

---

# Imitating Latent Policies from Observation

---

Ashley D. Edwards<sup>1</sup> Himanshu Sahni<sup>1</sup> Yannick Schroecker<sup>1</sup> Charles L. Isbell<sup>1</sup>

## Abstract

In this paper, we describe a novel approach to imitation learning that infers latent policies directly from state observations. We introduce a method that characterizes the causal effects of latent actions on observations while simultaneously predicting their likelihood. We then outline an action alignment procedure that leverages a small amount of environment interactions to determine a mapping between the latent and real-world actions. We show that this corrected labeling can be used for imitating the observed behavior, even though no expert actions are given. We evaluate our approach within classic control environments and a platform game and demonstrate that it performs better than standard approaches. Code for this work is available at <https://github.com/ashedwards/ILPO>.

## 1. Introduction

Humans often learn from and develop experiences through mimicry. Notably, we are capable of mirroring behavior through only the observation of state trajectories without direct access to the underlying actions (*e.g.*, the exact kinematic forces) and intentions that yielded them (Rizzolatti & Sinigaglia, 2010). In order to be general, artificial agents should also be equipped with the ability to quickly solve problems after observing the solution; however, imitation learning approaches typically require both observations and actions to learn policies along with extensive interaction with the environment.

A recent approach for overcoming these issues is to learn an initial self-supervised model for *how* to imitate by collecting experiences within the environment and then using this learned model to infer policies from expert observations (Pathak et al., 2018; Torabi et al., 2018a). However, unguided exploration can be risky in many real-world sce-

narios and costly to obtain. Thus, we need a mechanism for learning policies from observation alone without requiring access to expert actions *and* with only a few interactions within the environment.

In order to tackle this challenge, we hypothesize that predictable, though unknown, causes may describe the classes of transitions that we observe. These causes could be natural phenomena in the world, or the consequences of the actions that the agent takes. This work aims to demonstrate how an agent can predict and then imitate these latent causes, even though the ground truth environmental actions are unknown.

We follow a two-step approach, where the agent first learns a policy offline in a latent space that best describes the observed transitions. Then it takes a limited number of steps in the environment to ground this latent policy to the true action labels. We liken this to learning to play a video game by observing a friend play first, and then attempting to play it ourselves. By observing, we can learn the goal of the game and the types of actions we should be taking, but some interaction may be required to learn the correct mapping of controls on the joystick.

We first make the assumption that the transitions between states can be described through a discrete set of latent actions. We then learn a forward dynamics model that, given a state and latent action, predicts the next state and prior, supervised only by {state, next state} pairs. We use this model to greedily select the latent action that leads to the most probable next state. Because these latent actions are initially mislabeled, we use a few interactions with the environment to learn a relabeling that outputs the probability of the true action.

We evaluate our approach in four environments: classic control with cartpole, acrobot, and mountain car, and a recent platform game by OpenAI, CoinRun (Cobbe et al., 2018). We show that our approach is able to perform as well as the expert after just a few steps of interacting with the environment, and performs better than a recent approach for imitating from observations, Behavioral Cloning from Observation (Torabi et al., 2018a).

---

<sup>1</sup>Georgia Institute of Technology, Atlanta, GA, USA. Correspondence to: Ashley D. Edwards <[aedwards8@gatech.edu](mailto:aedwards8@gatech.edu)>.

## 2. Related work

Imitation learning approaches aim to train artificial and real-world agents to imitate expert behavior by providing a set of expert demonstrations. This approach has an extensive breadth of applications, ranging from early successes in autonomous driving (Pomerleau, 1989), to applications in robotics (Schaal, 1997; Chernova & Thomaz, 2014) and software agents (e.g. (Silver et al., 2016; Nair et al., 2017)). However, traditional approaches typically assume that the expert’s actions are known. This often requires the data to be specifically recorded for the purpose of imitation learning and drastically reduces the amount of data that is readily available. Recent approaches that do not require expert actions typically must *first* learn behaviors in the agent’s environment through extensive interactions. Our approach first learns latent behaviors from the demonstration data only, followed by only a few necessary interactions with the environment. We now describe classic approaches to imitation learning along with more modern approaches.

### 2.1. Classic approaches

Arguably, the most straight-forward approach to imitation learning is behavioral cloning (Pomerleau, 1989), which treats imitation learning as a supervised learning problem. More sophisticated methods achieve better performance by reasoning about the state-transitions explicitly, but often require extensive information about the effects of the agent’s actions on the environment. This information can come either in the form of a full, often unknown, dynamics model, or through numerous interactions with the environment. Inverse Reinforcement Learning (IRL) achieves this by using the demonstrated state-action pairs to explicitly derive the expert’s intent in the form of a reward function (Ng & Russell, 2000; Abbeel & Ng, 2004).

