

# Bayesian Constraint Inference from User Demonstrations Based on Margin-Respecting Preference Models

Dimitris Papadimitriou<sup>1</sup> and Daniel S. Brown<sup>2</sup>

**Abstract**—It is crucial for robots to be aware of the presence of constraints in order to acquire safe policies. However, explicitly specifying all constraints in an environment can be a challenging task. State-of-the-art constraint inference algorithms learn constraints from demonstrations, but tend to be computationally expensive and prone to instability issues. In this paper, we propose a novel Bayesian method that infers constraints based on preferences over demonstrations. The main advantages of our proposed approach are that it 1) infers constraints without calculating a new policy at each iteration, 2) uses a simple and more realistic ranking of groups of demonstrations, without requiring pairwise comparisons over all demonstrations, and 3) adapts to cases where there are varying levels of constraint violation. Our empirical results demonstrate that our proposed Bayesian approach infers constraints of varying severity, more accurately than state-of-the-art constraint inference methods. Code and videos: <https://sites.google.com/berkeley.edu/pbicrl>.

## I. INTRODUCTION

Frequently, robots and other autonomous agents are required to act in environments under the presence of constraints. In such cases, policies can be obtained by using constrained reinforcement learning (RL) algorithms [1], [2]. However, hand-specifying all constraints can be time consuming and error prone and some constraints may be user-dependent, precluding pre-specification of all relevant constraints. In environments in which some or all of the constraints are unknown, one might naturally aim at inferring them, a topic closely related to inverse reinforcement learning (IRL) [3]–[11].

Constraint inference in RL has been studied in both discrete [12], [13] and continuous [14] state spaces. However, the majority of these methods suffer from high computational complexity. The reason for this is that these methods are iterative, and a new policy must be optimized at each iteration, incurring a significant computational cost. To alleviate this bottleneck, in this paper we formulate the constraint inference problem as a preference-based learning problem based on a Markov Chain Monte Carlo (MCMC) algorithm. More specifically, we utilize a preference-based likelihood function

that can be evaluated efficiently, in order to calculate the likelihood of any proposed constraint. We further introduce a ranking over groups of preferences as a more natural framework for preferences in robotics applications. Finally, we wish to emphasize that not all constraints are of equal importance. For example, avoiding pedestrians vs. avoiding tree branches have very different levels of priority in autonomous driving. We borrow elements from maximum margin classification [15]–[17] that enable our method to discriminate between constraints whose violations have varying degrees of consequences.

## II. BACKGROUND AND RELATED WORK

Our work on constraint inference borrows elements from IRL and preference-based learning.

**Constrained Reinforcement Learning:** The typical constrained RL formulation obtains a policy  $\pi$  by maximizing the expected reward while satisfying constraints in expectation. Policies can be obtained by solving a constraint optimization problem [13]. One way to solve such an optimization problem, is by formulating the Lagrangian and solving the resulting min-max problem

$$\min_{r_p \leq 0} \max_{\pi} \mathbb{E}_{a \sim \pi} \left[ \sum_{t=1}^T \gamma^t r_{\theta}(s, a) \right] + r_p \left( \mathbb{E}_{a \sim \pi} \left[ \sum_{t=1}^T \mathbb{I}_C(s, a) \right] \right), \quad (1)$$

where  $r_{\theta}(s, a)$  is the reward function, parameterized by  $\theta$  and obtained at state  $s$  for action  $a$ ,  $\mathbb{I}_C$  is the indicator function of constraint violation for state  $s$  and action  $a$  and  $\gamma$  is the discount factor. The variable  $r_p \in (-\infty, 0]$  denotes the Lagrange multiplier. This Lagrange multiplier can also be viewed as the average penalty that the agent incurs for violating constraints. We draw inspiration from this formulation to model the environment constraints. More specifically, we assume that constraints are “soft” and that each time the agent violates one, the agent incurs a penalty reward  $r_p$ .

**Constraint Inference:** The topic of constraint inference is a relatively new area of research in control and RL. From a control perspective, a common approach in literature is to infer constraint parameters by optimizing some metric based on the KKT conditions [18], [19]. In the RL community, maximum likelihood based methods

<sup>1</sup>Dimitris Papadimitriou is with the Department of Mechanical Engineering, University of California at Berkeley, [dimitri@berkeley.edu](mailto:dimitri@berkeley.edu).

<sup>2</sup>Daniel S. Brown is with the School of Computing, University of Utah, [daniel.s.brown@utah.edu](mailto:daniel.s.brown@utah.edu).

have been developed that infer constraints in discrete environments [12]. A Bayesian approach similar to ours, that learns the posterior distribution of constraints allowing for active learning methods, is developed in [13]. The authors in [14], [20] propose alternating methods, of policy updates and constraint inference, that infer soft and hard constraints in continuous state-spaces. A maximum entropy approach for constraint inference is developed in [21] and an approach based on multi-task demonstrations in developed in [22]. Finally, the authors in [23] construct safe sets after observing expert demonstrations in environments with varying underlying reward functions, while in [24] the authors propose a constraint inference approach based only on sparse interventions from an expert. However, the previously mentioned approaches cannot infer constraints of varying magnitude, a capability that our approach provides.

