Curriculum Design for Teaching via Demonstrations: Theory and Applications

Gaurav Yengera[1,2] gyengera@mpi-sws.org
Rati Devidze[1] rdevidze@mpi-sws.org
Parameswaran Kamalaruban[3] kparameswaran@turing.ac.uk
Adish Singla[1] adishs@mpi-sws.org

[1]Max Planck Institute for Software Systems (MPI-SWS), Saarbrucken, Germany
[2]Saarland University, Saarland Informatics Campus (SIC), Saarbrucken, Germany
[3]The Alan Turing Institute, London, UK

Abstract

We consider the problem of teaching via demonstrations in sequential decision-making settings. In particular, we study how to design a personalized curriculum over demonstrations to speed up the learner's convergence. We provide a unified curriculum strategy for two popular learner models: Maximum Causal Entropy Inverse Reinforcement Learning (MaxEnt-IRL) and Cross-Entropy Behavioral Cloning (CrossEnt-BC). Our unified strategy induces a ranking over demonstrations based on a notion of difficulty scores computed w.r.t. the teacher's optimal policy and the learner's current policy. Compared to the state of the art, our strategy doesn't require access to the learner's internal dynamics and still enjoys similar convergence guarantees under mild technical conditions. Furthermore, we adapt our curriculum strategy to the setting where no teacher agent is present using task-specific difficulty scores. Experiments on a synthetic car driving environment and navigation-based environments demonstrate the effectiveness of our curriculum strategy.

1 Introduction

Imitation learning is a paradigm in which a learner acquires a new set of skills by imitating a teacher's behavior. The importance of imitation learning is realized in real-world applications where the desired behavior cannot be explicitly defined but can be demonstrated easily. These applications include the settings involving both human-to-machine interaction [1–4], and human-to-human interaction [5, 6]. The two most popular approaches to imitation learning are Behavioral Cloning (BC) [7] and Inverse Reinforcement Learning (IRL) [8]. BC algorithms aim to directly match the behavior of the teacher using supervised learning methods. IRL algorithms operate in a two-step approach: first, a reward function explaining the teacher's behavior is inferred; then, the learner adopts a policy corresponding to the inferred reward.

In the literature, imitation learning has been extensively studied from the learner's point of view to design efficient learning algorithms [9–15]. However, much less work is done from the teacher's point of view to reduce the number of demonstrations required to achieve the learning objective. In this paper, we focus on the problem of Teaching via Demonstrations (TvD), where a helpful teacher assists the imitation learner in converging quickly by designing a personalized curriculum [16–20]. Despite a substantial amount of work on curriculum design for reinforcement learning agents [21–27], curriculum design for imitation learning agents is much less investigated.

Prior work on curriculum design for IRL learners has focused on two concrete settings: non-interactive and interactive. In the non-interactive setting [17, 18], the teacher provides a near-optimal set of demonstrations as a single batch. These curriculum strategies do not incorporate any feedback

35th Conference on Neural Information Processing Systems (NeurIPS 2021), Sydney, Australia.

from the learner, hence unable to adapt the teaching to the learner's progress. In the interactive setting [28], the teacher can leverage the learner's progress to adaptively choose the next demonstrations to accelerate the learning process. We focus on designing an interactive curriculum algorithm with theoretical guarantees that is agnostic to learner dynamics, i.e., the gradient update rule of the learner. This makes our algorithm applicable to a broad range of learner models and in practical settings, such as tutoring systems involving human learners, where the learner's internal model is unknown. In contrast, the existing state-of-the-art work [28] has proposed interactive curriculum algorithms that are based on the dynamics of a specific IRL learner. A detailed comparison between our curriculum algorithm and the prior state of the art is presented in Section 1.1.

Our approach is motivated by works on curriculum design for supervised learning and reinforcement learning algorithms that use a ranking over the training examples using a difficulty score [29–35]. In particular, our work is inspired by theoretical results on curriculum learning for linear regression models [32]. We define difficulty scores for any demonstration based on the teacher's optimal policy and the learner's current policy. We then study the differential effect of the difficulty scores on the learning progress for two popular imitation learners: Maximum Causal Entropy Inverse Reinforcement Learning (MaxEnt-IRL) [10] and Cross-Entropy loss-based Behavioral Cloning (CrossEnt-BC) [36]. Our main contributions are as follows:[1]

1. Our analysis for both MaxEnt-IRL and CrossEnt-BC learners leads to a unified curriculum strategy, i.e., a preference ranking over demonstrations. This ranking is obtained based on the ratio between the demonstration's likelihood under the teacher's optimal policy and the learner's current policy. Experiments on a synthetic car driving environment validate our curriculum strategy.
2. For the MaxEnt-IRL learner, we prove that our curriculum strategy achieves a linear convergence rate (under certain mild technical conditions), notably without requiring access to the learner's internal dynamics.
3. We adapt our curriculum strategy to the learner-centric setting where a teacher agent is not present through the use of task-specific difficulty scores. As a proof of concept, we show that our strategy accelerates the learning process in synthetic navigation-based environments.

1.1 Comparison to Existing Approaches on Curriculum Design for Imitation Learning

In the non-interactive setting, [18] have proposed a batch teaching algorithm (SCOT) by showing that the teaching problem can be formulated as a set cover problem. In the interactive setting, [28] have proposed two teaching algorithms: (i) Omniscient algorithm (OMN), a principled approach based on the iterative machine teaching (IMT) framework from [37], and (ii) Blackbox algorithm (BBOX), a heuristic approach built by considering the gradient functional form of the linear MaxEnt-IRL learner model and the objective of the OMN algorithm. Our work is more closely related to the interactive teaching algorithms from [28]. However, compared to the IMT framework, our proposed curriculum strategy takes an entirely different perspective – it picks the next demonstration based on a preference ranking induced by the teacher's and learner's policies. This difference allows our curriculum strategy to be agnostic to the learner's internal dynamics, and applicable in both the teacher-centric and learner-centric settings as discussed below.

The teaching algorithms from [28] are very much tied to the MaxEnt-IRL learner model. In contrast, our curriculum strategy does not assume the knowledge of the learner's internal update rule (i.e., the gradient functional form). Furthermore, our curriculum strategy is applicable to the learner-centric setting where no teacher agent is present and the learner designs a self-curriculum (i.e., automatically orders demonstrations). The teaching algorithms from [28] cannot be directly extended to the learner-centric setting as they require knowledge of either the learner's internal dynamics or the true reward function of the environment.

1.2 Additional Related Work on Curriculum Design and Teaching

Curriculum design. Curriculum design for supervised learning settings has been extensively studied in the literature. Early works present the idea of designing a curriculum comprising of tasks with increasing difficulty to train a machine learning model [29–31]. However, these approaches require task-specific knowledge for designing heuristic difficulty measures. Recent works have tackled the

[1] Github repo: https://github.com/adishs/neurips2021_curriculum-teaching-demonstrations_code.

problem of automating curriculum design [38, 39]. There is also an increasing interest in theoretically analyzing the impact of a curriculum (ordering) of training tasks on the convergence of supervised learner models [32, 40, 41]. In particular, our work builds on the idea of difficulty scores of the training examples studied in [32].

