
Emergent Social Learning via Multi-agent Reinforcement Learning

Kamal Ndousse¹ Douglas Eck² Sergey Levine^{2,3} Natasha Jaques^{2,3}

Abstract

Social learning is a key component of human and animal intelligence. By taking cues from the behavior of experts in their environment, social learners can acquire sophisticated behavior and rapidly adapt to new circumstances. This paper investigates whether independent reinforcement learning (RL) agents in a multi-agent environment can learn to use social learning to improve their performance. We find that in most circumstances, vanilla model-free RL agents do not use social learning. We analyze the reasons for this deficiency, and show that by imposing constraints on the training environment and introducing a model-based auxiliary loss we are able to obtain generalized social learning policies which enable agents to: i) discover complex skills that are not learned from single-agent training, and ii) adapt online to novel environments by taking cues from experts present in the new environment. In contrast, agents trained with model-free RL or imitation learning generalize poorly and do not succeed in the transfer tasks. By mixing multi-agent and solo training, we can obtain agents that use social learning to gain skills that they can deploy when alone, even out-performing agents trained alone from the start.

1. Introduction

Social learning—learning by observing the behavior of other agents in the same environment—enables both humans and animals to discover useful behaviors that would be difficult to obtain through individual exploration, and to adapt rapidly to new circumstances (Henrich & McElreath, 2003; Laland, 2004). For example, fish are able to locate safe sources of food in new environments by observing where other members of their species congregate (Laland, 2004).

¹OpenAI, San Francisco, CA, USA ²Google Research – Brain team, Mountain View, CA, USA ³UC Berkeley, Berkeley, CA, USA. Correspondence to: Kamal Ndousse <kamal.ndousse@gmail.com>.

Humans are able to learn from one another without direct access to the experiences or memories of experts. Rather, the interactions that constitute social learning are mediated directly by the environment. A novice may gain the ability to solve a task merely by watching a self-interested expert complete the same task, and without an explicit incentive to mimic the expert. Beyond simple imitation, social learning allows individual humans to “stand on the shoulders of giants” and develop sophisticated behaviors or innovations that would be impossible to discover from scratch within one lifetime. In fact, social learning may be a central feature of humans’ unprecedented cognitive and technological development (Humphrey, 1976; Boyd et al., 2011).

Contemporary techniques for transmitting skills between agents in AI systems are much more limited. For example, Imitation Learning (IL) depends on manually curated expert demonstration trajectories. This requires a system extrinsic to the learning agent (*i.e.*, a developer) to (1) identify tasks of interest, (2) find and instrument task experts, and (3) collect demonstration data. The resulting agent is brittle, only able to replicate skills within the collected demonstrations, and otherwise vulnerable to compounding error (Ross et al., 2011; Reddy et al., 2019). In contrast, humans and animals capable of social learning are able to identify experts at tasks relevant to their own interests and learn from those experts without the intervention of an external system. In novel circumstances, social learning enables *adapting online* to the new task, which IL cannot do.

In this work, we investigate whether RL agents can *learn* to learn from the cues of other agents. We focus on two hypotheses related to the benefits of social learning: **H1**) Social learning can discover more **complex policies** that would be difficult to discover through costly individual exploration, and **H2**) Social learning can be used to **adapt online** to a new environment, when the social learning policy encodes how to learn from the cues of experts. We study a setting representative of real-world human environments: multi-agent reinforcement learning (MARL) where completely independent agents pursue their own rewards with no direct incentive to teach or learn from one another (decentralized training and decentralized execution, with no shared parameters or state). This setting is representative of real-world multi-agent problems such as autonomous driving, in which there are other cars on the road but the drivers

are not motivated to teach a learning agent, and may or may not have relevant expertise. Novice agents are able to see the experts in their partial observations, but must learn the social learning behavior that would allow them to make use of that information. In these challenging scenarios, we find that vanilla model-free RL agents fail to use social learning to improve their performance.

Since social learning does not occur by default, we thoroughly investigate the environment conditions and agent abilities which are needed for it to emerge, and propose both a training environment and model architecture that facilitate social learning.

We show that when social learning does occur, it enables agents to learn more complex behavior than agents trained alone and even the experts themselves.

Within each episode, social learners selectively follow experts only when doing so is rewarding — which pure imitation learning cannot do. To prevent social learners from becoming reliant on the presence of experts, we interleave training with experts and training alone. Social learners trained in this manner make use of what they learned from experts to improve their performance in the solo environment, even out-performing agents that were only ever trained solo. This demonstrates that agents can acquire skills through social learning that they could not discover alone, and which are beneficial even when experts are no longer present.

We show for the first time that social learning improves generalization for RL agents. When evaluated in novel environments with different experts, social learners use cues from the new experts to adapt online, achieving high zero-shot transfer performance. They significantly out-perform agents trained alone, agents that do not learn social learning, and imitation learning. Imitation learning leads to similar transfer performance as the original experts, whereas social learning enables agents to actively seek information about the new environment from a new set of experts and adjust their policy accordingly. The difference between imitation learning and social learning can best be understood using the following analogy: Teach someone to fish, and they will learn how to fish. Teach someone to acquire skills from other agents, and they can learn how to fish, hunt, swim, and any other skills they require.

The contributions of this work are to show how learning social learning—specifically, learning how to learn from other agents using RL—can lead to improved generalization to novel environments, reduce the cost of individual exploration, and enhance performance. We show how both algorithm and environment design are critical to facilitating social learning, and provide details for how to achieve it. Since deep RL often generalizes poorly (Cobbe et al.,

2019; Farebrother et al., 2018; Packer et al., 2018), we believe these results are a promising first step on a new path for enhancing generalization. In real-world tasks, such as autonomous driving or robotics, there are many human experts who know how to perform useful tasks, but individual exploration can be costly, inefficient, error-prone, or unsafe. Social learning could enable agents to adapt online using cues from these experts, without resorting to clumsy or costly exploration.

2. Related Work

There is a rich body of work on imitation learning, which focuses on training an agent to closely approximate the behavior of a curated set of expert trajectories (Pomerleau, 1989; Schaal, 1999; Billard et al., 2008; Argall et al., 2009). For example, Behavioral Cloning (e.g. Bain & Sammut (1995); Torabi et al. (2018)) uses supervised learning to imitate how the expert maps observations to actions. Because imitation learning can be brittle, it can also be combined with RL (e.g. Guo et al. (2019); Lerer & Peysakhovich (2019)), or make use of multiple experts (e.g. Cheng et al. (2020)). Third-person imitation learning (e.g. Shon et al. (2006); Calinon et al. (2007); Stadie et al. (2017)) enables agents to learn to copy an expert’s policy when observing the expert from a third-person perspective.

Our work shares some similarities with third-person imitation learning, in that our agents are never given access to expert’s observations and therefore only have a third-person view. However, we do not assume access to a curated set of expert trajectories. Rather, our agents co-exist in a multi-agent environment with other agents with varying degrees of expertise, who are not always within their partially observed field of view. We do not assume experts are incentivized to teach other agents (as in e.g. Omidshafiei et al. (2019); Christiano et al. (2017)). Since agents are not given privileged access to information about other agents, they do not explicitly model other agents (as in e.g. Albrecht & Stone (2018); Jaques et al. (2018)), but must simply treat them as part of the environment. Finally, unlike in imitation learning, we do not force our agents to copy the experts’ policy; in fact, we show that our agents eventually learn to out-perform the experts with which they learned.

