Joint Goal and Strategy Inference across Heterogeneous Demonstrators via Reward Network Distillation

Letian Chen

letian.chen@gatech.edu
Georgia Institute of Technology
Atlanta, GA

Muyleng Ghuy

mghuy3@gatech.edu
Georgia Institute of Technology
Atlanta, GA

Rohan Paleja

rpaleja3@gatech.edu
Georgia Institute of Technology
Atlanta, GA

Matthew Gombolay

matthew.gombolay@cc.gatech.edu
Georgia Institute of Technology
Atlanta, GA

ABSTRACT

Reinforcement learning (RL) has achieved tremendous success as a general framework for learning how to make decisions. However, this success relies on the interactive hand-tuning of a reward function by RL experts. On the other hand, inverse reinforcement learning (IRL) seeks to learn a reward function from readily-obtained human demonstrations. Yet, IRL suffers from two major limitations: 1) *reward ambiguity* - there are an infinite number of possible reward functions that could explain an expert's demonstration and 2) *heterogeneity* - human experts adopt varying strategies and preferences, which makes learning from multiple demonstrators difficult due to the common assumption that demonstrators seek to maximize the same reward. In this work, we propose a method to jointly infer a task goal and humans' strategic preferences via network distillation. This approach enables us to distill a robust task reward (addressing reward ambiguity) and to model each strategy's objective (handling heterogeneity). We demonstrate our algorithm can better recover task reward and strategy rewards and imitate the strategies in two simulated tasks and a real-world table tennis task.

CCS CONCEPTS

• **Theory of computation** → **Inverse reinforcement learning**; • **Computing methodologies** → **Learning from demonstrations**; *Neural networks*.

KEYWORDS

Inverse Reinforcement Learning; Heterogeneous Demonstration; Preference Learning

ACM Reference Format:

Letian Chen, Rohan Paleja, Muyleng Ghuy, and Matthew Gombolay. 2020. Joint Goal and Strategy Inference across Heterogeneous Demonstrators via Reward Network Distillation. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction (HRI '20)*, March 23–26,

2020, Cambridge, United Kingdom. ACM, New York, NY, USA, 10 pages.
<https://doi.org/10.1145/3319502.3374791>

1 INTRODUCTION

Despite the success reinforcement learning (RL) has achieved in recent years [12, 17, 24, 35], this optimization technique still requires human RL experts to carefully design reward functions to be effective. Researchers, with the best intentions, can naively specify a reward function that results in an RL policy that catastrophically fails to meet the researcher's latent expectations [3]. Unfortunately, the field of RL lacks a standard, repeatable process by which we can specify robotic tasks and have the robot autonomously synthesize sensible policies via RL. Prior research has further shown the difficulty human experts have in defining a desired policy for the robots even though they can readily demonstrate the task [6, 29]. As such, researchers have often turned to Imitation Learning (IL), an IRL techniques within Learning from Demonstration (LfD).

Although IL is a relatively straightforward approach that seeks to directly learn a mapping from states to actions, the learned policy is highly intertwined with environment dynamics. Slight changes in dynamics or highly stochastic environments could cause IL to easily fail due to a "covariate" shift [32]. In contrast, the demonstrated reward function is a more transferable and robust definition of the task, representing latent objectives a behavior tries to accomplish. Effective reward learning holds utility even after environment dynamics change [9]. Recently, researchers have made strides to disentangle reward functions from the environment's dynamics [13].

Despite progress, IRL still suffers from fundamental shortcomings [22]. The first drawback is known as *reward ambiguity*, which is defined as a problem where there exists an infinite number of possible reward functions that could explain an expert's demonstration. The second major drawback is caused by limitations in reasoning about heterogeneous demonstrations. Experts typically adopt heuristics (i.e., "mental shortcuts") to solve challenging optimization problems, but these highly-refined strategies can present a heterogeneity in behavior across task demonstrators that prior methods fail to effectively reason about. Previous work in IRL has tried to address heterogeneity [28], but only focused on simplified examples which do not typify the complexity of the real world.

To overcome these limitations, we propose Multi-Style Reward Distillation (MSRD, see Figure 1), a novel IRL technique for learning

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HRI '20, March 23–26, 2020, Cambridge, United Kingdom

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6746-2/20/03...\$15.00

<https://doi.org/10.1145/3319502.3374791>

Figure 1: Illustration of Multi-Style Reward Distillation (MSRD). R_0 represents task reward, $\tilde{R}_i$ represents strategy-only reward, and R_i represents strategy-combined reward (see Section 4.1 for explanation). Task reward is shared across different strategies, while strategy-only reward is specific to each strategy.

to tease out task and strategy rewards from heterogeneous demonstration. In MSRD, we model the demonstrated reward function as a combination of the task goal and the demonstrator’s strategic preferences. For example, while tennis players seek to win (i.e., the task reward), players exhibit various strategies based on their skills vis-à-vis those of the opponent (e.g., serve-and-volley, counter-punch, etc). We argue that human experts have specific strategic preferences, especially in complex tasks where it is intractable for humans to compute the global optimum [22]. In these cases, humans use various heuristics, which results in different types of demonstrations. Unlike prior work in IRL where demonstrators are assumed to optimize a single task reward (i.e., homogeneity), we take on a more human-aware approach by optimizing a combination of the task reward and an intrinsic strategy (i.e., affording heterogeneity).

Our method assumes that expert demonstrators prioritize task completion over conforming to a specific strategy. Under this assumption, heterogeneous demonstrations provide more information to help us tease out the latent task reward versus having only a demonstration of a single strategy. This approach allows us to better model the task reward function by distilling shared components of each demonstrators’ trajectories. The task reward recovered from our approach is thus robust and has less ambiguity.

Our contribution in this paper is three-fold:

- (1) We propose a novel IRL framework where we jointly learn task reward and strategic rewards to gain a better estimation of the task reward as well as an estimation for separate strategy reward component for each strategy.
- (2) We show our algorithm’s success on two virtual robot control tasks and one real-world physical robot table tennis task. *Nota bene: we are one of few cases that successfully applies IRL on a robotic continuous control problem.*

- (3) Our results show that MSRD’s learned task reward function achieves high correlation with the ground-truth task reward and that a policy learned from the specific strategy reward recovers the strategic preference.

