




Towards Generalisable Imitation Learning Through Conditioned Transition Estimation and Online Behaviour Alignment

Nathan Gavenski 
King’s College London
London, United Kingdom
nathan.schneider_gavenski@kcl.ac.uk

Matteo Leonetti 
King’s College London
London, United Kingdom
matteo.leonetti@kcl.ac.uk

Odinaldo Rodrigues 
King’s College London
London, United Kingdom
odinaldo.rodrigues@kcl.ac.uk




ABSTRACT

State-of-the-art imitation learning from observation methods (ILfO) have recently made significant progress, but they still have some limitations: they need action-based supervised optimisation, assume that states have a single optimal action, and tend to apply teacher actions without full consideration of the actual environment state. While the truth may be out there in observed trajectories, existing methods struggle to extract it without supervision. In this work, we propose Unsupervised Imitation Learning from Observation (UfO) that addresses all of these limitations. UfO learns a policy through a two-stage process, in which the agent first obtains an approximation of the teacher’s true actions in the observed state transitions, and then refines the learned policy further by adjusting agent trajectories to closely align them with the teacher’s. Experiments we conducted in five widely used environments show that UfO not only outperforms the teacher and all other ILfO methods but also displays the smallest standard deviation. This reduction in standard deviation indicates better generalisation in unseen scenarios.

KEYWORDS

Imitation Learning, Generalisation

ACM Reference Format:

Nathan Gavenski , Matteo Leonetti , and Odinaldo Rodrigues . 2026. Towards Generalisable Imitation Learning Through Conditioned Transition Estimation and Online Behaviour Alignment. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 9 pages. <https://doi.org/10.65109/VJZU8032>

1 INTRODUCTION

One way to learn how to perform a task is to observe someone else execute it and then attempt to imitate their actions. In computer science, this is called *imitation learning* (IL), where an agent observes a “teacher”. Unlike other learning approaches, IL leverages the teacher’s guidance, thus avoiding blindly learning by trial-and-error.

Learning from demonstration (ILfD) is the most common approach for IL, where the agent uses sequences of states and actions performed by the teacher. In practice, it suffers from two significant drawbacks: poor generalisation in environments with multiple solutions and the scarcity of action-labelled data [9]. *Learning from observation* (ILfO) mitigates the latter by learning a task with access to only sequences

of states via self-supervision. This allows agents to use much more easily obtainable data, such as video recordings of game-playing.


ILfO typically employs one of two approaches: *behavioural cloning* (BC) or *adversarial imitation learning* (AIL). BC seeks to match a teacher’s actions by minimising a supervised loss, while AIL aims to make the agent’s behaviour indistinguishable from that of the teacher. To avoid ILfD’s generalisation problem, BC assumes each state in the environment can be reached by only one optimal action, which is not always the case. To address this and correct the agent’s behaviour, AIL collects samples via *online play* (direct environment interaction). Yet, these strategies still have shortcomings. BC tends to overfit, and AIL can fall into *behaviour-seeking* mode (blindly mimicking the teacher’s behaviour, irrespective of the actual environment state). This leads to poor performance in environments that require the underlying action intent or the causal relationships between actions and outcomes for success.

To tackle these shortcomings, recent works [10, 19] use self-supervision to approximate the teacher’s behaviour, with others employing an artificial reward function [17, 28]. However, these works still assume that states can be reached only by a single action, rely on an action-based loss function, and exhibit behaviour-seeking due to reliance on the quality and coverage of the data samples.

In this work, we introduce Unsupervised Imitation Learning from Observation (UfO), a novel unsupervised approach to imitation learning from observation. Unlike prior ILfO methods [10, 17, 19, 26], which rely on supervised losses to infer or reconstruct teacher actions, UfO learns without action-based supervision. Instead, it derives its understanding of behaviour from the causal structure of state transitions. The main contribution of UfO is the reconstruction stage, where a policy and a generative model are jointly trained to approximate the environment’s dynamics through mutual optimisation: the policy seeks actions that best condition the generative model to reproduce the teacher’s transitions, and the generative model refines its predictions based on environment transitions. This unsupervised coupling eliminates dependence on action-level supervision while constraining learning to behaviourally meaningful transitions, mitigating the behavioural-seeking bias common in ILfO. A subsequent adversarial refinement stage further improves policy generalisation through online play, enabling UfO to surpass the teacher’s performance and outperform all competing methods across five benchmark environments. Together, these results establish UfO as the first method to achieve effective unsupervised imitation purely from observation, without relying on supervised action inference or external rewards.

2 BACKGROUND

In this work, we assume the environment to be a Markov Decision Process (MDP) $M = \langle S, A, T, r, \gamma \rangle$, in which S is the state space, A is

 This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/VJZU8032>

the action space, T is a transition function $T : S \times A \rightarrow S$, r is the immediate reward function, and γ is the discount factor [23, Ch. 3].

Solving an MDP yields a policy π with a probability distribution P over actions. In each state s , $P(a | s)$ is the probability of taking action a in s . For simplicity, we assume that $\pi(s)$ will determine the action the agent takes in state s , i.e., the action a such that $P(a | s)$ is the highest. We will differentiate policies via appropriate subscripts.

The objective of IL is to generate an agent policy π_θ that approximates the teacher’s policy π_ψ from a set of ‘trajectories’ representing the behaviour of the teacher solving the task. Trajectories may consist of sequences of state-action pairs or simply sequences of state transitions. We will use the symbol \mathcal{T} to denote the set of all trajectories, irrespective of their type. Notice that IL does not assume access to the teacher’s reward function or discount factor.

If the trajectories include the actions performed by the teacher, imitation learning becomes a supervised learning problem, known as *behavioural cloning* (BC), optimising:

$$\min_{\theta} \sum_{\tau \in \mathcal{T}} \sum_{(s,a) \in \tau} \mathcal{L}(a, \pi_{\theta}(s)). \quad (1)$$

Above, we minimise the supervised loss function \mathcal{L} (such as the mean square error) between the action a the teacher chooses at state s (annotated in a trajectory τ) and the action $\pi_{\theta}(s)$ the agent policy thinks it should take. Yet, relying on annotated actions from a teacher might not be feasible. In some cases, such as when the teacher’s data consists of video recordings, it might be impossible to retrieve which action was actually taken. Therefore, ILfO assumes that trajectories contain only sequences of states and generate agent policies based solely on these sequences. The two most common ILfO approaches are behavioural cloning and adversarial imitation learning (see Gavenski et al. [11] for details).

2.1 Behavioural Cloning Methods

Behavioural Cloning for ILfO tries to generate the agent policy using *inverse dynamics models* (IDM) and *forward dynamics models* (FDM). IDM tries to predict the action responsible for a transition (s, s') , where FDM tries to model T directly. The disadvantage of these dynamic models is that although the MDP consists of a single function T , they rely on learning one of many possible transition functions $T' \subseteq T$. IL methods that rely on IDM [10, 19] use teacher’s trajectories τ_ψ to create *self-supervised* actions $\tilde{a} \in \{\text{IDM}(s, s') \mid (s, s') \in \tau_\psi\}$ to train their policies using Eq. 1. On the other hand, methods that rely on FDM use the teacher’s trajectories coupled with action [5] or temporal [20] information to condition the next state generation and optimise the agent using:

$$\min_{\theta} \min_{\phi} \sum_{\tau \in \mathcal{T}} \sum_{(s,s') \in \tau} \mathcal{L}(s', \mathcal{G}_{\phi}(s, \pi_{\theta}(s))), \quad (2)$$

where \mathcal{G}_{ϕ} is a generative model approximating T , and \mathcal{L} is still a supervised loss function used to minimise the loss between the actions chosen by the teacher and agent.