Our work is distinct from Inverse RL (IRL) (e.g. Ng et al. (2000); Ramachandran & Amir (2007); Ziebart et al. (2008); Hadfield-Menell et al. (2016)), because our agents share a single environment with the experts which has a consistent reward function, and thus do not need to infer the reward function from expert trajectories.¹ However, we do consider learning from sub-optimal experts, which was studied by

¹Note that novices and experts do not share rewards; if an expert receives a reward the novice does not benefit.

Jacq et al. (2019) in the context of IRL.

Learning online from other agents has been studied in the context of autonomous driving, where Landolfi & Dragan (2018) trained cars which copied the policy of human drivers on the road when the agent had a high degree of uncertainty. However, they assume access to privileged information about other cars’ states and policies. Sun et al. (2019) make similar assumptions, but use other cars to update agents’ beliefs; e.g. about the presence of pedestrians in occluded regions.

Most closely related to our work is that of (Borsa et al., 2019), who train model-free RL agents in the presence of a scripted expert. In the partially observed exploration tasks they consider, they find that that novice agents learning from experts outperform solitary novices only when experts use information that is hidden from the novices. In our work, experts do not have access to privileged information about the environment; rather they are distinguished by their task skill. We analyze why it is difficult for RL agents to benefit from expert demonstrations in sparse reward environments, and propose a method to solve it. With our proposed learning algorithm, novices learning from unprivileged experts do outperform solitary learners. Unlike Borsa et al. (2019), we show for the first time that social learning allows agents to transfer effectively to unseen environments with new experts, resulting in better generalization than solo learners.

Finally, our approach uses a model-based auxiliary loss, which predicts the next state given the current state, to better enable agents to learn transition dynamics from trajectories in which they received no reward. The idea of using auxiliary losses in deep RL has been explored in a variety of works (e.g. Ke et al. (2019); Weber et al. (2017); Shelhamer et al. (2016); Jaderberg et al. (2016)). Model-based RL (MBRL) has also been applied to the multi-agent context (e.g. Krupnik et al. (2020)). Hernandez-Leal et al. (2019) use predicting the actions of other agents as an auxiliary task, and show that it improves performance on multi-agent tasks. However, they do not study social learning, transfer or generalization, and assume privileged access to other agents’ actions.

3. Learning Social Learning

We focus on Multi-Agent Partially Observable Markov Decision Process (MA-POMDP) environments defined by the tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \mathcal{I}, N \rangle$, where N is the number of agents. The shared environment state is $s \in \mathcal{S}$. \mathcal{I} is an inspection function that maps s to each agent’s partially observed view of the world, s^k . At each timestep t , each agent k chooses a discrete action $a_t^k \in \mathcal{A}$. Agents act simultaneously and there is no notion of turn-taking. Let \mathcal{A}^N be the joint action space. The transition function depends on the joint action

space: $\mathcal{T} : \mathcal{S} \times \mathcal{A}^N \times \mathcal{S} \rightarrow [0, 1]$. The reward function is the same across agents, but each agent receives its own individual reward $r_t^k = \mathcal{R}(s_t, a_t^k)$. Each agent k is attempting to selfishly maximize its own, individual reward by learning a policy π^k that optimizes the total expected discounted future reward: $J(\pi^k) = \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t r_{t+1}^k \mid s_0 \right]$, given starting state s_0 and discount factor $\gamma \in [0, 1]$. Note that agents are trained independently, *cannot directly observe other agents’ observations or actions, and do not share parameters*. To simplify notation, when we discuss the learning objectives for a single novice we forego the superscript notation.

3.1. Model-Free Sparse Reward Social Learning

To understand how experts present in the same environment could provide cues that enhance learning—and how model-free RL could fail to benefit from those cues—consider an example environment in which agents must first pick up a key, and use it to pass through a door to reach a goal. In this sparse-reward, hard exploration environment, agents only receive a reward for reaching the goal; the rest of the time, their reward is 0. Assume there is a boundless supply of keys, and after an agent passes through the door it closes — then other agents cannot modify the environment to make the task easier. Instead, an expert agent can demonstrate a novel state \tilde{s} that is difficult to produce through random exploration: opening the door. Ideally, we would like novice agents to learn from this demonstration by updating their internal representation to model \tilde{s} , and eventually learn how to produce this state.

However, if the novice does not receive any reward as a result of observing \tilde{s} , model-free RL receives little benefit from the expert’s behavior. Assume that the novice agent is learning policy π_θ with a policy-gradient objective on a trajectory including the demonstrated state $\tilde{s}_k, k \in (0, T - 1)$:

$$\nabla_\theta J(\theta) = \sum_{t=0}^{T-1} \nabla_\theta \log \pi_\theta(a_t | s_t) R_t \quad (1)$$

where $R_t = \sum_{t'=t+1}^T \gamma^{t'-t-1} r_{t'}$ is the total reward received over the course of the trajectory. If the agent receives 0 reward during the episode in which the demonstration occurred (e.g. it does not reach the goal), we can see that $\forall t \in (0, T), R_t = \sum_{t'=t+1}^T \gamma^{t'-t-1} (0) = 0$. Therefore $\nabla_\theta J(\theta) = \sum_{t=0}^{T-1} \nabla_\theta \log \pi_\theta(a_t | s_t) (0) = 0$, and the novice receives no gradients from the expert’s demonstration which allow it to update its policy.

Temporal Difference (TD) Learning could temporarily mitigate this issue, but as we show in the Appendix, this ability quickly deteriorates. Consider Q-learning, in which $Q(a, s) = \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t r_{t+1} \mid a, s \right]$ models the total ex-

pected reward from taking action a in state s . As the agent continues to receive 0 rewards during training, all Q values will be driven toward zero. Even if the agent observes a useful novel state such as \tilde{s} , as $Q(a, s) \rightarrow 0, \forall a \in \mathcal{A}, s \in \mathcal{S}$, the Q-learning objective becomes:

$$Q(\tilde{s}, a) = r + \gamma \max_{a'} Q(s', a') = 0 + \gamma 0 = 0 \quad (2)$$

Thus, the objective forces the value of $Q(\tilde{s}, a)$ to be zero. In this case, modeling transitions into or out of state \tilde{s} is not required in order to produce an output of zero, since all other Q-values are already zero. In both cases, we can see that before a model-free RL agent receives a reward from the environment, it will have difficulty modeling novel states or transition dynamics. This makes learning from the cues of experts particularly difficult.

Learning to model the expert’s policy $\pi^E(a^E|s^E)$ would likely help the novice improve performance on the task. However, the novice does not have explicit access to the expert’s states or actions. From its perspective, the other agent’s policy is simply a part of the environment transition dynamics. While the true state transition function $s_{t+1} = \mathcal{T}(s_t, a_t^N, a_t^E)$ depends on both the novice’s own action a_t^N , and the expert’s policy (since $a_t^E = \pi^E(s_t^E)$), the novice is only able to observe $p(s_{t+1}^N | s_t^N, a_t^N)$. Therefore, the novice can only obtain knowledge of the expert’s policy through correctly modeling the state transitions it observes. Since, as we have argued above, the novice will struggle to model state transitions in the absence of external reward, it will also have difficulty modeling another agent’s policy.

