

---

# B-Pref: Benchmarking Preference-Based Reinforcement Learning

---

Kimin Lee, Laura Smith, Anca Dragan, Pieter Abbeel  
UC Berkeley

## Abstract

Reinforcement learning (RL) requires access to a reward function that incentivizes the right behavior, but these are notoriously hard to specify for complex tasks. Preference-based RL provides an alternative: learning policies using a teacher’s preferences without pre-defined rewards, thus overcoming concerns associated with reward engineering. However, it is difficult to quantify the progress in preference-based RL due to the lack of a commonly adopted benchmark. In this paper, we introduce B-Pref: a benchmark specially designed for preference-based RL. A key challenge with such a benchmark is providing the ability to evaluate candidate algorithms quickly, which makes relying on real human input for evaluation prohibitive. At the same time, simulating human input as giving perfect preferences for the ground truth reward function is unrealistic. B-Pref alleviates this by simulating teachers with a wide array of irrationalities, and proposes metrics not solely for performance but also for robustness to these potential irrationalities. We showcase the utility of B-Pref by using it to analyze algorithmic design choices, such as selecting informative queries, for state-of-the-art preference-based RL algorithms. We hope that B-Pref can serve as a common starting point to study preference-based RL more systematically. Source code is available at <https://github.com/rll-research/B-Pref>.

## 1 Introduction

Deep reinforcement learning (RL) has emerged as a powerful method to solve a variety of sequential decision-making problems, including board games [58, 59], video games [10, 44, 68], autonomous control [9, 52], and robotic manipulation [5, 32, 35, 36]. However, scaling RL to many applications is difficult due to the challenges associated with defining a suitable reward function, which often requires substantial human effort. Specifying the reward function becomes harder as the tasks we want the agent to achieve become more complex (e.g., cooking or self-driving). In addition, RL agents are prone to exploit reward functions by discovering ways to achieve high returns in ways the reward designer did not expect nor intend. It is important to consider this phenomenon of reward exploitation, or reward hacking, since it may lead to unintended but dangerous consequences [28]. Further, there is nuance in how we might want agents to behave, such as obeying social norms that are difficult to account for and communicate effectively through an engineered reward function [4, 57, 67].

Preference-based RL [19, 31, 39] provides an alternative: a (human) teacher provides preferences between the two agent behaviors, and the agent then uses this feedback to learn desired behaviors (see Figure 1). This framework enables us to optimize the agent using RL without hand-engineered rewards by learning a reward function, which is consistent with the observed preferences [11, 12, 50]. Because a teacher can interactively guide agents according to their progress, preference-based RL has shown promising results (e.g., solving a range of RL benchmarks [19, 31], teaching novel behaviors [61, 71], and mitigating the effects of reward exploitation [39]).

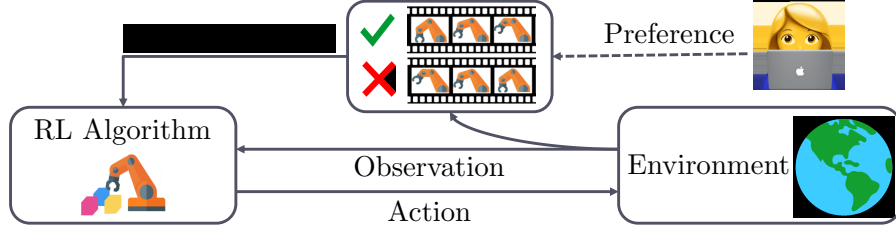


Figure 1: Illustration of preference-based RL. Instead of assuming that the environment provides a (hand-engineered) reward, a teacher provides preferences between the agent’s behaviors, and the agent uses this feedback in order to learn the desired behavior.

Despite significant progress on RL benchmarks designed for various purposes (e.g., offline RL [24, 26], generalization [20, 21], meta RL [76], and safe RL [48]), existing benchmarks are not tailored towards preference-based RL. The lack of a standard evaluation benchmark makes it hard to quantify scientific progress. Indeed, without consistent evaluation, it is not easy to understand the effects of algorithmic and design decisions or compare them across papers.

In this paper, we introduce B-Pref: a benchmark for preference-based RL consisting of various locomotion and robotic manipulation tasks from DeepMind Control Suite [65, 66] and Meta-world [76]. While utilizing real human input is ideal, this is prohibitive because it is hard to evaluate candidate algorithms quickly using real human input. Prior works [19, 31, 39] address this issue by simulating human input as giving perfect preferences with respect to an underlying ground truth reward function. However, evaluation on such ideal teachers is unrealistic because actual humans can exhibit various irrationalities [18] in decision making. So, in our benchmark, we design simulated human teachers with a wide array of irrationalities and propose evaluation metrics not solely for performance but also for robustness to these potential irrationalities.

To serve as a reference, we benchmark state-of-the-art preference-based RL algorithms [19, 39] in B-Pref and showcase the utility of B-Pref by using it to analyze algorithmic design choices for preference-based RL. Although existing methods provide fairly efficient performance on perfectly rational teachers, the poor performance on more realistic, irrational teachers calls for new algorithms to be developed.

The benchmark and reference implementations are available at <https://github.com/rll-research/B-Pref>. We believe that systematic evaluation and comparison will not only further our understanding of the strengths of existing algorithms, but also reveal their limitations and suggest directions for future research.

## 2 Preliminaries

We consider an agent interacting with an environment in discrete time [63]. At each timestep  $t$ , the agent receives a state  $\mathbf{s}_t$  from the environment and chooses an action  $\mathbf{a}_t$  based on its policy  $\pi$ .

In traditional reinforcement learning, the environment also returns a reward  $r(\mathbf{s}_t, \mathbf{a}_t)$  and the goal of agent is to maximize the discounted sum of rewards. However, for many complex domains and tasks, it is difficult to construct a suitable reward function. We consider the preference-based RL framework, where a (human) teacher provides preferences between the agent’s behaviors and the agent uses this feedback to perform the task [19, 31, 39, 41].

Formally, a segment  $\sigma$  is a sequence of observations and actions  $\{(\mathbf{s}_1, \mathbf{a}_1), \dots, (\mathbf{s}_H, \mathbf{a}_H)\}$ . Given a pair of segments  $(\sigma^0, \sigma^1)$ , a teacher indicates which segment is preferred, i.e.,  $y = (0, 1)$  or  $(1, 0)$ , that the two segments are equally preferred  $y = (0.5, 0.5)$ , or that two segments are incomparable, i.e., discarding the query. The goal of preference-based RL is to train an agent to perform behaviors desirable to a human teacher using as few queries as possible.

---

**Algorithm 1** SimTeacher: Simulated human teachers

---

**Require:** Discount factor  $\gamma$ , rationality constant  $\beta$ , probability of making a mistake  $\epsilon$

**Require:** Skip threshold  $\delta_{\text{skip}}$ , equal threshold  $\delta_{\text{equal}}$

**Require:** Pair of segments  $\sigma^0, \sigma^1$

```
1: if  $\max_{i \in \{0,1\}} \sum_t r(\mathbf{s}_t^i, \mathbf{a}_t^i) < \delta_{\text{skip}}$  then // SKIPPING QUERY
2:    $y \leftarrow \emptyset$ 
3: else if  $|\sum_t r(\mathbf{s}_t^1, \mathbf{a}_t^1) - \sum_t r(\mathbf{s}_t^0, \mathbf{a}_t^0)| < \delta_{\text{equal}}$  then // EQUALLY PREFERABLE
4:    $y \leftarrow (0.5, 0.5)$ 
5: else if  $\sigma^0 \succ \sigma^1 \sim P[\sigma^0 \succ \sigma^1; \beta, \gamma]$  then // SAMPLING PREFERENCES FROM (1)
6:    $y \leftarrow (1, 0)$  with probability of  $1 - \epsilon$ 
7:    $y \leftarrow (0, 1)$  otherwise // MAKING A MISTAKE
8: else
9:    $y \leftarrow (0, 1)$  with probability of  $1 - \epsilon$ 
10:   $y \leftarrow (1, 0)$  otherwise // MAKING A MISTAKE
11: end if
12: return  $y$ 
```

---

### 3 B-Pref: Benchmarks environments for preference-based RL

#### 3.1 Design factors

While ideally we would evaluate algorithms' real-world efficacy using real human feedback, designing a standardized and broadly available benchmark becomes challenging because we do not have ground truth access to the human's reward function. Instead, we focus on solving a range of existing RL tasks (see Section 3.4) using a simulated human, whose preferences are based on a ground truth reward function. Because simulated human preferences are immediately generated by the ground truth reward, we are able to evaluate the agent quantitatively by measuring the true average return and do more rapid experiments. A major challenge with simulating human input is that real humans are not perfectly rational and will not provide perfect preferences. To alleviate this challenge, we propose to simulate human input using a wide array of irrationalities (see Section 3.2), and measure an algorithm's robustness in handling such input (see Section 3.3).

#### 3.2 Simulated human teachers

We start from a (perfectly rational) deterministic teacher, which generates preferences as follows:

$$y = \begin{cases} (1, 0) & \text{If } \sum_{t=1}^H r(\mathbf{s}_t^0, \mathbf{a}_t^0) > \sum_{t=1}^H r(\mathbf{s}_t^1, \mathbf{a}_t^1) \\ (0, 1) & \text{otherwise,} \end{cases}$$

where  $H > 0$  is a length of segment  $\sigma$  and  $r$  is the ground truth reward. We remark that prior works [19, 31, 39] evaluated their methods using this ideal teacher. However, evaluating the performance of preference-based RL only using the ideal teacher is unrealistic because there are many possible irrationalities [17, 18] affecting a teacher's preferences (and expression of preferences) in different ways.

To design more realistic models of human teachers, we consider a common stochastic model [11, 19, 50] and systematically manipulate its terms and operators (see Algorithm 1):

**Stochastic preference model.** Because preferences from the human can be noisy, we generate preferences using a stochastic model defined as follows (Line 5):

$$P[\sigma^i \succ \sigma^j; \beta, \gamma] = \frac{\exp\left(\beta \sum_{t=1}^H \gamma^{H-t} r(\mathbf{s}_t^i, \mathbf{a}_t^i)\right)}{\exp\left(\beta \sum_{t=1}^H \gamma^{H-t} r(\mathbf{s}_t^i, \mathbf{a}_t^i)\right) + \exp\left(\beta \sum_{t=1}^H \gamma^{H-t} r(\mathbf{s}_t^j, \mathbf{a}_t^j)\right)}, \quad (1)$$

where  $\gamma \in (0, 1]$  is a discount factor to model myopic behavior,  $\beta$  is a rationality constant, and  $\sigma^i \succ \sigma^j$  denotes the event that segment  $i$  is preferable to segment  $j$ . This follows the Bradley-Terry model [13], which can be interpreted as assuming the probability of preferring a segment depends exponentially on the sum over the segment of an underlying reward. Note that this teacher becomes a perfectly rational and deterministic as  $\beta \rightarrow \infty$ , whereas  $\beta = 0$  produces uniformly random choices.

**Myopic behavior.** Humans are sometimes myopic (short-sighted), so a human teacher may remember and focus more on the behavior at the end of the clip they watched, for example. We model myopic behavior by introducing a weighted sum of rewards with a discount factor  $\gamma$  in (1), i.e.,  $\sum_{t=1}^H \gamma^{H-t} r(\mathbf{s}_t^i, \mathbf{a}_t^i)$ , so that our simulated teacher places more weight on recent timesteps.

**Skipping queries.** If both segments do not contain a desired behavior, a teacher would like to mark them as incomparable and discard the query. We model this behavior by skipping a query if the sum over the segment of an underlying reward is smaller than skip threshold  $\delta_{\text{skip}}$ , i.e.,  $\max_{i \in \{0,1\}} \sum_t r(\mathbf{s}_t^i, \mathbf{a}_t^i) < \delta_{\text{skip}}$  (Line 1).

**Equally preferable.** If the two segments are equally good, instead of selecting one segment as preferable, a teacher would like to mark the segments as equally preferable. Motivated by this, we provide an uniform distribution (0.5, 0.5) as a response (Line 3) if both segments have similar sum of rewards, i.e.,  $|\sum_t r(\mathbf{s}_t^1, \mathbf{a}_t^1) - \sum_t r(\mathbf{s}_t^0, \mathbf{a}_t^0)| < \delta_{\text{equal}}$ .

