

# DIRECT: Learning from Sparse and Shifting Rewards using Discriminative Reward Co-Training

Philipp Altmann  
LMU Munich  
philipp.altmann@ifi.lmu.de

Thomy Phan  
LMU Munich

Fabian Ritz  
LMU Munich

Claudia Linnhoff-Popien  
LMU Munich

Thomas Gabor  
LMU Munich

## ABSTRACT

We propose *discriminative reward co-training* (DIRECT) as an extension to deep reinforcement learning algorithms. Building upon the concept of *self-imitation learning* (SIL), we introduce an imitation buffer to store beneficial trajectories generated by the policy determined by their return. A discriminator network is trained concurrently to the policy to distinguish between trajectories generated by the current policy and beneficial trajectories generated by previous policies. The discriminator’s verdict is used to construct a reward signal for optimizing the policy. By interpolating prior experience, DIRECT is able to act as a surrogate, steering policy optimization towards more valuable regions of the reward landscape thus learning an optimal policy. Our results show that DIRECT outperforms state-of-the-art algorithms in sparse- and shifting-reward environments being able to provide a surrogate reward to the policy and direct the optimization towards valuable areas.

## KEYWORDS

Reinforcement Learning, Curriculum Learning, Exploration

## 1 INTRODUCTION

Imitation plays a central role in the human learning process being deeply rooted in the human brain [21, 37]. Therefore, imitation can be considered essential for the development of human intelligence [14]. Turing even suggested imitation to be a central capability on which to test an artificial intelligence in the *Imitation Game* [50]. Cognitive science suggests that the ability to access memories voluntarily enabled humans to act out past or future events and take control over their output by voluntarily rehearsing and refining, gaining skills like reenactive play and imitation learning [14]. This has also inspired new methods of reinforcement learning (RL), where recent positive experience is used to guide the optimization of an agent’s policy, also referred to as self-imitation learning (SIL) [33]. Similar to human behavior, imitation is incorporated to improve the exploratory capabilities [38]. This is especially important for RL in real-world scenarios where precisely constructed and smooth reward functions are not available [9]. In some instances the reward distribution might even shift between training and execution, a phenomena referred to as distributional shift which is crucial to tackle in order to safely deploy RL models in the real world [2]. The exploration of both sparse- and shifting rewards have been shown to be challenging problems, especially for deep (reinforcement) learning algorithms [49].

In this paper, we address policy optimization in such challenging reward scenarios and propose discriminative reward co-training (DIRECT). DIRECT incorporates a self-imitation buffer to store beneficial trajectories generated by former policies determined by their return. A discriminator, inspired by generative adversarial networks (GANs) [15], is trained to distinguish between trajectories experienced by the current policy and trajectories from this self-imitation buffer. If the discriminator can easily discern between the current policy and its former best deeds, we assume that the current policy does not live up to its full potential and has still room to improve. On the other hand, if the policy’s current actions easily match with its best recorded experience, then the discriminator needs to examine the experience samples more closely. This discriminative instance thus can serve as an additional or even a surrogate reward signal aiding the policy optimization towards self-imitation of previously beneficial experiences. This guidance is especially helpful for sparse reward landscapes and hard exploration problems. Overall, we provide the following contributions:

- We propose discriminative reward co-training (DIRECT) to shape rewards using a self-imitation buffer and the corresponding update rule, a discriminator and the corresponding training loss, and an accommodating training procedure.
- We show the directive capabilities and gained adaptability to environmental changes of DIRECT within varying safety gridworld environments.
- We provide benchmark comparisons against PPO, A2C, DQN and SIL and show that DIRECT achieves superior performance especially in sparse- and shifting reward settings.

## 2 BACKGROUND

*Reinforcement Learning.* We formulate our problem as *Markov decision process (MDP)* [39]  $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R} \rangle$ , where  $\mathcal{S}$  is a (finite) set of states  $s_t \in \mathcal{S}$ ,  $\mathcal{A}$  is the (finite) set of actions  $a_t$ ,  $\mathcal{P}(s_{t+1}|s_t, a_t)$  is the transition probability function, and  $\mathcal{R}(s_t, a_t)$  is the reward. The *reinforcement learning (RL)* goal is to find a policy  $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ , which maximizes the expectation of the return  $G(\tau)$  for all trajectories  $\tau \sim \pi$ .

$$G(\tau) = \sum_{t=0}^{\infty} \gamma^t \cdot \mathcal{R}(s_t, a_t) \quad (1)$$

where  $\gamma \in [0, 1]$  is the discount factor. Furthermore, let the action-value function  $Q^\pi(s_t, a_t)$  and the value function  $V^\pi(s_t)$  of any given action  $a_t$  and state  $s_t$  be the expected return following  $\pi$ , starting with action  $a_t$  in state  $s$  and starting in state  $s_t$  respectively. [46]

*Policy Optimization.* For the following, we consider  $\pi$ ,  $V^\pi$  and  $Q^\pi$  to be deep neural networks parameterized by  $\theta$ . We furthermore consider interactive training data in form of rollouts generated by the current policy acting within a finite-horizon environment, i.e. an environment that terminates episodes upon reaching a goal, termination state, or a maximum of steps, or executed actions. To learn an optimal policy  $\pi^*$  for high-dimensional state spaces, applying Q-learning updates [46] to  $Q^\pi$ , referred to as Deep Q-Network (DQN), has been proposed [30]. Policy gradient methods directly learn the policy parameters  $\theta$ , where the gradient is given by  $\nabla_\theta J(\pi_\theta) = \nabla_\theta \mathbb{E}_{\pi_\theta} [G(\tau)]$  for trajectories  $\tau$  generated by  $\pi_\theta$  and  $\theta$  can be updated via gradient ascent [47]. To enable updates using policy rollouts of incomplete episodes, the incorporation of the value network  $V^\pi$  to build the advantage  $A_t^\pi = Q^\pi(s_t, a_t) - V^\pi(s_t)$  has been proposed and is referred to as Advantageous Actor-Critic (A2C), where  $\pi$  represents the actor  $V^\pi$  represents the critic [29]. Further improving upon this concept, Proximal Policy Optimization (PPO) builds upon Trust Region Policy Optimization (TRPO) [44] introducing a surrogate to optimize the update size in order to yield faster and more stable convergence [45]. The main PPO objective is defined as  $\mathcal{L}_t^{\text{CLIP}}(\theta) = \mathbb{E}_t [\min(r_t A_t^\pi, \text{clip}(r_t, 1 - \epsilon, 1 + \epsilon) A_t^\pi)]$ , where  $r_t$  is the action probability ratio between the current policy  $\pi_\theta$  and the previous policy  $\pi_{\theta_{old}}$ , taken from TRPO, and  $\epsilon$  is a hyperparameter controlling the clipping [45].

*Reward Shaping.* Motivated by the assumption that a reward signal might not be optimally defined for an agent to learn the intended behavior, Ng et al. [31] proposed to extend the given reward by supplemental metrics, like the distance to a given target, or, the estimated remaining time to reach the goal. This can be applied to the extent, where the environmental reward signal is completely replaced by the enhanced metric [31]. However, in contrast to the discriminative reward we propose, the construction of reward shaping metrics often requires domain knowledge.