3.2. Social Learning with Auxiliary Losses

To mitigate this issue, we propose augmenting the novice agent with a model-based prediction loss. Specifically, we append additional layers θ_A to the policy network’s encoding of the current state, $f_\theta(s_t)$, as shown in Figure 1. We introduce an unsupervised mean absolute error (MAE) auxiliary loss to train the network to predict the next state s_{t+1} , given the current state s_t and action a_t :

$$\hat{s}_{t+1} = f_{\theta_A}(a_t, s_t); \quad J = \frac{1}{T} \sum_{t=0}^T |s_{t+1} - \hat{s}_{t+1}| \quad (3)$$

This architecture allows gradients from the auxiliary loss to contribute to improving $f_\theta(s_t)$. Figure 2 shows example state predictions generated by the auxiliary layers of a trained agent, demonstrating that this architecture enables effective learning of the transition dynamics.

We can now see that if the novel demonstration state is in a trajectory, $\tilde{s}_k \in (0, T)$, the term $|\tilde{s}_k - \hat{s}_k|$ will be part of the objective. It will not be 0 unless the agent learns to perfectly predict the novel demonstration state. Therefore, cues from the expert will provide gradients that allow the novice to




Figure 1. PPO + auxiliary loss deep neural network architecture. Convolution layers extract information about the state from pixels, which is fed into a Fully Connected (FC) layer, and a recurrent LSTM. The yellow shaded box shows the components of the model which learn the RL policy π and value function $V(s)$. The green shaded box shows the components of the model dedicated to computing the auxiliary loss, which predicts the next state from the current state.

improve its representation of the world, even if it does not receive any reward from the demonstration. This architecture also implicitly improves the agent’s ability to model other agents’ policies, since it must correctly predict other agents’ actions in order to accurately predict the next state. It is able to do this without ever being given explicit access to the other agent’s states or actions, or even any labels indicating what elements of the environment constitute other agents. We refer to our approach as social learning with an auxiliary predictive loss, and hypothesize that it will improve agents’ ability to learn from the cues of experts. Note, however, that simply modeling the other agent’s policy does not force these independent agents to copy the actions of other agents (unlike in imitation learning). When an independent social learner receives a low reward for following the policy of another agent, it can learn to avoid following that agent in future.

To optimize our agents, we test both TD-learning (deep Q-learning) and policy-gradient methods, and find that Proximal Policy Optimization (PPO) (Schulman et al., 2017) provides better performance and stability, and is most able to benefit from social learning. We use Generalized Advantage Estimation (GAE) (Schulman et al., 2016) to train the PPO value function. As shown in Figure 1, our agents use convolution layers to learn directly from a pixel representation of the state s_t . Because the environments under investigation are partially observed and non-Markov, we use a recurrent policy parameterized by a Long Short-Term Memory (LSTM) network (Hochreiter & Schmidhuber, 1997) to model the history of observations in each episode. LSTM hidden states are stored in the experience replay buffer. Following the recommendations of Andrychowicz et al. (2020), we recalculate the stored state advantage values used to compute the value function loss between mini-batch gradient updates. In addition, we recalculate and update the stored




Figure 2. Examples of future states \hat{s}_{t+1} predicted by the network given state s_t (“observed state”), and each of the possible movement actions a_t . Most predictions are highly accurate, indicating the network has learned to effectively model transition dynamics. The transition for the ‘do nothing’ action is less accurate because it is infrequently chosen by the agent.

hidden states as suggested by Kapturowski et al. (2018) for off-policy RL. A detailed description of the network architecture and hyperparameters are given in an appendix.

3.3. Social Learning Environments

Humans and animals are most likely to rely on social learning when individual learning is difficult or unsafe (Henrich & McElreath, 2003; Laland, 2004). Further, individuals prefer to learn from others that they perceive to be highly successful or competent, which is known as *prestige bias* (Jiménez & Mesoudi, 2019). Cues or signals associated with prestige have shown to be important to both human and animal social learning (Barkow et al., 1975; Horner et al., 2010).

Motivated by these two ideas, we introduce a novel environment specifically designed to encourage social learning by making individual exploration difficult and expensive, and introducing prestige cues. The code for both the environment and social learning agents is available at: <https://github.com/kandouss/marlgird>. In Goal Cycle (Figure 3a), agents are rewarded for navigating between several goal tiles in a certain order, and penalized for deviating from that order. The goal tiles are placed randomly and are visually indistinguishable, so it is not possible for an agent to identify the correct traversal order without potentially incurring an exploration penalty. When the penalty is large, this becomes a hard exploration task.

Agents can achieve high rewards by efficiently traversing between the goals in the correct order, and skilled agents are able to identify the correct order while incurring minimal penalties. Agents can accomplish this either by trial and error, or by observing other agents. In practice, since the behavior of other agents can be unpredictable and potentially non-stationary, agents more easily learn to solve the task directly through trial and error. But by adjusting the penalty for navigating to goals in the wrong order, we can penalize individual exploration and thereby encourage social learning. In all the Goal Cycle variants discussed here,

agents receive a reward of +1 for navigating to the first goal they encounter in an episode, and +1 for any navigating to any subsequent goals if in the correct order. They receive a penalty of -1.5 for navigating to the incorrect goal. In Appendix 7.2 we conduct experiments to assess the effect of the exploration penalty, and find that with no exploration penalty agents do not learn to rely on social learning.

Prestige cues are implemented in Goal Cycle signal through agents changing color as they collect rewards over the course of each episode. At time t , the color of an agent is a linear combination of the red and blue RGB color vectors: $\text{COLOR}_t = \text{BLUE} \cdot \tilde{c}_t + \text{RED} \cdot (1 - \tilde{c}_t)$, where $\tilde{c}_t = \text{sigmoid}(c_t)$, and c_t is a reward-dependent prestige value given by:

$$c_t = \begin{cases} \alpha_c c_{t-1} + r_t, & r_t \geq 0 \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

Thus the color of each agent changes from red to blue as it collects rewards (slowly decaying with constant α_c , but reverts to red if it incurs a penalty.

The Goal Cycle environments are challenging to master because there are many more incorrect than correct goal traversal sequences. When the penalty is large, the returns for most traversal orders are negative, so mistake-prone agents learn to avoid goal tiles altogether (after the first goal in an episode, which always gives a reward of +1). We can obtain agents that perform well in high penalty environments using a curriculum: we train them initially in low-penalty environments and gradually increase the penalty.