These methods rely on the assumption that the teacher transition function is an *injective function*, i.e., $T(s_i, a_i) \neq T(s_j, a_j)$, whenever $(s_i, a_i) \neq (s_j, a_j)$. However, this rarely holds in practice, resulting in the learning process encouraging the agent to favour a single action during a transition (even when multiple alternatives are viable). Hence, the agent tends to learn disproportionately skewed policies. The inability to give the correct weights to action alternatives in

unseen scenarios often leads to failure [24]. We will show how to tackle this issue in Section 3.1.

2.2 Adversarial Imitation Learning Methods

Adversarial imitation learning employs a *discriminator* that attempts to distinguish between the actions chosen by the agent and those chosen by the teacher for states encountered during online play. Instead of learning an action mapping function similar to BC, it tries to approximate the teacher policy by how the agent behaves. AIL can be further divided into two categories: (i) *inverse reinforcement learning*, and (ii) *generative adversarial learning*.

Most AIL approaches learn by minimising a variation of the adversarial loss from Goodfellow et al. [13]:

$$\min_{\omega} \max_{\theta} \sum_{\tau \in \mathcal{T}} \sum_{s \in \tau} \log(\mathcal{D}_{\omega}(s, \pi_{\psi}(s))) + \log(1 - \mathcal{D}_{\omega}(s, \pi_{\theta}(s))) \quad (3)$$

where \mathcal{D}_{ω} is a *discriminator model* that identifies whether a pair (s, a) is from the teacher or from the agent. It is important to note that Eq. 3 also does not assume access to π_{ψ} . It only requires the agent’s policy to follow the behaviour of the teacher’s. Therefore, to “fool” the discriminator, the agent attempts to follow the teacher’s trajectories as closely as possible. For inverse reinforcement learning approaches, \mathcal{D}_{ω} trains to predict a value between $[0, 1]$, where 0 means the model thinks (s, a) is from the agent and 1 means it is from the teacher. Generative adversarial learning methods use the gradient from \mathcal{D}_{ω} to teach the policy how to behave as the teacher. For ILfO methods, Eq. 3 uses state transitions (s', s) rather than state-action pairs to discriminate.

Dependence on adversarial losses requires IL approaches to rely heavily on the teacher samples’ coverage and quality. By forcing the agent to behave precisely as its teacher, adversarial methods often fall into behaviour-seeking mode [28] — trying to emulate the teacher’s behaviour too closely and forcing a trajectory that might not be appropriate in the current state of the environment. To make things worse, IL methods attempt to reduce the number of teacher samples to lower the learning cost (for greater *efficiency*), amplifying this issue further. Section 3.2 shows how UfO employs AIL sparingly to refine the agent’s policy without causing it to blindly mimic the teacher’s behaviour.

3 UNSUPERVISED IL FROM OBSERVATION

In this work, we propose Unsupervised Imitation Learning from Observation, a novel IL technique that avoids the need for action-based supervised optimisation, the reliance on the assumption of injective transition functions and the behaviour-seeking mode. UfO is an unsupervised imitation learning method composed of: (i) a policy model π_{θ} to learn to predict the best action given a state s ; (ii) a *conditioned* generative model \mathcal{G}_{ϕ} to generate the next state $\mathcal{G}_{\phi}(s, \pi_{\theta}(s))$, given a state s and the action $\pi_{\theta}(s)$; and (iii) a discriminator model \mathcal{D}_{ω} to discriminate between agent and teacher trajectories $\mathcal{D}_{\omega}(\tau_{\pi_{\theta}} || \tau_{\pi_{\psi}})$. Algorithm 1 provides an overview of UfO’s learning process. UfO initialises all models with random weights and collects teacher trajectories $\mathcal{T}_{\pi_{\psi}}$ (Lines 1 and 2). The learning process is divided into two stages. In the reconstruction stage (Lines 5-14), UfO learns how to approximate the teacher’s policy from the teacher observations collected, whilst learning an approximation of the MDP’s transition function (removing the need

for supervised losses). In the subsequent adversarial stage (Lines 16-22), UfO refines the policy to more closely resemble the teacher’s.

3.1 Reconstruction Stage

The reconstruction stage (illustrated in Fig. 1a) trains both generator and policy based on the teacher transition function and the agent transition function \mathcal{T}_{π_θ} . It does so, by iteratively *freezing* \mathcal{G}_ϕ and π_θ , and training each model by using \mathcal{T}_{π_θ} and \mathcal{T}_{π_ψ} , respectively. This part of the training runs for epochs_r, which is significantly higher than the second stage.

For the policy portion of the training (Lines 6-8), UfO uses the teacher samples to learn how to properly condition the generator based on the teacher’s transitions following:

$$\min_{\theta} \sum_{\tau \in \mathcal{T}_{\pi_\psi}} \sum_{(s, s') \in \tau} |s' - \mathcal{G}_\phi(s, \pi_\theta(s))|. \quad (4)$$

This part aims to enable the policy to learn how to accurately predict an action that allows the generator to determine the teacher’s next state based on its learned transition function. In other words, given that \mathcal{G}_ϕ learns a sufficiently approximate transition function $T' \subseteq T$, if the agent predicts the same action that the teacher would choose, the next state should also be similar – $T(s, \pi_\psi(s)) \approx \mathcal{G}_\phi(s, \pi_\theta(s))$.

For the generator part of the training (Lines 11-13), UfO uses the agent samples collected from the environment to learn how to approximate the MDP’s transitions using:

$$\min_{\phi} \sum_{\tau \in \mathcal{T}_{\pi_\theta}} \sum_{s \in \tau} |T(s, \pi_\theta(s)), \mathcal{G}_\phi(s, \pi_\theta(s))|. \quad (5)$$

Due to the lack of knowledge about actions, it is impossible to approximate the transition function based on the teacher’s observations alone. This part aims to enable the generator to learn how the MDP’s transition behaves through transitions caused by the agent in the environment.

