

Programmatic Reward Design by Example

Weichao Zhou and Wenchao Li

Boston University
{zwc662,wenchao}@bu.edu

Abstract

Reward design is a fundamental problem in reinforcement learning (RL). A misspecified or poorly designed reward can result in low sample efficiency and undesired behaviors. In this paper, we propose the idea of *programmatic reward design*, i.e. using programs to specify the reward functions in RL environments. Programs allow human engineers to express sub-goals and complex task scenarios in a structured and interpretable way. The challenge of programmatic reward design, however, is that while humans can provide the high-level structures, properly setting the low-level details, such as the right amount of reward for a specific sub-task, remains difficult. A major contribution of this paper is a probabilistic framework that can infer the best candidate programmatic reward function from expert demonstrations. Inspired by recent generative-adversarial approaches, our framework searches for the most likely programmatic reward function under which the optimally generated trajectories cannot be differentiated from the demonstrated trajectories. Experimental results show that programmatic reward functions learned using this framework can significantly outperform those learned using existing reward learning algorithms, and enable RL agents to achieve state-of-the-art performance on highly complex tasks.

Introduction

Reward signals are an integral part of reinforcement learning (RL). Most conventional reward functions are goal-driven – they reward the agent only at the end of each episode. The sparsity of such reward signals, however, can make RL algorithms highly inefficient. As the complexity of the task increases, it becomes difficult for the agent to grasp the intricacies of the task solely from a goal-driven reward (Amodei et al. 2016).

Inverse Reinforcement Learning (IRL) is a general paradigm that aims at recovering the intrinsic reward function of human experts from their demonstrations (Ng and Russell 2000; Ziebart et al. 2008). Earlier works of IRL require the provision of multiple feature functions to construct the reward function. More recent attempts use function approximation by means of deep neural networks to alleviate this limitation and have considerable success (Fu, Luo, and

Levine 2018; Finn, Levine, and Abbeel 2016; Finn et al. 2016). However, due to the lack of interpretability of the approximated reward functions, it is difficult to enforce specific correctness constraints in the reward learning process.

Recent works proposed logic-based reward functions (Li, Vasile, and Belta 2017; Camacho et al. 2019) to endow an RL agent with high-level knowledge of the task via logical specifications. The logic-constrained reward (Hasanbeig, Abate, and Kroening 2019) and reward machines (Icarte et al. 2018) explicitly represents reward functions as automata. However, it is still cumbersome to design the automata and they can be difficult to understand when the size of the automata are large. In this paper, we propose *programmatic reward functions*, i.e. reward functions represented as programs expressed in human-readable domain specific language (DSL). There are several benefits of using programmatic reward functions for RL policy training. First, programs allow human engineers to express sub-goals, complex task scenarios in a structural and interpretable way. The inclusion of such domain knowledge forms inductive biases that help improve the sample efficiency and performance of RL agents. Second, engineers can take advantage of the rich semantics of DSLs to explicitly memorize, manipulate and leverage hindsight experiences of the RL agent. Lastly, programs are amenable to the specification of symbolic constraints over the holes. In a typical design routine of a programmatic reward function, we assume that an engineer can provide human insights in the form of a partial program, or a *sketch* (Solar-Lezama 2008), analogous to human providing feature functions in IRL and logical specifications in logic-based reward designs, to express the high-level structures of the task specification or reward function. The *sketch* in essence defines the set of events and certain interactions among them that the human engineer deem relevant to the task. The low-level details, such as the right amount of reward for a specific event or sub-task, are left as *holes*.

Similar to Programming by Example (PBE) (Menon et al. 2013), we propose to infer the holes in a programmatic reward sketch from expert demonstrated trajectories. A key difference of our approach from PBE is that the demonstrated trajectories do not directly correspond to inputs or (intermediate) outputs of the program, but instead are assumed to be generated by an unknown expert policy that is optimal under some realization of the programmatic reward

function. A major contribution of this paper is a probabilistic learning algorithm that can complete a given programmatic reward sketch based on expert demonstrations. Our overall framework, called *Programmatic Reward Design by Example (PRDBE)*, consists of three components: a set of *example trajectories* demonstrated by a human expert, a *program sketch* and a *symbolic constraint* that the complete program should satisfy. Directly searching in the program space is a combinatorial problem and can easily become intractable. Our approach is to search for the *most likely* program that matches the expert’s intrinsic reward function. Our solution is inspired by generative adversarial approaches (Finn et al. 2016; Jeon, Seo, and Kim 2018) which introduce a discriminator to distinguish agent trajectories from the expert’s. However, instead of formulating an agent’s policy as a generative model, we sample trajectories that are optimal under a candidate programmatic reward function and iteratively improve the candidate program to maximize the chance of the discriminator making false predictions on those trajectories. To circumvent the issue of non-differentiability of programs, we employ a sampler to sample candidate programs from the space of valid programs. In particular, we use self-normalized importance sampling to sample trajectories from an agent’s policy. We summarize our contributions below.

- We propose Programmatic Reward Design by Example (PRDBE), a novel paradigm to design and learn program-like reward functions for RL problems.
- We develop a probabilistic learning framework that can infer the most likely candidate reward program from expert demonstrations.
- Our approach enables RL agents to achieve state-of-the-art performance on highly complex environments with only a few demonstrations. In addition, we show that programmatic reward functions generalize across different environment configurations of the same task.

Related Work

Inverse Reinforcement Learning. IRL (Ng and Russell 2000; Abbeel and Ng 2004) instantiates a learning-from-demonstrations (LfD) framework to effectively infer the intrinsic reward functions of human experts. It is also notable for having infinite number of solutions. To resolve the ambiguity of IRL, max-entropy (Ziebart et al. 2008), max-margin (Abbeel and Ng 2004; Ratliff, Bagnell, and Zinkevich 2006) and Bayesian (Ramachandran and Amir 2007) methods have been proposed. However, those IRL methods assume linear rewards on the basis of non-linear state features, and also call RL in a loop to repeatedly solve the entire environments. The recent works (Fu, Luo, and Levine 2018; Ho and Ermon 2016; Jeon, Seo, and Kim 2018; Finn et al. 2016), drawing a connection between IRL and Generative Adversarial Networks (GANs) (Goodfellow et al. 2014), have achieved substantially improved the scalability of IRL by using deep neural networks and data-driven approaches. Our work, while embracing a data-driven ideology, represents the reward function in an interpretable way.