*Intrinsic Rewards.* Especially in sparse reward setting, e.g. if the agent is only rewarded at the episode end instead of receiving constant incentives, intrinsic rewards have been shown to yield improved exploitative capabilities. Intrinsic rewards refer to a special form of reward shaping, where the environmental reward signal is complemented by an intrinsically motivated reward supplement [4, 35, 43]. In contrast to conventional reward shaping techniques, this intrinsic motivation can mostly be constructed without the need for domain knowledge, e.g. by learning an internal model.

*RL Safety.* Applying RL methods to real-world problems while considering safety requirements within the environment often plays an equally important role as merely maximizing the expected return. Leike et al. [24] therefore developed a suite of RL environments that require an agent to act safely with regard to previously defined safety problems. Current approaches for safe RL resemble the use of a shaped reward, where domain knowledge is incorporated to maximize safety in addition to the environmental return [1]. Distributional shifts represent a subset of said problems and a critical issue within various machine learning applications [40]. They refer to applying a learned model to a novel set of data containing changes not included in the training data. [2, 24, 40]. To produce policies that are robust to distributional shift, extending the training

samples by a more diverse pool of synthetic experience has been to be shown helpful [27]. Within supervised learning, adversarial methods have been shown to yield improved robustness against out-of-distribution samples [25].

*Adversarial Learning.* Similar to the Imitation Game [50], adversarial methods introduce an opponent or adversary, attempting to fool the decision-making instance, typically a classifier, by injecting adversarial samples. Adversarial samples refer to malicious or noisy samples specifically designed to attack a certain model [16, 23]. As many deep learning techniques have been shown to be prone to adversarial samples, not performing as expected, failing to classify them correctly, the adversarial counterpart has been attempted to be integrated into the learning process [7, 19, 26]. In general, adversarial methods have been applied to a wide variety of learning tasks like co-evolution [13, 52], or testing [3, 51] and recently gained further interest in data or image generation via generative adversarial networks (GANs) [15]. The GAN objective is given by the two-player minimax game

$$\min_{\mathcal{G}} \max_{\mathcal{D}} C(\mathcal{D}, \mathcal{G}) = \mathbb{E}_{x \sim p_{data}(x)} [\log \mathcal{D}(x)] + \mathbb{E}_{z \sim p_z(z)} [\log(1 - \mathcal{D}(\mathcal{G}(z)))] \quad (2)$$

where  $p_z(z)$  is a random noise variable used as an input to the Generator to learn its distribution  $p_g$ . Adding labels to the training data, thus turning the procedure into a semi-supervised learning process has also been shown to benefit the learning process and quality of generated samples, as shown with conditional GANs (cGANs) [28]. Applying adversarial architectures has been shown to improve a learned model’s robustness against changing circumstances [25, 48]. In addition to unsupervised tasks, like data generation, and the methodological connections to Actor-Critic methods [36], GANs have also been shown to be closely related to inverse reinforcement learning (IRL) and thus being applicable to those tasks as well [11]. In contrast to RL, aiming to find an optimal policy  $\pi^*$ , IRL aims to find a reward function that best describes a given trajectory of demonstrations and can be used to learn from experts, using imitation learning [32].

### 3 RELATED WORK

Even though many recent approaches tackling the difficulties of learning from a sparse reward signal have been proposed, real-world scenarios exist where a reward signal is not available at all, or the goal is to learn from dedicated expert behavior. For these scenarios, imitation learning approaches have been developed [18, 20, 33, 42].

*Generative Adversarial Imitation Learning (GAIL).* Proposed by Ho et al. [18], GAIL attempts to simplify the two-step IRL procedure [32, 54] to imitate expert demonstrations. They introduce a discriminator  $\mathcal{D}_\phi$ , parameterized by  $\phi$ , similar to the discriminator defined for GANs, but discriminating between actions taken by the expert and actions taken by the policy  $\pi_\theta$  to be optimized. From a GAN’s perspective, the learned policy can be seen as the generator. The GAIL objective is to find a saddle point of

$$\mathcal{L}^{\text{GAIL}}(\theta, \phi) = \mathbb{E}_{\pi_\theta} [\log \mathcal{D}_\phi(s_t, a_t)] + \mathbb{E}_{\pi_E} [\log(1 - \mathcal{D}_\phi(s_t, a_t))] - \lambda \mathcal{H}(\pi_\theta),$$

where  $\mathcal{H}$  is the current policy’s entropy, as already used for IRL,  $\pi_\theta$  is the learned policy and  $\pi_E$  is the expert policy. In practice, the expert policy is often represented by a dataset consisting of trajectories  $\tau$ . The authors suggest alternating gradient steps using Adam on  $\phi$  to increase Eq. 3 w.r.t.  $\mathcal{D}$  and TRPO [44] on  $\theta$  to decrease Eq. 3 w.r.t.  $\pi$ . [18]. To further reduce sample-complexity, [22] proposed to incorporate the discriminator off-policy. Furthermore, the incorporation of a discriminator network has been shown to provide a reward surrogate to be recovered using IRL, that is robust to variations in the underlying domain [12]. DIRECT builds upon this potential and uses the discriminator as a surrogate reward signal for improved policy optimization of challenging environmental reward signals (like sparse or shifting) in safety requiring environments. As with humans, imitation is not only applicable to learn the behavior from an expert, but also to improve one’s own behavior by reaffirming rewarding experiences made in the past and trying to achieve those again. However, besides the shared inspiration drawn from GANs, GAIL is dependent on (expert) demonstrations in order to learn a policy, while DIRECT is able to employ the concept of self-imitation on former beneficial trajectories in order to learn a policy, that maximizes the return of a given environmental reward function.

*Self-Imitation Learning (SIL)*. has been developed as an extension of A2C, adding a buffer  $\mathcal{B}$  that stores trajectories  $\tau$  previously encountered by the policy  $\pi_\theta$ . Thereby, the policy is additionally optimized towards beneficial recent behavior, i.e. trajectories with positive advantage values, or, where the actual return is greater than the estimated value. This extension has been shown to yield a deeper exploration of the reward landscape than comparable approaches by exploiting past experiences [33]. Similarly to experience replay [53], Oh et al. [33] propose to extend the A2C framework by a buffer  $\mathcal{B}$  storing the policy’s previous experiences to be revisited. After each rollout, the buffer is extended by the recently encountered states, actions and cumulative rewards  $\mathcal{B} = \{(s_t, a_t, G_t)\}$ . Policy and value estimate are updated sequentially towards both the RL objective by applying the on-policy A2C losses and past beneficial trajectories stored in the experience buffer by applying the following adapted A2C losses:

$$\begin{aligned} \mathcal{L}_{policy}^{SIL} &= -\log \pi_\theta(a_t | s_t) (G_t - V_\theta(s_t))_+ \\ \mathcal{L}_{value}^{SIL} &= \frac{1}{2} \|(G_t - V_\theta(s_t))_+\|^2 \end{aligned} \quad (3)$$

where  $(\cdot)_+ = \max(\cdot, 0)$ , is constraining updates of both the policy and the value estimate to be only made for beneficial trajectories (i.e., where the actual return is greater than the estimated value). Despite being designed for A2C, SIL also applies to PPO [33]. Furthermore, the concepts of SIL have been combined with advantage learning paradigms and shown applicable to off-policy methods [10]. Overall, both DIRECT and SIL use beneficial previous experience to improve exploration. However, DIRECT constructs a single surrogate reward function that can be used for policy optimization. In contrast, SIL incorporates those previous experiences in addition to current experience. Thus, the policy is trained both with new and good previous samples, but there is no intermediate reward to incentivize the exploration of a sparse reward signal.

