

# Interactive Teaching Algorithms for Inverse Reinforcement Learning

Parameswaran Kamalaruban<sup>1\*</sup>, Rati Devidze<sup>2\*</sup>, Volkan Cevher<sup>1</sup> and Adish Singla<sup>2</sup>

<sup>1</sup>LIONS, EPFL

<sup>2</sup>Max Planck Institute for Software Systems (MPI-SWS)

{kamalaruban,parameswaran,volkan.cevher}@epfl.ch, {rdevidze,adishs}@mpi-sws.org

## Abstract

We study the problem of inverse reinforcement learning (IRL) with the added twist that the learner is assisted by a helpful teacher. More formally, we tackle the following algorithmic question: How could a teacher provide an informative sequence of demonstrations to an IRL learner to speed up the learning process? We present an interactive teaching framework where a teacher adaptively chooses the next demonstration based on learner’s current policy. In particular, we design teaching algorithms for two concrete settings: an omniscient setting where a teacher has full knowledge about the learner’s dynamics and a blackbox setting where the teacher has minimal knowledge. Then, we study a sequential variant of the popular MCE-IRL learner and prove convergence guarantees of our teaching algorithm in the omniscient setting. Extensive experiments with a car driving simulator environment show that the learning progress can be speeded up drastically as compared to an uninformative teacher.

## 1 Introduction

Imitation Learning, also known as Learning from Demonstrations, enables a learner to acquire new skills by observing a teacher’s behavior. It plays an important role in many real-life learning settings, including human-to-human interaction [Buchsbbaum *et al.*, 2011; Shafto *et al.*, 2014], and human-to-robot interaction [Schaal, 1997; Billard *et al.*, 2008; Argall *et al.*, 2009; Chernova and Thomaz, 2014].

Inverse reinforcement learning (IRL) is one of the popular approaches to imitation learning: IRL algorithms operate by first inferring an intermediate reward function explaining the demonstrated behavior, and then obtaining a policy corresponding to the inferred reward [Russell, 1998; Abbeel and Ng, 2004]. IRL has been extensively studied in the context of designing efficient learning algorithms for a given set of demonstrations [Abbeel and Ng, 2004; Ratliff *et al.*, 2006; Ziebart *et al.*, 2008; Boularias *et al.*, 2011; Wulfmeier *et al.*, 2015; Finn *et al.*, 2016]. There has also been

some recent work on designing active/interactive IRL algorithms that focus on reducing the number of demonstrations that needs to be requested from a teacher [Amin *et al.*, 2017; Dorsa Sadigh *et al.*, 2017]. Despite these advances, the problem of generating an optimal sequence of demonstrations to teach an IRL agent is still not well understood.

Motivated by applications of intelligent tutoring systems to teach sequential decision-making tasks, such as surgical training<sup>1</sup> or car driving<sup>2</sup>, we study the IRL framework from the viewpoint of a “teacher” in order to best assist an IRL “learner”. Cakmak and Lopes; Brown and Niekum (2012; 2019) have studied the problem of teaching an IRL learner in a batch setting, i.e., the teacher has to provide a near-optimal set of demonstrations at once. Their teaching algorithms are *non-interactive*, i.e., they construct their teaching sequences without incorporating any feedback from the learner and hence unable to adapt the teaching to the learner’s progress.

In real-life pedagogical settings, it is evident that a teacher can leverage the learner’s feedback in adaptively choosing next demonstrations/tasks to accelerate the learning progress. For instance, consider a scenario where a driving instructor wants to teach a student certain driving skills. The instructor can easily identify the mistakes/weaknesses of the student (e.g., unable to do rear parking), and then carefully choose tasks that this student should perform, along with specific demonstrations to rectify any mistakes. In this paper, we study the problem of designing an *interactive* teaching algorithm for an IRL learner.

### 1.1 Overview of Our Approach

We consider an interactive teaching framework where at any given time: (i) the teacher observes the learner’s current policy, (ii) then, the teacher provides the next teaching task/demonstration, and (iii) the learner performs an update. We design interactive teaching algorithms for two settings:

- an “omniscient” teaching setting where the teacher has full knowledge of the learner’s dynamics and can fully observe the learner’s current policy.
- a “blackbox” teaching setting where the teacher doesn’t know learner’s dynamics and has only a noisy estimate of the learner’s current policy.

<sup>1</sup>Simulators for surgical training: <https://www.virtamed.com/en/>

<sup>2</sup>Simulators for car driving: <https://www.driverinteractive.com/>

\* Authors contributed equally to this work.

In the omniscient teaching setting, we study a sequential variant of the popular IRL algorithm, namely Maximum Causal Entropy (MCE) IRL algorithm [Ziebart *et al.*, 2008; Rhinehart and Kitani, 2017]). Our main idea in designing our omniscient teaching algorithm, OMNITEACHER, is to first reduce the problem of teaching a target policy to that of teaching a corresponding hyperparameter (see Section 3.2); then, the teacher greedily steers the learner towards this hyperparameter. We then prove convergence guarantees of the OMNITEACHER algorithm and show that it can significantly reduce the number of demonstrations required to achieve a desired performance of the learner (see Theorem 1, Section 4).

While omniscient teacher yields strong theoretical guarantees, its applicability is limited given that the learner’s dynamics are unknown and difficult to infer in practical applications. Based on insights from the omniscient teacher, we develop a simple greedy teaching algorithm, BBOXTEACHER, for a more practical blackbox setting (see Section 5).

We perform extensive experiments in a synthetic learning environment (with both linear and non-linear reward settings) inspired by a car driving simulator [Ng and Russell, 2000; Levine *et al.*, 2010]. We demonstrate that our teaching algorithms can bring significant improvements in speeding up the learning progress compared to an uninformative teaching strategy that picks demonstrations at random. Furthermore, the performance of the BBOXTEACHER algorithm is close to the OMNITEACHER algorithm even though it operates with limited information about the learner.

## 2 Problem Setup

We now formalize the problem addressed in this paper.

### 2.1 Environment

The environment is formally represented by an MDP  $\mathcal{M} := (\mathcal{S}, \mathcal{A}, T, \gamma, P_0, R^E)$ . The sets of possible states and actions are denoted by  $\mathcal{S}$  and  $\mathcal{A}$  respectively.  $T : \mathcal{S} \times \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$  captures the state transition dynamics, i.e.,  $T(s' | s, a)$  denotes the probability of landing in state  $s'$  by taking action  $a$  from state  $s$ . Here  $\gamma$  is the discounting factor, and  $P_0 : \mathcal{S} \rightarrow [0, 1]$  is an initial distribution over states  $\mathcal{S}$ . We denote a policy  $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$  as a mapping from a state to a probability distribution over actions. The underlying reward function is given by  $R^E : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ .

### 2.2 Interaction between Learner and Teacher

In our setting, we have two entities: a teacher and a sequential IRL learner. The teacher has access to the full MDP  $\mathcal{M}$  and has a *target policy*  $\pi^E$  computed as an optimal policy w.r.t.  $R^E$ . The learner knows the MDP  $\mathcal{M}$  but not the reward function  $R^E$ , i.e., has only access to  $\mathcal{M} \setminus R^E$ . The teacher’s goal is to provide an informative sequence of demonstrations to teach the policy  $\pi^E$  to the learner. Here, a teacher’s demonstration  $\xi = \{(s_\tau, a_\tau)\}_{\tau=0,1,\dots}$  is obtained by first choosing an initial state  $s_0 \in \mathcal{S}$  (where  $P_0(s_0) > 0$ ) and then choosing a trajectory of state-action pairs obtained by executing the policy  $\pi^E$  in the MDP  $\mathcal{M}$ .