Algorithm 1 UfO

```

1: Initialise parameters  $\theta$  for  $\pi_\theta$ ;  $\phi$  for  $\mathcal{G}_\phi$ ; and  $\omega$  for  $\mathcal{D}_\omega$ 
2: Collect teacher  $\pi_\psi$  trajectories  $\mathcal{T}_{\pi_\psi}$ 
3: Let  $\mathcal{T}_{\pi_\psi} = \{(s, s') \in \tau \mid \tau \in \mathcal{T}_{\pi_\psi}\}$ 
4: Stage 1: Reconstruction
5: for  $i \leftarrow 1$  to epochsr do
6:   for all  $(s, s') \in \mathcal{T}_{\pi_\psi}$  do
7:     Update  $\theta$  with  $\mathcal{L}_R(s', \mathcal{G}_\phi(s, \pi_\theta(s)))$ 
8:   end for
9:   Collect agent  $\pi_\theta$  trajectories  $\mathcal{T}_{\pi_\theta}$  in the environment
10:  Let  $\mathcal{T}_{\pi_\theta} = \{(s, s') \in \tau \mid \tau \in \mathcal{T}_{\pi_\theta}\}$ 
11:  for all  $(s, s') \in \mathcal{T}_{\pi_\theta}$  do
12:    Update  $\phi$  with  $\mathcal{L}_R(s', \mathcal{G}_\phi(s, \pi_\theta(s)))$ 
13:  end for
14: end for
15: Stage 2: Adversarial
16: for  $i \leftarrow 1$  to epochsa do
17:   Collect agent  $\pi_\theta$  trajectories  $\mathcal{T}_{\pi_\theta}$  in the environment
18:   for all  $\tau \in \mathcal{T}_{\pi_\theta}^\Delta \cup \mathcal{T}_{\pi_\psi}^\Delta$  do
19:      $y \leftarrow \{1 \mid \tau \in \mathcal{T}_{\pi_\psi}\} \cup \{0 \mid \tau \in \mathcal{T}_{\pi_\theta}\}$ 
20:     Update  $\omega$  and  $\theta$  with  $\mathcal{L}_D(y, \mathcal{D}_\omega(\tau))$ 
21:   end for
22: end for

```

Therefore, by intercalating policy and generator training, UfO learns an approximate transition function that is leveraged to learn a policy similar to the teacher’s. As the agent approximates its policy to the teacher, so do its trajectories, leading the generator to fine-tune itself to similar transitions. In other words, since the agent replicates trajectories close to the teacher and the generator only updates its weights via online play (i.e., by interacting with the environment), UfO ensures a positive correlation between the teacher and agent models. Specifically, to minimise Eq. 4, the policy must predict an action that generates the same next state $\mathcal{G}_\phi(s, \pi_\theta(s)) = \mathcal{G}_\phi(s, \pi_\psi(s))$, and to minimise Eq. 5, the generator has to correctly predict the next state based on the current state of the environment and the action predicted by the policy $\mathcal{G}_\phi(s, \pi_\theta(s)) = T(s, \pi_\theta(s))$. By not having any supervised losses, the policy must learn for itself which actions align more with the teacher’s next expected state and better understand how different values affect the transition function from \mathcal{G}_ϕ . Hence, the loss function does not introduce bias into the policy’s action predictions, thereby addressing the supervised loss limitations of ILfO approaches. Moreover, UfO circumvents the reliance on injective transition functions by learning from state transitions $P(s' \mid s, \pi_\theta(s))$ rather than direct state-action mappings $P(a \mid s)$. This is because the gradients updating θ originate from the generator itself (purple arrows in Fig. 1a, and the generator learns from state-action pairs.

Nevertheless, at the end of the reconstructive stage, the policy will have learned only from teacher transitions, and these may not be sufficient to determine how the agent should act in unseen cases. To mitigate the absence of training data for these cases, we add an adversarial stage with online play to further align the policy with the teacher’s and improve generalisation.

3.2 Adversarial Stage

The adversarial stage (Fig. 1b) trains both the policy, based on the agent samples collected from the environment \mathcal{T}_{π_θ} , and the discriminator \mathcal{D}_ω based on both the teacher samples \mathcal{T}_{π_ψ} and the agent’s. First, UfO freezes \mathcal{G}_ϕ so that the model does not learn to “fool” the discriminator by diverging from the learned approximate transition function from the reconstructive stage. In this stage, only the policy π_θ and discriminator \mathcal{D}_ω are trained. We freeze the generator \mathcal{G}_ϕ to preserve the dynamics learned during the reconstruction phase. This ensures that the policy alone is fine-tuned to better reproduce the teacher’s behavioural patterns without altering the environment approximation. If UfO were not to freeze its generator, given there are more optimisation functions binding \mathcal{G}_ϕ ’s optimisation to the MDP’s transition function, the generator would focus on learning how to force the discriminator to misclassify the states it generates, diverging from the learnt transition function $\mathcal{G}_\phi = T' \not\subseteq T$.

Afterwards, UfO collects samples using π_θ in the environment (Line 17) and trains both models by optimising a minmax framework:

$$\min_{\omega} \max_{\theta} \mathbb{E}_{\tau \sim \mathcal{T}_{\pi_\psi}^\Delta} \log(\mathcal{D}_\omega(\tau)) + \mathbb{E}_{\tau \sim \mathcal{T}_{\pi_\theta}^\Delta} \log(1 - \mathcal{D}_\omega(\tau)) \quad (6)$$

where $\mathcal{T}_{\pi_\psi}^\Delta$ is the sequence of the differences between subsequent states in teacher trajectories, and $\mathcal{T}_{\pi_\theta}^\Delta$ is the sequence of the differences between subsequent states generated by \mathcal{G}_ϕ . Hence, the i th entry from $\mathcal{T}_{\pi_\psi}^\Delta$ is $|s_i - T(s_i, \pi_\psi(s_i))|$, and from $\mathcal{T}_{\pi_\theta}^\Delta$ is $|s_i - \mathcal{G}_\phi(s_i, \pi_\theta(s_i))|$. UfO uses the difference between states to avoid memorising the teacher’s states since they are fixed during training. Moreover, \mathcal{D}_ω

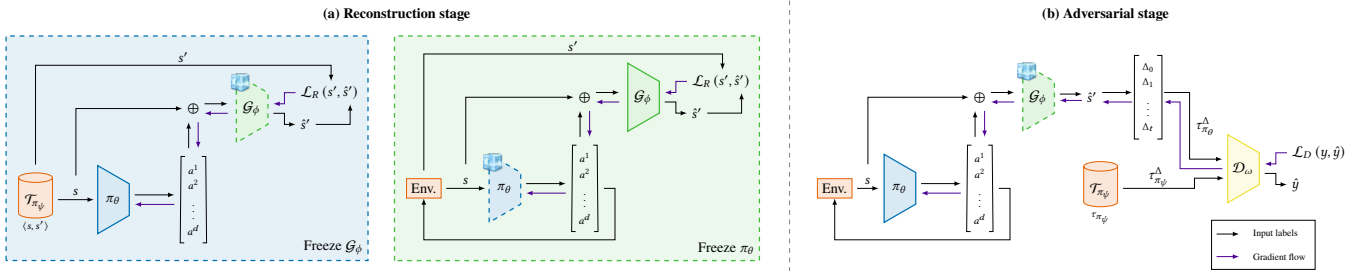


Figure 1: Two-stage training cycle. (a) Each model is iteratively frozen and trained using samples from the teacher \mathcal{T}_{π_ψ} and the agent, respectively. (b) \mathcal{G}_ϕ is frozen, and π_θ and \mathcal{D}_ω are trained with agent samples from the environment and the teacher’s.

can be more applicable to unseen states by using the difference since the differences between states will hold.