*Generative Adversarial Self-Imitation Learning (GASIL)*. The combination of GAIL and SIL, has been previously proposed with GASIL to regularize the environmental reward signal and evaluated across various control tasks [34]. Also, this slightly adapted reward function has been shown to be applicable to independent learning in cooperative multiagent systems [17]. However, even though sharing the inspiration, DIRECT presents a more elaborate combination and integration of a discriminator into the policy optimization. Therefore, we also propose additional hyperparameters to tune the discriminative reward signal, that enable training with a purely surrogate discriminative reward and no direct environmental feedback, which fundamentally separates our work from GASIL. This more advanced architecture also yields the opportunity for evaluating DIRECT more precisely. In addition to a proof-of-concept, we are therefore able to show the directive capabilities of the discriminator as a surrogate reward function in safety critical environments with sparse reward landscapes. Furthermore, we consider the impact of DIRECT to the robustness and adaptability of the resulting policies with regard to distributional shifts in the environment.

#### 4 DISCRIMINATIVE REWARD CO-TRAINING

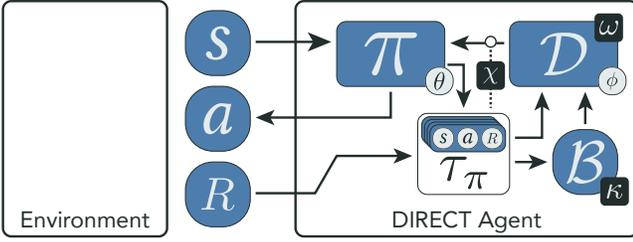
We propose an architecture called **Discriminative Reward Co-Training (DIRECT)**, as an extension of GAIL applied to self-imitation. Figure 1 illustrates the components of the architecture and their connection, which are formally introduced in the following section. The overall DIRECT objective is defined as

$$\operatorname{argmin}_\theta \operatorname{argmax}_\phi \mathcal{L}^{\text{DIRECT}}(\theta, \phi), \text{ with} \quad (4)$$

$$\begin{aligned} \mathcal{L}^{\text{DIRECT}}(\theta, \phi) &= \mathbb{E}_{\tau_E \sim \mathcal{B}} \left[ \log \left( 1 - \mathcal{D}_\phi(s, a, G) \right) \right] + \\ &+ \mathbb{E}_{\tau_\pi} \left[ \log \mathcal{D}_\phi(s, a, G) \right] - \lambda \mathcal{H}(\pi_\theta) \end{aligned} \quad (5)$$

*Self-Imitation Buffer  $\mathcal{B}$* . As suggested for self-imitation learning, we use a buffer to store beneficial past trajectories  $\tau_\pi^*$  that the agent should reiterate. Beneficial state-action transitions are determined by their discounted return they achieved when executed in the environment. As illustrated in Figure 1, the self-imitation buffer  $\mathcal{B}$  is updated continuously while the policy  $\pi_\theta$  is optimized and thus generates new experience  $\tau_\pi$ . Thus, the trajectories in the buffer  $\mathcal{B}$  are enabled to evolve, pushing the policy  $\pi_\theta$  towards promising areas. To survey this evolution of  $\mathcal{B}$  and ensure its improvement towards optimal behavior, we introduce the *buffer reward* as an additional reward, reflecting the mean reward of the trajectories  $\tau_{\mathcal{B}}$  currently stored in the self-imitation buffer  $\mathcal{B}$ . Also, we record the number of successful buffer updates, where old experience is replaced by improved and more recent behavior, we refer to as *buffer momentum*.

*Discriminator  $\mathcal{D}$* . Inspired by GAIL and GANs, the discriminator  $\mathcal{D}$  is trained to distinguish between samples generated by the current policy  $\tau_\pi$  and buffer samples  $\tau_E$ . Applying these concepts to SIL, we do not require any expert behavior to be imitated. Thus, expert samples are replaced by samples drawn from the self-imitation buffer  $\tau_{\mathcal{B}}$ , as specified in Equation 5, the discrimination in DIRECT happens between former good experiences  $\tau_{\mathcal{B}}$  and experiences by the current policy  $\tau_\pi$ . As for GAIL, the discriminator’s target is



**Figure 1: DIRECT Architecture** Following policy  $\pi$ , the agent executes actions  $a$  in the environment as a result of observing its state  $s$ , receiving a reward  $r$  for the actions’ consequences. Thereby, the policy generates trajectories  $\tau_\pi$  containing states, actions and rewards. After every rollout, the return  $G$  for each trajectory is calculated, and the self-imitation buffer  $\mathcal{B}$  is updated. Storing the best  $\kappa$  trajectories  $\tau_\pi$ , the discriminator  $\mathcal{D}$ , parameterized by  $\phi$ , is then updated to differentiate between trajectories sampled from the buffer  $\tau_{\mathcal{B}}$  and trajectories sampled from the current policy  $\tau_\pi^*$ . The update frequency of the discriminator in relation to the policy is given by  $\omega$ . Ideally, the trajectories in  $\mathcal{B}$  converge towards expert-like behavior during the training process. Therefore, the discriminator output can be used as a reward signal to update the policy parameters  $\theta$ , optimizing it towards good or even optimal past experiences. Additionally, the hyperparameter  $\chi$  can be used to tune the mixture between the environmental reward and the discriminative reward.

to label beneficial (expert) samples as “1” and samples generated by the current policy  $\pi_\theta$  to be optimized as “0”. Therefore, beneficial behavior is rewarded by the discriminator  $\mathcal{D}$ . As the samples in the self-imitation buffer  $\mathcal{B}$  are supposed to converge towards optimal behavior throughout the training process, the discriminator  $\mathcal{D}$  should be able to guide the policy optimization towards an optimal solution. In contrast to GAIL which imitates an expert, we attempt to maximize an environmental reward signal using self-imitation. We therefore propose adding the accumulated environmental return, already stored in the self-imitation buffer  $\mathcal{B}$ , as an input to the discriminator  $\mathcal{D}$ . From a GAN’s perspective, this addition is comparable to conditional GANs, adding labels to the training data [28]. In contrast to GANs, however, we do not wish to have a fast converging discriminator  $\mathcal{D}$ , as the self-imitation buffer  $\mathcal{B}$ , where target samples are drawn from, gradually evolves (hopefully towards optimal behavior) during the training process. Thus, a discriminator that converged prematurely, i.e., towards discriminating suboptimal samples as expert-like behavior, is not intended. Therefore, balancing the inner and outer optimization loops plays a more prominent role for training DIRECT than it does for GANs or GAIL.

*Policy Optimization.* Overall, DIRECT can be seen as an extension of the components described above, applicable to any model-free RL algorithm, providing an intermediate instance for processing the reward signal passed to the policy optimization instance (cf. Figure 1). Similar to reward shaping approaches, we expect this extension to benefit the optimization by providing a smoother

reward landscape that is easier to exploit by the agent. This guidance is especially helpful for sparse reward landscapes and hard exploration problems. For optimizing the policy, we chose to use Proximal Policy Optimization (PPO) [45], an approach shown to be successfully applicable to a wide variety of RL problems. The reward signal used by PPO to calculate the return after every policy rollout as is calculated as follows:

$$r_t^*(s_t, a_t) = \chi \mathcal{D}(s_t, a_t, G_t) + (1 - \chi) r_{t+1}, \quad (6)$$

where  $r_{t+1}$  is the real reward (*environmental reward*) received by the agent from the environment for executing action  $a_t$  and  $\chi$  is a hyperparameter determining the reward mixture, such that for  $\chi = 0$ , no discriminative reward is used and the algorithm exactly resembles PPO. With  $\chi = 1$ , PPO only receives a reward originating from the discriminator without any access to the real environmental reward signal.