**Preference-Based Learning:** When demonstrations of varying degrees of optimality are available, it might be appropriate to utilize a preference-based model which allows the agent to infer reward parameters by just observing preferences over trajectories [25], [26]. The authors in [27] propose a preference-based RL approach in which the agents utilize human feedback on their demonstrations in the form of ranking, to improve their policies. Based on the Bradley-Terry model [28], many deep preference learning frameworks have been developed and studied [10], [29]–[33]. Preference learning has also been extended to multimodal reward functions [34], offline RL [35], and model-based RL [36]. Generalizing from strict pairwise comparisons, the authors in [37] explore the benefits of providing feedback in the form of a scalar value, but do not consider constraint inference.

### III. LEARNING FROM PREFERENCES

This section presents the preference learning framework on which our model is based. We outline cases in which learning from preferences, even in the case when one has access to all possible pairwise comparisons, can be ineffective. We also introduce a less demanding group-wise comparison framework. Finally, we utilize tools from the large margin classification literature to train our models to mirror preferences of varying margins.

Throughout this paper, we assume that we have access to a dataset  $\mathcal{D}$  of  $N$  demonstrations  $\tau_i$ ,  $i = 1, \dots, N$ , of varying quality. To simplify notation, we assume that the reward function and constraints, only depend on the state. We further assume that the reward function can be modeled as the sum of a known nominal reward and an unknown penalty reward that is associated with the constraints. Following prior work [13], [14], we assume that the nominal reward is known and it is given by the inner product of the known nominal weights  $\mathbf{w}_n \in \mathbb{R}^{N_\phi}$

and the  $N_\phi$  features  $\phi(s)$

$$r_n(s) = \mathbf{w}_n^\top \phi(s), \quad (2)$$

and, similarly, the penalty reward is associated with the unknown penalty weights  $\mathbf{w}_p \in \mathbb{R}^{N_\phi}$  as follows

$$r_p(s) = \mathbf{w}_p^\top \phi(s). \quad (3)$$

We further associate each constraint with an unknown binary indicator variable  $\mathbf{c}_p \in \mathbb{R}^{N_\phi}$  on top of the weight value  $\mathbf{w}_p$ . This binary indicator captures whether a feature is a constraint feature or not. The inspiration behind our constraint modeling approach with a binary variable, comes from integer programming techniques in constraint optimization problems, where it is common to model non-convex constraints as a set of individual affine constraints [38], [39]. We denote the cumulative environment reward as  $r_\theta(s)$ ,  $\theta = \{\mathbf{c}_p, \mathbf{w}_p\}$ , to designate the dependence on the unknown penalty weight  $\mathbf{w}_p$  and unknown binary parameters  $\mathbf{c}_p$ . Thus, we have

$$r_\theta(s) = r_n(s) + r_p(s) = (\mathbf{w}_n + \mathbf{c}_p \circ \mathbf{w}_p)^\top \phi(s), \quad (4)$$

with  $\circ$  denoting the Hadamard product.

Providing feedback in the form of pairwise comparisons is a natural way for humans to teach robots and other AI agents. The most frequently used model to express pairwise comparisons is the Bradley-Terry model. This model assumes that the probability of demonstration  $\tau_i$  being preferred to demonstration  $\tau_j$  is modeled as

$$P(\tau_i \succ \tau_j) = \frac{e^{\beta \sum_{s \in \tau_i} r_\theta(s)}}{e^{\beta \sum_{s \in \tau_i} r_\theta(s)} + e^{\beta \sum_{s \in \tau_j} r_\theta(s)}}, \quad (5)$$

where  $\beta$  denotes the inverse temperature parameter of the softmax function. The parameters  $\theta$  of the reward function are then learned in such a way so that the preferences among the demonstrations are satisfactorily preserved. Given that the demonstrations can be of varying length, we propose to consider the mean reward per demonstration, with  $T_i$  denoting the length of demonstration  $i$ . Consequently, the log-likelihood of a certain  $\mathbf{w}_p$  and  $\mathbf{c}_p$  realization is given by

$$\begin{aligned} \mathcal{L}(\mathbf{w}_p, \mathbf{c}_p) &= \sum_{\tau_i \succ \tau_j} \log P(\tau_i \succ \tau_j) \\ &= \sum_{\tau_i \succ \tau_j} \log \frac{e^{\frac{\beta}{T_i} \sum_{s \in \tau_i} r_\theta(s)}}{e^{\frac{\beta}{T_i} \sum_{s \in \tau_i} r_\theta(s)} + e^{\frac{\beta}{T_j} \sum_{s \in \tau_j} r_\theta(s)}}. \end{aligned} \quad (6)$$

The log-likelihood objective function can then be used as a metric to infer the unknown penalty weight vector  $\mathbf{w}_p$  and binary variables  $\mathbf{c}_p$  that best justify the user provided preferences. Although this formulation has been used extensively in the RL community [10], [27], [29], the Bradley-Terry choice model can also have its limitations.

### A. Limitations of Learning from Preferences

A drawback of the Bradley-Terry model is that it fails to capture “how much more preferable” one choice is from another. To illustrate this, consider a simple preference learning task on the grid environment depicted in Figure 1a. The agent must navigate from the start to the goal state in an environment in which the orange and the red states must be avoided, with the red considered slightly worse. The grey states can be considered good states. The expert provides the agent with the following complete list of pairwise comparisons,  $\tau_1 \succ \tau_2$ ,  $\tau_2 \succ \tau_3$  and  $\tau_1 \succ \tau_3$ .

