

Imitation Learning from Video by Leveraging Proprioception

Faraz Torabi¹, Garrett Warnell² and Peter Stone¹

¹The University of Texas at Austin

²Army Research Laboratory

{faraztrb, pstone}@cs.utexas.edu, garrett.a.warnell.civ@mail.mil

Abstract

Classically, imitation learning algorithms have been developed for idealized situations, e.g., the demonstrations are often required to be collected in the exact same environment and usually include the demonstrator’s actions. Recently, however, the research community has begun to address some of these shortcomings by offering algorithmic solutions that enable imitation learning from observation (*IfO*), e.g., learning to perform a task from *visual* demonstrations that may be in a different environment and do not include actions. Motivated by the fact that agents often also have access to their own internal states (i.e., *proprioception*), we propose and study an *IfO* algorithm that leverages this information in the policy learning process. The proposed architecture learns policies over proprioceptive state representations and compares the resulting trajectories visually to the demonstration data. We experimentally test the proposed technique on several MuJoCo domains and show that it outperforms other imitation from observation algorithms by a large margin.

1 Introduction

Imitation learning [Schaal, 1997; Argall *et al.*, 2009; Osa *et al.*, 2018] is a popular method by which artificial agents learn to perform tasks. In the imitation learning framework, an expert agent provides demonstrations of a task to a learning agent, and the learning agent attempts to mimic the expert. Unfortunately, many existing imitation learning algorithms have been designed for idealized situations, e.g., they require that the demonstrations be collected in the exact same environment as the one that the imitator is in and/or that the demonstrations include the demonstrator’s actions, i.e., the internal control signals that were used to drive the behavior. These limitations result in the exclusion of a large amount of existing resources, including a large number of videos uploaded to the internet. For example, 300 hours of video are uploaded to YouTube every minute¹, many of which include

different types of tasks being performed. Without new imitation learning techniques, none of this video can be used to instruct artificial agents.

Fortunately, the research community has recently begun to focus on addressing the above limitations by considering the specific problem of imitation from observation (*IfO*) [Liu *et al.*, 2018; Torabi *et al.*, 2019d]. *IfO* considers situations in which agents attempt to learn tasks by observing demonstrations that contain only state information (e.g., videos). Among *IfO* algorithms that learn tasks by watching videos, most attempt to learn imitation policies that rely solely on self-observation through video, i.e., they use a convolutional neural network (CNN) that maps images of themselves to actions. However, in many cases, the imitating agent also has access to its own *proprioceptive* state information, i.e., direct knowledge of itself such as the joint angles and torques associate with limbs. In this paper, we argue that *IfO* algorithms that ignore this information are missing an opportunity that could potentially improve the performance and the efficiency of the learning process. Therefore, we are interested here in *IfO* algorithms that can make use of both visual and proprioceptive state information.

In this paper, we build upon our previous work [Torabi *et al.*, 2019b] proposing an algorithm that uses a *GAN*-like [Goodfellow *et al.*, 2014] architecture to learn tasks perform *IfO* directly from videos. Unlike our prior work, however, our method *also* uses proprioceptive information from the imitating agent during the learning process. We hypothesize that the addition of such information will improve both learning speed and the final performance of the imitator, and we test this hypothesis experimentally in several standard simulation domains. We compare our method with other, state-of-the-art approaches that do not leverage proprioception, and our results validate our hypothesis, i.e., the proposed technique outperforms the others by a large margin.

The rest of this paper is organized as follows. In Section 2, we review related work in imitation from observation. In Section 3, we review technical details surrounding Markov decision processes, imitation learning, and *IfO*. The proposed algorithm is presented in Section 4, and we describe the experiments that we have performed in Section 5.

¹<https://bit.ly/2quPG6O>

2 Related Work

In this section, we review research in imitation learning, plan/goal recognition by mirroring, and recent advances in imitation from observation (*IfO*).

Conventionally, imitation learning is used in autonomous agents to learn tasks from demonstrated *state-action* trajectories. The algorithms developed for this task can be divided into two general categories, (1) behavioral cloning [Bain and Sommut, 1999; Ross *et al.*, 2011; Daftry *et al.*, 2016] in which the agents learn a direct mapping from the demonstrated states to the actions, and (2) inverse reinforcement learning (IRL) [Abbeel and Ng, 2004; Bagnell *et al.*, 2007; Baker *et al.*, 2009] in which the agents first learn a reward function based on the demonstrations and then learn to perform the task using a reinforcement learning (*RL*) [Sutton and Barto, 1998] algorithm.