3.4. Social Learning in New Environments

A central thesis of this work is that social learning can help agents adapt more rapidly to novel environments. This could be particularly useful because deep RL often fails to generalize to even slight modifications of the training environment. Therefore, we test the zero-shot transfer performance of agents pre-trained in 3-Goal Cycle in three new environments. The first is Goal Cycle with 4 goals. This

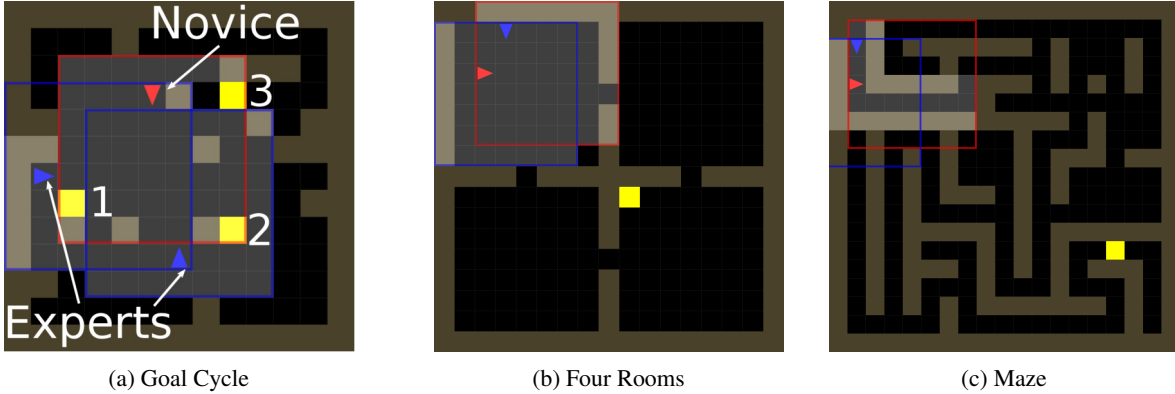


Figure 3. Goal Cycle (a) is a 13x13 environment in which the positions of goal tiles (yellow) and obstacles (brown) are randomized at the start of each episode. Agents receive rewards for repeatedly traversing the goal tiles in a certain order and penalties for deviations from the correct order. However, the correct order is not visible. Agents always receive a reward from the first goal tile they encounter in an episode. We test whether agents trained in Goal Cycle can generalize to Four Rooms (b), and Maze (c); larger environments with walls hiding the location of a single goal.

variant is significantly harder, because while there are two possible cycles between three goals, there are six ways to cycle between four goals (there are $(n-1)!$ distinct traversal orders, since the cycle can be entered at any goal). Thus, even optimal agents must incur higher penalties to identify the correct traversal order in each episode.

We also test transfer to the 17x17 Four Rooms (in Figure 3b), and 19x19 Maze environments (Figure 3c). In both environments, the goal and obstacles are randomly generated each episode, and agents must navigate around walls to find the goal within a time limit. These environments represent challenging transfer tasks, because navigating the maze is significantly more complex than navigation in Goal Cycle, and agents are unlikely to have encountered walls in Four Rooms while training in Goal Cycle (Dennis et al., 2020).

4. Experiments

To establish that independent agents with the auxiliary predictive loss described in Sec. 3.2 can use social learning to improve their own performance, we compare the performance of agents with auxiliary predictive losses learning in an environment shared with experts (SOCIAL PPO + AUX PRED) to that of agents with the same architecture but trained alone (Solo PPO + aux pred).

We also compare social learning performance (in the presence of experts) across two ablations. The first uses the same architecture with no auxiliary loss, which corresponds to standard PPO. The second replaces the next-state-prediction objective of Equation 3 with a simple autoencoder reconstruction objective that predicts the *current* state, $\hat{s}_t = f_{\theta_A}(a_t, s_t)$. The loss is thus $J = \frac{1}{T} \sum_{t=0}^T |s_t - \hat{s}_t|$. We refer to this as social learning with an auxiliary reconstruction loss (Solo PPO + aux rec). Unlike the auxiliary

prediction loss, this reconstruction loss does not need to learn about transition dynamics or other agents’ policies. Finally, we compare to imitation learning (IL), which is an upper bound on the performance that can be obtained when learning from other agents with direct access to their ground truth state and action information (which is not provided to the other methods).

4.1. H1: Learning Social Learning

Figure 4 compares Goal Cycle performance of agents trained alone (solo) to agents trained in the presence of experts (social). Solo agents, even with the predictive auxiliary loss (SOLO PPO + AUX PRED), are not able to discover the strategy of reaching the goals in the correct order; instead, all random seeds converged to a strategy of stopping after reaching a single goal (with a maximum return of 1). Even when trained in the presence of experts, agents without an auxiliary loss (SOCIAL PPO (VANILLA)) failed in the same way, with all seeds receiving a maximum reward of 1. This supports our hypothesis that model-free RL by itself has difficulty benefiting from the cues of experts in sparse reward environments.

In contrast, social agents trained with auxiliary predictive or reconstructive losses were able to achieve higher performance. This confirms hypothesis **H1**, that when successful, social learning can enable agents to learn more complex behavior than when learning alone. While the reconstruction loss helps agents learn from experts, still higher performance is obtained with the model-based prediction loss. In contrast to the reconstruction loss, the prediction loss helps agents learn representations that capture the transition dynamics and (implicitly) other agents’ policies. Note that in our experiments, solo agents with auxiliary predictive losses

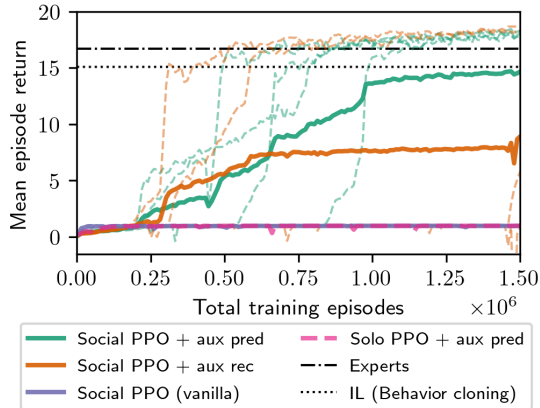


Figure 4. While learning in the presence of experts, only agents with an auxiliary loss that enables them to model expert cues (with either a predictive (pred) or reconstructive (rec) auxiliary loss) succeed in using social learning to improve performance. None of the seeds for solo agents are able to solve the task. Further, no vanilla PPO seeds solve the task, showing that social learning does not occur automatically with RL. In contrast, 4 of 5 seeds with an auxiliary reconstructive loss (SOCIAL PPO + AUX REC) are able to exceed the performance of the experts present in their environment. Faded dotted lines show the performance of individual random seeds and bolded lines show the mean of 5 seeds. The final performance is bimodal, with some novice seeds achieving expert-level performance and others failing to learn entirely. Since the normality assumption is violated, we refrain from using confidence intervals.

performed much worse than social learners, showing that good performance can depend on learning effectively from expert cues in sparse reward hard-exploration environments.

Interestingly, the majority of random seeds with predictive losses are actually able to exceed the performance of both the experts and the imitation learning (IL) methods. Unlike IL, which requires direct access to experts’ states and actions, social learners reach equivalent performance only through learning from their partially observed view of experts via RL. The reason that social learners can actually exceed the performance of the experts themselves is that they learn an improved policy that leverages social information. During each episode, social learners wait and observe while the expert explores the goals to find the correct order. Only then do they follow the expert’s cues. Thus, they allow the expert to incur the exploration penalties of Goal Cycle and use social learning to follow when it is safe. Despite being a non-cooperative strategy, this is reminiscent of social learning in the animal kingdom (Laland, 2018).