### 2.2. Direct policy optimization methods

Recently, more direct approaches have been introduced that aim to match the state-action visitation frequencies observed by the agent to those seen in demonstrations. GAIL (Ho & Ermon, 2016) learns to imitate policies from demonstrations and uses adversarial training to distinguish if a state-action pair comes from following the agent or expert’s policy while simultaneously minimizing the difference between the two. SAIL (Schroecker & Isbell, 2017) achieves a similar goal by using temporal difference learning to estimate the gradient of the normalized state-action visitation frequency directly. However, while these approaches are efficient in the amount of expert data necessary for training, they typically require a substantial amount of interactions within the environment.

### 2.3. Learning from state observations

Increasingly, works have aspired to learn from observation alone without utilizing expert actions. Imitation from Observation (Liu et al., 2017), for example, learns to imitate from videos without actions and translates from one context to another. However, this approach requires using learned features to compute rewards for reinforcement learning, which will thus require many environment samples to learn a policy. Similarly, time-contrastive networks (Sermanet et al., 2017) and unsupervised perceptual rewards (Sermanet et al., 2016) train robots to imitate from demonstrations of humans performing tasks, and recent work used audio to align different YouTube videos to train an agent to learn Montezuma’s revenge and Pitfall (Aytar et al., 2018). But these approaches also learn features for a reward signal that is later used for reinforcement learning. Finally, both third-person imitation learning (Stadie et al., 2017) and GAIfo (Torabi et al., 2018b) extend GAIL for use with demonstration data that lacks actions, but these approaches also utilize a reward signal in a similar manner as GAIL. Therefore, while each of these approaches learn policies from state observations, they require an intermediary step of using a reward signal, whereas we learn the policy directly without performing reinforcement learning.

A recent approach aimed to learn from observations by first learning how to imitate in a self-supervised manner, then given a task, attempt it zero-shot (Pathak et al., 2018). However, this approach requires learning in the agent’s environment first rather than initially learning from the observations. Another approach utilizes learned inverse dynamics to train agents from observation (Torabi et al., 2018a). A problem with such an approach is that learning a dynamics model usually requires a substantial number of interactions with the environment. Our work aims to first learn policies from demonstrations *offline*, and then only use a few interactions with the environment to learn the true action labels.

### 2.4. Multi-modal predictions

Our approach predicts forward dynamics given a state and latent action. This is similar to recent works that have learned action-conditional predictions for reinforcement learning environments (Oh et al., 2015; Chiappa et al., 2017), but those approaches utilize ground truth action labels. Rather, our approach learns a latent multi-modal distribution over future predictions. Other related works have utilized latent information to make multi-modal predictions. For example, BicycleGAN (Zhu et al., 2017) learns to predict a distribution over image-to-image translations, where the modes are sampled given a latent vector. InfoGAN uses latent codes for learning interpretable representations (Chen et al., 2016), and InfoGAIL (Li et al., 2017) uses that approach to capture latent factors of variation between different demonstrations.




Figure 1: The latent policy network learns a latent policy,  $\pi(z|s)$ , and a forward dynamics model,  $G$ . The action remapping network learns  $\pi(a|s_t, z)$  to align the latent actions  $z$  with ground-truth actions  $a$ . We train embeddings,  $E_a$  and  $E_p$ , concurrently with each network.

These works, however, do not attempt to learn direct priors over the modes, which is crucial in our formulation for deriving policies.

As such, our approach is more analogous to online clustering, as it predicts multiple expected next states and priors over them. However, we do not have direct access to the clusters or means. Other works have aimed to cluster demonstrations, but these approaches have traditionally segmented different types of trajectories, which represent distinct preferences, rather than next-state predictions (Hausman et al., 2017; Babes et al., 2011).

### 3. Approach

We now describe our approach, Imitating Latent Policies from Observation (ILPO), where we train an agent to imitate behaviors from expert state observations.

#### 3.1. Problem formulation

We aim to use ILPO to solve problems specified through a Markov Decision Process (MDP)  $\langle S, A, R, T \rangle$  (Sutton & Barto, 1998). Here,  $s \in S$  denotes the states in the environment,  $a_t \in A$  corresponds to actions,  $r_t \in R$  are the rewards the agent receives in each state, and  $T(s_t, a_t, s_{t+1})$  is the transition model, which we assume is unknown. Reinforcement learning approaches aim to learn policies  $\pi(a|s_t)$  that determine the probability of taking an action  $a$  in some state  $s_t$ . We use imitation learning to directly learn the policy and use the reward  $r_t$  only for evaluation purposes.

We are given a set of expert demonstrations described through noisy state observations  $\{s_1^* \dots s_n^*\} \in D$ . In our approach, we will use these observations to predict a multimodal forward dynamics model. As such, noise is necessary for ensuring that state transitions are properly modeled.

Given two consecutive observations  $\{s_t, s_{t+1}\}$ , we define  $z$  as a *latent action* that caused this transition to occur. As

such, the action spaces that we consider are discrete with deterministic transitions. Because our problems are specified through MDPs, we assume that the number of actions,  $|A|$ , is known. Hence, we can define  $\{z_1 \dots z_{|A|}\} \in Z$  latent actions, where  $|Z| = |A|$  is used as an initial guess for the number of latent actions. However, there may be more or less types of transitions that appear in the demonstration data. For example, if an agent has an action to move left but always moves right, then the "left" transition will not be observed. Or if the agent moves right and bumps into a wall, this stationary transition may appear to be another type of action. As such, we will empirically study the effect of using latent actions when  $|Z| \neq |A|$ .

