
Inverse Constrained Reinforcement Learning

Shehryar Malik^{*1} Usman Anwar^{*1} Alireza Aghasi² Ali Ahmed¹

Abstract

In real world settings, numerous constraints are present which are hard to specify mathematically. However, for the real world deployment of reinforcement learning (RL), it is critical that RL agents are aware of these constraints, so that they can act safely. In this work, we consider the problem of learning constraints from demonstrations of a constraint-abiding agent’s behavior. We experimentally validate our approach and show that our framework can successfully learn the most likely constraints that the agent respects. We further show that these learned constraints are *transferable* to new agents that may have different morphologies and/or reward functions. Previous works in this regard have either mainly been restricted to tabular (discrete) settings, specific types of constraints or assume the environment’s transition dynamics. In contrast, our framework is able to learn arbitrary *Markovian* constraints in high-dimensions in a completely model-free setting. The code is available at: <https://github.com/shehryar-malik/icrl>.

1. Introduction

Reward functions are a critical component in reinforcement learning settings. As such, it is important that reward functions are designed accurately and are well-aligned with the intentions of the human designer. This is known as agent (or value) alignment (see, e.g., Leike et al. (2018; 2017); Amodei et al. (2016)). Misspecified rewards can lead to unwanted and unsafe situations (see, e.g. Amodei & Clark (2016)). However, designing accurate reward functions remains a challenging task. Human designers, for example, tend to prefer simple reward functions that agree well with their intuition and are easily interpretable. For example, a human designer might choose a reward function that en-

^{*}Equal contribution ¹Information Technology University, Lahore, Pakistan ²Georgia State University, Atlanta, GA, USA. Correspondence to: Usman Anwar <usman.anwar@itu.edu.pk>.




Figure 1. A simple visualization. The agent’s reward is proportional to how close it is to the goal at the bottom-left corner. However, the real world (right) has a lion occupying the lower bridge. The expert is aware of this and so takes the upper bridge. However, the lion is not present in the simulator (middle and left). In a bid to maximize its reward the *nominal* policy chooses to take the lower bridge. Clearly, such a policy will result in the agent being eaten up by the lion in the real world. Our method, on the other hand, tries to reconcile the reward function in the simulator with the expert’s behavior, and ultimately learns to avoid the lower bridge.

courages an RL agent driving a car to minimize its traveling time to a certain destination. Clearly, such a reward function makes sense in the case of a human driver since inter-human communication is contextualized within a framework of unwritten and unspoken constraints, often colloquially termed as ‘common-sense’. That is, while a human driver will try to minimize their traveling time, they will be careful not to break traffic rules, take actions that endanger passersby, and so on. However, we cannot assume such behaviors from RL agents since they are not imbued with common-sense constraints.

Constrained reinforcement learning provides a natural framework for maximizing a reward function subject to some constraints (we refer the reader to Ray et al. (2019) for a brief overview of the field). However, in many cases, these constraints are hard to specify explicitly in the form of mathematical functions. One way to address this issue is to automatically extract constraints by observing the behavior of a constraint-abiding agent. Consider, for example, the cartoon in Figure 1. Agents start at the bottom-left corner and are rewarded according to how quickly they reach the goal at the bottom-right corner. However, what this reward scheme misses out is that in the real world the lower bridge is occupied by a lion which attacks any agents attempting

to pass through it. Therefore, agents that are naively trained to maximize the reward function will end up performing poorly in the real world. If, on the other hand, the agent had observed that the expert (in the leftmost figure) actually performed suboptimally with respect to the stipulated reward scheme by taking a longer route to the goal, it could have concluded that (for some unknown reason) the lower bridge must be avoided and consequently would have not been eaten by the lion!

Scobee & Sastry (2020) formalizes this intuition by casting the problem of recovering constraints in the maximum entropy framework for inverse RL (IRL) (Ziebart et al., 2008) and proposes a greedy algorithm to infer the smallest number of constraints that best explain the expert behavior. However, Scobee & Sastry (2020) has two major limitations: it assumes (1) tabular (discrete) settings, and (2) the environment’s transition dynamics. In this work, we aim to address both of these issues by *learning* a constraint function instead through a sample-based approximation of the objective function of Scobee & Sastry (2020). Consequently, our approach is model-free, admits continuous states and actions and can learn arbitrary Markovian constraints¹. Furthermore, we empirically show that our method scales well to high-dimensions.

Typical inverse RL methods only make use of expert demonstrations and do not assume any knowledge about the reward function at all. However, most reward functions can be expressed in the form “do this task while not doing these other things” where *other things* are generally constraints that a designer wants to impose on an RL agent. The main task (“do this”) is often quite easy to encode in the form of a simple *nominal* reward function. In this work, we focus on learning the constraint part (“do not do that”) from provided expert demonstrations and using it in conjunction with the nominal reward function to train RL agents. In this perspective, our work can be seen as a principled way to inculcate prior knowledge about the agent’s task in IRL. This is a key advantage over other IRL methods which also often end up making assumptions about the agent’s task in the form of regularizers such as in Finn et al. (2016).

The main contribution of this work is that it formulates the problem of inferring constraints from a set of expert demonstrations as a *learning* problem which allows it to be used in continuous, high-dimensional settings. To the best of our knowledge, this is the first work in this regard. This also eliminates the need to assume, as Scobee & Sastry (2020) does, the environment’s transition dynamics. Finally, we demonstrate the ability of our method to train constraint-

¹Markovian constraints are of the form $c(\tau) = \prod_{t=1}^T c(s_t, a_t)$, i.e., the constraint function c is independent of the past states and actions in a trajectory. (The notation used here is introduced in the next section.)

abiding agents in high-dimensions and show that it can also be used to prevent reward hacking.