Reward Design. Reward shaping (Ng, Harada, and Russell 1999) highly speeds up RL training by modifying a sparse

reward functions with state-based potential functions. Intrinsic reward generation (Bellemare et al. 2016; Pathak et al. 2017; Alshiekh et al. 2018) and adversarially-guided (Flet-Berliac et al. 2021) techniques aims at motivating the agent to exhaustively explore the environments. Unlike these approaches, our work aims at capturing human knowledge in the reward function and does not generate uninterpretable reward signals densely ranging over the entire state space. Logic-based reward designs (Li, Vasile, and Belta 2017; Hasanbeig, Abate, and Kroening 2019; Camacho et al. 2019) present human knowledge in reward functions with specification languages such as linear temporal logic (LTL)(Baier and Katoen 2008). Reward machine theories (Icarte et al. 2022) further directly represent the reward functions as finite state automata which can be translated into logic formulas. Our work distinguishes itself by 1) using programming languages instead of logics to expressively represent human insights in the reward functions; 2) adopting LfD to implement low-level details in the reward functions. Regarding LfD, Inverse reward design (IRD) (Hadfield-Menell et al. 2017) design reward functions in a manner similar to IRL. Safety-aware apprenticeship learning (Zhou and Li 2018) incorporate formal specification and formal verification with IRL. However, those works restrict the reward function to be a linear combination of state features. Our work does not have such limitations.

Interpretable Reinforcement Learning. Learning interpretable RL policies has drawn continual attention (Andre and Russell 2001, 2002; Verma et al. 2018; Zhu et al. 2019; Yang et al. 2021; Tian et al. 2020). Our work focuses on programmatic reward functions instead of policies because well-designed reward functions can benefit diverse RL training methods. There have been a variety of works on learning-based program synthesis (Ellis et al. 2021; Parisotto et al. 2016; Itzhaky et al. 2010; Gulwani and Jain 2017; Ellis, Solar-Lezama, and Tenenbaum 2015). Our work is inspired by the concept of sketch synthesis (Solar-Lezama 2008). Realizing sketch synthesis from example for reward function is our main contribution in this paper.

Background

Reinforcement Learning. RL problems model an environment as a Markov Decision Process (MDP) $\mathcal{M} := \langle \mathcal{S}, \mathcal{A}, \mathcal{P}, d_0 \rangle$ where \mathcal{S} is a state space, \mathcal{A} is an action space, $\mathcal{P}(s'|s, a)$ is the probability of reaching a state s' by performing an action a at a state s , d_0 is an initial state distribution. A *trajectory* $\tau = s^{(0)}a^{(0)}s^{(1)}a^{(1)} \dots s^{(T)}a^{(T)}$ is produced by sequentially performing actions for T steps after starting from an initial state $s^{(0)} \sim d_0$. An RL agent can select actions according to a control *policy* $\pi(a|s)$ which determines the probability of performing action a at any state s . A *reward function* is a mapping $f : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ from state-action pairs to the real space. With a slight abuse of notations, we also represent the total reward and the joint probability along a trajectory τ as $f(\tau) = \sum_{t=0}^T f(s^{(t)}, a^{(t)})$ and $\pi(\tau) = \prod_{t=0}^T \pi(a^{(t)}|s^{(t)})$ respectively. Let $H(\pi)$ be the entropy of π . The objective of entropy-regularized RL (Levine 2018) is to minimize $J_{RL}(\pi) = \mathbb{E}_{\tau \sim \pi}[f(\tau)] - H(\pi)$.

Learning from Demonstrations. When the reward function is not given but a set E of expert trajectories τ_E 's is demonstrated by some unknown expert policy π_E , GAIL(Ho and Ermon 2016) trains an agent policy π_A to imitate π_E via a generative adversarial objective as in (2) where a discriminator $D : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is trained to maximize (2) so that it can identify whether an (s, a) is sampled from τ_E 's; π_A as the generator is optimized to minimize (2) so that its trajectories τ_A 's are indistinguishable from τ_E 's. Bayesian GAIL (Jeon, Seo, and Kim 2018) labels any expert trajectory τ_E with 1_E and 0_E to respectively indicate classifying τ_E as being sampled from E or not. Likewise, 1_A and 0_A indicate classifying any agent trajectory τ_A as from E or not. Bayesian GAIL learns the most likely D that makes the correct classifications $0_A, 1_E$ in terms of $p(D|0_A, 1_E; \pi_A; E) \propto p(D)p(0_A, 1_E|\pi_A, D; E) \propto \sum_{\tau_A} p(\tau_A|\pi_A)p(0_A|\tau_A; D) \sum_{\tau_E} p(\tau_E|E)p(1_E|\tau_E; D)$ where $p(1|\tau; D) = \prod_{t=0}^T D(s^{(t)}, a^{(t)})$. Its logarithm (1) is lower-bounded by (2) due to Jensen's inequality. It is shown in (Fu, Luo, and Levine 2018) that by representing $D(s, a) := \frac{\exp(f(s, a))}{\exp(f(s, a)) + \pi_A(a|s)}$ with a neural network f , (2) is maximized when $f \equiv \log \pi_E$, implying that f is a reward function under which π_E is optimal. Hence, by representing D in (1) and (2) with f , an IRL objective of solving the most likely expert reward function f is obtained.

$$\begin{aligned} & \log \sum_{\tau_E} p(\tau_E|E)p(1_E|\tau_E; D) \sum_{\tau_A} p(\tau_A|\pi_A)p(0_A|\tau_A; D) \quad (1) \\ & \geq \mathbb{E}_{\tau_E \sim E} \left[\log \prod_{t=0}^T D(s_E^{(t)}, a_E^{(t)}) \right] + \\ & \quad \mathbb{E}_{\tau_A \sim \pi_A} \left[\log \prod_{t=0}^T 1 - D(s_A^{(t)}, a_A^{(t)}) \right] \quad (2) \end{aligned}$$

Program Synthesis by Sketching. In the sketch synthesis problem, a human designer provides a *sketch* e , i.e., an incomplete program wherein certain details are left empty. Each unknown detail is called a *hole* and denoted as $?_{id}$, indexed by id in order of its appearance in e . The formalism of the sketches follows a general grammar Λ as in (3) and (4) where \mathcal{G} is a family of functions customized with domain-specific knowledge for the tasks; x represents the input argument of the program. The grammar Λ induces a set \mathcal{E} of syntactically allowable sketches. When a sketch $e \in \mathcal{E}$ is given by the designer, $?_e = \{?_1, ?_2, \dots\}$ denotes an ordered set of the holes appearing in e . Let \mathcal{H} be a set of possible assignments to $?_e$. The sketch e and \mathcal{H} induce a set $\mathcal{L} = \{l \triangleq e[\mathbf{h}/?_e] \mid \mathbf{h} \in \mathcal{H}\}$ of complete programs where $e[\mathbf{h}/?_e]$ means substituting $?_e$ in e with an assignment \mathbf{h} . Besides the syntax, the program l is required to be a function with a list type $[(\mathcal{S}, \mathcal{A})]$ argument. A valid assignment to the argument can be a list-represented trajectory, i.e., by writing $\tau = s^{(0)}a^{(0)} \dots s^{(t)}a^{(t)}$ as $[(s^{(0)}a^{(0)}), \dots, (s^{(t)}a^{(t)})]$. Hereinafter we refer to τ either as a sequence of state-action pairs or as a list of state-action tuples depending on the contexts. The rules in (5) and (6) define the semantics of Λ wherein (6) replaces x with τ in $g(e_1, \dots, e_n)$. The result of directly applying input τ to a sketch e is written as $\llbracket e \rrbracket(\tau)$ which induces a partial program with $?_e$ as free variables.

