IMPOSSIBLY GOOD EXPERTS AND HOW TO FOLLOW THEM

Anonymous authors
Paper under double-blind review

ABSTRACT

We consider the sequential decision making problem of learning from an expert that has access to more information than the learner. For many problems this extra information will enable the expert to achieve greater long term reward than any policy without this privileged information access. We call these experts "Impossibly Good" because no learning algorithm will be able to reproduce their behavior. However, in these settings it is reasonable to attempt to recover the best policy possible given the agent's restricted access to information. We provide a set of necessary criteria on the expert that will allow a learner to recover the optimal policy in the reduced information space from the expert's advice alone. We also provide a new approach called ELF Distillation (**E**xplorer **L**earning from **F**ollower) that can be used in cases where these criteria are not met and environmental rewards must be taken into account. We show that this algorithm performs better than a variety of strong baselines on a challenging suite of minigrid environments.

1 INTRODUCTION

Sequential decision making is one of the most important problems in modern machine learning theory and practice. Reinforcement learning from an environmental reward signal is a powerful but unwieldy tool for attacking these problems. In contrast, imitation learning can be much more sample efficient and empirically easier to train than reinforcement learning, but requires a powerful expert that can either provide an offline dataset of instructions or online supervision. In many practical settings, these experts may have access to more information than the learning agent. This can occur, for example, when using human demonstrators to train robots that have inferior sensors, or in simulated environments where a synthetic expert uses hidden simulator information to train an agent. In these settings, it is possible that the expert's additional information makes it more powerful than any learning agent that does not have access to the hidden information. We call these experts "Impossibly Good" and show that learning from them using techniques that do not incorporate environmental rewards can cause the agent to drastically under perform the optimal policy in the reduced information space.

For example, consider a simulated robot tasked with retrieving a cell phone in an unknown apartment consisting of multiple rooms. The robot observes the world using a camera and is not given the location of the phone in advance, so it must explore each room in order to find it. Because this is a simulated environment, we can construct an expert that knows not only the location of the phone, but also the exact layout of each room and can compute the shortest path from the robot to the phone. We can then use this expert to construct a large corpus of training data across any number of apartments and phone locations. While these demonstrations may be optimal according to the expert that knows the phone's location, they crucially do not provide any demonstrations of the exploratory behavior that is necessary for the robot which must rely on its more limited sensors. At test time, the robot may need to explore many empty rooms before finding the one that contains the phone, but the expert has always walked directly to the goal and so it has never shown the robot what to do when encountering an empty room. In this case the expert is impossibly good because on average, it can reach the phone much faster than any agent that does not have access to the map, but must explore each room one by one. While we may be able to learn some important skills from this expert, we are crucially missing demonstrations of other necessary behavior, and so learning from this expert's advice alone may cause the robot to fail.

Our goal in these settings is to find an algorithm that retains the efficiency of imitation learning, while incorporating just enough reward feedback from the environment to achieve success. To address this, we introduce a new technique called ELF Distillation (**E**xplorer **L**earning from **F**ollower). The key insight of this approach is to train one **follower** policy using the advice of the impossibly good expert alone, and then use the relative success of this policy in different states to drive exploration of a second **explorer** policy using reward shaping. These two policies are trained jointly so that the explorer policy can be used to inform the distribution over states that the follower must learn from.

In order to study these problems, we have constructed a suite of MiniGrid(Chevalier-Boisvert et al., 2018) environments that clearly demonstrate the challenges of learning from impossibly good experts. While these are toy problems, they are quite challenging for many strong baselines, and allow us to clearly demonstrate the necessary concepts in a setting that avoids confounding implementation details.

2 Related Work

Recently many authors have identified the problem of learning from experts that have access to more information than the learning agent. This can occur in the self-driving domain, where a human expert may be able to see more than the car's perception system (de Haan et al., 2019; Bansal et al., 2018; Chen et al., 2019). It can also occur robot exploration, where the expert may already know a map of the world (Choudhury et al., 2017; Jain et al., 2021), or in heuristic search where the expert may know which edges are valid (Bhardwaj et al., 2017). In some cases it is possible to overcome these issues by considering a history of recent observations instead a single step, but unfortunately this exacerbates the "latching" behavior identified by several practitioners in both self-driving (Muller et al., 2006; Kuefler et al., 2017; Bansal et al., 2018; Codevilla et al., 2019) and natural language processing (Ortega et al., 2021) in which an agent becomes overly fixated on repeating recent actions.

In these partial information settings, Choudhury et al. (2018) showed that interactive imitation learning converges to the QMDP approximation of the expert's policy (Littman et al., 1995). This likely explains the empirical success (Chen et al., 2019; Lee et al., 2020) of such techniques in settings where information naturally reveals itself over time.

These interactive imitation learning techniques were originally designed to address covariate shift, a condition in which compounding single-step errors drive the learning agent into states not seen during test time. This effect was first noted by (Pomerleau, 1989) in the self-driving domain and has long been recognized as one of the most important challenges in imitation learning (Bagnell, 2015). Several approaches have been proposed to address this such as SEARN Daumé et al. (2009), DAgger (Ross et al., 2011) and AggreVaTe (Ross & Bagnell, 2014; Sun et al., 2017). Recent work (Spencer et al., 2021) has shown that covariate shift can be broken into *realizable* settings where off-policy methods such as behavior cloning work well with increasing data, and *non-realizable* settings where on-policy methods with an interactive expert (Ross et al., 2010) or an interactive simulator(Ziebart et al., 2008; Swamy et al., 2021) are necessary.

Some authors (Zhang et al., 2020; Kumor et al., 2021; Ortega et al., 2021) have proposed to address learning from experts with more information than the learner using the causal reasoning framework of Pearl et al. (2016), while Swamy et al. (2022) has recently shown that on-policy imitation learning methods can be more effective at recovering the expert's behavior than off-policy approaches in these situations, and has provided conditions under which an agent can asymptotically recover the expert's behavior.