**Making a mistake.** Humans can make accidental errors when they respond. To reflect this, we flip the preference with probability of  $\epsilon$  (Line 7 and Line 10).

### 3.3 Evaluation metrics

We evaluate two key properties of preference-based RL: *performance of the RL agent under a fixed budget of feedback* and *robustness to potential irrationalities*. Because the simulated human teacher’s preferences are generated by a ground truth reward, we measure the true average return of trained agents as evaluation metric. To facilitate comparison across different RL algorithms, we normalize returns with respect to the baseline of RL training using the ground truth reward:

$$\text{Normalized returns} = \frac{\text{Average returns of preference-based RL}}{\text{Average returns of RL with ground truth reward}}.$$

To evaluate the feedback-efficiency of preference-based RL algorithms, we compare normalized returns by varying the maximum budget of queries.

To evaluate the robustness, we evaluate against the following simulated human teachers with different properties:

- Oracle:  $\text{SimTeacher}(\beta \rightarrow \infty, \gamma = 1, \epsilon = 0, \delta_{\text{skip}} = 0, \delta_{\text{equal}} = 0)$
- Stoc:  $\text{SimTeacher}(\beta = 1, \gamma = 1, \epsilon = 0, \delta_{\text{skip}} = 0, \delta_{\text{equal}} = 0)$
- Mistake:  $\text{SimTeacher}(\beta \rightarrow \infty, \gamma = 1, \epsilon = 0.1, \delta_{\text{skip}} = 0, \delta_{\text{equal}} = 0)$
- Skip:  $\text{SimTeacher}(\beta \rightarrow \infty, \gamma = 1, \epsilon = 0, \delta_{\text{skip}} > 0, \delta_{\text{equal}} = 0)$
- Equal:  $\text{SimTeacher}(\beta \rightarrow \infty, \gamma = 1, \epsilon = 0, \delta_{\text{skip}} = 0, \delta_{\text{equal}} > 0)$
- Myopic:  $\text{SimTeacher}(\beta \rightarrow \infty, \gamma = 0.9, \epsilon = 0, \delta_{\text{skip}} = 0, \delta_{\text{equal}} = 0)$

In our evaluations, we consider one modification (i.e., irrationality) to the oracle teacher at a time, which allows us to isolate the individual effects. While each individually may not exactly model real human behavior, it would be straightforward to use our benchmark to create more complex teachers that combine multiple irrationalities.

### 3.4 Tasks

We consider two locomotion tasks (Walker-walk and Quadruped-walk) from DeepMind Control Suite (DMControl) [65, 66] and two robotic manipulation tasks (Button Press and Sweep Into) from Meta-world [76]. We focus on learning from the proprioceptive inputs and dense rewards because learning from visual observations and sparse rewards can cause additional issues, such as representation learning [37, 55, 60, 62, 74] and exploration [56]. However, we think it is an interesting and important direction for future work to consider visual observations and sparse rewards.

## 4 B-Pref: Algorithmic baselines for preference-based RL

Throughout this paper, we mainly focus on two of the most prominent preference-based RL algorithms [19, 39], which involve reward learning from preferences. Formally, a policy  $\pi_\phi$  and reward function  $\hat{r}_\psi$  are updated as follows (see Algorithm 3 in the supplementary material):

- *Step 1 (agent learning)*: The policy  $\pi_\phi$  interacts with environment to collect experiences and we update it using existing RL algorithms to maximize the sum of the learned rewards  $\hat{r}_\psi$ .
- *Step 2 (reward learning)*: We optimize the reward function  $\hat{r}_\psi$  via supervised learning based on the feedback received from a teacher.
- Repeat *Step 1* and *Step 2*.

### 4.1 Deep reinforcement learning from human preferences

In order to incorporate human preferences into deep RL, Christiano et al. [19] proposed a framework that learns a reward function  $\hat{r}_\psi$  from preferences [50, 70]. Specifically, we first model a preference predictor using the reward function  $\hat{r}_\psi$  as follows:

$$P_\psi[\sigma^1 \succ \sigma^0] = \frac{\exp \sum_t \hat{r}_\psi(\mathbf{s}_t^1, \mathbf{a}_t^1)}{\sum_{i \in \{0,1\}} \exp \sum_t \hat{r}_\psi(\mathbf{s}_t^i, \mathbf{a}_t^i)}, \quad (2)$$

where  $\sigma^i \succ \sigma^j$  denotes the event that segment  $i$  is preferable to segment  $j$ . We remark that this corresponds to assume a stochastic teacher following the Bradley-Terry model [13] but we do not assume that the type and degree of irrationality or systematic bias is available in our experiments. Because this could lead to a poor preference inference [17], future work may be able to further improve the efficiency of learning by approximating the teacher’s irrationality.

To align our preference predictor with the teacher’s preferences, we consider a binary classification problem using the cross-entropy loss. Specifically, given a dataset of preferences  $\mathcal{D}$ , the reward function, modeled as a neural network with parameters  $\psi$ , is updated by minimizing the following loss:

$$\mathcal{L}^{\text{Reward}} = - \mathbb{E}_{(\sigma^0, \sigma^1, y) \sim \mathcal{D}} \left[ y(0) \log P_\psi[\sigma^0 \succ \sigma^1] + y(1) \log P_\psi[\sigma^1 \succ \sigma^0] \right]. \quad (3)$$

Once we learn a reward function  $\hat{r}_\psi$ , we can update the policy  $\pi_\phi$  using any RL algorithm. A caveat is that the reward function may be non-stationary because we update it during training. To mitigate the effects of a non-stationary reward function, Christiano et al. [19] used on-policy RL algorithms, such as TRPO [52] and A2C [45]. We re-implemented this method using the state-of-the-art on-policy RL algorithm: PPO [54]. We refer to this baseline as PrefPPO.

### 4.2 PEBBLE

PEBBLE [39] is a state-of-the-art preference-based RL algorithm that improved the framework of Christiano et al. [19] using the following ideas:

**Unsupervised pre-training.** In the beginning of training, a naive agent executing a random policy does not provide good state coverage nor coherent behaviors. Therefore, the agent’s queries are not diverse and a teacher can not convey much meaningful information. As a result, it requires many samples (and thus queries) for these methods to show initial progress. Ibarz et al. [31] has addressed this issue by assuming that demonstrations are available at the beginning of the experiment. However, this is not ideal since suitable demonstrations are often prohibitively expensive to obtain in practice. Instead, PEBBLE pre-trains the policy only using intrinsic motivation [46, 51] to learn how to generate diverse behaviors. Specifically, by updating the agent to maximize the state entropy  $\mathcal{H}(\mathbf{s}) = -\mathbb{E}_{\mathbf{s} \sim p(\mathbf{s})} [\log p(\mathbf{s})]$ , it encourages the agent to efficiently explore an environment and collect diverse experiences (see the supplementary material for more details).

**Off-policy RL with relabeling.** To overcome the poor sample-efficiency of on-policy RL algorithms, PEBBLE used the state-of-the-art off-policy RL algorithm: SAC [27]. However, the learning process can be unstable because previous experiences in the replay buffer are labeled with previous learned rewards. PEBBLE stabilizes the learning process by relabeling all of the agent’s past experience every time it updates the reward model.

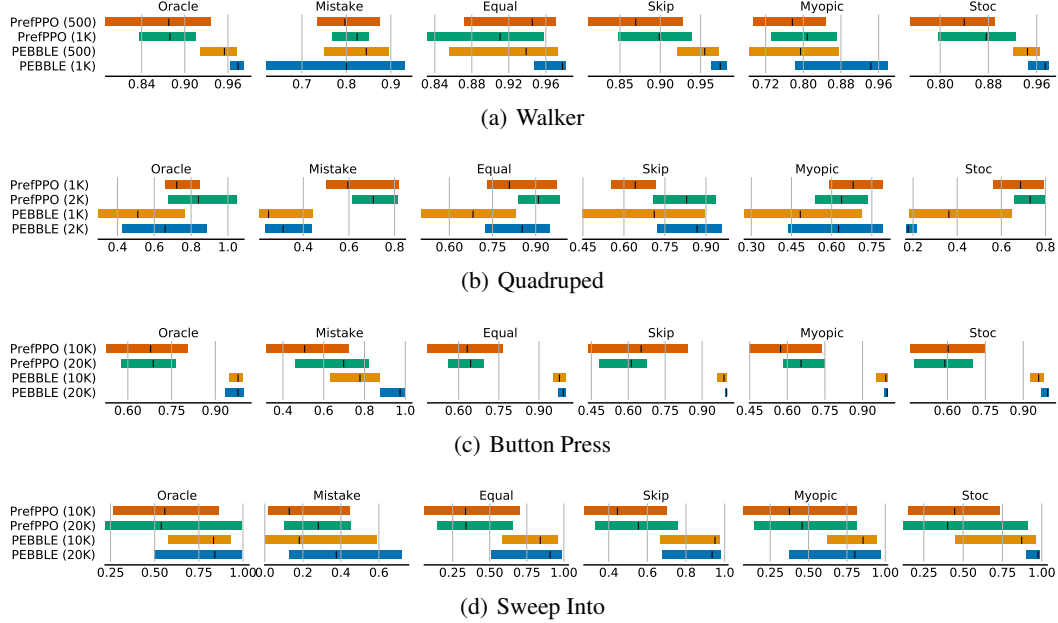


Figure 2: IQM normalized returns with 95% confidence intervals across ten runs. Learning curves and other metrics (median, mean, optimality gap) are in the supplementary material.

## 5 Using B-Pref to analyze algorithmic design decisions

We design our experiments to investigate the following:

- How do existing preference-based RL methods compare against each other across environments with different complexity?
- How to use B-Pref to analyze algorithmic design decisions for preference-based RL?

### 5.1 Training details

We implement PEBBLE and PrefPPO using publicly released implementations of SAC<sup>1</sup> and PPO.<sup>2</sup> All hyperparameters of all algorithms are optimized independently for each environment. All of the experiments were processed using a single GPU (NVIDIA GTX 1080 Ti) and 8 CPU cores (Intel Xeon Gold 6126). For reliable evaluation Agarwal et al. [1], we measure the normalized returns<sup>3</sup> and report the interquartile mean (IQM) across ten runs using an open-source library *rlible*.<sup>4</sup> More experimental details (e.g., model architectures and the final hyperparameters) and all learning curves with standard deviation are in the supplementary material.

### 5.2 Benchmarking prior methods

Figure 2 shows the IQM normalized returns of PEBBLE and PrefPPO at convergence on various simulated teachers listed in Section 3.2 (see the supplementary material for experimental details). For a fair comparison, we apply unsupervised pre-training and disagreement-based sampling to all methods (including SAC and PPO). PEBBLE outperforms PrefPPO in most of the environments (especially achieving large gains on robotic manipulation tasks). Interestingly, providing uniform labels to equally preferable segments (Equal) or skipping the queries with similar behaviors (Skip) is more useful than relying only on perfect labels (Oracle) on hard environments like Quadrupe (Figure 2(b)). While both PEBBLE and PrefPPO achieve fairly efficient performance on correct

<sup>1</sup>[https://github.com/denisyarats/pytorch\\_sac](https://github.com/denisyarats/pytorch_sac)

<sup>2</sup><https://github.com/DLR-RM/stable-baselines3>

<sup>3</sup>On robotic manipulation tasks, we measure the task success rate as defined by the Meta-world authors [76].