We consider an interactive teaching framework with three key steps formally described in Algorithm 1. At any time  $t$ ,

---

### Algorithm 1 Interactive Teaching Framework

---

- 1: **for**  $t = 1, 2, \dots, T$  **do**
  - 2:   teacher observes an estimate of learner’s policy  $\pi_t^L$
  - 3:   teacher provides a demonstration  $\xi_t$  to the learner
  - 4:   learner updates the policy to  $\pi_{t+1}^L$  using  $\xi_t$
- 

the teacher observes an estimate of learner’s current policy  $\pi_t^L$  and accordingly provides an informative demonstration  $\xi_t$  to assist the learner. Then, the learner updates its policy to  $\pi_{t+1}^L$  based on demonstration  $\xi_t$  and its internal learning dynamics.

### 2.3 Occupancy Measure and Expected Reward

We introduce the following two notions to formally define the teaching objective. For any policy  $\pi$ , the occupancy measure  $\rho$  and the total expected reward  $\nu$  of  $\pi$  in the MDP  $\mathcal{M}$  are defined as follows respectively:

$$\rho^\pi(s, a) := (1 - \gamma) \pi(a | s) \sum_{\tau=0}^{\infty} \gamma^\tau \mathbb{P}\{S_\tau = s | \pi, \mathcal{M}\} \quad (1)$$

$$\nu^\pi := \frac{1}{1 - \gamma} \sum_{s,a} \rho^\pi(s, a) R^E(s, a) \quad (2)$$

Here,  $\mathbb{P}\{S_\tau = s | \pi, \mathcal{M}\}$  denotes the probability of visiting the state  $s$  after  $\tau$  steps by following the policy  $\pi$ . Similarly, for any demonstration  $\xi = \{(s_\tau, a_\tau)\}_{\tau=0,1,\dots}$ , we define

$$\rho^\xi(s, a) := (1 - \gamma) \sum_{\tau=0}^{\infty} \gamma^\tau \mathbb{I}\{s_\tau = s, a_\tau = a\}$$

Then for a collection of demonstrations  $\Xi = \{\xi_t\}_{t=1,2,\dots}$ , we have  $\rho^{\Xi}(s, a) := \frac{1}{|\Xi|} \sum_t \rho^{\xi_t}(s, a)$ .

### 2.4 Teaching Objective

Let  $\pi^L$  denote the learner’s final policy at the end of teaching. The performance of the policy  $\pi^L$  (w.r.t.  $\pi^E$ ) in  $\mathcal{M}$  can be evaluated via the following measures (for some fixed  $\epsilon > 0$ ):

1.  $|\nu^{\pi^E} - \nu^{\pi^L}| \leq \epsilon$ , ensuring high reward [Abbeel and Ng, 2004; Ziebart, 2010].
2.  $D_{\text{TV}}(\rho^{\pi^E}, \rho^{\pi^L}) \leq \epsilon$ , ensuring that learner’s behavior induced by the policy  $\pi^L$  matches that of the teacher [Ho and Ermon, 2016]. Here  $D_{\text{TV}}(p, q)$  is the total variation distance between two probability measures  $p$  and  $q$ .

IRL learner’s goal is to  $\epsilon$ -approximate the teacher’s policy w.r.t. one of these performance measures [Ziebart, 2010; Ho and Ermon, 2016]. In this paper, we study this problem from the viewpoint of a teacher in order to provide a near-optimal sequence of demonstrations  $\{\xi_t\}_{t=1,2,\dots}$  to the learner to achieve the desired goal. The teacher’s performance is then measured by the number of demonstrations required to achieve the above objective.

## 3 Omniscient Teaching Setting

In this section, we consider the omniscient teaching setting where the teacher knows the learner’s dynamics including the learner’s current parameters at any given time. We begin by introducing a specific learner model that we study for this setting.

---

**Algorithm 2** Sequential MCE-IRL
 

---

- 1: **Initialization:** parameter  $\lambda_1$ , policy  $\pi_1^L := \pi_{\lambda_1}^L$
  - 2: **for**  $t = 1, 2, \dots, T$  **do**
  - 3:   receives a demonstration  $\xi_t$  with starting state  $s_{t,0}$
  - 4:   update  $\lambda_{t+1} \leftarrow \Pi_{\Omega} \left[ \lambda_t - \eta_t (\mu^{\pi_t^L, s_{t,0}} - \mu^{\xi_t}) \right]$
  - 5:   compute  $\pi_{t+1}^L \leftarrow \text{Soft-Value-Iter}(\mathcal{M} \setminus R^E, R_{\lambda_{t+1}}^L)$
  - 6: **Output:** policy  $\pi^L \leftarrow \pi_{T+1}^L$
- 

### 3.1 Learner Model

We consider a learner implementing an IRL algorithm based on Maximum Causal Entropy approach (MCE-IRL) [Ziebart *et al.*, 2008; Ziebart, 2010; Wulfmeier *et al.*, 2015; Rhinehart and Kitani, 2017; Zhou *et al.*, 2018]. First, we discuss the parametric reward function of the learner and then introduce a sequential variant of the MCE-IRL algorithm used in our interactive teaching framework.

**Parametric reward function.** We consider the learner model with parametric reward function  $R_{\lambda}^L: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$  where  $\lambda \in \mathbb{R}^d$  is a parameter. The reward function also depends on the learner’s feature representation  $\phi^L: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$ . For linear reward function  $R_{\lambda}^L(s, a) := \langle \lambda, \phi^L(s, a) \rangle$ ,  $\lambda$  represents the weights. As an example of non-linear rewards, the function  $R_{\lambda}^L$  could be high-order polynomial in  $\lambda$  (see Section 6.2). As a more powerful non-linear reward model,  $\lambda$  could be the weights of a neural network with  $\phi^L(s, a)$  as input layer and  $R_{\lambda}^L(s, a)$  as output [Wulfmeier *et al.*, 2015].

**Soft Bellman policies.** Given a fixed parameter  $\lambda$ , we model the learner’s behavior via the following soft Bellman policy<sup>3</sup>:

$$\begin{aligned} \pi_{\lambda}^L(a | s) &= \exp(Q_{\lambda}(s, a) - V_{\lambda}(s)) & (3) \\ V_{\lambda}(s) &= \log \sum_a \exp(Q_{\lambda}(s, a)) \\ Q_{\lambda}(s, a) &= R_{\lambda}^L(s, a) + \gamma \sum_{s'} T(s' | s, a) V_{\lambda}(s'). \end{aligned}$$

For any given  $\lambda$ , the corresponding policy  $\pi_{\lambda}^L$  can be efficiently computed via Soft-Value-Iteration procedure (see [Ziebart, 2010, Algorithm. 9.1], [Zhou *et al.*, 2018]).

