
Imitation by Predicting Observations

Andrew Jaegle¹ Yury Sulsky¹ Arun Ahuja¹ Jake Bruce¹ Rob Fergus¹ Greg Wayne¹

Abstract

Imitation learning enables agents to reuse and adapt the hard-won expertise of others, offering a solution to several key challenges in learning behavior. Although it is easy to observe behavior in the real-world, the underlying actions may not be accessible. We present a new method for imitation solely from observations that achieves comparable performance to experts on challenging continuous control tasks while also exhibiting robustness in the presence of observations unrelated to the task. Our method, which we call FORM (for “Future Observation Reward Model”) is derived from an inverse RL objective and imitates using a model of expert behavior learned by generative modelling of the expert’s observations, without needing ground truth actions. We show that FORM performs comparably to a strong baseline IRL method (GAIL) on the DeepMind Control Suite benchmark, while outperforming GAIL in the presence of task-irrelevant features.

1. Introduction

The goal of imitation is to learn to produce behavior that matches that of an expert on unseen data, given demonstrations of the expert’s behavior (Abbeel & Ng, 2004; Osa et al., 2018). The field of imitation learning offers tools for learning behavior when programmed rewards cannot be provided, or when rewards can only be partially or sparsely specified. Imitation learning has been at the heart of several breakthroughs in building AI agents (Pomerleau, 1989; Abbeel et al., 2010; Silver et al., 2016; Vinyals et al., 2019; OpenAI et al., 2019), allowing agents to learn even when faced with hard exploration problems (Gulcehre et al., 2020).

There is widespread evidence that imitation (among other forms of social learning) is a core mechanism by which humans and other animals learn to acquire a sophisticated behavioral repertoire (Tomasello, 1996; Laland, 2008; Byrne,

2009; Huber et al., 2009). While most algorithms for imitation learning assume that demonstrations contain the actions the expert executed, animals must imitate without directly observing what actions the expert took (i.e. without knowing exactly what commands were issued to produce the observable changes). In the context of machine learning, solving the problem of *imitation from observation* is a key step towards the tantalizing possibility of learning behavior from unlabeled and easy to collect data, such as raw video footage of human activity. Many recent algorithms for imitation have focused on addressing the problem of imitation in very small data regimes, but the challenge in imitating from these abundant sources of data is not primarily one of quantity. The challenge is rather how to learn models for imitation that are general enough to learn and generalize from data that depicts a rich (and unknown) reward structure. In this work, we show how predictive generative models can be used to learn a general reward model from observations alone.

Current state-of-the-art approaches to imitation (including from observation) pose learning as an adversarial game: a classifier estimates the probability that a state is visited by the expert or imitator, and the policy seeks to maximize the classifier error (Merel et al., 2017; Torabi et al., 2019a). Because these methods are based on matching the expert’s occupancy using a fixed dataset of demonstrations, they tend to be very sensitive to the precise details of the demonstrations and to the representation used. This property makes learning with adversarial methods difficult when using raw, noisy observations without extensive tuning and careful use of strong forms of regularization (Peng et al., 2019), domain or task knowledge (Zolna et al., 2020), or a combination of behavioral cloning and careful representation design (Abramson et al., 2020).

In this work, we introduce the future observation reward model (FORM) (see Figure 1), which address the problem of imitation from observation while exhibiting both (1) generality and expressiveness by coupling predictive generative models with inverse RL (IRL) and (2) improved robustness by foregoing an adversarial formulation. In FORM, the imitator tries to match the probability of observation sequences in the expert data. It does so using a learned generative model of expert observation sequences and a learned generative model of its own observation sequences. In other words,

¹DeepMind. Correspondence to: Andrew Jaegle <drewjaegle@deepmind.com>.

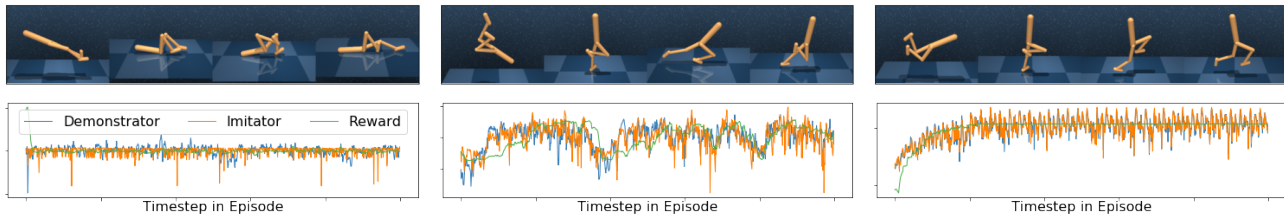


Figure 1. FORM learns to imitate expert behavior using sequences of internal state observations, without access to the expert’s actions. Visualizations of agent behavior (top) and reward curves for a single episode (bottom) are shown after 0 (left), 50k (middle) and 5M update steps. FORM imitates using two learned models: both the **demonstrator** model (trained offline) and the **imitator** model (trained online) log-likelihoods track the unseen **task reward** as the imitation agent learns. Agent behavior is shown as images, but we use lower-dimensional internal state observations in this work.

FORM casts the problem of learning from demonstrations as a sequence prediction problem, using a generative model of expert sequences to guide RL. Because FORM builds separate models of expert and imitator sequences, rather than using a single classifier to discriminate expert and imitator states, it is less prone to focus on irrelevant differences between the expert and imitator demonstrations. The structure of the FORM objective makes it theoretically straightforward to optimize using standard policy optimization tools and, as we show, empirically competitive on the DeepMind Control Suite continuous control benchmark domain.

This stands in contrast to adversarial methods, such as Generative Adversarial Imitation Learning (GAIL), whose objectives are known to be ill-posed (without additional regularization) and challenging to optimize both in theory and in practice (Arjovsky et al., 2017; Gulrajani et al., 2017; Mescheder et al., 2017). This property makes it difficult to apply adversarial techniques to imitation in settings with even small differences between expert and imitator settings (Zolna et al., 2020). Robustness to distractors is an important part of behavior learning, as recently illustrated by Stone et al. 2021 in the context of RL with image background distractors. These situations are common in practice: the lab environment where expert data is collected for a robot will be quite different to where it might be deployed. While it may be possible to collect a large number of demonstrations, it is impossible to exhaustively sample all possible sources of differences between the two domains (such as the surface texture, robot physical parameters, or environment appearance). These differences confound the signal that must be imitated, leading to the risk of spurious dependencies between the two being learned. As we will show, FORM exhibits greater robustness than a well-tuned adversarial imitation method, GAIL from Observation, or GAIfo (Torabi et al., 2019a) in presence of task-independent features.

We make the following technical contributions in this work:

1. We derive the FORM reward from an objective for inverse reinforcement learning from observations. We show that this reward can be maximized using gen-

erative models of expert and imitator behavior with standard policy optimization techniques.

2. We develop a practical algorithm for imitation learning using the FORM reward and demonstrate that it performs competitively with a well-tuned GAIfo model on the DeepMind Control Suite benchmark.
3. We show that FORM is more robust than GAIfo in the presence of extraneous, task-irrelevant features, which simulate domain shift between expert and imitator settings.

2. Background and related work

RL, IRL, and imitation Reinforcement learning is concerned with learning a policy that maximizes the expected return, which is given as the expected sum of all future discounted rewards (Sutton & Barto, 2018), which are typically observed. In imitation learning, on the other hand, we are not given a reward function, but we do have access to demonstrations produced by a demonstrator (or expert) policy, which maximizes some (unobserved) expected return.