*Training procedure.* To elucidate the proposed components, previously introduced and illustrated in Figure 1, Algorithm 1 shows the training procedure of the proposed architecture for Discriminative Reward Co-Training.

---

#### Algorithm 1 Discriminative Reward Co-Training

---

```

Initialize policy parameters  $\theta$ 
Initialize discriminator parameters  $\phi$ 
Initialize self-imitation buffer  $\mathcal{B} \leftarrow \emptyset$ 
for each iteration do
  Sample policy trajectories  $\tau_\pi$  by performing a PPO rollout
   $\tau_\pi \leftarrow [(s_0, a_0, R_1), (s_1, a_1, R_2), \dots, (s_t, a_t, R_{t+1})]$ 
  Calculate partial returns  $G(\tau)$  following Eq. 1
  Update self-imitation buffer  $\mathcal{B}$ , prioritizing samples via  $G(\tau)$ :
   $\mathcal{B} \leftarrow [(s_0, a_0, G(\tau)), \dots, (s_t, a_t, G(\tau))]$ 
  Sample beneficial trajectories  $\tau_{\mathcal{B}}$  from  $\mathcal{B}$ 
  Update the discriminator parameters  $\phi$ , to maximize Eq. 5
  Calculate discriminative reward  $r_t^*$  following Eq. 6
  Update policy parameters  $\theta$  to minimize Eq. 5
end for

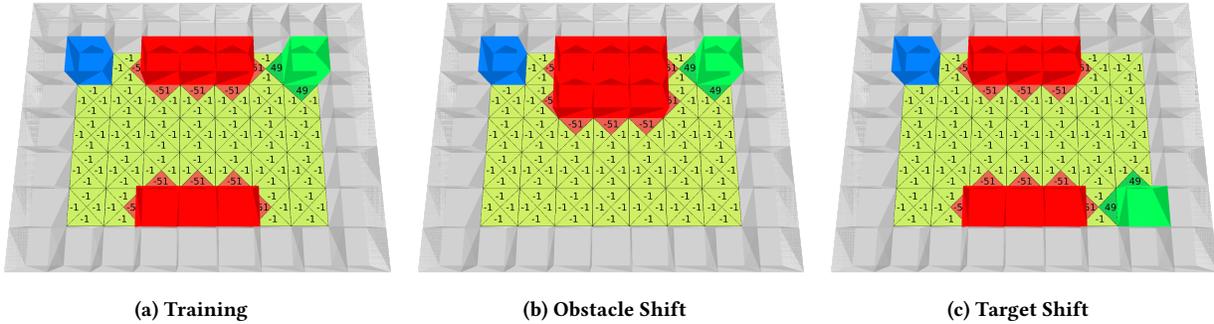
```

---

## 5 EXPERIMENTAL SETUP

*Hyperparameters.* Unless denoted otherwise we use DIRECT with the following hyperparameters that have shown the best results in preliminary experiments:

- $\kappa = 2048$ : a self-imitation buffer  $\mathcal{B}$  size, that resembles the rollout horizon of the used PPO implementation,
- $\omega = 1/2$ : in sparse reward settings resulting in concurrently training  $\mathcal{D}_\phi$  and the  $\pi_\theta$ , as the discriminator  $\mathcal{D}$  is updated with the samples from both the policy  $\tau_\pi$  and the buffer  $\tau_{\mathcal{B}}$ ,
- $\omega = 2/1$ : in dense reward settings, due to the well-defined reward signal a higher update rate  $\omega$  allows for faster adaption of the discriminator  $\mathcal{D}$ , thus faster overall training,
- $\chi = 0.5$  for dense reward settings, to showcase the ability of DIRECT to create an homogeneous reward signal, and
- $\chi = 1.0$  for sparse reward setting, to assess the prospects of a purely discriminative reward signal.



**Figure 2: Evaluation Environment Variations with dense reward landscape adapted from [24]** comprising an  $(7, 9)$  discrete observation space with cells  $\in \{Wall, Field, Lava, Target, Agent\}$  and the action space  $\mathcal{A} = \{Up, Right, Down, Left\}$ . The goal of the agent (blue) is to reach the target (green) rewarded without touching the lava (red). Both contact to lava and target fields terminate the episode with a reward of +50 and -50 respectively. To incentivize shorter paths each step is rewarded -1 which is propagated to the agent directly in the dense reward setting and upon episode termination in the sparse reward setting.

To benchmark DIRECT we use the PPO, A2C and DQN implementations by [41] with default PPO and DQN hyperparameters and A2C hyperparameters as suggested by [24], as well as a SIL implementation according to [33]. All implementations can be found here <sup>1</sup>.

*Domain.* We chose to evaluate the DIRECT using environments from the AI Safety Gridworlds [24] adapted to OpenAI gym [5]. The environments used are shown in Figure 2. Training a policy in the training environment (cf. Figure 2a) and evaluating the policy within the Obstacle- and Target Shift environments (cf. Figure 2b and Figure 2c) reveals the approaches’ robustness against distributional shift, which is considered one central problem to be solved for safe RL [2, 24]. Beside this concisely crafted challenges, we chose to evaluate DIRECT using the safety gridworlds, as they allow for better interpretability and visualizations of resulting policies and the course of training. To assess the ability of DIRECT to provide a surrogate direction to the policy optimization, we furthermore added a sparse reward setting, where the agent receives zero rewards throughout the episode and a final non-zero reward at terminal states. Besides testing an agents explorative capabilities and robustness against changes in the environment, the given environments are also suitable for evaluating an agent’s *safe exploration* capabilities, with the lava fields representing a safety risk to be avoided, by analyzing the cause of termination. Agents that cause more failure terminations by running into lava fields, either during training or within unseen environments are arguably less safe compared to agents terminating their episodes as a result of reaching the maximum number of steps [2]. Terminating in the target state using the optimal path yields respective returns of 42, 40 and 40. In any case, reaching the target state yields a return of at least -52.

*Metrics.* As a performance metric for a policy’s behavior, we use the average environmental return  $G(\tau)$  over 100 episodes (mean return). In order to reduce the necessary compute, this metric is also used to terminate the training upon reaching a threshold of 95% of the optimal solution, i.e. a return of 40 in environment (2a), and

a return of 38 in environments (2b) and (2c). To assess the resulting policies, we run single episode, deterministic evaluations and record the mean return, as well as the episode termination reason of the trained policies. All of the following results are averaged over eight runs with randomly selected seeds, trained in parallel in four identical environments for a maximum of 1 million timesteps.