**Gradient update for sequential MCE-IRL.** We consider a sequential MCE-IRL learner for our interactive setting where, at time step  $t$ , the learner receives next demonstration  $\xi_t = \{(s_{t,\tau}, a_{t,\tau})\}_{\tau=0,1,\dots}$  with starting state  $s_{t,0}$ . Given the learner’s current parameter  $\lambda_t$  and policy  $\pi_t^L := \pi_{\lambda_t}^L$ , the learner updates its parameter via an online gradient descent update rule given by  $\lambda_{t+1} = \lambda_t - \eta_t g_t$  with  $\eta_t$  as the learning rate. The gradient  $g_t$  at time  $t$  is computed as

$$g_t = \sum_{s,a} \left[ \rho^{\pi_t^L, s_{t,0}}(s, a) - \rho^{\xi_t}(s, a) \right] \nabla_{\lambda} R_{\lambda}^L(s, a) \Big|_{\lambda=\lambda_t},$$

---

<sup>3</sup>For the case of linear reward functions, soft Bellman policies are obtained as the solution to the MCE-IRL optimization problem [Ziebart, 2010; Zhou *et al.*, 2018]. Here, we extend this idea to model the learner’s policy with the general parametric reward function. However, we note that for our learner model, it is not always possible to formulate a corresponding optimization problem that leads to the policy form mentioned above.

---

**Algorithm 3** OMNITEACHER for sequential MCE-IRL
 

---

- 1: **for**  $t = 1, 2, \dots, T$  **do**
  - 2:   teacher picks  $\xi_t$  by solving Eq. (6)
  - 3:   learner receives  $\xi_t$  and updates using Algorithm 2
- 

where  $\rho^{\pi_t^L, s_{t,0}}(s, a)$  is given by Eq. (1) and computed with  $s_{t,0}$  as the only initial state, i.e.,  $P_0(s_{t,0}) = 1$ . As shown in the longer version of the paper [Kamalaruban *et al.*, 2019],  $g_t$  is the gradient of the loss function

$$c(\lambda; \xi_t) = - \sum_{\tau} \gamma^{\tau} \log \pi_{\lambda}^L(a_{t,\tau} | s_{t,\tau})$$

capturing the discounted negative log-likelihood of teacher’s demonstration  $\xi_t$ .

**Sequential MCE-IRL algorithm.** For brevity, we define the following quantities:

$$\mu^{\pi_t^L, s_{t,0}} = \sum_{s,a} \rho^{\pi_t^L, s_{t,0}}(s, a) \nabla_{\lambda} R_{\lambda}^L(s, a) \Big|_{\lambda=\lambda_t} \quad (4)$$

$$\mu^{\xi_t} = \sum_{s,a} \rho^{\xi_t}(s, a) \nabla_{\lambda} R_{\lambda}^L(s, a) \Big|_{\lambda=\lambda_t}, \quad (5)$$

and we write the gradient compactly as  $g_t = \mu^{\pi_t^L, s_{t,0}} - \mu^{\xi_t}$ . Algorithm 2 presents the proposed sequential MCE-IRL learner. In particular, we use the online projected gradient descent update given by  $\lambda_{t+1} \leftarrow \Pi_{\Omega} [\lambda_t - \eta_t g_t]$ , where  $\Omega = \{\lambda : \|\lambda\|_2 \leq z\}$  for large enough  $z$  (cf. Section 4.3).

### 3.2 Omniscient Teacher

Next, we present our omniscient teaching algorithm, OMNITEACHER, for sequential MCE-IRL learner.

**Policy to hyperparameter teaching.** The main idea in designing our algorithm, OMNITEACHER, is to first reduce the problem of teaching a target policy  $\pi^E$  to that of teaching a corresponding hyperparameter, denoted below as  $\lambda^*$ . Then, under certain technical conditions, teaching this  $\lambda^*$  to the learner ensures that the learner’s policy  $\epsilon$ -approximate the target policy  $\pi^E$ . We defer the technical details to Section 4.

**Omniscient teaching algorithm.** Now, we design a teaching strategy to greedily steer the learner’s parameter  $\lambda_t$  towards the target hyperparameter  $\lambda^*$ . The main idea is to pick a demonstration which minimizes the distance between the learner’s current parameter  $\lambda_t$  and  $\lambda^*$ , i.e., minimize  $\|\lambda_{t+1} - \lambda^*\| = \|\lambda_t - \eta_t g_t - \lambda^*\|$ . We note that, Liu *et al.*; Liu *et al.* (2017; 2018) have used this idea in their iterative machine teaching framework for regression and classification tasks. In particular, we consider the following optimization problem for selecting an informative demonstration at time  $t$ :

$$\min_{s, \xi} \eta_t^2 \left\| \mu^{\pi_t^L, s} - \mu^{\xi} \right\|^2 - 2\eta_t \left\langle \lambda_t - \lambda^*, \mu^{\pi_t^L, s} - \mu^{\xi} \right\rangle, \quad (6)$$

where  $s \in \mathcal{S}$  is an initial state (where  $P_0(s_0) > 0$ ) and  $\xi$  is a trajectory obtained by executing the policy  $\pi^E$  starting from  $s$ .<sup>4</sup> The resulting teaching strategy is given in Algorithm 3.

---

<sup>4</sup>Note that  $\xi$  is not constructed synthetically by the teacher, but is obtained by executing policy  $\pi^E$  starting from  $s$ . However,  $\xi$  is not just a random rollout of the policy  $\pi^E$ : When there are multiple possible trajectories from  $s$  using  $\pi^E$ , the teacher can choose the most desirable one as per the joint optimization problem in Eq. (6).

## 4 Omniscient Teaching Setting: Analysis

In this section, we analyze the teaching complexity of our algorithm OMNITEACHER; the proofs are provided in the longer version of the paper [Kamalaruban *et al.*, 2019].

### 4.1 Convergence to Hyperparameter $\lambda^*$

In this section, we analyze the teaching complexity, i.e., the total number of time steps  $T$  required, to steer the learner towards the target hyperparameter. For this analysis, we quantify the ‘‘richness’’ of demonstrations (i.e.,  $(s_{t,0}, \xi_t)$ ) as solution to Eq. (6) in terms of providing desired gradients to the learner.

For the state  $s_{t,0}$  picked by the teacher at time step  $t$ , let us first define the following objective function given by

$f_t(\mu) := \left\| \lambda_t - \lambda^* - \eta_t \mu^{\pi^L, s_{t,0}} + \eta_t \mu \right\|^2$ . Note that the optimal solution  $\mu_t^{\text{syn}} := \arg \min_{\mu \in \mathbb{R}^d} f_t(\mu)$  takes the closed form as  $\mu_t^{\text{syn}} = \mu^{\pi^L, s_{t,0}} - \frac{1}{\eta_t} (\lambda_t - \lambda^*)$ .

Ideally, the teacher would like to provide a demonstration  $\xi_t$  (starting from  $s_{t,0}$ ) for which  $\mu^{\xi_t} = \mu_t^{\text{syn}}$ . In this case, the learner’s hyperparameter converges to  $\lambda^*$  after this time step. However, in our setting, the teacher is restricted to only provide demonstrations obtained by executing the teacher’s policy  $\pi^E$  in the MDP  $\mathcal{M}$ . We say that a teacher is  $\Delta_{\max}$ -rich, if for every  $t$ , the teacher can provide a demonstration  $\xi_t$  that satisfies the following:

$$\mu^{\xi_t} = \mu^{\pi^L, s_{t,0}} - \beta_t (\lambda_t - \lambda^*) + \delta_t,$$

for some  $\delta_t$  s.t.  $\|\delta_t\|_2 \leq \Delta_{\max}$ , and  $\beta_t \in [0, \frac{1}{\eta_t}]$ . The main intuition behind these two quantities is the following: (i)  $\beta_t$  bounds the magnitude of the gradient in the desired direction of  $(\lambda_t - \lambda^*)$  and (ii)  $\delta_t$  accounts for the deviation of the gradient from the desired direction of  $(\lambda_t - \lambda^*)$ .