IRL has the related goal of recovering the unobserved reward function from expert behavior. IRL offers a general formula for imitation: estimate the reward function underlying the demonstration data (a “reward model”) and maximize this reward by RL (Ng & Russell, 2000), possibly iterating multiple times until convergence. Alternative approaches to imitation, such as behavioral cloning (BC) (Pomerleau, 1989) or BC from observations (BCO) (Torabi et al., 2018), typically have difficulty producing reliable behavior away from configurations seen in the expert demonstrations. This is because small errors in predicting actions or mimicking short-term agent behavior accumulates over long behavioral timescales.¹ IRL methods like FORM avoid this problem: because they perform RL on a learned reward, they can learn through experience to recover from mistakes by focusing on long-term consequences of each action.

¹The standard solution to this problem for BC assumes access to an expert policy that can be repeatedly queried (Ross et al., 2011), which is not always feasible.

GAIL and occupancy-based IRL Most contemporary IRL-based approaches to imitation – as exemplified by GAIL – use a strategy of state-action occupancy matching, typically by casting imitation as an adversarial game and learning a classifier to discriminate states and actions sampled uniformly from the expert demonstrations from those encountered by the imitator (Ho & Ermon, 2016; Torabi et al., 2019a; Fu et al., 2018; Kostrikov et al., 2019; Ghasemipour et al., 2019). In contrast, rather than classifying states as belonging to the expert or imitator, FORM learns to imitate using separate generative models of expert and imitator behavior. This means that FORM is built on predictive models of the form $p(x_t|x_{t-1})$, where x s are observations, rather than a single model of the form $p(\text{expert}|x)$ that tries to classify observations as generated by the expert or not. FORM’s objective is similar in spirit to classical feature-matching and maximum-entropy formulations of imitation (Ng & Russell, 2000; Abbeel & Ng, 2004; Ziebart et al., 2008), while also providing a fully probabilistic interpretation and making minimal assumptions about the environment (the FORM objective does not require an MDP or deterministic transitions).

Other related methods for imitation Other recent work has used generative models in the context of imitation learning: this work typically retains GAIL’s occupancy-based perspective (Baram et al., 2016; Jarrett et al., 2020; Liu et al., 2021) or introduces a generative model to provide a heuristic reward (Yu et al., 2020). Unlike FORM, which uses effect models (see Figure 2) that are suitable for imitation from observations, this work models quantities that are useful primarily in conjunction with actions (modeling state-action densities and/or dynamics models for GAIL augmentation). Other recently proposed methods learn reward models either purely or partially offline (Kostrikov et al., 2020; Jarrett et al., 2020; Arenz & Neumann, 2020). This approach leans on the presence of actions in the demonstrator data. Although FORM’s demonstrator effect model is learned offline, FORM’s online phase is essential to the process of distilling an effect model (which doesn’t use actions) into a policy (which does).

Learning to act from observations Many methods have been proposed for imitation from observation (Torabi et al., 2019b), but most methods that do so using IRL are based around GAIL (Wang et al., 2017; Torabi et al., 2019a; Sun et al., 2019). Recent work has obtained interesting results using solutions based around tracking or matching trajectories in learned feature spaces (Peng et al., 2018; Merel et al., 2019), by matching imitator actions to learned models of expert trajectories (interpreted as inverse models of the expert action) (Schmeckpeper et al., 2020; Zhu et al., 2020; Edwards et al., 2018; Pathak et al., 2018), and by learning to match features in learned invariant spaces (Sermanet et al., 2017). Finally, we note that much recent work has observed

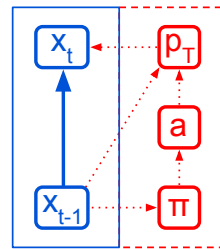


Figure 2. FORM’s demonstrator and imitator effect models are *effect models*, generative models $p(x_t|x_{t-1})$ of the change in observation (*observed*) produced by a policy π in an environment with transition dynamics p_T (*unobserved*). The models used in model-based RL are usually of the form $p(x_t|x_{t-1}, a_{t-1})$ and aim to model transition dynamics rather than the full distribution of outcomes given a policy.

that the structure of observation sequences can be exploited to generate behavior, whether in the context of language modeling (Brown et al., 2020), 3D navigation (Dosovitskiy & Koltun, 2017), few-shot planning (Rybkin et al., 2019), or value-based RL (Edwards et al., 2020). FORM uses generative models of future observations to exploit this property of observation transitions and connect it to inverse reinforcement learning to produce a practical algorithm for imitation.

3. Approach

3.1. Inverse reinforcement learning from observations

Our goal is to learn a policy that produces behavior like an expert (or demonstrator) by IRL, using only observation. Historically, the IRL procedure has been framed as matching the expected distribution over states and actions (or their features) along the imitator and demonstrator paths (Ng & Russell, 2000; Abbeel & Ng, 2004; Ziebart et al., 2008). As also noted in (Arenz & Neumann, 2020), we can express this as a divergence minimization problem:

$$\min_{\theta} \mathcal{D}_{KL}[p_{\theta}^I(\tau) || p^D(\tau)], \quad (1)$$

where $\tau = \{x_0, a_0, x_1, a_1, \dots, a_{T-2}, x_{T-1}\}$ is a trajectory consisting of actions $A = \{a_0, \dots, a_{T-1}\}$ and states $X = \{x_0, \dots, x_{T-1}\}$. We use x rather than o (for observation) or s (for state) because FORM does not assume that its inputs are Markovian – FORM applies to generic observations – but speaking in terms of states simplifies the comparison to other methods (like GAIL) that assume Markovian states are given or inferred. $p^D(\tau)$ is the distribution over trajectories induced by the demonstrator’s policy and the environment dynamics, while $p_{\theta}^I(\tau)$ is the corresponding distribution induced by an imitator with learnable parameters θ .

In imitation learning from observation, the imitator must reason about the demonstrator’s behavior without supervised access to the expert’s actions (its control signals). Accord-

ingly, we focus on distributions over observation sequences, which amounts to integrating out the imitator’s actions:

$$\begin{aligned} p_\theta^I(X) &= \int_A p_\theta^I(\tau) = \int_A p_\theta^I(A, X) \\ &= \int_A \prod_{t \geq 0} p(x_t | x_{<t}, a_{<t}) \pi_\theta^I(a_{t-1} | x_{<t}, a_{<t-1}). \end{aligned} \quad (2)$$

This density reflects both the environment transition dynamics $p(x_t | x_{t-1}, a_{t-1})$ and the imitator policy $\pi_\theta(a_t | x_t)$, whose parameters θ we seek to learn. Similarly, we can write the probability of a demonstrator trajectory in terms of the unobserved expert policy as

$$\begin{aligned} p^D(X) &= \prod_{t \geq 0} p^D(x_t | x_{<t}) \\ &= \int_A \prod_{t \geq 0} p(x_t | x_{<t}, a_{<t}) \pi^D(a_{t-1} | x_{<t}, a_{<t-1}). \end{aligned} \quad (3)$$

Our objective is to minimize the KL-divergence² between these two densities:

$$\begin{aligned} \min_\theta \mathcal{D}_{KL}[p_\theta^I(X) || p^D(X)] \\ = \min_\theta \mathbb{E}_{p_\theta^I(X)} \left[\log p_\theta^I(X) - \log p^D(X) \right]. \end{aligned} \quad (4)$$

Minimizing the divergence corresponds to maximizing the following expression in expectation:

$$\rho_{\text{FORM}} = \log p^D(X) - \log p_\theta^I(X). \quad (5)$$

In this work, we propose to imitate by treating ρ_{FORM} as a return and maximizing it using RL.