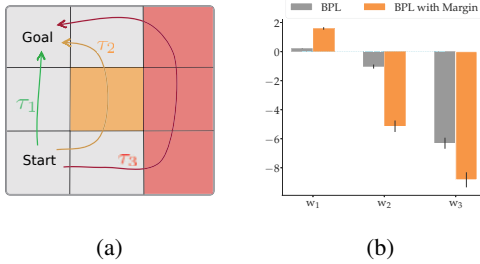


Fig. 1: (a) Example of three different types of trajectories. A preferable (green), a bad (orange) and a slightly worse than orange (red). (b) Weight values obtained from BPL and BPL with margins.

For simplicity, we assume that each state on the  $3 \times 3$  grid is associated with three binary features  $(x, y, z)$ . These are  $(1, 0, 0)$  for the grey,  $(0, 1, 0)$  for the orange and  $(0, 0, 1)$  for the red states. We estimate  $\mathbf{w}_p \in \mathbb{R}^3$  and  $\mathbf{c}_p \in \mathbb{R}^3$  using Bayesian preference learning (BPL) [30], the details of which can be seen in Algorithm 3 in the supplemental material on our website. We run BPL 100 independent times and report the average values of  $w_{pi}, i = 1, 2, 3$  in Figure 1b colored in grey. Conceptually, the Bradley-Terry model is a more appropriate modeling choice when the preferences among trajectories have similar significance. More specifically, the model fails to capture the fact that in our example, trajectory  $\tau_1$  is significantly preferable to the others, i.e.  $\tau_1 \succ \tau_2$  and  $\tau_1 \succ \tau_3$ . In this paper, we draw inspiration from the large margin classification literature [40]–[42], to propose a variation of the Bradley-Terry model that can be applied to preferences with non-uniform margins.

### B. Margin-Respecting Preference Learning

The likelihood function based on the Bradley-Terry model in Eq. (6) treats every preference equally and does not enable us to enforce the relative strength of preferences or desired margins over predicted rewards. This limitation, showcased in Figure 1a, is especially problematic in the constraint inference case when there

are multiple constraints of varying priority or importance. These constraints can be hard to learn even if there is a complete set of pairwise expert demonstration comparisons or a weighted likelihood function [25].

In the context of traditional  $K$ -class CNN classification, where the last dense layer of a neural network is denoted with  $\mathbf{W}$ , the inputs in that layer are denoted with  $\mathbf{f}_i, i = 1, \dots, N$  and the total number of data points is  $N$ , the softmax loss objective can be written as

$$\mathcal{L}(\mathbf{W}) = \frac{1}{N} \sum_{i=1}^N \log \frac{e^{\|\mathbf{W}_{y_i}\| \|\mathbf{f}_i\| \cos(\theta_{y_i,i})}}{\sum_{j=1}^K e^{\|\mathbf{W}_j\| \|\mathbf{f}_i\| \cos(\theta_{j,i})}}, \quad (7)$$

where  $\mathbf{W}_{y_i}$  denotes the row corresponding to the label  $y_i$  and  $\theta_{y_i,i}$  is the angle between vectors  $\mathbf{W}_{y_i}$  and  $\mathbf{f}_i$ . Using this expression, the following objective function modification has been proposed in the literature [15]–[17], that can encourage the creation of inter-class margins

$$\mathcal{L}(\mathbf{W}) = \frac{1}{N} \sum_{i=1}^N \log \frac{e^{\|\mathbf{W}_{y_i}\| \|\mathbf{f}_i\| \psi(\theta_{y_i,i})}}{e^{\|\mathbf{W}_{y_i}\| \|\mathbf{f}_i\| \psi(\theta_{y_i,i})} + \sum_{j \neq y_i}^K e^{\|\mathbf{W}_j\| \|\mathbf{f}_i\| \cos(\theta_{j,i})}}, \quad (8)$$

with the use of an appropriate function  $\psi(\cdot)$ . In the classification literature, the purpose of the  $\psi(\cdot)$  function is to alter the angle between the rows of the last layer weight matrix so that the margin among classes is enlarged. In our work, we follow a similar approach to [17] which assumes that  $\psi(\theta_{y_i,i}) = \cos(\theta_{y_i,i}) - m/(\|\mathbf{W}_{y_i}\| \|\mathbf{f}_i\|)$ , where  $m$  is a tunable hyperparameter, leading to a modified Bradley-Terry model that can be seen in Eq. (9).

In order to tune  $\psi(\cdot)$ , some measure of how much more preferable a demonstration is to another is required. Our approach allows the demonstrator to specify the desired relationship of the margins between groups. For our example in Figure 1a, the human specifies that the margin between  $\tau_1$  and  $\tau_2$  is approximately twice as large as that between  $\tau_2$  and  $\tau_3$ . We then tune  $\psi(\cdot)$  so that this relationship between the average rewards of the demonstrations is achieved under the learned constraint parameters. The improved results for an appropriate choice of  $m$  are shown in Figure 1b colored in orange.