The existing results on curriculum design for sequential decision-making settings are mostly empirical in nature. Similar to the supervised learning settings, the focus on curriculum design for reinforcement learning settings has been shifted from hand-crafted approaches [34, 35] to automatic methods [21–23, 25]. We refer the reader to a recent survey [42] on curriculum learning for reinforcement learning. Curriculum learning has also been studied in psychology literature [43–47]. One key aspect in these works has been to design algorithms which account for the pedagogical intentions of a teacher, which often aims to explicitly demonstrate specific skills rather than just provide the most optimal demonstration for a task. We see our work as complementary to these.

Machine teaching. The algorithmic teaching problem considers 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 [37, 48–50]. Most of the work in machine teaching for the supervised learning setting is on batch teaching 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 in the supervised learning setting [51–56].

Furthermore, [16–18, 28, 57, 58] have studied algorithmic teaching for sequential decision-making tasks. In particular, [17, 18] have proposed batch teaching algorithms for an IRL agent, where the teacher decides the entire set of demonstrations to provide to the learner before any interaction. These teaching algorithms do not utilize the feedback from the learner. In contrast, as discussed in Section 1.1, [28] have proposed interactive teaching algorithms (namely OMN and BBOX) for an IRL agent, where the teacher takes into account how the learner progresses. The works of [57, 58] are complementary to ours and study algorithmic teaching when the learner has a differet worldview than the teacher or has its own specific preferences.

2 Formal Problem Setup

Here, we formalize our problem setting which is based on prior work on sequential teaching [28, 37].

Environment. We consider an environment defined as a Markov Decision Process (MDP) $\mathcal{M} := \left(\mathcal{S}, \mathcal{A}, \mathcal{T}, \gamma, P_0, R^E\right)$, where the state and action spaces are denoted by $\mathcal{S}$ and $\mathcal{A}$, respectively. $\mathcal{T} : \mathcal{S} \times \mathcal{S} \times \mathcal{A} \to [0, 1]$ is the transition dynamics, γ is the discounting factor, and $P_0 : \mathcal{S} \to [0, 1]$ is an initial distribution over states $\mathcal{S}$. A policy $\pi : \mathcal{S} \times \mathcal{A} \to [0, 1]$ is a mapping from a state to a probability distribution over actions. The underlying reward function is given by $R^E : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$.

Teacher-learner interaction. We consider a setting with two agents: a teacher and a sequential learner. The teacher has access to the full MDP $\mathcal{M}$ and has a *target policy* π^E (e.g., a near-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 π^E to the learner. Here, a teacher's demonstration $\xi = \left\{\left(s^\xi_\tau, a^\xi_\tau\right)\right\}_{\tau=0,1,\ldots}$ is obtained by first choosing an initial state $s^\xi_0 \in \mathcal{S}$ (where $P_0(s^\xi_0) > 0$) and then choosing a trajectory, sequence of state-action pairs, obtained by executing the policy π^E in the MDP $\mathcal{M}$. The interaction between the teacher and the learner is formally described in Algorithm 1. For simplicity, we assume that the teacher directly observes the learner's policy π^L_t at any time t. In practice, the teacher could approximately infer the policy π^L_t by probing the learner and using Monte Carlo methods.

Generic learner model. Here, we describe a generic learner update rule for Algorithm 1. Let $\Theta \subseteq \mathbb{R}^d$ be a parameter space. The learner searches for a policy in the following parameterized policy space: $\Pi_\Theta := \{\pi_\theta : \mathcal{S} \times \mathcal{A} \to [0, 1]\,,\ \text{where } \theta \in \Theta\}$. For the policy search, the learner sequentially minimizes a loss function ℓ that depends on the policy parameter θ and the demonstration $\xi = \left\{\left(s^\xi_\tau, a^\xi_\tau\right)\right\}_\tau$ provided by the teacher. More concretely, we consider $\ell\left(\xi, \theta\right) := -\log \mathbb{P}\left(\xi|\theta\right)$, where $\mathbb{P}\left(\xi|\theta\right) = P_0(s^\xi_0) \cdot \prod_\tau \pi_\theta\left(a^\xi_\tau|s^\xi_\tau\right) \cdot \mathcal{T}\left(s^\xi_{\tau+1}|s^\xi_\tau, a^\xi_\tau\right)$ is the likelihood (probability) of the demonstration ξ under policy π_θ in the MDP $\mathcal{M}$. At time t, upon receiving a demonstration ξ_t pro-

Algorithm 1 Teacher-Learner Interaction

1: **Initialization:** Initial knowledge of learner π_1^L.
2: **for** $t = 1, 2, \ldots$ **do**
3: Teacher observes the learner's current policy π_t^L.
4: Teacher provides demonstration ξ_t to the learner.
5: Learner updates its policy to π_{t+1}^L using ξ_t.

vided by the teacher, the learner performs the following online projected gradient descent update: $\theta_{t+1} \leftarrow \text{Proj}_\Theta\left[\theta_t - \eta_t g_t\right]$, where η_t is the learning rate, and $g_t = \left[\nabla_\theta \ell\left(\xi_t, \theta\right)\right]_{\theta=\theta_t}$. Note that the parameter θ_1 reflects the initial knowledge of the learner. Given the learner's current parameter θ_t at time t, the learner's policy is defined as $\pi_t^L := \pi_{\theta_t}$.

Teaching objective. For any policy π, the value (total expected reward) of π in the MDP $\mathcal{M}$ is defined as $V^\pi := \sum_{s,a}\sum_{\tau=0}^{\infty}\gamma^\tau \cdot \mathbb{P}\left\{S_\tau = s \mid \pi, \mathcal{M}\right\} \cdot \pi\left(a \mid s\right) \cdot R^E\left(s, a\right)$, where $\mathbb{P}\left\{S_\tau = s \mid \pi, \mathcal{M}\right\}$ denotes the probability of visiting the state s after τ steps by following the policy π. Let π^L denote the learner's final policy at the end of teaching. The performance of the policy π^L (w.r.t. π^E) in $\mathcal{M}$ can be evaluated via $\left|V^{\pi^E} - V^{\pi^L}\right|$ [9, 59]. The teaching objective is to ensure that the learner's final policy ϵ-*approximates* the teacher's policy, i.e., $\left|V^{\pi^E} - V^{\pi^L}\right| \leq \epsilon$. The teacher aims to provide an optimized sequence of demonstrations $\{\xi_t\}_{t=1,2,\ldots}$ to the learner to achieve the teaching objective. The teacher's performance is then measured by the number of demonstrations required to achieve this objective. Based on existing work [28, 37], we assume that $\exists\, \theta^* \in \Theta$ such that $\pi^E = \pi_{\theta^*}$ (we refer to θ^* as the *target teaching parameter*). Similar to [28], we assume that a smoothness condition holds in the policy parameter space: $|V^{\pi_\theta} - V^{\pi_{\theta'}}| \leq \mathcal{O}\left(f\left(\|\theta - \theta'\|\right)\right) \forall \theta, \theta' \in \Theta$. Then, the teaching objective in terms of V^π convergence can be reduced to the convergence in the parameter space, i.e., we can focus on the quantity $\|\theta^* - \theta_t\|$.

3 Curriculum Design using Difficulty Scores

In this section, we introduce our curriculum strategy which is based on the concept of *difficulty scores* and is agnostic to the dynamics of the learner.

