

DAPPER: Discriminability-Aware Policy-to-Policy Preference-Based Reinforcement Learning for Query-Efficient Robot Skill Acquisition

Yuki Kadokawa¹, Jonas Frey², Takahiro Miki², Takamitsu Matsubara¹, and Marco Hutter²

Abstract—Preference-based Reinforcement Learning (PbRL) enables policy learning through simple queries comparing trajectories from a single policy. While human responses to these queries make it possible to learn policies aligned with human preferences, PbRL suffers from low query efficiency, as policy bias limits trajectory diversity and reduces the number of discriminable queries available for learning preferences. This paper identifies *preference discriminability*, which quantifies how easily a human can judge which trajectory is closer to their ideal behavior, as a key metric for improving query efficiency. To address this, we move beyond comparisons within a single policy and instead generate queries by comparing trajectories from multiple policies, as training them from scratch promotes diversity without policy bias. We propose **Discriminability-Aware Policy-to-Policy Preference-Based Efficient Reinforcement Learning (DAPPER)**, which integrates preference discriminability with trajectory diversification achieved by multiple policies. DAPPER trains new policies from scratch after each reward update and employs a discriminator that learns to estimate preference discriminability, enabling the prioritized sampling of more discriminable queries. During training, it jointly maximizes the preference reward and preference discriminability score, encouraging the discovery of highly rewarding and easily distinguishable policies. Experiments in simulated and real-world legged robot environments demonstrate that DAPPER outperforms previous methods in query efficiency, particularly under challenging preference discriminability conditions.

Index Terms—Reinforcement Learning, Human Factors and Human-in-the-Loop, Legged Robots.

I. INTRODUCTION

LEGGED locomotion has received considerable attention, with growing interest in enabling robots to acquire walking behaviors that reflect human preferences rather than relying solely on task-specific objectives [1, 2]. Existing approaches mainly rely on reinforcement learning [3, 4] or imitation learning [2, 5]. Reinforcement learning requires task-specific reward parameter tuning that depends heavily on expert knowledge, while imitation learning relies on demonstrations that demand operator proficiency with robot interfaces and motion strategies. Both introduce substantial costs for users, making it difficult to realize locomotion behaviors aligned with individual preferences. In this paper, we focus

This work was supported by JSPS KAKENHI Grant Numbers JP23KJ1585 and JP24K03018. (Corresponding author: Yuki Kadokawa.) ¹ Yuki Kadokawa and Takamitsu Matsubara are with Nara Institute of Science and Technology, Nara 630-0192, Japan, kadokawa.yuki@naist.ac.jp, takam-m@is.naist.jp. ² Jonas Frey, Takahiro Miki, and Marco Hutter are with ETH Zurich, Zurich 8092, Switzerland, jonfrey@ethz.ch, tamiki@ethz.ch, mahutter@ethz.ch.

©2026 IEEE

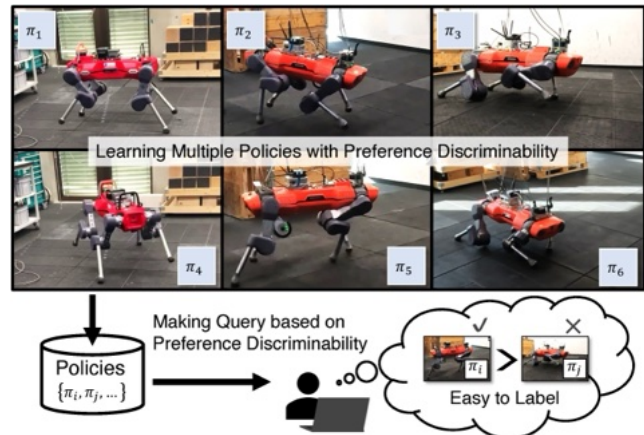


Fig. 1: Overview of the proposed learning framework: Multiple policies with diverse feature representations are sequentially learned. Queries are generated by sampling from these policies based on human preference discriminability, ensuring easier labeling. This approach enhances query efficiency by increasing the number of discriminable queries.

on quadruped locomotion and propose a preference-based learning framework that reduces user effort while shaping policies according to human preferences.

Preference-based Reinforcement Learning (PbRL) has gained attention as it enables the learning of control policies through simple query responses [6, 7]. In typical PbRL frameworks, two state-action trajectories are sampled from interactions between a single policy and the environment, and a human selects the more preferred one. Based on these preference labels, a reward model that approximates human preferences is learned, and the control policy is updated through reinforcement learning. By iteratively updating the reward model with new query responses and optimizing the policy accordingly, PbRL gradually learns a policy that reflects human preferences [8, 9].

A key challenge in PbRL is that diverse state-action trajectories are required to efficiently learn human preferences via queries [10, 11]. Ideally, each update of the reward model should lead to the acquisition of new policies that can generate novel state-action trajectories. However, when using a single policy, policy bias [12, 13] becomes a critical problem. Policy bias refers to the limited behavioral diversity that arises when learning from a single policy, as new policies are trained as incremental updates of previous ones. Because each policy is a continuation of the last, their behaviors tend to be strongly correlated, making it difficult to explore new actions and fully

IEEE Robotics & Automation Magazine (RAM) paper, presented at ICRA 2026, Vienna, Austria. Cite as RAM paper.

exploit updated reward functions. As a result, the collected trajectories tend to be similar, and behavioral diversity is significantly limited.

To address this issue, we focus on the ease of labeling and introduce *preference discriminability* as a key metric. This metric quantifies whether a human can determine how much two trajectories in a query are closer to their ideal behavior. When one trajectory is clearly closer to the ideal, labeling is straightforward. However, when both trajectories are similarly undesirable or equally good, humans may struggle to provide a clear preference, not because the trajectories themselves are indistinguishable, but because their underlying preferences are indistinguishable. By generating queries with high preference discriminability, we present clearer differences between trajectories, making them easier to label and thereby improving query efficiency, as a greater number of discriminable queries can be used to learn the human preference.

Building on this perspective, we focus on moving beyond comparisons within a single policy and instead generate queries by comparing trajectories from different policies as shown in Fig. 1. Training multiple policies encourages the emergence of more diverse behaviors, since learning different policies from scratch is not constrained by the policy bias. This approach avoids the tendency of gradual policy updates to produce similar behaviors and supports the generation of richer and more varied state-action trajectories. However, simply increasing diversity does not guarantee that the resulting queries will be easy for humans to label. To ensure that diversity leads to meaningful and discriminable comparisons, we further introduce a preference discriminability-aware query generation mechanism that estimates the discriminability of trajectory pairs.

To this end, we propose a novel framework, Discriminability-Aware Policy-to-Policy Preference-Based Efficient Reinforcement Learning (DAPPER), which integrates two key ideas: preference discriminability and trajectory diversification achieved by multiple policies. DAPPER approximates preference discriminability as a learnable function by training a discriminator that estimates whether a given query is likely to be distinguishable by humans, based on the history of queries labeled as indistinguishable. During training, DAPPER probabilistically samples queries based on their predicted preference discriminability, promoting the selection of discriminable queries. To enhance trajectory diversity, DAPPER trains multiple policies. Each time the reward model is updated, a new policy is retrained from scratch. Finally, during policy learning, we not only maximize for the reward of human preference but also incorporate the preference discriminability score as an auxiliary reward. This dual-objective approach encourages the discovery of policies that yield both high-reward and easily distinguishable behaviors.

We demonstrated the effectiveness of DAPPER in a legged robot simulator and real-world robot environment. DAPPER achieves successful policy learning with fewer queries than previous methods. Furthermore, compared to previous methods, only DAPPER consistently learn policies under the conditions that preference discriminability becomes increasingly

challenging. We also applied the learned policies to real-world environments.

II. RELATED WORK

A. Learning Control Policies

This section describes previous methods for learning control policies and the differences from our approach.

Imitation Learning: Learning from human demonstrations has been explored in previous studies [2, 5]. However, creating demonstrations for complex mechanisms such as quadrupedal robots is challenging. While recent efforts have focused on improving demonstration efficiency by refining demonstration methods, these approaches still require demonstrators to have prior knowledge of the policy features being taught. Moreover, effective data collection depends on a clear understanding of the desired policy features beforehand [2]. For non-experts, predicting and defining the precise behavior of a quadrupedal robot in advance is difficult, making it challenging to acquire the intended policy features.

Reinforcement Learning: Reinforcement Learning (RL) enables robots to autonomously acquire arbitrary policies through trial-and-error interactions with the environment by designing a reward function [3, 4]. However, designing an appropriate reward function that results in the desired behavior requires a carefully designed reward, often incorporating multiple auxiliary terms and considerable manual effort. Moreover, reward function design requires expert knowledge, making it difficult for non-experts to define rewards that effectively capture their intended requirements. As a result, acquiring policies resulting in specific motion patterns through RL remains difficult for non-experts.