The following lemma provides a bound on the teaching complexity of OMNITEACHER to steer the learner’s parameter to the target  $\lambda^*$ .

**Lemma 1.** *Given  $\epsilon' > 0$ , let the OMNITEACHER be  $\Delta_{\max}$ -rich with  $\Delta_{\max} = \frac{\epsilon'^2 \beta^2}{4\eta_{\max}[4(1-\beta)z+1]}$ , where  $\beta = \min_t \eta_t \beta_t$ , and  $\eta_{\max} = \max_t \eta_t$ . Then for the OMNITEACHER algorithm with  $T = \mathcal{O}(\log \frac{1}{\epsilon'})$ , we have  $\|\lambda_T - \lambda^*\| \leq \epsilon'$ .*

Note that the above lemma only guarantees the convergence to the target hyperparameter  $\lambda^*$ . To provide guarantees on our teaching objective, we need additional technical conditions which we discuss below.

### 4.2 Convergence to Policy $\pi_{\lambda^*}^L$

We require a smoothness condition on the learner’s reward function to ensure that the learner’s output policy  $\pi^L$  is close to the policy  $\pi_{\lambda^*}^L$  in terms of the total expected reward  $\nu^\pi$  (see Eq. (2)). Formally  $R_\lambda^L: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$  is  $m$ -smooth if the following holds:

$$\max_{s,a} |R_\lambda^L(s,a) - R_{\lambda'}^L(s,a)| \leq m \|\lambda - \lambda'\|. \quad (7)$$

Note that in the linear reward case, the smoothness parameter is  $m = \sqrt{d}$ . Given that the reward function  $R_\lambda^L$  is smooth, the following lemma illustrates the inherent smoothness of the soft Bellman policies given in Eq. (3).

**Lemma 2.** *Consider an MDP  $\mathcal{M}$ , and a learner model with parametric reward function  $R_\lambda^L$  which is  $m$ -smooth. Then for any  $\lambda, \lambda' \in \mathbb{R}^d$ , the associated soft Bellman policies  $\pi_\lambda^L$  and  $\pi_{\lambda'}^L$  satisfy the following:*

$$\left| \nu^{\pi_\lambda^L} - \nu^{\pi_{\lambda'}^L} \right| \leq R_{\max}^E \cdot \sqrt{\frac{8m}{(1-\gamma)^5}} \cdot \|\lambda - \lambda'\|,$$

where  $R_{\max}^E := \max_{s,a} |R^E(s,a)|$ .

The above lemma suggests that convergence to the target hyperparameter  $\lambda^*$  also guarantees convergence to the policy  $\pi_{\lambda^*}^L$  associated with the target hyperparameter in terms of total expected reward.

### 4.3 Convergence to Policy $\pi^E$

Finally, we need to ensure that the learner’s model  $(R_\lambda^L, \phi^L)$  is powerful enough to capture the teacher’s reward  $R^E$  and policy  $\pi^E$ . Intuitively, this would imply that there exists a target hyperparameter  $\lambda^*$  for which the soft Bellman policy  $\pi_{\lambda^*}^L$  has a total expected reward close to that of the teacher’s policy  $\pi^E$  w.r.t. the reward function  $R^E$ . Formally, we say that a learner is  $\ell$ -learnable (where  $\ell > 0$ ) if there exists a  $\lambda \in \Omega$ , such that the following holds:

$$\left| \nu^{\pi_\lambda^L} - \nu^{\pi^E} \right| \leq \ell.$$

In particular, for achieving the desired teaching objective (see Theorem 1), we require that the learner is  $\frac{\epsilon}{2}$ -learnable. This in turn implies that there exists a  $\lambda^* \in \Omega$  such that  $\left| \nu^{\pi_{\lambda^*}^L} - \nu^{\pi^E} \right| \leq \frac{\epsilon}{2}$ .<sup>5</sup> Then, combining Lemma 1 and Lemma 2, we obtain our main result:

**Theorem 1.** *Given  $\epsilon > 0$ . Define  $R_{\max}^E := \max_{s,a} |R^E(s,a)|$ . Let  $R_\lambda^L$  be  $m$ -smooth. Set  $\epsilon' = \frac{(1-\gamma)^5 \epsilon^2}{32m(R_{\max}^E)^2}$ . In addition, we have the following:*

- Teacher is  $\Delta_{\max}$ -rich with  $\Delta_{\max} = \frac{\epsilon'^2 \beta^2}{4\eta_{\max}[4(1-\beta)z+1]}$ , where  $\beta = \min_t \eta_t \beta_t$ , and  $\eta_{\max} = \max_t \eta_t$ .
- Learner is  $\frac{\epsilon}{2}$ -learnerable.
- Teacher has access to  $\lambda^*$  such that  $\left| \nu^{\pi_{\lambda^*}^L} - \nu^{\pi^E} \right| \leq \frac{\epsilon}{2}$ .

Then, for the OMNITEACHER algorithm with  $T = \mathcal{O}(\log \frac{1}{\epsilon})$ , we have  $\left| \nu^{\pi^E} - \nu^{\pi^L} \right| \leq \epsilon$ .

The above result states that the number of demonstrations required to achieve the desired objective is only  $\mathcal{O}(\log \frac{1}{\epsilon})$ . In practice, this can lead to a drastic speed up in teaching compared to an ‘‘agnostic’’ teacher (referred to as AGNOSTIC) that provides demonstrations at random, i.e., by randomly choosing the initial state  $s_0 \sim P_0$  and then picking a random rollout of the policy  $\pi^E$  in the MDP  $\mathcal{M}$ . In fact, for teaching regression and classification tasks, Liu *et al.* (2017) showed that an omniscient teacher achieves an exponential improvement in teaching complexity as compared to AGNOSTIC teacher who picks the examples at random.

<sup>5</sup>For the case of linear rewards, such a  $\lambda^*$  is guaranteed to exist and it can be computed efficiently; further details are provided in the longer version of the paper [Kamalaruban *et al.*, 2019].

## 5 Blackbox Teaching Setting

In this section, we study a more practical setting where the teacher (i) cannot directly observe the learner’s policy  $\pi_t^L$  at any time  $t$  and (ii) does not know the learner’s dynamics. The resulting teaching strategy, namely BBOXTEACHER, is given in Algorithm 4; details are provided below.

**Limited observability.** We first address the challenge of limited observability. The main idea is that in real-world applications, the teacher could approximately infer the learner’s policy  $\pi_t^L$  by probing the learner, for instance, by asking the learner to perform certain tasks or “tests”. Here, the notion of a test is to pick an initial state  $s \in \mathcal{S}$  (where  $P_0(s) > 0$ ) and then asking the learner to execute the current policy  $\pi_t^L$  from  $s$ . Formally, we characterize this probing via two parameters  $(B, k)$ : After an interval of  $B$  time steps of teaching, the teacher asks learner to perform  $k$  “tests” for every initial state. Then, based on observed demonstrations of the learner’s policy, the teacher can approximately estimate the occupancy measure  $\widehat{\rho^{\pi_t^L, s}} \approx \rho^{\pi_t^L, s}, \forall s \text{ s.t. } P_0(s) > 0$ .