It is important to note that the discriminator in UfO is a recurrent neural network. Thus, it receives the differences between all pairs of states in a trajectory and updates both the policy and its weights. Using the whole trajectory allows UfO to have a global view of the behaviour of the teacher and agent through the sequence of states and since the discriminator is recurrent, it can correct the policy when it acts optimally locally but not globally. Nevertheless, the use of an adversarial stage poses some challenges: (i) it is susceptible to entering behaviour-seeking mode, and (ii) heavy updates to the policy can lead to *catastrophic forgetting*, i.e., loss of previously learned information by a pre-trained model [16]. These challenges may arise because the adversarial stage encourages the agent to take actions that produce differences resembling those made by the teacher. However, these differences may not be optimal across all possible trajectories. This can result in poorer trajectories, causing a compounding error between the policy and the discriminator, significantly altering θ . As a consequence, the policy may lose its previously acquired knowledge. To address these challenges, UfO uses only 10 epochs in the adversarial stage (epochs_a), a drastically lower learning rate, and clips the gradients when updating θ . By avoiding prolonged adversarial training, UfO manages to avoid behavioural-seeking mode, while keeping the acquired knowledge the policy learned from its reconstruction step.

We note that UfO achieves state-of-the-art results even without the adversarial stage. However, this stage fine-tunes the policy further, by allowing it to perform online play in the environment. This leads to a slight increase in the reward and a decrease in the standard deviation between trajectories (as described in Sec. 6.2).

4 EXPERIMENTAL RESULTS

We compared UfO’s results against the following five key related methods. Behavioural Cloning from Observation (BCO) [25] and Generative adversarial Imitation from Observation (GAIfO) [26], which are usually used as baselines, and three of the better ILfO methods: Continuous Imitation Learning from Observation (CILO) [10], Modality-agnostic Adversarial Hypothesis Adaptation for Learning from Observations (MAHALO) [17], and Off-Policy Imitation Learning from Observations (OPOLO) [28]. We do not compare to other baselines, such as DICE-based methods, such as SMODICE [18] and

IQ-Learn [7], since MAHALO outperforms them in their experimentation. We experimented with five commonly used environments: Ant, Half Cheetah, Hopper, Swimmer, and Inverted Pendulum.

We trained each method using the hyperparameters provided in the supplementary material and ran for 10,000 randomly generated seeds not present in the training data. Each environment trajectory consists of 1,000 steps or the number of steps before the agent falls.

4.1 Implementation and Metrics

We used PyTorch to implement our agent and optimise the loss functions in Eq. 4-6 via Adam [15]. We used the datasets provided by Imitation Datasets [8], and Optuna [1] for all hyperparameter searches. It is important to note that we use Imitation Datasets for two reasons: (i) different from other datasets (e.g., D4RL [6] and Minari [27]), Imitation Datasets uses the newest version of these benchmarks (v4 instead of v2), allowing for easier reproducibility and longer support; and (ii) allows us to preserve the characteristics of the experiments across all evaluations by running them with the same seeds, guaranteeing replicability and having access to the teacher weights. Yet, we note that Imitation Datasets only provides the teacher datasets and weights, and all benchmarks are the same in this work as others (e.g., MuJoCo). We also provided all learning rates in the supplementary material with a link to the official implementation.

We evaluate all methods using the *Average Episodic Reward (AER)* metric and *Performance \mathcal{P}* . $AER(\pi)$ is the average accumulated reward for a policy π over n trajectories with t steps each: $AER(\pi) = \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^t \gamma^j r(s_{i,j}, \pi(s_{i,j}))$. $\mathcal{P}_\tau(\pi)$ is the reward of the policy π in a trajectory τ , normalised between random and teacher policies, where 0 corresponds to the reward of a random policy π_ξ , and 1 to the reward of the teacher policy π_ψ : $\mathcal{P}(\pi) = \frac{AER(\pi) - AER(\pi_\xi)}{AER(\pi_\psi) - AER(\pi_\xi)}$. Finally, we do not report accuracy since achieving high accuracy does not necessarily translate into a high reward for the agent.

To avoid any confusion, we note that although IL methods might achieve \mathcal{P} higher than one, the overall objective is not to be as close as possible to the teacher (e.g., $AER(\pi_\theta) \approx AER(\pi_\psi)$), but to measure whether the agent is learning a policy that is relatively better than both policies. Results in Tab. 1 are the average and standard deviation in the five environments. We do not report accuracy since achieving high accuracy does not necessarily translate into a high reward for the agent (seen in Sec. 6.1).

Methods	Metrics	InvertedPendulum-v4	Hopper-v4	Ant-v4	Swimmer-v4	HalfCheetah-v4
Random (π_{ξ})	AER	5.7 ± 3.26	18.8985 ± 18.4452	-57.2408 ± 97.069	0.73 ± 11.44	-284.446 ± 79.2153
Teacher (π_{ψ})	AER	1000 ± 0	3530.2857 ± 12.2681	5765.3842 ± 621.1187	355.4238 ± 1.832	9512.2995 ± 538.5918
UfO	AER	1000 ± 0	3573.4266 ± 8.5877	5904.2506 ± 398.1026	361.2001 ± 0.9960	9959.5745 ± 589.3388
	\mathcal{P}	1	1.0127	1.0238	1.0163	1.0457
CILO	AER	1000 ± 0	3559.3540 ± 10.3743	5848.3754 ± 458.6725	356.8693 ± 1.495	9510.299 ± 765.2461
	\mathcal{P}	1	1.0078	1.0143	1.0041	0.9998
MAHALO	AER	1000 ± 0	1033.9924 ± 0.5959	903.7594 ± 37.1806	87.4172 ± 10.9691	154.6271 ± 4.0698
	\mathcal{P}	1	0.2891	0.1650	0.0248	0.0448
OPOLO	AER	990.6254 ± 13.8079	3505.4686 ± 40.9337	5791.2084 ± 786.6135	352.2456 ± 2.1989	9442.2904 ± 869.6332
	\mathcal{P}	0.9906	0.9934	1.0044	0.9910	0.9929
GAIFO	AER	870.3248 ± 31.9845	2586.4747 ± 977.1989	2656.604 ± 963.7789	360.5495 ± 0.9515	1988.5288 ± 790.4208
	\mathcal{P}	0.8696	0.7315	0.4661	1.0145	0.2320
BCO	AER	521.826 ± 178.9	1156.643 ± 298.1488	2324.0222 ± 1573.7205	347.1842 ± 2.2823	6776.0547 ± 1755.5855
	\mathcal{P}	0.5191	0.3241	0.4090	0.9768	0.7207

Table 1: UfO and baselines AER and \mathcal{P} results for benchmarks as the average of 10,000 seeds.

4.2 Results

Tab. 1 shows the results from UfO and all baselines in 10,000 different seeds. Each seed is selected based on whether it appeared in the training data or during the online play for each method. We select each seed based on its first state because it is impossible to guarantee that all methods will not predict actions to match its current trajectory to one of the teacher’s. Hence, we reset the environment and checked whether the first state appeared during the online play or in the entire training dataset; if it did not, we stored that seed for use during test time. The first and second rows from Tab. 1 show a random policy π_{ξ} , and the teacher π_{ψ} used to generate each dataset in the same 10,000 seeds. Although the dataset’s original AER might differ from those shown in Tab. 1 for the teacher, we believe comparing each agent’s results under the same conditions is fairer whenever possible. In particular, they should employ the same seeds and weights, which Imitation Datasets allows since it shares the teacher weights. Therefore, the teacher row shows the AER for the teacher in the same seeds as all other methods.