However, in settings where information does not reveal itself, the learner has to actively gather information (Lee et al., 2021). Tennenholtz et al. (2021) examine similar settings but give the learner access to the confounder at test time. We build on these approaches by developing a new distillation technique that learns when to follow expert advice and when to explore to gather new information. Our technique uses the tools of policy distilliation (Rusu et al., 2015; Czarnecki et al., 2019) where an agent may learn from a mix of single-step imitation and multi-step reward signals.

3 IMPOSSIBLY GOOD EXPERTS

In order to study learning from an expert with more information than the learning agent, we use the framework of Partially Observable Markov Decision Processes (POMDPs), (Kaelbling et al., 1998), a generalization of Markov Decision Processes (MDPs) (Sutton & Barto, 2018) to situations where an agent must make decisions using limited observations. The goal in these settings is to sequentially make decisions in discrete time based on feedback from the environment in order to maximize a reward signal. As in an MDP, a POMDP contains an underlying state space $\mathcal{S}$ and action space $\mathcal{A}$. In an MDP, the agent receives the state s from the environment at each time step, but in a POMDP, the learning agent instead receives some possibly lossy and non-Markovian observation o that is a function of the state $o = O(s)$. Because the observations themselves are non-Markovian, the agent must reason over the history of past observations, actions and rewards up until the current time step $\tau_i = \{(o_1, a_1, r_1) \dots (o_{i-1}, a_{i-1}, r_{i-1}), o_i)\}$ referred to as a trajectory.

A learning agent's policy π_θ is a differentiable function with parameters θ mapping trajectories to a normalized distribution over actions. We define $\pi_\theta(\tau)$ as the agent's action distribution after witnessing the history τ. We also define the model class Π_θ to be the set of policies corresponding to all possible values of the parameters θ.

In order to reason about agents and experts that have different information, we consider a POMDP with two separate observation functions. The first $o_\theta = O_\theta(s)$ produces an observation that the agent sees, while the second $o^{**} = O^{**}(s)$ produces an observation that the expert sees. This allows us to reason about a trajectory in the underlying state space $\tau^{\mathcal{S}}$ and map it to a corresponding trajectory of observations for the agent τ_θ and a separate trajectory of observations for the expert τ^{**}

$$\tau^{\mathcal{S}} = \{(s_1, a_1, r_1) \dots (s_N, a_N, r_N)\}$$
$$\tau_\theta = \{(O_\theta(s_1), a_1, r_1) \dots (O_\theta(s_N), a_N, r_N)\}$$
$$\tau^{**} = \{(O^{**}(s_1), a_1, r_1) \dots (O^{**}(s_N), a_N, r_N)\}$$

The observational trajectory τ_θ represents what the learning agent observes while interacting with the environment, while the state trajectory $\tau^{\mathcal{S}}$ represents the ground truth state of the world which is unknown. For non-deterministic observation functions we use $O_\theta(o_\theta|s)$ and $O_\theta(\tau_\theta|\tau^{\mathcal{S}})$ to represent the probability of a state s or state space trajectory $\tau^{\mathcal{S}}$ producing an observation o_θ or observation trajectory τ_θ. Note that in our notation above the subscript θ is used to describe variables associated with the learner and the superscript $**$ to describe variables associated with the expert, but this is not meant to imply that O_θ is a function of the parameters θ.

We assume a model-free setting in which the states, observation functions, transition dynamics and reward dynamics are all unknown to the learning algorithm. Instead during training, the learning algorithm interacts with the environment, and receives observations and a scalar reward in the form $\{o_\theta, o^{**}, r\}$ with the assumption that o_θ and o^{**} were generated from the same unknown state s. The learning algorithm also has access to an expert π^{**} that can produce advice in the form of a suggested action $a^{**} \sim \pi^{**}(\tau^{**})$. At test time the agent will not receive the expert's observations o^{**} or have access to the expert's advice a^{**}, and so it must learn to make decisions using trajectories of o_θ alone.

We consider the episodic setting with a set of terminal states s^{term}. When a terminal state is reached, the environment informs the agent that the episode has ended using a special observation o^{term} and resets s using an unknown initial state distribution S^{init}. We write $\tau_\theta \sim \pi_\theta$ to refer to a trajectory generated by initializing according to S^{init} and repeatedly sampling actions from π_θ until a terminal state is reached. We also write $\rho_{\pi_\theta}(o_\theta)$ and $\rho_{\pi_\theta}(\tau_\theta)$ as the probability of observing o_θ and τ_θ respectively when acting according to π_θ from the initial state distribution. We will also use P_{π_θ} to refer to the set of trajectories τ that have nonzero probability when acting according to π_θ from the initial state distribution.

The return of a trajectory $G(\tau) = \sum_{i=1}^{|\tau|} r_i$ is the sum of all rewards encountered during that trajectory. The value function of a policy $V_{\pi_\theta}(s) = \mathbb{E}_{\tau \sim \pi_\theta | s_1 = s} G(\tau)$ is the expected value of the return of a trajectory that starts in state s and continues by sampling actions according to π_θ. We write $V_{\pi_\theta}^{init}$ to be $\mathbb{E}_{s \sim S^{init}} V_{\pi_\theta}(s)$, the expected value when acting according to π_θ after starting from states sampled from the initial state distribution. The goal of the learning algorithm is to learn a policy $\pi_\theta(\tau_\theta)$ using the reward signal r and expert advice a^{**} during training in order to maximize $V_{\pi_\theta}^{init}$ during

test time. An agent-optimal policy π^*_θ is one that achieves the greatest value when starting from the initial state distribution compared to all other policies in the model class $V^{init}_{\pi^*_\theta} \geq V^{init}_{\pi_\theta} \; \forall \; \pi_\theta \in \Pi_\theta$. Note that while this agent-optimal policy is the best policy achievable with limited information, it may still underperform the expert π^{**} that has more information.