In addition to MSRD, we also develop a method that helps generate heterogeneous policies of different strategies in simulated environments. This technique is leveraged by our virtual experiments to synthesize heterogeneous task strategies.

2 RELATED WORK

IRL is an LfD approach that aims to infer a demonstrator’s objective (i.e., reward function) given a set of performed trajectories. The problem is typically formulated within a Markov Decision Process (MDP) framework. IRL is known to be an ill-posed problem as there are infinitely many reward functions that could explain expert demonstration as optimal including degenerated cases (e.g., $R = 0$ [26]). Two primary methods exist in the IRL literature that aim to solve this ambiguity: maximum margin approaches [1, 31] and probabilistic approaches [30, 44, 45]. Maximum margin approaches try to find a reward function that explains an expert’s trajectory not only as being optimal, but also as being better than all other trajectories by a margin. Probabilistic approaches (e.g. Bayesian IRL [30], Maximum-Entropy IRL [45], and Maximum Causal Entropy IRL [44]) assume trajectories with higher rewards have an exponentially higher probability of being generated and apply a maximum likelihood framework to solve for such a reward function. Advantages of this framework include its tolerance of non-optimal demonstrations and ease of applying gradient-based optimization. Maximum-entropy based IRL methods are thus suitable for deep neural network models to find the reward function [41], obviating

the need for feature engineering. Despite the advantages of a maximum entropy IRL (ME-IRL) framework, perfect information about the system dynamics and the ability to calculate the exact state feature counts are required. Guided Cost Learning (GCL [12]) and its successor, Adversarial Inverse Reinforcement Learning (AIRL [13]) have been proposed to tackle ME-IRL problems. AIRL poses the reward learning problem in an adversarial setting, where the reward function tries to assign a high reward to expert demonstrations while assigning a low reward to generated trajectories. Meanwhile, the policy is optimized over the learned reward function to maximize its expected reward.

Alternatively, preference learning is another well-explored research area that aims to learn individuals' proclivities. In robotics, preference learning typically involves learning a human's preference and adapting a robot's behavior to better collaborate with human-beings [16, 28]. There are several methods that could infer preference information without asking the human to explicitly provide preference information [14, 15]. For example, Schafer et al. [34] utilizes collaborative filtering for a recommender system and Xu et al. [42] formulates a meta-IRL problem that could learn a prior over preferences. A more recent direction of preference learning lies in the multi-task setting. Dimitrakakis and Rothkopf [10] models reward-policy pairs (R_i, π_i) drawn from an unknown prior. Choi and eung Kim [7] integrates Dirichlet process mixture model into Bayesian IRL as a task prior. Repeated inverse reinforcement learning (RIRL) [2] formalizes the setting in which a user is observed performing different tasks. The goal of RIRL is to infer a task-independent preference. However, one of the biggest limitations of RIRL is the method assumes full knowledge of each task's reward function and the relative weights between task rewards and preference rewards, which is unrealistic in most real-world applications. Observational Repeated IRL (ORIRL) relaxes the assumption by introducing a learnable relative weight yet still assumes perfect knowledge of the task reward Woodworth et al. [40]. Furthermore, both RIRL and ORIRL work is carried out in discrete finite state-action space, which is a significant simplification of real-world, continuous, state-action space. While preference learning in a multi-task setting tries to learn shared preferences across tasks, we propose to learn a shared task reward from observing different strategies for the same task.

Learning from heterogeneous demonstrations poses a significant challenge in LfD. The typical approach is to assume homogeneity over demonstrators; however, [32] found that commercial line pilots exhibit significant heterogeneity in performing a well-posed task (e.g., executing a pre-specified flight plan) as to make it more practical to learn from a single trajectory and disregard the remaining data. We are motivated by such settings to develop a method that might still be capable of generating a robust task specification while also leveraging all available data to help mitigate the curse of dimensionality. Few have addressed multiple-reward-function learning. One such approach by Nikolaidis et al. [27] clusters different kinds of behaviors and applies inverse reinforcement learning on each cluster. Another approach, Option GAN [19], learns a division in demonstration state space and a separate reward function and policy for each subspace. However, neither considers the relationship between the learned reward functions. In contrast, our work poses a common task reward function in each of the reward functions,

together with a separate strategy reward function for each strategy, which enables us to tease out task reward and strategy rewards.

Another line of work that tries to model heterogeneous expert data lies in imitation learning. Generative Adversarial Imitation Learning (GAIL) [21] is a popular imitation learning algorithm in which a discriminator tries to distinguish between expert trajectories and generated trajectories, while a generator tries to deceive the discriminator. Some extensions to GAIL utilize a latent variable to model multi-modal (multi-strategy) demonstration data [18, 25]. Despite the good performance of imitation learning algorithms in some contexts, such methods are severely sensitive to a change of environments (e.g., a small disturbance could result in total failure of intended task). By contrast, the reward function is more robust to a change of dynamics in the environment, as the function represents the intention of the behavior. In this work, we separate the task reward and the strategy reward, thereby better extracting a transferable signal describing the demonstrators' latent intent.

In contrast, our approach proposes a novel problem setup that considers multi-strategy demonstrations for a single task. We propose a distillation-regularization structure between the task reward and strategy rewards, shown in Figure 1, which aims to distill common task rewards and separate strategy preferences.

3 PRELIMINARIES

3.1 Markov Decision Process

A Markov Decision Process (MDP) M is a 6-tuple $(S, A, R, T, \gamma, \rho_0)$, where S and A are the state and action spaces, respectively. $\gamma \in (0, 1)$ is the temporal discount factor. $R(s, a)$ represents the reward after executing action a in state s . In some cases, $R(s, a)$ could be simplified as $R(s)$. $T(s, a, s')$ is the transition probability to s' from state s after taking action a . $\rho_0(s)$ is the initial state distribution. The standard goal of reinforcement learning is to find the optimal policy π^* that maximizes the discounted future reward, $J(\pi) = \mathbb{E}_{\tau \sim \pi} [\sum_{t=0}^T \gamma^t R(s_t, a_t)]$, where $\tau = (s_0, a_0, \dots, s_T, a_T)$ denotes a sequence of states and actions induced by the policy and dynamics. We instead consider a more general maximum entropy objective introduced by Ziebart [44], which augments the standard objective with an entropy bonus to favor stochastic policies and to encourage exploration during optimization, $J(\pi) = \mathbb{E}_{\tau \sim \pi} [\sum_{t=0}^T \gamma^t R(s_t, a_t) + \alpha H(\pi(\cdot|s_t))]$.