<sup>4</sup><https://github.com/google-research/rlible>



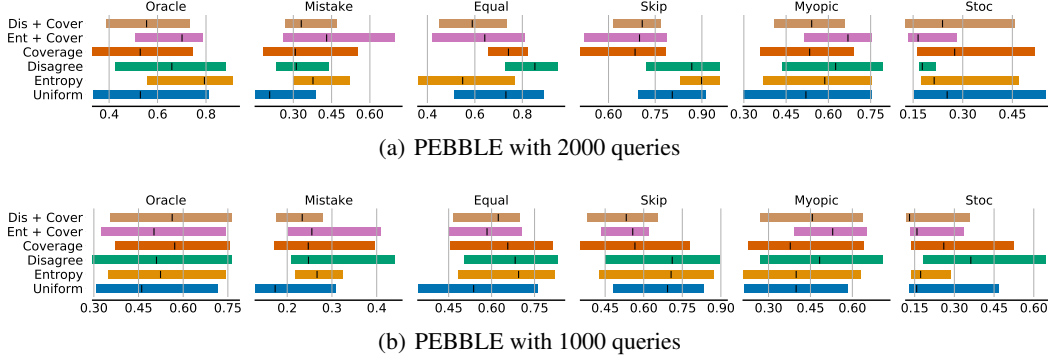


Figure 3: IQM normalized returns of PEBBLE with various sampling schemes across ten runs on Quadrupted. Learning curves and other metrics (median, mean, optimality gap) are in the supplementary material.

labels (Oracle, Equal and Skip), they often suffer from poor performance when teachers can provide the wrong labels (Mistake and Stoc). This suggests opportunities for further investigations and development of techniques that can improve the robustness to corrupted labels.<sup>5</sup>

### 5.3 Impact of design decisions in reward learning

Reward learning from preferences involves several design decisions, which can affect the performance of the overall framework. We showcase the utility of B-Pref by using it to analyze the following algorithmic design choices in depth:

**Selecting informative queries.** During training, all experiences are stored in an annotation buffer  $\mathcal{B}$  and we generate  $N_{\text{query}}$  pairs of segments<sup>6</sup> to ask teacher’s preferences from this buffer at each feedback session. To reduce the burden on the human, we should solicit preferences so as to maximize the information received. While finding optimal queries is computationally intractable [2], several sampling schemes [11, 12, 50] have been explored to find queries that are likely to change the reward model. Specifically, we consider the following sampling schemes, where more details are in the supplementary material:

- *Uniform sampling:* We pick  $N_{\text{query}}$  pairs of segments uniformly at random from the buffer  $\mathcal{B}$ .
- *Uncertainty-based sampling:* We first generate the initial batch of  $N_{\text{init}}$  pairs of segments  $\mathcal{G}_{\text{init}}$  uniformly at random, measure the uncertainty (e.g., variance across ensemble of preference predictors [19] or entropy of a single preference predictor [39]), and then select the  $N_{\text{query}}$  pairs of segments with high uncertainty.
- *Coverage-based sampling:* From the initial batch  $\mathcal{G}_{\text{init}}$ , we choose  $N_{\text{query}}$  center points such that the largest distance between a data point and its nearest center is minimized using a greedy selection strategy.
- *Hybrid sampling:* Similar to Yu et al. [75], we also consider hybrid sampling, which combines uncertainty-based sampling and coverage-based sampling. First, we select the  $N_{\text{inter}}$  pairs of segments  $\mathcal{G}_{\text{un}}$ , using uncertainty-based sampling, where  $N_{\text{init}} > N_{\text{inter}}$ , and then choose  $N_{\text{query}}$  center points from  $\mathcal{G}_{\text{un}}$ .

Figure 3 shows the IQM normalized returns of PEBBLE with various sampling schemes on Quadrupted. We find that the uncertainty-based sampling schemes (i.e., ensemble disagreement and entropy) are superior to other sampling schemes, while coverage-based sampling schemes do not improve on uniform sampling and slow down the sampling procedures. To analyze the effects of sampling schemes, we measure the fraction of equally preferable queries (i.e.,  $y = (0.5, 0.5)$ ) on

<sup>5</sup>We find that label smoothing [64] is not effective in handling corrupted labels in our experiments (see the supplementary material for supporting results). However, other regularization techniques, like label flipping, L2 regularization and weight decay, would be interesting for further study in future work.

<sup>6</sup>We do not compare segments of different lengths.

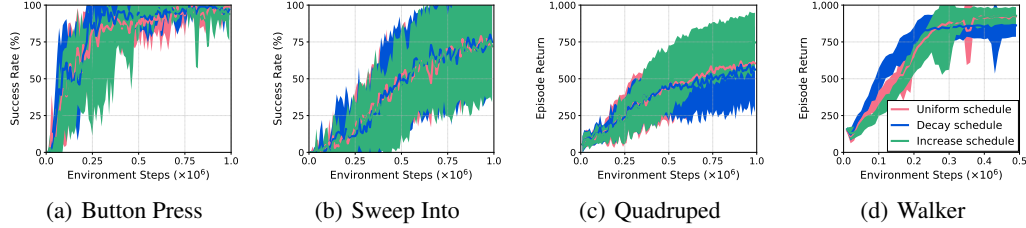


Figure 4: Learning curves of PEBBLE with different feedback schedules on the oracle teacher. The solid line and shaded regions represent the mean and standard deviation, respectively, across ten runs.

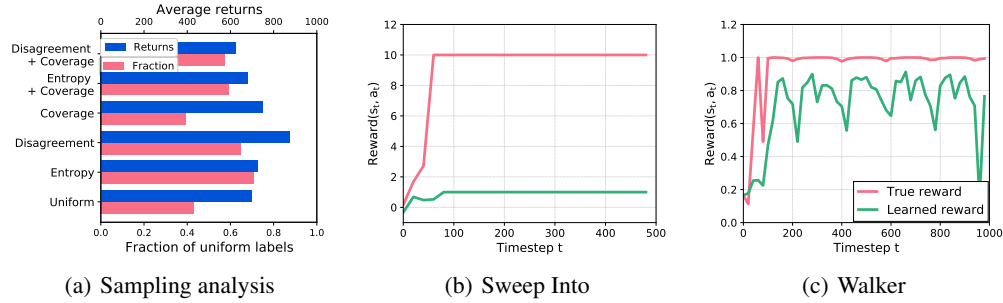


Figure 5: (a) Fraction of equally preferable queries (red) and average returns (blue) on the Equal teacher. We use PEBBLE with different sampling schemes on Quadrupted given a budget of 2000 queries. Even though a teacher provides more uniform labels, i.e.,  $y = (0.5, 0.5)$ , to uncertainty-based sampling schemes, they achieve higher returns than other sampling schemes. (b/c) Time series of learned reward function (green) and the ground truth reward (red) using rollouts from a policy optimized by PEBBLE. Learned reward functions align with the ground truth rewards in (b) Sweep Into and (c) Walker.

the Equal teacher. Figure 5(a) shows that uncertainty-based sampling schemes achieve high returns even though other sampling schemes receive more (non-uniform) perfect labels. We expect that this is because queries with high uncertainty provide significant information to the reward model. This also suggests opportunities for further investigations on uncertainty estimates like Bayesian methods [15, 25].

**Feedback schedule.** We also investigate the impact of the feedback schedule, which decides the number of queries at each feedback session. Lee et al. [39] used a uniform schedule, which always asks the same number of queries, and Christiano et al. [19], Ibarz et al. [31] used a decay schedule, which decreases the number of queries, roughly proportional to  $\frac{T}{t+T}$ , where  $t$  is the current timestep and  $T$  is the episode length. We additionally consider an increase schedule, which increases the number of queries, roughly proportional to  $\frac{T+t}{T}$ .

Figure 4 shows the learning curves of PEBBLE with different feedback schedule on the oracle teacher. Given the same total number of queries, increase and decay schedules change the size of the initial queries by a factor of 0.5 and 2, respectively. One can note that there is no big gain from rule-based schedules in most of the environments. Even though rule-based schedules are less effective than uniform scheduling, using an adaptive schedule like meta-gradient [72, 73] would be interesting for further study in future work.

**Reward analysis.** To investigate the quality of the learned reward function, we compare the learned reward function with the ground truth reward. Figure 5(b) and Figure 5(c) show the learned reward function optimized by PEBBLE on the oracle teacher in Sweep Into and Walker, where more evaluation results on other environments are also available in the supplementary material. Because we bound the output of the reward function using tanh function, the scale is different with the ground truth reward but the learned reward function is reasonably well-aligned.



## 6 Related work

**Benchmarks for deep reinforcement learning.** There is a large body of work focused on designing benchmarks for RL [7, 8, 14, 20–22, 24, 26, 30, 48, 65, 66, 76]. The Arcade Learning Environment [8] has become a popular benchmark to measure the progress of RL algorithms for discrete control tasks. For continuous control tasks, Duan et al. [22] presented a benchmark with baseline implementations of various RL algorithms, which in turn led to OpenAI Gym [14]. These benchmarks have significantly accelerated progress and have been strong contributors towards the discovery and evaluation of today’s most widely used RL algorithms [27, 44, 52–54].

Recently, researchers proposed more targeted RL benchmarks that have been designed for specific research purposes. Cobbe et al. [21] presented a suite of game-like environments where the train and test environments differ for evaluating generalization performance of RL agents. Ray et al. [48] provided a Safety Gym for measuring progress towards RL agents that satisfy the safety constraints. D4RL [24] and RL Unplugged [26] have been proposed to evaluate and compare offline RL algorithms. Yu et al. [76] proposed Meta-world to study meta- and multi-task RL. URLB [38] benchmarks performance of unsupervised RL methods. However, none of the existing RL benchmarks are tailored towards preference-based RL.

Freire et al. [23] proposed DERAILED, a benchmark suite for preference-based learning, but they focused on simple diagnostic tasks. In B-Pref, we consider learning a variety of complex locomotion and robotic manipulation tasks. Additionally, we design teachers with a wide array of irrationalities and benchmark state-of-the-art preference-based RL algorithms [19, 39] in depth.

**Human-in-the-loop reinforcement learning.** Several works have successfully utilized feedback from real humans to train RL agents [6, 19, 31, 34, 39, 43, 69]. MacGlashan et al. [43] proposed a reward-free method, which utilizes a human feedback as an advantage function and optimizes the agents via a policy gradient. Knox & Stone [34] trained a reward model via regression using unbounded real-valued feedback. However, these approaches are difficult to scale to more complex learning problems that require substantial agent experience.

Another promising direction has focused on utilizing the human preferences [3, 19, 31, 39, 41, 47, 61, 70, 71]. Christiano et al. [19] scaled preference-based learning to utilize modern deep learning techniques, and Ibarz et al. [31] improved the efficiency of this method by introducing additional forms of feedback such as demonstrations. Recently, Lee et al. [39] proposed a feedback-efficient RL algorithm by utilizing off-policy learning and pre-training. Stiennon et al. [61] and Wu et al. [71] showed that preference-based RL can be utilized to fine-tune GPT-3 [16] for hard tasks like text and book summarization, respectively. We benchmark these state-of-the-art preference-based RL algorithms in this paper.

## 7 Conclusion

In this paper, we present B-Pref, a benchmark specially designed for preference-based RL, covering a wide array of a teacher’s irrationalities. We empirically investigate state-of-the-art preference-based RL algorithms in depth and analyze the effects of algorithmic design decisions on our benchmark. We find that existing methods often suffer from poor performance when teachers provide wrong labels, and the effects of design decisions are varied depending on the task setups. These observations call for new algorithms in active learning [11, 12, 50] and meta-learning [72, 73] to be developed. By providing an open-source release of the benchmark, we encourage other researchers to use B-Pref as a common starting point to study preference-based RL more systematically.

**Limitations.** There are several important properties that are not explored in-depth in B-Pref. One is robustness of learned reward functions to new environments with different dynamics or initial states [49]. Also, we focus on tasks with proprioceptive inputs and dense rewards, but extensions to visual observations and sparse rewards are interesting directions to explore.

**Potential negative impacts.** Preference-based RL has several advantages (e.g., teaching novel behaviors, and mitigating the effects of reward exploitation); however, it also has potential drawbacks. Malicious users might teach the bad behaviors/functionality using this framework. Therefore, researchers should consider the safety issues with particular thought.

## Acknowledgements

This research is supported in part by ONR PECASE N000141612723, NSF NRI #2024675, ONR YIP, and Berkeley Deep Drive. Laura Smith was supported by NSF Graduate Research Fellowship. We also thank anonymous reviewers for critically reading the manuscript and suggesting substantial improvements.