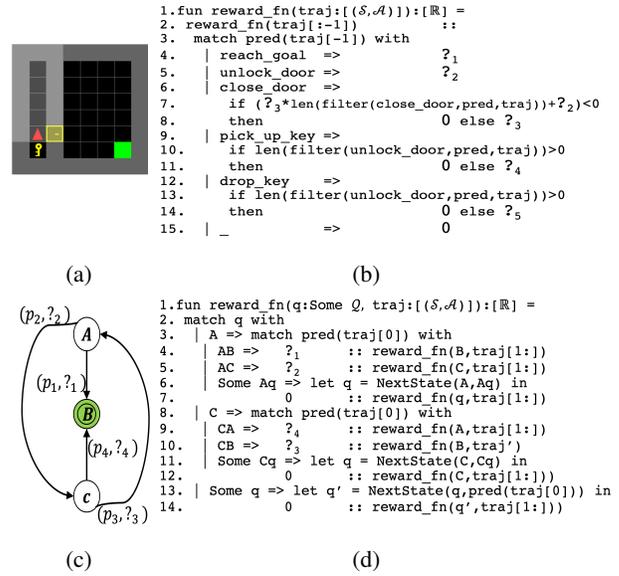


Figure 1: (a) MiniGrid 8x8 Door-Key; (b) A programmatic reward sketch for (a);(c) A reward function as a finite state automaton; (d) a programmatic reward sketch for (c)

For any complete program l , given a trajectory input τ , the output $\llbracket l \rrbracket(\tau)$ is required to be a $|\tau|$ -length real-valued list.

$$\text{Sketch } e := u \mid g(e_1, \dots, e_n) \quad g \in \mathcal{G} \quad (3)$$

$$\text{Term } u := \text{const} \mid ?_{id} \mid x \quad (4)$$

$$\llbracket \text{const} \rrbracket(\tau) := \text{const} \quad \llbracket ?_{id} \rrbracket(\tau) := ?_{id} \quad \llbracket x \rrbracket(\tau) := \tau \quad (5)$$

$$\llbracket g(e_1, \dots, e_n) \rrbracket(\tau) := g(e_1, \dots, e_n)[\tau/x] \quad (6)$$

Motivating Example

In this section, we motivate the problem of programmatic reward design with the pseudo-code of two programmatic reward function sketches, one for a navigation task in a grid-world and the other one for representing a reward function that has been formulated as a finite state automaton (FSA). For those two tasks, we assume that the domain expert provides a set of tokens such as `reach_goal`, `unlock_door` representing the behavior of the agent, and A, B, C representing the FSA states. A predicate function $\text{pred}(\cdot) \in \mathcal{G}$ of trajectory can output those tokens according to the last state-action in the input trajectory.

Example 1. (Door-Key Task) Fig.1a is an 8×8 Door-Key task in a Mini-Grid environment (Chevalier-Boisvert, Willems, and Pal 2018). An agent needs to pick up a key, unlock the yellow door on the grey wall and reach the green goal tile. In every step, the agent can observe at most 7×7 area in front if the area is not blocked by walls and doors. By default, the environment only returns a reward when the agent reaches the goal tile.

A sketch `reward_fn` of programmatic reward function for Example 1 is shown in Figure 1b. The holes $\{?_{id}\}_{id=1}^5$ are unknown numerical values. The input argument is written

as `traj` instead of `x` for readability. The output is a real-valued list concatenating the results of each recursive call to `reward_fn`. The elements in the output list are the very reward for the state-action pairs in the input trajectory in successive steps. This sketch responds to events such as reaching the goal and unlocking the door. We note that this programmatic reward function, whereas sparse in appearance considering the large state space, differs from the default goal-driven reward by informing the agent the stage-wise completion of the task. Line 7, 10 and 13 explicitly leverage hindsight experience to determine reward for the current step. Critically, $?_{id}$'s in Figure 1b ought not to be arbitrarily assigned. Suppose that a penalty $?_5 \leq 0$ for dropping the key is less than a award $?_4 \geq 0$ for picking up the key. The agent would repeatedly pick up and drop the key to obtain high net gain $?_5 + ?_4 \geq 0$ instead of reaching for the goal state. Besides, we observe that the penalties for redundant actions such as closing the door ought not to be overwhelmingly high. Otherwise an under-trained agent may be incentivized to reside away from the door for good.

Example 2. (Reward Function as A Finite State Automaton (FSA)) In Fig. 1c, a reward function for some RL problem is represented as an FSA comprised of at least 3 states A, B, C among which A indicates task initialization, B indicates task completion, C indicates the occurrence of some other event of interests. Each directed edge represents a state-transition triggered by a step in the environment. Other states and transitions are omitted. Each transition is annotated with $(p_{id}, ?_{id})$ where p_{id} is an unknown transition probability dependent on both the environment and the RL policy; $?_{id}$ is the reward sent to the agent at the moment of the transition. An RL agent is supposed to accomplish the task with minimal amount of steps in the environment. Note that the states in Fig. 1c are not to be confused with the MDP states in the environment; one step in the environment does not necessarily trigger the transitions drawn in Fig. 1c either.

FSA's such as the one in Example 2 are explicitly or implicitly adopted in several logic-based reward designs such as reward machines (Icarte et al. 2018). Such an FSA can be represented by the sketch in Fig. 1d by introducing a local variable q as the FSA state. The key problem again is to determine appropriate values for the $?_{id}$'s. We discuss the challenges facing this problem in the next section.

Problem Formulation

In this section we augment the concept of sketching. Then we characterize the problem of Programmatic Reward Design (PRD) and discuss the challenges in PRD.

Sketching with Symbolic Constraints

Given a sketch e , a designer can put constraints on $?_e$. A symbolic constraint c is a combination of multiple predicates $c := \mu \mid \neg\mu \mid c_1 \wedge c_2 \mid c_1 \vee c_2$ where the atomic predicate $\mu : \mathcal{H} \rightarrow \{\top, \perp\}$ is the basic building block. Under Boolean operations, the predicates follow the semantics (7) where the expression $\mu[\mathbf{h}/?_e]$ substitutes $?_e$ in μ with \mathbf{h} to output Boolean values. A satisfying implementation of the mapping $\text{eval} : \{\top, \perp\} \rightarrow \mathbb{R}$ is $\text{eval}(\cdot) := 2\mathcal{I}(\cdot) - 1$ where

$\mathcal{I} : \{\top, \perp\} \rightarrow \{0, 1\}$ is an indicator function. A symbolic constraint c is satisfied if $\llbracket c \rrbracket(\mathbf{h}) \geq 0$. Table 1 instantiates two predicates c_1, c_2 for Example 1. Then $c = c_1 \wedge c_2$ can be a symbolic constraint for Example 1. For simplicity, we only consider conjunctions of atomic predicates in this paper. Now we are now ready to state the Programmatic Reward Design problem as in Definition 1.