**Unknown learner’s dynamics.** To additionally deal with the second challenge of unknown learner’s dynamics, we propose a simple greedy strategy of picking an informative demonstration. In particular, we pick the demonstration at time  $t$  by solving the following optimization problem (see the corresponding equation Eq. (6) for the omniscient teacher):

$$\max_{s, \xi} \left| \sum_{s', a'} \left\{ \widehat{\rho^{\pi_t^L, s}}(s', a') - \rho^\xi(s', a') \right\} R^E(s', a') \right|, \quad (8)$$

where  $s \in \mathcal{S}$  is an initial state (where  $P_0(s_0) > 0$ ) and  $\xi$  is a trajectory obtained by executing the policy  $\pi^E$  starting from  $s$ . Note that we have used the estimate  $\widehat{\rho^{\pi_t^L, s}}$  instead of  $\rho^{\pi_t^L, s}$ . This strategy is inspired from insights of the omniscient teaching setting and can be seen as picking a demonstration  $(s, \xi)$  with maximal discrepancy between the learner’s current policy and the teacher’s policy in terms of expected reward.

## 6 Experimental Evaluation

In this section, we demonstrate the performance of our algorithms in a synthetic learning environment (with both linear and non-linear reward settings) inspired by a car driving simulator [Ng and Russell, 2000; Levine *et al.*, 2010].

### Environment Setup

Fig. 1 illustrates a car driving simulator environment consisting of 9 different lane types (henceforth referred to as tasks), denoted as T0, T1, . . . , T8. Each of these tasks is associated with different driving skills. For instance, task T0 corresponds to a basic setup representing a traffic-free highway—it has a very small probability of the presence of another car. However, task T1 represents a crowded highway with 0.25 probability of encountering a car. Task T2 has stones on the right lane, whereas task T3 has a mix of both cars and stones. Similarly, tasks T4 has grass on the right lane, and T5 has a mix of both grass and cars. Tasks T6, T7, and T8 introduce more complex features such as pedestrians, HOV, and police.

### Algorithm 4 BBOXTEACHER for a sequential IRL learner

- 1: **Initialization:** probing parameters ( $B$  frequency,  $k$  tests)
- 2: **for**  $t = 1, 2, \dots, T$  **do**
- 3:     **if**  $t \% B = 1$ : teacher estimates  $\pi_t^L$  using  $k$  tests
- 4:     teacher picks  $\xi_t$  by solving Eq. (8)
- 5:     learner receives  $\xi_t$  and updates using it’s algorithm

The agent’s goal is to navigate starting from an initial state at the bottom left to the top of the lane. The agent can take three different actions given by  $\mathcal{A} = \{\text{left}, \text{straight}, \text{right}\}$ . Action `left` steers the agent to the left of the current lane. If agent is already in the leftmost lane when taking action `left`, then the lane is randomly chosen with probability 0.5. We define similar dynamics for taking action `right`; action `straight` means no change in the lane. Irrespective of the action taken, the agent always moves forward. W.l.o.g., we consider that only the agent moves in the environment.

### 6.1 Linear Reward Setting

First, we study a linear reward setting, and use the notation of `lin` in subscript to denote the MDP  $\mathcal{M}_{\text{lin}}$ , the teacher’s reward  $R^{E_{\text{lin}}}$ , and the teacher’s policy  $\pi^{E_{\text{lin}}}$ .

#### MDP and Teacher’s Policy

We consider lanes corresponding to the first 8 tasks in the environment, namely T0, T1, . . . , T7. We have  $n$  number of lanes of a given task, each generated randomly according to the tasks’ characteristics described above. Hence, the total number of states  $|\mathcal{S}|$  in our MDP  $\mathcal{M}_{\text{lin}}$  is  $n \times 8 \times 20$  where each cell represents a state, and each lane is associated with 20 states (see Fig. 1). There is one initial state for each lane corresponding to the bottom left cell of that lane.

Next, we define the reward function  $R^{E_{\text{lin}}}$  (i.e., the teacher’s reward function). We consider a state-dependent reward that depends on the underlying features of a state  $s$  given by the vector  $\phi(s)$  as follows:

- features indicating the type of the current grid cell as `stone`, `grass`, `car`, `ped`, `HOV`, and `police`.
- features providing some look-ahead information such as whether there is a car or pedestrian in the immediate front cell (denoted as `car-in-f` and `ped-in-f`).

Given this, we define the teacher’s reward function of linear form as  $R^{E_{\text{lin}}}(s) = \langle w, \phi(s) \rangle$ , where the  $w$  values are given in Fig. 2. Teacher’s policy  $\pi^{E_{\text{lin}}}$  is then computed as the optimal policy w.r.t. this reward function and is illustrated via arrows in Fig. 1 (T0 to T7) representing the path taken by the teacher when driving in this environment.

#### Learner Model

We consider the learner model with linear reward function that depends only on state, i.e.,  $R_\lambda^L(s) = \langle \lambda, \phi^L(s) \rangle$ , where  $\phi^L(s) = \phi(s)$  as given in Fig. 2. The learner is implementing the sequential MCE-IRL in Algorithm 2, where the learner’s prior knowledge is captured by the initial policy  $\pi_1^L$  (corresponding to hyperparameter  $\lambda_1$ ).

In the experiments, we consider the following prior knowledge of the learner:  $\pi_1^L$  is initially trained based on demonstrations of  $\pi^{E_{\text{lin}}}$  sampled only from the lanes associated with the

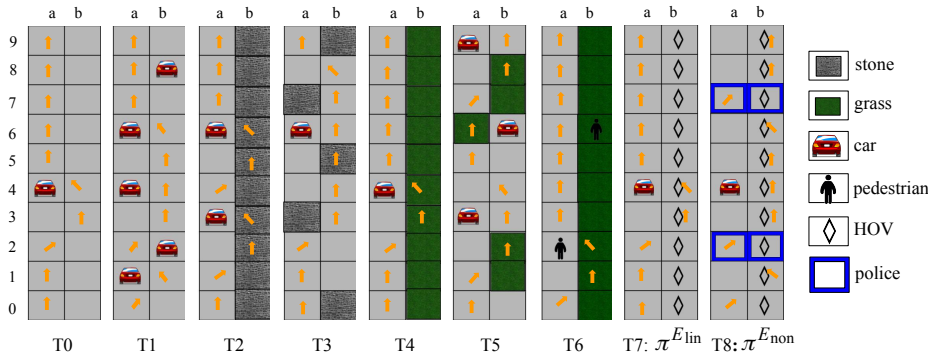


Figure 1: Car environment with 9 different lane types (tasks). In any given lane, an agent starts from the bottom-left corner and the goal is to reach the top of the lane. Arrows represent the path taken by the teacher’s policy.

| $\phi(s)$ | $w$  |
|-----------|------|
| stone     | -1   |
| grass     | -0.5 |
| car       | -5   |
| ped       | -10  |
| HOV       | -1   |
| police    | 0    |
| car-in-f  | -2   |
| ped-in-f  | -5   |

Figure 2:  $\phi(s)$  represents 8 features for a state  $s$ . Weight vector  $w$  is used to define the teacher’s reward function  $R^{E_{lin}}$ .