3.2 Inverse Reinforcement Learning

IRL considers an MDP sans reward function $(M \setminus R)$ with the goal being to infer reward function $R(s, a)$ given a set of demonstrated trajectories $\mathcal{D} = \{\tau_1, \tau_2, \dots, \tau_N\}$. A typical assumption for IRL is that demonstrated trajectories are optimal, or at least near-optimal. In the maximum entropy IRL (Max-Ent IRL) framework, inference of the reward function is turned to a maximum likelihood optimization problem by assigning occurrence probability of a trajectory proportional to the exponential of discounted cumulative reward, $p_\theta(\tau) \propto e^{\sum_{t=0}^T \gamma^t r_\theta(s_t, a_t)}$. Therefore, Max-Ent IRL aims to find the reward function under which the demonstrated trajectories have the highest likelihood, as shown in Equation 1.

$$\max_{\theta} \mathbb{E}_{\tau \sim \mathcal{D}} [\log p_\theta(\tau)] \quad (1)$$

AIRL [13] casts the optimization in a generative adversarial framework, learning a discriminator, D , to distinguish between experts and a deceiving generator, π , that learns to imitate the expert. This framework follows Max-Ent IRL’s assumption that trajectory likelihood is proportional to the exponential of rewards. D is defined in Equation 2, where $f_\theta(\tau)$ is the learnable reward function and π is the current policy. D is updated by minimizing its cross entropy loss to distinguish expert trajectories from generator policy rollouts (Equation 3). π is trained to maximize the pseudo-reward function given by $f(s, a)$.

$$D_\theta(s, a) = \frac{\exp(f_\theta(s, a))}{\exp(f_\theta(s, a)) + \pi(a|s)} \quad (2)$$

$$L_D = -\mathbb{E}_{\tau \sim \mathcal{D}, (s, a) \sim \tau} [\log D(s, a)] - \mathbb{E}_{\tau \sim \pi, (s, a) \sim \tau} [1 - \log D(s, a)] \quad (3)$$

3.3 Neural Network Distillation

Neural network distillation applies supervised regression to train a student network to produce the same output distribution as a trained teacher network. The method was first proposed by Hinton et al. [20] and has been applied in RL mainly for performing policy distillation [8, 33, 37]. Particularly, Teh et al. [37] proposes that instead of distilling each policy π_i to a general policy π_0 , we could gain a faster convergence with a two-column architecture by defining $\pi_i = \pi_0 + \tilde{\pi}_i$. Accordingly, π_i only needs to learn a near-zero difference between a common policy and the task-specific policy.

4 METHOD

4.1 Problem Setup

We consider a setup in which there is only one task (one MDP $M = (S, A, R^{(0)}, T, \gamma, \rho_0)$), but the demonstrations are generated by employing varying strategies. Different strategies may come from different experts who have personalized strategic preferences or one expert that has mastered several. Therefore, despite heterogeneity within demonstrations, all trajectories in the dataset should still be near-optimal in terms of task reward $R^{(0)}$. We denote each strategy’s demonstration dataset as $\mathcal{D}^{(i)} = \{\tau_1^{(i)}, \tau_2^{(i)}, \dots, \tau_M^{(i)}\}$, where $i \in \{1, 2, \dots, N\}$ is strategy index, M is the number of demonstration trajectories for one strategy, and N is the number of strategies. Our first objective is to infer a shared task reward function $R^{(0)}$ despite there being different strategies in the demonstration dataset. The second objective is to infer the strategy-only reward $\widetilde{R}^{(i)}$. Combining strategy-only reward with the task reward function will result in a strategy-combined reward $R^{(i)}$, which should induce the observed expert strategic behaviors.

4.2 Task and Strategy Reward

We first propose to model the strategy-combined reward function that is optimized by a demonstrator to be a linear combination of the task and the strategy-only reward as given by Equation 4.

$$R^{(i)}(\cdot) = R^{(0)}(\cdot) + \alpha_i \widetilde{R}^{(i)}(\cdot) \quad (4)$$

Despite the simplicity, we argue Equation 4 makes a reasonable assumption. Several previous IRL works also apply linear, if not

more constraining, assumptions to combine reward functions: Amin et al. [2] create a combined reward by adding a task reward with a cross-task shared preference, while we add shared task reward with each strategy reward; Woodworth et al. [40] propose a similar formulation, except their task reward is known and they try to learn task-independent preference in a multi-task setting. In the RL literature, there is also substantial work combining rewards linearly (e.g., adding intrinsic reward, such as curiosity or entropy to the original task reward for the sake of exploration). In fact, many engineered reward functions are also a linear combination of several reward components (e.g., for OpenAI Gym [4] MuJoCo [38] hopper environment, its reward function is defined as a linear combination of forward speed, living bonus, and action penalty). Additionally, assuming a linear combination could also provide interpretability (see Figure 4). Our approach is unique in that it shares the task reward while having the flexibility on each strategy reward, which we argue is more realistic while allowing us to apply joint inference of the task reward and strategy reward.

4.3 Reward Network Distillation

To infer the shared task reward function between different strategies, we propose utilizing network distillation to distill common knowledge from each separately learned strategy-combined reward $R^{(i)}$ to the task reward function $R^{(0)}$. We also want to regularize $R^{(i)}$ to be close to $R^{(0)}$, since we have the prior knowledge that despite heterogeneous strategic preferences, all experts should still be prioritizing optimizing the task reward $R^{(0)}$ to achieve at least near-optimal performance. Previous distillation methods mainly focus on distilling classification results, and therefore KL-divergence between teacher and student outputs could be a good choice for regularization. However, reward functions are real-valued functions and therefore probabilistic distance metrics do not fit. Thus, we propose to regularize the expected L2-norm of the difference between the reward functions, as shown in Equation 5, in which $\pi^{(i)}$ is the optimal policy under reward function $R_{\theta_i}^{(i)}$.

$$L_{\text{reg}} = \mathbb{E}_{(s, a) \sim \pi^{(i)}} \left(\left\| R_{\theta_i}^{(i)}(s, a) - R_{\theta_0}^{(0)}(s, a) \right\|_2 \right) \quad (5)$$

Note that we are using an index both on θ and R to denote that each strategy-combined reward $R^{(i)}$ has its own reward parameters, and that these are approximated by separate neural networks with parameters θ_i for each strategy and θ_0 for the task reward. There is no parameter sharing between strategy and task reward.