2. Preliminaries

2.1. Unconstrained RL

A finite-horizon Markov Decision Process (MDP) \mathcal{M} is a tuple $(\mathcal{S}, \mathcal{A}, p, r, \gamma, T)$, where $\mathcal{S} \in \mathbb{R}^{|\mathcal{S}|}$ is a set of states, $\mathcal{A} \in \mathbb{R}^{|\mathcal{A}|}$ is a set of actions, $p : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \mapsto [0, 1]$ is the transition probability function (where $p(s'|s, a)$ denotes the probability of transitioning to state s' from state s by taking action a), $r : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$ is the reward function, γ is the discount factor and T is the time-horizon. A trajectory $\tau = \{s_1, a_1, \dots, s_T, a_T\}$ denotes a sequence of states-action pairs such that $s_{t+1} \sim p(\cdot|s_t, a_t)$. A policy $\pi : \mathcal{S} \mapsto \mathcal{P}(\mathcal{A})$ is a map from states to probability distributions over actions, with $\pi(a|s)$ denoting the probability of taking action a in state s . We will sometimes abuse notation to write $\pi(s, a)$ to mean the joint probability of visiting state s and taking action a under the policy π and similarly $\pi(\tau)$ to mean the probability of the trajectory τ under the policy π .

Define $r(\tau) = \sum_{t=1}^T \gamma^t r(s_t, a_t)$ to be the total discounted reward of a trajectory. Forward RL algorithms try to find an optimal policy π^* that maximizes the expected total discounted reward $J(\pi) = \mathbb{E}_{\tau \sim \pi}[r(\tau)]$. On the other hand, given a set of trajectories sampled from the optimal (also referred to as expert) policy π^* , inverse RL (IRL) algorithms aim to recover the reward function r , which can then be used to learn the optimal policy π^* via some forward RL algorithm.

2.2. Constrained RL

While normal (unconstrained) RL tries to find a policy that maximizes $J(\pi)$, constrained RL instead focuses on finding a policy that maximizes $J(\pi)$ *while* respecting explicitly-defined constraints. A popular framework in this regard is the one presented in Altman (1999) which introduces the notion of a constrained MDP (CMDP). A CMDP \mathcal{M}^c is a simple MDP augmented with a cost function $c : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$ and a budget $\alpha \geq 0$. Define $c(\tau) = \sum_{t=1}^T \gamma^t c(s_t, a_t)$ to be the total discounted cost of the trajectory τ and $J^c(\pi) = \mathbb{E}_{\tau \sim \pi}[c(\tau)]$ to be the expected total discounted cost. The forward constrained RL problem is to find the policy π_c^* that maximizes $J(\pi)$ subject to $J^c(\pi) \leq \alpha$.

In this work, given a set \mathcal{D} of trajectories sampled from π_c^* , the corresponding unconstrained MDP \mathcal{M} (i.e., \mathcal{M}^c without the cost function c) and a budget α , we are interested in recovering a cost function which when augmented with \mathcal{M} has an optimal policy that generates the same set of trajectories as in \mathcal{D} . We call this as the inverse constrained reinforcement learning (ICRL) problem.

If the budget α is strictly greater than 0, then the cost function c defines *soft* constraints over all possible state-action pairs. In other words, a policy is allowed to visit states and take actions that have non-zero costs as long as the expected total discounted cost remains less than α . If, however, α is 0 then the cost function translates into hard constraints over all state-action pairs that have a non-zero cost associated with them. A policy can thus never visit these state-action pairs. In this work, we restrict ourselves to this hard constraint setting. Note that this is not particularly restrictive since, for example, safety constraints are often hard constraints as well as constraints imposed by physical laws.

Since we restrict ourselves to hard constraints, we can rewrite the constrained RL problems as follows: define $\mathcal{C} = \{(s, a) | c(s, a) \neq 0\}$ to be the constraint set induced by c . The forward constrained RL problem is to find the optimal constrained policy $\pi_{\mathcal{C}}^*$ that maximizes $J(\pi)$ subject to $\pi_{\mathcal{C}}^*(s, a) = 0 \forall (s, a) \in \mathcal{C}$. The inverse constrained RL problem is to recover the constraint set \mathcal{C} from trajectories sampled from $\pi_{\mathcal{C}}^*$.

Finally, we will refer to our unconstrained MDP as the nominal MDP hereinafter. The nominal MDP represents the nominal environment (simulator) in which we train our agent.

3. Formulation

3.1. Maximum Likelihood Constraint Inference

We take [Scobee & Sastry](#) as our starting point. Suppose that we have a set of trajectories $\mathcal{D} = \{\tau^{(i)}\}_{i=1}^N$ sampled from an expert $\pi_{\mathcal{C}}^*$ navigating in a constrained MDP $\mathcal{M}^{\mathcal{C}^*}$ where \mathcal{C}^* denotes the (true) constraint set. Furthermore, we are also given the corresponding nominal MDP \mathcal{M} . Our goal is to recover a constraint set which when augmented with \mathcal{M} results in a CMDP that has an optimal policy that respects the same set of constraints as $\pi_{\mathcal{C}}^*$ does. [Scobee & Sastry](#) pose this as a maximum likelihood problem. That is, if we let $p_{\mathcal{M}}$ denote probabilities given that we are considering MDP \mathcal{M} and assume a uniform prior on all constraint sets, then we can choose \mathcal{C}^* according to

$$\mathcal{C}^* \leftarrow \arg \max_{\mathcal{C}} p_{\mathcal{M}}(\mathcal{D} | \mathcal{C}). \quad (1)$$

Under the maximum entropy (MaxEnt) model presented in [Ziebart et al. \(2008\)](#), the probability of a trajectory under a deterministic MDP \mathcal{M} can be modelled as

$$\pi_{\mathcal{M}}(\tau) = \frac{\exp(\beta r(\tau)) \mathbb{1}^{\mathcal{M}}(\tau)}{Z_{\mathcal{M}}}, \quad (2)$$

where $Z_{\mathcal{M}} = \int \exp(\beta r(\tau)) \mathbb{1}^{\mathcal{M}}(\tau) d\tau$ is the partition function, $\beta \in [0, \infty)$ is a parameter describing how close the agent is to the optimal distribution (as $\beta \rightarrow \infty$ the agent

becomes a perfect optimizer and as $\beta \rightarrow 0$ the agent simply takes random actions) and $\mathbb{1}$ is an indicator function that is 1 for trajectories feasible under the MDP \mathcal{M} and 0 otherwise.