Preference-based Reinforcement Learning: Potentially, even non-experts can evaluate and distinguish desirable robot behaviors by responding to video-based comparison queries, enabling the learning of policies that align with human preferences [6, 11]. However, PbRL requires collecting a large number of queries, making policy learning infeasible within a practical timeframe. To the best of our knowledge, this paper is the first to apply PbRL to learning quadrupedal robot locomotion. To achieve this, we aim to improve query efficiency, enabling policy learning within a reasonable timeframe.

B. Improving Query Efficiency in PbRL

This section reviews previous works for improving query efficiency in PbRL and discusses the differences from our work.

Data Augmentation: In this approach, in addition to human annotations, the preference reward model under training is also used to automatically assign labels. Specifically, when a query sample is fed into the model and the predicted probability (ranging from 0 to 1) that one trajectory is preferred over the other shows a clear margin, the model's prediction is used as the label instead of querying a human. This enables automatic dataset expansion and reduces the need for human responses, thereby improving query efficiency [7, 11].

Robustifying to Incorrect Human Feedback: Incorrect human feedback can mislead the learning process, reducing the

effectiveness of each query. To improve query efficiency under such conditions, robust learning techniques are employed to reduce sensitivity to label noise. These include filtering out inconsistent responses and using robust loss functions that reduce the influence of potentially incorrect labels. By mitigating the negative impact of noisy data, these methods enable more reliable policy updates from fewer queries, thereby enhancing overall query efficiency [8, 9].

Uncertainty Sampling: Uncertainty sampling aims to improve query efficiency by prioritizing queries near the decision boundary of the preference reward model, where the output probability is close to 0.5. These queries are expected to provide the most informative feedback for refining the reward function. By focusing on ambiguous cases, the model can collect targeted training data that reduces uncertainty and sharpens its decision boundaries [10, 11].

Active Reward Learning: This approach collects queries from a pre-generated trajectory dataset and trains a policy only once using the reward function learned from those queries. Following information-theoretic objective, it implicitly favors trajectory pairs that are easier for humans to distinguish, since indistinguishable queries yield little information gain. However, because it depends entirely on a fixed dataset of candidate trajectories, it assumes that optimal and easily comparable pairs already exist in the dataset. Without a mechanism to generate new trajectories through policy optimization, it cannot expand behavioral variations for more discriminable queries [10, 14].

Reward Re-calculation: This approach improves query efficiency by recalculating the reward values of previously collected trajectories, enabling the reuse of past dataset without requiring additional queries. The agent can effectively expand its training dataset and refine the reward model more rapidly. This helps reduce the reliance on new annotations and contributes to more stable policy learning [10, 15].

Ours: While previous approaches improve query efficiency by leveraging model confidence [7], addressing annotation noise [10], actively selecting informative queries [11], selecting queries based on information gain [14], or reusing past data through reward re-calculation [15], they do not explicitly account for the labeling difficulty experienced by human annotators. In contrast, our framework improves query efficiency by modeling and leveraging human preference discriminability, which helps reduce indistinguishable queries and enhances query efficiency. This is particularly important in tasks where distinguishing task features is inherently difficult. If a query is not distinguishable to the human, a preference label cannot be obtained.

III. PRELIMINARIES

A. Reinforcement Learning

Reinforcement Learning (RL) assumes a scenario in which an agent takes actions through observed state [3, 4]. The agent receives feedback about its actions in the form of a reward. The goal of the agent is to learn actions to maximize total reward. The standard formalization of a RL problem builds on the notion of a Markov Decision Process (MDP) and consists

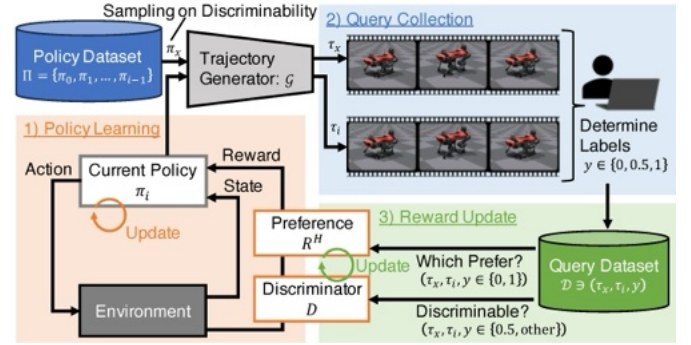


Fig. 2: Learning process of DAPPER: Our framework consists of three steps. (1) Learning current policy π_i with reward r . (2) Query collection. This step makes queries from current policy π_i and sampled policy from policy dataset Π based on discriminability. Then, trajectories of each policy are generated τ_x and τ_i , and queries are labeled y . (3) Reward Update. This step trains the preference reward model R^H and discriminator D from query dataset \mathcal{D} .

of: A continuous state space $S \subseteq \mathbb{R}^n$ which represents an observation from the environment; A continuous action space $A \subseteq \mathbb{R}^m$ representing the set of actions available to the agent; state transition probability $\mathcal{P} : S \times A \rightarrow P(S)$, where $P(S)$ denotes the probability that action a in state s leads the agent to state s' ; $\gamma \in [0, 1)$ is the discount factor. The agent interacts with the environment by taking an action $a \in A$ based on the current state $s \in S$, receiving a reward $r(s, a)$, and transitioning to a new state $s' \in S$ according to the probability distribution $\mathcal{P}(\cdot | s, a)$.

The agent moves through the state space by repeatedly taking actions, thereby generating a trajectory $\tau = (s_0, a_0, s_1, \dots, s_t, a_t, s_{t+1})$, where s_0 is the state the agent starts with, a_t is the action it takes in state s_t , and s_{t+1} is the next state produced by this action.

The most common task in reinforcement learning is to learn a policy $\pi : S \rightarrow A$ that prescribes the optimal action for the agent in each state. More specifically, the objective is often defined as maximizing the total rewards: $V_\pi(s) = \mathbb{E}_{\pi, \mathcal{P}} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \mid s_0 = s \right]$. Accordingly, common approaches have focused on learning an optimal policy by maximizing the value function V . Among various policy optimization algorithms, Proximal Policy Optimization (PPO) is recognized as an effective means to achieve this objective [3, 4].

B. Preference-based Reinforcement Learning

In Preference-based Reinforcement Learning (PbRL), the agent learns a reward function based on human feedback and then uses this reward function to learn the optimal policies [6]. The process is as follows.

First, the agent collects state-action trajectories τ . This transition is provided as preferences over pairs of trajectories, expressed as $P(\tau_i \succ \tau_j)$, where $P(\tau_i \succ \tau_j)$ represents the probability that trajectory τ_i is preferred over trajectory τ_j . Basically, for each trajectory of a query pair, PbRL utilizes an interval clipped from the trajectory to a specific step length.

Human preference is modeled by logistic regression as follows:

$$P(\tau_i \succ \tau_j) = \frac{\exp(r(\tau_i))}{\exp(r(\tau_i)) + \exp(r(\tau_j))} \quad (1)$$

where $r(\tau)$ is the cumulative reward function for trajectory τ .

Given the feedback data $\{(\tau_i, \tau_j)\}_{k=1}^N$, the reward function $r(\tau)$ is optimized by minimizing the binary cross-entropy loss function:

$$\mathcal{L} = -\sum_{k=1}^N [y_k \log P(\tau_{i_k} \succ \tau_{j_k}) + (1 - y_k) \log P(\tau_{j_k} \succ \tau_{i_k})] \quad (2)$$

where y_k is the query label from the human preference, with $y_k = 1$ if τ_{i_k} is preferred over τ_{j_k} , and $y_k = 0$ otherwise.

Using the estimated reward function $r(\tau)$, the agent learns a policy $\pi(a|s)$ within the reinforcement learning framework. The agent continuously receives human feedback as it takes new actions, updating the reward function $r(\tau)$ iteratively. This iterative process involves re-estimating the reward function based on new feedback data and refining the agent's policy.

IV. PROPOSED FRAMEWORK

We propose Discriminability-Aware Policy-to-Policy Preference-Based Efficient Reinforcement Learning (DAPPER), a framework for query-efficient policy learning through easy-to-label queries. An overview is shown in Fig. 2. Queries compare trajectories from different policies. Each query includes a preference label and a discriminability label assessing whether superiority can be determined. To enhance policy distinction, a discriminator is trained with the preference reward model to guide query sampling and promote diverse policies.