#### 3.2. Behavioral cloning

Given expert states and actions  $\{s_1, a_1 \dots s_n, a_n\}$ , behavioral cloning uses supervised learning to approximate  $\pi(a|s_t)$ . That is, given a state  $s_t$ , this approach predicts the probability of taking each action, i.e., the policy. However, imitation by observation approaches do not have access to expert actions. To address this, behavioral cloning from observation (BCO) (Torabi et al., 2018a) first learns an inverse dynamics model  $f(a|s_t, s_{t+1})$  by first collecting samples in the agent's environment. Then, the approach uses this model to label the expert observations and learn  $\pi(a|s_t)$ . However learning dynamics models online can require a large amount of data, especially in high-dimensional problems.

We make the observation that we do not need to know action labels to make an initial hypothesis of the policy. Rather, our approach learns a *latent* policy  $\pi_\omega(z|s_t)$  that estimates the probability that a latent action  $z$  would be taken when observing  $s_t$ . This process can be done *offline* and hence more efficiently utilizes the demonstration data. We then use a limited number of interactions with the environment to learn an action-remapping network that efficiently associates the true actions the agent can take with the latent policy identified by our learned model.

**Algorithm 1** Imitating Latent Policies from Observation

---

```

1: function ILPO( $s_0^*, s_1^*, \dots, s_N^*$ )
2:   Step 1: Learning latent policies
3:   for  $k \leftarrow 0 \dots \#Epochs$  do
4:     for  $i \leftarrow 0 \dots N - 1$  do ▷ (Omitting batching for clarity)
5:       Train latent dynamics parameters  $\theta \leftarrow \theta - \nabla_{\theta} \min_z \|G_{\theta}(E_{p\theta}(s_i^*), z) - s_{i+1}^*\|_2^2$ 
6:       Train latent policy parameters  $\omega \leftarrow \omega - \nabla_{\omega} \|\sum_z \pi_{\omega}(z|s_i^*)G_{\theta}(E_{p\theta}(s_i^*), z) - s_{i+1}^*\|_2^2$ 

7:   Step 2: Action remapping
8:   Observe state  $s_0$ 
9:   for  $t \leftarrow 0 \dots \#Interactions$  do
10:    Choose latent action  $z_t \leftarrow \arg \max_z \pi_{\omega}(z|E_{a\xi}(s_t))$ 
11:    Take  $\epsilon$ -greedy action  $a_t \leftarrow \arg \max_a \pi_{\xi}(a|z_t, E_{a\xi}(s_t))$ 
12:    Observe state  $s_{t+1}$ 
13:    Infer closest latent action  $z_t = \arg \min_z \|E_{p\theta}(s_{t+1}) - E_{p\theta}(G_{\theta}(E_{p\theta}(s_t), z))\|_2$ 
14:    Train action remapping parameters  $\xi \leftarrow \xi + \nabla_{\xi} \log \frac{\pi_{\xi}(a_t|z_t, E_{a\xi}(s_t))}{\sum_a \pi_{\xi}(a|z_t, E_{a\xi}(s_t))}$ 
    
```

---

### 3.3. Step 1: Learning latent policies

In order to learn  $\pi_{\omega}(z|s_t)$ , we introduce a latent policy network with two key components: a latent forward dynamics model  $G$  that learns to predict  $\hat{s}_{t+1}$ , and a prior over  $z$  given  $s_t$ , which gives us the latent policy, as shown in figure 1.

#### 3.3.1. LATENT FORWARD DYNAMICS

We first describe how to learn a latent forward dynamics model from expert state observations. Given an expert state  $s_t$  and latent action  $z$ , our approach trains a generative model  $G_{\theta}(E_p(s_t), z)$  to predict the next state  $s_{t+1}$ , where  $E_p$  is an embedding that is trained concurrently. Similar to recent works that predict state dynamics (Edwards et al., 2018; Goyal et al., 2018), our approach predicts the differences between states  $\Delta_t = s_{t+1} - s_t$ , rather than the absolute next state, and computes  $s_{t+1} = s_t + \Delta_t$ .

When learning to predict forward dynamics, a single prediction,  $f(s_{t+1}|s_t)$ , will not account for the different modes of the distribution, i.e., the effects of each action, and will thus predict the mean over all transitions. When using an action-conditional model (Oh et al., 2015; Chiappa et al., 2017), learning each mode is straightforward, as we can simply make predictions based on the observed next state after taking each action,  $f(s_{t+1}|s_t, a)$ . However, in our approach, we do not know the ground truth actions that yielded a transition. Instead, our approach trains a generative model  $G$  to make predictions based on each of the latent actions  $z \in Z$ ,  $f(s_{t+1}|s_t, z)$ .