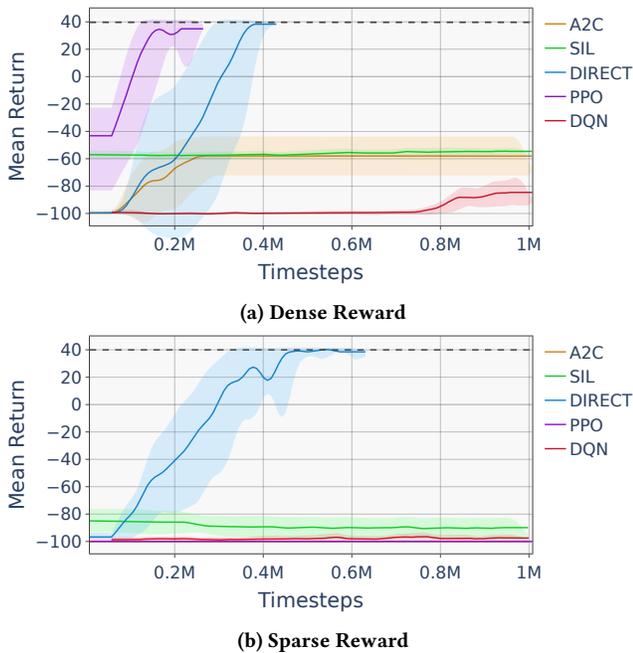
## 6 TRAINING

According to Figure 3a, both PPO and DIRECT are able to exploit the dense reward signal and find optimal policies within 200k and 350k timesteps on average respectively. The delay of convergence between PPO and DIRECT is most likely attributed to the double optimization loop architecture, where first, the policy needs to generate valuable experience to fill the buffer, before the discriminator can be conditioned to produce a valuable reward signal. Even though the reward signal is well-defined and dense, both A2C and SIL get stuck in a local optimum around a return of -60. Policies within this reward range are not able to reach the target, but cause an early episode termination by running into a near lava field. Thus, both SIL and A2C are not able to produce any optimal policies, which might be attributed to their short rollout horizon. DQN shows even worse performance, only reaching average return areas of -90 in the late optimization. The progress suggests that DQN might need a higher number of training steps in order to learn an optimal policy.

Figure 4 shows the average discriminative returns after training for all possible state-action combinations as in form of a heatmap. In contrast to the environmental reward structure (c.f. Figure 2a), that only rewards reaching the goal while penalizing steps and contact to lava, DIRECT is able to learn an additional intrinsic reward signal  $D$  that provides further incentive towards the target in the top right while still avoiding lava fields. More specifically, starting from the top left and following a policy that selects the action with the highest immediate discriminative reward, the goal is reached via the shortest possible path.

Overall, DIRECT shows competitive performance in comparison to the selected baselines. However, the real advantage of DIRECT really becomes visible looking at the training process of the sparse

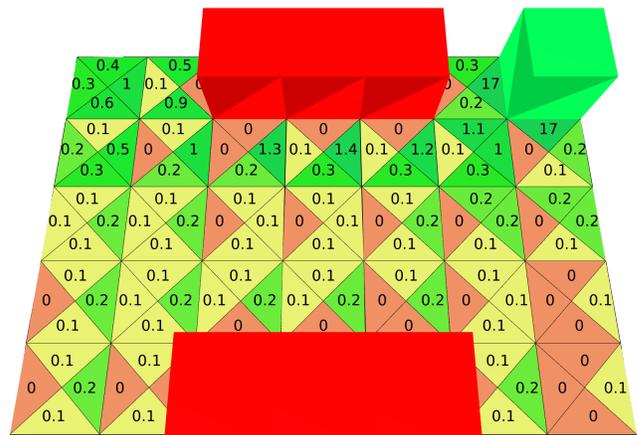
<sup>1</sup><https://github.com/philippaltnmann/DIRECT>



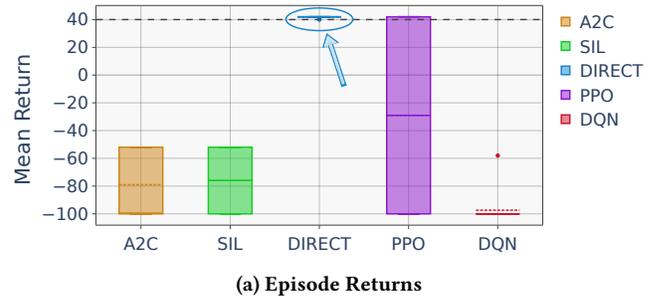
**Figure 3: Training Progress within the dense (a) and sparse (b) reward environments:** the optimization progress for A2C (orange), SIL (green), DIRECT (blue), PPO (purple), and DQN (red) averaged over eight runs, with the number of steps taken in the environment on the x-axis and the mean return on the y-axis. The shaded areas mark the 95% confidence intervals, the reward threshold of 40 is displayed by the dashed line.

environment in Figure 3b. While all other algorithms do not even converge to the locally optimal behavior of terminating the episode by running into the closest lava field, but produce policies that yield average returns of about -100, **all** Discriminative Reward Co-Trained policies converge to a near optimal behavior within under 650k timesteps of training. We attribute this superiority to the directive abilities of the discriminative reward signal fostering improved exploration capabilities of valuable experience. By employing the concept of self-imitation DIRECT improves the optimization’s sample efficiency, i.e. requiring to experience fewer good episodes in order to exploit optimal behavior.

The overall training performance of all evaluated approaches in both sparse and dense reward settings is summarized in Figure 5. With all returns above the reward threshold, i.e. reaching the target via the optimal path or at most two more steps, Discriminative Reward Co-Trained policies outperform all all compared approaches. PPO-trained policies perform second-best, with half of the policies reaching into reward regions, where the target is reached. The worse performance of the remaining policies is attributed to PPO’s bad ability to learn from sparse reward signals. A2C-, SIL- and DQN-trained policies did not converge to a target-finding behavior at all. This superiority of DIRECT policies is also reflected in their safety where all (100%) of the evaluation episodes terminate in the target state and, in contrast to all compared benchmarks, no unsafe behavior, i.e. failure terminations, is encountered.



**Figure 4: DIRECT Reward Heatmap:** Discriminative reward (cf. Equation 6) on average over all runs after training in the Dense Reward Environment.



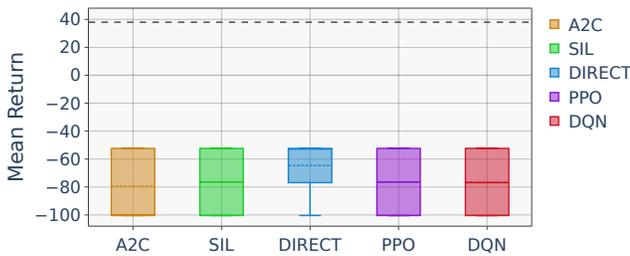
**(b) Episode Termination:** Percentage of episodes in which the agent reaches the target (green), hits an obstacle (red) or runs into a timeout (yellow).

**Figure 5: Training Evaluation:** Episode Returns (a) and termination reasons (b) of policies trained with A2C, SIL, DIRECT, PPO, and DQN during deterministic evaluation in both the sparse and dense reward settings of the training environments (cf. Figure 2a)

## 7 ADAPTATION

In the following section we will evaluate the approaches’ robustness to distributional shift. Therefore, we analyze their performance environments where, in comparison to the training environment (cf. Figure 2a), either the obstacles (cf. Figure 2b), or the position of the target (cf. Figure 2c) has shifted. Again, with real-world applications in mind, we will consider both dense- and sparse reward settings. Due to its directing modifications to the reward signal, we expect DIRECT to yield policies with improved stability and adaptability even to shifting reward signals.

*Shifted Obstacles.* In this evaluation scenario, the target position is the same as during training. However, the path that was shortest during training now is blocked by lava. Thus, the agent needs to adapt its behavior to reach the target successfully. To assess the policies' robustness to a distributional shift of its surrounding environment, we evaluated all previously trained policies deterministically in the shifted obstacle environment (cf. Figure 2b). The according results are shown in Figure 6. All approaches achieve mean returns between -100 and -50. This again highlights the susceptibility of RL algorithms to distributional shifts. Overall, no trained policy is able to find a way around the re-positioned obstacles ad-hoc. Therefore, we decided to perform a refinement-training on the previously trained policies to assess the underlying algorithms' adaptability to changes within the environment.