Learning from Sub-Optimal Experts. In real-world environments, not all other agents will have optimal performance. Therefore, we test whether it is still possible to learn from imperfect experts. The results in Figure 5 show that while expert performance does affect the novices’ final

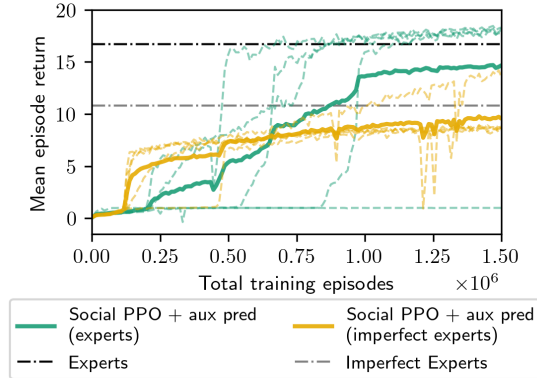


Figure 5. The effect of expert skill on social learning. With auxiliary predictive losses, novice agents benefit from social learning in the presence of near-optimal experts (green) or imperfect experts (yellow), surpassing the performance of solo agents in both cases. However, agents learning from imperfect experts achieve lower performance, and only 1/5 random seeds exceed the performance of the imperfect experts.

performance, even when learning from imperfect experts novices can use social learning to exceed the performance of agents trained alone. However, novices trained with optimal experts more frequently learn to exceed the performance of the experts with which they are trained.

4.2. H2: Adapting Online to New Environments

We hypothesize that using social learning to adapt online within an episode will allow agents to achieve higher performance when transferred to a new environment with experts. To evaluate this hypothesis, we investigate how well agents pre-trained in 3-Goal Cycle can generalize to three new environments: a 4 goal variant of the Goal Cycle environment, Four Rooms, and Maze. Figure 6 shows the zero-shot transfer performance of agents trained in 3-Goal Cycle to each new environment. Agents trained alone (which have not learned to rely on social learning) perform poorly, likely because they were never able to obtain good exploration policies during training, due to the high cost of exploration in the training environment.

Imitation learning agents perform poorly across all three transfer environments, under-performing the original 3-Goal experts in every case. These results demonstrate the brittleness of imitation learning, which can only fit the expert trajectories, and generalizes poorly to states outside of the training distribution (Ross et al., 2011; Reddy et al., 2019). In contrast, agents trained with social PPO + aux pred loss employ a generalized social learning policy that encodes how to adapt online based on expert cues. Thus, these agents use social learning to achieve high performance in the new environments with a new set of experts. For example, in the Maze environment they can follow the cues of the Maze ex-

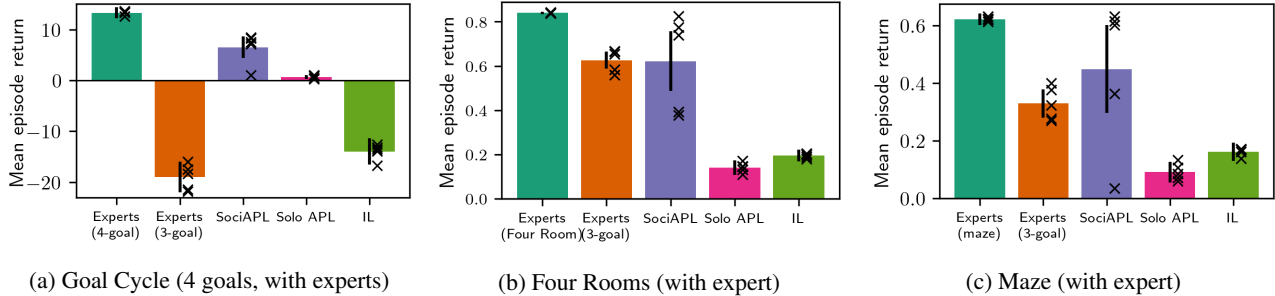


Figure 6. Zero-shot transfer performance in novel environments with expert demonstrators. Each bar shows the mean of the multiple seeds, with the performance of each seed shown as a black ‘x’. The experts and 3-goal solo learners were also trained with auxiliary predictive losses. Error bars show 95% confidence intervals for the estimated means. Agents which have learned to rely on social learning benefit from the cues of novel experts in the new environment., and easily out-perform agents trained alone, agents training with imitation learning, and in some cases, the experts from which they originally learned. Videos: tinyurl.com/SociAPL

perts to more closely approximate the optimal policy, rather than continuing to rely on learning from 3-Goal experts.

It is worth noting that the original 3-Goal experts transfer poorly to both 4-Goal Cycle and Maze. In 4-Goal, they receive large negative scores, repeatedly incurring a penalty for navigating to the goals in the incorrect order. This illustrates the brittleness of an RL policy which is overfit to the training environment. In both environments, social learning agents achieve better performance on the transfer task than the original experts, suggesting social learning may be a generally beneficial strategy for improving transfer in RL. Taken together, these results provide support for hypothesis **H2**, and suggest social learning can enable RL agents to adapt online (within an episode) to new environments by taking cues from other agents.

4.3. Social Learning Helps Solo Learning

A potential downside to learning to rely on the cues of experts is that agents could fail to perform well when experts are removed from the environment. For novice agents always trained with experts, there appears to be a trade-off between social learning and individual learning. As shown in Figure 7, novice SOCIAL PPO + AUX PRED agents initially learn to solve the Goal Cycle task with individual exploration, but eventually they overfit to the presence of experts and become reliant on expert cues to the detriment of solo performance. We can see this because the performance in the 3-goal environment in the presence of experts increases (7a), while the performance in the 3-goal environment when alone drops (7b). Observing the agent’s behavior reveals that it has learned to follow the cues of experts when they are present in the environment, but refrain from individual exploration when experts are not present. Given the high cost of individual exploration in this environment, this a safe but conservative strategy.

We find that we can improve the solo performance of social

learning agents by changing the distribution of training tasks to include episodes in which the experts are not present. The green curves in Figure 7 show the performance of a checkpoint switched from entirely social training environments to a mix of solo and social environments as training proceeds. The agent is able to retain good performance in solo 3-goal environments as well as 4-goal environments with experts, indicating that it is learning to opportunistically take advantage of expert cues while building individual expertise in the 3-goal environment. In fact, agents trained with SOCIAL PPO + AUX PRED perform better in solo transfer environments than agents exclusively trained in the solo environment (see Figure 7b), because agents trained alone never discover the optimal strategy due to the difficulty of exploring in this environment. This demonstrates that not only does social learning enable agents to discover skills that they could not learn by themselves (again supporting hypothesis **H1**), but that they are able to retain this knowledge to improve their individual performance even when experts are not present.