To train  $G$ , we compute the loss as:

$$\mathcal{L}_{min} = \min_z \|\Delta_t - G_{\theta}(E_p(s_t), z)\|_2^2. \quad (1)$$

To allow predictions to converge to the different modes, we only penalize the one closest to the true next observation,

$s_{t+1}$ . Hence the generator must learn to predict the closest mode within the multi-modal distribution. This approach essentially allows each generator to learn *transition clusters* for each type of transition that is represented through  $\Delta_t$ . If we penalized each of the next state predictions simultaneously, the generator would learn to always predict the *expected* next state, rather than each distinct state observed after taking a latent action  $z$ .

We use  $\Delta_t$  to better guide the generator to learn distinct transitions. For example, if we have an agent moving in discrete steps along the  $x$ -axis, then moving right would yield positive transitions  $\Delta = 1$  and moving left would yield negative transitions  $\Delta = -1$ . Our approach aims to train the generator to learn these different types of transitions.

Note that since  $G$  is learning to predict  $\Delta_t$ , we will need to add each prediction to  $s_t$  in order to obtain a prediction for  $s_{t+1}$ . For simplicity, in further discussion we will refer to  $G$  directly as the predictions summed with the state input  $s_t$ .

#### 3.3.2. LATENT POLICY LEARNING

Crucially, ILPO concurrently learns the latent policy  $\pi_{\omega}(z|E_p(s_t))$ . This represents the probability that given a state  $s_t$ , a latent transition of the type  $z$  will be observed in the expert data. We train this by computing the expectation of the generated predictions under this distribution, i.e., the expected next state, as:

$$\hat{s}_{t+1} = \mathbb{E}_{\pi_{\omega}}[s_{t+1}|s_t] \quad (2)$$

$$= \sum_z \pi_{\omega}(z|s_t)G_{\theta}(E_p(s_t), z). \quad (3)$$

We then minimize the loss as:

$$\mathcal{L}_{exp} = \|s_{t+1} - \hat{s}_{t+1}\|_2^2 \quad (4)$$

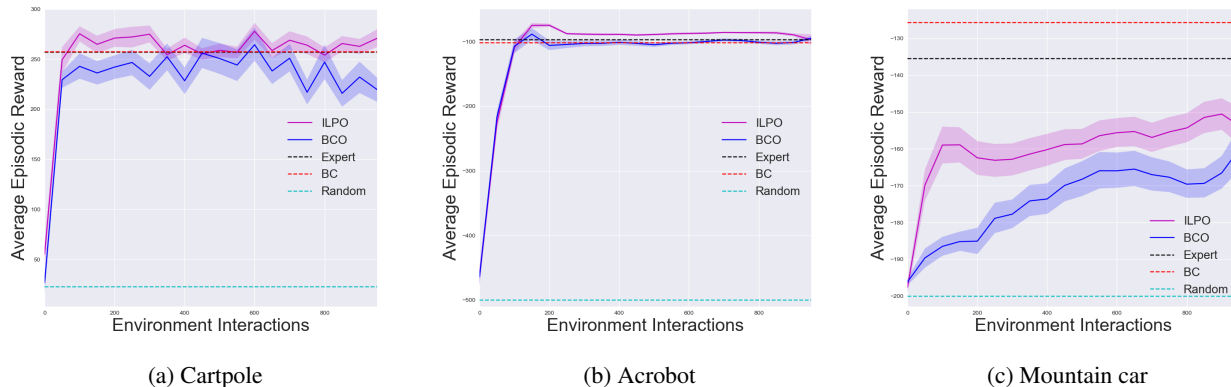


Figure 2: Classic control imitation learning results. The trials were averaged over 50 runs for ILPO and the policy was evaluated every 50 steps and averaged out of 10 policy runs. The reward used for training the expert and evaluation was +1 for every step that the pole was upright in cartpole, and a -1 step cost for acrobot and mountain car.

while holding the individual predictions fixed. This approach predicts the probability of each transition occurring. In other terms, this determines the most likely transition cluster.

With this loss, the latent policy is encouraged to make predictions that yield the most likely next state. For example, if an agent always moves right, then we should expect, given some state, for the next state to reflect the agent moving right. Any other type of transition should have a low probability so that it is not depicted within the next state.

The network is trained using the combined loss:

$$\mathcal{L}_{policy} = \mathcal{L}_{min} + \mathcal{L}_{exp}. \quad (5)$$

We outline the training procedure for this step in lines 3-6 in algorithm 1.

### 3.4. Step 2: Action Remapping

In order to imitate from expert observations, the agent needs to learn a mapping from the latent policy learned in the previous step to the true action space:  $\pi_{\xi}(a_t|z, E_a(s_t))$ , where  $E_a$  is an embedding that is trained concurrently. As such, it is invariantly necessary for the agent to explore the effect of its own actions within its environment. However, unlike BCO and other imitation from observation approaches, ILPO only needs to learn a mapping from  $a$  to  $z$  rather than a full dynamics model. The mapping  $\pi_{\xi}$  also depends on the current state  $s_t$  because latent actions are not necessarily invariant across states. The actions being predicted by each generator might change in different parts of the state-space. If an agent is flipped upside-down, for example, then the action "move up" would then look like "move down".