The following sections describe the proposed framework. Section IV-A introduces its main components. Using the reward functions described in Section IV-B, DAPPER learns policies that reflect human preferences through the iterative execution of three steps: (1) Policy Learning, (2) Query Collection, and (3) Reward Update. Starting from Section IV-C, the subsequent sections provide details of each learning step. The overall procedure is summarized in the pseudo-code shown in Algorithm 1.

A. Framework Components

DAPPER consists of the following components.

Policy $\pi_i(s; \theta_i)$: A policy used for learning human preferences is parameterized by θ_i . Multiple policies with distinct parameters θ_i are trained for use in query comparisons, and the resulting policies are collected in the policy dataset $\Pi \ni (\pi_1, \pi_2, \dots, \pi_i, \dots)$.

Trajectory Generator $\mathcal{G}(\pi)$: A function that samples a trajectory τ from a policy π . This function is used to make policy-based queries and also serves as input to the discriminator for comparing policies. We define the generator as: $\mathcal{G}(\pi) = \tau \sim \pi$, where $\tau = (s_0, a_0, s_1, \dots, a_{T-1}, s_T)$. Here, T denotes the episode length.

Preference Reward Model $R^H(s, a; \phi)$: A reward function model for learning human preferences parameterized by ϕ , defined identically to PbRL as described in Section III-B. Based

Algorithm 1: DAPPER

```

Set parameters described in Table I
Set policy dataset  $\Pi$  and query dataset  $\mathcal{D}$ 
while  $i = 1, 2, \dots, \text{until converge}$  do
  Set current policy  $\pi_i$  and rollout buffer  $\mathcal{B}_i$ 
  for  $j = 1, 2, \dots, J$  do
    for  $e = 1, 2, \dots, E$  do
      for  $t = 1, 2, \dots, T$  do
        Get observation  $s_t$  from environment
        Take action  $a_t$  by  $\pi_i(s_t)$ 
        Calculate reward  $r_t$  by Eq. (4)
        Push  $(s_t, a_t, r_t)$  to  $\mathcal{B}_i$ 
      Update  $\pi_i$  by RL update scheme (e.g. [3]) using  $\mathcal{B}_i$ 
  Store learned policy to policy dataset  $\pi_i \rightarrow \Pi$ 
  for  $n = 1, 2, \dots, N$  do
    Choose policy  $\pi_x$  by Eq. (6)
    Generate trajectories  $\tau_x, \tau_i$  and conduct queries
    Store label  $y$  to query dataset  $(\tau_x, \tau_i, y) \rightarrow \mathcal{D}$ 
  Update preference reward model  $R^H$  by Eq. (7)
  Update discriminator  $D$  by Eq. (8)

```

TABLE I: Learning Parameters in Experiments

Para.	Meaning	Value
α	Coefficient of Query Sampling	10^{-3}
β	Coefficient of Policy Exploration	0.6
γ	Discount factor of RL	0.99
N	Query number per iteration	10
T	Number of steps per episode	128
E	Number of episodes per one policy update cycle (Parallel environment number)	128
J	Number of policy learning iterations to train π_i	300

on this reward function, human preferences are estimated as shown in Eq. (1).

Discriminator $D(\mathcal{G}(\pi_i), \mathcal{G}(\pi_j); \psi)$: A model that estimates the probability of whether preferences of policy pairs are discriminable by humans parameterized by ψ . By learning from previously trained policies, it functions as a reward estimation, ensuring that new policy remain distinguishable. Additionally, it is used to prioritize policy sampling when selecting policies for queries.

B. Reward Formulation

Our framework proposes policy-based queries, in contrast to previous works that employ trajectory-based queries from a single policy as shown in Eq. (1). Accordingly, we define human preference as follows:

$$P(\pi_i \succ \pi_j) = \frac{\exp[R^H(\mathcal{G}(\pi_i))]}{\exp[R^H(\mathcal{G}(\pi_i))] + \exp[R^H(\mathcal{G}(\pi_j))]} \quad (3)$$

In DAPPER, in addition to aligning with human preferences, the framework aims to learn policies whose preferences can be discriminated in queries. To this end, the following reward function r is defined from preference reward model R^H and discriminability reward R^D , as shown below:

$$r = (1 - \beta)R^H + \beta R^D \quad (4)$$

The hyperparameter β trades off alignment with preference reward learned from human queries and policy diversity.

Let (s_t, a_t) be a one sample set of current step t obtained from the interaction between the current policy π_i and the environment, the discriminability reward r^D is calculated from the discriminator D as follows:

$$R^D(s_t, a_t) = \frac{1}{N'} \sum_{\pi_x \in \Pi} D((\bar{\tau}_i, s_t, a_t), \mathcal{G}(\pi_x)), \quad (5)$$

where $\bar{\tau}_i = (s_0, a_0, s_1, \dots, a_{t-1}, s_{t-1})$ denotes the partial trajectory of policy π_i up to step t , and N' is the number of previously learned policies $(\pi_1, \pi_2, \dots, \pi_{i-1}) \in \Pi$ used for comparison.

C. Policy Learning

In this process, the policy π_i for the current iteration i is initialized and trained until the reward r defined in Eq. (4) converges. Once training converges, the learned policy is added to the policy dataset Π and used for query generation.

D. Query Collection

In this process, queries are generated from the learned policies, and labels are collected through the queries. For query generation, DAPPER utilizes the discriminator D to perform prioritized sampling of policies used in queries, aiming to enhance query discriminability. Specifically, one of the policy pairs is selected from the current policy π_i , which has high discriminability against past policies learned through R^D . The other policy π_x is selected from the policy dataset $\Pi \ni (\pi_1, \pi_2, \dots, \pi_{i-1})$. At this stage, the discriminability probability $D(\mathcal{G}(\pi_x), \mathcal{G}(\pi_i))$ between the current policy π_i and a past policy $\pi_x \in \Pi$ is estimated. Based on this probability, policies for queries are sampled according to the sampling probability p_x defined as follows:

$$p_x = \frac{\exp[\alpha D(\mathcal{G}(\pi_x), \mathcal{G}(\pi_i))]}{\sum_{j=1}^N \exp[\alpha D(\mathcal{G}(\pi_x), \mathcal{G}(\pi_j))]} \quad (6)$$

where α is a temperature parameter that controls sampling sharpness: higher values of α favor more discriminable policies, while lower values result in more uniform sampling. Sampling is performed without duplication of any policy pairs already present in the query dataset \mathcal{D} .

Then, using the sampled policies, a state-action trajectory τ is generated from the environment, and the label $y = P(\pi_x \succ \pi_i)$ is collected. In this process, the label y takes values from $\{0, 0.5, 1\}$, where $y = 1$ indicates that π_x is preferred over π_i , $y = 0$ indicates that π_i is preferred over π_x , and $y = 0.5$ signifies that both policies are considered equally preferable or that ranking them is difficult. The obtained labeled data (π_x, π_i, y) is added to the query dataset \mathcal{D} .

E. Reward Update

1) *Execution Flow*: Using the query dataset \mathcal{D} , the preference reward model R^H for calculating r^H and the discriminator D for calculating r^D are updated. After the models R^H and D have converged, they are utilized in training the next iteration's policy π_{i+1} . This process enables learning policies

TABLE II: Constraint parameters: Following the definitions in Section V-A2, the constraint target parameters for each task are defined. These parameters are estimated by feature extractor f from state-action trajectories τ . The threshold is set to a common value of $d = 10$ for all parameters.

Parameter	Target: x^*	Marge: g
Body Velocity of X Direction [m/s]	1	0.2
Body Velocity of Y Direction [m/s]	0	0.2
Body Rotation of Z Axis [degree]	0	10
Collisions on Bodies and Legs	0	0

that reflect human preferences from queries while also promoting the acquisition of policies that remain distinguishable from past policies Π .

2) *Preference Reward*: Using the separable label dataset $y \in \{0, 1\}$ from the query dataset \mathcal{D} , the preference reward model R^H is updated. The loss function is defined as follows:

$$\begin{aligned} \mathcal{L}_\phi &= -\mathbb{E}_{(\pi_i, \pi_x, y) \sim \mathcal{D}} \left[y \log P(\pi_x \succ \pi_i) \right. \\ &\quad \left. + (1 - y) \log P(\pi_i \succ \pi_x) \right]. \end{aligned} \quad (7)$$

3) *Discriminability Reward*: The discriminator D is trained on the query dataset \mathcal{D} , where the label y can take inseparable values $y \in \{0.5, \text{other}\}$. The loss function is defined as follows:

$$\begin{aligned} \mathcal{L}_\psi &= -\mathbb{E}_{(\pi_i, \pi_x, y) \sim \mathcal{D}} \left[y \log D(\mathcal{G}(\pi_x), \mathcal{G}(\pi_i)) \right. \\ &\quad \left. + (1 - y) \log (1 - D(\mathcal{G}(\pi_x), \mathcal{G}(\pi_i))) \right], \end{aligned} \quad (8)$$

where D is defined as a symmetric function, i.e., $D(\pi_x, \pi_i) = D(\pi_i, \pi_x)$. This formulation is equivalent to the standard binary cross-entropy loss, encouraging the discriminator D to assign high probability when two policy embeddings are distinguishable, and low probability otherwise.