We now have the tools to formally state the impossibly good criterion. We say an expert π^{**} is impossibly good iff

$$V^{init}_{\pi^{**}} > V^{init}_{\pi^*_\theta} \tag{1}$$

This means that an expert is impossibly good if in expectation it achieves a greater value from the initial state distribution than the best policy in the model class Π_θ.

4 How Not To Follow Them

In this section we consider learning from an impossibly good expert using **Behavior Cloning** and **DAgger**, two common techniques that use expert demonstrations as a learning signal. We first describe these methods, then show examples designed to demonstrate their success and failure modes when learning from impossibly good experts. We then provide formal criteria for determining when these methods are capable of finding the agent-optimal policy π^*_θ from demonstrations provided by π^{**}. Note that Swamy et al. (2022) have also recently provided conditions under which an agent with restricted information may achieve expert level performance from demonstrations. Our goal is instead to give conditions for when we are able to learn the best policy possible under the restricted observations, given demonstrations from an expert with greater performance than we can ever hope to achieve. Finally we show that in some environments, the agent-optimal policy π^*_θ cannot be recovered from expert advice alone.

Behavior Cloning is an off-policy technique in which the expert π^{**} generates a static dataset of trajectories prior to training. The algorithm then iteratively samples batches of state-action transitions from this dataset and trains the learning agent to match the expert's demonstrations. Behavior Cloning is known to suffer from covariate shift (Ross et al., 2011), a condition in which errors made by the model during test time can quickly take the agent outside of the distribution of states encountered during training. Despite this limitation, Behavior Cloning enjoys widespread popularity due to its simplicity and ability to perform well on some problems (Spencer et al., 2021).

DAgger Ross et al. (2011) is an on-policy method designed to overcome the covariate shift of Behavior Cloning. Recently Swamy et al. (2022) has shown that on-policy methods such as DAgger can improve an agent's performance when learning from an expert with more information. Instead of generating a static dataset by acting according to the expert π^{**}, DAgger iteratively collects batches of data by acting according to the learning agent π_θ and simultaneously recording the expert's advice in each visited state. After collecting a batch of data, the agent is trained to replicate the expert's advice on all data collected so far. This allows the agent to train on states that may not be visited by the expert, but are necessary to recovering from mistakes made by the learner.

Diagrams of the examples in this section are found in Figure 1. In these environments, the state space, action space and state-action transition function consist of a directed graph of nodes representing the location of the agent. The state space also contains a random variable X that is uniformly sampled from $\{0, 1\}$ at the start of each episode and remains fixed until termination. The expert observes the value of X at the start of the episode, while the agent only observes X after reaching certain nodes in the graph. Transitions between nodes are deterministic, and both expert and agent observe the label of the node they currently occupy. In all examples there are two terminal nodes, one which produces a reward of 1 when $X = 0$ and the other which produces a reward of 1 when $X = 1$. The agent also receives a small reward penalty $-\epsilon$ at each step to encourage faster solutions.

Due to the deterministic transition function and the fact that X is fixed for each episode, we can simplify π_θ to be a function of the current node and the remembered value of $X \in \{0, 1, ?\}$ rather than an entire trajectory of observations. We also assume the expert π^{**} acts optimally according to its observations.

First consider Example I. Here, the agent sees the value of X as soon as it visits B. This allows the agent-optimal policy π^*_θ to always know which terminal reward to transition to from B. An agent for this example must learn distributions for observations $\pi_\theta(A, ?)$, $\pi_\theta(B, 0)$ and $\pi_\theta(B, 1)$. Because the

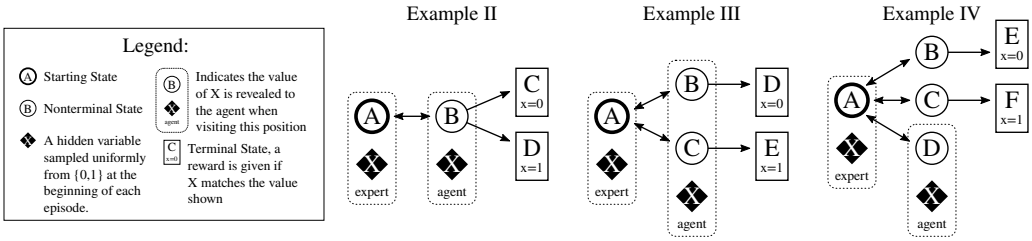

Figure 1: Two environments that can be solved with Behavior Cloning when learning from an impossibly good expert.

expert π^{**} takes the same path as the agent-optimal policy, it will generate correct training examples for these distributions. Supervising π_θ to match these examples will recover π^*_θ.

Example II in Figure 1 shows a case where Behavior Cloning fails. To see this, note that when generating data, the expert will travel from A to B when $X = 0$ and from A to C when $X = 1$. Under the agent's observation function O_θ, these correspond to $(A, ?)$, $(B, 0)$ and $(C, 1)$. Now consider the agent-optimal behavior in this example. Because of its limited information, the agent must take the first step blindly, and so it will sometimes reach either $(B, 1)$ or $(C, 0)$. In these cases, the optimal behavior is to backtrack to A where it will encounter either $(A, 0)$ or $(A, 1)$, remembering the value of X. From these observations, the agent now has enough information to guarantee reaching $(B, 0)$ or $(C, 1)$ and achieve the terminal reward. Note though that data generated by acting according to the expert did not provide demonstrations for what to do when observing $(B, 1)$ or $(C, 0)$ so we cannot hope to recover this optimal behavior.

On the other hand, we can show that DAgger can recover π^*_θ in this example. Now data is generated by the learning agent π_θ, and as we just saw, it has no way to decide whether to transition to B or C at the first step. This means that unlike Behavior Cloning, the observations $(B, 1)$ and $(C, 0)$ will exist in the training set. For these observations, the optimal behavior according to the expert π^{**} is to backtrack to A and continue down the opposite branch so these labels will be provided along with the corresponding observations. These labels are also optimal according to π^*_θ, so this extra data will allow DAgger to recover the agent-optimal policy.