Due to the computational cost of optimizing $\pi^{(i)}$ using RL, we seek to avoid fully optimizing it inside the IRL loop. Therefore, we apply an iterative reward function and policy training schedule, similar to AIRL. Through such way, we could learn reward function and policy simultaneously. Combining AIRL’s objective (Equation 3) with the distillation objective, we want to maximize L_D in Equation 6.

$$L_D = \sum_{i=1}^N \left[\mathbb{E}_{(s, a) \sim \tau_j^{(i)} \sim \mathcal{D}^{(i)}} \log D_{\theta_i}(s, a) + \mathbb{E}_{(s, a) \sim \pi^{(i)}} \log(1 - D_{\theta_i}(s, a)) - \mathbb{E}_{(s, a)} \left(\left\| R_{\theta_i}^{(i)}(s, a) - R_{\theta_0}^{(0)}(s, a) \right\|_2 \right) \right] \quad (6)$$

D is dependent on θ_i via Equation 2 ($R_{\theta_i}^{(i)}$ corresponds to f_θ). Each $\pi^{(i)}$ optimizes¹ $R_{\theta_i}^{(i)}$. Yet, while Equation 6 should be able to distill the shared reward into $R_{\theta_0}^{(0)}$, the distillation is inefficient as $R_{\theta_0}^{(0)}$ will work as a strong regularization for $R_{\theta_i}^{(i)}$ before successful distillation.

Instead, our structure in Equation 4 allows for a two-column reparameterization, speeding up knowledge transfer and making the learning process easier [37]. Combining Equation 4 and Equation 6, we arrive at Equation 7.

$$L_D = \sum_{i=1}^N \left[\mathbb{E}_{(s,a) \sim \tau_j^{(i)} \sim \mathcal{D}^{(i)}} \log D_{\theta_i, \theta_0}(s, a) + \mathbb{E}_{(s,a) \sim \pi^{(i)}} \log (1 - D_{\theta_i, \theta_0}(s, a)) - \alpha_i \mathbb{E}_{(s,a)} \left(\left\| \widetilde{R_{\theta_i}^{(i)}}(s, a) \right\|_2 \right) \right] \quad (7)$$

The key difference between Equation 7 and Equation 6 is that D depends on both $R_{\theta_0}^{(0)}$ and $\widetilde{R_{\theta_i}^{(i)}}$ instead of separate $R_{\theta_i}^{(i)}$. Thus, $R_{\theta_0}^{(0)}$ directly updates from the discriminator's loss rather than waiting for knowledge to be learned by a strategy-combined reward and subsequently distilled into a task reward. Further, the last term of Equation 7 reduces to a simple L2-regularization on strategy-only reward's output, weighted by α_i . This formulation provides us with a new view to interpret the relative weights of the strategy-only reward α_i : the larger α_i is, the more the strategy-only reward will influence the strategy-combined reward. Therefore, we will have higher regularization to account for possible overwhelming of the task reward function. Comparing Equation 7 and 3, we could interpret MSRD in another view: optimizing θ_i only via IRL objective results in a combination of task and strategy reward, and adding regularization on strategy reward will encourage to encode only necessary information in θ_i and share more knowledge in θ_0 .

4.4 Multi-Style Reward Distillation

We summarize our algorithm in Algorithm 1. In the algorithm we first collect the heterogeneous-strategy expert dataset and initialize network parameters as well as relative weights. For each training epoch, we will iterate over all strategies (line 5). For each strategy, we first collect K trajectories generated by its corresponding policy $\pi^{(i)}$ (line 6). We also sample expert trajectories for strategy i from the dataset (line 7). We then train the Discriminator (reward function) with loss given by Equation 7 and data τ_j^{gen} and τ_j^{exp} (line 8). After training the reward function, we could assign pseudo-reward to trajectories we generate with reward function $R_{\theta_i}^{(i)}$ (line 9). Finally, we update the policy according to the trajectory generated and pseudo-reward signal (line 10). In practice, we could also postpone the gradient update for $R^{(0)}$ at the end of one sweep of strategies to stabilize the learning of the task reward.

5 EXPERIMENT SETUP

We tested MSRD on both virtual and real-world environments. Here, we describe our experiments, showing results in Section 6.

¹We choose Trust Region Policy Optimization (TRPO) [13]

Algorithm 1 Multi-Strategy Reward Distillation

```

1: Obtain heterogeneous-strategy expert dataset  $\mathcal{D}^{(i)} = \{\tau_1^{(i)}, \dots, \tau_M^{(i)}\} \forall i \in \{1, \dots, N\}$ 
2: Initialize  $R^{(0)}, R^{(i)}, \pi^{(i)} \forall i \in \{1, 2, \dots, N\}$ 
3: Determine relative weights  $\alpha_i \forall i \in \{1, 2, \dots, N\}$ 
4: while not converged do
5:   for  $i = 1$  to  $N$  do
6:     Collect  $K$  ( $K \leq M$ ) trajectories  $\tau_j^{\text{gen}}$  by executing  $\pi^{(i)}$ 
7:     Sample  $K$  trajectories  $\tau_j^{\text{exp}}$  from  $\mathcal{D}^{(i)}$ 
8:     Train  $\theta_i$  and  $\theta_0$  with Equation 7,  $\tau_j^{\text{gen}}$  and  $\tau_j^{\text{exp}}$ 
9:     Assign reward for all transitions in  $\tau_j^{\text{gen}}$ :  $r(s, a) = R_{\theta_i}^{(i)}(s, a)$ 
10:    Update policy  $\pi_i$  using trajectories  $\tau_j^{\text{gen}}$  via TRPO with entropy bonus to encourage exploration
11:  end for
12: end while
13: return  $R^{(0)}, R^{(i)}, \pi^{(i)}$ 

```

5.1 Virtual Experiments

We first tested MSRD on two simulated environments: a simpler inverted pendulum control task and a more difficult hopper locomotion task (see Figure 6 and supplementary materials for illustration). The goal of the inverted pendulum task is to balance a pendulum on a cart by moving the cart left/right, making it a single degree of freedom (DoF) problem, based on a 2D observation. The reward for inverted pendulum is defined as the negative absolute value of the pendulum angle from upright position. The objective of hopper is to control 3-DoF joints to move forward based on its 11-dimensional observation. The reward for hopper is defined as the speed at which it moves forward. We used the OpenAI Gym [4] MuJoCo [38] implementation for both environments but made the following changes to fit our application: 1) remove termination judgements to gain flexibility in behaviors; 2) add timeout constraint of 1,000 steps.