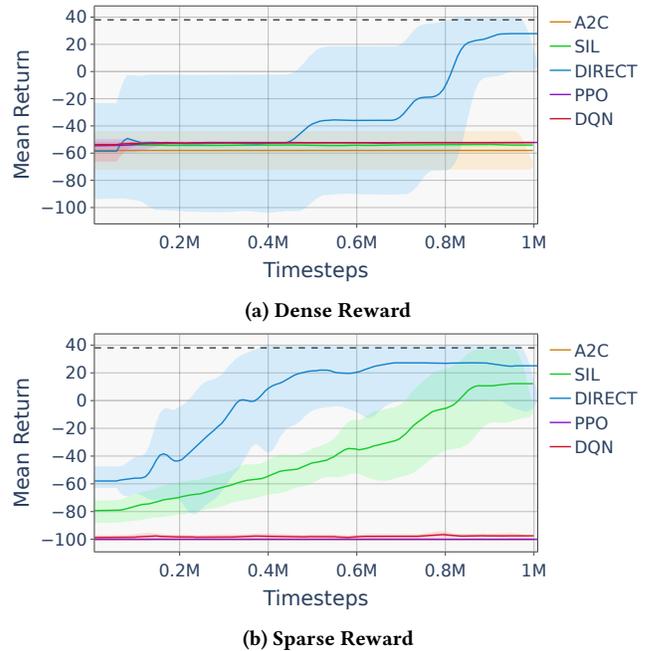


**Figure 6: Shifted Obstacle Evaluation:** Episode Returns of policies trained with A2C, SIL, DIRECT, PPO, and DQN during deterministic evaluation in both the sparse and dense reward settings of the shifted obstacle environments (cf. Figure 2b)

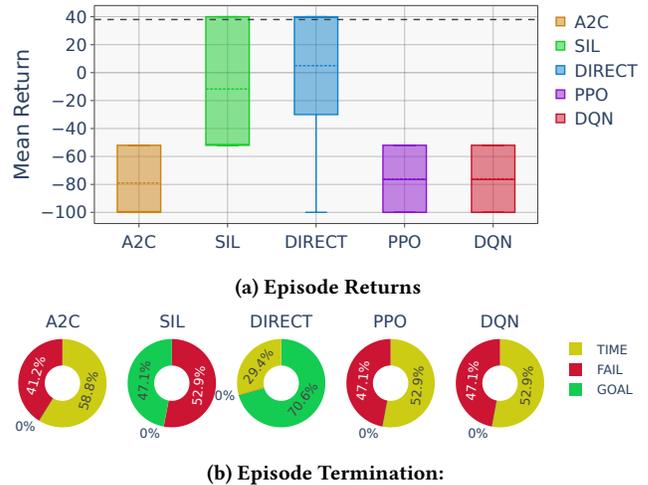
Trying to adapt to the dense reward signals (cf. Figure 7a), all algorithms converge to sub-optimal returns of around -50 within the first 400k timesteps. However, in contrast to all other approaches, DIRECT is able to explore valuable regions of the reward landscape and exploit them via self-imitation, converging to a mean return of around 30 within the 1M timesteps. The mean return indicates that DIRECT is not able to find an optimal policy during the refinement training. However, in contrast to all other approaches, DIRECT is able to solve the scenario by reliably finding the target within 100 timesteps and without touching the lava.

Switching to the sparse reward setting (cf. Figure 7b), the performance of most approaches drops to -100, where the episodes most likely end with a timeout after 100 steps. Besides DIRECT, SIL shows surprisingly good adaptability to the shifted obstacle, converging to an average return of around ten throughout the training process. This is especially surprising as the originating policies were never able to adequately solve the training environment. However, DIRECT adapted policies again show even faster convergence towards high reward regions. Interestingly, despite the increased difficulty due to the sparse reward signal, DIRECT shows even better performance than within the less challenging dense reward setting.

Overall, evaluating the adapted policies deterministically (cf. Figure 8), only approaches incorporating self-imitation are able to adapt to the environmental shift. In concurrence to previous work, this finding demonstrates the superiority of self-imitation based



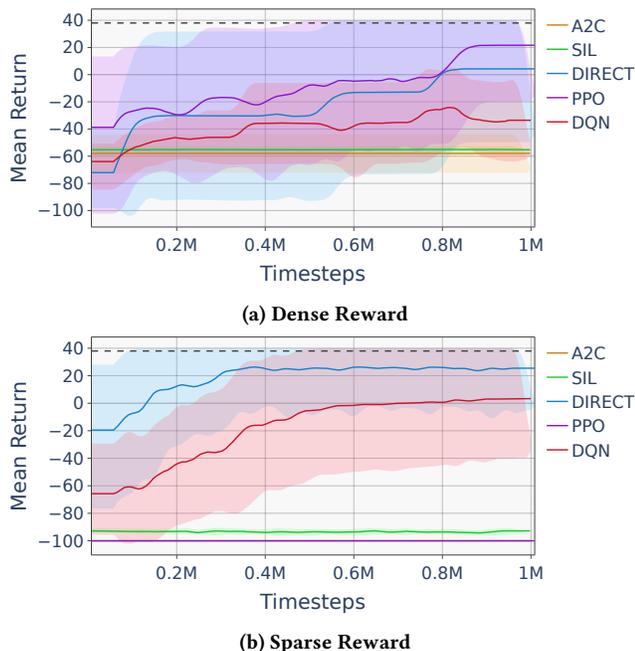
**Figure 7: Adaption Progress Shifted Obstacle:** the adaption progress to dense (a) and sparse (b) reward settings averaged over eight runs, with the number of additional timesteps taken in the shifted environment on the x-axis and the mean return on the y-axis. The shaded areas mark the 95% confidence intervals, the reward threshold of 38 is displayed by the dashed line.



**Figure 8: Adaption Evaluation:** Episode Returns (a) and termination reasons (b) of policies trained with A2C, SIL, DIRECT, PPO, and DQN during deterministic evaluation in both the sparse and dense reward settings of the obstacle shift environments (cf. Figure 2b).

approaches in sample efficient exploitation of valuable experience [33]. Again, in comparison, DIRECT policies yield superior safety avoiding lava fields overall. All remaining approaches seem to suffer from base-policies that are already set to the prior training environment, and therefore are less adaptable to the changes within the environment.

*Shifted target.* Complementary to the environment inspected above, the shifted target environment (c.f. Figure 2c) contains a different target position compared to the original training environment but the obstacle positions remain unaltered. Thus, the agent also needs to adapt its behavior. Again, we evaluated the policies trained in section 6 within the shifted environment. As the results resemble the findings of Figure 6, this plot is omitted.



**Figure 9: Adaption Progress to the dense (a) and sparse (b) shifted target (cf. Figure 2c) environments:** the optimization progress averaged over eight runs, with the number of timesteps taken in the environment on the x-axis and the mean return on the y-axis. The shaded areas mark the 95% confidence intervals, the reward threshold of 40 is displayed by the dashed line.