Finally Example III in Figure 1 shows a case where DAgger will fail to recover π^*_θ. Now the problem is not only data coverage, but also the instructions received by the expert. Similar to Example II, the expert will always tell the learner to transition to B or C from the initial observation $(A, ?)$. However, due to the information asymmetry, the agent-optimal policy π^*_θ must first visit D in order to discover the value of X before backtracking to A and continuing to B or C as appropriate. Even if some uncertainty in the agent's policy causes it to visit D while generating data, the labels for $(A, ?)$ will never tell the agent that this is the correct behavior.

These examples have given us two intuitive conditions that must be met in order for these methods to recover π^*_θ. The first is that all observations that are encountered when acting according to the agent-optimal policy π^*_θ must be encountered during training. The second condition is that the expert π^{**} must instruct the agent to take correct actions according to π^*_θ. We now generalize these intuitions into a formal set of conditions that must be met in order for these methods to recover the agent-optimal policy π^*_θ in a simplified learning model.

Consider a dataset D generated by acting according to an arbitrary generation policy π_D and sampling labels a^{**} from π^{**}. Also consider a learning policy $\widetilde{\pi}_\theta$ that memorizes the empirical distribution of expert recommendations for each trajectory τ_θ in the dataset D.

$$\widetilde{\pi}_\theta(a|\tau_\theta) = \frac{\sum_i^{|D|} \mathbb{1}_{\tau_{\theta i}=\tau_\theta, a_i^{**}=a}}{\sum_i^{|D|} \mathbb{1}_{\tau_{\theta i}=\tau_\theta}} \tag{2}$$

Theorem 1. *The empirical policy $\widetilde{\pi}$ will recover π^*_θ as $|D| \to \infty$ iff the following conditions hold:*

1. *(Coverage)* $\rho_{\pi_D}(\tau^*_\theta) \neq 0 \; \forall \; \tau^*_\theta \in P_{\pi^*_\theta}$

2. *(Correctness)* $\mathbb{E}_{\tau^{\mathcal{S}} \sim \pi_D} \pi^{**}(O^{**}(\tau^{\mathcal{S}}))O_\theta(\tau^*_\theta|\tau^{\mathcal{S}}) = \pi^*_\theta(\tau^*_\theta) \; \forall \; \tau^*_\theta \in P_{\pi^*_\theta}$

The proof is in Appendix A.

The first condition states that all policy input trajectories that have nonzero probability of being visited under the agent-optimal policy π^*_θ must also have non-zero probability under the dataset sampling policy π_D. This corresponds to the intuition developed from the examples above that during training we must visit all states that the agent-optimal policy needs to learn about.

The second condition states that the distribution of expert advice for the state trajectories $\tau^{\mathcal{S}}$ which map to agent observation trajectories τ^*_θ must be equal to $\pi^*_\theta(\tau^*_\theta)$ for all trajectories with nonzero probability of being visited under the agent-optimal policy π^*_θ. This corresponds to the intuition developed earlier that an expert must tell the agent to act according to π^*_θ in order to recover π^*_θ.

Next we show it may not be possible to learn the agent-optimal policy π^*_θ from expert demonstrations alone regardless of the dataset generation policy π_D.

Theorem 2. *There exist environments with impossibly good experts that violate the correctness condition in Theorem 1 regardless of the dataset distribution policy π_D.*

Proof. Example III in Figure 1 is a proof by example.

In this environment $\pi^*_\theta(A \rightarrow D|(A, ?)) = 1.0$. However, $\pi^{**}(A \rightarrow D|\tau^{**}) = 0 \; \forall \; \tau^{**}$ so there is no π_D that can sample states mapping to τ^{**} that will cause $\pi^{**}(\tau^{**})$ to produce the labels required to learn π^*_θ ■

This result shows that in some cases it is impossible to learn the agent-optimal policy π^*_θ from the labels generated by an impossibly good expert alone, regardless of the rollout policy π_D that generates the dataset.

5 How to Follow Them

In the previous section, we have shown that it is possible to construct an environment with an impossibly good expert such that the learner cannot recover the agent-optimal policy when learning from the expert's advice, regardless of the dataset generating distribution π_D. Paradoxically, this means we must treat advice from these experts as sub-optimal from the agent's perspective, even though they come from an expert that can achieve higher reward than any policy we can learn. In light of this result, it is natural to ask if we can improve this situation by incorporating environmental reward into our learning algorithm.

The literature on imitation learning (Rajeswaran et al., 2017) and policy distillation (Rusu et al., 2015) provide many useful tools that can be used to learn from sub-optimal experts. We describe these methods using the unifying distillation framework of Czarnecki et al. (2019). In this setting, distillation algorithms consist of two steps. The first step rolls out a set of N state-action-reward transitions $D = \{(\tau_1, a_1, r_1) \ldots (\tau_N, a_N, r_N)\}$ using a dataset generation policy π_D (q_θ in Czarnecki et al. (2019)). In the second step, the parameters θ of the learner's policy π_θ are then updated in proportion to

$$\mathbb{E}_{\pi_D}\left[\sum_{i=1}^{|D|} -\nabla_\theta \log \pi_\theta(a_i|\tau_i)\widehat{R}_t + \nabla_\theta l(\pi_\theta(a_i|\tau_i), \pi^{**}(a_i|\tau_i))\right] \tag{3}$$

These steps are repeated until performance converges or some other stopping criterion is met. The update rule in Equation 3 can be thought of as mixing a reinforcement learning objective based on a non-differentiable multi-step return $\widehat{R}_t = \sum_{i=t}^{|\tau|} \hat{r}_i$ with a single-step differentiable loss function l. The reward term $\hat{r}_i$ can be the environmental reward r_i, or any other term we choose that assigns greater value to some behavior that we want to encourage. The loss l is more restrictive because it

must be differentiable, but provides a powerful way to inject direct supervision into the training process. By varying our choice of π_D, $\widehat{r}_i$ and l we can arrive at several different distillation algorithms, shown in Table 1, again using structure and methods similar to Czarnecki et al. (2019).