5.1.1 Heterogeneous Demonstration Collection. Our algorithm can utilize heterogeneous-strategy demonstrations. Therefore, we first need to generate a variety of demonstrations to emulate heterogeneous strategies that humans will apply in solving problems for our virtual experiments. Typical RL algorithms can only generate single-modal policy for each task in attempt to maximize the task reward. Some previous work has tried to generate diverse behaviors [11, 43]. Among them, Diversity is All You Need (DIAYN) trains a discriminator to distinguish different behaviors and trains policies utilizing the discriminator's output as a pseudo-reward. The pseudo-reward is shown in Equation 8, where z is the strategy index, q is the posterior decoding z from s , and $p(z)$ is the prior distribution. However, DIAYN only discovers different behaviors without task specification, so we augment DIAYN to incorporate task reward by the linear form of Equation 4.

$$r_z(s, a) = \log q_\phi(z|s) - \log p(z), \quad (8)$$

We also propose a method to encourage different strategies taking different actions in the same state via a diversity reward in which k is the strategy index and π_k is the policy for strategy k

Figure 2: Correlation between ground-truth and estimated task reward, normalized for each strategy to $[0, 1]$, for inverted pendulum (left) and hopper (right) environments. Reward is invariant to shift/scale. r is correlation coefficient.

Figure 3: Correlation between ground-truth vs. estimated strategy reward by MSRD and Airl on Inverted Pendulum.

(Equation 9). Equation 9 encourages the KL-divergence between different strategy's policies to be large. Linearly combining KL-encouraged diversity reward with the task reward, we can train strategies that optimize both the task goal and the diversity goal.

$$r_{\text{KL}}^k(s) = \sum_{i=1}^N \text{KL}(\pi_k(\cdot|s) || \pi_i(\cdot|s)), \quad (9)$$

We trained both "DIAYN + Extrinsic Reward" and "KL-Encouraged + Extrinsic Reward" policies to collect heterogeneous trajectories that applied different strategies to solve the task. From all the strategies generated, we chose 20 strategies for inverted pendulum and the two most significant strategies for hopper. Generally, different strategies in inverted pendulum encourage the cart to stay at different angles, but some strategies maintain dynamic balance by periodically moving the cart left and right. Two different strategies in hopper are "Hop" and "Crawl" as illustrated in Figure 6. More details are available in the supplementary materials.

5.2 Real-World Experiments

For our second environment, we tested MSRD on its capability to learn various table tennis strokes from expert human demonstration. Participants were asked to kinetically teach a robot arm to hit an incoming ping pong ball using three different stroke strategies: push, slice, and top spin (see supplementary videos for illustration), from both a forehand and backhand position. Detailed environment setup could be found in supplementary.

5.2.1 Experiment Design and Subjects. We adopt a within-subjects design for demonstration collection, requiring subjects finish all six combinations of positions and strategies. We pseudo-randomized the order of forehand and backhand, as well as the order of three strategies. We recruited 10 subjects from a population of college graduate students. Subjects were compensated.

5.2.2 Experiment Procedure. When the participants arrived, they read and signed a consent form detailing the purpose, duration, and study procedure. For each trial, we began by showing a video tutorial on how to move the robot arm to hit the incoming ball with a specific strategy. After the tutorial, we allowed participant

to practice with the automatic ping pong feeder. Once the subject felt ready, we began the recording process. Saved trajectories were those in which the participant was able to successfully return the ball to the opponent's side and the movement of striking closely resembled the strategy assigned. We collected three recordings for each conditions, position (forehand/backhand), and strategy combination. We record at each timestep of the trajectory the robot joint angles, angle rates and the ball position.

6 RESULTS AND DISCUSSION

In this section, we report and analyze MSRD's results on three environments and benchmark against Airl to elucidate MSRD's advantage on recovering both the latent task reward (the essential goal of the demonstrators) and the means by which the task is accomplished (i.e. the strategy). We explore two hypotheses:

H1: *The task reward learned by MSRD has a higher correlation with the true task reward than Airl.*

H2: *Strategy-only reward learned by MSRD has a higher correlation with true strategic preferences than Airl.*

We assessed both hypotheses quantitatively and qualitatively for the simulation environments only as the ground-truth reward functions are available. In the physical robot experiment, users are instructed to execute the task instead of optimizing an objective function, meaning that we do not have access to the underlying task reward or strategy-only rewards. Therefore, **H1** and **H2** were assessed qualitatively in the physical robot experiment. We did not compare MSRD method with more traditional IRL methods (such as Max-Margin IRL, Feature-Matching IRL, Max-Ent IRL, etc.) because they rely on sophisticated feature designs. We leverage end-to-end neural network architectures to avoid manual feature extraction. Furthermore, we wanted to provide a fair comparison between methods; given MSRD leverages Airl, the most appropriate choice is thus Airl. In future work, we propose considering integrated other IRL methods into MSRD. In future work, our method could be combined with any existing strategy inference method (e.g., [27]) to make it applicable to cases where strategy labels are unavailable.

Figure 4: Task/Strategy Reward Functions of Inverted Pendulum vs Pendulum Angle and Corresponding Behaviors.

6.1 Simulated Environments

To test **H1**, we constructed a dataset of trajectories that have various task performances utilizing noise injection [5]. We note that this dataset was also generated with various-strategy policies to be representative of the entire trajectory space. We then evaluated the reward function learned by AIRL and MSRD on the trajectories, comparing estimated vs. ground-truth rewards. We show a correlation of estimated rewards and ground-truth task rewards in Fig.2. The task reward function learned through MSRD has a higher correlation with the ground-truth reward function (0.998 and 0.943) versus AIRL (0.51 and 0.89) for each domain, respectively). AIRL’s reward function overfits to some strategies and mixes the task reward with that strategy-only reward, making its estimation unreliable for other strategies’ trajectories.