Nevertheless, generalization capabilities of neural networks should encourage a strong correlation between  $a$  and  $z$ . The

dynamics in two states are often more similar for the same action than they are for two different ones, thus assigning the latent actions to the same type of transition should allow the network to generalize more easily. This intuition will allow us to learn such a mapping from only a few interactions with the environment, but is not a requirement for learning, and the algorithm will be able to learn to imitate the expert's policy regardless.

#### 3.4.1. COLLECTING EXPERIENCE

To obtain training data for the remapped policy  $\pi_{\xi}$ , we allow the agent to interact with the environment to collect experiences in the form of  $\{s_t, a_t, s_{t+1}\}$  triples. This interaction can follow any policy, such as a random policy or one that is updated in an online way. The only stipulation is that a diverse section of the state space is explored to facilitate generalization. We choose to iteratively refine the remapped policy  $\pi_{\xi}$  and collect experiences by following this current estimate, in addition to an  $\epsilon$ -greedy exploration strategy.

#### 3.4.2. ALIGNING ACTIONS

While collecting experiences  $\{s_t, a, s_{t+1}\}$  in the agent's environment, we proceed in two steps to train the remapped policy. First, we identify the latent action that corresponds to the environmental state transitions  $\{s_t, s_{t+1}\}$  and then we use the environmental action  $a$  taken as a label to train  $\pi_{\xi}(a_t|z_t, E_a(s_t))$  in a supervised manner.

To do this, given state  $s_t$ , our method uses the latent dynamics model,  $G$ , trained in step 1, to predict each possible next state  $\hat{s}_{t+1}$  after taking a latent action  $z$ . Then it identifies the latent action that corresponds to the predicted next state

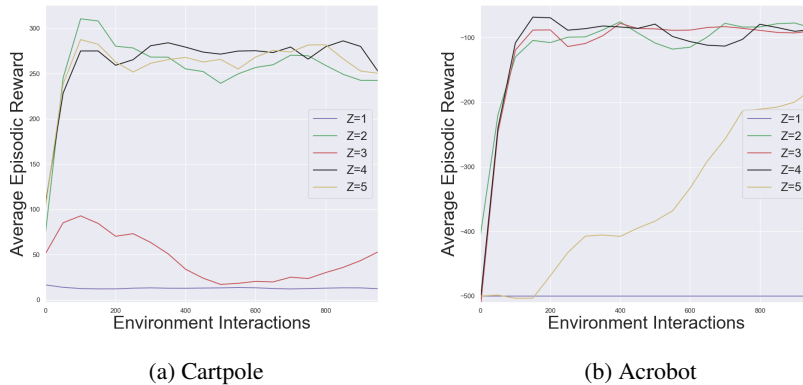


Figure 3: Cartpole and Acrobot results for selecting  $|Z|$ . The trials were averaged over 5 runs for ILPO and the policy was evaluated every 50 steps and averaged out of 10 policy runs. The reward used for training the expert and evaluation was +1 for every step that the pole was upright in cartpole, and a -1 step cost for acrobot.

that is the most similar to the observed next state  $s_{t+1}$ :

$$z_t = \arg \min_z \|s_{t+1} - G_\theta(E_p(s_t), z)\|_2. \quad (6)$$

To extend this approach to situations where euclidean distance is not meaningful (such as high-dimensional visual domains), we may also measure distance in the space of the embedding  $E_p$  learned in the previous step. In these domains, the latent action is thus given by:

$$z_t = \arg \min_z \|E_p(s_{t+1}) - E_p(G_\theta(E_p(s_t), z))\|_2. \quad (7)$$

Having obtained the latent actions  $z_t$  most closely corresponding to the environmental action  $a_t$ , we then train  $\pi(a_t|z, s_t)$  as a straight forward classification problem using a cross-entropy loss:

$$\mathcal{L}_{map} = \log \frac{\pi_\xi(a_t|z_t, E_a(s_t))}{\sum_a \pi_\xi(a|z_t, E_a(s_t))}. \quad (8)$$

### 3.4.3. IMITATING LATENT POLICIES FROM OBSERVATION

Combining the two steps into a full imitation learning algorithm, given a state  $s_t$ , we use the latent policy outlined in step 1 to identify the latent cause that is most likely to have the effect that the expert intended,  $z^* = \arg \max_z \pi_\omega(z|s_t)$ , and subsequently identify the action that is most likely to cause this effect,  $a^* = \arg \max_a \pi_\xi(a|z^*, s_t)$ . The agent can then follow this policy to imitate the expert’s behavior without having seen any expert actions. We outline the training procedure for this step in lines 7-14 in algorithm 1.