We observed that UfO was the only method that consistently achieved performance higher than 1. It was the only method to outperform its teacher in HalfCheetah, and by a large margin (447.275), whereas two of the other methods had difficulty even approximating the teacher’s performance. Moreover, UfO’s standard deviation was lower than all other methods in all environments, only slightly higher than the teacher’s in HalfCheetah and smaller or equal to the teacher’s in all other environments. We argue that this demonstrates that UfO can outperform the teacher and, more importantly, it deviates less from the desired behaviour. UfO achieves an average performance of 1.0197, the highest across all benchmarks, with CILO being the second with 1.0054. Note that \mathcal{P} needs to be analysed in tandem with coefficients of variation (CVs — the ratio of the standard deviation to the mean average), because methods may be more sensitive to variations in trajectories, with lower CVs suggesting greater generalisability. In addition to its superior performance in all environments, UfO also displays the lowest CV of all methods.

CILO and OPOLO achieved the second and third best performances, 1.0054 and 0.9945, and their CVs were 3.32 and 5.18, respectively. This was not surprising, given that CILO has been shown to

outperform OPOLO in previous experiments [10]. We observed that CILO lost performance when learning policies in environments with more strict clusters of actions (i.e., InvertedPendulum and HalfCheetah), confirming observations in [10]. As for OPOLO, we noticed that performance was lost when the number of teacher samples was increased. We argue that this is due to its behaviour-seeking nature and its reliance on temporal information. As the dataset grows in size, the dimensionality problem becomes more apparent. Increasing the size also increases the chance of the same state appearing in different timesteps and being considered different (due to the temporal information). Yet, requiring fewer samples to learn is theoretically better if the model remains precise. Unfortunately, this is not the case, as pointed out by Gavenski et al. [10] when comparing their method and OPOLO with the same number of sampled trajectories. To further test this, we assessed UfO’s ability to generalise in Sec. 6.3, by evaluating its performance in trajectories with varying initial state distances and its sample efficiency.

GAIFO, BCO, and MAHALO had the worst performances, 0.6627, 0.5899, and 0.3047, respectively. This was expected for GAIFO and BCO, since they have more rudimentary mechanisms for handling trajectory variations, use action-based supervised losses, and adopt behaviour-seeking mode more often. We did not expect their higher CV, despite their lower performance. Usually, as ILfO methods fail to learn the teacher policy, their results also display lower CV (as indeed happened for MAHALO). HalfCheetah was the only environment in which BCO outperformed GAIFO. We interpret this as BC approaches generally being better at solving the environment task, as evidenced by CILO (a BC approach) outperforming OPOLO (which is not). We argue that MAHALO’s poor results are due to its reliance on external data with specific characteristics (see the Supplementary Material for a discussion, including an ablation study to validate our results).

5 THEORETICAL UNDERPINNING FOR CONVERGENCE

Sec. 4 presents the empirical evidence for the convergence of UfO over five benchmarks. We now provide the theoretical explanation of convergence for UfO under the reconstruction and adversarial stages.

5.1 UfO Setup

We first start by considering: (i) π_θ is the policy model, computed by a neural network parametrised by θ . Given a state, it selects the action to be taken; (ii) the generative model \mathcal{G}_ϕ is a neural network parametrised by ϕ . It outputs the state reached by the execution of an action in a state; (iii) the discriminator model \mathcal{D}_ω is a neural network parametrised by ω . Given a trajectory, it outputs a value in the unit interval, estimating the likelihood that a trajectory belongs to the teacher. Furthermore, let $\tau_\pi \sim \mathcal{T}_\pi$ be trajectories generated by a policy π (either π_θ or π_ψ), and \mathcal{L} be the squared loss: $\mathcal{L}(s', \mathcal{G}_\phi(s, \pi_\theta(s))) = \|s' - \mathcal{G}_\phi(s, \pi_\theta(s))\|^2$ (as stated in Sec. 3).

5.2 Reconstruction Stage

We begin by analysing the first stage from UfO: the reconstruction stage. Lemma 5.1 provides the intuition behind local convergence.

LEMMA 5.1. *Let π_θ be the policy model, \mathcal{G}_ϕ be the generative model, and \mathcal{L} be the squared loss, which interleaves every k iteration the parameter updates as follows:*

- $\theta_k \leftarrow \min_\theta \mathbb{E}_{(s,s') \sim \tau_{\pi_\psi}} \mathcal{L}(s', \mathcal{G}_{\phi_{k-1}}(\pi_{\theta_k}(s)))$; and
- $\phi_k \leftarrow \min_\phi \mathbb{E}_{(s,s') \sim \tau_{\pi_\theta}} \mathcal{L}(s', \mathcal{G}_{\phi_k}(\pi_{\theta_k}(s)))$.

We assume that π_θ and \mathcal{G}_ϕ are continuously differentiable with Lipschitz-continuous gradients, and that learning rates are bounded so that parameter updates remain within compact sets.

Then the parameters θ_k and ϕ_k , where k is the k th iteration from the reconstruction stage, converge to a local stationary point (θ^*, ϕ^*) . Thus, the alternating update scheme corresponds to block coordinate descent with differentiable and Lipschitz-continuous objectives. Standard convergence results (e.g., Bertsekas and Tsitsiklis [2] and Razaviyayn et al. [22]) ensure convergence to a stationary point under these assumptions.

Lemma 5.1 ensures convergence of the joint optimisation for the first stage of both neural networks. We now present Lemma 5.2, which justifies the role of \mathcal{G}_ϕ individually within this setup.

LEMMA 5.2. *Let π_{θ_k} be a fixed neural network trained up to the k iteration (θ_k), and $\mathcal{G}_{\phi_{k-1}}$ be a neural network parametrised by ϕ_{k-1} up to the $k-1$ iteration using $\mathbb{E}_{(s,s') \sim \tau_{\pi_\theta}} \mathcal{L}(s', \mathcal{G}_{\phi_{k-1}}(s, \pi_{\theta_k}(s)))$ with sufficient model capacity and training data, the k th iteration of the generative update leads to $\mathcal{G}_{\phi_k}(s, \pi_{\theta_k}(s)) \rightarrow \mathbb{E}[s' | s, \pi_{\theta_k}(s)]$*

In other words, this setup is a supervised regression problem where the target is the next state under the transition dynamics $\mathcal{G}_\phi(s, \pi_\theta(s)) \approx T(s, \pi_\theta(s))$. We note that the reconstruction stage only updates \mathcal{G}_ϕ weights from the environment since it is impossible to predict $\mathcal{G}_\phi(s, \pi_\theta(s))$ with no knowledge of how $T(s, \pi_\theta(s))$ occurs. Therefore, given universal approximation capacity, \mathcal{G}_ϕ converges to the conditional expectation of the MDP's dynamics (T). Moreover, as π_θ improves and generates trajectories closer to those in \mathcal{T}_{π_ψ} , the distribution τ_{π_θ} becomes increasingly aligned with τ_{π_ψ} , improving the generator's data quality.

Now, we present Lemma 5.3, which characterises the role of π_θ individually for the reconstruction stage.