V. EXPERIMENTS

The following experiments investigate whether the proposed framework can achieve higher query efficiency than existing methods by considering preference discriminability. In comprehensive comparative experiments, it is impractical for an actual human to manually label all queries due to the enormous time and human cost. To address this, we define a simulated annotator model and perform automatic labeling. We first formulate the simulated annotator's preference model and preference discriminability as mathematical expressions. Subsequently, we conduct a human subject experiment to validate the simulated annotator model by examining whether actual human responses to trajectory comparison queries can be explained by the preference discriminability defined in the model (at Section V-C). From the experimental results, we quantitatively determine an approximate threshold for preference discriminability and define multiple levels around it. These thresholds are then used in the subsequent PbRL experiments (from Section V-D) to automate query labeling with the simulated annotator model, enabling efficient evaluation of the proposed framework.

In our experiments, we aim to address the following research questions:

TABLE III: Terms of feature extraction: The second and third columns show the parameter ranges, while the fourth column presents the target feature values f^* . The remaining columns indicate the walking patterns and the features used in each pattern, where a checkmark means the feature follows f^* , and a numerical value means it follows the specified value. ‘‘Gait Cycle’’ and ‘‘Phase Difference’’ are calculated from the timing of foot-ground contact. ‘‘FL’’ (Front Left), ‘‘FR’’ (Front Right), ‘‘RL’’ (Rear Left), and ‘‘RR’’ (Rear Right) refer to the respective legs of a quadrupedal robot. ‘‘Gait Cycle’’ is the duration of a foot’s contact cycle, while ‘‘Phase Difference’’ denotes the delay between two legs’ contact timings. Both are updated at each foot contact, with the most recent value used at step t . ‘‘Phase Difference’’ is given as the delay ratio to the current gait cycle. ‘‘Body Height’’ and ‘‘Body Incline Angle’’ are computed at each step t .

Parameter	Min	Max	Target f^*	Posture	Trot Only	Normal	Crawl
Body Height [cm]	25	65	45	45	-	55	30
Body Incline Angle [degree]	-10	10	0	✓	-	✓	✓
Gait Cycle (FL) [s]	0.2	0.4	0.5	-	✓	✓	✓
Phase Difference (FL-FR) [%]	0	100	50	-	✓	✓	✓
Phase Difference (FL-RL) [%]	0	100	50	-	✓	✓	✓
Phase Difference (FL-RR) [%]	0	100	0	-	✓	✓	✓

- To what extent does the preference discriminability defined in the simulated annotator model explain actual human responses to trajectory comparison queries? (Section V-C)
- Does our framework improve query efficiency compared to previous methods? (Section V-D)
- How does our framework perform when the threshold of preference discriminability is set substantially above or below the range observed in actual human responses? (Section V-E)
- How does our framework affect computational time efficiency? (Section V-F)
- How does hyperparameter selection affect query efficiency and preference discriminability in our framework? (Section V-G)
- Can our framework effectively learn from VLM-based preference labels, which rely on world knowledge aggregated from large-scale data, when used as substitutes for human queries that may include labeling errors or indistinguishable cases? (Section V-H)
- Can the proposed method maintain its effectiveness as the number of feature dimensions that parameterize the target robot behaviors increases? (Section V-I)
- Is our framework applicable to sim-to-real policy transfer? (Section V-J)

A. Common Settings

All experiments are conducted using the ANYmal quadrupedal robot and the *legged_gym* RL-framework based on the IsaacGym physics simulator [1]. The trained policy is evaluated at 50 Hz predicting target joint positions, which are tracked by per joint PID controller running at 400 Hz. In all experiments, the learning parameter settings follow those of previous works, including simulation parameters, specifically state-action definition and network architectures [1]. Policy learning utilized Lagrangian Proximal Policy Optimization (LPPO) [3, 4] with parallel environments, and all experiments are conducted using a Nvidia GeForce RTX 4090 GPU. To stabilize the R^D signal particularly during the early training phase, we introduced Monte Carlo dropout [16] into the discriminator D . Furthermore, to prevent the discriminator D from overfitting to earlier policies when continuously updated

across iterations, their parameters were reset and retrained from scratch at each iteration.

1) *Feature Extraction*: To quantify the degree of alignment, we utilize the policy feature representation $f(s, a)$ extracted by a feature extractor introduced in previous PbRL works [17, 18]. In the feature space defined by f , we specify a target feature f^* that represents the desired behavior in accordance with human preferences. The target features are listed in Table III. In this paper, feature extraction is also performed when comparing different policies. In such cases, trajectory τ is generated by a policy π , then the average feature representation is calculated as $f(\pi) = \frac{1}{T} \sum_{(s,a) \in \tau} f(s, a)$, where T denotes the episode length.

2) *Constraint Terms of LPPO*: This paper adopts LPPO to learn a policy that satisfies constraints while suppressing unrealistic behaviors that may hinder accurate query labeling. Constraint term is defined by indicator function \mathbb{I} , allowing the cost function c to be expressed as follows: $c(s_t, a_t) = \mathbb{I}[|x_t - x^*| > g]$, where x_t denotes a current value of constraint target parameter, x^* is the target value, and g is the permissible margin of error. This formulation penalizes each timestep in which the deviation from the target exceeds the margin g . To enforce the constraint over an episode of length T , LPPO defines an threshold d on the total number of penalty-inducing steps: $\sum_{t=1}^T c(s_t, a_t) \leq d$. This constraint ensures that the number of steps in which the system deviates significantly from the desired behavior remains within an acceptable bound. The constraint terms are listed in Table II.

3) *Labeling Queries from Simulated Annotator Model*: To reduce the human cost associated with large-scale evaluation, we employ a simulated annotator model. Specifically, we define a simulated annotator’s preference model and preference discriminability as mathematical expressions, enabling automatic labeling. The simulated annotator assigns labels $y \in \{0, 0.5, 1\}$ based on the extracted policy features $f(\pi)$ and a predefined target feature f^* . All feature values are normalized to the range $[0, 1]$ according to the parameter specifications listed in Table III.

To estimate preference and discriminability labels, we define a unified discriminability metric based on the difference in feature errors between two trajectories τ_i and τ_j with respect to a target feature f^* . Specifically, we use the following scale-



Fig. 3: Analysis of actual human query responses for determining the threshold of preference discriminability d^{disc} in the simulated annotator model: A questionnaire was conducted using a console system, as shown in (b). At the top of the screen, a reference video was displayed, showing a policy that matched a predefined target behavior. Below the reference video, two candidate policy videos were shown for comparison. Participants were asked to choose the trajectory closer to the reference by selecting one of three options: “Left is better”, “Right is better”, or “Can’t decide”. The candidate policies were prepared by exhaustively combining the listed values of Body Height [cm] $\in \{29, 37, 45, 53, 61\}$ and Body Incline Angle [degree] $\in \{-8, -4, 0, 4, 8\}$. For each combination of these parameters, a policy was trained using DAPPER under an idealized setting where $d^{disc} = 0$, by optimizing the behavior to match the specified Body Height and Body Incline Angle values until convergence. Participants were instructed to compare only two features defined by f^* : Body Height and Body Incline Angle. Gait patterns were excluded. They could prioritize the features freely, but were told to consider both. Policy pairs were sampled from (a), and 105 predefined queries were presented sequentially. The queries were selected to ensure diversity in preference discriminability d^{disc} . All participants evaluated the same set of queries, and the order of the queries was randomized for each participant. The results shown in (c) represent the mean and standard deviation across five participants.

normalized L1-based distance:

$$d^{pref}(\tau, f^*) = \frac{1}{\sqrt{|f|}} \sum_{k=1}^{|f|} |f_k(\tau) - f_k^*|, \quad (9a)$$

$$d^{disc}(\tau_i, \tau_j, f^*) = |d^{pref}(f(\tau_i), f^*) - d^{pref}(f(\tau_j), f^*)|, \quad (9b)$$

where $|f|$ denotes the number of feature dimensions. Here, the trajectory with the smaller feature error d^{pref} is labeled as more preferred ($y \in \{0, 1\}$), while the discriminability label ($y \in \{0.5, \text{other}\}$) is determined by whether preference discriminability d^{disc} exceeds a predefined threshold. If the value is above the threshold, the query is considered distinguishable;

otherwise, it is treated as indistinguishable.

