*, abs/1912.06680, 2019.

[4] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.

[5] Sen Wang, Daoyuan Jia, and Xinshuo Weng. Deep reinforcement learning for autonomous driving. *arXiv preprint arXiv:1811.11329*, 2018.

[6] Xiao-Yang Liu, Hongyang Yang, Jiechao Gao, and Christina Dan Wang. Finrl: deep reinforcement learning framework to automate trading in quantitative finance. In Anisoara Calinescu and Lukasz Szpruch, editors, *International Conference on AI in Finance (ICAIF)*, pages 1:1–1:9. ACM, 2021.

[7] Chris Gamble and Jim Gao. Safety-first ai for autonomous data centre cooling and industrial control. In *DeepMind*. 2018.

[8] Dmitry Kalashnikov, Alex Irpan, Peter Pastor, Julian Ibarz, Alexander Herzog, Eric Jang, Deirdre Quillen, Ethan Holly, Mrinal Kalakrishnan, Vincent Vanhoucke, and Sergey Levine. Qt-opt: Scalable deep reinforcement learning for vision-based robotic manipulation. *CoRR*, abs/1806.10293, 2018.

[9] Shixiang Gu, Ethan Holly, Timothy Lillicrap, and Sergey Levine. Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3389–3396, 2017.

[10] Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems (NeurIPS)*, volume 29. Curran Associates, Inc., 2016.

[11] Kimin Lee, Laura M Smith, and Pieter Abbeel. Pebble: Feedback-efficient interactive reinforcement learning via relabeling experience and unsupervised pre-training. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning (ICML)*, volume 139 of *Proceedings of Machine Learning Research*, pages 6152–6163. PMLR, 18–24 Jul 2021.

[12] Jongjin Park, Younggyo Seo, Jinwoo Shin, Honglak Lee, Pieter Abbeel, and Kimin Lee. SURF: Semi-supervised reward learning with data augmentation for feedback-efficient preference-based reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2022.

[13] Xinran Liang, Katherine Shu, Kimin Lee, and Pieter Abbeel. Reward uncertainty for exploration in preference-based reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2022.

[14] Hieu Pham, Zihang Dai, Qizhe Xie, and Quoc V. Le. Meta pseudo labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11557–11568, June 2021.

[15] Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Karol Hausman, Chelsea Finn, and Sergey Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In Leslie Pack Kaelbling, Danica Kragic, and Komei Sugiura, editors, *Proceedings of the Conference on Robot Learning (CoRL)*, volume 100 of *Proceedings of Machine Learning Research*, pages 1094–1100. PMLR, 30 Oct–01 Nov 2020.

[16] Yuval Tassa, Yotam Doron, Alistair Muldal, Tom Erez, Yazhe Li, Diego de Las Casas, David Budden, Abbas Abdolmaleki, Josh Merel, Andrew Lefrancq, et al. Deepmind control suite. *arXiv preprint arXiv:1801.00690*, 2018.

[17] Saran Tunyasuvunakool, Alistair Muldal, Yotam Doron, Siqi Liu, Steven Bohez, Josh Merel, Tom Erez, Timothy Lillicrap, Nicolas Heess, and Yuval Tassa. dm_control: Software and tasks for continuous control. *Software Impacts*, 6:100022, 2020.

[18] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2016.

[19] A. V. Bernstein and E. V. Burnaev. Reinforcement learning in computer vision. In Antanas Verikas, Petia Radeva, Dmitry Nikolaev, and Jianhong Zhou, editors, *Tenth International Conference on Machine Vision (ICMV)*, volume 10696, pages 458 – 464. International Society for Optics and Photonics, SPIE, 2018.

[20] Victor Zhong, Caiming Xiong, and Richard Socher. Seq2sql: Generating structured queries from natural language using reinforcement learning. *arXiv preprint arXiv:1709.00103*, 2017.

[21] Lixin Zou, Long Xia, Zhuoye Ding, Jiaxing Song, Weidong Liu, and Dawei Yin. Reinforcement learning to optimize long-term user engagement in recommender systems. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, page 2810–2818, New York, NY, USA, 2019. Association for Computing Machinery.

[22] Guanjie Zheng, Fuzheng Zhang, Zihan Zheng, Yang Xiang, Nicholas Jing Yuan, Xing Xie, and Zhenhui Li. Drn: A deep reinforcement learning framework for news recommendation. In *Proceedings of the 2018 World Wide Web Conference (WWW)*, page 167–176, Republic and Canton of Geneva, CHE, 2018. International World Wide Web Conferences Steering Committee.