## References

- [1] Agarwal, Rishabh, Schwarzer, Max, Castro, Pablo Samuel, Courville, Aaron, and Bellemare, Marc G. Deep reinforcement learning at the edge of the statistical precipice. In *Advances in Neural Information Processing Systems*, 2021.
- [2] Ailon, Nir. An active learning algorithm for ranking from pairwise preferences with an almost optimal query complexity. *Journal of Machine Learning Research*, 13(1), 2012.
- [3] Akrou, Riad, Schoenauer, Marc, and Sebag, Michele. Preference-based policy learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 2011.
- [4] Amodei, Dario, Olah, Chris, Steinhardt, Jacob, Christiano, Paul, Schulman, John, and Mané, Dan. Concrete problems in ai safety. *arXiv preprint arXiv:1606.06565*, 2016.
- [5] Andrychowicz, OpenAI: Marcin, Baker, Bowen, Chociej, Maciek, Jozefowicz, Rafal, McGrew, Bob, Pachocki, Jakub, Petron, Arthur, Plappert, Matthias, Powell, Glenn, Ray, Alex, et al. Learning dexterous in-hand manipulation. *The International Journal of Robotics Research*, 39(1):3–20, 2020.
- [6] Arumugam, Dilip, Lee, Jun Ki, Saskin, Sophie, and Littman, Michael L. Deep reinforcement learning from policy-dependent human feedback. *arXiv preprint arXiv:1902.04257*, 2019.
- [7] Beattie, Charles, Leibo, Joel Z, Teplyashin, Denis, Ward, Tom, Wainwright, Marcus, Küttler, Heinrich, Lefrancq, Andrew, Green, Simon, Valdés, Víctor, Sadik, Amir, et al. Deepmind lab. *arXiv preprint arXiv:1612.03801*, 2016.
- [8] Bellemare, Marc G, Naddaf, Yavar, Veness, Joel, and Bowling, Michael. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013.
- [9] Bellemare, Marc G, Candido, Salvatore, Castro, Pablo Samuel, Gong, Jun, Machado, Marlos C, Moitra, Subhodeep, Ponda, Sameera S, and Wang, Ziyu. Autonomous navigation of stratospheric balloons using reinforcement learning. *Nature*, 588(7836):77–82, 2020.
- [10] Berner, Christopher, Brockman, Greg, Chan, Brooke, Cheung, Vicki, Dębiak, Przemysław, Dennison, Christy, Farhi, David, Fischer, Quirin, Hashme, Shariq, Hesse, Chris, et al. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*, 2019.
- [11] Biyik, Erdem and Sadigh, Dorsa. Batch active preference-based learning of reward functions. In *Conference on Robot Learning*, 2018.
- [12] Biyik, Erdem, Huynh, Nicolas, Kochenderfer, Mykel J, and Sadigh, Dorsa. Active preference-based gaussian process regression for reward learning. In *Robotics: Science and Systems*, 2020.
- [13] Bradley, Ralph Allan and Terry, Milton E. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.
- [14] Brockman, Greg, Cheung, Vicki, Pettersson, Ludwig, Schneider, Jonas, Schulman, John, Tang, Jie, and Zaremba, Wojciech. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- [15] Brown, Daniel S and Niekum, Scott. Deep bayesian reward learning from preferences. *arXiv preprint arXiv:1912.04472*, 2019.

- [16] Brown, Tom B, Mann, Benjamin, Ryder, Nick, Subbiah, Melanie, Kaplan, Jared, Dhariwal, Prafulla, Neelakantan, Arvind, Shyam, Pranav, Sastry, Girish, Askell, Amanda, et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- [17] Chan, Lawrence, Critch, Andrew, and Dragan, Anca. The impacts of known and unknown demonstrator irrationality on reward inference, 2021. URL <https://openreview.net/forum?id=CzRSsOG6JDw>.
- [18] Chipman, Susan EF. *The Oxford handbook of cognitive science*. Oxford University Press, 2016.
- [19] Christiano, Paul F, Leike, Jan, Brown, Tom, Martic, Miljan, Legg, Shane, and Amodei, Dario. Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems*, 2017.
- [20] Cobbe, Karl, Klimov, Oleg, Hesse, Chris, Kim, Taehoon, and Schulman, John. Quantifying generalization in reinforcement learning. In *International Conference on Machine Learning*, 2019.
- [21] Cobbe, Karl, Hesse, Chris, Hilton, Jacob, and Schulman, John. Leveraging procedural generation to benchmark reinforcement learning. In *International conference on machine learning*, 2020.
- [22] Duan, Yan, Chen, Xi, Houthooft, Rein, Schulman, John, and Abbeel, Pieter. Benchmarking deep reinforcement learning for continuous control. In *International conference on machine learning*, 2016.
- [23] Freire, Pedro, Gleave, Adam, Toyer, Sam, and Russell, Stuart. Derail: Diagnostic environments for reward and imitation learning. *arXiv preprint arXiv:2012.01365*, 2020.
- [24] Fu, Justin, Kumar, Aviral, Nachum, Ofir, Tucker, George, and Levine, Sergey. D4rl: Datasets for deep data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219*, 2020.
- [25] Gal, Yarin and Ghahramani, Zoubin. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *International Conference on Machine Learning*, 2016.
- [26] Gulcehre, Caglar, Wang, Ziyu, Novikov, Alexander, Paine, Tom Le, Colmenarejo, Sergio Gomez, Zolna, Konrad, Agarwal, Rishabh, Merel, Josh, Mankowitz, Daniel, Paduraru, Cosmin, et al. RL unplugged: A suite of benchmarks for offline reinforcement learning. In *Advances in Neural Information Processing Systems*, 2020.
- [27] Haarnoja, Tuomas, Zhou, Aurick, Abbeel, Pieter, and Levine, Sergey. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International Conference on Machine Learning*, 2018.
- [28] Hadfield-Menell, Dylan, Milli, Smitha, Abbeel, Pieter, Russell, Stuart, and Dragan, Anca. Inverse reward design. In *Advances in Neural Information Processing Systems*, 2017.
- [29] Hazan, Elad, Kakade, Sham, Singh, Karan, and Van Soest, Abby. Provably efficient maximum entropy exploration. In *International Conference on Machine Learning*, 2019.
- [30] Henderson, Peter, Islam, Riashat, Bachman, Philip, Pineau, Joelle, Precup, Doina, and Meger, David. Deep reinforcement learning that matters. In *AAAI Conference on Artificial Intelligence*, 2018.
- [31] Ibarz, Borja, Leike, Jan, Pohlen, Tobias, Irving, Geoffrey, Legg, Shane, and Amodei, Dario. Reward learning from human preferences and demonstrations in atari. In *Advances in Neural Information Processing Systems*, 2018.
- [32] Kalashnikov, Dmitry, Irpan, Alex, Pastor, Peter, Ibarz, Julian, Herzog, Alexander, Jang, Eric, Quillen, Deirdre, Holly, Ethan, Kalakrishnan, Mrinal, Vanhoucke, Vincent, et al. Qt-opt: Scalable deep reinforcement learning for vision-based robotic manipulation. In *Conference on Robot Learning*, 2018.
- [33] Kingma, Diederik P and Ba, Jimmy. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.

- [34] Knox, W Bradley and Stone, Peter. Interactively shaping agents via human reinforcement: The tamer framework. In *International Conference on Knowledge Capture*, 2009.
- [35] Kober, Jens and Peters, Jan. Policy search for motor primitives in robotics. *Machine learning*, 84(1-2):171–203, 2011.
- [36] Kober, Jens, Bagnell, J Andrew, and Peters, Jan. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11):1238–1274, 2013.
- [37] Laskin, Michael, Lee, Kimin, Stooke, Adam, Pinto, Lerrel, Abbeel, Pieter, and Srinivas, Aravind. Reinforcement learning with augmented data. In *Advances in Neural Information Processing Systems*, 2020.
- [38] Laskin, Michael, Yarats, Denis, Liu, Hao, Lee, Kimin, Zhan, Albert, Lu, Kevin, Cang, Catherine, Pinto, Lerrel, and Abbeel, Pieter. URLB: Unsupervised reinforcement learning benchmark. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- [39] Lee, Kimin, Smith, Laura, and Abbeel, Pieter. Pebble: Feedback-efficient interactive reinforcement learning via relabeling experience and unsupervised pre-training. In *International Conference on Machine Learning*, 2021.
- [40] Lee, Lisa, Eysenbach, Benjamin, Parisotto, Emilio, Xing, Eric, Levine, Sergey, and Salakhutdinov, Ruslan. Efficient exploration via state marginal matching. *arXiv preprint arXiv:1906.05274*, 2019.
- [41] Leike, Jan, Krueger, David, Everitt, Tom, Martic, Miljan, Maini, Vishal, and Legg, Shane. Scalable agent alignment via reward modeling: a research direction. *arXiv preprint arXiv:1811.07871*, 2018.
- [42] Liu, Hao and Abbeel, Pieter. Behavior from the void: Unsupervised active pre-training. *arXiv preprint arXiv:2103.04551*, 2021.
- [43] MacGlashan, James, Ho, Mark K, Loftin, Robert, Peng, Bei, Roberts, David, Taylor, Matthew E, and Littman, Michael L. Interactive learning from policy-dependent human feedback. In *International Conference on Machine Learning*, 2017.
- [44] Mnih, Volodymyr, Kavukcuoglu, Koray, Silver, David, Rusu, Andrei A, Veness, Joel, Bellemare, Marc G, Graves, Alex, Riedmiller, Martin, Fidjeland, Andreas K, Ostrovski, Georg, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.
- [45] Mnih, Volodymyr, Badia, Adria Puigdomenech, Mirza, Mehdi, Graves, Alex, Lillicrap, Timothy, Harley, Tim, Silver, David, and Kavukcuoglu, Koray. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, 2016.
- [46] Oudeyer, Pierre-Yves, Kaplan, Frdric, and Hafner, Verena V. Intrinsic motivation systems for autonomous mental development. *IEEE transactions on evolutionary computation*, 11(2): 265–286, 2007.
- [47] Pilarski, Patrick M, Dawson, Michael R, Degris, Thomas, Fahimi, Farbod, Carey, Jason P, and Sutton, Richard S. Online human training of a myoelectric prosthesis controller via actor-critic reinforcement learning. In *International Conference on Rehabilitation Robotics*, 2011.
- [48] Ray, Alex, Achiam, Joshua, and Amodei, Dario. Benchmarking safe exploration in deep reinforcement learning. *arXiv preprint arXiv:1910.01708*, 2019.
- [49] Reddy, Siddharth, Dragan, Anca, Levine, Sergey, Legg, Shane, and Leike, Jan. Learning human objectives by evaluating hypothetical behavior. In *International Conference on Machine Learning*, 2020.
- [50] Sadigh, Dorsa, Dragan, Anca D, Sastry, Shankar, and Seshia, Sanjit A. Active preference-based learning of reward functions. In *Robotics: Science and Systems*, 2017.
- [51] Schmidhuber, Jürgen. Formal theory of creativity, fun, and intrinsic motivation (1990–2010). *IEEE Transactions on Autonomous Mental Development*, 2(3):230–247, 2010.