3.2. Optimizing FORM with effect models

To see how we will capture this expression, first note that each term of ρ_{FORM} is a log-density over the states encountered in an episode $\log p(X)$, which we can rewrite as $\log p(x_0) + \sum_{t>0} \log p(x_t | x_{<t})$ using the chain rule for probability. As the initial state is independent of the policy, we can simplify the expression used in each reward term to $\sum_{t \geq 0} \log p(x_t | x_{t-1})$. This means the return can be expressed solely in terms of next-step conditional densities. To simplify the discussion, we present all results from here forward in terms of one-step predictive models $\log p(x_t | x_{t-1})$, but the FORM derivation and algorithm applies equally well to generic sequence models $\log p(x_t | x_{<t})$.

We propose to learn a reward model by introducing models of the state transition densities under (1) the demonstrator

²We use the reverse KL because the policy learns on its own trajectories, as in RL (Levine, 2018).

$p^D(x_t | x_{t-1})$ and (2) the imitator $p_\theta^I(x_t | x_{t-1})$. We refer to these as *effect models* to differentiate them from how “model” is used elsewhere in the RL literature to refer to models of transition dynamics (Figure 2). Unlike transition models, which are typically action-conditional and are assumed to model policy-independent transition dynamics, effect models are not conditioned on actions and attempt to capture the effects of policy and environment dynamics. A similar class of models was used to model an expert’s behavior in recent work (Rhinehart et al., 2020).

Algorithm 1 Imitation learning with FORM

Input: A fixed dataset \mathbb{D} of expert state transitions, a replay buffer to fill with imitator data, an environment.

Init: Randomly initialize demonstrator effect model $p_\omega^D(x_t | x_{t-1})$, imitator effect model $p_\phi^I(x_t | x_{t-1})$, and imitator policy $\pi_\theta^I(a_t | x_t)$.

while p_ω^D not converged **do**

 # Train demonstrator effect model

 Sample batch of trajectories from the expert dataset \mathbb{D} .

 Update p_ω^D by taking a gradient step (e.g. with Adam) on:

$$\max_\omega \mathbb{E}_D \left[\sum_{t \geq 0} \log p_\omega^D(x_t | x_{t-1}) \right].$$

end

while $\pi_\theta^I(a_t | x_t)$ not converged **do**

 # Train imitator effect model and policy

 Sample trajectories from the environment using π_θ^I and add them to the replay buffer.

 Sample batch of trajectories from the replay buffer.

 Label the reward of each sampled transition (x_{t-1}, a_{t-1}, x_t) using $p_\omega^D(x_t | x_{t-1})$ and $p_\phi^I(x_t | x_{t-1})$:

$$r_t = \log p_\omega^D(x_t | x_{t-1}) - \log p_\phi^I(x_t | x_{t-1}).$$

 Update $\pi_\theta^I(a_t | x_t)$ with a step of a policy improvement algorithm (e.g. with MPO) using returns computed from the reward-labeled trajectories (e.g. with Re-trace).

 Update $p_\phi^I(x_t | x_{t-1})$ by taking a gradient step (e.g. with Adam) on:

$$\max_\phi \mathbb{E}_I \left[\sum_{t \geq 0} \log p_\phi^I(x_t | x_{t-1}) \right].$$

end

We wish to maximize this return using standard tools for policy optimization. We can do this without introducing bias only if the policy gradients do not depend on gradients of any term in the reward (which aren’t accounted for by

standard policy optimizers). As we derive in Sec. A of the appendix, this assumption holds, and we can write the policy gradient as:

$$\nabla_{\theta} \mathcal{J}_{\text{FORM}}(\pi_{\theta}^I) = \mathbb{E}_{\tau \sim \pi_{\theta}^I} [\rho_{\text{FORM}} \sum_{t \geq 0} \nabla_{\theta} \log \pi_{\theta}^I(a_t | x_t)]. \quad (6)$$

Intuitively, the policy gradient does not involve gradients of either p_{θ}^I or p^D because neither of these densities are conditioned on the actions sampled from the policy (in effect, the contribution of the density to the policy gradient is integrated out). Because the demonstrator effect model is independent of the imitator, we can train it offline on expert demonstrations using a maximum likelihood objective:

$$\max_{\omega} \mathbb{E}_{p^D(X)} \left[\sum_{t \geq 0} \log p_{\omega}^D(x_t | x_{t-1}) \right]. \quad (7)$$

The model of the imitator density $\log p_{\theta}^I(X)$, on the other hand, needs to capture the transition density under the current policy (it acts as a self-model). Accordingly, we train it by taking stochastic gradient descent steps on the following objective at the same time as the imitator policy is training:

$$\max_{\phi} \mathbb{E}_{p_{\theta}^I(X)} \left[\sum_{t \geq 0} \log p_{\phi}^I(x_t | x_{t-1}) \right]. \quad (8)$$

By incorporating both models, we obtain the full FORM policy objective:

$$\max_{\theta} \mathbb{E}_{\pi_{\theta}^I(X)} \left[\sum_{t \geq 0} \log p_{\omega}^D(x_t | x_{t-1}) - \log p_{\phi}^I(x_t | x_{t-1}) \right]. \quad (9)$$

Despite the inclusion of two terms with opposite signs, the FORM policy objective is **not** an adversarial loss: FORM is based on a KL-minimization objective, rather than an adversarial minimax objective, and is not formulated as a zero-sum game. The second term in the objective can be viewed as an entropy-like expression, similar to the one that arises in maximum-entropy RL (Levine, 2018).

This objective includes both an expectation with respect to the current imitator policy and a term that reflects the current imitator effect model. This suggests that this objective is easiest to optimize in an on-policy setting. Nonetheless, we find that it can be stably optimized in a moderately off-policy setting. In all experiments here, we sample transitions from a replay buffer, computing rewards as they are consumed. We compute returns using the Retrace algorithm on the raw rewards (Munos et al., 2016) (which corrects for mildly off-policy actions using importance sampling). We optimize the policy using the MPO algorithm (Abdolmaleki et al., 2018). We choose MPO because it is known to perform well in mildly off-policy settings: FORM itself does not make any MPO-specific assumptions, and we expect it to perform well with many other policy optimizers. We describe our full procedure in Algorithm 1.

3.3. GAIL, occupancy-based imitation, and robustness

GAIL and its variants are justified in terms of matching the state-action occupancy of an expert – GAIL attempts to unconditionally match the rates at which states and actions are visited – rather than directly matching a policy or its effects. In contrast, FORM’s reward is derived directly from an objective that matches a policy’s effects on a sequence (eq. 5). This has consequences for their robustness, as we will explain.