Table 1: Distillation Baselines.

Algorithm	π_D	l	$\widehat{r}_i$
Policy Gradient	π_θ	0	r_i
Teacher Distill	π^{**}	$H(\pi^{**}(s_i), \pi_\theta(s_i))$	0
On-Policy Distill	π_θ	$H(\pi^{**}(s_i), \pi_\theta(s_i))$	0
On-Policy Distill+R	π_θ	$H(\pi^{**}(s_i), \pi_\theta(s_i))$	r_i
N-Distill +R	π_θ	$H(\pi^{**}(s_i), \pi_\theta(s_i))$	-$H(\pi^{**}(s_{i+1}), \pi_\theta(s_{i+1})) + r_i$
Expert Matching Reward +R	π_θ	0	$\alpha \mathbb{1}_{a_i = a_i^{**}} - \beta \mathbb{1}_{a_i \neq a_i^{**}} + r_i$

Policy Gradient is a general class of reinforcement learning algorithms (Williams, 1992; Kakade, 2001; Schulman et al., 2015; 2017). Applying policy gradient methods to our problems essentially throws away the expert advice and learns from environmental rewards alone.

Teacher Distill and **On-Policy Distill** are the names used by Czarnecki et al. (2019) for versions of the Behavior Cloning and DAgger Ross et al. (2011) algorithms adapted for the distillation framework where data is continually generated online. These algorithms do not include an $\hat{r}_i$ term, and so will not be able to recover π_θ^* in all cases.

On-Policy Distill+R Combines the single step loss term from On-Policy Distill with the environmental reward used in Policy Gradient. The motivation is that if On-Policy Distill attempts to follow the expert, and Policy Gradient seeks high environmental reward, then combining them seeks to do both.

N-Distill+R is also from Czarnecki et al. (2019) and is similar to On-Policy Distill+R except that it augments the reward using future agreement with the expert. This encourages the model to not only take actions that the expert immediately agrees with, but also to seek out states in which it is easy to agree with the expert in the future.

Expert matching Reward+R removes the single-step loss term and instead assigns a fixed scalar α to the reward r_i when the agent's actions agree with the expert and negative reward $-\beta$ when the actions disagree.

Teacher Distill + PPO is similar to DAPG (Rajeswaran et al., 2017) and first trains an agent with Teacher Distill for a fixed number of time steps, then refines this policy using reinforcement learning. This is a popular technique used to learn robot behavior from human demonstrations.

The methods above that attempt to learn from both the expert advice and the reward signal suffer from one common problem: they all attempt to balance reward seeking and expert following behavior uniformly across all sequences of observations, and do not attempt to explicitly discover where following the expert yields high reward and where it does not.

ELF Distillation is our new technique which overcomes this limitation. Pseudocode is shown in Algorithm 1. The key insight is to train two policies jointly: a *follower* π_ϕ which attempts to learn how to follow the expert, and an *explorer* π_θ that attempts to maximize environmental reward using the follower's value function as reward shaping.

The follower is trained using a distillation that samples data according to a switching policy that rolls out experience according to the explorer π_θ for a random number of time steps, then allows the follower π_ϕ to take over. This rollout behavior is also used in AggreVaTe Ross & Bagnell (2014) and allows us to learn to follow the expert from states visited by the explorer, while also training a value function v_ϕ on the second half of the trajectory that is faithful to the follower. The follower's distillation uses cross entropy to the expert's advice as the single-step loss function l with no reward term.

The explorer is trained using policy gradients with reward shaping that encourages the explorer to visit states with high value according to the follower's learned value function. The reward term

$\hat{r}_i = r_i + v_\phi(\tau_{i+1}) - v_\phi(\tau_i)$ uses the follower's learned value function v_ϕ as potential-based reward shaping Ng et al. (1999). This is referred to as Teacher V Reward in Czarnecki et al. (2019). The single-step loss function is 0 unless the follower's value function $v_\phi(\tau)$ is greater than a target value v in which case it uses cross-entropy to the follower's distribution. This accelerates learning in states where the follower has already learned to reach high returns.

Training proceeds in an alternating fashion, first training the follower on new data generated by the switching policy, then training the explorer on new data generated using the explorer (on-policy). It is necessary to keep training the follower because the explorer may reach states in the middle of training that would not be visited otherwise, and it is important to build accurate value estimates for them as well.

At test time, the follower is no longer necessary and can be discarded. The only purpose of the follower is to discover where the expert is reliable and where it isn't in order to overcome the limitation of trying to find good global mixing rules for reward seeking and expert following behavior.

Algorithm 1 ELF Distillation

Require: N_{steps}, Expert π^{**}, Horizon T, Dataset Sizes N_{follow}, $N_{explore}$, Target value v.
  Initialize follower π_ϕ.
  Initialize follower value estimator v_ϕ.
  Initialize explorer π_θ.
  for $i = 1$ **to** N_{steps} **do**
    Initialize $D_{follow} = \{\}$.
    for $j = 1$ **to** N_{follow} **do**
      Choose a uniform random switching time $t \in \{1, \ldots, T\}$.
      Add new trajectory τ to D_{follow} by sampling $\tau_{1\ldots t} \sim \pi_\theta$ and $\tau_{t\ldots T} \sim \pi_\phi$.
    end for
    Train π_ϕ to match π^{**} for all transitions in D_{follow} using cross entropy.
    Train v_ϕ to match the returns from all transitions in D_{follow} after the switching time t.
    Initialize $D_{explore} = \{\}$.
    for $j = 1$ **to** $N_{explore}$ **do**
      Add a new trajectory τ to $D_{explore}$ by sampling $\tau \sim \pi_\theta$.
      Reshape all rewards in τ to be $\hat{r} = r + v_\phi(\tau_{t+1}) - v_\phi(\tau_t)$
    end for
    Train π_θ using policy gradient with an additional cross entropy loss to π_ϕ when $v_\phi(\tau_t) \geq v$.
  end for

6 EXPERIMENTS

In order to evaluate ELF Distill, we compare it against the baselines in Section 5 on several challenging minigrid environments with partial information. Figure 2 shows the results as well as diagrams of the environments.