## 4. Experiments and results

In this section, we discuss the experiments used to evaluate ILPO. We aim to demonstrate that our approach is able to

imitate from state observations only and with little interactions with the environment. In addition to this, we aim to show that learning dynamics online through environment interactions is less sample efficient than learning a latent policy first. In these experiments, we compare the environment interactions in ILPO during the action remapping phase with online data collection in BCO used for learning the inverse dynamics model. These samples are obtained after following the policies of each respective method, and we evaluate each approach after the same number of interactions.

We evaluate ILPO within classic control problems as well as a more complex visual domain. We used OpenAI Baselines (Dhariwal et al., 2017) to obtain expert policies and generate demonstrations for each environment. We compare ILPO against this expert, a random policy, and Behavioral Cloning (BC), which is given ground truth actions, averaged over 50 trials, and Behavioral Cloning from Observation (BCO). More details can be found in the appendix.

### 4.1. Classic control environments

We first evaluated our approach within classic control environments (Sutton & Barto, 1998). We used the standard distance metric from equation 6 to compute the distances between observed and predicted next states for ILPO. We used the same network structure and hyperparameters across both domains, as described in the appendix. We used 50,000 expert state observations to train ILPO and BCO, and the corresponding actions to train Behavioral Cloning (BC).

**Cartpole** is an environment where an agent must learn to balance a pole on a cart by applying forces of  $-1$  and  $1$  on it. The state space consists of 4 dimensions:  $\{x, \dot{x}, \theta, \dot{\theta}\}$ , and the action space consists of the 2 forces. As such, ILPO must predict 2 latent actions and generate predicted next




Figure 4: CoinRun environment used in experiments. CoinRun consists of procedurally generated levels and the goal is to get a single coin at the end of a platform. We used an easy level (left) and hard level (middle and right) in our experiments. The middle image is the agent’s state observation in the hard level and the image on the right is a zoomed out version of the environment. This task is difficult because the gap in the middle of the platform can not be recovered from and there is a trap once the agent reaches the other side. When training, the images also include a block showing  $x$  and  $y$  velocities.

states with 4 dimensions.

**Acrobot** is an environment where an agent with 2 links must learn to swing its end-effector up by applying a torque of  $-1$ ,  $0$ , or  $1$  to its joint. The state space consists of 6 dimensions:  $\{\cos \theta_1, \sin \theta_1, \cos \theta_2, \sin \theta_2, \dot{\theta}_1, \dot{\theta}_2\}$ , and the action space consists of the 3 forces. As such, ILPO must predict 3 latent actions and generate a predicted next state with 6 dimensions.

**Mountain car** is an environment where an agent on a single-dimension track must learn to push a car up a mountain by applying a force of  $-1$ ,  $0$ , or  $1$  to it. The state space consists of 2 dimensions:  $\{x, \dot{x}\}$ , and the action space consists of the 3 forces. As such, ILPO must predict 3 latent actions and generate a predicted next state with 2 dimensions.

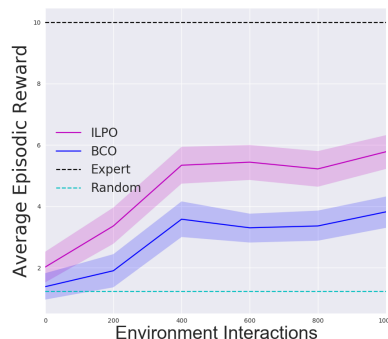
#### 4.1.1. RESULTS

Figure 2 (left) shows the imitation learning results in cartpole. ILPO learns the correct policy and is able achieve the same performance as the expert and behavioral cloning in less than 100 steps within the environment. Furthermore, ILPO performs much better than BCO, as it does not need to learn state-transitions while collecting experience, only a mapping from latent to real actions.

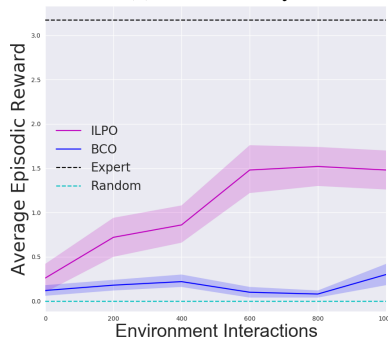
Figure 2 (middle) shows the imitation learning results in acrobot. ILPO again learns the correct policy after a few steps and is able achieve as good of performance as the expert and behavioral cloning, again within 100 steps. While BCO learns quickly, ILPO again performs better than it.

Finally, the results for mountain car are shown in Figure 2 (right). Neither ILPO nor BCO performed as well as the expert. Nevertheless, it is clear that ILPO outperforms BCO.

We were also interested in evaluating the effect of using a



(a) CoinRun easy



(b) CoinRun hard

Figure 5: CoinRun imitation learning results. The trials were averaged over 50 runs for ILPO and BCO the policy was evaluated every 200 steps and averaged out of 10 policy runs. The reward used for training the expert and evaluation was  $+10$  after reaching the coin.

different number of latent actions. These results are shown in figure 3. We see that choosing  $|Z| = |A|$  is a good initial guess for the size of  $Z$ , but the agent is also able to learn from other sizes.  $|Z| = 1$  performed poorly in both acrobot and cartpole. This is because every action will collapse to a single latent and the state predictions cannot be disentangled.