First, note that a policy is a local concept (it describes how to map states or observations to actions), while an occupancy is a global concept (it describes the rates at which an agent visits states and actions in expectation). To see why the occupancy is global, note that the occupancy (Ho & Ermon, 2016; Torabi et al., 2019a) of a state x_i by a policy π is given by $\rho_{\pi}(x_i) = \sum_{t=0}^{\infty} \gamma^t p(x_t = x_i | \pi)$, where in general:

$$p(x_t = x_i | \pi) = \int_{\{x, a\} < t} p(x_t = x_i | x_{<t}, a_{<t}) p(x_{<t}, a_{<t} | \pi) \quad (10)$$

In other words, to reason about the occupancy of a state is to reason about every possible way the policy might arrive there. In practice, for GAIL, the discriminator computes a state’s reward by comparing the frequency of x_t to the frequency of all other states that are seen in the data, whatever the conditions under which that state was produced. Because FORM relies on conditional probabilities and does not depend on long-horizon visitation in its derivation, the only relevant states are those that appear under similar conditions. Essentially, FORM’s reward involves comparisons to fewer observations because it takes a state’s context – namely, the transition that produced it – into account.

We expect this property to mitigate GAIL’s sensitivity to noise. It’s easiest to see why this should happen by comparing GAIfo and FORM for two-state inputs. Here, FORM maximizes $\log p^D(x_t | x_{t-1}) - \log p^I(x_t | x_{t-1})$ (each term estimated separately by maximum likelihood), while GAIfo maximizes $\log \frac{p^D(x_t, x_{t-1})}{p^I(x_t, x_{t-1})} = \log \frac{p^D(x_t | x_{t-1}) p^D(x_{t-1})}{p^I(x_t | x_{t-1}) p^I(x_{t-1})}$ (the entire log ratio is estimated in one go by a discriminator). If a feature is present in the imitator data but was never in the demonstrator data, then $p^D(x_{t-1})$ will be close to 0 on this data, driving the log ratio to $-\infty$ regardless of the probability of the transition that follows. The presence of noise makes spurious features like this inevitable. This makes it difficult for GAIL to focus on the meaningful controllable differences in the data, namely in the transition probabilities $p(x_t | x_{t-1})$. By estimating each term separately (avoiding a discriminator) and including only transition-related terms (using a conditional density), FORM reduces the susceptibility to sensitivity of this kind.

Imitation by Predicting Observations

	Expert	BC	BCO	GAIfO	GAIfO+GP	VAIfO	VAIfO+GP	FORM
Reacher Easy	974.6	970.3 ± 12.2	966.6 ± 7.9	869.9 ± 48.6	915.9 ± 37.8	861.6 ± 61.4	901.3 ± 30.4	950.2 ± 14.9
Reacher Hard	981.3	892.4 ± 19.1	940.1 ± 3.9	818.7 ± 11.3	783.7 ± 119.7	604.4 ± 426.3	891.0 ± 73.9	957.3 ± 6.1
Cheetah Run	930.5	227.5 ± 37.4	75.7 ± 4.2	607.6 ± 429.6	921.3 ± 6.9	820.0 ± 98.8	918.3 ± 6.4	827.9 ± 31.9
Quadruped Walk	972.4	752.1 ± 37.3	191.9 ± 33.6	672.6 ± 409.8	963.6 ± 4.8	927.8 ± 5.0	945.8 ± 15.5	963.6 ± 2.5
Quadruped Run	962.9	719.2 ± 14.0	271.4 ± 48.4	952.5 ± 7.5	952.3 ± 2.1	926.6 ± 38.3	950.0 ± 2.7	948.5 ± 1.5
Hopper Stand	965.8	534.1 ± 13.2	91.4 ± 8.5	400.0 ± 164.3	748.5 ± 224.1	835.8 ± 103.0	891.2 ± 42.1	815.7 ± 9.2
Hopper Hop	711.5	98.4 ± 4.8	9.1 ± 7.2	689.2 ± 10.0	694.4 ± 0.3	610.5 ± 74.8	683.6 ± 22.3	636.2 ± 38.9
Walker Stand	993.6	731.7 ± 29.7	385.9 ± 27.6	989.4 ± 1.5	985.4 ± 1.6	989.4 ± 0.5	986.0 ± 1.9	985.1 ± 2.6
Walker Walk	983.2	719.5 ± 50.0	61.9 ± 20.7	976.5 ± 2.8	981.6 ± 1.4	971.2 ± 5.6	975.2 ± 1.3	977.8 ± 1.0
Walker Run	952.1	108.5 ± 33.2	39.0 ± 7.8	949.5 ± 5.6	947.6 ± 5.5	949.0 ± 2.6	948.5 ± 2.1	942.0 ± 4.5
Humanoid Stand	905.9	780.7 ± 30.5	9.99 ± 2.51	4.9 ± 1.0	856.2 ± 15.5	257.5 ± 12.4	863.5 ± 7.7	704.6 ± 12.1
Humanoid Walk	809.5	293.9 ± 16.2	9.61 ± 5.73	1.2 ± 0.4	798.4 ± 1.0	658.2 ± 123.6	795.5 ± 3.4	783.0 ± 3.3
Humanoid Run	736.6	54.2 ± 5.1	1.04 ± 0.24	0.6 ± 0.0	683.4 ± 6.9	676.6 ± 25.8	691.6 ± 24.0	691.1 ± 7.8

Table 1. Asymptotic performance on 13 tasks from six DCS domains (mean ± standard deviation across three seeds) of our method (FORM) and baselines Behavioral Cloning from Observations (BCO) (Torabi et al., 2018), GAIL from Observations (GAIfO) (Torabi et al., 2019a), and regularized variants with a tuned gradient penalty (Gulrajani et al., 2017) (GAIfO+GP), a variational discriminator bottleneck (Peng et al., 2019) (VAIfO), or both forms of regularization (VAIfO+GP). Because BC (Pomerleau, 1989) uses expert actions it is not comparable to the other methods, but nevertheless performs poorly on many tasks, even with 1000 demonstrations. FORM performs competitively with well-regularized forms of GAIfO, while generally outperforming BCO and GAIfO. For each task, we highlight the method with **best** and second best mean performance.

Finally, we note that GAIL is typically justified by the observation that recovering an expert’s occupancy is equivalent to recovering its policy, but this is only true in Markov Decision Processes (MDPs) (Syed et al., 2008; Ho & Ermon, 2016) and not in general. In practice, imitation must often be done using noisy or high-dimensional observations rather than ground-truth MDP states, and matching occupancy in these spaces is problematic. In settings like this, relying on the global occupancy induced by a policy rather than on the immediate effects of a policy may lead to misleading results. For example, GAIL will attempt to match the occupancy of all noise dimensions, and this is usually possible. In practice, this means that the GAIL objective needs to be carefully regularized to avoid overfitting to irrelevant differences. These effects appear to be stronger when training an IRL agent from replay, as discussed in (Kostrikov et al., 2019), and they may be further exacerbated when imitating without actions. In our experiments, GAIfO fails completely on the Humanoid tasks of the Control Suite when unregularized. Even with strong regularization, GAIfO is very sensitive to the presence of irrelevant differences between demonstrator and imitator domains, as our experiments illustrate.

4. Experiments

We evaluate FORM against strong baselines on 13 tasks from six domains from the DeepMind Control Suite (DCS) (Tassa et al., 2018), a set of benchmarks for continuous control domains, chosen to match those frequently used in the imitation learning literature³. All approaches use internal

³We note that many imitation learning methods are evaluated on the superficially similar OpenAI Gym Mujoco benchmark (Brockman et al., 2016), but the Gym domains have essentially deterministic initial states and other properties that make them poorly

Mujoco state representations: these are smaller than e.g. image observations, and vary in size from 6- (reacher) to 67- (humanoid) and 78-dimensional (quadruped). As observed by (Zolna et al., 2020), GAIL struggles to imitate in the presence of a small number of differences between expert and imitator domains. We conduct a similar experiment to characterize the robustness of FORM and GAIfO to irrelevant, but undersampled, factors of variation in the demonstrator data. Because the focus of our evaluation concerns robustness to distractors, rather than the minimum number of demonstrations needed for successful imitation, we conduct all experiments using 1000 demonstrations, sufficient to ensure mostly satisfactory performance in the absence of distractors.