Looking at the adaptation progress within the dense reward setting in Figure 9a, similar to the training, both PPO and DIRECT reach areas of the return, where the target is reached. Again, presumably due to the delayed learning caused by the discriminator optimization, DIRECT converges slightly slower. Furthermore, DQN shows convergence toward target-reaching return regions, proving above hypothesis that its insufficient convergence within the training environment can be overcome by further training.

This is even strengthened within the sparse reward setting shown in Figure 9b, where DQN shows second-best performance. Again, the performance of stuck approaches drops to a mean return of -100, while DIRECT is able to provide guidance towards the shifted target

and yields the fastest convergence and best performing policies. Again, as seen in Figure 10, DIRECT policies excel the adaptation performance of all compared approaches in deterministic evaluation in the shifted target environment.

Overall DIRECT trained policies showed the best adaptability to both shifted targets and shifted obstacles. They exceed state-of-the-art approaches especially in sparse reward setting. However, in contrast to the initial training, none of the approaches yields an optimal policy.



**Figure 10: Adaption Evaluation:** Episode Returns of policies trained with A2C, SIL, DIRECT, PPO, and DQN during deterministic evaluation in both the sparse and dense reward settings of the shifted target environments (cf. Figure 2c)

## 8 CONCLUSION

In this paper we proposed discriminative reward co-training (DIRECT) to provide intermediate feedback in challenging sparse- and shifting reward scenarios. We first introduced the DIRECT architecture including a self-imitation buffer and the corresponding update rule, a discriminator and the corresponding training loss, and an accommodating training procedure. Our experimental results showed the directive capabilities and gained adaptability to environmental changes of DIRECT within varying safety gridworld environments. The provided benchmark comparisons against PPO, A2C, DQN and SIL showed that DIRECT achieves superior performance especially in sparse- and shifting reward settings.

Even though outperforming state-of-the-art approaches in such challenging scenarios, DIRECT’s advanced architecture, in some instances, limits fast convergence. Therefore, further research should consist of incorporating the discriminative reward more directly. Also, techniques to ensure proper convergence and reduce the risk of vanishing rewards to increase the reliability of DIRECT should be considered. Therefore, mechanisms to steer the momentum of the self-imitation buffer would be helpful. Also, an increased diversity of samples within the buffer could improve the overall performance. Employing a diversity measure could also incentivize self-conscious explorations of the environment without the need for a reward function at all [8]. Overall, building upon DIRECT, a broader machine learning framework based on imitation could enable generalized real-world applicability, following multiple psychological and philosophical movements that suggest that imitation is the basis for human intelligence [6].

## ACKNOWLEDGMENTS

This work was partially funded by the Bavarian Ministry for Economic Affairs, Regional Development and Energy as part of a project to support the thematic development of the Institute for Cognitive Systems and is part of the Munich Quantum Valley, which is supported by the Bavarian state government with funds from the Hightech Agenda Bayern Plus.