- [52] Schulman, John, Levine, Sergey, Abbeel, Pieter, Jordan, Michael, and Moritz, Philipp. Trust region policy optimization. In *International Conference on Machine Learning*, 2015.
- [53] Schulman, John, Moritz, Philipp, Levine, Sergey, Jordan, Michael, and Abbeel, Pieter. High-dimensional continuous control using generalized advantage estimation. In *International Conference on Learning Representations*, 2016.
- [54] Schulman, John, Wolski, Filip, Dhariwal, Prafulla, Radford, Alec, and Klimov, Oleg. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [55] Schwarzer, Max, Anand, Ankesh, Goel, Rishab, Hjelm, R Devon, Courville, Aaron, and Bachman, Philip. Data-efficient reinforcement learning with self-predictive representations. In *International Conference on Learning Representations*, 2021.
- [56] Seo, Younggyo, Chen, Lili, Shin, Jinwoo, Lee, Honglak, Abbeel, Pieter, and Lee, Kimin. State entropy maximization with random encoders for efficient exploration. In *International Conference on Machine Learning*, 2021.
- [57] Shah, Rohin, Krashennnikov, Dmitrii, Alexander, Jordan, Abbeel, Pieter, and Dragan, Anca. Preferences implicit in the state of the world. In *International Conference on Learning Representations*, 2019.
- [58] Silver, David, Schrittwieser, Julian, Simonyan, Karen, Antonoglou, Ioannis, Huang, Aja, Guez, Arthur, Hubert, Thomas, Baker, Lucas, Lai, Matthew, Bolton, Adrian, et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354, 2017.
- [59] Silver, David, Hubert, Thomas, Schrittwieser, Julian, Antonoglou, Ioannis, Lai, Matthew, Guez, Arthur, Lanctot, Marc, Sifre, Laurent, Kumaran, Dharshan, Graepel, Thore, et al. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419):1140–1144, 2018.
- [60] Srinivas, Aravind, Laskin, Michael, and Abbeel, Pieter. Curl: Contrastive unsupervised representations for reinforcement learning. In *International Conference on Machine Learning*, 2020.
- [61] Stiennon, Nisan, Ouyang, Long, Wu, Jeff, Ziegler, Daniel M, Lowe, Ryan, Voss, Chelsea, Radford, Alec, Amodei, Dario, and Christiano, Paul. Learning to summarize from human feedback. In *Advances in Neural Information Processing Systems*, 2020.
- [62] Stooke, Adam, Lee, Kimin, Abbeel, Pieter, and Laskin, Michael. Decoupling representation learning from reinforcement learning. In *International Conference on Machine Learning*, 2021.
- [63] Sutton, Richard S and Barto, Andrew G. *Reinforcement learning: An introduction*. MIT Press, 2018.
- [64] Szegedy, Christian, Vanhoucke, Vincent, Ioffe, Sergey, Shlens, Jon, and Wojna, Zbigniew. Rethinking the inception architecture for computer vision. In *Conference on Computer Vision and Pattern Recognition*, 2016.
- [65] Tassa, Yuval, Doron, Yotam, Muldal, Alistair, Erez, Tom, Li, Yazhe, Casas, Diego de Las, Budden, David, Abdolmaleki, Abbas, Merel, Josh, Lefrancq, Andrew, et al. Deepmind control suite. *arXiv preprint arXiv:1801.00690*, 2018.
- [66] Tassa, Yuval, Tunyasuvunakool, Saran, Muldal, Alistair, Doron, Yotam, Liu, Siqi, Bohez, Steven, Merel, Josh, Erez, Tom, Lillicrap, Timothy, and Heess, Nicolas. dm\_control: Software and tasks for continuous control. *arXiv preprint arXiv:2006.12983*, 2020.
- [67] Turner, Alexander Matt, Ratzlaff, Neale, and Tadepalli, Prasad. Avoiding side effects in complex environments. *arXiv preprint arXiv:2006.06547*, 2020.
- [68] Vinyals, Oriol, Babuschkin, Igor, Czarnecki, Wojciech M, Mathieu, Michael, Dudzik, Andrew, Chung, Junyoung, Choi, David H, Powell, Richard, Ewalds, Timo, Georgiev, Petko, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782): 350–354, 2019.

- [69] Warnell, Garrett, Waytowich, Nicholas, Lawhern, Vernon, and Stone, Peter. Deep tamer: Interactive agent shaping in high-dimensional state spaces. In *Conference on Artificial Intelligence*, 2018.
- [70] Wilson, Aaron, Fern, Alan, and Tadepalli, Prasad. A bayesian approach for policy learning from trajectory preference queries. In *Advances in Neural Information Processing Systems*, 2012.
- [71] Wu, Jeff, Ouyang, Long, Ziegler, Daniel M, Stiennon, Nissan, Lowe, Ryan, Leike, Jan, and Christiano, Paul. Recursively summarizing books with human feedback. *arXiv preprint arXiv:2109.10862*, 2021.
- [72] Xu, Zhongwen, van Hasselt, Hado, and Silver, David. Meta-gradient reinforcement learning. In *Advances in Neural Information Processing Systems*, 2018.
- [73] Xu, Zhongwen, van Hasselt, Hado, Hessel, Matteo, Oh, Junhyuk, Singh, Satinder, and Silver, David. Meta-gradient reinforcement learning with an objective discovered online. In *Advances in Neural Information Processing Systems*, 2020.
- [74] Yarats, Denis, Kostrikov, Ilya, and Fergus, Rob. Image augmentation is all you need: Regularizing deep reinforcement learning from pixels. In *International Conference on Learning Representations*, 2021.
- [75] Yu, Kai, Bi, Jinbo, and Tresp, Volker. Active learning via transductive experimental design. In *International Conference on Machine Learning*, 2006.
- [76] Yu, Tianhe, Quillen, Deirdre, He, Zhanpeng, Julian, Ryan, Hausman, Karol, Finn, Chelsea, and Levine, Sergey. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on Robot Learning*, 2020.
- [77] Ziebart, Brian D. Modeling purposeful adaptive behavior with the principle of maximum causal entropy. 2010.

## Checklist

1. For all authors...
  - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [\[Yes\]](#)
  - (b) Did you describe the limitations of your work? [\[Yes\]](#) See Section 7.
  - (c) Did you discuss any potential negative societal impacts of your work? [\[Yes\]](#) See Section 7.
  - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [\[Yes\]](#)
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 (e.g. for benchmarks)...
  - (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 Section 5 and the supplementary material.
  - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [\[Yes\]](#) See Section 5 and the supplementary material.
  - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [\[Yes\]](#) See Section 5 and the supplementary material.
  - (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 and the supplementary material.
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]
  - (b) Did you mention the license of the assets? [Yes]
  - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
  - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
  - (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]

# Appendix:

## B-Pref: Benchmarking Preference-Based Reinforcement Learning

### A Preliminaries: Reinforcement learning algorithms

**Proximal policy optimization.** Proximal policy optimization (PPO) [54] is a state-of-the-art on-policy algorithm for learning a continuous or discrete control policy,  $\pi_\phi(\mathbf{a}|\mathbf{s})$ . PPO forms policy gradients using action-advantages,  $A_t = A^\pi(\mathbf{a}_t, \mathbf{s}_t) = Q^\theta(\mathbf{a}_t, \mathbf{s}_t) - V^\pi(\mathbf{s}_t)$ , and minimizes a clipped-ratio loss over minibatches of recent experience (collected under  $\pi_{\bar{\phi}}$ ):

$$\mathcal{L}_\pi^{\text{PPO}} = -\mathbb{E}_{\tau_t \sim \pi} [\min(\rho_t(\phi)A_t, \text{clip}(\rho_t(\phi), 1 - \epsilon, 1 + \epsilon)A_t)], \quad \rho_t(\phi) = \frac{\pi_\phi(\mathbf{a}_t|\mathbf{s}_t)}{\pi_{\bar{\phi}_{old}}(\mathbf{a}_t|\mathbf{s}_t)}, \quad (4)$$

where  $\bar{\phi}$  are the delayed parameters and  $\epsilon$  is a clip ratio. Our PPO agents learn a state-value estimator,  $V_\theta(\mathbf{s})$ , which is regressed against a target of discounted returns and used with Generalized Advantage Estimation [53]:

$$\mathcal{L}_V^{\text{PPO}}(\theta) = \mathbb{E}_{\tau_t \sim \pi} [(V_\theta(\mathbf{s}_t) - V_{\bar{\theta}}(\mathbf{s}_t))^2]. \quad (5)$$

PPO is more robust to the non-stationarity in rewards caused by online learning.

**Soft actor-critic.** Soft actor-critic (SAC) [27] is an off-policy actor-critic method based on the maximum entropy RL framework [77], which encourages exploration and greater robustness to noise by maximizing a weighted objective of the reward and the policy entropy. To update the parameters, SAC alternates between a soft policy evaluation and a soft policy improvement. At the soft policy evaluation step, a soft Q-function, which is modeled as a neural network with parameters  $\theta$ , is updated by minimizing the following soft Bellman residual:

$$\mathcal{L}_Q^{\text{SAC}} = \mathbb{E}_{\tau_t \sim \mathcal{B}} [(Q_\theta(\mathbf{s}_t, \mathbf{a}_t) - r_t - \gamma \bar{V}(\mathbf{s}_{t+1}))^2], \quad (6)$$

with  $\bar{V}(\mathbf{s}_t) = \mathbb{E}_{\mathbf{a}_t \sim \pi_{\bar{\phi}}} [Q_{\bar{\theta}}(\mathbf{s}_t, \mathbf{a}_t) - \alpha \log \pi_{\bar{\phi}}(\mathbf{a}_t|\mathbf{s}_t)],$

where  $\tau_t = (\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1}, r_t)$  is a transition,  $\mathcal{B}$  is a replay buffer,  $\bar{\theta}$  are the delayed parameters, and  $\alpha$  is a temperature parameter. At the soft policy improvement step, the policy  $\pi_\phi$  is updated by minimizing the following objective:

$$\mathcal{L}_\pi^{\text{SAC}} = \mathbb{E}_{\mathbf{s}_t \sim \mathcal{B}, \mathbf{a}_t \sim \pi_\phi} [\alpha \log \pi_\phi(\mathbf{a}_t|\mathbf{s}_t) - Q_\theta(\mathbf{s}_t, \mathbf{a}_t)]. \quad (7)$$

SAC enjoys good sample-efficiency relative to its on-policy counterparts by reusing its past experiences. However, for the same reason, SAC is not robust to a non-stationary reward function.

---

#### Algorithm 2 EXPLORE: Unsupervised exploration

---

```

1: Initialize parameters of  $\pi_\phi$  and a buffer  $\mathcal{B} \leftarrow \emptyset$ 
2: for each iteration do
3:   for each timestep  $t$  do
4:     Collect  $\mathbf{s}_{t+1}$  by taking  $\mathbf{a}_t \sim \pi_\phi(\mathbf{a}_t|\mathbf{s}_t)$ 
5:     Compute intrinsic reward  $r_t^{\text{int}} \leftarrow r^{\text{int}}(\mathbf{s}_t)$  as in (8)
6:     Store transitions  $\mathcal{B} \leftarrow \mathcal{B} \cup \{(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1}, r_t^{\text{int}})\}$ 
7:   end for
8:   for each gradient step do
9:     Sample minibatch  $\{(\mathbf{s}_j, \mathbf{a}_j, \mathbf{s}_{j+1}, r_j^{\text{int}})\}_{j=1}^B \sim \mathcal{B}$ 
10:    Optimize RL objective function with respect to  $\phi$ 
11:   end for
12: end for
13: return  $\mathcal{B}, \pi_\phi$ 

```

---

---

**Algorithm 3** Preference-based RL with reward learning

---

**Require:** frequency of teacher feedback  $K$

**Require:** number of queries  $N_{\text{query}}$  per feedback session

```
1: Initialize parameters of  $\pi_\phi, \hat{r}_\psi$ , a dataset of preferences  $\mathcal{D} \leftarrow \emptyset$ , and a buffer  $\mathcal{B} \leftarrow \emptyset$ 
2: // EXPLORATION PHASE
3:  $\mathcal{B}, \pi_\phi \leftarrow \text{EXPLORE}()$  in Algorithm 2
4: for each iteration do
5:   // REWARD LEARNING
6:   if iteration %  $K == 0$  then
7:     for  $m$  in  $1 \dots N_{\text{query}}$  do
8:        $(\sigma^0, \sigma^1) \sim \text{SAMPLE}()$  (see Section C) and query SimTeacher in Algorithm 1 for  $y$ 
9:       Store preference  $\mathcal{D} \leftarrow \mathcal{D} \cup \{(\sigma^0, \sigma^1, y)\}$ 
10:    end for
11:    for each gradient step do
12:      Sample minibatch  $\{(\sigma^0, \sigma^1, y)_j\}_{j=1}^D \sim \mathcal{D}$  and optimize  $\mathcal{L}^{\text{Reward}}$  in (3) with respect to  $\psi$ 
13:    end for
14:  end if
15:  // POLICY LEARNING
16:  for each timestep  $t$  do
17:    Collect  $\mathbf{s}_{t+1}$  by taking  $\mathbf{a}_t \sim \pi(\mathbf{a}_t|\mathbf{s}_t)$  and store transitions  $\mathcal{B} \leftarrow \mathcal{B} \cup \{(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1}, \hat{r}_\psi(\mathbf{s}_t))\}$ 
18:  end for
19:  for each gradient step do
20:    Sample random minibatch  $\{(\tau_j)\}_{j=1}^B \sim \mathcal{B}$  and optimize RL objective function with respect to  $\phi$ 
21:  end for
22:  Reset  $\mathcal{B} \leftarrow \emptyset$  if on-policy RL algorithm is used
23: end for
```

---

## B Preference-based reinforcement learning

Algorithm 3 summarizes the full procedure of preference-based RL methods that we consider in this paper. We also utilize unsupervised pre-training which encourages our agent to visit a wider range of states by using the state entropy  $\mathcal{H}(\mathbf{s}) = -\mathbb{E}_{\mathbf{s} \sim p(\mathbf{s})} [\log p(\mathbf{s})]$  as an intrinsic reward [42, 29, 40, 56]. By following Lee et al. [39], we define the intrinsic reward of the current state  $\mathbf{s}_t$  as follows:

$$r^{\text{int}}(\mathbf{s}_t) = \log(\|\mathbf{s}_t - \mathbf{s}_t^k\|), \quad (8)$$

where  $\mathbf{s}_t^k$  is the  $k$ -NN of  $\mathbf{s}_t$  within a set  $\{\mathbf{s}_i\}_{i=1}^N$ . The full procedure of unsupervised pre-training is summarized in Algorithm 2.