Difficulty scores. We begin by assigning a difficulty score $\Psi_\theta\left(\xi\right)$ for any demonstration ξ w.r.t. a parameterized policy π_θ in the MDP $\mathcal{M}$. Inspired by difficulty scores for supervised learning algorithms [32], we consider a difficulty score which is directly proportional to the loss function ℓ, i.e., $\Psi_\theta\left(\xi\right) \propto g\left(\ell\left(\xi, \theta\right)\right)$, for a monotonically increasing function g. Setting $g(\cdot) = \exp(\cdot)$ leads to $\Psi_\theta\left(\xi\right) := \frac{1}{\prod_\tau \pi_\theta(a_\tau^\xi | s_\tau^\xi)}$ for MDPs with deterministic transition dynamics. Based on this insight, we define the following difficulty score which we use throughout our work.

Definition 1. *The difficulty score of a demonstration ξ w.r.t. the policy π_θ in the MDP $\mathcal{M}$ is given by* $\Psi_\theta\left(\xi\right) := \frac{1}{\prod_\tau \pi_\theta(a_\tau^\xi | s_\tau^\xi)}$.

Intuitively, the difficulty score of a demonstration ξ w.r.t. an agent's policy is inversely proportional to the preference of the agent to follow the demonstration. Demonstrations with a higher likelihood under the agent's policy (higher preference) have a lower difficulty score and vice versa. With the above definition, the difficulty scores for any demonstration ξ w.r.t. the teacher's and learner's policies (at any time t) are respectively given by $\Psi^E\left(\xi\right) := \Psi_{\theta^*}\left(\xi\right)$ and $\Psi_t^L\left(\xi\right) := \Psi_{\theta_t}\left(\xi\right)$.

General curriculum strategy. Our curriculum strategy picks the next demonstration ξ_t to provide to the learner based on a preference ranking induced by the teacher's and learner's difficulty scores. The difficulty score of a demonstration ξ w.r.t. the teacher and learner (at any time t) is denoted by Ψ^E and Ψ_t^L respectively. Specifically, our curriculum strategy is given by:

$$\xi_t \leftarrow \arg\max_\xi \frac{\Psi_t^L\left(\xi\right)}{\Psi^E\left(\xi\right)}. \tag{1}$$

Teacher-centric and learner-centric settings. In the teacher-centric setting formalized in Section 2, our curriculum strategy utilizes the difficulty scores induced by the learner's current policy π_t^L and the teacher's policy π^E. From Eq. (1) and Definition 1, we obtain the following teacher-centric curriculum strategy: $\xi_t \leftarrow \arg\max_\xi \prod_\tau \frac{\pi^E(a_\tau^\xi|s_\tau^\xi)}{\pi_t^L(a_\tau^\xi|s_\tau^\xi)}$.

Additionally, we also consider the learner-centric setting where a teacher agent is not present and the target policy π^E is unknown. Here, the learner can benefit from designing a self-curriculum based on its current policy π_t^L. We adapt our curriculum strategy to this setting by utilizing task-specific domain knowledge to define the teacher's difficulty score $\Psi^E(\xi)$ for any demonstration ξ. From Eq. (1), given the learner's current policy π_t^L and the teacher's difficulty score $\Psi^E(\xi)$, the learner-centric curriculum strategy is given as follows: $\xi_t \leftarrow \arg\max_\xi \frac{1}{\Psi^E(\xi)\prod_\tau \pi_t^L(a_\tau^\xi|s_\tau^\xi)}$.

Note that our curriculum strategy only requires access to the learner's and teacher's policies (π_t^L and π^E) and does not depend on the learner's internal dynamics (i.e, its update rule as mentioned in Section 1). This makes our approach more widely applicable to practical applications where it is often possible to infer an agent's policy, but the internal update rule is unkown.

4 Theoretical Analysis of Our Curriculum Strategy

In this section, we present the theoretical analysis of our curriculum strategy for two popular learner models, namely, MaxEnt-IRL and CrossEnt-BC. For our analysis, we consider the teacher-centric setting as introduced in Section 2. Our curriculum strategy obtains a preference ranking over the demonstrations to provide to the learner based on the difficulty scores (see Definition 1). To this end, we analyze the relationship between the difficulty scores (w.r.t. the learner and the teacher) of the provided demonstration and the teaching objective (convergence towards the target teaching parameter θ^*) during each sequential update step of the learner.

Given two difficulty values $\psi^E, \psi^L \in \mathbb{R}$, we define the feasible set of demonstrations at time t as $\mathcal{D}_t\left(\psi^E, \psi^L\right) := \left\{\xi : \Psi^E\left(\xi\right) = \psi^E \text{ and } \Psi_t^L\left(\xi\right) = \psi^L\right\}$. This set contains all demonstrations ξ for which the teacher's difficulty score $\Psi^E(\xi)$ is equal to the value ψ^E, and the learner's difficulty score $\Psi^L(\xi)$ is equal to the value ψ^L. Let $\Delta_t\left(\psi^E, \psi^L\right)$ denote the expected convergence rate of the teaching objective at time t, given difficulty values ψ^E and ψ^L:

$$\Delta_t\left(\psi^E, \psi^L\right) := \mathbb{E}_{\xi_t|\psi^E,\psi^L}[\|\theta^* - \theta_t\|^2 - \|\theta^* - \theta_{t+1}(\xi_t)\|^2], \tag{2}$$

where the expectation is w.r.t. the uniform distribution over the set $\mathcal{D}_t\left(\psi^E, \psi^L\right)$. Below, for MaxEnt-IRL and CrossEnt-BC learner models, we analyse the differential effect of ψ^E and ψ^L on $\Delta_t\left(\psi^E, \psi^L\right)$, i.e., the effect of picking demonstrations with higher or lower difficulty scores on the learning progress.

4.1 Analysis for MaxEnt-IRL Learner

Here, we consider the popular MaxEnt-IRL learner model [10, 59, 60] in an MDP $\mathcal{M}$ with deterministic transition dynamics, i.e., $\mathcal{T} : \mathcal{S} \times \mathcal{S} \times \mathcal{A} \to \{0, 1\}$. The MaxEnt-IRL learner model uses a parametric reward function $R_\theta : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$ where $\theta \in \mathbb{R}^d$ is a parameter. The reward function R_θ also depends on a feature mapping $\phi : \mathcal{S} \times \mathcal{A} \to \mathbb{R}^{d'}$ which encodes each state-action pair (s, a) by a feature vector $\phi\left(s, a\right) \in \mathbb{R}^{d'}$. For our theoretical analysis, we consider R_θ with a linear form, i.e., $R_\theta\left(s, a\right) := \langle\theta, \phi\left(s, a\right)\rangle$ and $d = d'$. In our experiments, we go beyond these simplifications and consider environments with stochastic transition dynamics and non-linear reward functions.

Under the MaxEnt-IRL learner model, the parametric policy takes the following soft-Bellman form: $\pi_\theta(a|s) = \exp\left(Q_\theta(s, a) - V_\theta(s)\right)$, where $V_\theta(s) = \log\sum_a \exp Q_\theta(s, a)$ and $Q_\theta(s, a) = R_\theta(s, a) + \gamma\sum_{s'} \mathcal{T}(s'|s, a) \cdot V_\theta(s')$. For any given θ, the corresponding policy π_θ can be efficiently computed via the Soft-Value-Iteration procedure with reward R_θ (see [59, Algorithm. 9.1]). For the above setting and a given parameter θ, the probability distribution $\mathbb{P}\left(\xi|\theta\right)$ over the demonstration ξ takes the closed-form $\mathbb{P}\left(\xi|\theta\right) = \frac{\exp\left(\langle\theta,\mu^\xi\rangle\right)}{Z(\theta)}$, where $\mu^\xi := \sum_{\tau=0}^{\infty} \gamma^\tau \phi(s_\tau^\xi, a_\tau^\xi)$ and $Z\left(\theta\right)$ is a normalization factor. Then, at time t, the gradient of the MaxEnt-IRL learner is given by $g_t = \mu^{\pi_{\theta_t}} - \mu^{\xi_t}$, where

$\mu^\pi := \sum_{s,a} \sum_{\tau=0}^{\infty} \gamma^\tau \cdot \mathbb{P}\{S_\tau = s \mid \pi, \mathcal{M}\} \cdot \pi(a \mid s) \cdot \phi(s,a)$ is the feature expectation vector of policy π. We note that our curriculum strategy in Eq. (1) is not using knowledge of g_t.