LEMMA 5.3. *Let $\mathcal{G}_{\phi_{k-1}}$ be the fixed generative neural network trained up to the $k-1$ iteration, and π_{θ_k} be the policy neural network parametrised by θ_k up to the $k-1$ iteration. Assume $\mathcal{G}_\phi + \epsilon = T$, where*

ϵ is some error that inevitably decreases as k increases on the support of π_θ . Then, updating θ_k via $\mathbb{E}_{(s,s') \sim \tau_{\pi_\psi}} \mathcal{L}(s', \mathcal{G}_{\phi_{k-1}}(s, \pi_{\theta_k}(s)))$ leads to π_θ producing teacher-like transitions.

If \mathcal{G}_ϕ accurately approximates the environment's transition model, then matching $\mathcal{G}_\phi(s, \pi_\theta(s))$ to the teacher-observed s' encourages π_θ to choose actions that align with the teacher's behaviour. As π_θ improves and produces more teacher-like state transitions, the generator receives better on-policy samples in future iterations, enhancing its approximation of the environment. This recursive structure forms a positive feedback loop that enables consistent *unsupervised* updates to θ . It is important to note that UfO learns the teacher's actions unsupervisedly in an unsupervised way via this process, and π_θ only learns from teacher transitions via $\mathcal{G}_\phi(s, \pi_\theta(s))$. Therefore, no direct action signal is backpropagated to the neural network at any time.

Lemmas 5.2 and 5.3 show that improvements to the generator reinforce the policy's ability to mimic teacher behaviour and vice-versa. We can now explore the convergence of this coupled process.

LEMMA 5.4. *Let (θ_k, ϕ_k) be the parameters at iteration k of the alternating optimisation previously defined. Under boundedness of updates and the assumptions of Lemmas 5.1-5.3, (θ_k, ϕ_k) converges to (θ^*, ϕ^*) such that $\tau_{\pi_\theta} + \epsilon = \tau_{\pi_\psi}$ where ϵ is an acceptable error assuming both τ_{π_θ} and τ_{π_ψ} start in the same state s .*

In other words, Lemma 5.4 follows from the two-timescale stochastic approximation theory [3]. The generator adapts quickly to the current policy, and the policy updates using a fixed generator, where the joint process converges to a consistent fixed point.

5.3 Adversarial Stage

After the reconstruction stage, an adversarial stage refines the policy π_θ using a discriminator \mathcal{D}_ω . Lemma 5.5 explores the discriminator optimisation under this stage.

LEMMA 5.5. *Let τ denote a trajectory, which is a finite sequence of states $\tau = (s_0, s_1, \dots, s_t)$, $p_\pi(\tau)$ denote the probability densities over trajectories by the teacher policy π_ψ and the agent's π_θ , and the discriminator \mathcal{D}_ω be trained under the loss function:*

$$\max_\omega \mathbb{E}_{\tau_{\pi_\psi} \sim \mathcal{T}_{\pi_\psi}} \log \mathcal{D}_\omega(\tau_{\pi_\psi}) + \mathbb{E}_{\tau_{\pi_\theta} \sim \mathcal{T}_{\pi_\theta}} \log(1 - \mathcal{D}_\omega(\tau_{\pi_\theta})).$$

Then, for a fixed policy π_θ and π_ψ , the optimal \mathcal{D}_ω is:

$$\mathcal{D}_\omega^*(\tau) = \frac{p_{\pi_\psi}(\tau)}{p_{\pi_\psi}(\tau) + p_{\pi_\theta}(\tau)}.$$

Thus, minimising the result value function for π_θ corresponds to minimising the Jensen-Shannon divergence, $JSD(\pi_\psi || \pi_\theta)$, between the teacher and agent trajectory distributions.

Lemma 5.5 follows standard results in generative adversarial learning [13], where the optimal discriminator for two distributions is known to be \mathcal{D}_ω^* . The generator (in our work, $\mathcal{G}_\phi(s, \pi_\theta(s))$) trains to match the teacher by minimising the Jensen-Shannon divergence.

While Lemma 5.5 establishes the objective of the adversarial stage, it does not guarantee that the policy remains stable after reconstruction. Therefore, Lemma 5.6 ensures that under mild constraints, fine-tuning does not lead to catastrophic forgetting or divergence.

LEMMA 5.6. Assume that during adversarial training, the generator \mathcal{G}_ϕ remains fixed after reconstruction, the discriminator \mathcal{D}_ω is near-optimal at each step, and the policy π_θ is updated using small, clipped gradients. Then the adversarial stage preserves the behavioural properties of the reconstructed policy and improves generalisation without destabilising earlier learning.

In other words, policy updates are clipped and learning rates are kept small to avoid destabilising the pre-trained policy. Therefore, the fixed generator ensures that the transition model used to compute gradient signals remains stable. With a near-optimal discriminator, the gradients provided to π_θ reliably reflect the discrepancy between agent and expert trajectories. Gradient clipping and a reduced learning rate guarantee that updates to π_θ are conservative, preventing large deviations from the pre-trained policy. Under these conditions, the adversarial stage acts as a regularised adjustment rather than a full retraining, mitigating the risk of catastrophic forgetting [12] while refining behavioural fidelity.

Finally, we note that these results only assume differentiability, Lipschitz continuity, and bounded updates. No convexity assumptions are made. Convergence should be interpreted as local convergence to stationary points, which are typical in deep learning settings.

6 DISCUSSION

In this section, we consider some key aspects of UfO’s behaviour: (i) the interaction between generator and policy in the first stage; (ii) how the adversarial step improves UfO overall results; and (iii) how sample efficient UfO is.

6.1 Teacher Action Approximation

It is vital for UfO to estimate action effects using solely the loss in Eq. 4 closely via the process described in Sec. 3.1. If the signal from the generator is not strong enough and the policy is unable to approximate its predictions, or the generator is unable to learn a transition function close enough to the MDP’s, the adversarial stage will not work, and UfO will not achieve the desired results.

In Fig. 2, we show the training behaviour for the Hopper-v4 environment during the reconstruction stage with a smoothing function for better visibility. It displays the error from π_θ according to the teacher’s dataset, the generative error from \mathcal{G}_ϕ according to the teacher’s dataset (split into training and evaluation) and the agent’s collected samples, and the reward accumulated for the agent during each training step. As expected, the generator learns a transition function that approximates the actual MDP transition function quite early in the training phase (purple line). However, for the generator to see transitions closer to the teacher’s, the policy has to improve its error, and we notice a direct correlation between the errors of π_θ and \mathcal{G}_ϕ . We also observe that as the errors of both models decrease, the rewards increase, which shows that Eq. 4 and 5 yield the expected outcome. Yet, we note that the policy error stagnates after a certain number of epochs, potentially indicating that the gradient signal from the generator is no longer enough to improve the policy’s predictions further. This behaviour is expected since the generator is not optimised to improve the policy’s predictions themselves, but rather to learn a transition function that can approximate the teacher’s and the MDP’s. Thus, the adversarial stage is crucial for further improving the policy’s predictions and for stabilising the




Figure 2: Training information for UfO in Hopper-v4.

model’s behaviour. Finally, we believe that this stage could be further improved by using a curriculum learning approach. In this approach, the \mathcal{G}_ϕ would receive more samples it predicted poorly, and the policy would receive more meaningful weight updates.