Properties	Predicates
$[c_1]$ Reward reaching the goal	$\bigwedge_{id=1}^5 (?_{id} \leq ?_1)$
$[c_2]$ Penalize dropping unused key	$?_5 + ?_4 \leq 0$

Table 1: The correspondence between two desired properties and the predicates for the sketch in Fig. 1b.

$$\begin{aligned}
\llbracket \mu \rrbracket(\mathbf{h}) &:= \text{eval}(\mu[\mathbf{h}/?_e]) & \llbracket \neg\mu \rrbracket(\mathbf{h}) &:= -\llbracket \mu \rrbracket(\mathbf{h}) \\
\llbracket c_1 \wedge c_2 \rrbracket(\mathbf{h}) &:= \min(\llbracket c_1 \rrbracket(\mathbf{h}), \llbracket c_2 \rrbracket(\mathbf{h})) \\
\llbracket c_1 \vee c_2 \rrbracket(\mathbf{h}) &:= \max(\llbracket c_1 \rrbracket(\mathbf{h}), \llbracket c_2 \rrbracket(\mathbf{h})) \quad (7)
\end{aligned}$$

Definition 1 (Programmatic Reward Design (PRD)). For an RL task, a PRD problem is a tuple $\langle e, ?_e, \mathcal{H}, c \rangle$ where e is a sketch with holes $?_e$ that takes values from set \mathcal{H} ; c is a symbolic constraint. The solution to a PRD problem $\langle e, ?_e, \mathcal{H}, c \rangle$ is any valid program $l \triangleq e[\mathbf{h}/?_e]$ subject to $\mathbf{h} \in \mathcal{H} \wedge \llbracket c \rrbracket(\mathbf{h}) \geq 0$.

Challenges in PRD

We note that solving the PRD problem does not guarantee that the resulting reward function will be effective. The challenge of assigning proper values to the holes is faced not only by PRD but also by other symbolic reward design approaches. We use Example 2 to illustrate two aspects of this challenge. *A) Goal-driven rewards.* The reward function specified via logically guided approaches (Hasanbeig, Abate, and Kroening 2019) can usually be translated into transition systems similar to the one in Fig. 1c. Many of those approaches end up only assigning non-trivial values to $?_1$ and $?_4$ while equalizing the rewards for all the other transitions. However, when p_1 and p_4 are extremely small, e.g. $p_1, p_4 \ll p_2, p_3$, such goal-driven reward assignment barely provide any guidance when the agent explores the environment. *B) Unknown dynamics.* The reward shaping approach from (Camacho et al. 2019) adopts reward structures similar to the one in Fig. 1c but allocates rewards to all transitions while ignoring the dynamics of the environment. This approach may result in inefficiency in large environments. For instance, if $p_2 p_4 < p_1 \ll p_2$, a global optimal policy would only aim at triggering one transition $A \rightarrow B$. However, the shaped reward may assign a non-trivial reward to $?_2$, causing the RL algorithm to spend excessive training episodes on a local-optimal policy that lingers over $A \rightarrow C \rightarrow B$.

Given a sketch such as the one in Fig. 1d, a PRD designer would also face the those challenges. Thus we assume that a PRD problem is accompanied by a set of demonstrations that show how an expert can accomplish the task, similar to the setting of IRL. These demonstrations will help narrow down the solutions to the PRD problem. We hereby propose Programmatic Reward Design by Example (PRDBE).

Programmatic Reward Design by Example

Similar to Bayesian GAIL introduced in the Background section, we consider a probabilistic inference perspective for formulating PRDBE. We first introduce a term to correlate programmatic reward functions with trajectory distributions.

Definition 2 (Nominal Trajectory Distribution). Given a programmatic reward function l , a **nominal trajectory distribution** of l is $\hat{p}(\tau|l) = p(\tau) \exp(l(\tau))$ where $p(\tau)$ is the probability of sampling τ under the passive dynamics of the environment; $l(\tau)$ is short for $\sum_t \llbracket l \rrbracket(\tau)[t]$. Furthermore, a **normalized nominal trajectory distribution** of l is $p(\tau|l) = p(\tau) \exp(l(\tau)) / Z_l$ where $Z_l = \sum_{\tau} p(\tau) \hat{p}(\tau|l)$.

The nominal trajectory distribution $\hat{p}(\tau|l)$ can be viewed as a proxy of a possibly non-Markov policy $\pi_l(a^{(t)}|s^{(t)}) \propto \exp(\llbracket l \rrbracket(\tau)[t])$. Such a policy trivially minimizes $J_{RL}(\pi_l)$ (Levine 2018), the RL loss described in the Background section. Intuitively, we search for an l^* such that π_{l^*} matches π_E . Given this intuition, we formally define the problem of Programmatic Reward Design by Example in Definition 3.

Definition 3 (Programmatic Reward Design by Example (PRDBE)). Given a set of expert demonstrated trajectories E and a PRD problem $\langle e, \mathcal{P}_e, \mathcal{H}, c \rangle$, the PRDBE problem $\langle e, \mathcal{P}_e, \mathcal{H}, c, E \rangle$ is to find a solution l^* to the PRD problem such that for any τ the nominal trajectory distribution satisfies $\hat{p}(\tau|l^*) = p(\tau|l^*) = p_E(\tau)$ where p_E is the probability of sampling τ from E .

However, solving the PRDBE problem requires addressing the following challenges: a) the set of solutions to the PRD problem may not contain a satisfying solution l^* for PRDBE, and b) the sketch may not be differentiable w.r.t the holes. In other words, there may not exist a perfect solution to the PRDBE problem and gradient-based optimizations may not be readily applicable to PRDBE. To overcome these issues, we propose a learning framework with a relaxed objective.