Assume that all trajectories in \mathcal{D} are i.i.d. and sampled from the MaxEnt distribution. We have

$$p(\mathcal{D} | \mathcal{C}) = \frac{1}{(Z_{\mathcal{M}^{\mathcal{C}}})^N} \prod_{i=1}^N \exp(\beta r(\tau^{(i)})) \mathbb{1}^{\mathcal{M}^{\mathcal{C}}}(\tau^{(i)}). \quad (3)$$

Note that $\mathbb{1}^{\mathcal{M}^{\mathcal{C}}}(\tau^{(i)})$ is 0 for all trajectories that contain any state-action pair that belongs to \mathcal{C} . To maximize this, [Scobee & Sastry](#) propose a greedy strategy wherein they start with an empty constraint set and incrementally add state-action pairs that result in the maximal increase in $p(\mathcal{D} | \mathcal{C})$.

3.2. Sample-Based Approximation

Since we are interested in more realistic settings where the state and action spaces can be continuous, considering all possible state-action pairs individually usually becomes intractable. Instead, we propose a learning-based approach wherein we try to approximate $\mathbb{1}^{\mathcal{M}^{\mathcal{C}}}(\tau)$ using a neural network. Consider the log likelihood

$$\begin{aligned} \mathcal{L}(\mathcal{C}) &= \frac{1}{N} \log p(\mathcal{D} | \mathcal{C}) \\ &= \frac{1}{N} \sum_{i=1}^N \left[\beta r(\tau^{(i)}) + \log \mathbb{1}^{\mathcal{M}^{\mathcal{C}}}(\tau^{(i)}) \right] - \log Z_{\mathcal{M}^{\mathcal{C}}} \\ &= \frac{1}{N} \sum_{i=1}^N \left[\beta r(\tau^{(i)}) + \log \mathbb{1}^{\mathcal{M}^{\mathcal{C}}}(\tau^{(i)}) \right] - \\ &\quad \log \int \exp(\beta r(\tau)) \mathbb{1}^{\mathcal{M}^{\mathcal{C}}}(\tau) d\tau. \end{aligned} \quad (4)$$

Note that $\mathbb{1}^{\mathcal{M}^{\mathcal{C}}}(\tau) = \prod_{t=0}^T \mathbb{1}^{\mathcal{M}^{\mathcal{C}}}(s_t, a_t)$ merely tells us whether the trajectory τ is feasible under the constraint set \mathcal{C} or not. Let us have a binary classifier ζ_{θ} parameterized with θ do this for us instead, i.e., we want $\zeta_{\theta}(s_t, a_t)$ to be 1 if (s_t, a_t) is not in \mathcal{C}^* and 0 otherwise. Using $\zeta_{\theta}(\tau)$ as a short hand for $\prod_{t=0}^T \zeta_{\theta}(s_t, a_t)$, we have

$$\begin{aligned} \mathcal{L}(\mathcal{C}) = \mathcal{L}(\theta) &= \frac{1}{N} \sum_{i=1}^N \left[\beta r(\tau^{(i)}) + \log \zeta_{\theta}(\tau^{(i)}) \right] \\ &\quad - \log \int \exp(\beta r(\tau)) \zeta_{\theta}(\tau) d\tau. \end{aligned} \quad (5)$$

Let $\mathcal{M}^{\bar{\zeta}_{\theta}}$ denote the MDP obtained after augmenting \mathcal{M} with the cost function $\bar{\zeta}_{\theta} := 1 - \zeta_{\theta}^2$, and $\pi_{\mathcal{M}^{\bar{\zeta}_{\theta}}}$ denote the

²Note that since we are assuming that α is 0, we can assign any non-zero (positive) cost to state-action pairs that we want to constrain. Here $1 - \zeta_{\theta}$ assigns a cost of 1 to all such pairs.

corresponding MaxEnt policy. Taking gradients of (5) with respect to θ gives us (see Section 8.1 in the supplementary material for derivation)

$$\nabla_{\theta} \mathcal{L}(\theta) = \frac{1}{N} \sum_{i=1}^N \nabla_{\theta} \log \zeta_{\theta}(\tau^{(i)}) - \mathbb{E}_{\tau \sim \pi_{\mathcal{M}\bar{\zeta}_{\theta}}} [\nabla_{\theta} \log \zeta_{\theta}(\tau)]. \quad (6)$$

Using a sample-based approximation for the right-hand term we can rewrite the gradient as

$$\nabla_{\theta} \mathcal{L}(\theta) \approx \frac{1}{N} \sum_{i=1}^N \nabla_{\theta} \log \zeta_{\theta}(\tau^{(i)}) - \frac{1}{M} \sum_{j=1}^M \nabla_{\theta} \log \zeta_{\theta}(\hat{\tau}^{(j)}), \quad (7)$$

where $\hat{\tau}$ are sampled from $\pi_{\mathcal{M}\bar{\zeta}_{\theta}}$ (discussed in the next section). Notice that making $\nabla_{\theta} \mathcal{L}(\theta)$ zero essentially requires matching the expected gradient of $\log \zeta_{\theta}$ under the expert (left hand term) and nominal (right hand term) trajectories. For brevity, we will write $\pi_{\mathcal{M}\bar{\zeta}_{\theta}}$ as π_{θ} from now onwards. We can choose ζ_{θ} to be a neural network with parameters θ and a sigmoid at the output. We train our neural network via gradient descent by using the expression for the gradient given above.