For the MaxEnt-IRL learner, we obtain the following theorem, which shows the differential effect of the difficulty scores (w.r.t. the teacher and the learner) on the expected rate of convergence of the teaching objective $\Delta_t\left(\psi^E, \psi^L\right)$. We note that [32] obtained a similar results for linear regression learner models in the supervised learning setting.

Theorem 1. *Assume that η_t is sufficiently small for all t s.t. $\eta_t \|g_t\|^2 \ll 2\,|\langle \theta^* - \theta_t, g_t\rangle|$, where g_t is the gradient of the MaxEnt-IRL learner. Then, for the MaxEnt-IRL learner, the expected convergence rate of the teaching objective $\Delta_t\left(\psi^E, \psi^L\right)$ is:*

- *monotonically decreasing with value ψ^E, i.e., $\frac{\partial \Delta_t}{\partial \psi^E} < 0$, and*
- *monotonically increasing with value ψ^L, i.e., $\frac{\partial \Delta_t}{\partial \psi^L} > 0$.*

Theorem 1 suggests that choosing demonstrations with lower difficulty score w.r.t. the teacher's policy and higher difficulty score w.r.t. the learner's policy would lead to faster convergence. Our curriculum strategy in Eq. (1) induces a preference ranking over demonstrations that aligns with these insights of Theorem 1. Furthermore, the following theorem states that the particular form of combining the two difficulty scores used in curriculum strategy, Eq. (1), achieves linear convergence to the teaching objective. This is similar to the state-of-the-art OMN algorithm based on the IMT framework for sequential learners [28, 37]. Importantly, unlike the OMN algorithm, our curriculum strategy does not rely on specifics of the learner model when selecting demonstrations.

Theorem 2. *Consider Algorithm 1 with the MaxEnt-IRL learner and our curriculum strategy in Eq. (1). Then, the teaching objective $\|\theta^* - \theta_t\| \leq \epsilon$ is achieved in $t = \mathcal{O}(\log \frac{1}{\epsilon})$ iterations.*

In the above theorem, the constant terms suppressed by the $\mathcal{O}(\cdot)$ notation depend on the learning rate of the learner (η_t), the distance between the learner's initial parameter/knowledge and the target teaching parameter ($\|\theta^* - \theta_1\|$), and the *richness* of the set of demonstrations obtained by executing the policy π^E in the MDP $\mathcal{M}$. The *richness* notion is formally discussed in the Appendix.

4.2 Analysis for CrossEnt-BC Learner

Next, we consider the CrossEnt-BC learner model [15, 36]. In this case, the learner's parametric policy takes the following softmax form: $\pi_\theta(a|s) = \frac{\exp(H_\theta(s,a))}{\sum_{a'} \exp(H_\theta(s,a'))}$, where $H_\theta : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$ is a parametric scoring function that depends on the parameter $\theta \in \mathbb{R}^d$ and a feature mapping $\phi : \mathcal{S} \times \mathcal{A} \to \mathbb{R}^{d'}$. For our theoretical analysis, we consider a linear scoring function H_θ of the form $H_\theta(s,a) := \langle \theta, \phi(s,a)\rangle$ (with $d = d'$). Then, at time step t, the gradient g_t of the CrossEnt-BC learner is given by: $g_t = \sum_{\tau=0}^{\infty} \left(\mathbb{E}_{a \sim \pi_{\theta_t}(\cdot|s_\tau^{\xi_t})} \left[\phi(s_\tau^{\xi_t}, a)\right] - \phi(s_\tau^{\xi_t}, a_\tau^{\xi_t}) \right)$. In the experiments, we also consider non-linear scoring functions parameterized by neural networks.

Similar to Theorem 1, we obtain the following theorem for the CrossEnt-BC learner, which also justifies our curriculum strategy in Eq. (1).

Theorem 3. *Assume that η_t is sufficiently small for all t s.t. $\eta_t \|g_t\|^2 \ll 2\,|\langle \theta^* - \theta_t, g_t\rangle|$, where g_t is the gradient of the CrossEnt-BC learner. Then, for the CrossEnt-BC learner, the expected convergence rate of the teaching objective $\Delta_t\left(\psi^E, \psi^L\right)$ is (after first-order approximation):*

- *monotonically decreasing with ψ^E, i.e., $\frac{\partial \Delta_t}{\partial \psi^E} < 0$, and*
- *monotonically increasing with ψ^L, i.e., $\frac{\partial \Delta_t}{\partial \psi^L} > 0$.*

We note that the proof of Theorem 3 relies on the first-order Taylor approximation of the term $\sum_\tau \log \sum_{a'} \exp\left(H_\theta\left(s_\tau^{\xi_t}, a'\right)\right)$ around θ_t (detailed in the Appendix). Due to this approximation, it is more challenging to obtain a convergence result analogous to Theorem 2.

5 Experimental Evaluation: Teacher-Centric Setting

Inspired by the works of [28, 61, 62], we evaluate the performance of our curriculum strategy, Eq. (1), in a synthetic car driving environment on MaxEnt-IRL and CrossEnt-BC learners. In particular, we consider the environment of [28] and the teacher-centric setting of Section 2.

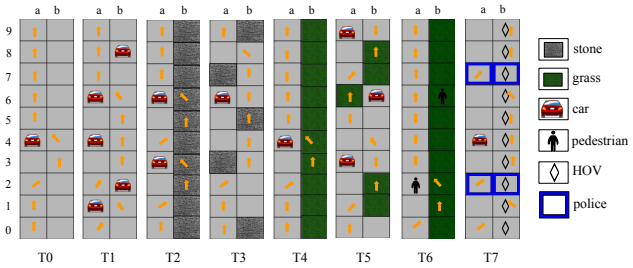

Figure 1: Car environment with 8 different types of tasks. Arrows represent the path taken by the teacher's policy.

$\phi^E(s)$	w^*
stone	-1
grass	-0.5
car	-5
ped	-10
car-front	-2
ped-front	-5
HOV	+1
police	0

Figure 2: Environment features $\phi^E(s)$ and reward weights w^*.

Car driving setup. Fig. 1 illustrates a synthetic car driving environment consisting of 8 different types of tasks, denoted as `T0`, `T1`, ..., `T7`. Each type is associated with different driving skills. For instance, `T0` corresponds to a basic setup representing a traffic-free highway. `T1` represents a crowded highway. `T2` has stones on the right lane, whereas `T3` has a mix of both cars and stones. Similarly, `T4` has grass on the right lane, and `T5` has a mix of both grass and cars. `T6` and `T7` introduce more complex features such as pedestrians, police, and HOV (high occupancy vehicles). The agent starts navigating from an initial state at the bottom of the left lane of each task, and the goal is to reach the top of a lane while avoiding cars, stones, and other obstacles. The agent's action space is given by $\mathcal{A} = \{$`left`, `straight`, `right`$\}$. Action `left` steers the agent to the left of the current lane. If the agent is already in the leftmost lane when taking action `left`, then the lane is randomly chosen with uniform probability. We define similar stochastic dynamics for taking action `right`; action `straight` means no change in the lane. Irrespective of the action taken, the agent always moves forward.