As we mentioned in section 3.1, ILPO requires stochastic demonstrations. We found that although the agent was capable of performing well with deterministic demonstrations, the performance decreased in this setting. See the appendix for more discussion.

## 4.2. CoinRun

We also evaluated our approach in a more complex visual environment, CoinRun (Cobbe et al., 2018). This environment consists of procedurally generated platform environments. In particular, the background, player, enemies, platforms, obstacles, and goal locations are all randomly instantiated. The agent can take actions left, right, jump, and down, jump-left, jump-right, and do-nothing. The game ends when the






Figure 6: Next state predictions computed by ILPO in the CoinRun easy task. The highlighted state represents the closest next state obtained from equation 7.

agent reaches a single coin in the game. We used 1000 episodes of expert demonstrations to train ILPO and BCO.

In these experiments, we evaluated each approach within a single easy and hard level, shown in figure 4. This environment is more difficult than classic control because it uses images as inputs and contains more actions. As such, the dynamics learning takes place over many more dimensions. In particular, the state space consists of  $128 \times 128 \times 3$  pixels and 7 actions. Thus, ILPO must predict  $128 \times 128 \times 3$  dimensions for the next-state predictions. We found that ILPO performed better when predicting  $|Z| = 5$  latent actions. This is likely because certain actions, such as moving left, were used less frequently. We use the embedded distance metric from equation 7 to compute the distances between observed and predicted next states.

#### 4.2.1. RESULTS

Figure 5 shows the results for imitation learning. In both the easy and hard tasks, ILPO was not able to perform as well as the expert, but performed significantly better than BCO. As this environment is high-dimensional, it takes more steps to learn the alignment policy than the previous experiments. However, ILPO often learned to solve the task almost immediately, but some random seeds led to bad initialization that resulted in the agent not learning at all. However, good initialization sometimes allows the agent to learn good initial policies zero-shot (see videos in the supplementary for an example). As such, we found that it was possible for the agent to sometimes perform as well as the expert. The results consist of all of the seeds averaged, including those that yielded poor results.

Figure 6 shows the predictions made by the model. ILPO is able to predict moving right and jumping. Because these are the most likely modes in the data, the other generators also predict different velocities. The distance metric is able to correctly select the closest state.

In general, it can often be difficult to learn dynamics from visual inputs. Unlike BCO, by learning a latent policy first, ILPO is able to reduce the number of environment interactions necessary to learn. BCO requires solving both an inverse dynamics model and behavioral cloning each time

it collects a batch of experience. As such, this approach is less efficient than ILPO and in many scenarios would be difficult to perform in realistic environments.

## 5. Discussion and conclusion

In this paper, we introduced ILPO and described how agents can learn to imitate latent policies from only expert state observations and very few environment interactions. We demonstrated that this approach recovered the expert behavior in four different domains consisting of classic control and vision based tasks. ILPO requires very few environment interactions compared to BCO, a recent dynamics-based imitation from observation approach. In many real world scenarios, unguided exploration in the environment may be very risky but expert observations can be readily made available. Such a method of learning policies directly from observation followed by a small number of action alignment interactions with the environment can be very useful for these types of problems.

There are many ways that this work can be extended. First, future work could address two assumptions in the current formulation of the problem: 1) that it requires that actions are discrete and 2) assumes that the state transitions are deterministic. Second, the action remapping step can be made even more efficient by enforcing stronger local consistencies between latent actions and generated predictions across different states. This will drastically reduce the number of samples required to train the action remapping network by decreasing variation between latent and real actions.

We believe this work will introduce opportunities for learning to observe not only from similar agents, but from other agents with different embodiments whose actions are unknown or do not have a known correspondence. Another contribution would be to learn to transfer across environments.

Finally, our work is complimentary to many of the related approaches discussed. Several algorithms utilize behavioral cloning as a pre-training step for more sophisticated imitation learning approaches. As such, we believe ILPO could also be used for pre-training imitation by observation.