To test **H2**, we calculated the correlations of MSRD’s strategy-only rewards with the true strategic preferences and compared that with the correlation of AIRL’s rewards when AIRL is trained on each individual strategy. In simulated domains, true strategic preferences are available as the pseudo-reward in Equation 8 and 9. Correlations of both methods for all strategy rewards in inverted pendulum are shown in Figure 3. A paired t-test shows that MSRD achieves a statistically significantly higher correlation ($M = 0.779$, $SD = 0.239$) for the strategy rewards versus AIRL ($M = 0.635$, $SD = 0.324$) trained separately for each strategy, $t(19) = 1.813$, $p = 0.0428$ (one-tailed). A Shapiro-Wilk test showed the residuals were normally distributed ($p = 0.877$). For the hopper domain, MSRD achieved 0.85 and 0.93 correlation coefficient for the hop and crawl strategy, compared with AIRL’s 0.80 and 0.82 respectively. We omit a t-test here due to the limited number of strategies. We could test the discrimination of strategy rewards by evaluating each strategy’s reward function on each strategy’s trajectory; we expect to observe that the strategy-only reward function of each strategy gives its corresponding trajectory the highest reward. We show in Figure 5 that this expectation holds. Out of 20 strategy-only rewards, 16 receive highest rewards in corresponding trajectories. A Binomial test shows we are significantly better than chance ($p < .001$).

We are unable to examine the reward landscape in the inverted pendulum environment as it is four dimensional. Thus, we choose to fix three dimensions (cart position, cart velocity and pendulum angular velocity) to zero and investigate the reward change within the one remaining dimension (pendulum angle). The relationship

Figure 5: Evaluation of strategy rewards on strategy trajectories (Inverted Pendulum); rows normalized to $[0, 1]$.

between rewards and pendulum angles in task and strategy reward functions are illustrated in Figure 4, in which the task reward function reaches its peak when the angle of the pendulum is near zero. This precisely recovers the task reward. For strategy-only reward functions, strategy 13 encourages the pendulum to lean left (demonstration behavior shown in bottom-left of Figure 4), while strategy 4 encourages the policy to tilt the pendulum right (demonstration behavior shown in bottom-right). Therefore, strategy-only rewards learned by MSRD captures specific preferences within demonstrations. Figure 4 also shows the magnitude of the task reward is larger than the strategy reward, which affirms our expectation that an emphasis is being put towards accomplishing the task.

In the hopper environment, it is harder to visualize the reward landscape due to a high-dimensional observation space and a lack of interpretability of states. Therefore, instead of visualizing a reward curve, we evaluate the estimated strategy-only reward on trajectories from both strategies to provide evidence for **H2**. Figure 6 shows that when given a hopping trajectory, the hop strategy-only reward function gives higher reward for that behavior than crawl strategy-only reward function. Similarly, in the crawl trajectory case (Figure 6), the crawling strategy-only reward gives a higher value than the hop strategy-only reward. Therefore, the strategy-only reward function recovered by MSRD gives a higher reward to the corresponding behavior than the other strategy-only reward

Figure 6: Hop and crawl strategy reward on Hop (top) and Crawl (bottom) trajectories. Blue and orange numbers are learned Hop and Crawl strategy rewards, respectively.

function, thus providing encouragement to the policy towards the intended behavior (H2).

These results across our simulated environments show our algorithms' success in both task reward recovery and strategy reward decomposition. This capability is a novel contribution to the field of LfD in that we are able to tease out strategies from the underlying task and effectively learn policies that can reproduce both the strategy and a well-performed policy for the underlying task.

6.2 Physical Environment

We utilize four deep neural networks (DNNs) consisting of three fully connected layers (32 nodes on each hidden-layer) to represent task, push, slice and topspin rewards. The ball states alongside robot arm joints serve as our inputs. The label of different types of demonstration (forehand-push, backhand-slice, etc.) is available to our algorithm. Figure 7 shows four frames of the learned trajectories for our defined strategies. The change in angle and the upward/downward motion of the paddle throughout the policy trajectory are key factors in the display of different strategies (as these induce spin). Push is associated with a small change in angle as it is not attempting to add any spin onto the ball. Slice places a backspin on the ball, and thus the angle of the paddle will quickly tilt up as shown in Figure 7. Conversely, topspin places a topspin on the ball; to do so, the associated trajectory has a quick upward motion. Figure 8 provides quantitative evidence that the strategy-only reward should be maximal given a demonstration utilizing the corresponding strategy. After just 30 minutes of training on each strategy, the robot was able to learn to strike 83% of the fed balls. The robot learned to perform all strategies, and the robot's best strategy, topspin, resulted in 90% of balls returned successfully.

To further verify that our task reward was learning the correct behavior, the task reward function was used to evaluate each of the demonstrated trajectories (three trajectories for each of the three strategies). In the case of the original demonstrations (i.e., where the ball was struck and landed in bounds), the average and standard deviation of the task reward across demonstrations and

Figure 7: Different strategies learned in robot table tennis task. Top to bottom: Push, Slice, and Topspin.

strategies was 1.476 ± 0.051 . We then virtually manipulated the trajectories to indicate that the ball was unsuccessfully returned, which achieved an average and standard deviation for the task reward of only 1.284 ± 0.042 . A Friedman Test shows this result is statistically significant ($\chi^2(1, 9) = 9, p < 0.05$), providing support that our task reward is learning to identify successful returns, as unsuccessful trajectories should be associated with lower reward when compared to successful. A video of the robot returning an incoming ball can be found in our supplementary materials. In this real-world robot control task, we see that MSRD can successfully recover all three strategies' behaviors using human data.

7 FUTURE WORK

Our approach to teasing out demonstrator-specific rewards from a common, task reward might be able to tease out sub-optimality. Researchers have discussed models for human sub-optimality, e.g., ϵ -greedy approaches [5, 15]. However, there is an argument to be made for sub-optimality being reflected via satisficing heuristics [23, 36] or a person-specific augmentation of a task reward [39], which is analogous to MSRD's model. In future work, we aim to explore models of human decision-making towards being able to recover a demonstrator's latent objective from sub-optimal demonstration. We also plan to utilize different inference methods to allow for unlabelled demonstrations (i.e., to infer the strategy label).