B. Comparing Methods

We compared the following methods to verify the effectiveness of the proposed method:

- Baseline [6]: A method that generates queries from pairs of state-action trajectories obtained from different episodes using a single policy.
- SURF [7]: Same as the Baseline, queries are generated from a single policy. However, this method automatically labels queries when the preference values calculated from

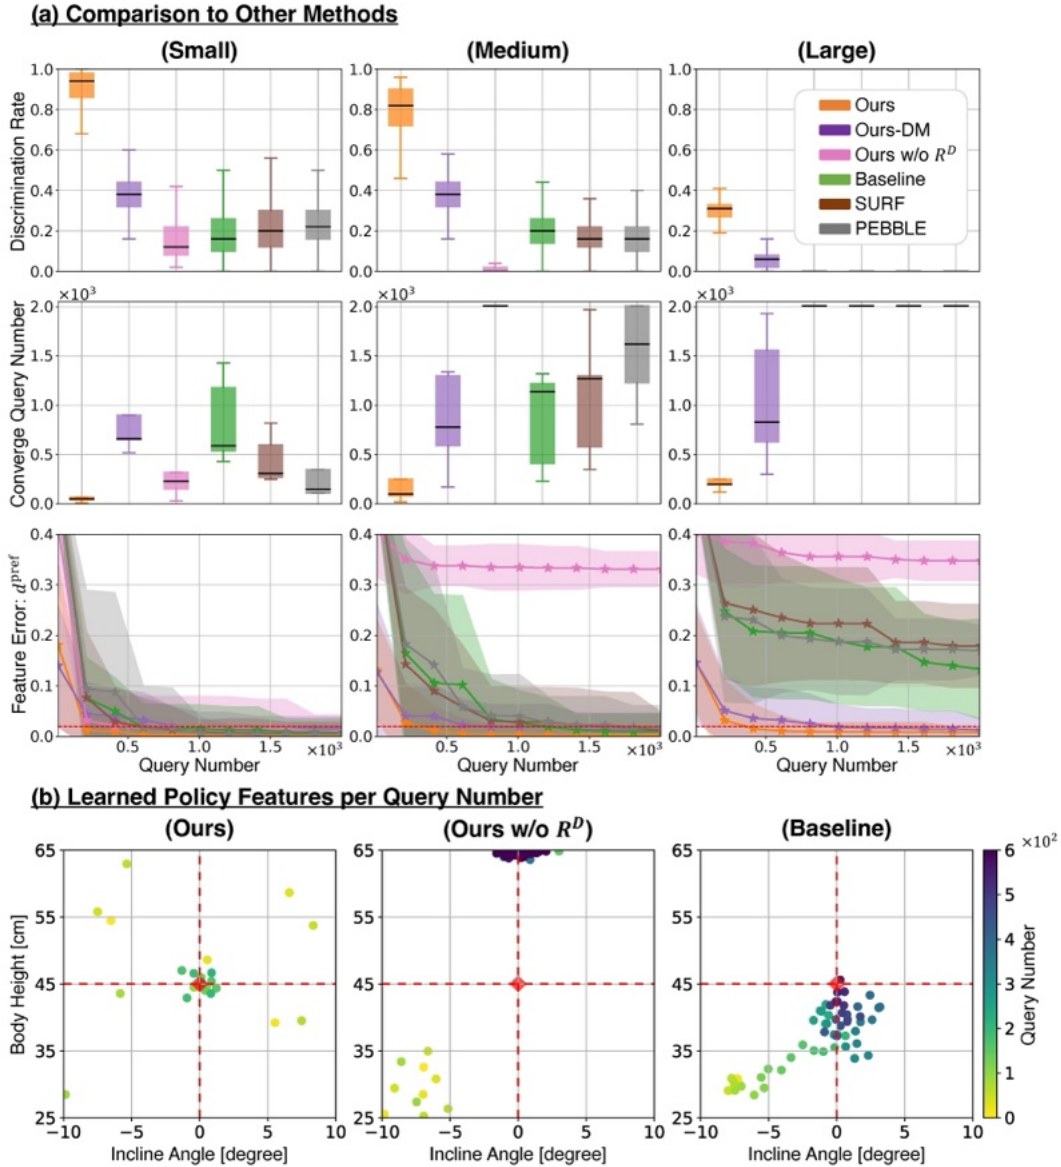


Fig. 4: Evaluation of query efficiency: **(a)** Comparison to previous works. **Top**: Discrimination rate of queries. **Middle**: Converge query number for reaching the feature error threshold ($d^{\text{pref}} < 0.02$), empirically defined by the convergence values of each method. If not converged within 2×10^3 queries, the value is clipped at 2×10^3 . The threshold is shown as a red dashed line in the bottom plot. **Bottom**: Minimum values of feature error d^{pref} achieved by the learned policies. “Small”, “Medium”, and “Large” denote levels of preference discriminability d^{disc} . Each curve shows the mean and standard deviation over five experiments. **(b)** Learned policy features per query number. This experiments use $d^{\text{disc}} = \text{“Medium”}$. The target feature f^* is located at the intersection of the red dashed lines. The number of plotted points corresponds to the queries collected until the feature error falls below the threshold ($d^{\text{pref}} < 0.02$), indicated by the red frame.

the learned reward function are sufficiently high, thereby increasing the number of training samples.

- PEBBLE [15]: Extends the Baseline with unsupervised pre-training and off-policy learning. The agent explores with intrinsic rewards maximizing state entropy to generate diverse behaviors before learning from queries. Past reward values are re-calculated and utilized after each reward model update, enabling sample-efficient policy learning.

We also verified the following two ablation methods of our framework:

- Ours w/o R^D : Same as the proposed framework, queries are generated from pairs of state-action trajectories ob-

tained from different policies. However, the reward function R^D , which encourages acquiring policies with distinct features, is not used.

- Ours with Distance Maximization (Ours-DM): This is the proposed framework with different R^D which is calculated based on the feature distance from previously learned policies using the following equation: $R^{D'} = \frac{-1}{N'} \sum_{\pi_x \sim \Pi} |f(\pi_i) - f(\pi_x)|$. This setting is included to evaluate whether simply maximizing the distance between policies is more effective than Ours.




Fig. 5: Performance comparison across different preference discriminability thresholds: The convergence query number denotes the point at which the learned policy reaches the feature error threshold ($d^{\text{pref}} < 0.02$). If convergence is not achieved within 2×10^3 queries, the value is clipped at 2×10^3 . Each curve shows the mean and standard deviation over five experiments.

C. Analysis of Actual-Human Preference Discriminability

1) *Settings*: This paper proposes a framework for generating queries that are easier for humans to distinguish, based on preference discriminability. This section examines whether actual humans can distinguish similar queries and how their ability declines with increasing similarity. To control for individual preference variability, we define a target feature f^* and treat the policy closer to f^* as the correct answer. We then analyze the proportion of indistinguishable responses ($y = 0.5$) as a function of d^{disc} to quantify actual-human preference discriminability. Additionally, we evaluate the accuracy of query labels when participants express a clear preference ($y \neq 0.5$). This experiment uses the “Posture” feature values defined in Table III.

2) *Results*: The experimental results are shown in Fig. 3. As seen in Fig. 3(c), when d^{disc} is low, around 60% of queries were judged as indistinguishable. As d^{disc} increases, distinguishability improves, reaching 100% for all participants when $d^{\text{disc}} > 0.4$. Label accuracy for distinguishable queries exceeded 80% at all d^{disc} levels, and reached 100% when $d^{\text{disc}} > 0.3$. These results suggest that $d^{\text{disc}} > 0.4$ ensures consistent discrimination, while lower values depend on individual ability. Based on this, we define three threshold levels as criteria for automatic labeling in the subsequent sections: “Small” (0.2), “Medium” (0.3), and “Large” (0.4).

D. Comparison of Query Efficiency with Other Methods

1) *Settings*: The proposed framework improves query efficiency by utilizing queries from multiple policy trajectories instead of a single policy trajectory, thereby enhancing query discriminability to improve query efficiency. To verify its effectiveness, we compare the query discrimination rate and the number of queries required to reach the predefined target feature f^* (as shown in Table III) against the other methods described in Section V-B. We evaluate all methods across three levels of preference discriminability d^{disc} to examine their robustness under increasingly difficult query discrimination conditions. This experiment uses target feature values of “Posture” as defined in Table III.