## References

- Abbeel, P. and Ng, A. Y. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, pp. 1. ACM, 2004.
- Aytar, Y., Pfaff, T., Budden, D., Paine, T. L., Wang, Z., and de Freitas, N. Playing hard exploration games by watching youtube. *arXiv preprint arXiv:1805.11592*, 2018.
- Babes, M., Marivate, V., Subramanian, K., and Littman, M. L. Apprenticeship learning about multiple intentions. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pp. 897–904, 2011.
- Chen, X., Duan, Y., Houthoofd, R., Schulman, J., Sutskever, I., and Abbeel, P. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in Neural Information Processing Systems*, pp. 2172–2180, 2016.
- Chernova, S. and Thomaz, A. L. Robot learning from human teachers. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 8(3):1–121, 2014. URL <http://www.morganclaypool.com/doi/abs/10.2200/S00568ED1V01Y201402AIM028>.
- Chiappa, S., Racaniere, S., Wierstra, D., and Mohamed, S. Recurrent environment simulators. *arXiv preprint arXiv:1704.02254*, 2017.
- Cobbe, K., Klimov, O., Hesse, C., Kim, T., and Schulman, J. Quantifying generalization in reinforcement learning. *arXiv preprint arXiv:1812.02341*, 2018.
- Dhariwal, P., Hesse, C., Klimov, O., Nichol, A., Plappert, M., Radford, A., Schulman, J., Sidor, S., and Wu, Y. Openai baselines. <https://github.com/openai/baselines>, 2017.
- Edwards, A. D., Downs, L., and Davidson, J. C. Forward-backward reinforcement learning. *arXiv preprint arXiv:1803.10227*, 2018.
- Goyal, A., Brakel, P., Fedus, W., Lillicrap, T., Levine, S., Larochelle, H., and Bengio, Y. Recall traces: Backtracking models for efficient reinforcement learning. *arXiv preprint arXiv:1804.00379*, 2018.
- Hausman, K., Chebotar, Y., Schaal, S., Sukhatme, G., and Lim, J. J. Multi-modal imitation learning from unstructured demonstrations using generative adversarial nets. In *Advances in Neural Information Processing Systems*, pp. 1235–1245, 2017.
- Ho, J. and Ermon, S. Generative adversarial imitation learning. In *Advances in Neural Information Processing Systems*, pp. 4565–4573, 2016.
- Li, Y., Song, J., and Ermon, S. Infogail: Interpretable imitation learning from visual demonstrations. In *Advances in Neural Information Processing Systems*, pp. 3815–3825, 2017.
- Liu, Y., Gupta, A., Abbeel, P., and Levine, S. Imitation from observation: Learning to imitate behaviors from raw video via context translation. *arXiv preprint arXiv:1707.03374*, 2017.
- Nair, A., McGrew, B., Andrychowicz, M., Zaremba, W., and Abbeel, P. Overcoming exploration in reinforcement learning with demonstrations. *arXiv preprint arXiv:1709.10089*, 2017.
- Ng, A. A. Y. and Russell, S. J. S. Algorithms for inverse reinforcement learning. In *Icml*, volume 0, pp. 663–670, 2000. URL <http://ai.stanford.edu/~ang/papers/icml00-irl.pdf>.
- Oh, J., Guo, X., Lee, H., Lewis, R. L., and Singh, S. Action-conditional video prediction using deep networks in atari games. In *Advances in Neural Information Processing Systems*, pp. 2863–2871, 2015.
- Pathak, D., Mahmoudieh, P., Luo, G., Agrawal, P., Chen, D., Shentu, Y., Shelhamer, E., Malik, J., Efros, A. A., and Darrell, T. Zero-shot visual imitation. *arXiv preprint arXiv:1804.08606*, 2018.
- Pomerleau, D. A. Alvin: An autonomous land vehicle in a neural network. In *Advances in neural information processing systems*, pp. 305–313, 1989.
- Rizzolatti, G. and Sinigaglia, C. The functional role of the parieto-frontal mirror circuit: interpretations and misinterpretations. *Nature reviews neuroscience*, 11(4):264, 2010.
- Schaal, S. Robot learning from demonstration. *Neural Information Processing System (NIPS)*, 1997. ISSN 1049-5258. URL <http://www.iam.ira.uka.de/users/rogalla/WebOrdnerMaterial/ml-robotlearning.pdf>.
- Schroeder, Y. and Isbell, C. L. State aware imitation learning. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30*, pp. 2911–2920. Curran Associates, Inc., 2017. URL <http://papers.nips.cc/paper/6884-state-aware-imitation-learning.pdf>.
- Sermanet, P., Xu, K., and Levine, S. Unsupervised perceptual rewards for imitation learning. *arXiv preprint arXiv:1612.06699*, 2016.
- Sermanet, P., Lynch, C., Hsu, J., and Levine, S. Time-contrastive networks: Self-supervised learning from multi-view observation. *arXiv preprint arXiv:1704.06888*, 2017.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Sander, D., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., and Hassabis, D. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587): 484–489, 2016.
- Stadie, B. C., Abbeel, P., and Sutskever, I. Third-person imitation learning. *arXiv preprint arXiv:1703.01703*, 2017.
- Sutton, R. S. and Barto, A. G. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge, 1998.
- Torabi, F., Warnell, G., and Stone, P. Behavioral cloning from observation. *arXiv preprint arXiv:1805.01954*, 2018a.
- Torabi, F., Warnell, G., and Stone, P. Generative adversarial imitation from observation. *arXiv preprint arXiv:1807.06158*, 2018b.
- Zhu, J.-Y., Zhang, R., Pathak, D., Darrell, T., Efros, A. A., Wang, O., and Shechtman, E. Toward multimodal image-to-image translation. In *Advances in Neural Information Processing Systems*, pp. 465–476, 2017.