The goal in these environments is to reach a door that has the same color as a set of balls hidden in various locations in the environment. Each environment contains two doors, one red and one blue and their locations and color are fixed for all episodes. The positions of the balls are also fixed, but their color is randomized at the start of each episode (all balls are assigned the same color). The agent gets a reward of $1 - 0.01 * |\tau|$ for entering the correct door for an episode. During training, demonstrations are provided by an expert that knows the color of the balls in advance and computes a shortest path to the correct door without bothering to visit locations where the balls are visible. In order to succeed at this task, the agent must learn to take information gathering actions (by seeking out a ball) and then proceed to the correct door. Visibility is strictly limited and the agent can only view the 3X3 tiles directly in front of them. While these environments are small, they are quite challenging due to the sparse reward and the need to automatically learn the color-matching behavior. As seen in Figure 2, ELF Distill outperforms all other methods in almost all cases. Descriptions of the individual environments, as well as a discussion of the individual methods' performance on them can be found in Appendix B.

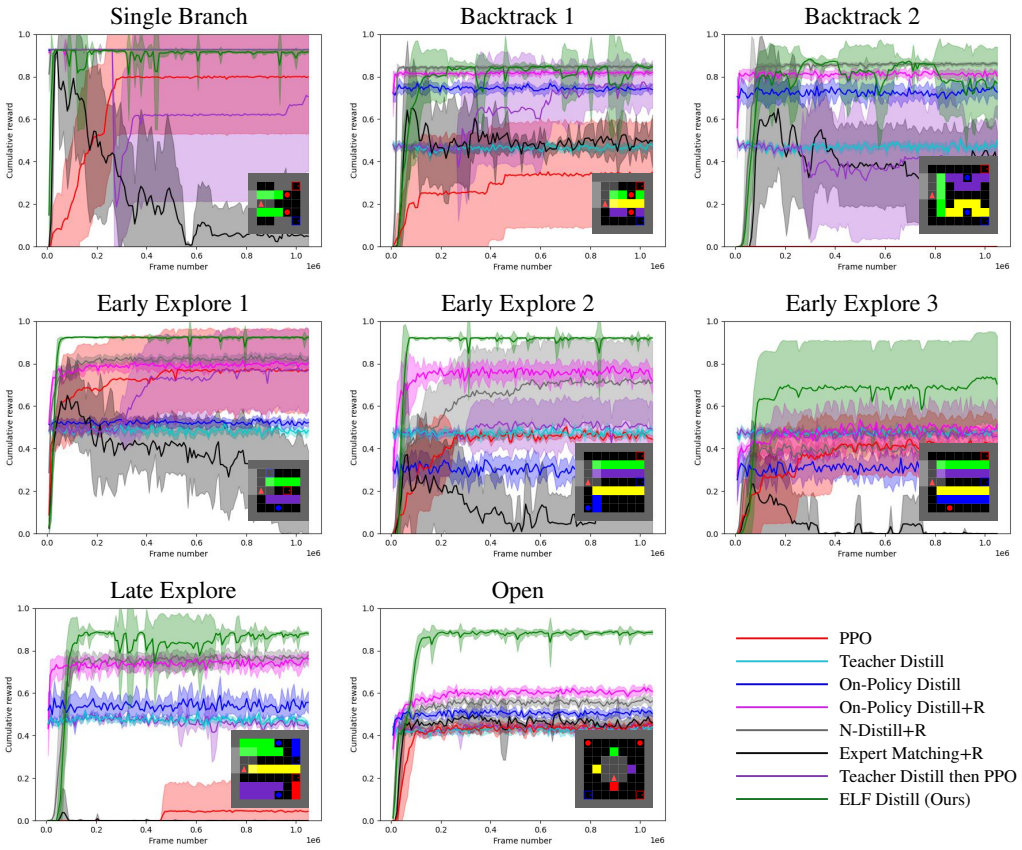

Figure 2: ELF Distill compared against seven baselines on eight environments designed to require various levels of deviation from an impossibly good expert's advice. The bottom right of each plot shows a diagram of the environment.

We use the same fully connected network for all methods. In order to avoid compounding issues of memory, our models automatically remember the target color as soon as they observe a ball. Model details are listed in Appendix C.

Each method was trained on $2^{20} \approx 1M$ frames. When training ELF Distill, both the follower and explorer were trained on $2^{19} \approx 500K$ frames, so that the total frames observed during training was equal to the other methods. All methods were trained with 10 different random seeds in each environment. PPO (Schulman et al., 2017) was used as the loss function to maximize reward in all algorithms with a $\widehat{r}_i$ term.

7 Conclusion

Imitation learning remains a powerful tool for sequential decision making problems. Unfortunately when experts have more information than the policies that learn from them, the experts may be impossibly good, and their behavior and performance cannot be replicated by any learning algorithm. We have shown that in these settings, it may still be possible to recover the optimal policy in the limited information space using only the expert's advice, but that this is not guaranteed as we have shown by providing counter-examples. To address this, we have introduced ELF Distillation, a training method that uses the estimated value function of a policy trained from expert demonstrations to guide exploration for a second reinforcement learning agent. We have shown that this algorithm outperforms several distillation baselines that incorporate both environmental reward and expert demonstrations on a set of challenging minigrid environments.

REFERENCES

J Andrew Bagnell. An invitation to imitation. Technical report, Carnegie-Mellon Univ Pittsburgh Pa Robotics Inst, 2015.

Mayank Bansal, Alex Krizhevsky, and Abhijit S. Ogale. Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst. *CoRR*, abs/1812.03079, 2018. URL http://arxiv.org/abs/1812.03079.

Mohak Bhardwaj, Sanjiban Choudhury, and Sebastian Scherer. Learning heuristic search via imitation. In *Conference on Robot Learning*, pp. 271–280. PMLR, 2017.