Environment MDP. Based on the above setup, we define the environment MDP, $\mathcal{M}_{\text{car}}$, consisting of 8 types of tasks, namely `T0`–`T7`, and 5 tasks of each type. Every location in the environment is associated with a state. Each task is of length 10 and width 2, leading to a state space of size $5 \times 8 \times 20$. We consider an action-independent reward function $R^{E_{\text{car}}}$ that is dependent on an underlying feature vector ϕ^E (see Fig. 2). The feature vector of a state s, denoted by $\phi^E(s)$, is a binary vector encoding the presence or absence of an object at the state. In this work we have two types of features: features indicating the type of the current cell as `stone`, `grass`, `car`, `ped`, `police`, and `HOV`, as well as features providing some look-ahead information such as whether there is a car or pedestrian in the immediate front cell (denoted as `car-front` and `ped-front`). For `T0`-`T6` the reward w.r.t. the underlying feature vector ϕ^E is given by $:= \langle w^*, \phi^E \rangle$ (see Fig. 2). Essentially there are different penalties (negative rewards) for colliding with specific obstacles such as `stone` and `car`. In tasks of type `T7`, there is a reward of value $+1$ for driving on `HOV`, however, if `police` is present while driving on `HOV`, a reward value of -5 is obtained. Overall, this results in a nonlinear reward function $R^{E_{\text{car}}}$.

5.1 Teaching Algorithms

Here, we introduce the teaching algorithms considered in our experiments. The teacher's near-optimal policy π^E is obtained via policy iteration [63]. The teacher selects demonstrations to provide to the learner using its teaching algorithm. We compare the performance of our proposed CUR teacher, which implements our strategy in Eq. (1), against the following baselines:

- CUR-T: A variant of our CUR teacher that samples demonstrations based on the difficulty score Ψ^E alone, and sets Ψ_t^L to constant.
- CUR-L: A similar variant of our CUR teacher that samples demonstrations based on the difficulty score Ψ_t^L alone, and sets Ψ^E to constant.
- AGN: an agnostic teacher that picks demonstrations based on random ordering [28, 32].
- OMN: The omniscient teacher is a state-of-the-art algorithm [28, 37], which is applicable only to MaxEnt-IRL learners. With complete knowledge of θ^*, the learner's internal parameter θ_t and the learner's gradient update rule, the teacher picks demonstrations to directly steer the learner towards θ^*, i.e., the teacher minimizes $\|\theta^* - (\theta_t - \eta_t g_t)\|$.

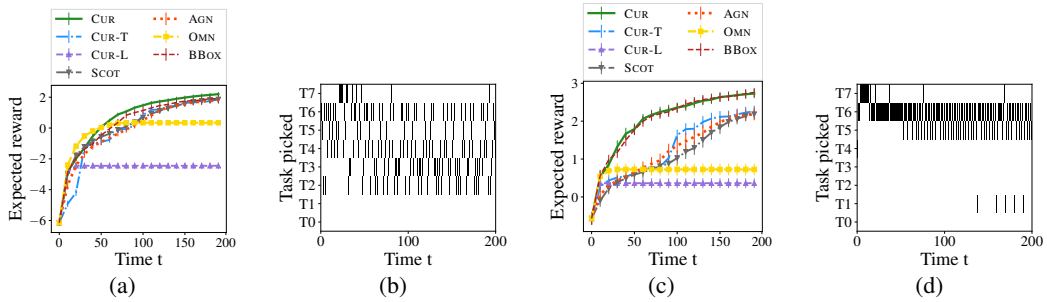

Figure 3: Learning curves and curriculum visualization for MaxEnt-IRL learners (with varying initial knowledge) trained on the car driving environment: (a) reward convergence plot and (b) curriculum generated by the CUR teacher for the learner with initial knowledge of `T0`; (c) reward convergence plot and (d) curriculum generated by the CUR teacher for the learner with initial knowledge of `T0`–`T3`.

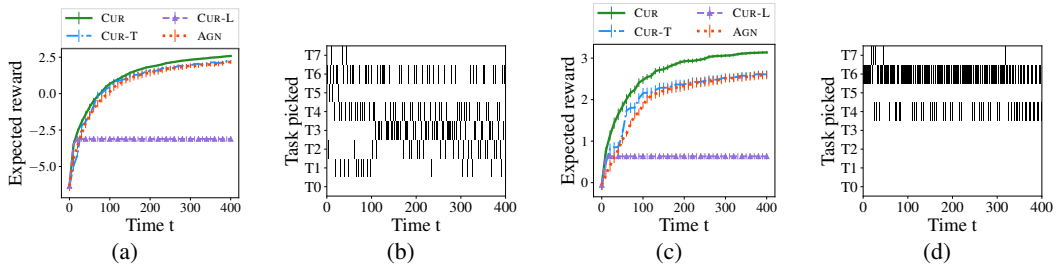

Figure 4: Learning curves and curriculum visualization for CrossEnt-BC learners (with varying initial knowledge) trained on the car driving environment: (a) reward convergence plot and (b) curriculum generated by the CUR teacher for the learner with initial knowledge of `T0`; (c) reward convergence plot and (d) curriculum generated by the CUR teacher for the learner with initial knowledge of `T0`–`T3`.

- BBOX: The blackbox teacher [28] is designed based on an approximate functional form of the gradient of the linear MaxEnt-IRL learner model[2]. Specifically the teacher picks a demonstration ξ which maximizes $|\sum_{s',a'}\{\rho^{\pi_t^L}(s',a') - \rho^{\xi}(s',a')\}R^E(s',a')|$. The BBOX teacher does not require access to θ^*or the learner's internal parameter θ_t, however it does require exact access to the environment MDP's true reward function, R^E.
- SCOT: The SCOT teacher uses a set cover algorithm [18] to pre-compute the entire curriculum as a batch, prior to training, and hence does not incorporate learner feedback. It aims to find the smallest set of demonstrations required to teach an optimal reward function to an IRL learner.

5.2 Learner Models

Next, we describe the MaxEnt-IRL and CrossEnt-BC learner models. All seven teaching algorithms are applicable to the MaxEnt-IRL learner, while for the CrossEnt-BC learner we evaluate the CUR, CUR-T, CUR-L, and AGN algorithms.