2) *Results*: As shown in Fig. 4(a), Ours (DAPPER) achieves the target features with the fewest number of queries. When queries become harder to distinguish due to increased preference discriminability d^{disc} , the number of required

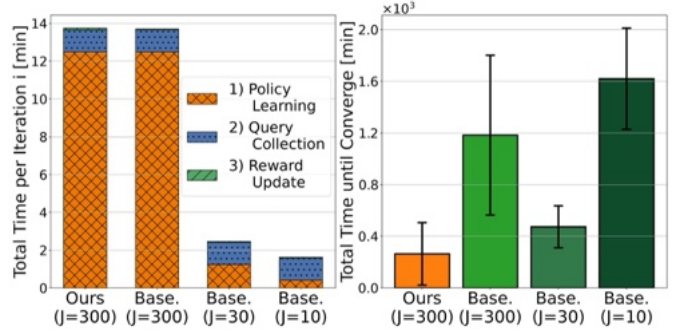


Fig. 6: Analysis of computational time: (Left) Total computational time per iteration i . As shown in Fig. 2 and Algorithm 1, each iteration i consists of three processes: “1) Policy Learning”, “2) Query Collection”, and “3) Reward Update”. In this plot, the stacked bar represents the total computational time, computed as the sum of the mean time spent on each process (averaged across the five trials). Note that the “3) Reward Update” process takes only a few seconds and is barely visible in the plot. (Right) Total computational time until convergence, defined as the point where the learned policy’s feature error satisfies the threshold ($d^{\text{pref}} < 0.02$). The bars show the mean time to converge, and the error bars denote the standard deviation across the five trials. As shown in Algorithm 1, the parameter J represents the number of policy learning iterations used to train π_i in “1) Policy Learning” within each cycle. The computation time for “2) Query Collection” is estimated at about 70 seconds based on Fig. 3, where one human response took around 7 seconds on average and each iteration included $N = 10$ queries.

queries increases. In particular, when d^{disc} is “Large”, Ours requires the fewest queries; by contrast, the ablation variant requires over three times more queries, and the trajectory-based method fails to learn the target features within the given number of queries. Ours consistently achieves the highest discrimination rate, and under the “Large” condition, only Ours and Ours-DM are able to generate a sufficient number of discriminable queries for learning a policy that closes the target features.

As shown in Fig. 4(b), the Baseline suffers from policy bias that constrains exploration to the features of previously learned policies and slows progress toward the target features f^* . In contrast, Ours w/o R^D initially learns policies with similar feature representations because it lacks R^D , but retraining policies from scratch at each reward model update reduces the influence of policy bias and occasionally allows larger jumps to new feature regions. However, without the discriminability reward R^D to actively encourage exploration, the policy tends to converge to local optima due to lack of discriminable queries. Meanwhile, Ours (i.e., with R^D) integrates training-from-scratch with the discriminability reward, consistently acquiring diverse and distinguishable policies that successfully converge toward the target feature value f^* .

E. Analysis of Sensitivity to Preference Discriminability Threshold

1) *Settings*: As shown in Fig. 4, our framework demonstrates improved query efficiency under preference discriminability thresholds corresponding to the range observed in actual human responses. In this section, we investigate how the framework behaves when the threshold is set beyond this human-observed range, either higher or lower, to observe its behavior under more challenging or relaxed discriminability

conditions. For comparison, the Baseline is used to examine how our framework behaves differently under these varying threshold conditions. This experiment uses the target feature value of “Posture”.

2) *Results*: As shown in Fig. 5, Ours consistently outperforms the Baseline across a wide and fine-grained range of thresholds. The results also show that the performance gap narrows when the preference discriminability threshold is very low. However, even when the threshold is set to zero, meaning that all queries are distinguishable, the Baseline still performs worse than Ours because policy bias prevents it from learning substantially different behaviors within a small number of iterations i . Conversely, as the threshold increases and discriminable queries become scarce, policy learning becomes increasingly difficult, and Ours ultimately fails to learn policies.

F. Analysis of Computational Time Efficiency

1) *Settings*: As shown in Fig. 4, this paper demonstrates improvements in preference discriminability and the resulting query efficiency. However, achieving these improvements introduces additional computational components. Specifically, a new policy is retrained from scratch at each iteration i to avoid policy bias, and a discriminator D is introduced as an additional component to estimate preference discriminability. In this experiment, we evaluate whether the gain in query efficiency achieved by our framework outweighs the additional computational cost in terms of total computation time until policy-learning converge, by examining whether it remains more time-efficient than the Baseline, which does not include these additional components. We evaluate the Baseline under different numbers of policy learning iterations ($J \in \{300, 30, 10\}$), since the Baseline does not initialize the policy at each iteration i and can continue learning with fewer policy learning steps. This experiment uses target feature values of “Posture” under the “Medium” preference discriminability condition.

2) *Results*: The results are shown in Fig. 6. Regarding total time per iteration i , the “(1) Policy Learning” process takes about 12–13 minutes, accounting for most of the total 13–14 minutes. In contrast, the “(3) Reward Update” requires only a few seconds, having a negligible effect on the total time. Reducing the number of policy learning iterations J naturally decreases the overall computation time per iteration i . Regarding the total time until convergence, with $J = 30$, the Baseline reduces the overall time to roughly 40% of that with $J = 300$, but it is still approximately twice as long as Ours. However, when J is reduced to 10, the time for “(1) Policy Learning” becomes smaller, but “(2) Query Collection” dominates, and more iterations are required for convergence, resulting in longer total training time. Overall, even when accounting for retraining from scratch and the additional discriminator D , Ours demonstrates superior time efficiency compared to the Baseline.

G. Influence of Different Query Exploration/Sampling

1) *Settings*: As shown in Eq. (4), the reward in DAPPER is adjusted by setting the ratio β between r^D and r^H to balance

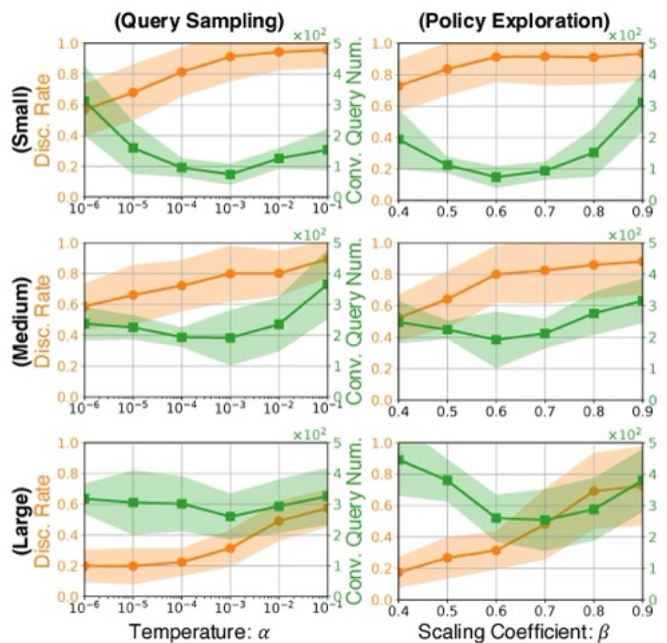


Fig. 7: Performance comparison of DAPPER components: (Left) different temperature patterns of query sampling, (Right) different scaling coefficient patterns of policy exploration. “Small”, “Medium”, and “Large” denote levels of preference discriminability d^{disc} . Each curve plots the mean and standard deviation over five experiments.

policy discriminability and convergence toward the target feature values. Additionally, the policy sampling from the policy dataset Π for query generation also affects discriminability performance (Eq. (6)). Therefore, this section comprehensively compares these parameters and evaluates their influence on convergence to target features. This experiment uses target feature values of “Posture” as defined in Table III.

2) *Results*: The results are shown in Fig. 7. Regarding query sampling, a higher temperature coefficient α increases the probability of selecting policies with high discriminability, thereby improving the overall discriminability rate. However, a large α results in biased sampling, as only easily distinguishable policies are selected by the discriminator D . Moderately reducing α improves query efficiency by encouraging diverse queries, which better support learning the preference reward R^H . Regarding policy exploration, a higher weight β promotes broader exploration, which is expected to facilitate reaching the optimal policy. However, excessive exploration degrades convergence, as also observed in the results of Ours-DM in Section V-D. Since $\alpha = 10^{-3}$ and $\beta = 0.6$ consistently yield good performance across different d^{disc} settings, we use them in all experiments.