A Generative Adversarial Learning Framework

Our learning framework realizes the probability matching between $\hat{p}(\tau|l)$, $p(\tau|l)$ and $p_E(\tau)$ in a generative-adversarial fashion. It searches for an l such that even the best discriminator represented with a reward function f as mentioned in the Background section cannot distinguish trajectories sampled by p_E from those by $p(\tau|l)$. Given an intermediate, learned reward function f , while Bayesian GAIL trains a π_A to minimize the log-likelihood (1) of correct classifications between agent trajectory τ_A and expert trajectory τ_E , we learn an l to maximize the log-likelihood of false classifications, i.e., $\log p(1_A, 0_E|l, f; E) = \log \sum_{\tau_A} p(\tau_A|l) p(1_A|\tau_A; l, f) \sum_{\tau_E} p_E(\tau_E) p(0_E|\tau_E; l, f)$ with $p(1_A|\tau_A; l, f) := \prod_{t=0}^T \frac{\exp(f(s_A^{(t)}, a_A^{(t)}))}{\exp(f(s_A^{(t)}, a_A^{(t)}) + \exp(\llbracket l \rrbracket(\tau_A)[t]))}$ and $p(0_E|\tau_E; l, f) := \prod_{t=0}^T \frac{\exp(\llbracket l \rrbracket(\tau_E)[t])}{\exp(f(s_E^{(t)}, a_E^{(t)}) + \exp(\llbracket l \rrbracket(\tau_E)[t]))}$ being the *discriminators*; $p(\tau_A|l)$ being the *generator* of τ_A 's. Since the l 's are non-differentiable, we do not directly optimize l . Recalling that \mathcal{L} is the program space induced by

\mathcal{H} , we optimize a sampler $q : \mathcal{L} \rightarrow [0, 1]$ to concentrate the distribution density on those candidate programs l 's that incur high values of $\log p(1_A, 0_E|l, f; E)$. Due to the introduction of q , the log-likelihood of false classification becomes $\log p(1_A, 0_E|q, f; E) \geq \sum_{l \in \mathcal{L}} q(l) \log p(1_A, 0_E|l, f; E)$ which is further lower-bounded by (8). We then introduce an agent policy π_A for importance sampling of $p(\tau_A|l)$ in (9). Thus q and π_A together constitute the *generator*. After cancelling out the passive-dynamics induced $p(\tau_A)$ in the nominator and denominator as in (10), we handle the normalization constant Z_l via self-normalized importance sampling (Bugallo et al. 2017) as in (11) by i.i.d. sampling a set $\{\tau_{A,i}\}_{i=1}^m$ of trajectories with π_A . We refer to (11) as the generative objective $J_{gen}(q)$.

$$\begin{aligned} & \sum_{l \in \mathcal{L}} q(l) \log p(1_A, 0_E|l, f; E) \\ & \geq \mathbb{E}_{l \sim q} \left\{ \mathbb{E}_{\tau_A \sim p(\tau_A|l)} \left[\log p(1_A|\tau_A; l, f) \right] + \right. \\ & \quad \left. \mathbb{E}_{\tau_E \sim p_E} \left[\log p(0_E|\tau_E; l, f) \right] \right\} \quad (8) \\ & = \mathbb{E}_{l \sim q} \left\{ \mathbb{E}_{\tau_A \sim \pi_A} \left[\frac{p(\tau_A|l)}{p(\tau_A|\pi_A)} \log p(1_A|\tau_A; l, f) \right] + \right. \\ & \quad \left. \mathbb{E}_{\tau_E \sim p_E} \left[\log p(0_E|\tau_E; l, f) \right] \right\} \quad (9) \\ & \geq \mathbb{E}_{l \sim q} \left\{ \mathbb{E}_{\tau_A \sim \pi_A} \left[\frac{\exp(l(\tau_A))}{Z_l \pi_A(\tau_A)} \log p(1_A|\tau_A; l, f) \right] + \right. \\ & \quad \left. \mathbb{E}_{\tau_E \sim p_E} \left[\log p(0_E|\tau_E; l, f) \right] \right\} \quad (10) \\ & \approx \mathbb{E}_{l \sim q} \left\{ \frac{\sum_{i=1}^m \left(\frac{\exp(l(\tau_{A,i}))}{\pi_A(\tau_{A,i})} \right) \log p(1_A|\tau_{A,i}; l, f)}{\sum_{i=1}^m \frac{\exp(l(\tau_{A,i}))}{\pi_A(\tau_{A,i})}} + \right. \\ & \quad \left. \mathbb{E}_{\tau_E \sim p_E} \left[\log p(0_E|\tau_E; l, f) \right] \right\} \quad (11) \\ & \triangleq J_{gen}(q) \end{aligned}$$

Next, we note that the existence of symbolic constraint c imposes a prior $p(l|c)$ over the search space of l . We let $p(l|c)$ be a uniform distribution among the programs that do not violate c , i.e., $\{l \triangleq e[\mathbf{h}/\mathcal{P}_e][\llbracket c \rrbracket(\mathbf{h})] \geq 0\} \subseteq \mathcal{L}$ while being zero anywhere else. Then the objective of our learning framework for q becomes minimizing $D_{KL}(q(l)||p(l|1_A, 0_E, f; E, c))$ where $p(l|1_A, 0_E, f; E, c) = \frac{p(1_A, 0_E|l, f; E) p(l|c)}{p(1_A, 0_E|f; E)}$. We minimize this KL-divergence by maximizing its evidence lower-bound (ELBO) as in (12) which equals the sum of $J_{gen}(q)$, an entropy term $H(q)$ and a $J_c(q) \triangleq \mathbb{E}_{l \sim q} [\log p(l|c)]$.

$$\begin{aligned} ELBO(q) & = J_{gen}(q) - D_{KL}(q(l)||p(l|c)) \\ & = H(q) + J_c(q) + J_{gen}(q) \quad (12) \end{aligned}$$

In our implementation, the holes in the designed sketches are all unknown reals, i.e., $\mathcal{H} = \mathbb{R}^{|\mathcal{e}|}$. Rather than sampling from \mathcal{L} , we can use a neural network q_φ to specify the mean

Algorithm 1: Generative Adversarial PRDBE

Input: PRDBE tuple $(e, \mathcal{?}_e, \mathcal{H}, c, E)$, iteration number N , sample number m, K , optimization parameters α, β

Output: l^*, π^*

- 1: **initialization:** program space $\mathcal{L} = \{e[\mathbf{h}/\mathcal{?}_e] | \mathbf{h} \in \mathcal{H}\}$; an agent policy π_{ϕ_0} ; neural reward function f_{θ_0} ; sampler $q_{\phi_0} : \mathcal{L} \rightarrow [0, 1]$
 - 2: **while** iteration number $n = 0, 1, \dots, N$ **do**
 - 3: Sample trajectory set $\{\tau_{A,i}\}_{i=1}^m$ by using policy π_{ϕ_n}
 - 4: Calculate rewards $\{[\![l_n^*]\!] (\tau_{A,i})\}_{i=1}^m$ with the most likely program $l_n^* = \arg \max_l q_{\phi_n}(l)$
 - 5: Update ϕ_n to ϕ_{n+1} using policy learning algorithm, e.g. PPO (Schulman et al. 2017)
 - 6: Sample K programs $\{l_k\}_{k=1}^K$ by using q_{ϕ_n}
 - 7: Calculate rewards $\{[\![l_k]\!] (\tau_{A,i})\}_{i=1}^m\}_{k=1}^K$
 - 8: Update $\theta_{n+1} \leftarrow \theta_n + \alpha \nabla_{\theta} J_{adv}(f_{\theta_n})$
 - 9: Update $\varphi_{n+1} \leftarrow \varphi_n + \beta \nabla_{\varphi} ELBO(q_{\varphi_n})$
 - 10: **end while**
 - 11: **return** $l^* := \arg \max_l q_{\phi_N}(l)$ and $\pi^* := \pi_{\phi_N}$
-