MaxEnt-IRL learner. For alignment with the prior state-of-the-art work on teaching sequential MaxEnt-IRL learners [28], we perform *teaching over states* in our experiments. More concretely, at time t the teacher picks a state s_t (where $P_0(s_t) > 0$) and provides all demonstrations starting from s_t to the learner given by $\Xi_{s_t} = \left\{\xi = \left\{\left(s_\tau^\xi, a_\tau^\xi\right)\right\}_\tau \text{ s.t. } s_0^\xi = s_t\right\}$. The gradient g_t of the MaxEnt-IRL learner is then given by $g_t = \mu^{\pi_{\theta_t}, s_t} - \mu^{\Xi_{s_t}}$, where (i) $\mu^{\Xi_{s_t}} := \frac{1}{|\Xi_{s_t}|}\sum_{\xi\in\Xi_{s_t}}\mu^\xi$, and (ii) μ^{π,s_t} is the feature expectation vector of policy π with starting state set to s_t (see Section 4.1). The learner's feature mapping is given by $\phi(s,a) = \phi^E(s)$ and a non-linear parametric reward function, $R_\theta^L(s,a) = \langle\theta_{1:d'}, \phi(s,a)\rangle + \langle\theta_{d'+1:2d'}, \phi(s,a)\rangle^2$ is chosen, where d' is dimension of $\phi(s,a)$. As explained in [28], a linear reward representation cannot capture the optimal behaviour for $\mathcal{M}_{\text{car}}$. We consider learners with varying levels of initial knowledge, i.e., the learner is trained on a subset

[2]Note that for a linear reward function $R^E = \langle\theta^*, \phi(s,a)\rangle$, the BBOX teacher's objective is equivalent to $|\langle\theta^*, g_t\rangle|$, where g_t is the gradient of the MaxEnt-IRL learner model (see Section 4.1).

of tasks before the teaching process starts. In this setting, for our curriculum strategy in Eq. (1) the difficulty score of a set of demonstrations is computed as the mean difficulty score (Definition 1) of the individual demonstrations in the set.

CrossEnt-BC learner. We consider the CrossEnt-BC learner model of Section 4.2 as our second learner model. The learner's feature mapping is given by $\phi(s,a) = \mathbb{E}_{s' \sim \mathcal{T}(\cdot|s,a)}[\phi^E(s')]$. A quadratic parametric form is selected for the scoring function, i.e. $H_\theta(s,a) = \langle \theta_{1:d'}, \phi(s,a) \rangle + \langle \theta_{d'+1:2d'}, \phi(s,a) \rangle^2$, where d' is the dimension of $\phi(s,a)$. We consider learners with varying initial knowledge and perform teaching over states similar to the MaxEnt-IRL learner.

5.3 Experimental results

Figs. 3a, 3c and 4a, 4c show the convergence of the total expected reward for the MaxEnt-IRL and CrossEnt-BC learners respectively, averaged over 10 runs. The CUR teacher outperforms OMN despite not requiring information about the learner's dynamics. For non-linear parametric reward functions, the MaxEnt-IRL learner no longer solves a convex optimization problem. As a result, forcing the learner to converge to a fixed parameter doesn't necessarily perform well, as seen by the poor performance of the OMN teacher in Fig. 3c. The CUR teacher is competitive with the BBOX teacher. Unlike our CUR teacher, the BBOX teacher does require exact access to the true reward function, R^E. The CUR teacher consistently outperforms the AGN and SCOT teachers, as well as both the CUR-T and CUR-L variants.

Figs. 3b, 3d and 4b, 4d visualize the curriculum generated by the CUR teacher for the MaxEnt-IRL and CrossEnt-BC learners respectively. Here, the curriculum refers to the type of task, `T0`–`T7`, associated with the demonstrations provided by the teacher to the learner at time step t. For both types of learners we see that at the beginning of training, the teacher focuses on tasks which teach skills the learner is yet to master. For example, in Fig. 4d, the teacher picks tasks `T4`, `T6`, and `T7`, which teaches the learner to avoid grass, pedestrians, and to navigate through police and HOV. We also notice that the CUR teacher can identify degradation in performance on previously mastered tasks, e.g., task `T1` in Fig. 3d, and corrects for this by picking them again later during training.

Additional results under limited observability. In the above experiments, we consider the learner's policy to be fully observable by the teacher at every time step. Here, we study the performance of our CUR teacher under the limited observability setting, similar to [28], where the learner's policy needs to be estimated by probing the learner. The probing process is formally characterized by two parameters, (B, k), where the learner's policy is probed after every B time steps and each probing step corresponds to querying the learner's policy π_t^L a total of k times from each state $s \in \mathcal{S}$ in the MDP. The learner's policy, $\pi_t^L(a|s)\ \forall a, s$, is then approximated based on the fraction of the k queries in which the learner performed action a from state s. In between every B time steps that the learner is probed, the CUR teacher does not update its estimate of the learner's policy. We note that the $(B = 1, k \to \infty)$ setting corresponds to full observability of the learner. Fig. 5 depicts the performance of the CUR teacher for different values of (B, k). Even under limited observability, the CUR teacher's performance is competitive with the full observability setting. The performance of $(B = 40, k \to \infty)$ is even slightly better at certain time steps during later stages of training compared to $(B = 1, k \to \infty)$, which is possibly due to the strategy of greedily picking demonstrations not being necessarily optimal. Also, for the limited observability setting it can be interesting to explore approaches that alleviate the need to query the full policy of the learner [64, 65].

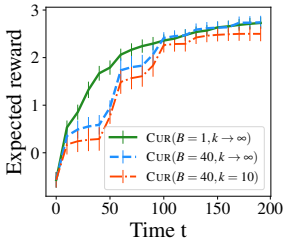

Figure 5: Learning curves for the MaxEnt-IRL learner under limited observability.

6 Experimental Evaluation: Learner-Centric Setting

In this section, we evaluate our curriculum strategy under a learner-centric setting, i.e., no teacher agent is present, and the teacher's difficulty $\Psi^E(\xi)$ is expressed by a task-specific difficulty score (see Section 3). We evaluate the effectiveness of our approach for training a multi-task neural policy to solve discrete optimization problems. Here, we provide an overview of the results with additional details in the Appendix for reproducibility.

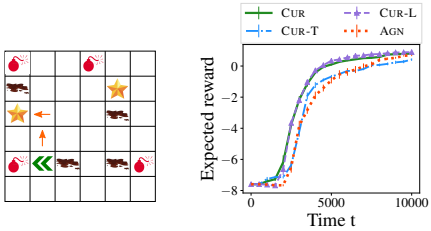

(a) (b)

Figure 6: Illustration of a shortest path navigation task (left) and convergence curves (right).

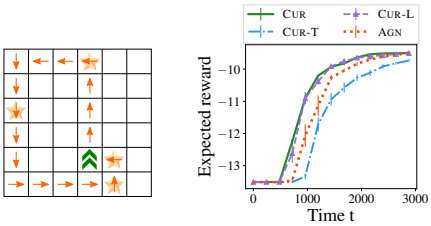

(a) (b)

Figure 7: Illustration of a travelling salesman navigation task (left) and convergence curves (right).

Experiment setup. We begin by describing the synthetic navigation-based environments considered in our experiments. Our first navigation environment comprises of tasks based on the shortest path problem [66]. We represent each task with a grid-world (see Fig. 6a) containing goal cells (depicted by stars), and cells with muds and bombs (shown in brown and red respectively). The agent aims to navigate to the closest goal cell while avoiding muds and bombs. Our second navigation environment comprises of tasks inspired by the travelling salesman problem (TSP) [67, 68] (see Fig. 7a). Again we represent each task with a grid-world, where the agent's goal is to find the shortest tour which visits all goals and returns to its initial location (see Fig. 7a). Orange arrows in Figs. 6a and 7a depict the optimal path for the agent. In our experimental setup, we begin by creating a pool of tasks and split them into training and test sets. The curriculum algorithms order the training tasks during the training phase based on their strategy to speed up the learner's progress. The aim of the learner is to learn a multi-task neural policy that can generalize to new unseen tasks in the test set.