Dian Chen, Brady Zhou, Vladlen Koltun, and Philipp Krähenbühl. Learning by cheating, 2019. URL https://arxiv.org/abs/1912.12294.

Maxime Chevalier-Boisvert, Lucas Willems, and Suman Pal. Minimalistic gridworld environment for gymnasium. https://github.com/Farama-Foundation/MiniGrid, 2018.

Sanjiban Choudhury, Ashish Kapoor, Gireeja Ranade, Sebastian Scherer, and Debadeepta Dey. Adaptive information gathering via imitation learning. *arXiv preprint arXiv:1705.07834*, 2017.

Sanjiban Choudhury, Mohak Bhardwaj, Sankalp Arora, Ashish Kapoor, Gireeja Ranade, Sebastian Scherer, and Debadeepta Dey. Data-driven planning via imitation learning. *The International Journal of Robotics Research*, 37(13-14):1632–1672, 2018.

Felipe Codevilla, Eder Santana, Antonio M. López, and Adrien Gaidon. Exploring the limitations of behavior cloning for autonomous driving. *CoRR*, abs/1904.08980, 2019. URL http://arxiv.org/abs/1904.08980.

Wojciech M Czarnecki, Razvan Pascanu, Simon Osindero, Siddhant Jayakumar, Grzegorz Swirszcz, and Max Jaderberg. Distilling policy distillation. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 1331–1340. PMLR, 2019.

Hal Daumé, John Langford, and Daniel Marcu. Search-based structured prediction. *Machine learning*, 75(3):297–325, 2009.

Pim de Haan, Dinesh Jayaraman, and Sergey Levine. Causal confusion in imitation learning. *Advances in Neural Information Processing Systems*, 32:11698–11709, 2019.

Unnat Jain, Iou-Jen Liu, Svetlana Lazebnik, Aniruddha Kembhavi, Luca Weihs, and Alexander G Schwing. Gridtopix: Training embodied agents with minimal supervision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 15141–15151, 2021.

Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2):99–134, 1998.

Sham M Kakade. A natural policy gradient. *Advances in neural information processing systems*, 14, 2001.

Alex Kuefler, Jeremy Morton, Tim Wheeler, and Mykel Kochenderfer. Imitating driver behavior with generative adversarial networks. In *2017 IEEE Intelligent Vehicles Symposium (IV)*, pp. 204–211. IEEE, 2017.

Daniel Kumor, Junzhe Zhang, and Elias Bareinboim. Sequential causal imitation learning with unobserved confounders. *Advances in Neural Information Processing Systems*, 34:14669–14680, 2021.

Gilwoo Lee, Brian Hou, Sanjiban Choudhury, and Siddhartha S Srinivasa. Bayesian residual policy optimization:: Scalable bayesian reinforcement learning with clairvoyant experts. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5611–5618. IEEE, 2021.

Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter. Learning quadrupedal locomotion over challenging terrain. *Science robotics*, 5(47):eabc5986, 2020.

Michael L Littman, Anthony R Cassandra, and Leslie Pack Kaelbling. Learning policies for partially observable environments: Scaling up. In *Machine Learning Proceedings 1995*, pp. 362–370. Elsevier, 1995.

Urs Muller, Jan Ben, Eric Cosatto, Beat Flepp, and Yann L Cun. Off-road obstacle avoidance through end-to-end learning. In *Advances in neural information processing systems*, pp. 739–746. Citeseer, 2006.

Andrew Y Ng, Daishi Harada, and Stuart Russell. Policy invariance under reward transformations: Theory and application to reward shaping. In *Icml*, volume 99, pp. 278–287, 1999.

Pedro A Ortega, Markus Kunesch, Grégoire Delétang, Tim Genewein, Jordi Grau-Moya, Joel Veness, Jonas Buchli, Jonas Degrave, Bilal Piot, Julien Perolat, et al. Shaking the foundations: delusions in sequence models for interaction and control. *arXiv preprint arXiv:2110.10819*, 2021.

Judea Pearl, Madelyn Glymour, and Nicholas P Jewell. *Causal inference in statistics: A primer*. John Wiley & Sons, 2016.

Dean A Pomerleau. Alvinn: An autonomous land vehicle in a neural network. 1989.

Aravind Rajeswaran, Vikash Kumar, Abhishek Gupta, Giulia Vezzani, John Schulman, Emanuel Todorov, and Sergey Levine. Learning complex dexterous manipulation with deep reinforcement learning and demonstrations. *arXiv preprint arXiv:1709.10087*, 2017.

Stephane Ross and J Andrew Bagnell. Reinforcement and imitation learning via interactive no-regret learning. *arXiv preprint arXiv:1406.5979*, 2014.

Stephane Ross, Geoffrey J. Gordon, and J. Andrew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning, 2010. URL `https://arxiv.org/abs/1011.0686`.

Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. 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.

Andrei A Rusu, Sergio Gomez Colmenarejo, Caglar Gulcehre, Guillaume Desjardins, James Kirkpatrick, Razvan Pascanu, Volodymyr Mnih, Koray Kavukcuoglu, and Raia Hadsell. Policy distillation. *arXiv preprint arXiv:1511.06295*, 2015.

John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pp. 1889–1897. PMLR, 2015.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

Jonathan Spencer, Sanjiban Choudhury, Arun Venkatraman, Brian Ziebart, and J. Andrew Bagnell. Feedback in imitation learning: The three regimes of covariate shift, 2021. URL `https://arxiv.org/abs/2102.02872`.

Wen Sun, Arun Venkatraman, Geoffrey J Gordon, Byron Boots, and J Andrew Bagnell. Deeply aggrevated: Differentiable imitation learning for sequential prediction. In *International conference on machine learning*, pp. 3309–3318. PMLR, 2017.

Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

Gokul Swamy, Sanjiban Choudhury, J Andrew Bagnell, and Steven Wu. Of moments and matching: A game-theoretic framework for closing the imitation gap. In *International Conference on Machine Learning*, pp. 10022–10032. PMLR, 2021.