## C Sampling schemes

We consider the following sampling schemes in this paper:

- *Uniform sampling*: We pick  $N_{\text{query}}$  pairs of segments uniformly at random from the buffer  $\mathcal{B}$ .
- *Disagreement*: We first generate the initial batch of  $N_{\text{inter}}$  pairs of segments  $\mathcal{G}_{\text{init}}$  uniformly at random, measure the variance across ensemble of preference predictors  $\{P_{\psi_i}[\sigma^1 \succ \sigma^0]\}_{i=1}^{N_{\text{en}}}$ , and then select the  $N_{\text{query}}$  pairs of segments with high uncertainty.
- *Entropy*: We first generate the initial batch of  $N_{\text{inter}}$  pairs of segments  $\mathcal{G}_{\text{init}}$  uniformly at random, measure the entropy of a single preference predictor  $\mathcal{H}(P_\psi)$ , and then select the  $N_{\text{query}}$  pairs of segments with high uncertainty.
- *Coverage*: From the initial batch  $\mathcal{G}_{\text{init}}$ , we choose center points, which increase the dissimilarity between the selected queries. Specifically, we concatenate the states of segments, i.e.,  $\mathbf{s}_{\text{concat}} = \text{Concat}(\mathbf{s}_{k+1}^0, \dots, \mathbf{s}_{k+H}^0, \mathbf{s}_{k+1}^1, \dots, \mathbf{s}_{k+H}^1)$ <sup>7</sup> and measure the Euclidean distance. Then, we choose  $N_{\text{query}}$  center points such that the largest distance between a data point and its nearest center is minimized using a greedy selection strategy.

---

<sup>7</sup>Concatenating states would not be an optimal choice because it is not permutation-invariant, which is also not handled in the prior work [11]. However, we expect that an issue from permutation-variance is not significant because a probability to sample two segments in a different order is very low. However, it is an interesting future direction to explore to address this limitation.

- *Disagreement + Coverage*: We first select the  $N_{\text{inter}}$  pairs of segments  $\mathcal{G}_{\text{un}}$ , using the disagreement sampling, where  $N_{\text{init}} > N_{\text{inter}}$ , and then choose  $N_{\text{query}}$  center points from  $\mathcal{G}_{\text{un}}$ .
- *Entropy + Coverage*: We first select the  $N_{\text{inter}}$  pairs of segments  $\mathcal{G}_{\text{un}}$ , using the entropy sampling, and then choose  $N_{\text{query}}$  center points from  $\mathcal{G}_{\text{un}}$ .

## D Experimental Details

**Training details.** We use PEBBLE and PrefPPO<sup>8</sup> with a full list of hyperparameters in Table 1 and Table 2, respectively. We pre-train an agent for 10K timesteps and 32K timesteps for PEBBLE and PrefPPO, respectively.

Hyperparameter	Value	Hyperparameter	Value
Initial temperature	0.1	Hidden units per each layer	1024 (DMControl), 256 (Meta-world)
Segment of length	50 (DMControl), 25 (Meta-world)	# of layers	2 (DMControl), 3 (Meta-world)
Learning rate	0.0003 (Meta-world)	Batch Size	1024 (DMControl), 512 (Meta-world)
	0.0001 (Quadruped), 0.0005 (Walker)	Optimizer	Adam [33]
Critic target update freq	2	Critic EMA $\tau$	0.005
$(\beta_1, \beta_2)$	(.9, .999)	Discount $\bar{\gamma}$	.99
Frequency of feedback	5000 (Meta-world), 20000 (Walker)	Maximum budget /	2000/200, 1000/100, 500/50 (DMControl)
	30000 (Quadruped)	# of queries per session	20K/100, 10K/50 (Meta-world)

Table 1: Hyperparameters of the PEBBLE algorithm.

Hyperparameter	Value	Hyperparameter	Value
GAE parameter $\lambda$	0.9 (Quadruped), 0.92 (otherwise)	Hidden units per each layer	256
Segment of length	50 (DMControl), 25 (Meta-world)	# of layers	3
Learning rate	0.0003 (Meta-world)	Batch Size	64 (Walker), 256 (Sweep Into)
	$5e^{-5}$ (DMControl)		128 (Quadruped, Button)
Discount $\bar{\gamma}$	.99	Frequency of feedback	32000 (DMControl)
# of environments per worker	8 (Button), 16 (Quadruped),	PPO clip range	0.4
	32 (Walker, Sweep Into)	Entropy bonus	0.0
# of timesteps per rollout	500 (DMControl)	Maximum budget /	2000/200, 1000/100 (DMControl)
	250 (Meta-world)	# of queries per session	

Table 2: Hyperparameters of the PrefPPO algorithm.

**Reward model.** For the reward model, we use a three-layer neural network with 256 hidden units each, using leaky ReLUs. To improve the stability in reward learning, we use an ensemble of three reward models, and bound the output using tanh function. Each model is trained by optimizing the cross-entropy loss defined in (3) using ADAM learning rule [33] with the initial learning rate of 0.0003.

**Simulated human teachers.** For all experiments, To evaluate the robustness, we evaluate against the following simulated human teachers with different hyperparameters:

- Oracle:  $\text{SimTeacher}(\beta \rightarrow \infty, \gamma = 1, \epsilon = 0, \delta_{\text{skip}} = 0, \delta_{\text{equal}} = 0)$
- Stoc:  $\text{SimTeacher}(\beta = 1, \gamma = 1, \epsilon = 0, \delta_{\text{skip}} = 0, \delta_{\text{equal}} = 0)$
- Mistake:  $\text{SimTeacher}(\beta \rightarrow \infty, \gamma = 1, \epsilon = 0.1, \delta_{\text{skip}} = 0, \delta_{\text{equal}} = 0)$
- Skip:  $\text{SimTeacher}(\beta \rightarrow \infty, \gamma = 1, \epsilon = 0, \delta_{\text{skip}} = \delta_{\text{adapt}}(\epsilon_{\text{adapt}}, t), \delta_{\text{equal}} = 0)$
- Equal:  $\text{SimTeacher}(\beta \rightarrow \infty, \gamma = 1, \epsilon = 0, \delta_{\text{skip}} = 0, \delta_{\text{adapt}} = \delta_{\text{skip}}(\epsilon_{\text{adapt}}, t))$
- Myopic:  $\text{SimTeacher}(\beta \rightarrow \infty, \gamma = 0.9, \epsilon = 0, \delta_{\text{skip}} = 0, \delta_{\text{equal}} = 0)$

Because each environment has a different scale of ground truth rewards, it is hard to design standardized Skip and Equal teachers using hard threshold. To address this issue, we use an adaptive

<sup>8</sup>For Meta-world, the frequency is chosen from  $\{8K, 16K, 32K, 64K\}$  and # of queries per session is chosen from  $\{50, 100, 250, 500, 1000\}$ .

threshold, which is defined as follows:

$$\delta_{\text{adapt}}(\epsilon_{\text{adapt}}, t) = \frac{H}{T} R_{\text{avg}}(\phi_t) \epsilon_{\text{adapt}}, \quad (9)$$

where  $t$  is the current timestep,  $\epsilon_{\text{adapt}} \in [0, 1]$  is hyperparameters to control the threshold,  $T$  is the episode length,  $H$  is a length of segment and  $R_{\text{avg}}(\pi_t)$  is the average returns of current policy with the parameters  $\pi_t$ . By adaptively rescaling the threshold based on the performance of agent, a teacher skips queries or provides uniform labels (i.e.,  $y = (0.5, 0.5)$ ). For all experiments, we choose  $\epsilon_{\text{adapt}} = 0.1$ .

## E Additional experimental results

**Reward analysis.** Figure 6 shows the learned reward function optimized by PEBBLE on the oracle teacher in all tested environments. Because we bound the output using tanh function, the scale is different with the ground truth reward but the learned reward function is reasonably well-aligned.

**Regularization for handling corrupted labels.** To improve the robustness to corrupted labels, we apply the label smoothing [64]. By following Christiano et al. [19], Ibarz et al. [31], we use a soft label  $\hat{y} = 0.9 * y + 0.05$  for the cross-entropy computation. As shown in Figure 7, we find that the gains from label smoothing are marginal.

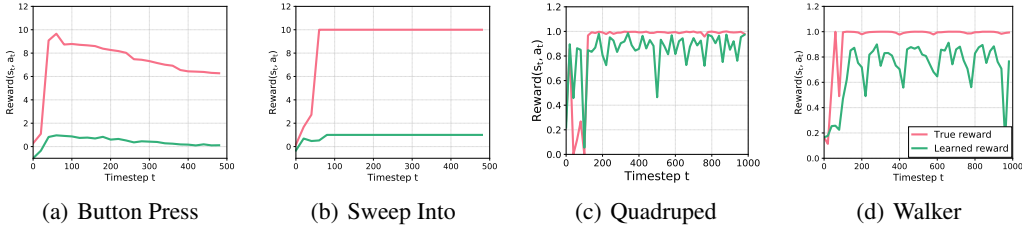


Figure 6: Time series of learned reward function (green) and the ground truth reward (red) using rollouts from a policy optimized by PEBBLE.

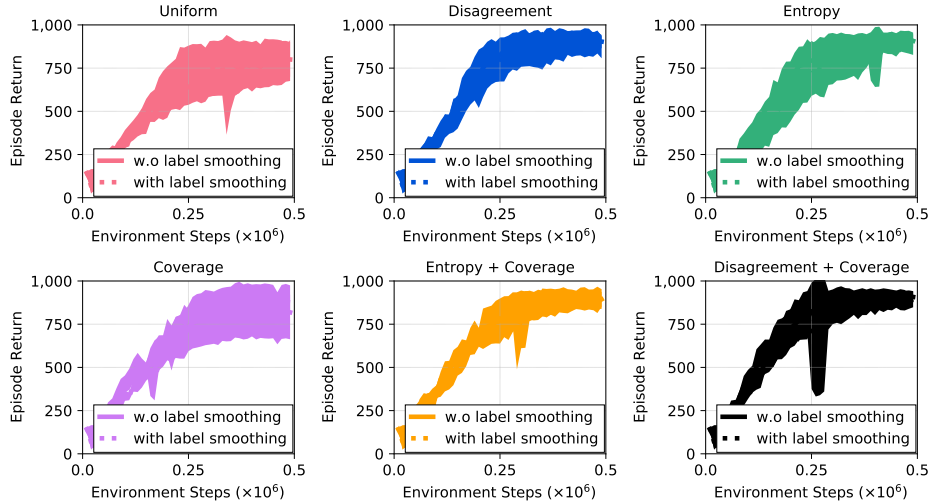


Figure 7: Learning curves of PEBBLE with 500 queries on Walker on the Mistake teacher. The solid line and shaded regions represent the mean and standard deviation, respectively, across five runs.

## F Learning curves

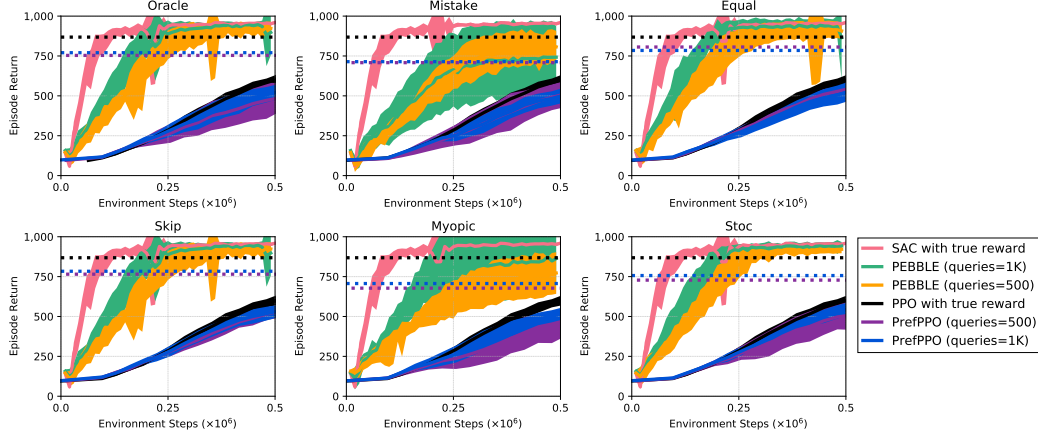


Figure 8: Learning curves of PEBBLE and PrefPPO on Walker-walk as measured on the ground truth reward. The solid line and shaded regions represent the mean and standard deviation, respectively, across ten runs. Asymptotic performance of PPO and PrefPPO is indicated by dotted lines of the corresponding color.