Expert data. For all domains, we train an expert via RL on the ground truth task reward. Experts are trained to convergence using MPO, with the same policy and value architecture used for imitation (under all imitation conditions). For imitation, we generate a fixed dataset of 1000 demonstration trajectories from each policy, each of which depicts a single episode 1000 timesteps in duration (i.e. 10^6 steps total). All imitation methods are trained using the same demonstrator data.

Distractor data. To probe robustness to a domain shift between the expert and imitation domains, we deliberately introduce spurious signals, unrelated to the task or agent state, into the state observation vector. During the demonstration phase, these take the form of binary noise patterns drawn from a fixed set which are appended to the state vector and held constant for the duration of the episode. During suited for evaluating imitation learning methods (see Sec. D of the Appendix for a discussion).

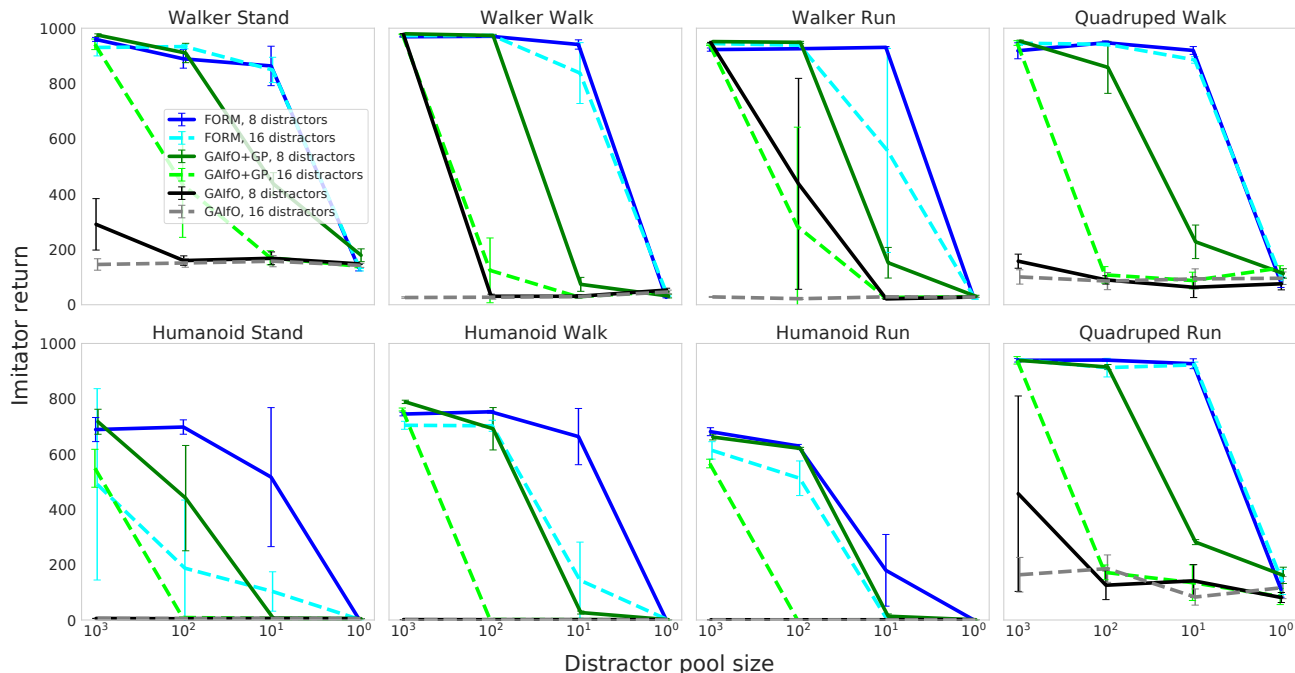


Figure 3. Performance of FORM, GAIfo+GP, and GAIfo in the presence of distractor features. The distractor pool size M (#unique points sampled in expert data) is varied from 10^3 down to 1 for both $N = 8$ and $N = 16$ distractor dimensions. FORM exhibits greater stability than GAIfo+GP in both settings, maintaining performance down to $M = 10$. Error bars indicate standard deviation across 3 seeds. For reasons of legibility and space, comparison to all other baselines are given in the appendix.

imitation binary noise vectors are also appended, but any binary pattern is permitted (i.e. no longer need come from the fixed set). The different patterns appended during the expert and imitation phases impose a domain shift between them.

Formally, during imitation we append the state vector x_t with a binary pattern $b \sim [0, 1]^N$ to form an augmented observation $\tilde{x}_t = [x_t; b]$, where N is the number of distractor dimensions. During demonstration, $b \sim \{b_1, \dots, b_M\}$, $b \in [0, 1]^N$, while M controls the the number of distinct patterns, known as the *pool size*. M and N control the magnitude of the domain shift: increasing N makes the task harder by reducing the fraction of state that contains signal, while increasing M makes the task easier by ensuring that all distractor features are present in both demonstrator and imitator data.

Due to the input normalization procedure (Sec. 4.1), the IL agent has no way of distinguishing noise dimensions from ones carrying state information. Ideally however, it should learn to ignore the extra dimensions since they are unrelated to the task, making it robust to changes in the distractor pattern. Our setup directly parallels situations encountered in practice involving undersampled factors of variation. For example, when performing IL using visual inputs with a robot, the background appearance of the rooms in which

the expert data collection and imitation during deployment are performed correspond to two distinct distractor patterns that are intermingled with task-relevant portions of the state. For IL to work in such settings the algorithm must be robust to changes in the background distractors. We can see how sensitive a model is to the presence of undersampled factors of variation by observing how stable its performance is as the pool size M decreases.

4.1. Details of the FORM implementation

Architecture. We use simple feedforward architectures to parameterize the density models (3 layer MLPs with 256 units, and tanh and ELU (Clevert et al., 2016) nonlinearities). We model the density as a mixture of 4 Gaussian components, with the network outputting GMM mixture coefficients and the means and standard deviations of each component. We use Gaussians with a diagonal covariance matrix. In all experiments, we clip the standard deviation to a minimum value of 0.0001. We use the same architecture and same hyperparameters for the imitator and demonstrator effect model in each setting.

Effect model training. All demonstrator models were trained offline for 2 million steps. We standardized effect model inputs using per-dimension means and variances estimated by by exponential moving average: we found that

this improved generative model training (it did not affect GAIfo training).

Three forms of regularization were used with the demonstrator and imitator generative models: (i) ℓ_2 weight-decay, (ii) training on data generated by agent rollouts, i.e. using the network output at a timestep as the input at the next during training (a common trick used in the recurrent neural network literature (Bengio & Shazeer, 2015)), (iii) prediction of observations at multiple future timesteps (Hafner et al., 2019). See the Appendix for more details. In all experiments, we share the hyperparameter settings of all regularizers between the demonstrator and imitator effect model (we do not tune them separately). We tuned ℓ_2 weight (sweeping values of [0.0, 0.01, 0.1, and 1.0]) and the fraction of each batch generated by agent rollouts (sweeping values of [0.0, 0.01, 0.1, 1.0]) per domain, but otherwise use identical hyperparameters for all FORM models.