### C. Learning from Grouped Preferences

Pairwise comparisons can be hard to obtain from humans for each possible pair of demonstrations. Hence, we propose to learn from coarsely grouped demonstrations as a more realistic alternative. In this case, a human can categorize demonstrations into distinct groups depending on their desirability. Given rankings among groups  $\mathcal{G}_i, i = 1, \dots, K$ , where the groups are ordered from most to least preferred, the modified Bradley-Terry model and the likelihood function, for  $\tau_i \in \mathcal{G}_k, \tau_j \in \mathcal{G}_\ell$



Fig. 2: Illustrative example of grouped rankings of demonstrations with different inter-class margins.

and  $k < \ell$ , can be written as

$$\begin{aligned} \mathcal{L}(\boldsymbol{\theta}) &= \mathcal{L}(\mathbf{w}_p, \mathbf{c}_p) = \sum_{\substack{\forall \tau_i \in \mathcal{G}_k, \forall \tau_j \in \mathcal{G}_\ell \\ \forall k < \ell}} \log P(\tau_i \succ \tau_j) \\ &= \sum_{\substack{\forall \tau_i \in \mathcal{G}_k, \forall \tau_j \in \mathcal{G}_\ell \\ \forall k < \ell}} \log \frac{e^{\frac{\beta}{T_i} \sum_{s \in \tau_i} r_{\boldsymbol{\theta}}(s) - m_{k\ell}}}{e^{\frac{\beta}{T_i} \sum_{s \in \tau_i} r_{\boldsymbol{\theta}}(s) - m_{k\ell}} + e^{\frac{\beta}{T_j} \sum_{s \in \tau_j} r_{\boldsymbol{\theta}}(s)}}}, \end{aligned} \quad (9)$$

where  $m_{k\ell}$  designates the margin parameter for groups  $k$  and  $\ell$ . Note that the soft-max expression is evaluated between all  $\tau_i \in \mathcal{G}_k$  and  $\tau_j \in \mathcal{G}_\ell$ . Thus, we leverage a small number of course-grained preference groupings over trajectories to automatically construct a large number of pairwise trajectory preferences.

#### IV. PREFERENCE-BASED CONSTRAINT INFERENCE

In this section, we outline the constraint inference approach for two cases, 1) when the environment constraint features are known and 2) when the constraint features assume a parametric form.

##### A. Constraint Inference with Known Features

Based on the likelihood in Eq. (9) we propose Preference-Based Bayesian Inverse Constraint Reinforcement Learning (PBICRL), an MCMC style algorithm that computes the posterior distribution of the constraints based on expert provided preferences over demonstrations. At each iteration, PBICRL randomly chooses a feature  $j$  and then either samples a new binary constraint value or a new penalty weight value (using a Gaussian proposal distribution). The frequency with which a binary or a penalty weight value is sampled is determined by the sampling frequency  $f_s$ . After sampling, the new likelihood of the demonstrations is evaluated and compared with its value at the previous iteration to determine whether the sample is accepted or not. Upon completion, the algorithm returns the maximum a posteriori (MAP) penalty weight vector, as well as the MAP indicators of whether each feature is a constraint.

##### B. Constraint Inference with Unknown Features

The previous section detailed the constraint inference approach in the case when the features were known. In this section, we propose a variation of PBICRL that further allows for inferring the parametric form of

---

#### Algorithm 1 PBICRL

---

- 1: **Parameters:** Number of iterations  $k$ , sampling frequency  $f_s$ ,  $\sigma$ , margins  $m_{k\ell}, \forall k \neq \ell \in \{1, \dots, K\}$
  - 2: Randomly sample: penalty weight  $\mathbf{w}_p \in \mathbb{R}^{N_\phi}$ , constraint feature  $\mathbf{c}_p \in \{0, 1\}^{N_\phi}$
  - 3: **for**  $i = 1, \dots, k$  **do**
  - 4:   Randomly sample feature  $j$  from  $\{1, \dots, N_\phi\}$
  - 5:   **if**  $i \bmod f_s \neq 0$  **then**
  - 6:     Set  $\mathbf{c}'_p[j] = -\mathbf{c}_p[j]$ ,  $\mathbf{w}'_p = \mathbf{w}_p$
  - 7:   **else**
  - 8:     Set  $\mathbf{w}'_p[j] = \mathbf{w}_p[j] + \mathcal{N}(0, \sigma)$ ,  $\mathbf{c}'_p = \mathbf{c}_p$
  - 9:   Compute Likelihood using (9)
  - 10:   **if**  $\log \mathcal{L}(\mathbf{w}'_p, \mathbf{c}'_p) \geq \log \mathcal{L}(\mathbf{w}_p, \mathbf{c}_p)$  **then**
  - 11:     Set  $\mathbf{w}_p = \mathbf{w}'_p$ ,  $\mathbf{c}_p = \mathbf{c}'_p$
  - 12:   **else**
  - 13:     Set  $\mathbf{w}_p = \mathbf{w}'_p$ ,  $\mathbf{c}_p = \mathbf{c}'_p$
  - 14:     w.p.  $\mathcal{L}(\mathbf{w}'_p, \mathbf{c}'_p) / \mathcal{L}(\mathbf{w}_p, \mathbf{c}_p)$
  - 15: **Return**  $\mathbf{w}_p, \mathbf{c}_p$
- 

the constraints. More specifically, we assume that the  $M_\phi$  constraint features  $\phi_\vartheta^j, j = 1, \dots, M_\phi$ , now have functional forms parameterized by  $\boldsymbol{\vartheta}$ . In this case, we assume that the environment reward is given by

$$\begin{aligned} r_{\boldsymbol{\theta}}(s) &= r_n(s) + r_{p,\boldsymbol{\vartheta}}(s) \\ &= \mathbf{w}_n^\top \boldsymbol{\phi}_n(s) + (\mathbf{c}_p \circ \mathbf{w}_p)^\top \boldsymbol{\phi}_\vartheta(s), \end{aligned} \quad (10)$$