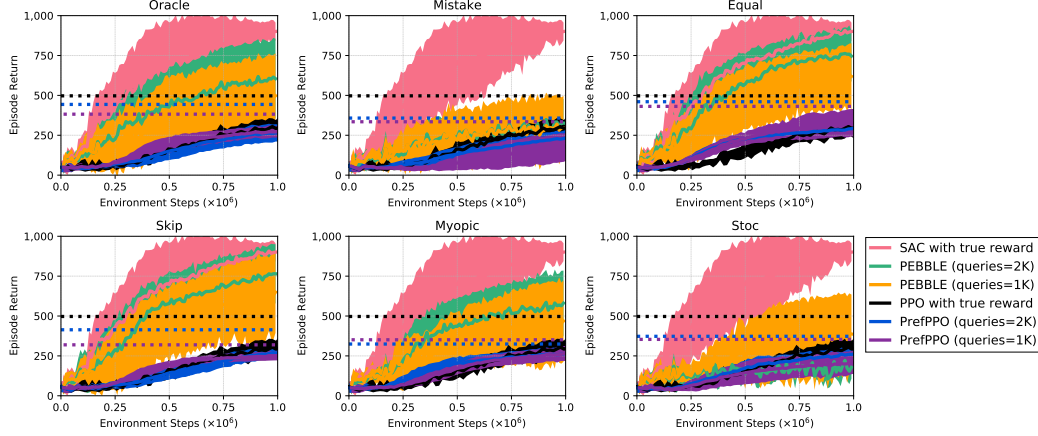


Figure 9: Learning curves of PEBBLE and PrefPPO on Quadruped-walk as measured on the ground truth reward. The solid line and shaded regions represent the mean and standard deviation, respectively, across ten runs. Asymptotic performance of PPO and PrefPPO is indicated by dotted lines of the corresponding color.



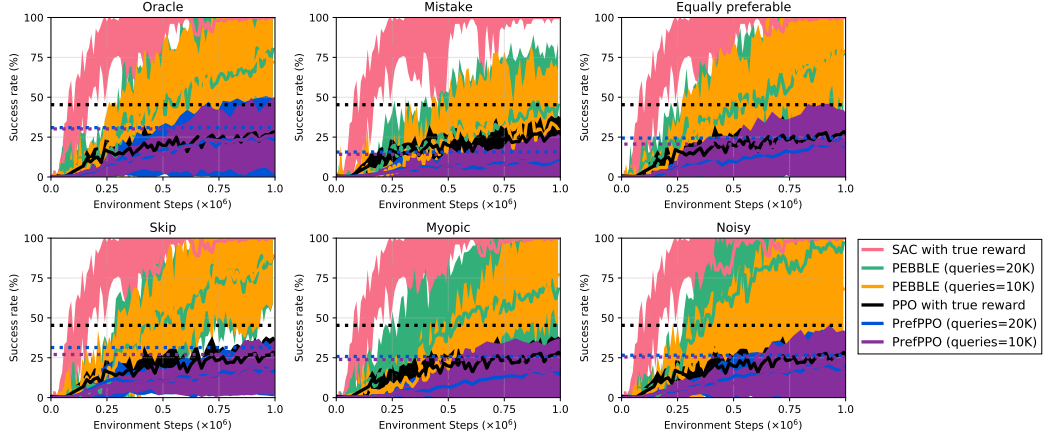


Figure 10: Learning curves of PEBBLE and PrefPPO on Sweep Into as measured on the success rate. The solid line and shaded regions represent the mean and standard deviation, respectively, across ten runs. Asymptotic performance of PPO and PrefPPO is indicated by dotted lines of the corresponding color.

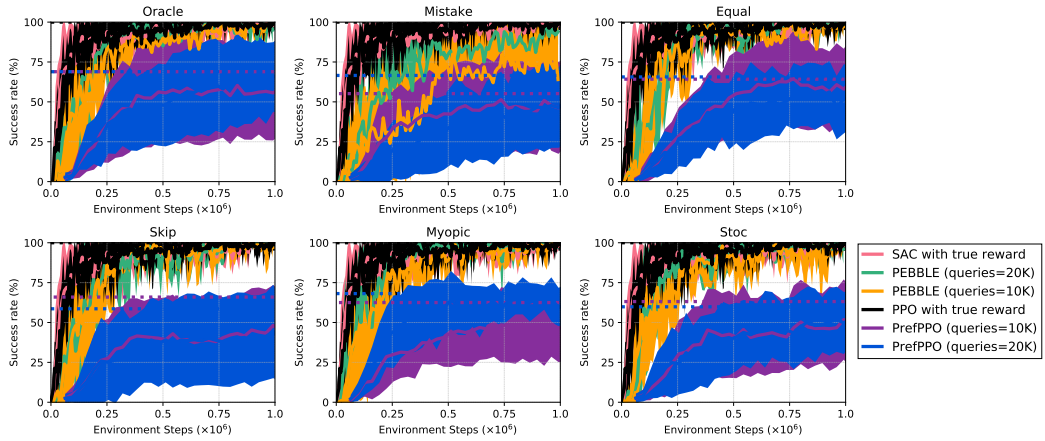


Figure 11: Learning curves of PEBBLE and PrefPPO on Button Press as measured on the success rate. The solid line and shaded regions represent the mean and standard deviation, respectively, across ten runs. Asymptotic performance of PPO and PrefPPO is indicated by dotted lines of the corresponding color.

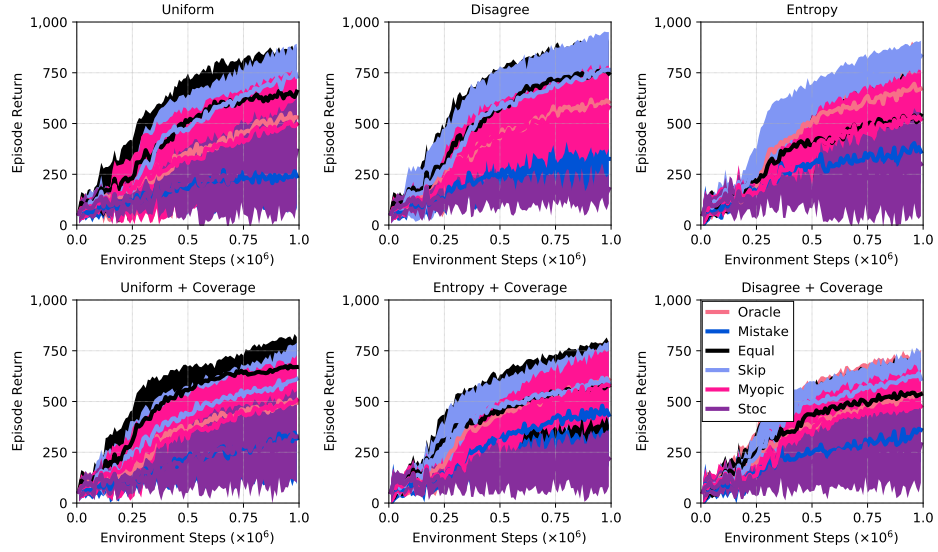


Figure 12: Learning curves of PEBBLE with 2000 queries on Quadruped-walk as measured on the ground truth reward. The solid line and shaded regions represent the mean and standard deviation, respectively, across ten runs.

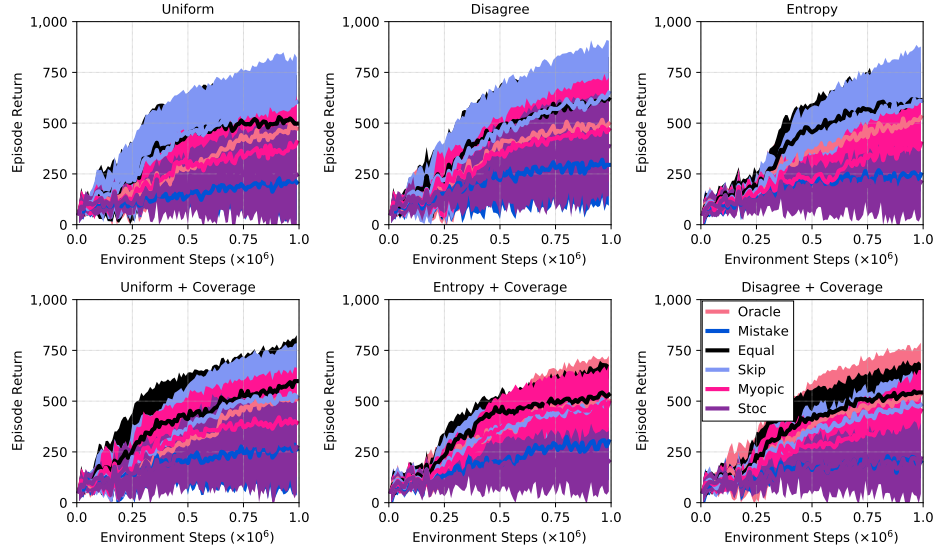


Figure 13: Learning curves of PEBBLE with 1000 queries on Quadruped-walk as measured on the ground truth reward. The solid line and shaded regions represent the mean and standard deviation, respectively, across ten runs.

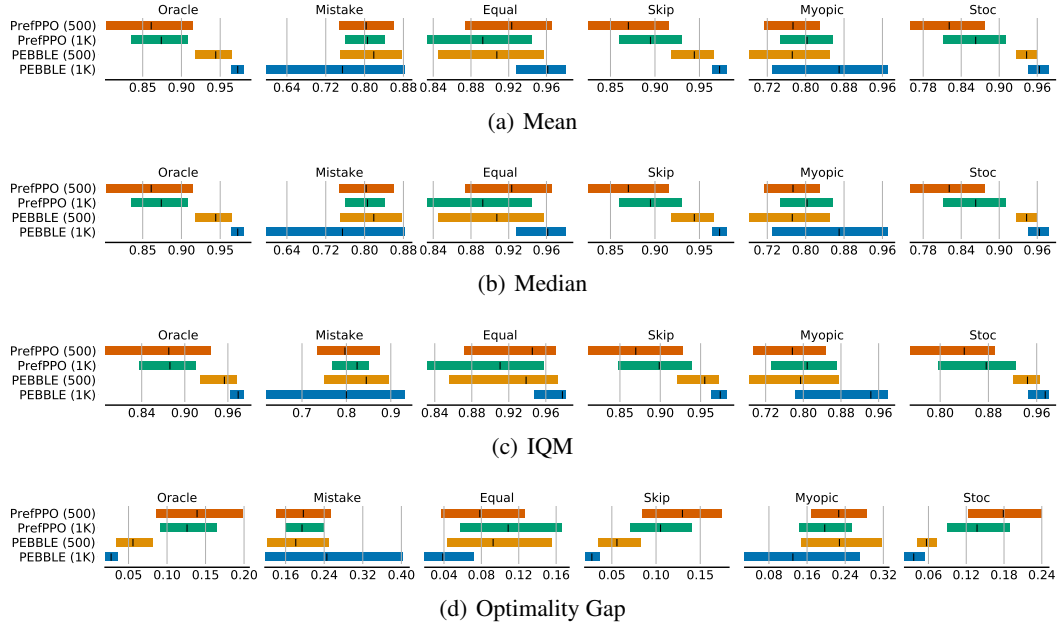


Figure 14: Aggregate metrics on Walker with 95% confidence intervals (CIs) across ten runs. Higher mean, median and IQM scores and lower optimality gap are better. The CIs are estimated using the percentile bootstrap with stratified sampling.