4.2. Baselines: GAIfo and BCO

To ensure DMCS experiments were fair and well-calibrated, we implemented and tuned a strong GAIfo baseline. The GAIfo discriminator is conditioned on the current observation. We found that there was no benefit to conditioning on pairs of subsequent observations (see Table 5 in the Appendix). This is likely because DCS observations include velocity observations as well as static positions. The discriminator network uses the same architecture as the FORM effect models except for the mixture-of-Gaussians head, which is replaced by a (scalar) classifier head. We additionally found that there was no benefit in standardizing the observations as we do for FORM. For GAIfo+GP, we apply a gradient penalty (Gulrajani et al., 2017) to the last two layers of the discriminator. For VAIfo (VAIL from observations), we introduce a variational bottleneck in the discriminator architecture and add a KL-constraint term to the loss, as in (Peng et al., 2019). Following (Kostrikov et al., 2019), we train both the policy and the discriminator using data sampled from a replay buffer.

In BCO, an inverse model $p(a_t|x_t, x_{t+1})$ is trained on imitator trajectories and then used to label the actions on demonstrator trajectories (Torabi et al., 2018). We train the inverse model using the same architecture as the FORM effector models and the GAIfo discriminator, replacing the output head with a Gaussian distribution (the same class of distributions used by the RL agent to produce the actions). The BCO agent is then trained in a supervised fashion on expert trajectories labeled by the inverse model. The BC agent is trained directly on expert trajectories with expert actions. Because BC is trained using expert actions, while the other imitation algorithms we evaluate are not, it is not strictly comparable. We include it to calibrate readers to the difficulty of these tasks and the relative performance of the

algorithms we evaluate for imitation from observation.

4.3. Policy architecture

For both IRL methods (FORM and GAIfo), the underlying policy is trained with MPO and experience replay. Both the policy and critic networks encode a concatenation of the environment’s observations that has been passed through a tanh activation. Both encode the observations with independent 3-layer MLPs using ELU activations. The policy network then projects to parameterize the mean and scale of a Gaussian action distribution. The critic concatenates the sampled action, applies layernorm (Ba et al., 2016) and a tanh, and applies another 3-layer MLP to produce the Q -value. All hidden layers have a width of 256 units.

4.4. Results

No distractors. In Table 1 we compare FORM to BCO and GAIfo, with various strong forms of regularization on the DeepMind Control Suite in the absence of distractors. The results are shown alongside the reward obtained by the expert RL agent. BCO succeeds only on the Reacher domain, performing poorly on the others. GAIfo in general performs well but fails completely on the Humanoid domain. The addition of a tuned gradient penalty or the introduction of a variational discriminator bottleneck allows GAIfo to also perform well on the Humanoid domain. FORM achieves competitive performance to strongly regularized forms of GAIfo (GAIfo+GP, VAIfo, VAIfo+GP). Despite access to 1000 demonstrations, no method is able to match expert performance on the Humanoid tasks, illustrating the challenge this domain poses due to the dimensionality of the state-space and highly variable initial conditions. This number of demonstrations (1000) may seem large when compared to the numbers used in work that uses the Gym benchmark (Brockman et al., 2016). But please note that tasks from the Gym benchmark are easier to imitate, requiring almost no generalization between demonstrator and imitator due to its essentially deterministic initialization.

With distractors. We now explore the robustness of the different approaches to settings where distractors are present in the observations. Figure 3 compares FORM with GAIfo and GAIfo+GP with $N = 8$ and $N = 16$ dimensions of distractor features, as the distractor pool size M is varied from 10^3 down to 1. With $N = 8$ distractor dimensions, FORM is consistently able to maintain performance as $M = 10$, by which point the performance of GAIfo+GP has dropped significantly. For $N = 16$ distractor dimensions the degradation for GAIfo+GP is more severe, even at the easier setting of $M = 100$. In contrast, FORM is still able to perform well on most tasks (Humanoid Stand being the exception). We compare FORM to all other baselines on this setting in Figs. 5 and 6 of the appendix, and this trend holds generally. We find that GAIfo with a variational bottleneck

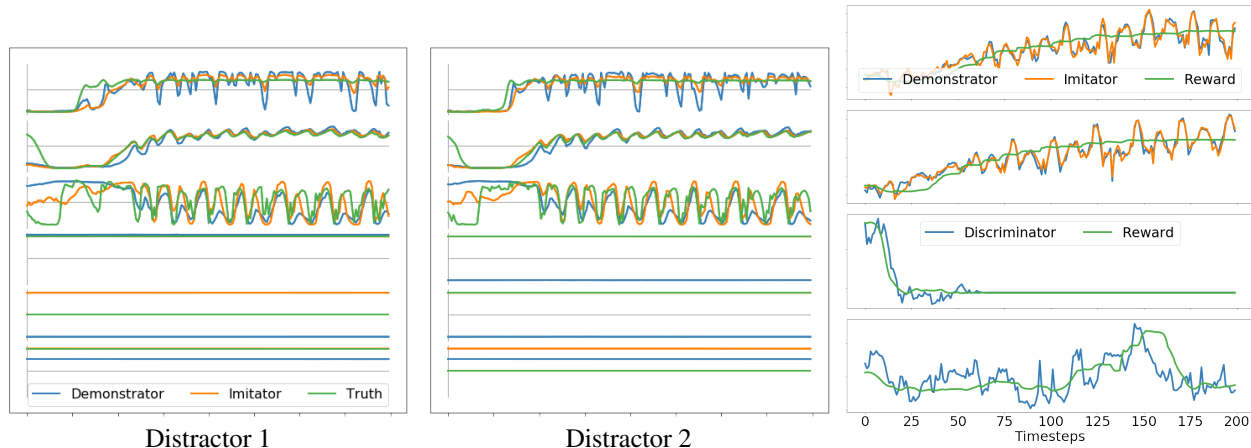


Figure 4. Left, Middle — subset of observation dimensions (3 internal states (top) and 3 distractor features (bottom)) from a FORM imitation agent on the Walker Run task. Distractor features appear as horizontal lines as they do not vary with time in an episode. **True observations**, with **demonstrator** and **imitator** predictions overlaid. Distractor 1 shows the agent learning with a pattern previously seen in the expert demonstrations. Distractor 2 uses a novel pattern not seen in expert data. The FORM agent behavior and model predictions are qualitatively unchanged, showing robustness to the distractor pattern. **Right** — reward model traces for both FORM and GAIfo+GP, alongside the ground truth **task reward**. *Top*: the log-likelihoods of the **demonstrator** and **imitator** components of the FORM reward model, for distractor 1 (top) and distractor 2 (bottom). *Bottom*: the expert probability output by the GAIfo+GP **discriminator** for distractor 1 (top) and distractor 2 (bottom). Both FORM and GAIfo+GP agents were trained with $N = 8$ distractor dimensions and a pool of $M = 10$ distractors in expert data. FORM is robust to distractor features in this setting even when its predictions are imperfect, and the imitation agent obtains good reward on the task. In contrast, the behavior of the GAIfo+GP agent depends significantly on the distractors and largely fails at the task.

regularizer (VAIfo and VAIfo+GP) performs similarly to GAIfo, and still exhibits sensitivity to noise. BC and BCO generally perform worse than FORM, but BC shows good noise resilience on Humanoid Stand in particular.