[23] Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems (NeurIPS)*, volume 30. Curran Associates, Inc., 2017.

[24] Borja Ibarz, Jan Leike, Tobias Pohlen, Geoffrey Irving, Shane Legg, and Dario Amodei. Reward learning from human preferences and demonstrations in atari. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems (NeurIPS)*, volume 31. Curran Associates, Inc., 2018.

[25] Dorsa Sadigh, Anca Dragan, Shankar Sastry, and Sanjit Seshia. Active preference-based learning of reward functions. In *Proceedings of Robotics: Science and Systems (RSS)*, Cambridge, Massachusetts, July 2017.

[26] Erdem Biyik and Dorsa Sadigh. Batch active preference-based learning of reward functions. In Aude Billard, Anca Dragan, Jan Peters, and Jun Morimoto, editors, *Proceedings of The 2nd Conference on Robot Learning (CoRL)*, volume 87 of *Proceedings of Machine Learning Research*, pages 519–528. PMLR, 29–31 Oct 2018.

[27] Erdem Biyik, Kenneth Wang, Nima Anari, and Dorsa Sadigh. Batch active learning using determinantal point processes. *CoRR*, abs/1906.07975, 2019.

[28] Malayandi Palan, Gleb Shevchuk, Nicholas Charles Landolfi, and Dorsa Sadigh. Learning reward functions by integrating human demonstrations and preferences. In *Proceedings of Robotics: Science and Systems (RSS)*, FreiburgimBreisgau, Germany, June 2019.

[29] Erdem Biyik, Nicolas Huynh, Mykel Kochenderfer, and Dorsa Sadigh. Active Preference-Based Gaussian Process Regression for Reward Learning. In *Proceedings of Robotics: Science and Systems (RSS)*, Corvalis, Oregon, USA, July 2020.

[30] Jun Shu, Qi Xie, Lixuan Yi, Qian Zhao, Sanping Zhou, Zongben Xu, and Deyu Meng. Meta-weight-net: Learning an explicit mapping for sample weighting. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems (NeurIPS)*, volume 32. Curran Associates, Inc., 2019.

[31] Guoqing Zheng, Ahmed Hassan Awadallah, and Susan T. Dumais. Meta label correction for noisy label learning. In *Thirty-Fifth Conference on Artificial Intelligence (AAAI)*, pages 11053–11061, 2021.

[32] Yali Du, Lei Han, Meng Fang, Ji Liu, Tianhong Dai, and Dacheng Tao. Liir: Learning individual intrinsic reward in multi-agent reinforcement learning. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems (NeurIPS)*, volume 32. Curran Associates, Inc., 2019.

[33] Songyuan Zhang, Zhangjie Cao, Dorsa Sadigh, and Yanan Sui. Confidence-aware imitation learning from demonstrations with varying optimality. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems (NeurIPS)*, volume 34, pages 12340–12350. Curran Associates, Inc., 2021.

[34] Ralph Allan Bradley and Milton E. Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.

[35] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning (ICML)*, volume 80 of *Proceedings of Machine Learning Research*, pages 1861–1870. PMLR, 10–15 Jul 2018.

[36] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.

[37] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

[38] Kimin Lee, Laura Smith, Anca Dragan, and Pieter Abbeel. B-pref: Benchmarking preference-based reinforcement learning. In J. Vanschoren and S. Yeung, editors, *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks (NeurIPS)*, volume 1, 2021.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] See abstract and Section 1.
   (b) Did you describe the limitations of your work? [Yes] See Section 6.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See Section 6.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes] We have read the ethics and make sure our work conforms to them.
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...

(a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See implementation details in Section 5.1.

(b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See implementation details in Section 5.1 and Appendix C.

(c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Figure 2, Figure 3, Figure 4 and Table 1 in Section 5.2 and Section 5.3.

(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Section 5.1.

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

(a) If your work uses existing assets, did you cite the creators? [Yes] Our code refers to the source code published by B-Pref and SURF, and their papers are cited in our paper. See Section 5.1.

(b) Did you mention the license of the assets? [Yes] All referenced code is open source and we notice that B-Pref provides MIT license and mujoco provides Apache License. As for DMControl and Meta-world, you can see references for license details.

(c) Did you include any new assets either in the supplemental material or as a URL? [Yes] We release the link of our source code in Appendix C.

(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes] MIT license permits us to copy and modify the source code of B-Pref. We also get the permission of the author of SURF.

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]