H. Applicability of VLM-Based Preference Labels as Human Query Substitutes

1) *Settings*: Vision-language models (VLMs) provide an attractive alternative to human annotators for generating preference labels [19, 20]. Using VLMs instead of humans could significantly reduce annotation cost, but it also introduces challenges such as noisy or overconfident predictions. Highlighting this alternative is particularly relevant for PbRL,

You are a vision and statistics assistant.

Your job is to analyze robot body incline and height across videos, compare statistical differences, and decide which video is closer to the reference or whether the result is uncertain.

(Analysis Target)

- Video1 is the reference, while Video2 and Video3 are comparison targets.
- Estimate body incline (relative to the forward direction) and body height for Video2 and Video3 with respect to Video1.
- Each video has multiple frames; compute frame-wise values and analyze distributions.
- The body region is marked in red.

(Statistical Comparison)

- Test incline and height separately using Welch's t-test, based on absolute mean values.

(Criteria)

- $p < 0.1 \rightarrow$ significant difference
- $p \geq 0.1 \rightarrow$ no difference

(Decision Rule)

- If both incline and height show no difference, the final answer is "Can't decide".
- Otherwise, select the closer video (smaller absolute mean):
 - "Left is better" if Video2 is closer to the reference,
 - "Right is better" if Video3 is closer to the reference.
- If evidence is weak or uncertain, output "Can't decide".

Fig. 8: Prompt for VLM-based preference labeling: We used OpenAI GPT-5 model with the instruction text shown in the figure to generate preference labels. Three videos were used in each comparison: two candidate videos and one reference video. Each video is 3 seconds long at 30 Hz, from which three frames were sampled at 0.5, 1.5, and 2.5 seconds and downsampled from the original resolution of $1,600 \times 900$ pixels (the same resolution as in Fig. 3) to 800×450 pixels to reduce computational and resource costs. These processed frames, together with the textual instructions, were used as inputs. The model was then asked to output one of three labels: "Left is better", "Right is better", or "Can't decide".

where both sample efficiency and robustness to noisy labels are key for practical deployment. In this experiment, we conduct a preliminary study to examine whether our framework can leverage VLM-based preferences effectively. The prompt and configuration provided to the VLM are shown in Fig. 8. While VLMs offer a potential substitute for human preference annotations, they often tend to be overconfident when choosing between options and rarely select the "Can't decide" label. This behavior may stem from their underlying training or preference-oriented fine-tuning. To address this, as illustrated in Fig. 8, we assumed that the VLM's frame-wise estimates are independent and identically distributed and applied Welch's t-test. If no statistically significant difference was found, the query was mapped to the "Can't decide" label. This experiment used the target feature value of "Posture" as defined in Table III.

2) *Results*: As shown in Table IV, VLM-based preference labels exhibit both discriminability and accuracy limitations. The Baseline achieved only about 60% label accuracy, which is extremely low considering that the labels are binary. Moreover, nearly half of the queries could not be assigned a label, leading to almost no learning in terms of the feature error d^{pref} . These results indicate that VLMs tend to generate overconfident outputs, avoiding the "Can't decide" option and thus producing unreliable supervision. In contrast, our framework leverages preference discriminability to reduce the number of indistinguishable queries, improving label accuracy by approximately 20% over the Baseline. Consequently, the

TABLE IV: Performance evaluation with VLM-based preference labeling: Each trial was conducted with up to 1,000 queries. Each value represents the mean and standard deviation over five experiments.

Evaluation Metric	Ours	Baseline
Indiscriminable Rate [%]	23.6 ± 15.7	46.9 ± 16.0
Label Accuracy [%]	84.7 ± 12.8	61.9 ± 15.2
Minimum Feature Error d^{pref}	0.005 ± 0.002	0.318 ± 0.061

learned policy reached the feature error threshold ($d^{\text{pref}} < 0.02$) defined in Fig. 4, demonstrating that useful signals can be extracted from VLM-generated labels when query design accounts for discriminability.

I. Scalability with Respect to Feature Dimensions

1) *Settings*: This section evaluates the applicability of the proposed framework to different tasks by testing its effectiveness on two additional task settings with a greater number of feature parameters. The experiment uses the target feature values of "Trot Only" and "Normal", as defined in Table III. In these settings, the number of features required for learning from queries is increased to four and six, respectively, compared to two in the tasks used in the previous section, making the learning process more challenging. All experiments were conducted under the "Medium" preference discriminability condition.

2) *Results*: As shown in Fig. 9, Ours achieves faster convergence to the target feature f^* across all tasks. As the number of features $|f|$ increases, the difficulty of discrimination also rises, and the number of required queries tends to increase accordingly. In particular, previous methods struggle significantly with the "Normal" setting, which requires learning six features, resulting in very few identifiable queries. While the learning difficulty may ultimately depend on how close the policy needs to get to the target feature f^* , it is possible that even our proposed framework may face limitations as the feature dimensionality increases further. Investigating this potential limitation remains an important direction for future work.

J. Real-World Application

In this section, we apply the learned control policies from simulation to real-world environments using a sim-to-real transfer approach [1]. The transfer policy is trained in a domain-randomized environment over sensor noise and physical parameters, using well-validated parameter sets from previous work [1]. During policy learning, the feature representation f used as input to the learned preference reward model R^H is treated as privileged information available only in simulation for reward computation, and therefore domain randomization is not applied to f . As shown in Fig. 10, the learned policies were applied in the real world. This validation follows the same sim-to-real setup as in previous work [1], without additional tunings. For details of the robot motion, please refer to the supplementary videos.




Fig. 9: Comparison of query efficiency in various tasks: “Trot only” means using feature extraction of trot patterns, “Normal” adds feature extraction of body posture as shown in Table III. This experiments use preference discriminability $d^{\text{disc}} = \text{“Medium”}$, meaning that if a policy’s features are closer than a specified threshold, the policy is considered non-identifiable ($y = 0.5$). Each curve plots the mean and standard deviation over five experiments.




Fig. 10: Real-world execution of policies trained in simulation: The “Crawl” walking pattern listed in Table III is trained in simulation, and the learned policies were deployed on the ANYmal quadrupedal robot.

VI. DISCUSSION

Robustness to Label Noise in Preference Annotations:

This paper focuses on preference discriminability and proposes a robust learning framework that performs effectively even under high discriminability thresholds (i.e., “Large”). In practice, however, some queries within these thresholds may not be labeled as indistinguishable by human annotators but instead receive incorrect superiority labels, as shown in Fig. 3(c). In addition, as shown in Table IV, we observed that the proposed framework can still operate effectively even when the generated labels include a certain amount of mislabeling (approximately 15%), suggesting robustness to imperfect annotations. This type of labeling noise has also been studied extensively in several prior works [8, 9]. Building on these insights, integrating such noise-robust methods with our framework presents a promising direction for future research.

Sensitivity to Hand-crafted Feature Design and Future Direction Toward Automated Extraction: Some consideration is necessary when defining the terms of feature extraction [18]. In particular, the success of our framework still relies on a

set of carefully designed features, and the overall performance is sensitive to the quality of this feature selection. If multiple similar features are selected and their optimization objectives happen to conflict, it may hinder policy learning. For example, if satisfying different features requires the robot to exhibit opposing behaviors, achieving the target may become difficult or even infeasible. Likewise, if the selected features are noisy or fail to capture informative aspects of user preferences, the learned policy can become suboptimal. While such cases are unlikely in well-structured tasks, they remain a possibility, and attention to feature design is crucial. A promising future direction is to reduce this dependence on hand-crafted features by integrating automated feature extraction or richer representations learned from high-dimensional sensor data.

Trade-off Between Computational Cost and Policy Diversity: An additional aspect to consider is the computational cost associated with retraining a new policy from scratch after each reward update in order to learn policies with different feature representations. While this approach is more expensive than continuing from previously trained policies, it was adopted to avoid policy bias and to ensure diversity among the learned policies. This highlights a trade-off between computational efficiency and policy diversity: retraining increases training time but provides more reliable exploration and reduces the risk of converging to biased behaviors. We recognize that this trade-off becomes a critical concern when extending the framework to more complex and computationally demanding domains, such as large-scale decision-making problems that require semantic or linguistic understanding. To mitigate this limitation, potential directions include warm-starting policy initialization within a range that avoids policy bias, or leveraging meta-learning techniques to initialize policies in a cost-effective manner. These strategies could help reduce computational overhead while maintaining sufficient diversity across learned policies, making the framework more scalable to large-scale or high-dimensional applications.