In practice, since we have a limited amount of data, ζ_{θ} , parameterized as a neural network, will tend to overfit. To mitigate this, we incorporate the following regularizer into our objective function.

$$R(\theta) = -\delta \sum_{\tau \in \{\mathcal{D}, \mathcal{S}\}} |1 - \zeta_{\theta}(\tau)| \quad (8)$$

where \mathcal{S} denotes the set of trajectories sampled from π_{θ} and $\delta \in [0, 1)$ is a fixed constant. R incentivizes ζ_{θ} to predict values close to 1 (recall that $\zeta_{\theta} \in (0, 1)$), thus encouraging it to choose the smallest number of constraints that best explain the expert data.

3.3. Forward Step

To evaluate (7) we need to sample from π_{θ} . Recall that π_{θ} needs to maximize $J(\pi)$ subject to $\pi_{\theta}(s, a) = 0 \forall (s, a) \in \mathcal{Z}$ where \mathcal{Z} is the constraint set induced by $\bar{\zeta}_{\theta}$. However, since ζ_{θ} outputs continuous values in the range $(0, 1)$, we instead solve the *soft* constrained RL version, wherein we try to find a policy π that maximizes $J(\pi)$ subject to $\mathbb{E}_{\tau \sim \pi} [\bar{\zeta}_{\theta}(\tau)] \leq \alpha$. In our experiments, we set α to a very small value. Note that if α is strictly set to 0 our optimization program will have an empty feasible set.

We represent our policy as a neural network with parameters ϕ and train it through the constrained policy optimization algorithm introduced in Tessler et al. (2019), which essentially tries to solve the following equivalent unconstrained min-max problem on the Lagrangian of the objective func-

tion

$$\min_{\lambda \geq 0} \max_{\phi} J(\pi^{\phi}) + \frac{1}{\beta} \mathcal{H}(\pi^{\phi}) - \lambda (\mathbb{E}_{\tau \sim \pi^{\phi}} [\bar{\zeta}_{\theta}(\tau)] - \alpha) \quad (9)$$

by gradient ascent on ϕ (via the Proximal Policy Optimization (PPO) algorithm (Schulman et al., 2017)) and gradient descent on the Lagrange multiplier λ . Note that we also add the entropy $\mathcal{H}(\pi^{\phi}) = -\mathbb{E}_{\tau \sim \pi^{\phi}} [\log \pi^{\phi}(\tau)]$ of π^{ϕ} to our objective function. Maximizing the entropy ensures that we recover the MaxEnt distribution given in (2) at convergence (see Section 8.4 in the supplementary material for proof).

3.4. Incorporating Importance Sampling

Running the forward step in each iteration is typically very time consuming. To mitigate this problem, instead of approximating the expectation in (6) with samples from π_{θ} , we approximate it with samples from an older policy $\pi_{\bar{\theta}}$ where $\bar{\theta}$ were the parameters of ζ at some previous iteration. We, therefore, only need to learn a new policy after a fixed number of iterations. To correct for the bias introduced in the approximation because of using a different sampling distribution, we add importance sampling weights to our expression for the gradient. In this case, the per step importance sampling weights can be shown to be given by (see Section 8.2 in the supplementary material for derivation)

$$\omega(s_t, a_t) = \frac{\zeta_{\theta}(s_t, a_t)}{\zeta_{\bar{\theta}}(s_t, a_t)}. \quad (10)$$

Note that the importance sampling weights at each timestep only depend on the state and action at that timestep (and not the past and/or future states and actions). The gradient can thus be rewritten as

$$\begin{aligned} \nabla_{\theta} \mathcal{L}(\theta) &\approx \frac{1}{N} \sum_{i=1}^N \sum_{s_t, a_t \in \tau^{(i)}} \nabla_{\theta} \log \zeta_{\theta}(s_t, a_t) - \\ &\frac{1}{M} \sum_{j=1}^M \sum_{\hat{s}_t, \hat{a}_t \in \tau^{(j)}} \omega(\hat{s}_t, \hat{a}_t) \nabla_{\theta} \log \zeta_{\theta}(\hat{s}_t, \hat{a}_t), \end{aligned} \quad (11)$$

where $\{\tau^{(j)}\}_{j=1}^M$ are sampled from $\pi_{\bar{\theta}}$.

3.5. KL-Based Early Stopping

Importance sampling allows us to run multiple updates of ζ_{θ} after each forward step. However, importance sampling weights may have high variance, which, in practice, may result in bad updates to ζ_{θ} especially when policies $\pi_{\bar{\theta}}$ and π_{θ} become very “different”. To mitigate this, we impose an upper bound on the forward and reverse KL divergences between $\pi_{\bar{\theta}}$ and π_{θ} . We do this by terminating updates of ζ_{θ} once the KL-divergences exceed some pre-defined threshold. The forward and reverse KL-divergences can be bounded in




Figure 2. The environments used in the experiments. Note that nominal agents are not aware of the constraints shown.