5. Conclusions

In this work we investigated how model-free deep RL can benefit from the presence of expert agents in the environment. We design an environment to elicit social learning by increasing the cost of individual exploration and introducing prestige cues. However, we find that even in this environment, model-free agents fail to use social learning to reach optimal performance. By adding a model-based auxiliary loss, which requires implicitly modeling other agents’ policies, we are able to train agents to use social learning. When deployed to novel environments, these agents use their social learning policies to adapt online using the cues of novel experts, and perform well in zero-shot transfer tasks. In contrast, agents trained alone or with imitation learning cannot adapt to the new task. Further, by mixing social and solo




Figure 7. ORANGE LINES: performance throughout training for one of the SOCIAL PPO + AUX PRED seeds from Figure 4 evaluated in three Goal Cycle environments: (a) the training environment, 3-goal Goal Cycle with experts; (b) a solo version of the training environment (no experts); and (c) 4-Goal Cycle, with experts. The agent initially relies mostly on individual exploration as indicated by good zero-shot transfer performance in the solo environment. After about 500k episodes the agent becomes reliant on cues from the 3-Goal experts, and its performance in solo 3-goal Goal Cycle degrades. **GREEN LINES:** performance of the same agent if the distribution of training environments is changed from 100% social to 75% solo and 25% social after about 500k episodes. As training proceeds the agent retains the capacity to solve the solo 3-goal environment while learning to use cues from expert behavior when they are available. The performance of this agent in the solo environment actually exceeds that of agents trained exclusively in the solo environment. The agent retains the capacity to use social learning when possible, with performance in the social training environment (a) and 4-goal Goal Cycle with experts (c) comparable to the orange agent (for which 100% of training episodes included experts).

training, we obtain social learning agents that actually have higher performance in the solo environment than agents trained exclusively in the solo environment. Our results demonstrate that social learning not only enables agents to learn complex behavior that they do not discover when trained alone, but that it can allow agents to generalize better when transferred to novel environments.

5.1. Limitations and Future Work

Our social learning experiments focus on exploratory navigation tasks. Future work includes extending this to other domains such as manipulation, to scenarios in which experts pursue different goals than the novices, and to scenarios with multiple experts that employ a variety of strategies. We found that training in a mixture of social and solitary environments can permit novice social learning agents to develop effective strategies for both social and individual task variants, and notably that the resulting individual skill far exceeds that of a solitary novice. However, in this work we do not thoroughly explore different strategies for augmenting solitary with social experience. Further research could clarify the circumstances in which adding social experiences could aid solitary task performance, and see the development of algorithms to facilitate this for arbitrary tasks. Finally, we assume that prestige cues which indicate how well experts perform in the environment are visible to social learners. Future research could investigate how to accurately predict another agent’s performance when prestige cues are not available.

6. Acknowledgements

We are grateful to the OpenAI Scholars program for facilitating the collaboration that led to this paper. We would also like to thank Ken Stanley, Joel Lehman, Bowen Baker, Yi Wu, Aleksandra Faust, Max Kleiman-Weiner, and DJ Strouse for helpful discussions related to this work, and Alethea Power for providing compute resources that facilitated our experiments.

References

- Albrecht, S. V. and Stone, P. Autonomous agents modelling other agents: A comprehensive survey and open problems. *Artificial Intelligence*, 258:66–95, 2018.
- Andrychowicz, M., Raichuk, A., Stańczyk, P., Orsini, M., Girgin, S., Marinier, R., Hussenot, L., Geist, M., Pietquin, O., Michalski, M., Gelly, S., and Bachem, O. What matters in on-policy reinforcement learning? a large-scale empirical study. *arXiv preprint arXiv:2006.05990*, 2020.
- Argall, B. D., Chernova, S., Veloso, M., and Browning, B. A survey of robot learning from demonstration. *Robotics and autonomous systems*, 57(5):469–483, 2009.
- Bain, M. and Sammut, C. A framework for behavioural cloning. In *Machine Intelligence 15*, pp. 103–129, 1995.
- Barkow, J. H., Akiwowo, A. A., Barua, T. K., Chance, M., Chapple, E. D., Chattopadhyay, G. P., Freedman, D. G., Geddes, W., Goswami, B., Isichei, P., et al. Prestige and culture: a biosocial interpretation [and comments and replies]. *Current Anthropology*, 16(4):553–572, 1975.

- Biewald, L. Experiment tracking with weights and biases, 2020. URL <https://www.wandb.com/>. Software available from wandb.com.
- Billard, A., Calinon, S., Dillmann, R., and Schaal, S. Survey: Robot programming by demonstration. *Handbook of robotics*, 59(BOOK_CHAP), 2008.
- Borsa, D., Heess, N., Piot, B., Liu, S., Hasenclever, L., Munos, R., and Pietquin, O. Observational learning by reinforcement learning. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pp. 1117–1124. International Foundation for Autonomous Agents and Multiagent Systems, 2019.
- Boyd, R., Richerson, P. J., and Henrich, J. The cultural niche: Why social learning is essential for human adaptation. *Proceedings of the National Academy of Sciences*, 108(Supplement_2):10918–10925, June 2011.
- Calinon, S., Guenter, F., and Billard, A. On learning, representing, and generalizing a task in a humanoid robot. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 37(2):286–298, 2007.
- Cheng, C.-A., Kolobov, A., and Agarwal, A. Policy improvement from multiple experts. *arXiv preprint arXiv:2007.00795*, 2020.
- Chevalier-Boisvert, M., Willems, L., and Pal, S. Minimalistic gridworld environment for openai gym. <https://github.com/maximecb/gym-minigrid>, 2018.
- Christiano, P. F., Leike, J., Brown, T., Martic, M., Legg, S., and Amodei, D. Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems*, pp. 4299–4307, 2017.
- Cobbe, K., Klimov, O., Hesse, C., Kim, T., and Schulman, J. Quantifying generalization in reinforcement learning. In *International Conference on Machine Learning*, pp. 1282–1289, 2019.
- Dennis, M., Jaques, N., Vinitsky, E., Bayen, A., Russell, S., Critch, A., and Levine, S. Emergent complexity and zero-shot transfer via unsupervised environment design. *arXiv preprint arXiv:2012.02096*, 2020.
- Farebrother, J., Machado, M. C., and Bowling, M. Generalization and regularization in dqn. *arXiv preprint arXiv:1810.00123*, 2018.
- Guo, X., Chang, S., Yu, M., Tesauro, G., and Campbell, M. Hybrid reinforcement learning with expert state sequences. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 3739–3746, 2019.
- Hadfield-Menell, D., Russell, S. J., Abbeel, P., and Dragan, A. Cooperative inverse reinforcement learning. In *Advances in neural information processing systems*, pp. 3909–3917, 2016.
- Henrich, J. and McElreath, R. The evolution of cultural evolution. *Evolutionary Anthropology: Issues, News, and Reviews*, 12(3):123–135, May 2003.
- Hernandez-Leal, P., Kartal, B., and Taylor, M. E. Agent modeling as auxiliary task for deep reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 15, pp. 31–37, 2019.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Horner, V., Proctor, D., Bonnie, K. E., Whiten, A., and de Waal, F. B. Prestige affects cultural learning in chimpanzees. *PloS one*, 5(5):e10625, 2010.
- Humphrey, N. K. The social function of intellect. In *Growing points in ethology*, pp. 303–317. Cambridge University Press, 1976.
- Jacq, A., Geist, M., Paiva, A., and Pietquin, O. Learning from a learner. In *International Conference on Machine Learning*, pp. 2990–2999, 2019.
- Jaderberg, M., Mnih, V., Czarnecki, W. M., Schaul, T., Leibo, J. Z., Silver, D., and Kavukcuoglu, K. Reinforcement learning with unsupervised auxiliary tasks. *arXiv preprint arXiv:1611.05397*, 2016.
- Jaques, N., Lazaridou, A., Hughes, E., Gulcehre, C., Ortega, P. A., Strouse, D., Leibo, J. Z., and de Freitas, N. Intrinsic social motivation via causal influence in multi-agent rl. *arXiv preprint arXiv:1810.08647*, 2018.
- Jiménez, Á. V. and Mesoudi, A. Prestige-biased social learning: current evidence and outstanding questions. *Palgrave Communications*, 5(1):1–12, 2019.
- Kapturowski, S., Ostrovski, G., Quan, J., Munos, R., and Dabney, W. Recurrent experience replay in distributed reinforcement learning. In *International conference on learning representations*, 2018.
- Ke, N. R., Singh, A., Touati, A., Goyal, A., Bengio, Y., Parikh, D., and Batra, D. Learning dynamics model in reinforcement learning by incorporating the long term future. *arXiv preprint arXiv:1903.01599*, 2019.
- Kingma, D. and Ba, J. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 12 2014.