Curriculum algorithms. We compare the performance of four different curriculum algorithms: (i) the CUR algorithm picks tasks from the training set using Eq. (1) (i.e., Ψ_t^L/Ψ^E), where Ψ^E is defined by a task-specific difficulty score (refer to the Appendix); (ii) the CUR-L algorithm picks tasks from the training set using Eq. (1) where the numerator is Ψ_t^L and the denominator is set to 1 (i.e., Ψ_t^L); (iii) the CUR-T algorithm picks tasks from the training set using Eq. (1) where the numerator is set to 1 and the denominator is set to Ψ^E (i.e., $1/\Psi^E$); (iv) the AGN algorithm picks tasks with a uniform distribution over the training set.

Learner model. We consider a neural CrossEnt-BC learner (see Section 4.2). The learner's scoring function H_θ is parameterized by a 6-layer Convolutional Neural Network (CNN). The CNN takes as input a feature mapping of the agent's current position in a task, and outputs a score for each action. The learner minimizes the cross-entropy loss between its predictions and the demonstrations.

Results. Figs. 6b and 7b, show the reward convergence curves on the test set for the different curriculum algorithms averaged over 5 runs. The CUR algorithm leads to faster reward convergence compared to the AGN algorithm, which is the common approach for training a neural policy. CUR-L is competitive with CUR in this setting which highlights the importance of the learner's difficulty.

7 Discussion and Conclusions

We presented a unified curriculum strategy, with theoretical guarantees, for the sequential MaxEnt-IRL and CrossEnt-BC learner models, based on the concept of difficulty scores. Our proposed strategy is independent of the learner's internal dynamics and is applicable in both teacher-centric and learner-centric settings. Experiments on a synthetic car driving environment and on navigation-based environments demonstrated the effectiveness of our curriculum strategy.

Our work provides theoretical underpinnings of curriculum design for teaching via demonstrations, which can be beneficial in educational applications such as tutoring systems and also for self-curriculum design for imitation learners. As such we do not see any negative societal impact of our work. Some of the interesting directions for future work include: obtaining convergence bounds for CrossEnt-BC and other learner models, designing curriculum algorithms for reinforcement learning agents based on the concept of difficulty scores, and designing approaches to efficiently approximate the learner's policy using less queries.

References

[1] S. Schaal. Learning from demonstration. In *NeurIPS*, pages 1040–1046, 1997.

[2] Brenna D Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot learning from demonstration. *Robotics and autonomous systems*, 57(5):469–483, 2009.

[3] Jens Kober, J Andrew Bagnell, and Jan Peters. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11):1238–1274, 2013.

[4] Maya Cakmak and Andrea L Thomaz. Eliciting good teaching from humans for machine learners. *Artificial Intelligence*, 217:198–215, 2014.

[5] 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*, 120(3):331–340, 2011.

[6] 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.

[7] D. A. Pomerleau. Efficient training of artificial neural networks for autonomous navigation. *Neural Computation*, 3(1):88–97, 1991.

[8] Stuart Russell. Learning agents for uncertain environments. In *COLT*, pages 101–103, 1998.

[9] Pieter Abbeel and Andrew Y Ng. Apprenticeship learning via inverse reinforcement learning. In *ICML*, page 1, 2004.

[10] Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, and Anind K Dey. Maximum entropy inverse reinforcement learning. In *AAAI*, 2008.

[11] Abdeslam Boularias, Jens Kober, and Jan Peters. Relative entropy inverse reinforcement learning. In *AISTATS*, pages 182–189, 2011.

[12] Markus Wulfmeier, Peter Ondruska, and Ingmar Posner. Maximum entropy deep inverse reinforcement learning. *arXiv preprint arXiv:1507.04888*, 2015.

[13] Chelsea Finn, Sergey Levine, and Pieter Abbeel. Guided cost learning: Deep inverse optimal control via policy optimization. In *ICML*, pages 49–58, 2016.

[14] Dorsa Sadigh, Anca D Dragan, Shankar Sastry, and Sanjit A Seshia. Active preference-based learning of reward functions. *RSS*, 2017.

[15] Takayuki Osa, Joni Pajarinen, Gerhard Neumann, J Andrew Bagnell, Pieter Abbeel, and Jan Peters. An algorithmic perspective on imitation learning. *Foundations and Trends® in Robotics*, 2018.

[16] Thomas J Walsh and Sergiu Goschin. Dynamic teaching in sequential decision making environments. *UAI*, 2012.

[17] Maya Cakmak and Manuel Lopes. Algorithmic and human teaching of sequential decision tasks. In *AAAI*, volume 26, 2012.

[18] Daniel S Brown and Scott Niekum. Machine teaching for inverse reinforcement learning: Algorithms and applications. In *AAAI*, 2019.

[19] Manuel Lopes, Francisco Melo, and Luis Montesano. Active learning for reward estimation in inverse reinforcement learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 31–46. Springer, 2009.

[20] Kareem Amin, Nan Jiang, and Satinder Singh. Repeated inverse reinforcement learning. In *NeurIPS*, pages 1815–1824, 2017.

[21] Maxwell Svetlik, Matteo Leonetti, Jivko Sinapov, Rishi Shah, Nick Walker, and Peter Stone. Automatic curriculum graph generation for reinforcement learning agents. In *AAAI*, volume 31, 2017.

[22] Carlos Florensa, David Held, Markus Wulfmeier, Michael Zhang, and Pieter Abbeel. Reverse curriculum generation for reinforcement learning. In *CORL*, pages 482–495, 2017.

[23] Carlos Florensa, David Held, Xinyang Geng, and Pieter Abbeel. Automatic goal generation for reinforcement learning agents. In *ICML*, 2018.

[24] Wojciech Czarnecki, Siddhant Jayakumar, Max Jaderberg, Leonard Hasenclever, Yee Whye Teh, Nicolas Heess, Simon Osindero, and Razvan Pascanu. Mix & match agent curricula for reinforcement learning. In *ICML*, 2018.

[25] Sanmit Narvekar and Peter Stone. Learning curriculum policies for reinforcement learning. In *AAMAS*, pages 25–33, 2019.

[26] Sainbayar Sukhbaatar, Zeming Lin, Ilya Kostrikov, Gabriel Synnaeve, Arthur Szlam, and Robert Fergus. Intrinsic motivation and automatic curricula via asymmetric self-play. In *ICLR*, 2018.

[27] Martin A Riedmiller, Roland Hafner, Thomas Lampe, Michael Neunert, Jonas Degrave, Tom Van de Wiele, Vlad Mnih, Nicolas Heess, and Jost Tobias Springenberg. Learning by playing solving sparse reward tasks from scratch. In *ICML*, 2018.

[28] Parameswaran Kamalaruban, Rati Devidze, Volkan Cevher, and Adish Singla. Interactive teaching algorithms for inverse reinforcement learning. In *IJCAI*, 2019.

[29] Jeffrey L Elman. Learning and development in neural networks: The importance of starting small. *Cognition*, 48(1):71–99, 1993.

[30] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *ICML*, pages 41–48, 2009.

[31] Wojciech Zaremba and Ilya Sutskever. Learning to execute. *arXiv preprint arXiv:1410.4615*, 2014.

[32] Daphna Weinshall, Gad Cohen, and Dan Amir. Curriculum learning by transfer learning: Theory and experiments with deep networks. In *ICML*, 2018.

[33] Daphna Weinshall and Dan Amir. Theory of curriculum learning, with convex loss functions. *arXiv preprint arXiv:1812.03472*, 2018.