Algorithm 1 ICRL

Input: Expert trajectories \mathcal{D} , iterations N , backward iterations B , maximum allowable KLs ϵ_F and ϵ_R

Initialize θ and ϕ randomly

for $i = 1$ **to** N **do**

Learn π^ϕ by solving (9) using current ζ_θ

for $j = 1$ **to** B **do**

Sample set of trajectories $\mathcal{S} = \{\tau^{(k)}\}_{k=1}^M$ from π^ϕ

Compute importance sampling weights $w(\tau^{(k)})$ using (10) for $k = 1, \dots, M$

Use \mathcal{S} and \mathcal{D} to update θ via SGD by using the gradient in (11)

Compute forward and reverse KLs using (12)

if forward KL $\geq \epsilon_F$ **or** reverse KL $\geq \epsilon_R$: **then**

break

end if

end for

end for

terms of the importance sampling weights as follows (see Section 8.3 in the supplementary material for derivation).

$$D_{\text{KL}}(\pi_{\bar{\theta}} || \pi_\theta) \leq 2 \log \bar{\omega}$$

$$D_{\text{KL}}(\pi_\theta || \pi_{\bar{\theta}}) \leq \frac{\mathbb{E}_{\tau \sim \pi_{\bar{\theta}}} [(\omega(\tau) - \bar{\omega}) \log \omega(\tau)]}{\bar{\omega}}. \quad (12)$$

Here $\bar{\omega} = \mathbb{E}_{\tau \sim \pi_{\bar{\theta}}} [\omega(\tau)]$ and $\omega(\tau) = \prod_{t=1}^T \omega(s_t, a_t)$.

Algorithm 1 summarizes our training procedure.

4. Experiments

We conduct experiments to validate our method and assess the quality of constraints it recovers. In the following discussion, we will refer to both the nominal and constrained environments. The nominal environment corresponds to the simulator we use to train our agent (also called as the nominal agent). It does not include any information about the constraints. The constrained (or true) environment refers to the real-world setting where episodes are terminated if the agent violates any constraints. We report two metrics for

each experiment that we conduct: (1) the true reward which corresponds to the reward accumulated by the agent in the constrained environment, and (2) the average number of constraint violations per each timestep. For the latter metric, we assume that the true constraints are known. However, the training algorithm does not require this information. Note that since episodes get terminated in the constrained environment as soon as an agent violates some constraint, agents simply trained to maximize the reward in the nominal environment will only be able to accumulate a small amount of reward in the constrained environment.

Furthermore, we construct the following two baselines for comparisons:

- **Binary Classifier (BC):** We simply train a binary classifier (using the cross-entropy loss) to classify between nominal and expert trajectories. We train the agent by solving, as before, (9) (note that the binary classifier corresponds to ζ).
- **GAIL-Constraint (GC):** This baseline is based on a well-known imitation learning method called GAIL (Ho & Ermon, 2016). Since GAIL does not assume knowledge of the reward function, while our method does, we construct the following variant for a fairer comparison. We begin by defining the following reward function: $\bar{r}(s, a) := r(s, a) + \log \zeta(s, a)$. Note that since for constrained state-action pairs \bar{r} will be $-\infty$, agents maximizing \bar{r} will try to satisfy all constraints. Now since we already know r , we only parameterize ζ as a neural network which, then, corresponds to the GAIL’s discriminator.

We design our experiments to answer the following two questions:

- Can ICRL train agents to abide by the constraints that an expert does?
- Are the constraints recovered by ICRL transferable to other agents with possibly different morphologies

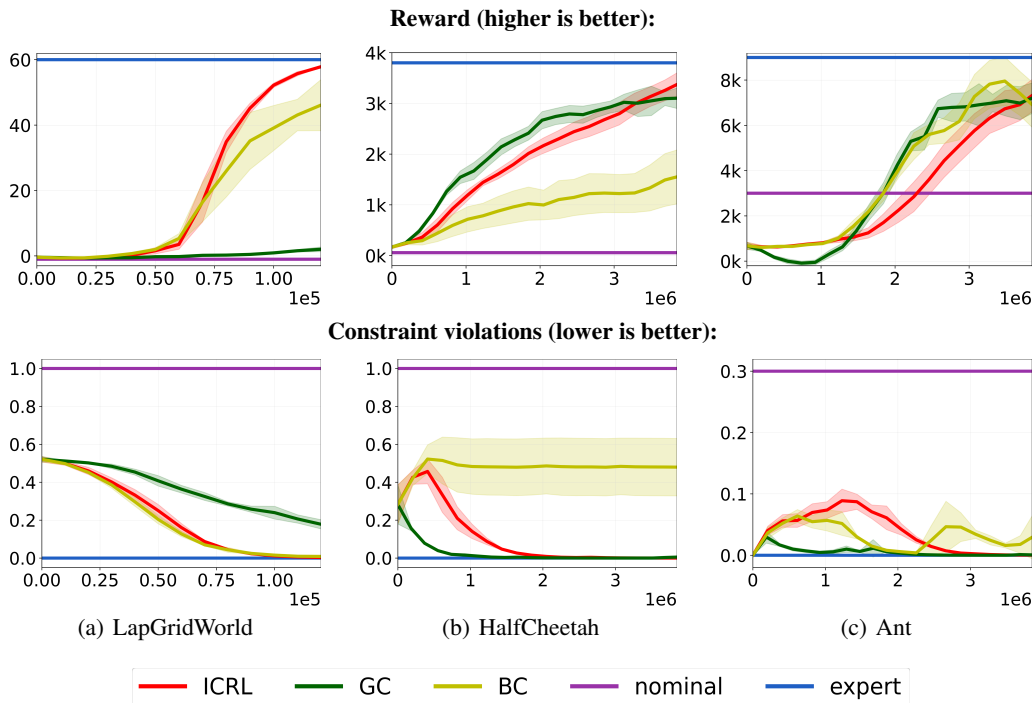


Figure 3. Performance of agents during training over several seeds (5 in LapGridWorld, 10 in others). The x-axis is the number of timesteps taken in the environment. The shaded regions correspond to the standard error.

and/or reward functions?

As we show below, the answer to both of these questions is in the affirmative. Furthermore, while the GC baseline performs equally well to ICRL in the first case, it completely fails in the second one.

Finally, all hyperparameters and additional details of the experiments can be found in Section 8.5 in the supplementary material.

4.1. Learning Constraints

In this section, we design experiments to test whether ICRL can train agents to abide by the constraints that an expert agent does.