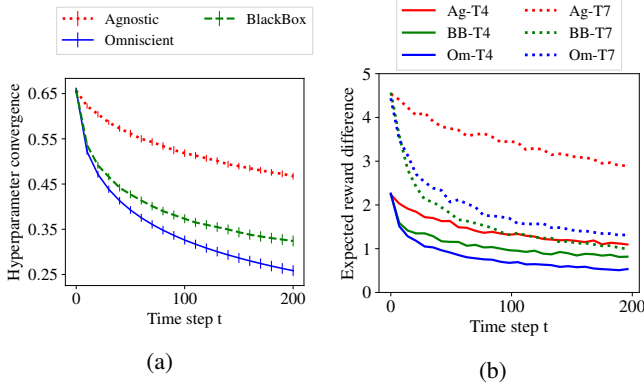


Figure 3: Linear setting of Section 6.1. (a) Convergence of the learner’s  $\lambda_t$  to the target  $\lambda^*$ . (b) Difference of total expected reward  $\nu$  of learner’s policy w.r.t. teacher’s policy in different tasks.

tasks T0, T1, T2, and T3. Intuitively, the learner initially possesses skills to avoid collisions with cars and to avoid hitting stones while driving. We expect to teach three major skills to the learner, i.e., avoiding grass while driving (task T4, T5, and T6), maintaining distance to pedestrians (task T6), and not to drive on HOV (task T7).

## Experimental Results

We evaluate the performance of different teaching algorithms, and report the results by averaging over 10 runs. We use  $n = 5$  lanes of each task (i.e., 40 lanes in total). OMNITEACHER in Algorithm 3 computes the target hyperparameter  $\lambda^*$  as per Footnote 5. For BBOXTEACHER in Algorithm 4, we use  $B = 5$  and  $k = 5$ .

Fig. 3a shows the convergence of  $\|\lambda_t - \lambda^*\|_2$ , where  $\lambda^*$  is computed by OMNITEACHER. As expected, OMNITEACHER outperforms other teaching algorithms (BBOXTEACHER and AGNOSTIC teacher) that do not have knowledge of the learner’s dynamics and are not directly focusing on teaching this hyperparameter  $\lambda^*$ . Interestingly, the convergence of BBOXTEACHER is still significantly faster than AGNOSTIC teacher. In Fig. 3b, we consider the teaching objective of total expected reward difference  $|\nu^{\pi^{E_{lin}}} - \nu^{\pi_t^L}|$ . We separately plot this objective with starting states limited to task T4 and task

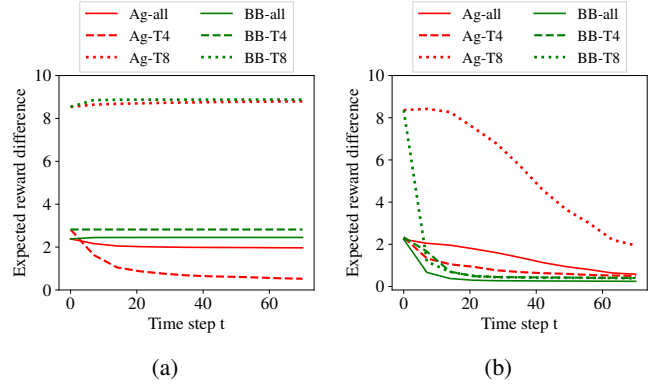


Figure 4: Non-linear setting of Section 6.2. (a) Results for a learner model with linear function  $R_\lambda^L$  unable to represent the teacher’s reward. (b) Results for a learner model with non-linear function  $R_\lambda^L$ .

T7. The results suggest that convergence to target  $\lambda^*$  leads to reduction in expected reward difference (one of the teaching objective). The performance of BBOXTEACHER (BB-T7) is even better than OMNITEACHER (Om-T7) for task T7. This is because BBOXTEACHER (Eq. (8)) picks tasks on which the learner’s total expected reward difference is highest.

In Fig. 5, we compare the teaching curriculum of BBOXTEACHER for three different settings, where the learner’s initial policy  $\pi_1^L$  is trained based on demonstrations of  $\pi^{E_{lin}}$  sampled only from the lanes associated with the tasks (a) T0, (b) T0–T3, and (c) T0–T5. The curriculum here refers to the task associated with the state  $s_{t,0}$  picked by the teacher at time  $t$  to provide the demonstration. In these plots, we can see specific structures and temporal patterns emerging in the curriculum. In particular, we can see that the teaching curriculum focuses on tasks that help the learner acquire new skills. For instance, in Fig. 5b, the teacher primarily picks tasks that provide new skills corresponding to the features *grass*, *HOV*, and *ped*. Recall that, for BBOXTEACHER (Algorithm 4), we use  $B = 5$  and  $k = 5$  for our experiments. As a result, the curriculum plots show blocks of length  $B = 5$  as the algorithm ends up picking the same task until new tests are performed to get a fresh estimate of the learner’s policy.




Figure 5: Teaching curriculum (i.e., the task associated with the picked state  $s_{t,0}$  by BBOXTEACHER in Algorithm 4) for three different settings depending on the learner’s initial skills trained on (a) T0, (b) T0–T3, and (c) T0–T5.

## 6.2 Non-linear Reward Setting

Here, we study a non-linear reward setting, and use the notation of non in subscript for  $\mathcal{M}_{\text{non}}$ ,  $R^{E_{\text{non}}}$ , and  $\pi^{E_{\text{non}}}$ .

### MDP and Teacher’s Policy

We consider the MDP  $\mathcal{M}_{\text{non}}$  which consists of lanes corresponding to 5 tasks: T0, T1, T2, T4, and T8 with the total number of states being  $n \times 5 \times 20$  as described above.

Here, we define the teacher’s reward function  $R^{E_{\text{non}}}$ . The key difference between  $R^{E_{\text{non}}}$  and  $R^{E_{\text{lin}}}$  is that the teacher in this setting prefers to drive on HOV (gets a +1 reward instead of  $-1$  reward). However, if police is present while driving on HOV, there is a penalty of  $-5$ . Teacher’s optimal policy  $\pi^{E_{\text{non}}}$  for T8 is given in Fig. 1. For other tasks (apart from T7 and T8), the teacher’s optimal policy  $\pi^{E_{\text{non}}}$  is the same as  $\pi^{E_{\text{lin}}}$ .

### Learner Model

We consider two different learner models: (i) with linear reward function  $R_{\lambda}^L(s) = \langle \lambda, \phi^L(s) \rangle$ , and (ii) with non-linear reward function  $R_{\lambda}^L(s) = \langle \lambda_{1:d'}, \phi^L(s) \rangle + \langle \lambda_{d'+1:2d'}, \phi^L(s) \rangle^2$ . Here,  $\phi^L(s) = \phi(s)$  as in Section 6.1 and  $d' = 8$  is the dimension of  $\phi(s)$ . As a prior knowledge to get  $\pi_1^L$  at time  $t = 1$ , we train the learner initially based on demonstrations of  $\pi^{E_{\text{non}}}$  sampled only from the lanes associated with the tasks T0, T1, and T2.

### Experimental Results

We use similar experimental settings as in Section 6.1 (i.e.,  $n = 5$ , averaging 10 runs, etc.). We separately report results for teaching a learner with linear  $R_{\lambda}^L$  (see Fig. 4a) and non-linear  $R_{\lambda}^L$  (see Fig. 4b).

Fig. 4a shows that both BBOXTEACHER and AGNOSTIC teacher are unable to make progress in teaching task T8 to the learner. Interestingly, the overall performance measuring the total expected reward difference on the whole MDP for BBOXTEACHER (BB-all) is worse compared to AGNOSTIC (Ag-all): This is an artifact of BBOXTEACHER’s strategy in Eq. (8) being stuck in always picking task T8 (even though the learner is not making any progress). Fig. 4b illustrates that the above mentioned limitations are gone when teaching a learner using a non-linear reward function. Here, the rate of reduction in total expected reward difference is significantly faster for BBOXTEACHER as compared to AGNOSTIC, as was also observed in Fig. 3b.