with  $\mathbf{w}_p \in \mathbb{R}^{M_\phi}$ ,  $\mathbf{c}_p \in \mathbb{R}^{M_\phi}$  and  $\boldsymbol{\vartheta} \in \mathbb{R}^{M_\phi}$  being unknown parameters,  $\boldsymbol{\theta} = \{\mathbf{c}_p, \mathbf{w}_p, \boldsymbol{\vartheta}\}$ , and  $\mathbf{w}_n \in \mathbb{R}^{N_\phi}$  still corresponding to the known nominal weight vector of the known features, which are now denoted with  $\boldsymbol{\phi}_n$ . This version of PBICRL, outlined in Algorithm 2 in the supplemental material, uses a modified sampling scheme in which the unknown constraint parameters  $\boldsymbol{\vartheta}$  are also sampled and included in the calculation of the likelihood function in Eq. (9).

#### V. EXPERIMENTS

We utilize four different environments to quantify the performance of PBICRL: 1) a 2D point mass (PM) navigation environment, 2) the Fetch-Reach (FR) environment from OpenAI Gym and 3) the HalfCheetah (HC) and 4) Ant environments provided in Gymnasium [43]. Demonstrations in all environments were gathered from policies obtained with the Soft Actor-Critic (SAC) algorithm [44].

##### A. Point Mass Environment

First, we study a two-dimensional environment, pictured in Figure 3a, with the agent following a point mass model that navigates from some starting state towards the goal set, at the top right corner, while avoiding certain obstacles. At each time step the agent accrues a reward

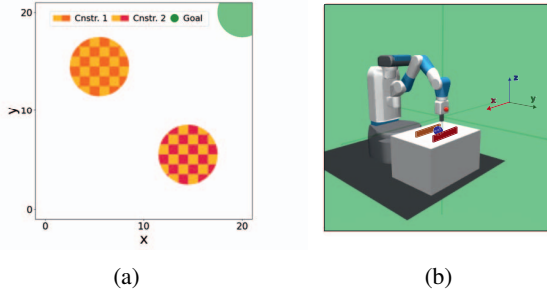


Fig. 3: 2D point mass navigational environment (a) and Fetch-Reach robot (b) environments.

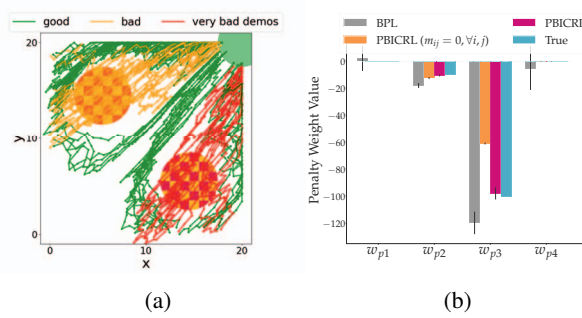


Fig. 4: (a) 2D point mass environment demonstrations. (b) Inference results for point mass environment. Results averaged over 5 seeds.

proportional to the inverse of their distance from the center of the goal set while at the same time accruing a living cost of  $-1$ . Once the point mass enters the goal region the episode is completed. There are a total of two constraints, the orange and the red one. The agent incurs a penalty of  $r_{p1}$  and  $r_{p2}$  upon entering the orange and red constraint regions, respectively. Furthermore, the actual unknown penalties satisfy  $r_{p2} \ll r_{p1}$ , as demonstrations that violate the red constraint are considered highly suboptimal. Although the penalties are unknown, we aim at recovering them accurately based on the expert provided pairwise comparisons and the relative margins among them. There are four features in the environment, namely: 1) the inverse distance from the target center, 2) an indicator variable on whether the orange constraint is violated, 3) an indicator variable on whether the red constraint is violated and 4) a binary variable indicating whether the target has not been reached yet. The ground-truth weight vector associated with the features is  $\mathbf{w} = \mathbf{w}_n + \mathbf{w}_p = [1, -10, -100, -1]$ .

Similar to [12], [13], we assume that we only have access to the nominal weight vector  $\mathbf{w}_n = [1, 0, 0, -1]$ . Using PBICRL we infer  $\mathbf{w}_p$  and  $\mathbf{c}_p$ . Figure 4a depicts some demonstrations that are categorized in three groups  $\mathcal{G}_1$ ,  $\mathcal{G}_2$  and  $\mathcal{G}_3$ , representing the good, bad and very bad

sets of demonstrations, respectively. We compare the performance of PBICRL with BPL [30], the details of which can be seen in Algorithm 3 in the supplemental material. In BPL we also assume knowledge of the nominal reward weight vector  $\mathbf{w}_n$  and the algorithm infers the penalty weight vector  $\mathbf{w}_p$ . The main difference between PBICRL and BPL is that the latter is missing the binary variables that signify the presence or absence of constraints.

We first run PBICRL, with all the margins  $m_{ij}, \forall i, j$  set to zero, and BPL using 60 demonstrations for each of the three groups and we average the results over 5 seeds. Figure 4b shows the inferred weights for PBICRL and BPL, respectively, along with the true values. Bayesian preference learning (BPL) fails to infer the correct weights, as for instance  $w_{p1}$  and  $w_{p4}$  should both be zero. On the other hand, the performance of PBICRL is significantly better, as it correctly identifies the two constrained features and the corresponding weight values are close to the true ones.