Reward hacking: Reward hacking refers to situations in which agents optimize their rewards but their behavior is misaligned with the intentions of the users, e.g., when a vacuum cleaner ejects dust so that it can collect even more (Russell & Norvig, 2010). We design a simple environment called LapGridWorld to see if our method can avoid this behavior. In this environment³ (shown in Figure 2(a)) an agent moves on a rectangular track. The agent is *intended* to sail clockwise around the track. Each time it drives onto

³This environment is similar to the boat race environment in Leike et al. (2017).

a golden dollar tile, it gets a reward of 3. However, the nominal agent “cheats” by stepping back and forth on one dollar tile, rather than going around the track, and ends up getting more reward than the expert (which goes around the track, as intended).

High dimensions: For this experiment, we use two simulated robots called HalfCheetah and Ant from OpenAI Gym (Brockman et al., 2016). The robots are shown in Figures 2(b) and 2(c) respectively. The HalfCheetah robot has 18 state and 6 action dimensions and is rewarded proportionally to the distance it covers in each step. The Ant robot has 113 state and 8 action dimensions and is rewarded proportionally to the distance it moves away from the origin. Since, for both agents, moving backwards is relatively easier than moving forwards, in the constrained environments, we constrain all states for which $x \leq -3$.

We plot the results of our experiments in Figures 3. The x-axis corresponds to the number of timesteps the agent takes in the environment. The top row corresponds to the reward of the agent during training when evaluated in the constrained environment. The bottom row shows the average number of constraints that the agent violates per timestep (when run in the nominal environment). In addition to our method (ICRL) and the two baselines, we also show the reward and average number of constraint violations of the nominal and expert agents. As can be observed, both ICRL

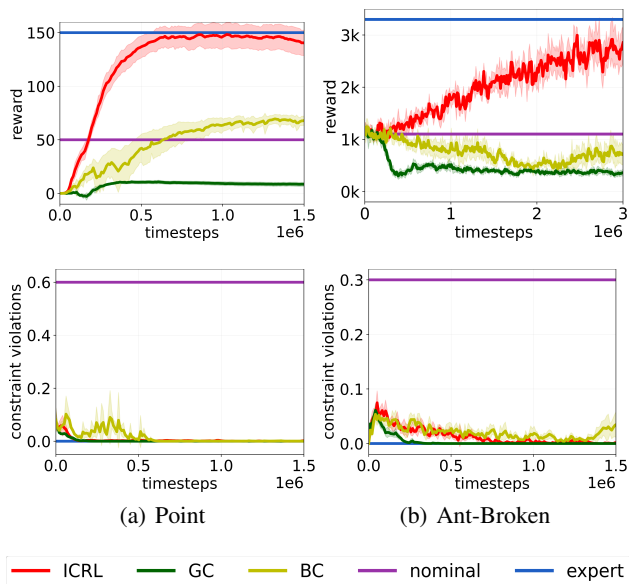


Figure 4. Transferring constraints. The x-axis is the number of timesteps taken in the environment. All plots were smoothed and averaged over 5 seeds. The shaded regions correspond to the standard error.

and GC perform equally well and are able to achieve expert-level performance. However, BC performs much more worse than these two.

4.2. Transferring constraints

In many cases, constraints are actually part of the environment and are the same for different agents (for example, **all** vehicles must adhere to the same speed limit). In such instances, it is useful to first learn the constraints using only one agent and then transfer them onto other agents. These new agents may have different morphologies and/or reward functions. In this section, we design experiments to assess the quality of constraints recovered by our method and the two baselines. To do so, we try to transfer constraints learnt in the Ant experiment in the previous section to two new agents. The first agent, shown in Figure 2(d), is the Point robot from OpenAI gym. The Point robot’s reward function encourages it to move counter-clockwise in a circle at a distance of 10 units from the origin.⁴ The second agent, called Ant-Broken, is the Ant robot with two of its legs disabled (colored in blue in Figure 2(e)).⁵ To transfer constraints, we simply solve (9) using the ζ learnt in the previous section.⁶

⁴This environment is similar to the PointCircle environment in Achiam et al. (2017).

⁵This is based on the broken ant environment in Eysenbach et al. (2020).

⁶Note that in this case, ζ must only be fed a subset of the state and action spaces that are common across the two agents. As such,

Figure 4 shows the results. The plots represent the same metrics as in the previous section. Note that only agents trained to respect constraints recovered by our method are able to achieve expert-level performance. Both BC and GC perform even worse in terms of the reward than the nominal agent, which is simply trained to maximize the reward in the nominal environment without worrying about any constraints. However, note that both BC and GC are able to achieve zero cost. This suggests that both BC and GC constrain too many state-action pairs, thus preventing agents from performing their tasks optimally.

5. Ablation Studies

In this section, we discuss ablation studies to assess the importance of importance sampling and the KL-based early stopping technique. We carry out four sets of experiments by disabling either one or both of importance sampling and early stopping. Additionally, we also plot results for the case when neither of importance sampling and early stopping are disabled. In each set of experiments, we vary the number of backward iterations, B (see Algorithm 1 and observe (as in the previous section) the true reward (on the constrained environment) and average number of constraint violations. All four sets of experiments are shown in Figure 5. For comparison, we also show the best true reward and lowest number of constraint violations that our method was able to achieve with all of its components enabled (these are the same results as in the previous section).

From the plots, it is clear that both of importance sampling and early stopping are important to the success of our method.

6. Related Work

Forward Constrained RL: Several approaches have been proposed in literature to solve the forward constrained RL problem in the context of CMDPs (Altman, 1999). Achiam et al. (2017) analytically solves trust region policy optimization problems at each policy update to enforce the constraints. Chow et al. (2018) uses a Lyapunov approach and also provide theoretical guarantees. Le et al. (2019) proposes an algorithm for cases when there are multiple constraints. Finally, a well-known approach centers around rewriting the constrained RL problem as an equivalent unconstrained min-max problem by using Lagrange multipliers (Zhang et al., 2019; Tessler et al., 2019; Bhatnagar, 2010) (see Section 3.3 for further details).