[34] Minoru Asada, Shoichi Noda, Sukoya Tawaratsumida, and Koh Hosoda. Purposive behavior acquisition for a real robot by vision-based reinforcement learning. *Machine learning*, 23(2-3):279–303, 1996.

[35] Yuxin Wu and Yuandong Tian. Training agent for first-person shooter game with actor-critic curriculum learning. In *ICLR*, 2016.

[36] Michael Bain. A framework for behavioural cloning. In *Machine Intelligence 15*, pages 103–129, 1995.

[37] 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.

[38] Alex Graves, Marc G Bellemare, Jacob Menick, Rémi Munos, and Koray Kavukcuoglu. Automated curriculum learning for neural networks. In *ICML*, pages 1311–1320, 2017.

[39] Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *ICML*, pages 2304–2313, 2018.

[40] Tianyi Zhou and Jeff Bilmes. Minimax curriculum learning: Machine teaching with desirable difficulties and scheduled diversity. In *ICLR*, 2018.

[41] Tianyi Zhou, Shengjie Wang, and Jeff Bilmes. Curriculum learning by optimizing learning dynamics. In *AISTATS*, pages 433–441, 2021.

[42] Sanmit Narvekar, Bei Peng, Matteo Leonetti, Jivko Sinapov, Matthew E Taylor, and Peter Stone. Curriculum learning for reinforcement learning domains: A framework and survey. *Journal of Machine Learning Research*, 21:1–50, 2020.

[43] Mark K Ho, Michael Littman, James MacGlashan, Fiery Cushman, and Joseph L Austerweil. Showing versus doing: Teaching by demonstration. In *NeurIPS*, 2016.

[44] Mark K Ho, Michael Littman, and Joseph L Austerweil. Teaching by intervention: Working backwards, undoing mistakes, or correcting mistakes? In *Conference of the Cognitive Science Society*, 2017.

[45] Mark K Ho, Michael Littman, Fiery Cushman, and Joseph L Austerweil. Effectively learning from pedagogical demonstrations. In *Conference of the Cognitive Science Society*, 2018.

[46] Mark K Ho, Fiery Cushman, Michael Littman, and Joseph L Austerweil. People teach with rewards and punishments as communication not reinforcements. In *Journal of Experimental Psychology: General*, page 520–549, 2019.

[47] Mark K Ho, Fiery Cushman, Michael Littman, and Joseph L Austerweil. Communication in action: Planning and interpreting communicative demonstrations. In *Journal of Experimental Psychology: General*, 2021.

[48] Sally A Goldman and Michael J Kearns. On the complexity of teaching. *Journal of Computer and System Sciences*, 50(1):20–31, 1995.

[49] Xiaojin Zhu, Adish Singla, Sandra Zilles, and Anna N. Rafferty. An overview of machine teaching. *CoRR*, abs/1801.05927, 2018.

[50] Scott Cheng-Hsin Yang, Yue Yu, arash Givchi, Pei Wang, Wai Keen Vong, and Patrick Shafto. Optimal cooperative inference. In *AISTATS*, pages 376–385, 2018.

[51] Francisco S Melo, Carla Guerra, and Manuel Lopes. Interactive optimal teaching with unknown learners. In *IJCAI*, pages 2567–2573, 2018.

[52] Weiyang Liu, Bo Dai, Xingguo Li, Zhen Liu, James Rehg, and Le Song. Towards black-box iterative machine teaching. In *ICML*, 2018.

[53] Yuxin Chen, Adish Singla, Oisin Mac Aodha, Pietro Perona, and Yisong Yue. Understanding the role of adaptivity in machine teaching: The case of version space learners. In *NeurIPS*, 2018.

[54] Teresa Yeo, Parameswaran Kamalaruban, Adish Singla, Arpit Merchant, Thibault Asselborn, Louis Faucon, Pierre Dillenbourg, and Volkan Cevher. Iterative classroom teaching. In *AAAI*, 2019.

[55] Anette Hunziker, Yuxin Chen, Oisin Mac Aodha, Manuel Gomez Rodriguez, Andreas Krause, Pietro Perona, Yisong Yue, and Adish Singla. Teaching multiple concepts to a forgetful learner. In *NeurIPS*, 2019.

[56] Farnam Mansouri, Yuxin Chen, Ara Vartanian, Jerry Zhu, and Adish Singla. Preference-based batch and sequential teaching: Towards a unified view of models. In *NeurIPS*, 2019.

[57] Luis Haug, Sebastian Tschiatschek, and Adish Singla. Teaching inverse reinforcement learners via features and demonstrations. In *NeurIPS*, 2018.

[58] Sebastian Tschiatschek, Ahana Ghosh, Luis Haug, Rati Devidze, and Adish Singla. Learner-aware teaching: Inverse reinforcement learning with preferences and constraints. In *NeurIPS*, 2019.

[59] Brian D Ziebart. *Modeling Purposeful Adaptive Behavior with the Principle of Maximum Causal Entropy*. PhD thesis, University of Washington, 2010.

[60] Brian D Ziebart, J Andrew Bagnell, and Anind K Dey. The principle of maximum causal entropy for estimating interacting processes. *IEEE Transactions on Information Theory*, 59(4):1966–1980, 2013.

[61] Andrew Y Ng and Stuart J. Russell. Algorithms for inverse reinforcement learning. In *ICML*, 2000.

[62] Sergey Levine, Zoran Popovic, and Vladlen Koltun. Feature construction for inverse reinforcement learning. *NeurIPS*, 23:1342–1350, 2010.

[63] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 2018.

[64] Alexis Jacq, Matthieu Geist, Ana Paiva, and Olivier Pietquin. Learning from a learner. In *ICML*, 2019.

[65] Daniel S Brown, Jordan Schneider, Anca Dragan, and Scott Niekum. Value alignment verification. In *ICML*, 2021.

[66] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numer. Math.*, 1:269–271, December 1959.

[67] Wouter Kool, Herke van Hoof, and Max Welling. Attention, learn to solve routing problems! In *ICLR*, 2019.

[68] Chaitanya K. Joshi, Thomas Laurent, and Xavier Bresson. On learning paradigms for the travelling salesman problem. In *NeurIPS graph representation learning workshop*, 2019.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] In Sections 3 and 4 we present our curriculum strategy and provide theoretical analysis as described in the introduction. The experimental results mentioned in the abstract and introduction are detailed in Sections 5 and 6.
   (b) Did you describe the limitations of your work? [Yes] We discuss the limitations of our work in Sections 5 and 7. We point out that convergence bounds for CrossEnt-BC learners have not been obtained yet and that it would beneficial to explore approaches to estimate the learner's policy π_t^L without fewer queries.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] As stated in Section 7, this work presents the theoretical underpinnings of curriculum design for teaching via demonstrations. This can be beneficial for education applications and faster training of IRL/BC learners. As such in the present form, we do not see any direct negative societal impacts of our work.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes] We confirm that our paper conforms with the ethics review guidelines.
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] In Section 4.1 we state our assumptions of a deterministic MDP and linear reward structure. Similarly, we state our assumption of a linear scoring function for the CrossEnt-BC learner in Section 4.2. The simplifications considered are reiterated in the proofs.
   (b) Did you include complete proofs of all theoretical results? [Yes] Complete proofs of all theoretical results are included in the appendix.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] The code and instructions are included as a URL.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] Training details are presented in the appendix and are also present in the code.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] Error bars are included in all the result graphs, as can be seen in figures.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] The details are provided in the appendix.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]