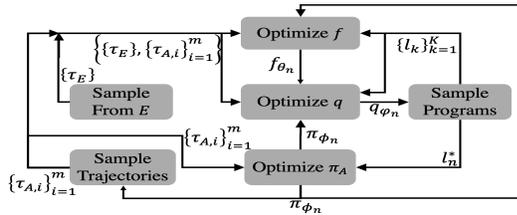


Figure 2: Flowchart for our learning framework

and variance of a $|\mathcal{?}_e|$ -dimensional multivariate Gaussian distribution in \mathcal{H} . We use $q_{\varphi}(l)$ to denote the probability of this Gaussian distribution producing an $\mathbf{h} \in \mathcal{H}$ such that $l = e[\mathbf{h}/\mathcal{?}_e]$. The mean of the Gaussian corresponds to the most likely l^* which we use to train a neural network policy π_{ϕ} as the agent policy π_A in (11). To calculate the gradient of J_{gen} , we use the logarithmic trick (Peters and Schaal 2008), i.e., $\nabla_{\varphi_n} \mathbb{E}_{l \sim q_{\varphi_n}}[\cdot] \approx \frac{1}{K} \sum_{k=1}^K \nabla_{\varphi_n} \log q_{\varphi_n}(l_k)[\cdot]$. We note that $J_c(q)$ is $-\infty$ once the support of q contains an l that violates c . Since the support of the Gaussian specified by q_{φ} is the real space, $J_c(q_{\varphi})$ can always be $-\infty$. Hence, we relax J_c by using a ReLU function and a cross-entropy loss to only penalize q_{φ} if l^* violates c . We also train a neural reward function f_{θ} as the f in (11) by maximize an adversarial objective $J_{adv}(f_{\theta})$ that differs from the generative objective J_{gen} by changing $1_A, 0_E$ into $0_A, 1_E$ such that $p(0_A | \tau_A; l, f_{\theta}) := \prod_{t=0}^T \frac{\exp([\![l]\!] (\tau_A)[t])}{\exp(f_{\theta}(s_A^{(t)}, a_A^{(t)})) + \exp([\![l]\!] (\tau_A)[t])}$ and $p(1_E | \tau_E; l, f_{\theta}) := \prod_{t=0}^T \frac{f_{\theta}(s_E^{(t)}, a_E^{(t)})}{\exp(f_{\theta}(s_E^{(t)}, a_E^{(t)})) + \exp([\![l]\!] (\tau_E)[t])}$. The algorithm is summarized in Algorithm 1 and depicted in Fig. 2. In a nutshell, this algorithm iteratively samples trajectories with π_{ϕ} , train π_{ϕ} with l^* , then update f_{θ} and q_{φ} respectively.

Experiments

In this section, we experimentally investigate: **A. Performance:** whether Algorithm 1 can efficiently train an agent policy to attain high performance; **B. Example Efficiency:** whether Algorithm 1 can work with only a few demonstrated trajectories; **C. Generalization:** whether the programs learned via Algorithm 1 in one environment can generalize across different environments of the same task.

Benchmark. We select from the MiniGrid environments (Chevalier-Boisvert, Willems, and Pal 2018) three challenging RL tasks in ascending order of difficulty: the basic setup and the default reward function for the *DoorKey* task have been described in Example 1; *KeyCorridor* shown in Fig.3e and *ObstructedMaze* in Fig.3i both require the agent to travel from room to room to find the key(s) to pick up a colored ball in some locked room. In *KeyCorridor*, all but one room are unlocked. In *ObstructedMaze*, most of the rooms are locked and obstructed by green balls and the keys are hidden in grey boxes which the agent must open first. Note that the agent can only carry one object at a time, which makes picking up and dropping the same objects multiple times almost inevitable. The environments can vary in size by changing the number of rooms and tiles (e.g. *DoorKey-8x8* vs. *DoorKey-16x16*). The placements of the objects and doors are randomized in each instance of an environment (e.g. *ObstructedMaze-Full*).

Baselines. We compare Algorithm 1 with IRL algorithms, GAN-GCL (Fu, Luo, and Levine 2018) and GAIL (Ho and Ermon 2016) to answer question A. We use PPO (Schulman et al. 2017), and AGAC (Flet-Berliac et al. 2021) for RL training in line 5 of Algorithm 1. We answer question B by varying the number of demonstrated trajectories when running Algorithm 1. We answer question C by using the programmatic reward functions learned via Algorithm 1 in small environments to train RL agents in larger environments (the results are annotated with AGAC/PPO.*prog*). In all three tasks, we provide the results of running the aforementioned RL algorithms as well as an intrinsic-reward augmented RL algorithm, RIDE (Raileanu and Rocktäschel 2020), with the default rewards. The sketches and symbolic constraints are in a similar form to those described in Example 1. We implement an LSTM reward function f_{θ} for Algorithm 1 and IRL baselines; an MLP q_{φ} ; a CNN version and an LSTM version of the actor-critic RL agent π_{ϕ} respectively but only report the version with the higher performance.

DoorKey. We use 10 example trajectories demonstrated in a *DoorKey-8x8* environment. The PRDBE problem and its solution are both annotated as *prog*. In Fig.3b, running Algorithm 1 by using PPO and AGAC in line 5 respectively produces a high-performance policy with fewer frames than by training PPO or AGAC with the default reward. On the other hand, RIDE, GAN-GCL and GAIL all fail with close-to-zero returns. In Fig.3c we reduce the number of examples from 10 to 1 and it does not affect the number of frames that Algorithm 1 needs to produce a policy with average return of at least 0.8, regardless of whether PPO or AGAC is used in line 5. This shows that Algorithm 1 is *example efficient*. In Fig.3h, we use the learned program to train PPO and AGAC