These results demonstrate that the learning progress can be sped up significantly by an adaptive teacher even with limited knowledge about the learner, as compared to an uninformative teacher. These results also signify the importance that the learner’s representation of feature space and reward function should be powerful enough to learn the desired behavior.

## 7 Related Work

**Imitation learning.** The two popular approaches for imitation learning include (i) behavioral cloning, which directly replicates the desired behavior [Bain and Sommut, 1999], and (ii) inverse reinforcement learning (IRL) which infers the reward function explaining the desired behavior [Russell, 1998]. We refer the reader to a recent survey by Osa *et al.* (2018) on imitation learning. Our approach is inspired by real-life pedagogical settings where carefully choosing the teaching demonstrations and tasks can accelerate the learning progress [Ho *et al.*, 2016]. Hadfield-Menell *et al.* (2016) have studied the value alignment problem in a game-theoretic setup, and provided an approximate scheme to generate instructive demonstrations for an IRL agent. In our work, we devise a systematic procedure (with convergence guarantees for the omniscient setting) to choose an optimal sequence of demonstrations, by taking into account the learner’s dynamics.

**Steering and teaching in reinforcement learning.** A somewhat different but related problem setting is that of *reward shaping* and *environment design* where the goal is to design the reward function to steer the behavior of a learning agent [Ng *et al.*, 1999; Zhang *et al.*, 2009; Sorg *et al.*, 2010]. Another related problem setting is considered in the *advice-based interaction* framework (e.g., [Torrey and Taylor, 2013; Amir *et al.*, 2016]), where the goal is to communicate advice to a suboptimal agent on how to act in the world.

**Algorithmic teaching.** Another line of research relevant to our work is that of algorithmic teaching. Here, one studies the interaction between a teacher and a learner where the teacher’s objective is to find an optimal training sequence to steer the learner towards a desired goal [Goldman and Kearns, 1995; Liu *et al.*, 2017; Zhu *et al.*, 2018]. Algorithmic teaching provides a rigorous formalism for a number of real-world applications such as personalized education and intelligent tutoring systems [Patil *et al.*, 2014; Rafferty *et al.*, 2016;

Hunziker *et al.*, 2018], social robotics [Cakmak and Thomaz, 2014], and human-in-the-loop systems [Singla *et al.*, 2014; Singla *et al.*, 2013]. Most of the work in machine teaching is in a batch setting where the teacher provides a batch of teaching examples at once without any adaptation. The question of how a teacher should adaptively select teaching examples for a learner has been addressed recently but only in the supervised learning setting [Melo *et al.*, 2018; Liu *et al.*, 2018; Chen *et al.*, 2018; Yeo *et al.*, 2019].

**Teaching sequential tasks.** Walsh and Goschin; Cakmak and Lopes (2012; 2012) have studied algorithmic teaching for sequential decision-making tasks. Cakmak and Lopes (2012) studied the problem of teaching an IRL agent in the batch setting. They considered the IRL algorithm by Ng and Russell (2000), which could only result in inferring an equivalent class of reward weight parameters for which the observed behavior is optimal. Brown and Niekum (2019) have extended the work of Cakmak and Lopes (2012) by showing that the teaching problem can be formulated as a set cover problem. However, their teaching strategy does not take into account how the learner progresses (i.e., it is non-interactive). In contrast, we study interactive teaching setting to teach a sequential MCE-IRL algorithm [Ziebart *et al.*, 2008; Rhinehart and Kitani, 2017]. This interactive setting, in turn, allows the teacher to design a personalized and adaptive curriculum important for efficient learning [Tadepalli, 2008]. Haug *et al.* (2018) have studied the problem of teaching an IRL agent adaptively; however, they consider a very different setting where the teacher and the learner have a mismatch in their worldviews.

## 8 Conclusions

We studied the problem of designing interactive teaching algorithms to provide an informative sequence of demonstrations to a sequential IRL learner. In an omniscient teaching setting, we presented OMNITEACHER which achieves the teaching objective with  $\mathcal{O}(\log \frac{1}{\epsilon})$  demonstrations. Then, utilizing the insights from OMNITEACHER, we proposed BBOXTEACHER for a more practical blackbox setting. We showed the effectiveness of our algorithms via extensive experiments in a learning environment inspired by a car driving simulator.

As future work, we will investigate extensions of our ideas to more complex environments; we hope that, ultimately, such extensions will provide a basis for designing teaching strategies for intelligent tutoring systems (see Footnote 1 and Footnote 2). It would also be interesting to benchmark active imitation methods for MCE-IRL learner using our omniscient teacher (see, e.g., [Brown and Niekum, 2019]). Our results are also important in getting a better understanding of the robustness of MCE-IRL learner against adversarial training-set poisoning attacks. Our fast convergence results in Theorem 1 suggests that the MCE-IRL learner is actually brittle to adversarial attacks, and designing a robust MCE-IRL is an important direction of future work.

## Acknowledgments

This work was supported in part by the Swiss National Science Foundation (SNSF) under grant number 407540\_167319.

## References

- [Abbeel and Ng, 2004] Pieter Abbeel and Andrew Y Ng. Apprenticeship learning via inverse reinforcement learning. In *ICML*, 2004.
- [Amin *et al.*, 2017] Kareem Amin, Nan Jiang, and Satinder P. Singh. Repeated inverse reinforcement learning. In *NIPS*, pages 1813–1822, 2017.
- [Amir *et al.*, 2016] Ofra Amir, Ece Kamar, Andrey Kolobov, and Barbara J. Grosz. Interactive teaching strategies for agent training. In *IJCAI*, pages 804–811, 2016.
- [Argall *et al.*, 2009] Brenna D Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot learning from demonstration. *Robotics and autonomous systems*, 2009.
- [Bain and Sommut, 1999] Michael Bain and Claude Sommut. A framework for behavioural cloning. *Machine Intelligence 15*, 15:103, 1999.
- [Billard *et al.*, 2008] Aude Billard, Sylvain Calinon, Ruediger Dillmann, and Stefan Schaal. Robot programming by demonstration. In *Springer handbook of robotics*. 2008.
- [Boularias *et al.*, 2011] Abdeslam Boularias, Jens Kober, and Jan Peters. Relative entropy inverse reinforcement learning. In *AISTATS*, pages 182–189, 2011.
- [Brown and Niekum, 2019] Daniel S. Brown and Scott Niekum. Machine teaching for inverse reinforcement learning: Algorithms and applications. In *AAAI*, 2019.
- [Buchsbaum *et al.*, 2011] Daphna Buchsbaum, Alison Gopnik, Thomas L Griffiths, and Patrick Shafto. Children’s imitation of causal action sequences is influenced by statistical and pedagogical evidence. *Cognition*, 2011.
- [Cakmak and Lopes, 2012] Maya Cakmak and Manuel Lopes. Algorithmic and human teaching of sequential decision tasks. In *AAAI*, 2012.
- [Cakmak and Thomaz, 2014] Maya Cakmak and Andrea L Thomaz. Eliciting good teaching from humans for machine learners. *Artificial Intelligence*, 217:198–215, 2014.
- [Chen *et al.*, 2018] Yuxin Chen, Adish Singla, Oisín Mac Aodha, Pietro Perona, and Yisong Yue. Understanding the role of adaptivity in machine teaching: The case of version space learners. In *NeurIPS*, 2018.
- [Chernova and Thomaz, 2014] Sonia Chernova and Andrea L Thomaz. Robot learning from human teachers. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 2014.
- [Dorsa Sadigh *et al.*, 2017] Anca D Dragan Dorsa Sadigh, Shankar Sastry, and Sanjit A Seshia. Active preference-based learning of reward functions. In *RSS*, 2017.
- [Finn *et al.*, 2016] Chelsea Finn, Sergey Levine, and Pieter Abbeel. Guided cost learning: Deep inverse optimal control via policy optimization. In *ICML*, pages 49–58, 2016.
- [Goldman and Kearns, 1995] Sally A Goldman and Michael J Kearns. On the complexity of teaching. *Journal of Computer and System Sciences*, 50(1):20–31, 1995.