Constraint Inference: Previous work done on inferring

for the Point experiment we re-run experiments in the previous section and train ζ only on the (x, y) position coordinates of the Ant agent, since the rest of elements in the state space are specific to Ant.

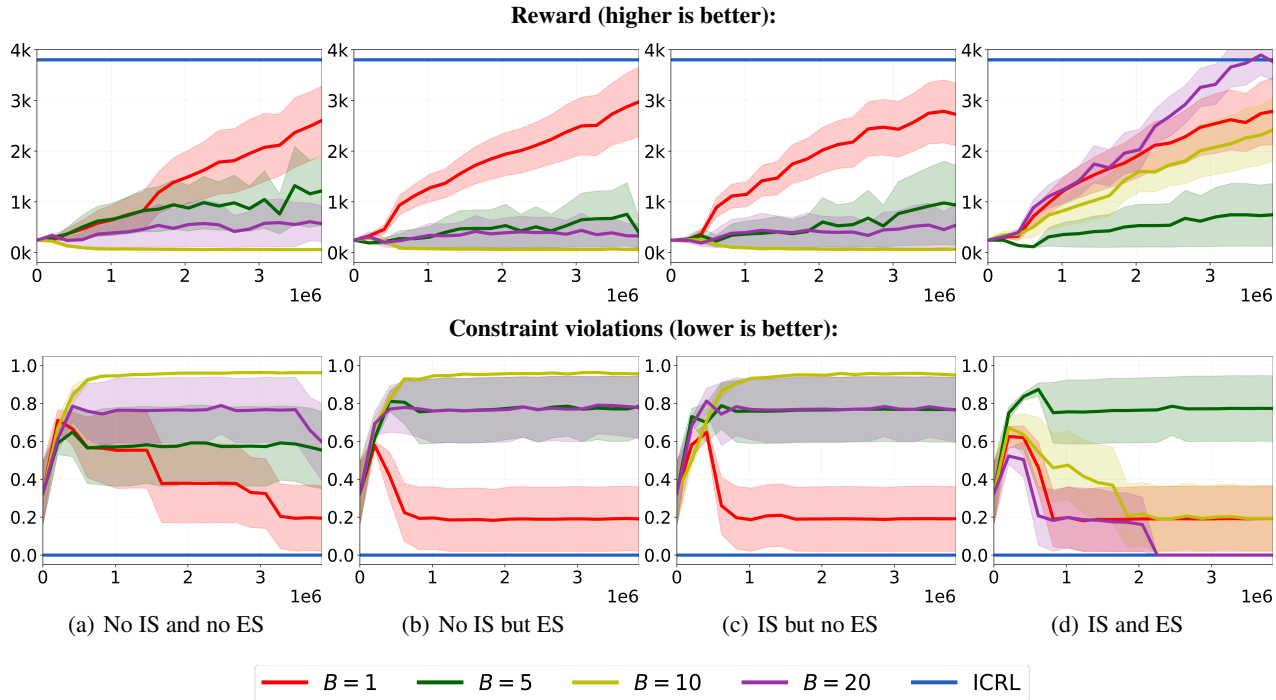


Figure 5. Ablation studies on the HalfCheetah environment. All plots were averaged over 5 seeds. IS refers to importance sampling and ES to early stopping. The x-axis corresponds to the number of timesteps the agent takes in the environment. Shaded regions correspond to the standard error.

constraints from expert demonstrations has either focused on either inferring specific type of constraints such as geometric (D’Arpino & Shah, 2017; Subramani et al., 2018), sequential (Pardowitz et al., 2005) or convex (Miryoosefi et al., 2019) constraints or is restricted to tabular settings (Scobee & Sastry, 2020; Chou et al., 2018) or assumes transition dynamics (Chou et al., 2020).

Preference Learning: Constraint inference also links to preference learning which aims to extract user preferences (constraints imposed by an expert on itself, in our case) from different sources such as ratings (Daniel et al., 2014), comparisons (Christiano et al., 2017; Sadigh et al., 2017), human reinforcement signals (MacGlashan et al., 2017) or the initial state of the agent’s environment (Shah et al., 2019). Preference learning also includes inverse RL, which aims to recover an agent’s reward function by using its trajectories. To deal with the inherent ill-posedness of this problem, inverse RL algorithms often incorporate regularizers (Ho & Ermon, 2016; Finn et al., 2016) or assume a prior distribution over the reward function (Jeon et al., 2018; Michini & How, 2012; Ramachandran & Amir, 2007).

7. Conclusion and Future Work

We have presented a method to learn constraints from an expert’s demonstrations. Unlike previous works, our method

both learns arbitrary constraints and can be used in continuous settings.

While we consider our method to be an important first step towards learning arbitrary constraints in real-world continuous settings, there is still considerable room for improvement. For example, as is the case with Scobee & Sastry, our formulation is also based on (2) which only holds for deterministic MDPs. Secondly, we only consider hard constraints. Lastly, one very interesting extension of this method can be to learn constraints *only* from logs data in an offline way to facilitate safe-RL in settings where it is difficult to even build nominal simulators such as is the case for plant controllers.

References

- Achiam, J., Held, D., Tamar, A., and Abbeel, P. Constrained policy optimization. In *International Conference on Machine Learning*, 2017.
- Altman, E. *Constrained Markov Decision Processes*. Chapman and Hall, 1999.
- Amodei, D. and Clark, J. Faulty reward functions in the wild. <https://openai.com/blog/faulty-reward-functions/>, 2016. Accessed: 2020-07-20.