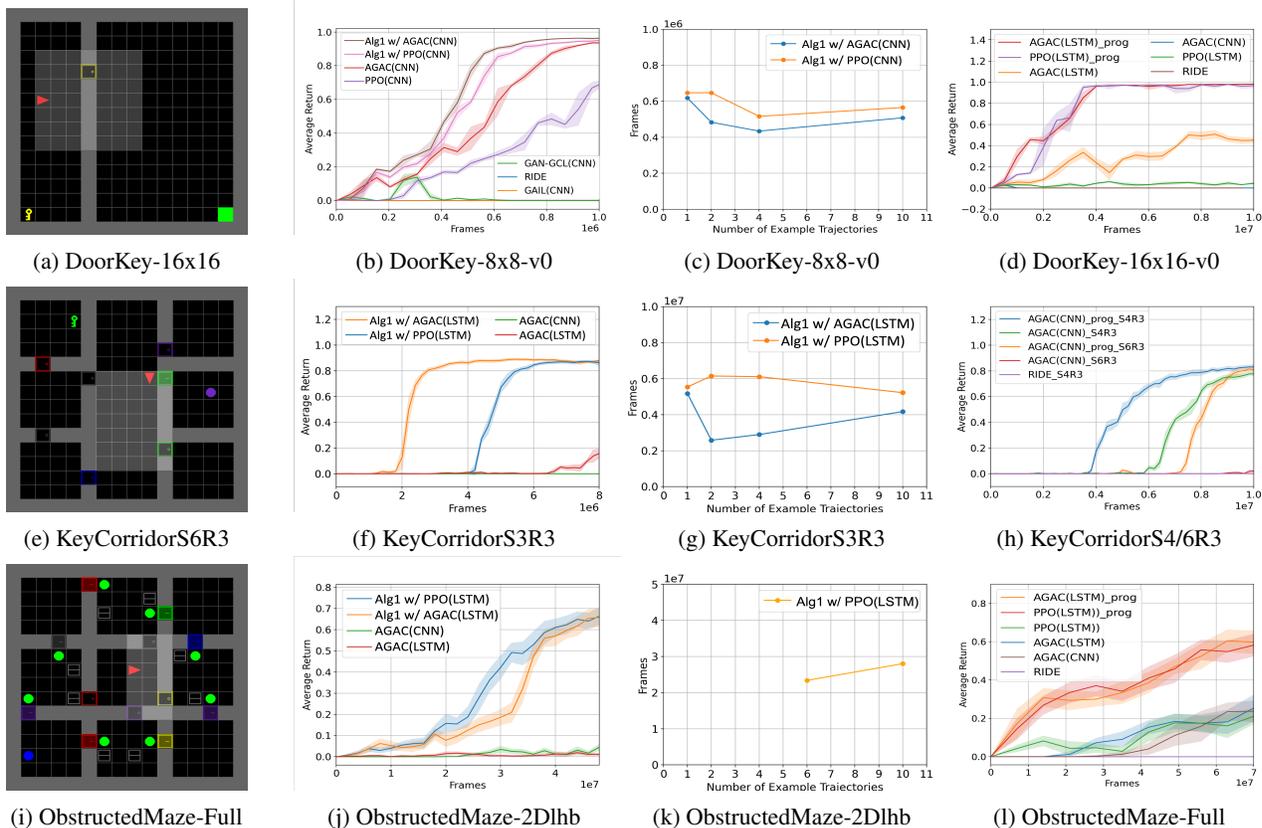


Figure 3: Frames are number of interactions with the environment. Average return is the average default reward achieved over a series of episodes and is no larger than 1. Alg1 w/ AGAC/PPO indicates using AGAC or PPO as the policy learning algorithm in line 5 of Algorithm 1. AGAC/PPO_{prog} indicates training an AGAC or PPO agent with reward provided by a programmatic reward function *prog*. CNN and LSTM indicate the structures of the actor-critic networks.

agents in a larger DoorKey environment and achieve high performances with fewer frames than training PPO, AGAC or RIDE with the default reward. This shows that the learned programmatic reward *generalizes* well across environments.

KeyCorridor. We use 10 example trajectories demonstrated in a 6×6 KeyCorridorS3R3 environment. In Fig.3f, by using PPO and AGAC in line 5 of Algorithm 1, we respectively obtain high performance with significantly fewer frames than by training AGAC with the default reward. We omit GAIL and GAN-GCL because both fail in this task. As shown in Fig.3g, reducing the number of examples (to 1) does not affect the performance of Algorithm 1. In Fig.3h, we use the learned program to train AGAC agents in two larger environments, 10×10 KeyCorridorS4R3 and 16×16 KeyCorridorS6R3, and achieve high performances with fewer frames than training AGAC with the default reward. Note that when using the default reward, AGAC has the prior SOTA performance for this task.

ObstructedMaze. We use 10 example trajectories demonstrated in a two-room ObstructedMaze-2Dh1b environment. When training the PPO agent, we discount PPO by episodic state visitation counts as in many exploration-driven approaches including AGAC. In Fig.3j we show that Algo-

1 can produce high-performance policies with fewer frames than AGAC trained with the default reward. Since AGAC(CNN) is the SOTA for this task, we do not show the results of other methods. In Fig.3h, we use the learned program to train AGAC agents in a larger environment with as many as 9 rooms, ObstructedMaze-Full, and achieve higher performances with fewer frames than other methods trained with the default reward. Regarding the number of demonstrations, as shown in Fig.3k, the performance of Algorithm 1 does not change much when the number is decreased to 6, but drops drastically when the number is further decreased possibly due to the higher complexity of this task.

Conclusion

We propose a novel paradigm for using programs to specify the reward functions in RL environments. We have developed a framework to complete a reward program sketch by learning from expert demonstrations. We experimentally validate our approach on challenging benchmarks and by comparing with SOTA baselines. Our future work will focus on reducing human efforts in the sketch creation process.