Figure 4 visualizes the effect on imitation performance when the distractor pattern ($M = 10$, $N = 8$) is changed between imitation training runs on the Walker Run task. Demonstrator and imitator model predictions from the FORM model show minimal change, with the agent achieving good reward for both patterns. In contrast, the GAIfo+GP model is highly sensitive to the change in distractor pattern and fails at the task. Collectively, these results show the fragility of GAIL-based IL methods to task-irrelevant features, and also illustrate the superior robustness of FORM in this setting.

5. Discussion

In this work we introduce the Future Observation Reward Model, or FORM, an approach to inverse reinforcement learning that can be used for imitation from observations without actions. FORM makes few assumptions about the data being modeled, which makes it a promising approach for learning behavior from data collected under realistic conditions. In particular, we show that FORM is competitive with GAIL from observations while exhibiting improved stability in the face of spurious features. FORM imitates using likelihood-based generative models, a family of mod-

els that has been extensively studied and that can be scaled to real-world, noisy data. These properties make FORM a good candidate for the development of sophisticated approaches to imitation that can handle high-dimensional data with domain shifts.

FORM currently has several limitations. The demonstrator model p^D must be trained off-line before learning the imitator model p^I and policy π^I . This two-stage training is inefficient in wall clock terms relative to the monolithic training procedure of GAIL. This is compounded by the difficulty in assessing the quality of p^D using training likelihood alone. In practice, we find that it is a poor predictor of subsequent imitator performance, necessitating both both training stages to be performed in order to ascertain if p^D was modeled effectively. A second issue is that we currently model proprioceptive state: moving to image pixel-based inputs will require larger and more complex generative models, which will likely lead to added difficulties.

Acknowledgements

We are grateful to Josh Abramson, Feryal Behbahani, Federico Carnevale, Ashley Edwards, Tom Erez, Karol Gregor, Raia Hadsell, Leonard Hasenclever, Nicolas Heess, Alden Hung, Josh Merel, Nikolay Savinov, Yuval Tassa, Konrad Zolna and others at DeepMind for insightful discussions and suggestions.

References

- Abbeel, P. and Ng, A. Y. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of International Conference on Machine Learning (ICML)*, 2004.
- Abbeel, P., Coates, A., and Ng, A. Y. Autonomous helicopter aerobatics through apprenticeship learning. *International Journal of Robotics Research*, 29(13): 1608–1639, 2010.
- Abdolmaleki, A., Springenberg, J. T., Tassa, Y., Munos, R., Heess, N., and Riedmiller, M. Maximum a posteriori policy optimisation. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2018.
- Abramson, J., Ahuja, A., Brussee, A., Carnevale, F., Cassin, M., Clark, S., Dudzik, A., Georgiev, P., Guy, A., Harley, T., Hill, F., Hung, A., Kenton, Z., Landon, J., Lillicrap, T., Mathewson, K., Muldal, A., Santoro, A., Savinov, N., Varma, V., Wayne, G., Wong, N., Yan, C., and Zhu, R. Imitating interactive intelligence. *arXiv preprint arXiv:2012.05672*, 2020.
- Arenz, O. and Neumann, G. Non-adversarial imitation learning and its connections to adversarial methods. *arXiv preprint arXiv:2008.03525*, 2020.
- Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein GAN. In *Proceedings of International Conference on Machine Learning (ICML)*, 2017.
- Ba, J. L., Kiros, J. R., and Hinton, G. E. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- Babuschkin, I., Baumli, K., Bell, A., Bhupatiraju, S., Bruce, J., Buchlovsky, P., Budden, D., Cai, T., Clark, A., Danihelka, I., Fantacci, C., Godwin, J., Jones, C., Hennigan, T., Hessel, M., Kapturowski, S., Keck, T., Kemaev, I., King, M., Martens, L., Mikulik, V., Norman, T., Quan, J., Papamakarios, G., Ring, R., Ruiz, F., Sanchez, A., Schneider, R., Sezener, E., Spencer, S., Srinivasan, S., Stokowiec, W., and Viola, F. The DeepMind JAX Ecosystem, 2020. URL <http://github.com/deepmind>.
- Baram, N., Anschel, O., and Mannor, S. Model-based adversarial imitation learning. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2016.
- Bengio, Samy, V. O. J. N. and Shazeer, N. M. Scheduled sampling for sequence prediction with recurrent neural networks. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2015.
- Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., and Zhang, Q. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. OpenAI gym. *arXiv preprint arXiv:1606.01540*, 2016.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. Language models are few-shot learners. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2020.
- Byrne, R. W. Animal imitation. *Current Biology*, 19(3):111 – 114, 2009.
- Clevert, D.-A., Unterthiner, T., and Hochreiter, S. Fast and accurate deep network learning by exponential linear units (ELUs). In *Proceedings of International Conference on Learning Representations (ICLR)*, 2016.
- Dosovitskiy, A. and Koltun, V. Learning to act by predicting the future. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2017.
- Edwards, A. D., Sahni, H., Schroecker, Y., and Isbell, C. L. Imitating latent policies from observation. In *Proceedings of International Conference on Machine Learning (ICML)*, 2018.
- Edwards, A. D., Sahni, H., Liu, R., Hung, J., Jain, A., Wang, R., Ecoffet, A., Miconi, T., Isbell, C., and Yosinski, J. Estimating $Q(s,s')$ with deep deterministic dynamics gradients. In *Proceedings of International Conference on Machine Learning (ICML)*, 2020.
- Fu, J., Luo, K., and Levine, S. Learning robust rewards with adversarial inverse reinforcement learning. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2018.
- Ghasemipour, S. K. S., Zemel, R., and Gu, S. A divergence minimization perspective on imitation learning methods. In *Conference on Robotic Learning (CoRL)*, 2019.
- Gulcehre, C., Paine, T. L., Shahriari, B., Denil, M., Hoffman, M., Soyer, H., Tanburn, R., Kapturowski, S., Rabinowitz, N., Williams, D., Barth-Maron, G., Wang, Z., de Freitas, N., and Team, W. Making efficient use of demonstrations to solve hard exploration problems. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2020.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. C. Improved training of wasserstein GANs. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2017.