8 CONCLUSION

We explored a new problem of apprenticeship via IRL when demonstrators provide heterogeneous, strategy-based demonstrations seeking to maximize a latent task reward. By explicitly modeling the task and strategy reward separately, and jointly learning both, our algorithm is able to recover more robust task rewards, discover unique strategy rewards, and imitate strategy-specific behaviors.

ACKNOWLEDGMENTS

We want to thank Dr. Sonia Chernova and the RAIL lab for letting us borrow their Sawyer robot and for Dr. Chernova's valuable feedback in revising our manuscript. This work was sponsored by ONR under grant N00014-18-S-B001, MIT Lincoln Laboratory grant 7000437192, and GaTech institute funding.

Figure 8: Evaluations of Strategy Rewards on Strategy Trajectories for Table Tennis. Each row is normalized to $[0, 1]$.

REFERENCES

- [1] Pieter Abbeel and Andrew Y Ng. 2004. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*. ACM, 1.
- [2] Kareem Amin, Nan Jiang, and Satinder Singh. 2017. Repeated Inverse Reinforcement Learning. In *Advances in Neural Information Processing Systems 30*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.). Curran Associates, Inc., 1815–1824. <http://papers.nips.cc/paper/6778-repeated-inverse-reinforcement-learning.pdf>
- [3] 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). <https://arxiv.org/pdf/1606.06565.pdf>
- [4] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. 2016. OpenAI Gym. *arXiv:arXiv:1606.01540*
- [5] Daniel S. Brown, Wonjoon Goo, and Scott Niekum. 2019. Ranking-Based Reward Extrapolation without Rankings. *CoRR abs/1907.03976* (2019). [arXiv:1907.03976](http://arxiv.org/abs/1907.03976)
- [6] Tsang-Hsiang Cheng, Chih-Ping Wei, and Vincent S. Tseng. 2006. Feature Selection for Medical Data Mining: Comparisons of Expert Judgment and Automatic Approaches. In *Proceedings of the 19th IEEE Symposium on Computer-Based Medical Systems (CBMS '06)*. IEEE Computer Society, Washington, DC, USA, 165–170. <https://doi.org/10.1109/CBMS.2006.87>
- [7] Jaedug Choi and Kee eung Kim. 2012. Nonparametric Bayesian Inverse Reinforcement Learning for Multiple Reward Functions. In *Advances in Neural Information Processing Systems 25*, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (Eds.). Curran Associates, Inc., 305–313. <http://papers.nips.cc/paper/4737-nonparametric-bayesian-inverse-reinforcement-learning-for-multiple-reward-functions.pdf>
- [8] Wojciech M. Czarnecki, Razvan Pascanu, Simon Osindero, Siddhant Jayakumar, Grzegorz Swirszcz, and Max Jaderberg. 2019. Distilling Policy Distillation. In *Proceedings of Machine Learning Research (Proceedings of Machine Learning Research)*, Kamalika Chaudhuri and Masashi Sugiyama (Eds.), Vol. 89. PMLR, 1331–1340. <http://proceedings.mlr.press/v89/czarnecki19a.html>
- [9] Nathaniel D Daw and Peter Dayan. 2014. The algorithmic anatomy of model-based evaluation. *Philosophical Transactions of the Royal Society B: Biological Sciences* 369, 1655 (2014), 20130478.
- [10] Christos Dimitrakakis and Constantin A. Rothkopf. 2012. Bayesian Multitask Inverse Reinforcement Learning (*EWRL '11*). Springer-Verlag, Berlin, Heidelberg, 273–284. https://doi.org/10.1007/978-3-642-29946-9_27
- [11] Benjamin Eysenbach, Abhishek Gupta, Julian Ibarz, and Sergey Levine. 2019. Diversity is All You Need: Learning Skills without a Reward Function. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=Sx63jRqFm>
- [12] Chelsea Finn, Sergey Levine, and Pieter Abbeel. 2016. Guided cost learning: Deep inverse optimal control via policy optimization. In *International Conference on Machine Learning*. 49–58.
- [13] Justin Fu, Katie Luo, and Sergey Levine. 2018. Learning Robust Rewards with Adversarial Inverse Reinforcement Learning. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=rkHywl-A->
- [14] Matthew Gombolay, Reed Jensen, Jessica Stigile, Toni Golen, Neel Shah, Sung-Hyun Son, and Julie Shah. 2018. Human-machine collaborative optimization via apprenticeship scheduling. *Journal of Artificial Intelligence Research* 63 (2018), 1–49.
- [15] Matthew Gombolay, Reed Jensen, Jessica Stigile, Sung-Hyun Son, and Julie Shah. 2016. Apprenticeship scheduling: learning to schedule from human experts. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*. AAAI Press, 826–833.
- [16] Matthew Gombolay, Xi Jessie Yang, Bradley Hayes, Nicole Seo, Zixi Liu, Samir Wadhwan, Tania Yu, Neel Shah, Toni Golen, and Julie Shah. 2018. Robotic assistance in the coordination of patient care. *The International Journal of Robotics Research* 37, 10 (2018), 1300–1316.
- [17] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. 2018. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Jennifer Dy and Andreas Krause (Eds.), Vol. 80. PMLR, Stockholm, Sweden, 1861–1870. <http://proceedings.mlr.press/v80/haarnoja18b.html>
- [18] Karol Hausman, Yevgen Chebotar, Stefan Schaal, Gaurav Sukhatme, and Joseph J Lim. 2017. Multi-Modal Imitation Learning from Unstructured Demonstrations using Generative Adversarial Nets. In *Advances in Neural Information Processing Systems 30*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.). Curran Associates, Inc., 1235–1245. <http://papers.nips.cc/paper/6723-multi-modal-imitation-learning-from-unstructured-demonstrations-using-generative-adversarial-nets.pdf>
- [19] Peter Henderson, Wei-Di Chang, Pierre-Luc Bacon, David Meger, Joelle Pineau, and Doina Precup. 2018. Optiongan: Learning joint reward-policy options using generative adversarial inverse reinforcement learning. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- [20] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. Distilling the Knowledge in a Neural Network. In *NIPS Deep Learning and Representation Learning Workshop*. <http://arxiv.org/abs/1503.02531>
- [21] 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>
- [22] Gary Klein. 1997. The recognition-primed decision (RPD) model: Looking back, looking forward. *Naturalistic decision making* (1997), 285–292.
- [23] Gary A Klein, Judith Ed Orasanu, Roberta Ed Calderwood, and Caroline E Zsombok. 1993. Decision making in action: Models and methods. In *This book is an outcome of a workshop held in Dayton, OH, Sep 25–27, 1989*. Ablex Publishing.
- [24] Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. 2016. End-to-end Training of Deep Visuomotor Policies. *J. Mach. Learn. Res.* 17, 1 (Jan. 2016), 1334–1373. <http://dl.acm.org/citation.cfm?id=2946645.2946684>
- [25] Yunzhu Li, Jiaming Song, and Stefano Ermon. 2017. InfoGAIL: Interpretable Imitation Learning from Visual Demonstrations. In *Advances in Neural Information Processing Systems 30*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.). Curran Associates, Inc., 3812–3822. <http://papers.nips.cc/paper/6971-infogail-interpretable-imitation-learning-from-visual-demonstrations.pdf>
- [26] Andrew Y Ng, Stuart J Russell, et al. 2000. Algorithms for inverse reinforcement learning. In *ICML*, Vol. 1. 2.
- [27] Stefanos Nikolaidis, Keren Gu, Ramya Ramakrishnan, and Julie A. Shah. 2014. Efficient Model Learning for Human-Robot Collaborative Tasks. *CoRR abs/1405.6341* (2014). [arXiv:1405.6341](http://arxiv.org/abs/1405.6341)
- [28] Stefanos Nikolaidis, Swaprava Nath, Ariel D. Procaccia, and Siddhartha Srinivasa. 2017. Game-Theoretic Modeling of Human Adaptation in Human-Robot Collaboration (*HRI '17*). ACM, New York, NY, USA, 323–331. <https://doi.org/10.1145/2909824.3020253>
- [29] Hema Raghavan, Omid Madani, and Rosie Jones. 2006. Active learning with feedback on features and instances. *Journal of Machine Learning Research* 7, Aug (2006), 1655–1686.
- [30] Deepak Ramachandran and Eyal Amir. 2007. Bayesian Inverse Reinforcement Learning. In *IJCAI*, Vol. 7. 2586–2591.
- [31] Nathan D Ratliff, J Andrew Bagnell, and Martin A Zinkevich. 2006. Maximum margin planning. In *Proceedings of the 23rd international conference on Machine learning*. ACM, 729–736.
- [32] Stephane Ross, Geoffrey Gordon, and Drew Bagnell. 2011. A Reduction of Imitation Learning and Structured Prediction to No-Regret Online Learning. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research)*, Geoffrey Gordon, David Dunson, and Miroslav Dudik (Eds.), Vol. 15. PMLR, Fort Lauderdale, FL, USA, 627–635. <http://proceedings.mlr.press/v15/ross11a.html>
- [33] Andrei A Rusu, Sergio Gomez Colmenarejo, Caglar Gulcehre, Guillaume Desjardins, James Kirkpatrick, Razvan Pascanu, Volodymyr Mnih, Koray Kavukcuoglu, and Raia Hadsell. 2015. Policy distillation. *arXiv preprint arXiv:1511.06295* (2015).
- [34] J Ben Schafer, Dan Frankowski, Jon Herlocker, and Shilad Sen. 2007. Collaborative filtering recommender systems. In *The adaptive web*. Springer, 291–324.
- [35] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal Policy Optimization Algorithms. *CoRR abs/1707.06347* (2017). [arXiv:1707.06347](http://arxiv.org/abs/1707.06347)
- [36] Herbert A Simon. 1972. Theories of bounded rationality. *Decision and organization* 1, 1 (1972), 161–176.
- [37] Yee Teh, Victor Bapst, Wojciech M. Czarnecki, John Quan, James Kirkpatrick, Raia Hadsell, Nicolas Heess, and Razvan Pascanu. 2017. Distral: Robust multitask reinforcement learning. In *Advances in Neural Information Processing Systems 30*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.). Curran Associates, Inc., 4496–4506. <http://papers.nips.cc/paper/7036-distral-robust-multitask-reinforcement-learning.pdf>
- [38] Emanuel Todorov, Tom Erez, and Yuval Tassa. 2012. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 5026–5033.
- [39] Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature* (2019), 1–5.
- [40] Bryce Woodworth, Francesco Ferrari, Teofilo E. Zosa, and Laurel D. Riek. 2018. Preference Learning in Assistive Robotics: Observational Repeated Inverse Reinforcement Learning. In *Proceedings of the 3rd Machine Learning for Healthcare Conference (Proceedings of Machine Learning Research)*, Finale Doshi-Velez, Jim Fackler, Ken Jung, David Kale, Rajesh Ranganath, Byron Wallace, and Jenna Wiens (Eds.), Vol. 85. PMLR, Palo Alto, California, 420–439. <http://proceedings.mlr.press/v85/woodworth18a.html>