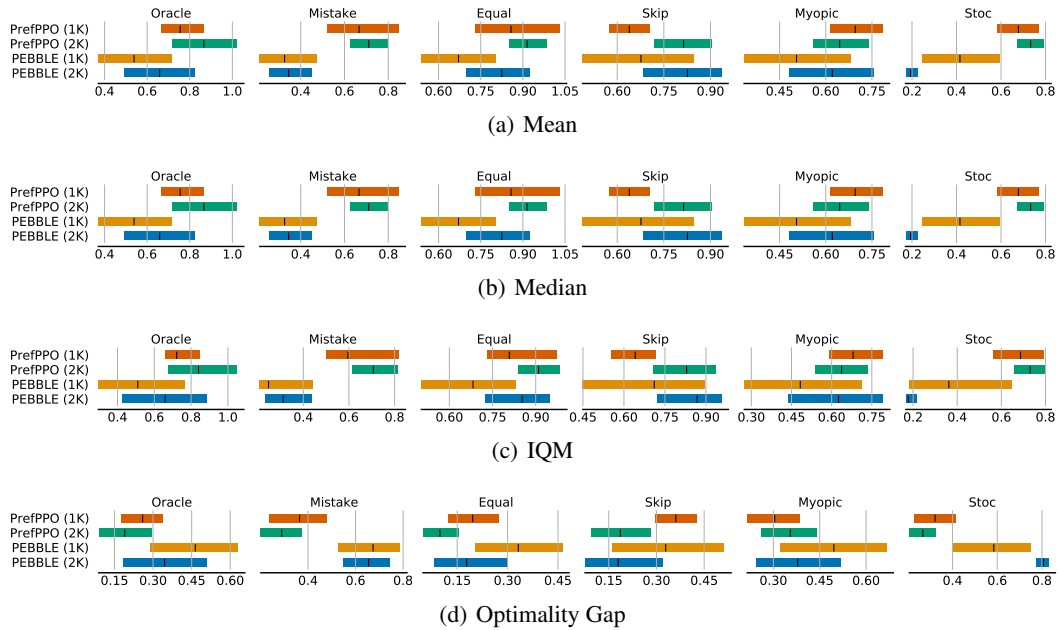


Figure 15: Aggregate metrics on Quadruped with 95% confidence intervals (CIs) across ten runs. Higher mean, median and IQM scores and lower optimality gap are better. The CIs are estimated using the percentile bootstrap with stratified sampling.

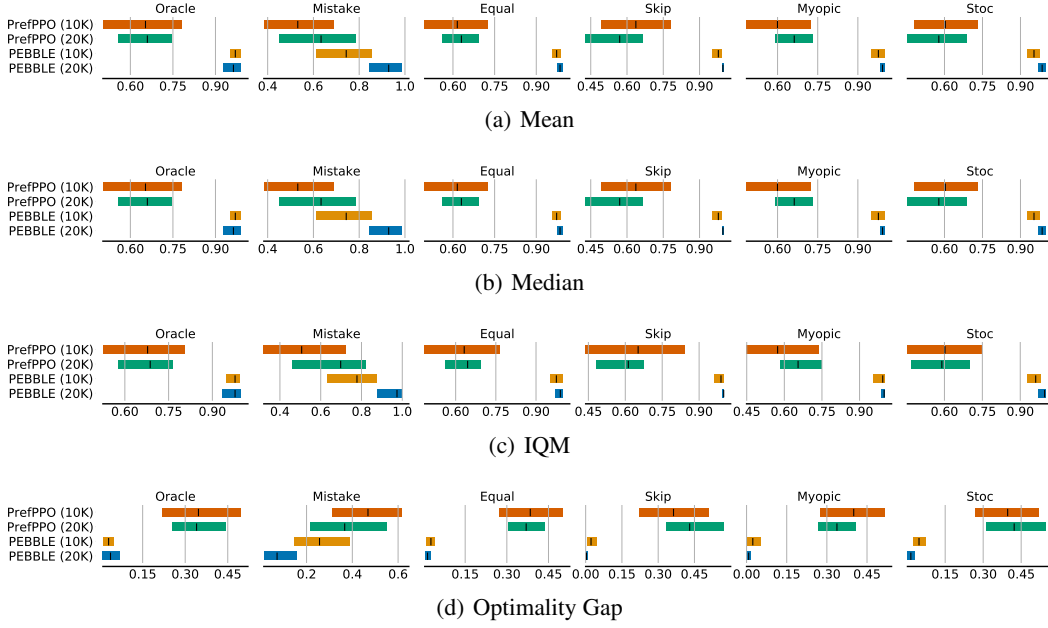


Figure 16: Aggregate metrics on Button Press with 95% confidence intervals (CIs) across ten runs. Higher mean, median and IQM scores and lower optimality gap are better. The CIs are estimated using the percentile bootstrap with stratified sampling.

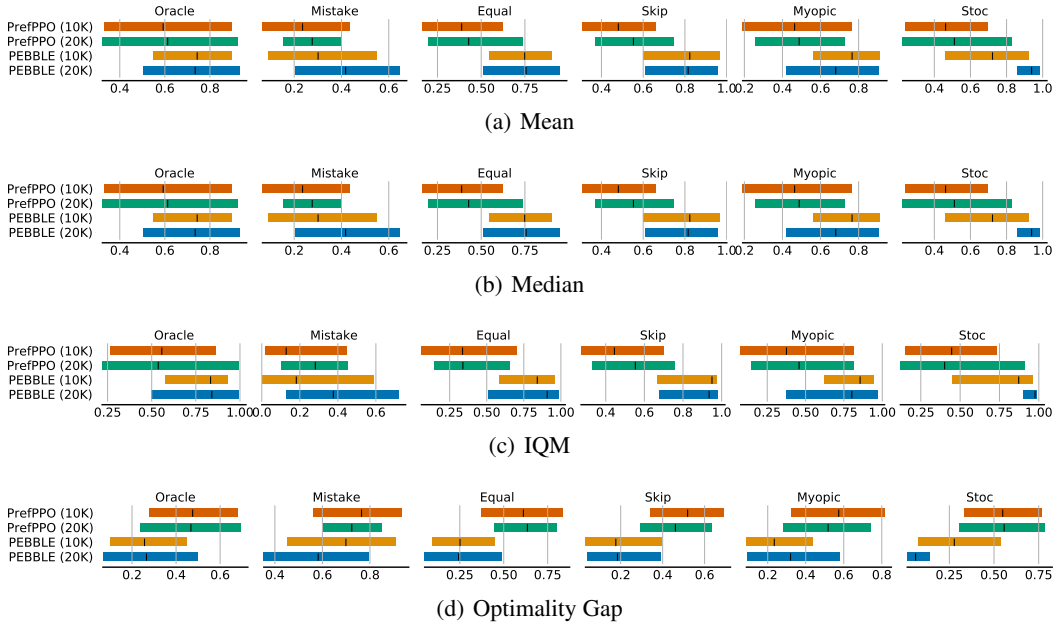


Figure 17: Aggregate metrics on Sweep Into with 95% confidence intervals (CIs) across ten runs. Higher mean, median and IQM scores and lower optimality gap are better. The CIs are estimated using the percentile bootstrap with stratified sampling.

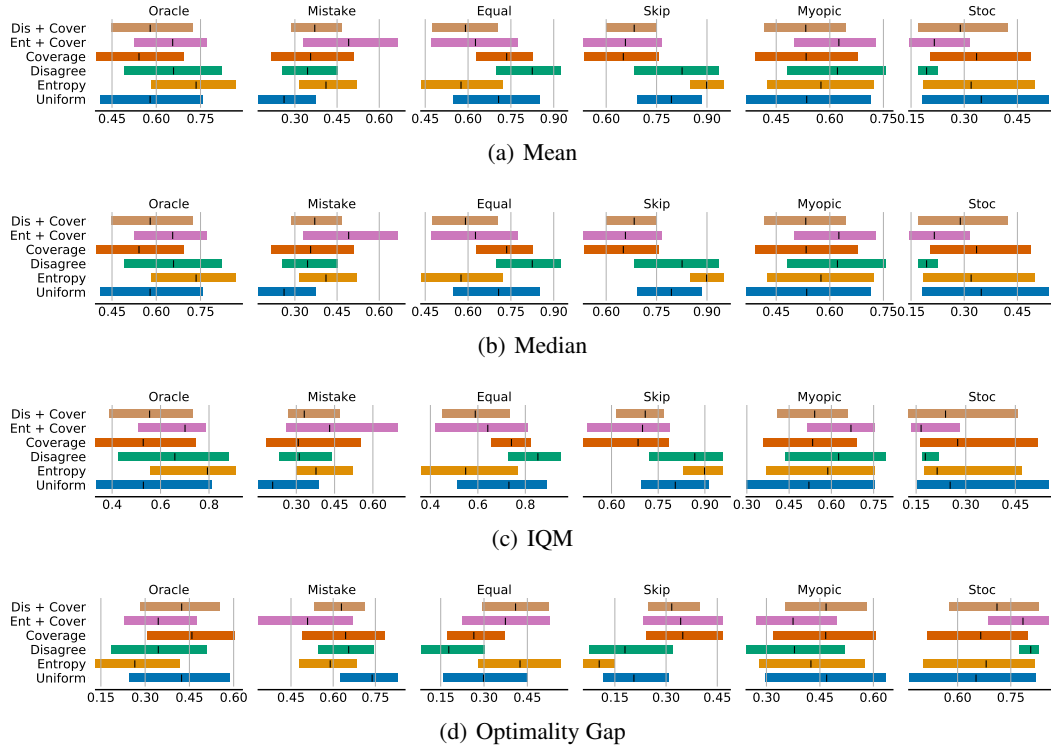


Figure 18: Aggregate metrics of PEBBLE on Quadruped with 2000 queries across ten runs. Higher mean, median and IQM scores and lower optimality gap are better. The CIs are estimated using the percentile bootstrap with stratified sampling.

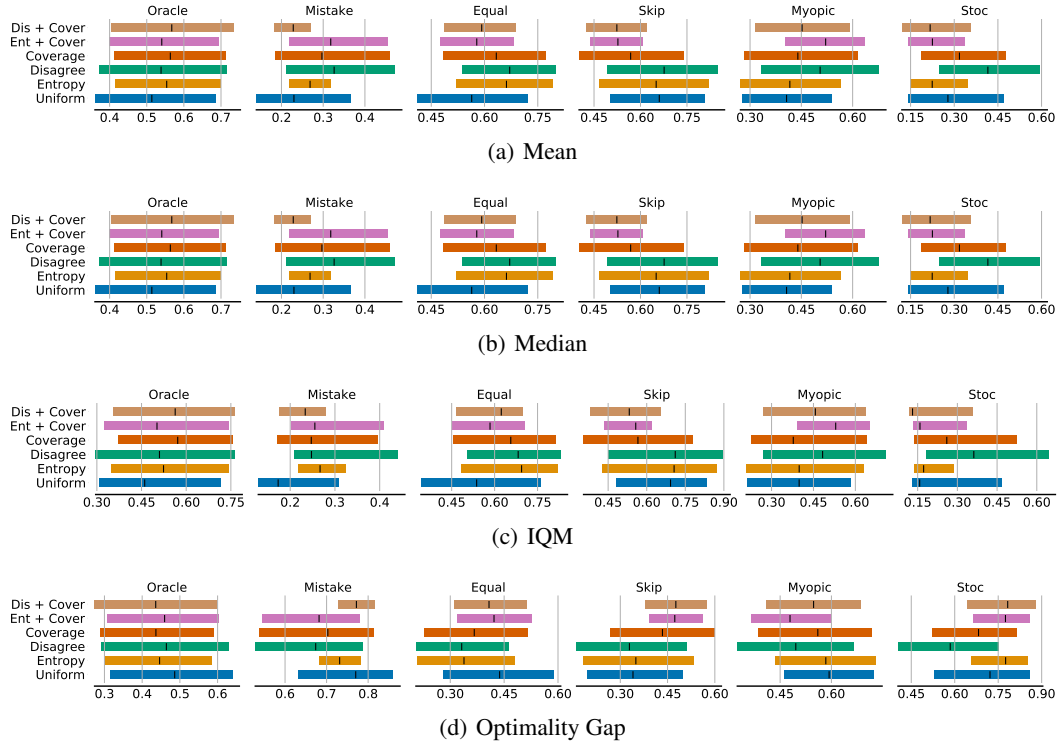


Figure 19: Aggregate metrics of PEBBLE on Quadraped with 1000 queries across ten runs. Higher mean, median and IQM scores and lower optimality gap are better. The CIs are estimated using the percentile bootstrap with stratified sampling.