- Hafner, D., Lillicrap, T., Fischer, I., Villegas, R., Ha, D., Lee, H., and Davidson, J. Learning latent dynamics for planning from pixels. In *Proceedings of International Conference on Machine Learning (ICML)*, 2019.
- Ho, J. and Ermon, S. Generative adversarial imitation learning. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2016.
- Huber, L., Range, F., Voelkl, B., Szucsich, A., Virányi, Z., and Miklosi, A. The evolution of imitation: what do the capacities of non-human animals tell us about the mechanisms of imitation? *Philosophical Transactions of the Royal Society of London B*, 364(1528):2299 – 2309, 2009.
- Jarrett, D., Bica, I., and van der Schaar, M. Strictly batch imitation learning by energy-based distribution matching. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2020.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2014.
- Kostrikov, I., Agrawal, K. K., Dwibedi, D., Levine, S., and Tompson, J. Discriminator-actor-critic: Addressing sample inefficiency and reward bias in adversarial imitation learning. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2019.
- Kostrikov, I., Nachum, O., and Tompson, J. Imitation learning via off-policy distribution matching. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2020.
- Laland, K. Animal cultures. *Current Biology*, 18(9):366 – 370, 2008.
- Levine, S. Reinforcement learning and control as probabilistic inference: Tutorial and review. *arXiv preprint arXiv:1805.00909*, 2018.
- Liu, M., He, T., Xu, M., and Zhang, W. Energy-based imitation learning. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2021.
- Mania, H., Guy, A., and Recht, B. Simple random search of static linear policies is competitive for reinforcement learning. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2018.
- Merel, J., Tassa, Y., TB, D., Srinivasan, S., Lemmon, J., Wang, Z., Wayne, G., and Heess, N. Learning human behaviors from motion capture by adversarial imitation. *arXiv preprint arXiv:1707.02201*, 2017.
- Merel, J., Hasenclever, L., Galashov, A., Ahuja, A., Pham, V., Wayne, G., Teh, Y. W., and Heess, N. Neural probabilistic motor primitives for humanoid control. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2019.
- Mescheder, L., Nowozin, S., and Geiger, A. The numerics of GANs. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2017.
- Munos, R., Stepleton, T., Harutyunyan, A., and Bellemare, M. G. Safe and efficient off-policy reinforcement learning. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2016.
- Ng, A. and Russell, S. Algorithms for inverse reinforcement learning. In *Proceedings of International Conference on Machine Learning (ICML)*, 2000.
- OpenAI, Berner, C., Brockman, G., Chan, B., Cheung, V., Debiak, P., Dennison, C., Farhi, D., Fischer, Q., Hashme, S., Hesse, C., Józefowicz, R., Gray, S., Olsson, C., Pachocki, J., Petrov, M., de Oliveira Pinto, H. P., Raiman, J., Salimans, T., Schlatter, J., Schneider, J., Sidor, S., Sutskever, I., Tang, J., Wolski, F., and Zhang, S. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1910.07113*, 2019.
- Osa, T., Pajarinen, J., Neumann, G., Bagnell, J. A., Abbeel, P., and Peters, J. An algorithmic perspective on imitation learning. *Foundations and Trends in Robotics*, 7(1–2): 1–179, 2018.
- Pathak, D., Mahmoodieh, P., Luo, G., Agrawal, P., Chen, D., Shentu, Y., Shelhamer, E., Malik, J., Efros, A. A., and Darrell, T. Zero-shot visual imitation. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2018.
- Peng, X. B., Abbeel, P., Levine, S., and van de Panne, M. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions on Graphics*, 37(4):143:1–143:14, 2018.
- Peng, X. B., Kanazawa, A., Toyer, S., Abbeel, P., and Levine, S. Variational discriminator bottleneck: Improving imitation learning, inverse RL, and GANs by constraining information flow. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2019.
- Pomerleau, D. A. ALVINN: An autonomous land vehicle in a neural network. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 1989.
- Rhinehart, N., McAllister, R., and Levine, S. Deep imitative models for flexible inference, planning, and control. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2020.

- Ross, S., Gordon, G. J., and Bagnell, J. A. A reduction of imitation learning and structured prediction to no-regret online learning. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2011.
- Rybkin, O., Pertsch, K., Derpanis, K. G., Daniilidis, K., and Jaegle, A. Learning what you can do before doing anything. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2019.
- Schmeckpeper, K., Xie, A., Rybkin, O., Tian, S., Daniilidis, K., Levine, S., and Finn, C. Learning predictive models from observation and interaction. In *Proceedings of European Conference on Computer Vision (ECCV)*, 2020.
- Sermanet, P., Lynch, C., Chebotar, Y., Hsu, J., Jang, E., Schaal, S., and Levine, S. Time-contrastive networks: Self-supervised learning from video. In *Proceedings of IEEE International Conference on Robotics and Automation*, 2017.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Driessche, G. v. d., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., 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.
- Stone, A., Ramirez, O., Konolige, K., and Jonschkowski, R. The distracting control suite – a challenging benchmark for reinforcement learning from pixels. *arXiv preprint 2101.02722*, 2021.
- Sun, W., Vemula, A., Boots, B., and Bagnell, J. A. Provably efficient imitation learning from observation alone. In *Proceedings of International Conference on Machine Learning (ICML)*, 2019.
- Sutton, R. S. and Barto, A. G. *Reinforcement Learning: An Introduction*. The MIT Press, second edition, 2018.
- Syed, U., Bowling, M., and Schapire, R. E. Apprenticeship learning using linear programming. In *Proceedings of International Conference on Machine Learning (ICML)*, 2008.
- Tassa, Y., Doron, Y., Muldal, A., Erez, T., Li, Y., de Las Casas, D., Budden, D., Abdolmaleki, A., Merel, J., Lefrancq, A., Lillicrap, T., and Riedmiller, M. DeepMind control suite. *arXiv preprint arXiv:1801.00690*, 2018.
- Tomasello, M. Do apes ape? In Heyes, C. M. and Jr., B. G. G. (eds.), *Social Learning in Animals: The Roots of Culture*, chapter 15, pp. 319–346. Academic Press, 1996.
- Torabi, F., Warnell, G., and Stone, P. Behavioral cloning from observation. In *Proceedings of International Joint Conference on Artificial Intelligence*, 2018.
- Torabi, F., Warnell, G., and Stone, P. Generative adversarial imitation from observation. In *Imitation, Intent, and Interaction (I3) (ICML Workshop)*, 2019a.
- Torabi, F., Warnell, G., and Stone, P. Recent advances in imitation learning from observation. In *Proceedings of International Joint Conference on Artificial Intelligence*, 2019b.
- Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, M., Dudzik, A., Chung, J., Choi, D. H., Powell, R., Ewalds, T., Georgie, P., Oh, J., Horgan, D., Kroiss, M., Danihelka, I., Huang, A., Sifre, L., Cai, T., Agapiou, J. P., Jaderberg, M., Vezhnevets, A. S., Leblond, R., Pohlen, T., Dalibard, V., Budden, D., Sulsky, Y., Molloy, J., Paine, T. L., Gulcehre, C., Wang, Z., Pfaff, T., Wu, Y., Ring, R., Yogatama, D., Wünsch, D., McKinney, K., Smith, O., Schaul, T., Lillicrap, T., Kavukcuoglu, K., Hassabis, D., Apps, C., and Silver, D. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575(7782): 350—354, 2019.
- Wang, Z., Merel, J., Reed, S., Wayne, G., de Freitas, N., and Heess, N. Robust imitation of diverse behaviors. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2017.
- Yu, X., Lyu, Y., and Tsang, I. W. Intrinsic reward driven imitation learning via generative model. In *Proceedings of International Conference on Machine Learning (ICML)*, 2020.
- Zhu, Z., Lin, K., Dai, B., and Zhou, J. Off-policy imitation learning from observations. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2020.
- Ziebart, B. D. *Modeling Purposeful Adaptive Behavior with the Principle of Maximum Causal Entropy*. PhD thesis, Carnegie Mellon University, 2010.
- Ziebart, B. D., Maas, A., Bagnell, J. A., and Dey, A. K. Maximum entropy inverse reinforcement learning. In *Proceedings of AAAI Conference on Artificial Intelligence*, 2008.
- Zolna, K., Reed, S., Novikov, A., Colmenarejo, S. G., Budden, D., Cabi, S., Denil, M., de Freitas, N., and Wang, Z. Task-relevant adversarial imitation learning. In *Conference on Robotic Learning (CoRL)*, 2020.