To capture how much undesirable the red constraint is, which should translate to a lower  $w_{p3}$  value, we further need to tune the  $m_{ij}$  margin parameters. To tune these margins, our method can utilize additional feedback from the expert. In this example, the expert provides feedback that the relative margins among demonstrations in  $\mathcal{G}_2$  and  $\mathcal{G}_3$  are approximately four times the magnitude of those between  $\mathcal{G}_1$  and  $\mathcal{G}_2$ . The parameters  $m_{ij}$  are automatically tuned so the cumulative rewards of the demonstrations under the inferred parameters approximately satisfy this condition. After tuning, the new inferred weights, as seen in Figure 4b, improve significantly.

### B. Fetch-Reach Robot

In this section, we utilize the Fetch-Reach robotic environment [43] and train it with a SAC policy to navigate towards the blue goal set, as shown in Figure 3b. The environment contains two rectangular constraint regions: the orange one is associated with a ground-truth violation penalty of  $-20$  and the considerably worse red one has a violation penalty of  $-100$ . The blue colored sphere designates the goal set towards which the robot should navigate. The environment has four features, namely: 1) the inverse distance from the center of the goal set, binary indicators for violating the 2) orange and 3) red constraints, and 4) a binary living indicator. The original ground-truth weight vector  $\mathbf{w} \in \mathbb{R}^4$  is  $[0.1, -20, -100, -5]$ , while the known nominal weight vector is  $\mathbf{w}_n = [0.1, 0, 0, -5]$  and the unknown constraint violation parameters  $\mathbf{w}_p$  and  $\mathbf{c}_p$  are to be inferred.

Figure 5b shows the results for the inferred constraint weights. The  $m_{ij}$  parameters were tuned under the assumption that the user provides feedback that the relative margins between the two pairs of groups  $\mathcal{G}_1, \mathcal{G}_2$

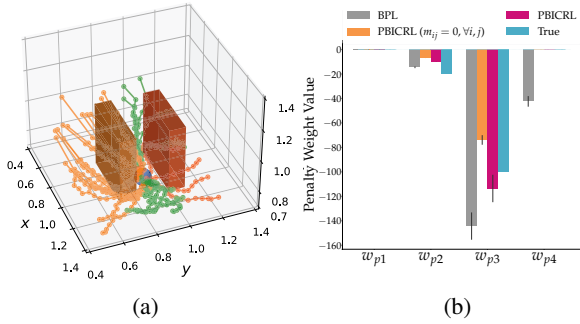


Fig. 5: (a) Fetch-Reach robot demonstrations. (b) Inference results for Fetch-Reach environment. Results averaged over 5 seeds.

and  $\mathcal{G}_2, \mathcal{G}_3$  are approximately the same. The original average rewards of the demonstrations, along with the ones for the inferred weights under  $m_{ij} = 0, \forall i, j$  and after tuning  $m_{ij}$ , can be seen in the supplemental material for both the 2D and Fetch-Reach environments.

### C. HalfCheetah and Ant with Parametric Constraints

In this section we consider parametric 1D halfspace constraint functions of the form  $z \leq \vartheta$ , where PBICRL, must also infer the parameter  $\vartheta$ . We focus on the  $z$  dimension as this is the primary axis along which the agent is incentivized to run forwards. This framework further allows us to compare our method to state-of-the-art constraint inference techniques such as [14], [20], where the single 1D constraint also assumes a parametric form. Given that in our experiments constraint violation is technically possible, we focus on comparing with [20], which is a soft constraint inference algorithm.

In the HalfCheetah and Ant environments [45] we test the performance of PBICRL in a scenario with a single constraint function, similar to [20]. More specifically, we assume that there exists an unobserved constraint at a particular location  $z$ . For the HalfCheetah we assume that the agent faces a constraint at the  $z = 8$  location and whenever  $z \geq 8$  the agent incurs a penalty reward of  $-50$ . For the Ant environment we assume that the states for which  $z \geq 10$  are constrained and associated with a penalty reward of  $-100$ .

The known features  $\phi_n$  of the HalfCheetah and Ant environments are the  $z$ -coordinate difference between two time steps and the square norm of the action vector. The corresponding nominal weights are  $\mathbf{w}_n = [20, -0.1]$ . In this setting, PBICRL infers the unknown constraint location  $\vartheta$  as well as the  $c_p, \mathbf{w}_p$  parameters associated with it. We compare the inferred constraints of PBICRL with Inverse Constraint Learning (ICL) [20] by reporting the Constraint Mean Squared Error (CMSE), which quantifies the error between the location of the actual and

the inferred constraints. For PBICRL, 100 good and 100 bad demonstrations were provided. In the ICL method, we used 200 good demonstrations some of which violated the constraint. Table I shows the CMSE values for both methods. PBICRL consistently estimates a constraint location very close to the true one, while ICL tends to infer a constraint close to the starting states of the demonstrations.

TABLE I: CMSE for PBICRL and ICL.

Env.	ICL	PBICRL
HC	$69.74 \pm 0.40$	<b><math>0.073 \pm 0.12</math></b>
Ant	$103.37 \pm 0.21$	<b><math>0.20 \pm 0.29</math></b>

Another benefit of our approach lies in its time complexity. Prior work alternates between policy learning and constraint inference [14], [20] leading to high time complexities. By contrast, our approach infers constraints in the above environments using an order of magnitude less time. Table II shows the time complexity of ICL and PBICRL for the HalfCheetah and Ant environments. All results have been averaged over 5 seeds.