- Krupnik, O., Mordatch, I., and Tamar, A. Multi-agent reinforcement learning with multi-step generative models. In *Conference on Robot Learning*, pp. 776–790, 2020.
- Laland, K. N. Social learning strategies. *Animal Learning & Behavior*, 32(1):4–14, February 2004.
- Laland, K. N. *Darwin’s unfinished symphony: How culture made the human mind*. Princeton University Press, 2018.
- Landolfi, N. C. and Dragan, A. D. Social cohesion in autonomous driving. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 8118–8125. IEEE, 2018.
- Lerer, A. and Peysakhovich, A. Learning existing social conventions via observationally augmented self-play. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 107–114, 2019.
- Ng, A. Y., Russell, S. J., et al. Algorithms for inverse reinforcement learning. In *Icml*, volume 1, pp. 2, 2000.
- Omidshafiei, S., Kim, D.-K., Liu, M., Tesauro, G., Riemer, M., Amato, C., Campbell, M., and How, J. P. Learning to teach in cooperative multiagent reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 6128–6136, 2019.
- Packer, C., Gao, K., Kos, J., Krähenbühl, P., Koltun, V., and Song, D. Assessing generalization in deep reinforcement learning. *arXiv preprint arXiv:1810.12282*, 2018.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style, high-performance deep learning library. In Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32*, pp. 8024–8035. Curran Associates, Inc., 2019.
- Pomerleau, D. A. Alvin: An autonomous land vehicle in a neural network. In *Advances in neural information processing systems*, pp. 305–313, 1989.
- Ramachandran, D. and Amir, E. Bayesian inverse reinforcement learning. In *IJCAI*, volume 7, pp. 2586–2591, 2007.
- Reddy, S., Dragan, A. D., and Levine, S. Sqil: Imitation learning via reinforcement learning with sparse rewards. *arXiv preprint arXiv:1905.11108*, 2019.
- Ross, S., Gordon, G., and Bagnell, D. 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*, pp. 627–635. JMLR Workshop and Conference Proceedings, 2011.
- Schaal, S. Is imitation learning the route to humanoid robots? *Trends in cognitive sciences*, 3(6):233–242, 1999.
- Schulman, J., Moritz, P., Levine, S., Jordan, M., and Abbeel, P. High-dimensional continuous control using generalized advantage estimation. In *International Conference on Learning Representations (ICLR)*, 2016.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Shelhamer, E., Mahmoudieh, P., Argus, M., and Darrell, T. Loss is its own reward: Self-supervision for reinforcement learning. *arXiv preprint arXiv:1612.07307*, 2016.
- Shon, A., Grochow, K., Hertzmann, A., and Rao, R. P. Learning shared latent structure for image synthesis and robotic imitation. In *Advances in neural information processing systems*, pp. 1233–1240, 2006.
- Stadie, B. C., Abbeel, P., and Sutskever, I. Third-person imitation learning. *arXiv preprint arXiv:1703.01703*, 2017.
- Sun, L., Zhan, W., Chan, C.-Y., and Tomizuka, M. Behavior planning of autonomous cars with social perception. In *2019 IEEE Intelligent Vehicles Symposium (IV)*, pp. 207–213. IEEE, 2019.
- Torabi, F., Warnell, G., and Stone, P. Behavioral cloning from observation. *arXiv preprint arXiv:1805.01954*, 2018.
- Weber, T., Racanière, S., Reichert, D. P., Buesing, L., Guez, A., Rezende, D. J., Badia, A. P., Vinyals, O., Heess, N., Li, Y., et al. Imagination-augmented agents for deep reinforcement learning. *arXiv preprint arXiv:1707.06203*, 2017.
- Ziebart, B. D., Maas, A. L., Bagnell, J. A., and Dey, A. K. Maximum entropy inverse reinforcement learning. In *Aaai*, volume 8, pp. 1433–1438. Chicago, IL, USA, 2008.

7. Appendix

7.1. SociAPL auxiliary prediction loss

Figure 8 shows the training curve for the next-state auxiliary prediction loss of Equation 3, for the five SociAPL agent seeds plotted in Figure 4. The curve shows that the agent is effectively learning to predict the next state with low mean absolute error. However, because the agent’s policy is changing at the same time, the prediction problem is non-stationary, which means that the loss does not always decrease. If the agent discovers a new behavior, the model will be required to predict new state transitions not previously experienced.




Figure 8. Next-state auxiliary prediction loss (in Mean Absolute Error (MAE)) over the course of training the SociAPL agents shown in Figure 4.

7.2. Training without an exploration penalty

The Goal Cycle environments include a penalty for reaching the goals in the incorrect order. We conducted an additional experiment to investigate whether social learning emerges effectively without the exploration penalty. Figure 9 compares the transfer results of agents trained in Goal Cycle with the standard exploration penalty of $p = -1.5$, to agents trained with no penalty ($p = 0$). Without the penalty, agents in Goal Cycle lack an incentive to learn from experts in the environment. The solitary policies they learn mirror those of the three-goal experts, reflected by similar performance in the four-goal transfer environment. This mirrors the finding that animals engage in social learning when individual exploration is most costly or dangerous (Laland, 2004).

7.3. Identifying the experts

In most of our experiments, we assume that social learners are trained in the presence of two high-performing expert agents. We conduct an additional experiment to determine if social learners can learn effectively in the presence of both high-performing experts, and non-experts. In this new




Figure 9. Next-state auxiliary prediction loss (in Mean Absolute Error (MAE)) over the course of training the SociAPL agents shown in Figure 4.

experiment, the transfer task experts are accompanied by two non-expert distractor agents that take random movement actions. As shown by the maze transfer results in Figure 10, our method is robust to the presence of these distractions and still outperforms the baselines.




Figure 10. Distractors

7.4. Imitation learning details

The imitation learning agents shown in Section 4 were trained with behavior cloning (Bain & Sammut, 1995). The model architecture is similar to that of the SociAPL agents, but layers specific to the value function and next state prediction were unused. The policy parameters θ were optimized with Adam (Kingma & Ba, 2014) to minimize the loss

$$L_{BC} = \mathbb{E}_D - \log \pi_{\theta}(a_{expert} | s_{demo}),$$

where D is a dataset of expert trajectories consisting of states s_{demo} and expert actions a_{expert} . Figure 11 shows the returns of imitation learning agents during training.