6.2 Role of Each Stage

Tab. 2 shows UfO and teacher results in the Half-Cheetah-v4 environment, the only in which no baseline achieved teacher performance. This environment clearly displays the role of UfO’s two stages. The reconstruction stage achieves performance above 1, with over 307 more reward points than the teacher. The AER of 9819.6491 here with standard deviation of 626.0181 are already superior to all other methods’ (see Tab. 1). However, the standard deviation is significantly higher than the teacher’s (by 12.59%). This is not ideal because we want a model that can perform at least as well as the teacher in all different environment settings. UfO’s adversarial stage improves \mathcal{P} and AER further, whilst reducing its standard deviation by 5.85%.

Step	AER (\uparrow)	CV (\downarrow)	\mathcal{P} (\uparrow)
Teacher (π_ψ)	9512.2995 \pm 538.5918	5.66%	1.0000
Reconstruction (Sec. 3.1)	9819.6491 \pm 626.0181	6.38%	1.0314
Adversarial (Sec. 3.2)	9959.5745 \pm 589.3388	5.92%	1.0457

Table 2: HalfCheetah-v4 environment results.

As mentioned in Sec. 3.2, the adversarial stage (Algorithm 1, Line 16–22) might lead to behavioural-seeking mode and a catastrophic forgetting of the policy. Therefore, it is preferable to run this stage for only a few epochs ($\text{epochs}_a \leq 10$) to reduce the policy’s standard deviation through online play and slightly improve the agent’s performance. Given that it only runs for a few epochs, it offers an excellent trade-off between time and policy improvement.

6.3 Sample Efficiency

If a large number of teacher samples is required for UfO to achieve teacher performance, using it in environments where sample collection is challenging or costly may be impractical. However, performance decay with an increase in the number of samples indicates an issue with the behaviour-seeking aspect. To see how UfO handles




Figure 3: UfO’s CV and AER for the HalfCheetah-v4.

varying amounts of samples, we experimented with 10 to 900 trajectories. Fig. 3 shows the reconstruction stage results for all experiments in the HalfCheetah-v4 environment.

Although UfO does not achieve teacher performance on 10 trajectories (unlike OPOLO and CILO), its AER and \mathcal{P} improve when increasing training data. It eventually surpasses teacher performance with 300 trajectories, albeit with a standard deviation higher than all other baselines. As we increase the number of trajectories further, we observe a diminishing return from UfO’s AER. This would be a significant drawback were the method not to become more generalisable. UfO’s CV keeps decreasing up to around 800 trajectories, when the method seems to overfit with a simultaneous increase in variation and AER. We conclude that the ideal number of samples is the one that maximises AER while minimising CV (i.e., maximising generalisability). Although other methods achieve performance comparable to their teacher on smaller datasets, UfO is the only one that consistently improves generalisation as the dataset size increases.

7 RELATED WORK

The simplest form of imitation learning is BC [21], which addresses the IL problem as a supervised learning problem. It uses teacher samples of the form (s, a, s') , where $s' = T(s, a)$, to learn how to approximate the teacher’s policy. However, supervised approaches are costly for more complex scenarios, requiring more samples about the actions’ effects in the environment. For example, solving Atari games requires approximately 100 times more samples than control tasks, such as CartPole. Generative Adversarial Imitation Learning (GAIL) [14] matches the state-action frequencies from the agent to those seen in the training data, creating a policy with action distributions closer to the teacher. GAIL leverages the data structure more efficiently, reducing its dependency on large quantities of data and making it more suitable for complex environments. It uses adversarial training to discriminate state-action pairs from the agent or the teacher while minimising the difference between them.

Recent self-supervised approaches [19, 25] that learn from observations use teacher’s transitions $(s_{\pi_{\psi}}, s'_{\pi_{\psi}})$ and leverage random transitions $(s_{\pi_{\xi}}, s'_{\pi_{\xi}})$ to learn the IDM of the environment, and afterwards, generate self-supervised labels for the teacher’s trajectories. Generative Adversarial Imitation Learning from Observation (GAIfo) [26] is an ILfO method based on GAIL. The discriminator

is trained using state transitions, and the policy is optimised using an inverse reinforcement learning strategy. Imitating Latent Policies from Observation (ILPO) [5] differs from such work by trying to estimate the probability of a latent action given a state. Within some environment steps, it remaps latent actions. More recently, Off-Policy Learning from Observations (OPOLO) [28] used a dual-form of the expectation function and an adversarial structure to achieve off-policy ILfO. Modality-agnostic Adversarial Hypothesis Adaptation for Learning from Observations (MAHALO) [17] unifies ILfO with offline reinforcement learning by learning pessimistic value estimates through adversarially trained reward and critic functions. It accommodates partially labelled datasets and generalises across several ILfO settings. Continuous Imitation Learning from Observation (CILO) [10] used the same IDM structure of earlier methods but introduced an exploration mechanism to further improve its IDM’s predictions. Path signatures were used as non-parametric encodings to identify trajectories that behave similarly to the teacher efficiently.

8 CONCLUSION

In this work, we proposed Unsupervised Imitation Learning from Observation (UfO), an unsupervised ILfO method addressing critical limitations of existing approaches. Firstly, UfO eliminates reliance on action-based supervised optimisation, reducing the risk of overfitting. Secondly, it mitigates behaviour-seeking that often lead to suboptimal performance. Lastly, UfO provides generalisation by more effectively handling environments where transitions can result from multiple actions. It learns to approximate the teacher’s policy and refine it through adversarial learning. A two-stage training process ensures the agent captures accurate state transitions and aligns its policy with the teacher’s through online interactions. By leveraging transitions rather than direct state-action mappings, UfO avoids common ILfO pitfalls and achieves superior results across diverse initial states.

Our experiments highlight the effectiveness of UfO, consistently surpassing the teacher in five benchmark environments. Unlike traditional methods, UfO achieves state-of-the-art results without sacrificing robustness, making it suitable for real-world scenarios where action labels are unavailable or costly. It demonstrates strong generalisation, maintaining low trajectory variance as evidenced by its coefficient of variation (2.64%). Beyond imitation learning, UfO offers a general framework for robust policy learning from indirect supervision. Its proxy-based learning approach is applicable to offline RL, sim-to-real transfer, and causal representation learning in domains such as healthcare and economics. Moreover, the refinement stage resembles curriculum learning and can be adapted to enhance sample efficiency and robustness in safety-critical systems.

Future work will explore UfO’s application to more complex and dynamic environments, further enhancing sample efficiency through advanced exploration strategies and expanding its applicability as a foundation for observational learning in autonomous systems.

ACKNOWLEDGMENTS

This work was supported by UK Research and Innovation [grant number EP/S023356/1], in the UKRI Centre for Doctoral Training in Safe and Trusted Artificial Intelligence (www.safeandtrusted.ai.org) and made possible via King’s Computational Research, Engineering and Technology Environment (CREATE) [4].