- Amodei, D., Olah, C., Steinhardt, J., Christiano, P. F., Schulman, J., and Mané, D. Concrete problems in AI safety, 2016. arXiv:1606.06565.
- Bhatnagar, S. An actor-critic algorithm with function approximation for discounted cost constrained markov decision processes. *Systems and Control Letters*, 59(12): 760–766, 2010.
- Biewald, L. Experiment tracking with weights and biases, 2020. URL <https://www.wandb.com/>. Software available from wandb.com.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. OpenAI Gym, 2016. arXiv:1606.01540.
- Chou, G., Berenson, D., and Ozay, N. Learning constraints from demonstrations, 2018. arXiv:1812.07084.
- Chou, G., Ozay, N., and Berenson, D. Learning constraints from locally-optimal demonstrations under cost function uncertainty. *IEEE Robotics Autom. Lett.*, 5(2):3682–3690, 2020.
- Chow, Y., Nachum, O., Duenez-Guzman, E., and Ghavamzadeh, M. A lyapunov-based approach to safe reinforcement learning. In *Advances in Neural Information Processing Systems*, 2018.
- Christiano, P. F., Leike, J., Brown, T. B., Martic, M., Legg, S., and Amodei, D. Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems*, 2017.
- Daniel, C., Viering, M., Metz, J., Kroemer, O., and Peters, J. Active reward learning. In *Proceedings of Robotics: Science and Systems*, 2014.
- D’Arpino, C. P. and Shah, J. A. C-learn: Learning geometric constraints from demonstrations for multi-step manipulation in shared autonomy. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2017.
- Eysenbach, B., Asawa, S., Chaudhari, S., Salakhutdinov, R., and Levine, S. Off-dynamics reinforcement learning: Training for transfer with domain classifiers, 2020. arXiv:2006.13916.
- Finn, C., Levine, S., and Abbeel, P. Guided cost learning: Deep inverse optimal control via policy optimization. In *International Conference on Machine Learning*, 2016.
- Hill, A., Raffin, A., Ernestus, M., Gleave, A., Kanervisto, A., Traore, R., Dhariwal, P., Hesse, C., Klimov, O., Nichol, A., Plappert, M., Radford, A., Schulman, J., Sidor, S., and Wu, Y. Stable baselines. <https://github.com/hill-a/stable-baselines>, 2018.
- Ho, J. and Ermon, S. Generative adversarial imitation learning. In *Advances in Neural Information Processing Systems*, 2016.
- Jeon, W., Seo, S., and Kim, K.-E. A bayesian approach to generative adversarial imitation learning. In *Advances in Neural Information Processing Systems*, 2018.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *International Conference on Learning Representations, ICLR*, 2015.
- Le, H., Voloshin, C., and Yue, Y. Batch policy learning under constraints. In *International Conference on Machine Learning, ICML*, 2019.
- Leike, J., Martic, M., Krakovna, V., Ortega, P. A., Everitt, T., Lefrancq, A., Orseau, L., and Legg, S. AI safety gridworlds, 2017. arXiv:1711.09883.
- Leike, J., Krueger, D., Everitt, T., Martic, M., Maini, V., and Legg, S. Scalable agent alignment via reward modeling: A research direction, 2018. arXiv:1811.07871.
- MacGlashan, J., Ho, M. K., Loftin, R., Peng, B., Wang, G., Roberts, D. L., Taylor, M. E., and Littman, M. L. Interactive learning from policy-dependent human feedback. In *International Conference on Machine Learning*, 2017.
- Michini, B. and How, J. P. Bayesian nonparametric inverse reinforcement learning. In *Machine Learning and Knowledge Discovery in Databases*, 2012.
- Miryoosefi, S., Brantley, K., III, H. D., Dudik, M., and Schapire, R. E. Reinforcement learning with convex constraints. In *Advances in Neural Information Processing Systems*, 2019.
- Pardowitz, M., Zöllner, R., and Dillmann, R. Learning sequential constraints of tasks from user demonstrations. In *5th IEEE-RAS International Conference on Humanoid Robots*, 2005.
- Ramachandran, D. and Amir, E. Bayesian inverse reinforcement learning. In *20th International Joint Conference on Artificial Intelligence*, 2007.
- Ray, A., Achiam, J., and Amodei, D. Benchmarking safe exploration in deep reinforcement learning, 2019. <https://cdn.openai.com/safexp-short.pdf>.
- Russell, S. and Norvig, P. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 3 edition, 2010.
- Sadigh, D., Dragan, A. D., Sastry, S., and Seshia, S. A. Active preference-based learning of reward functions. In *Robotics: Science and Systems XIII*, 2017.

- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms, 2017. arxiv:1707.06347.
- Scobee, D. R. and Sastry, S. S. Maximum likelihood constraint inference for inverse reinforcement learning. In *International Conference on Learning Representations*, 2020.
- Shah, R., Krasheninnikov, D., Alexander, J., Abbeel, P., and Dragan, A. D. Preferences implicit in the state of the world. In *International Conference on Learning Representations*, 2019.
- Subramani, G., Zinn, M., and Gleicher, M. Inferring geometric constraints in human demonstrations. In *2nd Conference on Robot Learning (CoRL)*, 2018.
- Tessler, C., Mankowitz, D. J., and Mannor, S. Reward constrained policy optimization. In *International Conference on Learning Representations*, 2019.
- Zhang, R., Yu, T., Shen, Y., Jin, H., and Chen, C. Text-based interactive recommendation via constraint-augmented reinforcement learning. In *Advances in Neural Information Processing Systems*, 2019.
- Ziebart, B. D., Maas, A., Bagnell, J. A., and Dey, A. K. Maximum entropy inverse reinforcement learning. In *23rd National Conference on Artificial Intelligence (AAAI)*, 2008.