Gokul Swamy, Sanjiban Choudhury, J Andrew Bagnell, and Zhiwei Steven Wu. Sequence model imitation learning with unobserved contexts. *arXiv preprint arXiv:2208.02225*, 2022.

Guy Tennenholtz, Assaf Hallak, Gal Dalal, Shie Mannor, Gal Chechik, and Uri Shalit. On covariate shift of latent confounders in imitation and reinforcement learning. *arXiv preprint arXiv:2110.06539*, 2021.

Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3):229–256, 1992.

Junzhe Zhang, Daniel Kumor, and Elias Bareinboim. Causal imitation learning with unobserved confounders. *Advances in neural information processing systems*, 33:12263–12274, 2020.

Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, Anind K Dey, et al. Maximum entropy inverse reinforcement learning. In *Aaai*, volume 8, pp. 1433–1438. Chicago, IL, USA, 2008.

APPENDIX

A PROOF OF THEOREM 1

Proof. We first show that the coverage condition is necessary, and then show that the correctness condition is necessary and sufficient when the coverage condition holds.

The coverage condition is necessary: If the coverage condition does not hold, the denominator of Equation 2 will be 0, resulting in an undefined distribution for $\widetilde{\pi}_\theta$ at trajectories that are encountered while acting according to π^*_θ. Undefined behavior is not optimal according π^*_θ so this shows that the coverage condition is necessary to recover π^*_θ.

The correctness condition is necessary and sufficient when the coverage condition holds: Because the dataset D is sampled according to π_D, as $|D| \to \infty$, the empirical distribution of expert recommendations in Equation 2 becomes:

$$\widetilde{\pi}_\theta(\tau_\theta) = \mathbb{E}_{\tau^{\mathcal{S}} \sim \pi_D} \pi^{**}(O^{**}(\tau^{\mathcal{S}}))O_\theta(\tau_\theta|\tau^{\mathcal{S}}) \; \forall \; \tau_\theta \in P_{\pi_D}$$

If the coverage condition holds, then $P_{\pi^*_\theta} \subseteq P_{\pi_D}$. This allows us to extend the expression above to $P_{\pi^*_\theta}$, the set of trajectories with nonzero probability under π^*_θ

$$\widetilde{\pi}_\theta(\tau_\theta) = \mathbb{E}_{\tau^{\mathcal{S}} \sim \pi_D} \pi^{**}(O^{**}(\tau^{\mathcal{S}}))O_\theta(\tau_\theta|\tau^{\mathcal{S}}) \; \forall \; \tau_\theta \in P_{\pi^*_\theta}$$

This means that if the correctness condition holds:

$$\widetilde{\pi}_\theta(\tau_\theta) = \mathbb{E}_{\tau^{\mathcal{S}} \sim \pi_D} \pi^{**}(O^{**}(\tau^{\mathcal{S}}))O_\theta(\tau^*_\theta|\tau^{\mathcal{S}}) = \pi^*_\theta(\tau_\theta) \; \forall \; \tau_\theta \in P_{\pi^*_\theta}$$

■

This shows that the agent-optimal policy can only be recovered if the distribution over labels generated by the expert π^{**} under the sampling distribution of the state space, match the agent-optimal policy π^*_θ wherever the states map to τ^*_θ.

B DISCUSSION OF RESULTS IN MINIGRID ENVIRONMENTS

Our eight MiniGrid Environments belong to five categories. In all environments we colored the walls so that the agents with their very limited field of view could find their way around.

Single Branch was designed to mimic Example I in Figure 1. In order to each any door, the agent must pass by the balls. This is the easiest case, and the only one in which Teacher Distill can do better than 0.5. PPO is often successful here, but often takes longer to reach the goal than methods that use the expert. Here we see a common failure mode of Expert Matching+R in that the agent learns a state-action loop for which the agent agrees with most of the expert's action, but disagrees with a few of them. This allows the agent to continuously cycle and gain infinite reward according to it's reshaped reward objective that attempts to maximize agreement with the expert while failing to complete the task.

Backtrack are two environments, one larger than the other designed to mimic Example II in Figure 1. This time a ball is placed near each door so that if the agent heads to the wrong one initially,

it can backtrack to the other. On-Policy Distill performs well here, but is suboptimal to reward seeking approaches due to ambiguous expert advice before the balls have been observed. We also see Teacher Distill then PPO perform well in the smaller version, but it takes a long time to learn. PPO performs much worse on this task. ELF Distill, On-Policy Distill+R and N-Distill+R all perform very well on these problems, with ELF Distill underperforming the others in some cases.

Early Explore are three environments of different sizes designed to mimic Example III in Figure 1. Here the ball is placed on a separate hallway that is not on the path to any door. An algorithm must incorporate environmental rewards to be successful here. Here ELF Distill clearly dominates all other methods, although the problem becomes more difficult the futher the agent has to explore from states suggested by the expert.

Late Explore is similar to Backtrack in that the balls are near the doors. However in this case, they are far enough away that they will not be observed by an agent that walks directly to the door. Here again ELF Distill dominates with On-Policy Distill+R and N-Distill+R also making progress.

Open places the balls and doors in opposite corners of the room with most of the space freely navigable. Again ELF Distill outperforms all others.

C Model Details

All methods train the same small model that takes a 3x3 grid with two channels representing object type (wall, door, etc) and a single color. The model automatically remembers the color of balls that it has seen in the past. The field of view of the agent is a 3x3 grid of tiles. The input to the network is provided as integer indices, so we use three embedding layers to construct a 3x3 16-channel representation of the observed object type, another 3x3 16-channel representation of the observed object color and a 1x1 16-channel representation of the remembered ball color (a third "grey" color is used when the ball has not been observed yet). The feature for the remembered ball color is tiled to 3x3 and these features are added together and flattened. The model then uses two fully connected layers with 256 channels each and ReLU activations. A policy head consists of another 256-channel linear layer followed by a Tanh activation and a linear projection to the number of actions 7. A value head is the same except it projects to a single value.