REFERENCES

- [1] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. 2019. Optuna: A Next-generation Hyperparameter Optimization Framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (Anchorage, AK, USA) (*KDD '19*). Association for Computing Machinery, New York, NY, USA, 2623–2631. <https://doi.org/10.1145/3292500.3330701>
- [2] Dimitri P. Bertsekas and John N. Tsitsiklis. 2000. Gradient Convergence in Gradient methods with Errors. *SIAM J. Optim.* 10, 3 (2000), 627–642. <https://doi.org/10.1137/S1052623497331063>
- [3] Vivek S. Borkar. 2008. *Stochastic Approximation: A Dynamical Systems Viewpoint*. Springer. <https://api.semanticscholar.org/CorpusID:118540541>
- [4] King’s College London e Research team. 2023. King’s computational research, engineering and technology environment (CREATE). <https://doi.org/10.18742/mvf-m076>
- [5] Ashley Edwards, Himanshu Sahni, Yannick Schroecker, and Charles Isbell. 2019. Imitating Latent Policies from Observation. In *Proceedings of the 36th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 97)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.). PMLR, 1755–1763. <https://proceedings.mlr.press/v97/edwards19a.html>
- [6] Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. 2020. D4RL: Datasets for Deep Data-Driven Reinforcement Learning. [arXiv:2004.07219](https://arxiv.org/abs/2004.07219) [cs.LG]
- [7] Divyansh Garg, Shuvam Chakraborty, Chris Cundy, Jiaming Song, and Stefano Ermon. 2021. Iq-learn: Inverse soft-q learning for imitation. *Advances in Neural Information Processing Systems* 34 (2021), 4028–4039.
- [8] Nathan Gavenski, Michael Luck, and Odinaldo Rodrigues. 2024. Imitation Learning Datasets: A Toolkit For Creating Datasets, Training Agents and Benchmarking. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems* (Auckland, New Zealand) (*AAMAS '24*). International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2800–2802.
- [9] Nathan Gavenski, Juarez Monteiro, Adilson Medronha, and Rodrigo C. Barros. 2022. How Resilient Are Imitation Learning Methods to Sub-optimal Experts?. In *Intelligent Systems*, João Carlos Xavier-Junior and Ricardo Araújo Rios (Eds.). Springer International Publishing, Cham, 449–463.
- [10] Nathan Gavenski, Juarez Monteiro, Felipe Meneguzzi, Michael Luck, and Odinaldo Rodrigues. 2024. Explorative Imitation Learning: A Path Signature Approach for Continuous Environments. In *Proceedings of the 27th European Conference on Artificial Intelligence* (Santiago de Compostela, Spain) (*ECAI '24, Vol. 392*). IOS Press, 1551–1558. <https://doi.org/10.3233/FAIA240660>
- [11] Nathan Gavenski, Odinaldo Rodrigues, and Michael Luck. 2024. Imitation Learning: A Survey of Learning Methods, Environments and Metrics. <https://doi.org/10.48550/ARXIV.2404.19456> [arXiv:2404.19456](https://arxiv.org/abs/2404.19456)
- [12] Ian J. Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. 2015. An Empirical Investigation of Catastrophic Forgetting in Gradient-Based Neural Networks. [arXiv:1312.6211](https://arxiv.org/abs/1312.6211) [stat.ML] <https://arxiv.org/abs/1312.6211>
- [13] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger (Eds.), Vol. 27. Curran Associates, Inc., 2672–2680. https://proceedings.neurips.cc/paper_files/paper/2014/file/f033ed80deb0234979a61f95710dbe25-Paper.pdf
- [14] Jonathan Ho and Stefano Ermon. 2016. Generative Adversarial Imitation Learning. In *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett (Eds.), Vol. 29. Curran Associates, Inc., 4565–4573. <https://proceedings.neurips.cc/paper/2016/file/cc7e2b878868cbae992d1fb743995d8f-Paper.pdf>
- [15] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. [arXiv preprint arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014).
- [16] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences* 114, 13 (March 2017), 3521–3526. <https://doi.org/10.1073/pnas.1611835114>
- [17] Anqi Li, Byron Boots, and Ching-An Cheng. 2023. MAHALO: Unifying Offline Reinforcement Learning and Imitation Learning from Observations. In *Proceedings of the 40th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 202)*, Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (Eds.). PMLR, 19360–19384. <https://proceedings.mlr.press/v202/li23b.html>
- [18] Yecheng Jason Ma, Andrew Shen, Dinesh Jayaraman, and Osbert Bastani. 2022. Smodice: Versatile offline imitation learning via state occupancy matching. [arXiv preprint arXiv:2202.02433](https://arxiv.org/abs/2202.02433) 1, 2 (2022), 3.
- [19] Juarez Monteiro, Nathan Gavenski, Felipe Meneguzzi, and Rodrigo C. Barros. 2023. Self-Supervised Adversarial Imitation Learning. In *International Joint Conference on Neural Networks, IJCNN 2023, Gold Coast, Australia, June 18-23, 2023*. IEEE, 1–8. <https://doi.org/10.1109/IJCNN54540.2023.10191197>
- [20] Antonio Paolillo, Paolo Robuffo Giordano, and Matteo Saveriano. 2023. Dynamical System-based Imitation Learning for Visual Servoing using the Large Projection Formulation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 755–761. <https://doi.org/10.1109/ICRA48891.2023.10160935>
- [21] Dean A. Pomerleau. 1988. ALVINN: An Autonomous Land Vehicle in a Neural Network. In *Advances in Neural Information Processing Systems*, D. Touretzky (Ed.), Vol. 1. Morgan-Kaufmann. https://proceedings.neurips.cc/paper_files/paper/1988/file/812b4ba287f5ee0bc9d43bbf5bbe87fb-Paper.pdf
- [22] Meisam Razaviyayn, Maziar Sanjabi, and Zhi-Quan Tom Luo. 2013. A Stochastic Successive Minimization Method for Nonsmooth Nonconvex Optimization with Applications to Transceiver Design in Wireless Communication Networks. *Mathematical Programming* 157 (2013), 515–545. <https://api.semanticscholar.org/CorpusID:2512169>
- [23] Richard S. Sutton and Andrew G. Barto. 2018. *Reinforcement Learning: An Introduction* (2nd ed.). MIT Press, Cambridge, MA. <https://search.worldcat.org/title/1043175824>
- [24] Gokul Swamy, Sanjiban Choudhury, J. Andrew Bagnell, and Steven Wu. 2021. Of Moments and Matching: A Game-Theoretic Framework for Closing the Imitation Gap. In *Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 10022–10032. <https://proceedings.mlr.press/v139/swamy21a.html>
- [25] Faraz Torabi, Garrett Warnell, and Peter Stone. 2018. Behavioral Cloning from Observation. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence* (Stockholm, Sweden) (*IJCAI '18*). AAAI Press, 4950–4957. <https://www.ijcai.org/proceedings/2018/0687.pdf>
- [26] Faraz Torabi, Garrett Warnell, and Peter Stone. 2018. Generative Adversarial Imitation from Observation. [arXiv:1807.06158v4](https://arxiv.org/abs/1807.06158v4) [cs.LG]
- [27] Omar G. Younis, Rodrigo Perez-Vicente, John U. Balis, Will Dudley, Alex Davey, and Jordan K Terry. 2024. *Minari*. Farama Foundation. <https://doi.org/10.5281/zenodo.13767625>
- [28] Zhuangdi Zhu, Kaixiang Lin, Bo Dai, and Jiayu Zhou. 2020. Off-Policy Imitation Learning from Observations. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 12402–12413. <https://proceedings.neurips.cc/paper/2020/file/92977ae4d2ba21425a59afb269c2a14e-Paper.pdf>