TABLE II: Runtime in seconds for PBICRL and ICL.

Env.	ICL	PBICRL
HC	$20650 \pm 2982$	<b><math>3298 \pm 659</math></b>
Ant	$22703 \pm 2961$	<b><math>3165 \pm 12</math></b>

## VI. CONCLUSION AND FUTURE DIRECTIONS

We presented a novel Bayesian constraint learning algorithm (PBICRL) that utilizes preferences over groups of demonstrations. We also extended the classic Bradley-Terry choice model to allow different margins between preferences. Our empirical results show that PBICRL can more accurately and more efficiently infer constraints in environments under both known and unknown features than prior state-of-the-art approaches for both constraint learning and Bayesian preference learning.

One of the benefits of our Bayesian inference approach is that the posterior distribution can provide information about uncertainty and hence, could be used to design active learning approaches as well as design policies that satisfy certain safety criteria. One particular direction we consider interesting is the utilization of the preference margins, which inherently quantify the severity of individual constraints, in an active learning scheme that uses a query function that takes into consideration those margins. Future work also includes extending our approach to higher dimensional environments using unsupervised learning [46] and human-guided representation alignment [30], [32], [47] to learn feature representations that can be used for constraint learning.

## REFERENCES

- [1] J. Achiam, D. Held, A. Tamar, and P. Abbeel, “Constrained policy optimization,” in *International conference on machine learning*. PMLR, 2017, pp. 22–31.
- [2] S. Miryoosefi and C. Jin, “A simple reward-free approach to constrained reinforcement learning,” in *International Conference on Machine Learning*. PMLR, 2022, pp. 15 666–15 698.
- [3] S. Arora and P. Doshi, “A survey of inverse reinforcement learning: Challenges, methods and progress,” *Artificial Intelligence*, vol. 297, p. 103500, 2021.
- [4] P. Abbeel and A. Y. Ng, “Apprenticeship learning via inverse reinforcement learning,” in *Proceedings of the twenty-first international conference on Machine learning*, 2004, p. 1.
- [5] A. Y. Ng, S. Russell, et al., “Algorithms for inverse reinforcement learning” in *Icml*, vol. 1, 2000, p. 2.
- [6] M. Wulfmeier, P. Ondruska, and I. Posner, “Maximum entropy deep inverse reinforcement learning,” *arXiv preprint arXiv:1507.04888*, 2015.
- [7] B. D. Ziebart, A. L. Maas, J. A. Bagnell, A. K. Dey, et al., “Maximum entropy inverse reinforcement learning,” in *Aaai*, vol. 8. Chicago, IL, USA, 2008, pp. 1433–1438.
- [8] M. Lopes, F. Melo, and L. Montesano, “Active learning for reward estimation in inverse reinforcement learning,” in *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2009, Bled, Slovenia, September 7-11, 2009, Proceedings, Part II 20*. Springer Berlin Heidelberg, 2009, pp. 31–46.
- [9] D. Ramachandran and E. Amir, “Bayesian inverse reinforcement learning,” in *IJCAI*, vol. 7, 2007, pp. 2586–2591.
- [10] D. Brown, W. Goo, P. Nagarajan, and S. Niekum, “Extrapolating beyond suboptimal demonstrations via inverse reinforcement learning from observations,” in *International conference on machine learning*. PMLR, 2019, pp. 783–792.
- [11] E. Klein, M. Geist, B. Piot, and O. Pietquin, “Inverse reinforcement learning through structured classification,” *Advances in neural information processing systems*, vol. 25, 2012.
- [12] D. R. Scobee and S. S. Sastry, “Maximum likelihood constraint inference for inverse reinforcement learning,” *arXiv preprint arXiv:1909.05477*, 2019.
- [13] D. Papadimitriou, U. Anwar, and D. S. Brown, “Bayesian methods for constraint inference in reinforcement learning,” *Transactions on Machine Learning Research*.
- [14] S. Malik, U. Anwar, A. Aghasi, and A. Ahmed, “Inverse constrained reinforcement learning,” in *International Conference on Machine Learning*. PMLR, 2021, pp. 7390–7399.
- [15] W. Liu, Y. Wen, Z. Yu, M. Li, B. Raj, and L. Song, “Sphereface: Deep hypersphere embedding for face recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 212–220.
- [16] W. Liu, Y. Wen, Z. Yu, and M. Yang, “Large-margin softmax loss for convolutional neural networks,” *arXiv preprint arXiv:1612.02295*, 2016.
- [17] F. Wang, J. Cheng, W. Liu, and H. Liu, “Additive margin softmax for face verification,” *IEEE Signal Processing Letters*, vol. 25, no. 7, pp. 926–930, 2018.
- [18] G. Chou, N. Ozay, and D. Berenson, “Learning constraints from locally-optimal demonstrations under cost function uncertainty,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3682–3690, 2020.
- [19] D. Papadimitriou and J. Li, “Constraint inference in control tasks from expert demonstrations via inverse optimization,” *arXiv preprint arXiv:2304.03367*, 2023.
- [20] A. Gaurav, K. Rezaee, G. Liu, and P. Poupart, “Learning soft constraints from constrained expert demonstrations,” *arXiv preprint arXiv:2206.01311*, 2022.