- [41] Markus Wulfmeier, Peter Ondruska, and Ingmar Posner. 2015. Maximum entropy deep inverse reinforcement learning. *arXiv preprint arXiv:1507.04888* (2015).
- [42] Kelvin Xu, Ellis Ratner, Anca Dragan, Sergey Levine, and Chelsea Finn. 2019. Learning a Prior over Intent via Meta-Inverse Reinforcement Learning. In *Proceedings of the 36th International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.), Vol. 97. PMLR, Long Beach, California, USA, 6952–6962. <http://proceedings.mlr.press/v97/xu19d.html>
- [43] Yunbo Zhang, Wenhao Yu, and Greg Turk. 2019. Learning Novel Policies For Tasks. In *Proceedings of the 36th International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.), Vol. 97. PMLR, Long Beach, California, USA, 7483–7492. <http://proceedings.mlr.press/v97/zhang19q.html>
- [44] Brian D. Ziebart. 2010. *Modeling Purposeful Adaptive Behavior with the Principle of Maximum Causal Entropy*. Ph.D. Dissertation. Pittsburgh, PA, USA. Advisor(s) Bagnell, J. Andrew. AAI3438449.
- [45] Brian D. Ziebart, Andrew Maas, J. Andrew Bagnell, and Anind K. Dey. 2008. Maximum Entropy Inverse Reinforcement Learning (*AAAI'08*). AAAI Press, 1433–1438. <http://dl.acm.org/citation.cfm?id=1620270.1620297>