Balancing Low-level Constraints and High-level Preferences: In our framework, policy constraints (Table II) are used to enforce basic requirements such as movement speed or safety, rather than to represent human preferences. These constraints do not need to be strictly maximized, but only sufficiently satisfied within a margin. This approach is particularly effective because such low-level objectives are easy to define and help improve query discriminability. In contrast, high-level preferences such as motion style are difficult to specify explicitly, and are therefore learned through PbRL. This separation allows the system to delegate well-defined tasks to constraints, while using PbRL to capture complex, preference-based behaviors.

Stability of the Discriminability Reward During Early Training: A potential concern in our framework is the stability of the discriminability reward R^D during the early stages of training. Since the discriminator is initially trained on only a small number of labels, its preference discriminability estimates may be noisy or biased, which could misguide policy exploration and lead to suboptimal behaviors. This sensitivity can become more noticeable as the input and output dimensionality of the discriminator increases, making it harder to

obtain reliable preference discriminability scores from limited data. In our experiments, however, such instability was not severe, partly because the feature dimensionality was relatively small and the discriminator D incorporated Monte Carlo dropout [16], which helped mitigate instability in its outputs. Nevertheless, instability remains a possible challenge in more complex scenarios, and developing mechanisms to further stabilize the R^D represents an important direction for future work, for example by incorporating dimensionality reduction techniques or attention mechanisms.

Toward Capturing Richer and Higher-dimensional Human Preferences: While our framework takes a step toward making PbRL more practical by explicitly addressing query discriminability, it still falls short of fully capturing human preferences. The feature dimensionality in our experiments (up to six) may overlook subtle aspects of human preference, such as smoothness, energy efficiency, or style. Moreover, learning a satisfactory policy still requires hundreds of queries, although Fig. 4 shows that our framework reduces this number to about one-tenth of that required by previous approaches. These findings indicate that further progress toward real human experiments will require richer, higher-dimensional preference features and additional techniques to reduce the query burden. In addition, when features become higher-dimensional, humans may perceive a wider threshold of indistinguishability, which corresponds to a larger threshold of d^{disc} . This phenomenon was also observed in Fig. 5, where increasing the threshold step by step beyond the real human range (Fig. 3) eventually caused the our framework to fail to learn policies. Developing mechanisms that can adaptively handle such broader indistinguishability ranges will be an important direction for future research.

Limitations and Opportunities of VLM-based Preference Annotation: Another important point emerging from our VLM-based experiments is the challenge of relying on VLMs as annotators. Limitations such as weak spatial understanding and overconfidence in discrete choices restrict their current reliability for preference labeling [19, 20]. Our evaluation should therefore be regarded as a basic proof-of-concept, representing an initial step to verify feasibility. While preliminary, these findings highlight both the potential and the current limitations of integrating VLM feedback, and they point to opportunities for future research on more robust methods in PbRL.

Domain Scope and Future Extensions Beyond Quadruped Locomotion: Our experiments have primarily focused on quadruped locomotion tasks, as this domain provides a representative and challenging benchmark for evaluating PbRL frameworks. Quadruped robots require learning diverse gait patterns under high-dimensional dynamics, making them a suitable testbed for studying query efficiency and policy diversity. While the proposed framework is not limited to locomotion and could, in principle, be extended to other robotic domains such as manipulation or humanoid control, a thorough exploration of these directions lies beyond the scope of this paper. We consider this an important avenue for future work, as demonstrating the applicability of our framework across a broader range of skills and platforms would further strengthen its generality.

VII. CONCLUSION

This paper proposed DAPPER, a framework that improves query efficiency by incorporating human preference discriminability and trajectory diversification achieved by multiple policies. DAPPER trains multiple policies and learns a discriminator to generate queries that are easier for humans to label, addressing the limited diversity problem caused by policy bias inherent in previous trajectory-based approaches. Experiments on simulated and real-world legged robot tasks demonstrated that DAPPER significantly improves query efficiency compared to previous methods, particularly under challenging discriminability conditions.

REFERENCES

- [1] N. Rudin, D. Hoeller, P. Reist, and M. Hutter, "Learning to Walk in Minutes Using Massively Parallel Deep Reinforcement Learning," in *Conference on Robot Learning*, 2022, pp. 91–100.
- [2] C. Li, M. Vlastelica, S. Blaes, J. Frey, F. Grimmering, and G. Martius, "Learning Agile Skills via Adversarial Imitation of Rough Partial Demonstrations," in *Conference on Robot Learning*, 2023, pp. 342–352.
- [3] L. Zhang, L. Shen, L. Yang, S.-Y. Chen, B. Yuan, X. Wang, and D. Tao, "Penalized Proximal Policy Optimization for Safe Reinforcement Learning," in *International Joint Conference on Artificial Intelligence*, 2022.
- [4] A. K. Jayant and S. Bhatnagar, "Model-based Safe Deep Reinforcement Learning via a Constrained Proximal Policy Optimization Algorithm," *Advances in Neural Information Processing Systems*, vol. 35, pp. 24 432–24 445, 2022.
- [5] W. Wang, C. Zeng, H. Zhan, and C. Yang, "A Novel Robust Imitation Learning Framework for Complex Skills with Limited Demonstrations," *IEEE Transactions on Automation Science and Engineering*, 2024.
- [6] P. F. Christiano, J. Leike, T. B. Brown, M. Martic, S. Legg, and D. Amodei, "Deep reinforcement learning from human preferences," in *Advances in Neural Information Processing Systems*, 2017, pp. 4299–4307.
- [7] J. Park, Y. Seo, J. Shin, H. Lee, P. Abbeel, and K. Lee, "SURF: Semi-Supervised Reward Learning with Data Augmentation for Feedback-Efficient Preference-Based Reinforcement Learning," in *International Conference on Learning Representations*, 2022.
- [8] J. Cheng, G. Xiong, X. Dai, Q. Miao, Y. Lv, and F.-Y. Wang, "RIME: Robust Preference-based Reinforcement Learning with Noisy Preferences," in *International Conference on Machine Learning*, 2024, pp. 8229–8247.
- [9] X. Hu, J. Li, X. Zhan, Q.-S. Jia, and Y.-Q. Zhang, "Query-Policy Misalignment in Preference-Based Reinforcement Learning," in *International Conference on Learning Representations*.
- [10] S. Casper, X. Davies, C. Shi, *et al.*, "Open Problems and Fundamental Limitations of Reinforcement Learning from Human Feedback," *Transactions on Machine Learning Research*, 2023.
- [11] A. Najar and M. Chetouani, "Reinforcement learning with human advice: a survey," *Frontiers in Robotics and AI*, vol. 8, p. 584075, 2021.
- [12] T. Osa, V. Tangkaratt, and M. Sugiyama, "Discovering Diverse Solutions in Deep Reinforcement Learning by Maximizing State-action-based Mutual Information," *Neural Networks*, vol. 152, pp. 90–104, 2022.
- [13] P. Swazinna, S. Udluft, and T. Runkler, "Overcoming Model Bias for Robust Offline Deep Reinforcement Learning," *Engineering Applications of Artificial Intelligence*, vol. 104, p. 104366, 2021.
- [14] E. Bıyık, M. Palan, N. C. Landolfi, D. P. Losey, and D. Sadigh, "Asking Easy Questions: A User-Friendly Approach to Active Reward Learning," in *Conference on Robot Learning*, 2020, pp. 1177–1190.
- [15] K. Lee, L. Smith, P. Abbeel, *et al.*, "PEBBLE: Feedback-Efficient Interactive Reinforcement Learning via Relabeling Experience and Unsupervised Pre-training," in *International Conference on Machine Learning*, 2021.
- [16] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning," in *International Conference on Machine Learning*, 2016, pp. 1050–1059.
- [17] K. Metcalfe, M. Sarabia, and B.-J. Theobald, "Rewards Encoding Environment Dynamics Improves Preference-based Reinforcement Learning," *arXiv preprint arXiv:2211.06527*, 2022.
- [18] C. Wirth and J. Fürtnkranz, "Preference-based Reinforcement Learning: A Preliminary Survey," in *Proceedings of the ECML/PKDD-13 Workshop*, 2013.

IEEE Robotics & Automation Magazine (RAM) paper, presented at ICRA 2026, Vienna, Austria. Cite as RAM paper.

- [19] J. Rocamonde, V. Montesinos, E. Nava, E. Perez, and D. Lindner, "Vision-Language Models are Zero-Shot Reward Models for Reinforcement Learning," in *International Conference on Learning Representations*, 2024.
- [20] B. Chen, Z. Xu, S. Kirmani, B. Ichter, D. Sadigh, L. Guibas, and F. Xia, "SpatialVLM: Endowing Vision-Language Models with Spatial Reasoning Capabilities," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 14 455–14 465.