- [21] M. Baert, P. Mazzaglia, S. Leroux, and P. Simoens, “Maximum causal entropy inverse constrained reinforcement learning,” *arXiv preprint arXiv:2305.02857*, 2023.
- [22] K. Kim, G. Swamy, Z. Liu, D. Zhao, S. Choudhury, and Z. S. Wu, “Learning shared safety constraints from multi-task demonstrations,” *arXiv preprint arXiv:2309.00711*, 2023.
- [23] D. Lindner, X. Chen, S. Tschiatschek, K. Hofmann, and A. Krause, “Learning safety constraints from demonstrations with unknown rewards,” *arXiv preprint arXiv:2305.16147*, 2023.
- [24] C. Basich, S. Mahmud, and S. Zilberstein, “Learning constraints on autonomous behavior from proactive feedback,” in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 3680–3687.
- [25] C. Wirth, R. Akrou, G. Neumann, J. Fürnkranz, et al., “A survey of preference-based reinforcement learning methods,” *Journal of Machine Learning Research*, vol. 18, no. 136, pp. 1–46, 2017.
- [26] K. Lee, L. Smith, A. Dragan, and P. Abbeel, “B-pref: Benchmarking preference-based reinforcement learning,” in *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021.
- [27] R. Akrou, M. Schoenauer, and M. Sebag, “Preference-based policy learning,” in *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2011, Athens, Greece, September 5-9, 2011. Proceedings, Part I 11*. Springer, 2011, pp. 12–27.
- [28] R. A. Bradley and M. E. Terry, “Rank analysis of incomplete block designs: I. the method of paired comparisons,” *Biometrika*, vol. 39, no. 3/4, pp. 324–345, 1952.
- [29] P. F. Christiano, J. Leike, T. Brown, M. Martic, S. Legg, and D. Amodei, “Deep reinforcement learning from human preferences,” *Advances in neural information processing systems*, vol. 30, 2017.
- [30] D. Brown, R. Coleman, R. Srinivasan, and S. Niekum, “Safe imitation learning via fast bayesian reward inference from preferences,” in *International Conference on Machine Learning*. PMLR, 2020, pp. 1165–1177.
- [31] K. Lee, L. Smith, and P. Abbeel, “Pebble: Feedback-efficient interactive reinforcement learning via relabeling experience and unsupervised pre-training,” *arXiv preprint arXiv:2106.05091*, 2021.
- [32] A. Bobu, Y. Liu, R. Shah, D. S. Brown, and A. D. Dragan, “Sirl: Similarity-based implicit representation learning,” in *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, 2023, pp. 565–574.
- [33] R. Rafailov, A. Sharma, E. Mitchell, C. D. Manning, S. Ermon, and C. Finn, “Direct preference optimization: Your language model is secretly a reward model,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [34] V. Myers, E. Biyik, N. Anari, and D. Sadigh, “Learning multi-modal rewards from rankings,” in *Conference on Robot Learning*. PMLR, 2022, pp. 342–352.
- [35] D. Shin, A. D. Dragan, and D. S. Brown, “Benchmarks and algorithms for offline preference-based reward learning,” *arXiv preprint arXiv:2301.01392*, 2023.
- [36] Y. Liu, G. Datta, E. Novoseller, and D. S. Brown, “Efficient preference-based reinforcement learning using learned dynamics models,” *International Conference on Robotics and Automation (ICRA)*, 2023.
- [37] N. Wilde, E. Biyik, D. Sadigh, and S. L. Smith, “Learning reward functions from scale feedback,” *arXiv preprint arXiv:2110.00284*, 2021.
- [38] A. Richards and J. How, “Mixed-integer programming for control,” in *Proceedings of the 2005, American Control Conference, 2005*. IEEE, 2005, pp. 2676–2683.
- [39] P. Belotti, P. Bonami, M. Fischetti, A. Lodi, M. Monaci, A. Nogales-Gómez, and D. Salvagnin, “On handling indicator constraints in mixed integer programming,” *Computational Optimization and Applications*, vol. 65, pp. 545–566, 2016.
- [40] M. Montazery and N. Wilson, “Dominance and optimisation based on scale-invariant maximum margin preference learning,” *International Joint Conferences on Artificial Intelligence*, 2017.
- [41] S. Teso, A. Passerini, and P. Viappiani, “Constructive preference elicitation by setwise max-margin learning,” *arXiv preprint arXiv:1604.06020*, 2016.
- [42] X. Yuan, R. Henao, E. Tsalik, R. Langley, and L. Carin, “Non-gaussian discriminative factor models via the max-margin rank-

- likelihood,” in *International Conference on Machine Learning*. PMLR, 2015, pp. 1254–1263.
- [43] R. de Lazcano, K. Andreas, J. J. Tai, S. R. Lee, and J. Terry, “Gymnasium robotics,” 2023. [Online]. Available: <http://github.com/Farama-Foundation/Gymnasium-Robotics>
- [44] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel, *et al.*, “Soft actor-critic algorithms and applications,” *arXiv preprint arXiv:1812.05905*, 2018.
- [45] E. Todorov, T. Erez, and Y. Tassa, “Mujoco: A physics engine for model-based control,” in *2012 IEEE/RSJ international conference on intelligent robots and systems*. IEEE, 2012, pp. 5026–5033.
- [46] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, “A simple framework for contrastive learning of visual representations,” in *International conference on machine learning*. PMLR, 2020, pp. 1597–1607.
- [47] C. Mattson and D. S. Brown, “Leveraging human feedback to evolve and discover novel emergent behaviors in robot swarms,” *Genetic and Evolutionary Computation Conference (GECCO)*, 2023.