## REFERENCES

- [1] Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. 2017. Constrained policy optimization. In *International conference on machine learning*. PMLR, 22–31.
- [2] Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. 2016. Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565* (2016).
- [3] Tao Bai, Jinqi Luo, Jun Zhao, Bihan Wen, and Qian Wang. 2021. Recent Advances in Adversarial Training for Adversarial Robustness. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, Zhi-Hua Zhou (Ed.). International Joint Conferences on Artificial Intelligence Organization, 4312–4321. <https://doi.org/10.24963/ijcai.2021/591> Survey Track.
- [4] Marc Bellemare, Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Remi Munos. 2016. Unifying count-based exploration and intrinsic motivation. *Advances in neural information processing systems* 29 (2016).
- [5] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. 2016. OpenAI Gym. *arXiv:1606.01540* [cs.LG]
- [6] Daniel C Dennett. 2017. *From bacteria to Bach and back: The evolution of minds*. WW Norton & Company.
- [7] Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. 2018. Boosting adversarial attacks with momentum. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 9185–9193.
- [8] Benjamin Eysenbach, Abhishek Gupta, Julian Ibarz, and Sergey Levine. 2018. Diversity is All You Need: Learning Skills without a Reward Function. In *International Conference on Learning Representations*.
- [9] Benjamin Eysenbach and Sergey Levine. 2021. Maximum Entropy RL (Probably) Solves Some Robust RL Problems. In *International Conference on Learning Representations*.
- [10] Johan Ferret, Olivier Pietquin, and Matthieu Geist. 2021. Self-Imitation Advantage Learning. In *AAMAS 2021-20th International Conference on Autonomous Agents and Multiagent Systems*.
- [11] Chelsea Finn, Paul Christiano, Pieter Abbeel, and Sergey Levine. 2016. A Connection between Generative Adversarial Networks, Inverse Reinforcement Learning, and Energy-Based Models. *arXiv e-prints arXiv:1611.03852*, Article arXiv:1611.03852 (Nov. 2016). *arXiv:1611.03852* [cs.LG]
- [12] Justin Fu, Katie Luo, and Sergey Levine. 2018. Learning Robust Rewards with Adversarial Inverse Reinforcement Learning. In *International Conference on Learning Representations*.
- [13] Thomas Gabor, Andreas Sedlmeier, Marie Kiermeier, Thomy Phan, Marcel Henrich, Monika Pichlmair, Bernhard Kempter, Cornel Klein, Horst Sauer, Reiner SchmidSiemens AG, et al. 2019. Scenario co-evolution for reinforcement learning on a grid world smart factory domain. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 898–906.
- [14] Liane Gabora and Anne Russon. 2011. The Evolution of Intelligence. *The Cambridge Handbook of Intelligence* (2011), 328–350.
- [15] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2020. Generative adversarial networks. *Commun. ACM* 63, 11 (2020), 139–144.
- [16] Chuan Guo, Jacob Gardner, Yurong You, Andrew Gordon Wilson, and Kilian Weinberger. 2019. Simple black-box adversarial attacks. In *International Conference on Machine Learning*. PMLR, 2484–2493.
- [17] Xiaotian Hao, Weixun Wang, Jianye Hao, and Yaodong Yang. 2019. Independent Generative Adversarial Self-Imitation Learning in Cooperative Multiagent Systems. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems* (Montreal QC, Canada) (AAMAS ’19). International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1315–1323.
- [18] Jonathan Ho and Stefano Ermon. 2016. Generative Adversarial Imitation Learning. In *Advances in Neural Information Processing Systems* 29, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Eds.). Curran Associates, Inc., 4565–4573. <http://papers.nips.cc/paper/6391-generative-adversarial-imitation-learning.pdf>
- [19] Sandy H. Huang, Nicolas Papernot, Ian J. Goodfellow, Yan Duan, and Pieter Abbeel. 2017. Adversarial Attacks on Neural Network Policies. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Workshop Track Proceedings*. OpenReview.net. <https://openreview.net/forum?id=ryvIRyBKI>
- [20] Ahmed Hussein, Mohamed Medhat Gaber, Eyad Elyan, and Chrisina Jayne. 2017. Imitation learning: A survey of learning methods. *ACM Computing Surveys (CSUR)* 50, 2 (2017), 1–35.
- [21] Marco Iacoboni, Roger P Woods, Marcel Brass, Harold Bekkering, John C Mazziotta, and Giacomo Rizzolatti. 1999. Cortical mechanisms of human imitation. *Science* 286, 5449 (1999), 2526–2528.
- [22] Ilya Kostrikov, Kumar Krishna Agrawal, Debidatta Dwibedi, Sergey Levine, and Jonathan Tompson. 2019. Discriminator-Actor-Critic: Addressing Sample Inefficiency and Reward Bias in Adversarial Imitation Learning. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=Hk4fpoA5Km>
- [23] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. 2016. Adversarial examples in the physical world. *arXiv preprint arXiv:1607.02533* (2016).
- [24] Jan Leike, Miljan Martic, Victoria Krakovna, Pedro A Ortega, Tom Everitt, Andrew Lefrancq, Laurent Orseau, and Shane Legg. 2017. AI safety gridworlds. *arXiv preprint arXiv:1711.09883* (2017).
- [25] Jiashuo Liu, Zheyang Shen, Peng Cui, Linjun Zhou, Kun Kuang, Bo Li, and Yishi Lin. 2021. Stable adversarial learning under distributional shifts. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 8662–8670.
- [26] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. 2018. Towards Deep Learning Models Resistant to Adversarial Attacks. In *International Conference on Learning Representations*.
- [27] Russell Mendonca, Xinyang Geng, Chelsea Finn, and Sergey Levine. 2020. Meta-reinforcement learning robust to distributional shift via model identification and experience relabeling. *arXiv preprint arXiv:2006.07178* (2020).
- [28] Mehdi Mirza and Simon Osindero. 2014. Conditional Generative Adversarial Nets. *arXiv e-prints arXiv:1411.1784* (Nov. 2014). *arXiv:1411.1784* [cs.LG]
- [29] Volodymyr Mnih, Adrià Puigdomènech Badia, Mehdi Mirza, Alex Graves, Tim Harley, Timothy P. Lillicrap, David Silver, and Koray Kavukcuoglu. 2016. Asynchronous Methods for Deep Reinforcement Learning. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48* (New York, NY, USA) (ICML ’16). JMLR.org, 1928–1937.
- [30] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602* (2013).
- [31] Andrew Y Ng, Daishi Harada, and Stuart Russell. 1999. Policy invariance under reward transformations: Theory and application to reward shaping. In *ICML*, Vol. 99. 278–287.
- [32] Andrew Y. Ng and Stuart J. Russell. 2000. Algorithms for Inverse Reinforcement Learning. In *Proceedings of the Seventeenth International Conference on Machine Learning (ICML ’00)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 663–670.
- [33] Junhyuk Oh, Yijie Guo, Satinder Singh, and Honglak Lee. 2018. Self-Imitation Learning. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer Dy and Andreas Krause (Eds.). PMLR, Stockholm, Sweden, 3878–3887. <http://proceedings.mlr.press/v80/oh18b.html>
- [34] Junhyuk Oh, Yijie Guo, Satinder Singh, and Honglak Lee. 2019. Generative Adversarial Self-Imitation Learning. <https://openreview.net/forum?id=HJeABnCcKQ>
- [35] Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. 2017. Curiosity-driven exploration by self-supervised prediction. In *International conference on machine learning*. PMLR, 2778–2787.
- [36] David Pfau and Oriol Vinyals. 2016. Connecting generative adversarial networks and actor-critic methods. *arXiv e-prints arXiv:1610.01945* (Oct. 2016).
- [37] Jean Piaget. 1951. *Play, dreams and imitation in childhood*. Vol. 25. Routledge. <https://doi.org/10.4324/9781315009698>
- [38] Alexander Pritzel, Benigno Uria, Sriram Srinivasan, Adria Puigdomenech Badia, Oriol Vinyals, Demis Hassabis, Daan Wierstra, and Charles Blundell. 2017. Neural episodic control. In *International Conference on Machine Learning*. PMLR, 2827–2836.
- [39] Martin L Puterman. 2014. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.
- [40] Joaquin Quionero-Candela, Masashi Sugiyama, Anton Schwaighofer, and Neil D. Lawrence. 2009. *Dataset Shift in Machine Learning*. The MIT Press.
- [41] Antonin Raffin, Ashley Hill, Adam Gleave, Anssi Kanervisto, Maximilian Ernestus, and Noah Dormann. 2021. Stable-Baselines3: Reliable Reinforcement Learning Implementations. *Journal of Machine Learning Research* 22, 268 (2021), 1–8. <http://jmlr.org/papers/v22/20-1364.html>
- [42] Stefan Schaal. 1999. Is imitation learning the route to humanoid robots? *Trends in cognitive sciences* 3, 6 (1999), 233–242.
- [43] Lukas Schäfer, Filippos Christianos, Josiah P Hanna, and Stefano V Albrecht. 2022. Decoupled Reinforcement Learning to Stabilise Intrinsically-Motivated Exploration. In *AAMAS*. 1146–1154.
- [44] John Schulman, Sergey Levine, Philipp Moritz, Michael Jordan, and Pieter Abbeel. 2015. Trust Region Policy Optimization. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37* (Lille, France) (ICML ’15). JMLR.org, 1889–1897.

- [45] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal Policy Optimization Algorithms. *arXiv e-prints arXiv:1707.06347* (July 2017). arXiv:1707.06347 [cs.LG]
- [46] Richard S. Sutton and Andrew G. Barto. 2018. *Reinforcement Learning: An Introduction*. MIT Press Ltd, Cambridge, MA, USA. [https://www.ebook.de/de/product/32966850/richard\\_s\\_university\\_of\\_alberta\\_sutton\\_andrew\\_g\\_co\\_director\\_autonomous\\_learning\\_laboratory\\_barto\\_reinforcement\\_learning.html](https://www.ebook.de/de/product/32966850/richard_s_university_of_alberta_sutton_andrew_g_co_director_autonomous_learning_laboratory_barto_reinforcement_learning.html)
- [47] Richard S Sutton, David McAllester, Satinder Singh, and Yishay Mansour. 1999. Policy gradient methods for reinforcement learning with function approximation. *Advances in neural information processing systems* 12 (1999).
- [48] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. 2013. Intriguing properties of neural networks. *arXiv e-prints arXiv:1312.6199* (2013).
- [49] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. 2017. Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 23–30. <https://doi.org/10.1109/IROS.2017.8202133>
- [50] Alan M Turing. 1950. Computing machinery and intelligence. *Mind* 59, 236 (1950), 433.
- [51] Jonathan Uesato\*, Ananya Kumar\*, Csaba Szepesvari\*, Tom Erez, Avraham Ruderman, Keith Anderson, Krishnamurthy (Dj) Dvijotham, Nicolas Heess, and Pushmeet Kohli. 2019. Rigorous Agent Evaluation: An Adversarial Approach to Uncover Catastrophic Failures. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=B1xhQhRcK7>
- [52] Rui Wang, Joel Lehman, Jeff Clune, and Kenneth O Stanley. 2019. Poet: open-ended coevolution of environments and their optimized solutions. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 142–151.
- [53] Ziyu Wang, Victor Bapst, Nicolas Heess, Volodymyr Mnih, Remi Munos, Koray Kavukcuoglu, and Nando de Freitas. 2016. Sample Efficient Actor-Critic with Experience Replay. *arXiv e-prints arXiv:1611.01224*, Article arXiv:1611.01224 (Nov. 2016). arXiv:1611.01224 [cs.LG] <https://ui.adsabs.harvard.edu/abs/2016arXiv161101224W>
- [54] Brian D. Ziebart, Andrew Maas, J. Andrew Bagnell, and Anind K. Dey. 2008. Maximum Entropy Inverse Reinforcement Learning. In *Proceedings of the 23rd National Conference on Artificial Intelligence - Volume 3 (Chicago, Illinois) (AAAI'08)*. AAAI Press, 1433–1438.