7.5. Q-learning learning in sparse reward environments

Before a Temporal Difference (TD) learning agent has received any reward, it will be difficult for it to learn




Figure 11. Returns over the course of training imitation learning agents in the 3 goal Goal Cycle environment with experts.

to model transition dynamics. Consider as an example deep Q-learning, in which the Q-function is parameterized by a neural network which encodes the state using a function $f_\theta(s)$. Assume the network is randomly initialized such that all Q-values are small, random values; *i.e.* $\forall a \in \mathcal{A}, s \in \mathcal{S}, Q(a, s) = \epsilon \sim \mathcal{N}(0, 0.1)$. Assume that the agent has not yet learned to navigate to the goal, and has received zero rewards so far during training. Therefore, when the agent observes the experience $(s, a, r = 0, s')$, the Q-learning objective is:

$$J(\theta) = (r + \gamma \max_{a'} Q(s', a') - Q(s, a))^2 \quad (5)$$

$$= (0 + \gamma \epsilon_1 - \epsilon_2)^2 \quad (6)$$

In effect, this induces a correlation between $Q(s', a')$ and $Q(s, a)$, and consequently $f_\theta(s')$ and $f_\theta(s)$, as a result of observing the state transition $s \rightarrow s'$. However, as the agent continues to receive zero reward, all Q-values will be driven toward zero. Once this occurs, even if the agent observes a useful novel state such as \tilde{s} , our equation becomes:

$$(r + \gamma \max_{a'} Q(s', a') - Q(\tilde{s}, a))^2 = (0 - Q(\tilde{s}, a))^2 \quad (7)$$

such that the objective forces the value of $Q(\tilde{s}, a)$ to be zero. In this case, modeling transitions into or out of state \tilde{s} is not required in order to produce an output of zero, since all other Q-values are already zero.

7.6. Environment details

The environments used in this paper were originally based on Minigrid (Chevalier-Boisvert et al., 2018). The partial states constituting agent observations are 27×27 RGB images corresponding to 7×7 grid tiles. There are 7 possible actions, though only three actions (rotate left, rotate right, move forward) are relevant in the experiments discussed in this paper. Agents are unable to see or move through the obstructive tiles that clutter their environments, and obstructed regions of their partial views appear as purple. However, agents can both see and move through goal tiles as well as other agents.




Figure 12. We use a penalty curriculum to obtain experts in environments where exploration is expensive. In the scenario visualized here, an agent trained for 81920 episodes in Goal Cycle environments with penalty of 0.5, then continued with a penalty of 1.0




Figure 13. Sample 19x19 maze environment.

When the penalty for individual exploration in Goal Cycle environments is large, agents are unable to learn effective strategies. We used a penalty curriculum to obtain experts for such tasks as shown in Figure 12.

7.7. Network architecture details

The value, policy, and auxiliary task networks share three convolution layers, a fully connected layer, and an LSTM layer. Values and policies are computed with two fully connected layers, and the prediction-based auxiliary branch has a fully connected layer followed by transposed convolution layers that mirror the input convolution layers. The convolution and transposed convolution use leaky ReLU activation functions; all other layers use tanh activation functions.

- Shared input layers:
 - Conv (3, 32) 3x3 filters, stride 3, padding 0,
 - Conv (32, 64) 3x3 filters, stride 1, padding 0,
 - Conv (64, 64) 3x3 filters, stride 1, padding 0,
 - FC (576, 192),
 - LSTM (192, 192)
- Value MLP:
 - FC (192, 64),
 - FC (64, 64),
 - FC (64, 1)
- Policy MLP:
 - FC (192, 64),
 - FC (64, 64),
 - FC (64, 7)
- Auxiliary state prediction layers:
 - FC (192 + 7, 576),
 - DeConv (64, 64) 3x3 filters, stride 1, padding 0,
 - DeConv (64, 32) 3x3 filters, stride 1, padding 0,
 - DeConv (32, 3) 3x3 filters, stride 3, padding 0

Altogether there are 668555 parameters. We experimented with smaller (357291-parameter) networks but did not observe a significant performance difference. The networks were sized to roughly saturate the available (desktop) compute resources.

7.8. Hyperparameters

Each agent uses a single Adam optimizer (Kingma & Ba, 2014) to update its parameters. Each of the novice agents was trained with a learning rate of $1e-4$. For SociAPL, the expert agents were trained with a learning rate of $1e-5$. Weights for all agents in the generalization experiments as well as the imperfect experts in SociAPL-IE were kept frozen and not updated.

Each parameter update consists of 20 sequential mini-batch updates with the same batch of rollouts (128 episodes). Each mini-batch consists of a uniform random sample of trajectories from that batch. Hidden states are stored alongside the trajectories, and the initial hidden state for each mini-batch trajectory is retrieved from these stored values. Hidden states and advantage values for the entire batch are re-calculated every 2 mini-batches. The mini-batch iteration ceases if $KL(\pi, \pi^{rollout})$ exceeds a target of 0.01. If for any mini-batch the estimated divergence exceeds a hard limit of 0.03, the update terminates and all changes to the network parameters and optimizer state are reverted.

batch size	128 episodes
mini-batches per batch	20
mini-batch num trajectories	512
mini-batch trajectory length	16
hidden state/advantage update interval	2 minibatches
return discount γ	0.993
GAE- λ	0.97
PPO clip ratio	0.2
KL target	0.01
KL hard limit	0.03

For each mini-batch iteration, the loss used to update agent parameters is

$$L^{\text{total}} = L^{\pi}(\theta) + c_V \cdot L^V(\theta) + c_{\text{aux}} \cdot L^{\text{aux}}(\theta) - c_{\text{ent}} \cdot L^{\text{ent}}(\theta),$$

where the policy loss L^{π} is computed with PPO-clip (Schulman et al., 2017) and GAE (Schulman et al., 2016), the value loss L^V is the mean squared error of the values estimated for each step in the trajectory, $L^{\text{ent}}(\theta)$ is the policy entropy bonus, and $L^{\text{aux}}(\theta)$ is the auxiliary prediction loss (Equation 3). The loss scaling coefficients used in our experiments are $c_V = 0.1$, $c_{\text{ent}} = 1e-5$, and $c_{\text{aux}} = 3$.

The prestige decay constant α_c used for the Goal Cycle environments (i.e. Equation 4) was 0.99.

In general we sought hyperparameters that enable stable training. We experimented with mini-batch sizes varying from 32 to 1025 trajectories and found training to be more stable with larger mini-batches. Training was less stable with learning rates higher than $1e-4$.

We randomized seeds for both the network parameter initialization and environment generation for each trial of each experiment.

7.9. Compute

The experiments in this paper were performed primarily on a desktop computer with an AMD Ryzen 3950x CPU and two Nvidia GTX 1080TI GPUs, as well as g4dn.8xlarge instances provisioned on Amazon AWS. Either system can run two or three trials simultaneously, each consisting of three agents training together in a shared environment. Collecting experience and updating parameters were comparably time consuming, and a single 1.5M episode (375M step) 3-agent training run took about 30 hours. We used Ubuntu 18.04 with python3.8 and all neural networks are implemented in PyTorch v1.6 (Paszke et al., 2019). Training metrics were logged with Weights and Biases (Biewald, 2020).