References

- Abbeel, P.; and Ng, A. Y. 2004. Apprenticeship Learning via Inverse Reinforcement Learning. In *Proceedings of the Twenty-first International Conference on Machine Learning, ICML '04*, 1–. New York, NY, USA: ACM. ISBN 1-58113-838-5.
- Alshiekh, M.; Bloem, R.; Ehlers, R.; Könighofer, B.; Niekum, S.; and Topcu, U. 2018. Safe reinforcement learning via shielding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Amodi, D.; Olah, C.; Steinhardt, J.; Christiano, P. F.; Schulman, J.; and Mané, D. 2016. Concrete Problems in AI Safety. arXiv:1606.06565.
- Andre, D.; and Russell, S. J. 2001. Programmable reinforcement learning agents. In *Advances in neural information processing systems*, 1019–1025.
- Andre, D.; and Russell, S. J. 2002. State abstraction for programmable reinforcement learning agents. In *AAAI/IAAI*, 119–125.
- Baier, C.; and Katoen, J.-P. 2008. *Principles of Model Checking (Representation and Mind Series)*. The MIT Press. ISBN 026202649X.
- Bellemare, M.; Srinivasan, S.; Ostrovski, G.; Schaul, T.; Saxton, D.; and Munos, R. 2016. Unifying count-based exploration and intrinsic motivation. *Advances in neural information processing systems*, 29: 1471–1479.
- Bugallo, M. F.; Elvira, V.; Martino, L.; Luengo, D.; Miguez, J.; and Djuric, P. M. 2017. Adaptive Importance Sampling: The past, the present, and the future. *IEEE Signal Processing Magazine*, 34(4): 60–79.
- Camacho, A.; Toro Icarte, R.; Klassen, T. Q.; Valenzano, R.; and McIlraith, S. A. 2019. LTL and Beyond: Formal Languages for Reward Function Specification in Reinforcement Learning. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, 6065–6073. International Joint Conferences on Artificial Intelligence Organization.
- Chevalier-Boisvert, M.; Willems, L.; and Pal, S. 2018. Minimalistic Gridworld Environment for OpenAI Gym. <https://github.com/maximecb/gym-minigrid>. Accessed: 2022-04-25.
- Ellis, K.; Solar-Lezama, A.; and Tenenbaum, J. 2015. Unsupervised Learning by Program Synthesis. In Cortes, C.; Lawrence, N. D.; Lee, D. D.; Sugiyama, M.; and Garnett, R., eds., *Advances in Neural Information Processing Systems* 28, 973–981. Curran Associates, Inc.
- Ellis, K.; Wong, C.; Nye, M.; Sablé-Meyer, M.; Morales, L.; Hewitt, L.; Cary, L.; Solar-Lezama, A.; and Tenenbaum, J. B. 2021. DreamCoder: bootstrapping inductive program synthesis with wake-sleep library learning. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, 835–850.
- Finn, C.; Christiano, P. F.; Abbeel, P.; and Levine, S. 2016. A Connection between Generative Adversarial Networks, Inverse Reinforcement Learning, and Energy-Based Models. arXiv:1611.03852.
- Finn, C.; Levine, S.; and Abbeel, P. 2016. Guided cost learning: Deep inverse optimal control via policy optimization. In *International conference on machine learning*, 49–58. PMLR.
- Flet-Berliac, Y.; Ferret, J.; Pietquin, O.; Preux, P.; and Geist, M. 2021. Adversarially Guided Actor-Critic. In *International Conference on Learning Representations*.
- Fu, J.; Luo, K.; and Levine, S. 2018. Learning Robust Rewards with Adversarial Inverse Reinforcement Learning. In *International Conference on Learning Representations*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Advances in neural information processing systems*, 2672–2680.
- Gulwani, S.; and Jain, P. 2017. Programming by Examples: PL meets ML. Springer.
- Hadfield-Menell, D.; Milli, S.; Abbeel, P.; Russell, S. J.; and Dragan, A. 2017. Inverse Reward Design. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Hasanbeig, M.; Abate, A.; and Kroening, D. 2019. Logically-constrained neural fitted q-iteration. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, 2012–2014*. International Foundation for Autonomous Agents and Multiagent Systems.
- Ho, J.; and Ermon, S. 2016. Generative adversarial imitation learning. In *Advances in Neural Information Processing Systems*, 4565–4573.
- Icarte, R. T.; Klassen, T.; Valenzano, R.; and McIlraith, S. 2018. Using Reward Machines for High-Level Task Specification and Decomposition in Reinforcement Learning. In Dy, J.; and Krause, A., eds., *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 2107–2116. PMLR.
- Icarte, R. T.; Klassen, T. Q.; Valenzano, R.; and McIlraith, S. A. 2022. Reward machines: Exploiting reward function structure in reinforcement learning. *Journal of Artificial Intelligence Research*, 73: 173–208.
- Itzhaky, S.; Gulwani, S.; Immerman, N.; and Sagiv, M. 2010. A Simple Inductive Synthesis Methodology and its Applications. In *OOPSLA/SPLASH'10, October 17-21, 2010, Reno/Tahoe, Nevada, USA*.
- Jeon, W.; Seo, S.; and Kim, K.-E. 2018. A Bayesian Approach to Generative Adversarial Imitation Learning. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Levine, S. 2018. Reinforcement Learning and Control as Probabilistic Inference: Tutorial and Review. .
- Li, X.; Vasile, C.-I.; and Belta, C. 2017. Reinforcement learning with temporal logic rewards. In *2017 IEEE/RSJ*

- International Conference on Intelligent Robots and Systems (IROS)*, 3834–3839. IEEE.
- Menon, A. K.; Tamuz, O.; Gulwani, S.; Lampson, B.; and Kalai, A. T. 2013. A Machine Learning Framework for Programming by Example. In *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28, ICML'13*, 1–187–1–195. JMLR.org.
- Ng, A. Y.; Harada, D.; and Russell, S. J. 1999. Policy Invariance Under Reward Transformations: Theory and Application to Reward Shaping. In *Proceedings of the Sixteenth International Conference on Machine Learning, ICML '99*, 278–287. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc. ISBN 1-55860-612-2.
- Ng, A. Y.; and Russell, S. J. 2000. Algorithms for Inverse Reinforcement Learning. In *Proceedings of the Seventeenth International Conference on Machine Learning, ICML '00*, 663–670. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc. ISBN 1-55860-707-2.
- Parisotto, E.; Mohamed, A.; Singh, R.; Li, L.; Zhou, D.; and Kohli, P. 2016. Neuro-Symbolic Program Synthesis. arXiv:1611.01855.
- Pathak, D.; Agrawal, P.; Efros, A. A.; and Darrell, T. 2017. Curiosity-driven exploration by self-supervised prediction. In *International conference on machine learning*, 2778–2787. PMLR.
- Peters, J.; and Schaal, S. 2008. Reinforcement learning of motor skills with policy gradients. *Neural networks*, 21(4): 682–697.
- Raileanu, R.; and Rocktäschel, T. 2020. RIDE: Rewarding Impact-Driven Exploration for Procedurally-Generated Environments. In *International Conference on Learning Representations*.
- Ramachandran, D.; and Amir, E. 2007. Bayesian inverse reinforcement learning. *Urbana*, 51(61801): 1–4.
- Ratliff, N. D.; Bagnell, J. A.; and Zinkevich, M. A. 2006. Maximum Margin Planning. In *Proceedings of the 23rd International Conference on Machine Learning, ICML '06*, 729–736. New York, NY, USA: ACM. ISBN 1-59593-383-2.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal Policy Optimization Algorithms. arXiv:1707.06347.
- Solar-Lezama, A. 2008. *Program synthesis by sketching*. University of California, Berkeley.
- Tian, L.; Ellis, K.; Kryven, M.; and Tenenbaum, J. 2020. Learning abstract structure for drawing by efficient motor program induction. *Advances in Neural Information Processing Systems*, 33.
- Verma, A.; Murali, V.; Singh, R.; Kohli, P.; and Chaudhuri, S. 2018. Programmatically interpretable reinforcement learning. In *International Conference on Machine Learning*, 5045–5054. PMLR.
- Yang, Y.; Inala, J. P.; Bastani, O.; Pu, Y.; Solar-Lezama, A.; and Rinard, M. 2021. Program Synthesis Guided Reinforcement Learning. arXiv:2102.11137.
- Zhou, W.; and Li, W. 2018. Safety-aware apprenticeship learning. In *International Conference on Computer Aided Verification*, 662–680. Springer.
- Zhu, H.; Xiong, Z.; Magill, S.; and Jagannathan, S. 2019. An inductive synthesis framework for verifiable reinforcement learning. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation*, 686–701.
- Ziebart, B. D.; Maas, A.; Bagnell, J. A.; and Dey, A. K. 2008. Maximum Entropy Inverse Reinforcement Learning. In *Proceedings of the 23rd National Conference on Artificial Intelligence - Volume 3, AAAI'08*, 1433–1438. AAAI Press. ISBN 978-1-57735-368-3.