- [Hadfield-Menell *et al.*, 2016] Dylan Hadfield-Menell, Stuart J Russell, Pieter Abbeel, and Anca Dragan. Cooperative inverse reinforcement learning. In *NIPS*, 2016.
- [Haug *et al.*, 2018] Luis Haug, Sebastian Tschiatschek, and Adish Singla. Teaching Inverse Reinforcement Learners via Features and Demonstrations. In *NeurIPS*, 2018.
- [Ho and Ermon, 2016] Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. In *NIPS*, 2016.
- [Ho *et al.*, 2016] Mark K Ho, Michael Littman, James MacGlashan, Fiery Cushman, and Joseph L Austerweil. Showing versus doing: Teaching by demonstration. In *NIPS*, 2016.
- [Hunziker *et al.*, 2018] Anette Hunziker, Yuxin Chen, Oisín Mac Aodha, Manuel Gomez-Rodriguez, Andreas Krause, Pietro Perona, Yisong Yue, and Adish Singla. Teaching multiple concepts to a forgetful learner. *CoRR*, abs/1805.08322, 2018.
- [Kamalaruban *et al.*, 2019] Parameswaran Kamalaruban, Rati Devidze, Volkan Cevher, and Adish Singla. Interactive teaching algorithms for inverse reinforcement learning. *CoRR*, abs/1905.11867, 2019.
- [Levine *et al.*, 2010] Sergey Levine, Zoran Popovic, and Vladlen Koltun. Feature construction for inverse reinforcement learning. In *NIPS*, pages 1342–1350, 2010.
- [Liu *et al.*, 2017] Weiyang Liu, Bo Dai, Ahmad Humayun, Charlene Tay, Chen Yu, Linda B. Smith, James M. Rehg, and Le Song. Iterative machine teaching. In *ICML*, 2017.
- [Liu *et al.*, 2018] Weiyang Liu, Bo Dai, Xingguo Li, Zhen Liu, James Rehg, and Le Song. Towards black-box iterative machine teaching. In *ICML*, 2018.
- [Melo *et al.*, 2018] Francisco S Melo, Carla Guerra, and Manuel Lopes. Interactive optimal teaching with unknown learners. In *IJCAI*, pages 2567–2573, 2018.
- [Ng and Russell, 2000] Andrew Y Ng and Stuart J Russell. Algorithms for inverse reinforcement learning. In *ICML*, 2000.
- [Ng *et al.*, 1999] Andrew Y Ng, Daishi Harada, and Stuart Russell. Policy invariance under reward transformations: Theory and application to reward shaping. In *ICML*, 1999.
- [Osa *et al.*, 2018] Takayuki Osa, Joni Pajarinen, Gerhard Neumann, J Andrew Bagnell, Pieter Abbeel, Jan Peters, et al. An algorithmic perspective on imitation learning. *Foundations and Trends® in Robotics*, 7(1-2):1–179, 2018.
- [Patil *et al.*, 2014] Kaustubh R Patil, Xiaojin Zhu, Łukasz Kopeć, and Bradley C Love. Optimal teaching for limited-capacity human learners. In *NIPS*, pages 2465–2473, 2014.
- [Rafferty *et al.*, 2016] Anna N Rafferty, Emma Brunskill, Thomas L Griffiths, and Patrick Shafto. Faster teaching via pomdp planning. *Cognitive science*, 2016.
- [Ratliff *et al.*, 2006] Nathan D Ratliff, J Andrew Bagnell, and Martin A Zinkevich. Maximum margin planning. In *ICML*, pages 729–736, 2006.
- [Rhinehart and Kitani, 2017] Nicholas Rhinehart and Kris M Kitani. First-person activity forecasting with online inverse reinforcement learning. In *ICCV*, pages 3696–3705, 2017.
- [Russell, 1998] Stuart Russell. Learning agents for uncertain environments. In *COLT*, pages 101–103, 1998.
- [Schaal, 1997] Stefan Schaal. Learning from demonstration. In *NIPS*, pages 1040–1046, 1997.
- [Shafto *et al.*, 2014] Patrick Shafto, Noah D Goodman, and Thomas L Griffiths. A rational account of pedagogical reasoning: Teaching by, and learning from, examples. *Cognitive psychology*, 71:55–89, 2014.
- [Singla *et al.*, 2013] Adish Singla, Ilija Bogunovic, G Bartók, A Karbasi, and A Krause. On actively teaching the crowd to classify. In *NIPS Workshop on Data Driven Education*, 2013.
- [Singla *et al.*, 2014] Adish Singla, Ilija Bogunovic, Gábor Bartók, Amin Karbasi, and Andreas Krause. Near-optimally teaching the crowd to classify. In *ICML*, 2014.
- [Sorg *et al.*, 2010] Jonathan Sorg, Satinder P. Singh, and Richard L. Lewis. Reward design via online gradient ascent. In *NIPS*, pages 2190–2198, 2010.
- [Tadepalli, 2008] Prasad Tadepalli. Learning to solve problems from exercises. *Computational Intelligence*, 2008.
- [Torrey and Taylor, 2013] Lisa Torrey and Matthew Taylor. Teaching on a budget: Agents advising agents in reinforcement learning. In *AAMAS*, pages 1053–1060, 2013.
- [Walsh and Goschin, 2012] Thomas J. Walsh and Sergiu Goschin. Dynamic teaching in sequential decision making environments. In *UAI*, 2012.
- [Wulfmeier *et al.*, 2015] Markus Wulfmeier, Peter Ondruska, and Ingmar Posner. Maximum entropy deep inverse reinforcement learning. *CoRR*, abs/1507.04888, 2015.
- [Yeo *et al.*, 2019] Teresa Yeo, Parameswaran Kamalaruban, Adish Singla, Arpit Merchant, Thibault Asselborn, Louis Faucon, Pierre Dillenbourg, and Volkan Cevher. Iterative classroom teaching. In *AAAI*, 2019.
- [Zhang *et al.*, 2009] Haoqi Zhang, David C Parkes, and Yiling Chen. Policy teaching through reward function learning. In *EC*, pages 295–304, 2009.
- [Zhou *et al.*, 2018] Zhengyuan Zhou, Michael Bloem, and Nicholas Bambos. Infinite time horizon maximum causal entropy inverse reinforcement learning. *IEEE Trans. Automat. Contr.*, 63(9):2787–2802, 2018.
- [Zhu *et al.*, 2018] Xiaojin Zhu, Adish Singla, Sandra Zilles, and Anna N. Rafferty. An overview of machine teaching. *CoRR*, abs/1801.05927, 2018.
- [Ziebart *et al.*, 2008] Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, and Anind K Dey. Maximum entropy inverse reinforcement learning. In *AAAI*, 2008.
- [Ziebart, 2010] Brian D Ziebart. *Modeling purposeful adaptive behavior with the principle of maximum causal entropy*. Carnegie Mellon University, 2010.