In contrast, imitation from observation (*IfO*) is a framework for learning a task from *state-only* demonstrations. This framework has recently received a great deal of attention from the research community. The *IfO* algorithms that have been developed can be categorized as either (1) model-based, or (2) model-free. Model-based algorithms require the agent to learn an explicit model of its environment as part of the imitation learning process. One algorithm of this type is behavioral cloning from observation (*BCO*) [Torabi *et al.*, 2018], in which the imitator learns a dynamics model of its environment using experience collected by a known policy, and then uses this model to infer the missing demonstrator actions. Using the inferred actions, the imitator then computes an imitation policy using behavioral cloning [Bain and Sammut, 1995]. Another approach of this type is reinforced inverse dynamics modeling (*RIDM*) [Pavse *et al.*, 2019] which also learns a model of its environment using an exploration policy and then it further optimizes the model using a sparse reward function. It is shown that in some experiments the algorithm can even outperform the expert. Another model-based approach to *IfO* is imitating latent policies from observation (*ILPO*) [Edwards *et al.*, 2019]. Given the current state of the expert, this approach predicts the next state using a latent policy and a forward dynamics model. It then uses the difference between the predicted state and the actual demonstrator next state to update both the model and the imitation policy. Afterwards, the imitator interacts with its environment to correct the action labels.

Model-free algorithms, on the other hand, do not require any sort of model to learn imitation policies. One set of approaches of this type learns a time-dependent representation of tasks and then relies on hand-designed, time-aligned reward functions to learn the task via *RL*. For example, Sermanet *et al.* [2018] propose an algorithm that learns an embedding function using a triplet loss function that seeks to push states that are close together in time closer together in the embedded space, while pushing other states further away. Liu *et al.* [2018] also propose a new architecture to learn a state representation—specifically, one that is capable of handling viewpoint differences. Gupta *et al.* [2017] also propose a neural network architecture to try to learn a state representation that can overcome possible embodiment mismatch be-

tween the demonstrator and the imitator. Each of these approaches requires multiple demonstrations of the same task to be time-aligned, which is typically not a realistic assumption. Aytar *et al.* [2018] propose an *IfO* algorithm that first learns an embedding using a self-supervised objective, and then constructs a reward function based on the embedding representation difference between the current state of the imitator and a specific checkpoint generated by the visual demonstration. Goo and Niekum [2019] propose an algorithm that uses a shuffle-and-learn style [Misra *et al.*, 2016] loss in order to train a neural network that can predict progress in the task which can then be used as the reward function.

Another set of model-free algorithms follow a more end-to-end approach to learning policies directly from observations. An algorithm of this type is generative adversarial imitation from observation (*GAIfo*) [Torabi *et al.*, 2019b], which uses a *GAN*-like architecture to bring the state transition distribution of the imitator closer to that of the demonstrator. Another approach of this type is the work of Merel *et al.* [2017], which is concerned instead with single state distributions. Stadie *et al.* [2017] also propose an algorithm in this space that combines adversarial domain confusion methods [Ganin *et al.*, 2016] with adversarial imitation learning algorithms in an attempt to overcome changes in viewpoint. The method we propose in this paper also belongs to the category of end-to-end model-free imitation from observation algorithms. However, it is different from the algorithms discussed above in that we explicitly incorporate the imitator’s proprioceptive information in the learning process in order to study the improvement such information can make with respect to the performance and speed of the learning process.

A method that is closely related to imitation from observation is plan/goal recognition through mirroring [Vered *et al.*, 2016; Vered *et al.*, 2018] in that it attempts to infer higher-level variables such as the goal or the future plan by observing other agents. However, in plan and goal recognition the observer already has fixed controllers, and then uses these controllers to match/explain the observed agent in order to infer their goal/plan. In imitation from observation, on the other hand, the agent seeks to learn a controller that the observer can use to imitate the observed agent.

3 Background

In this section, we establish notation and provide background information about Markov decision processes (*MDPs*) and adversarial imitation learning.

3.1 Notation

We consider artificial learning agents operating in the framework of Markov decision processes (*MDPs*). An *MDP* can be described as a 6-tuple $\mathcal{M} = \{\mathcal{S}, \mathcal{A}, P, r, \gamma\}$, where \mathcal{S} and \mathcal{A} are state and action spaces, $P(s_{t+1}|s_t, a_t)$ is a function which represents the probability of an agent transitioning from state s_t at time t to s_{t+1} at time $t+1$ by taking action a_t , $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{R}$ is a function that represents the reward feedback that the agent receives after taking a specific action at a given state, and γ is a discount factor. In the context of the notation established above, we are interested here in learning

a policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$ that can be used to select an action at each state.

In this paper, we refer to s as the *proprioceptive* state, i.e., s is the most basic, internal state information available to the agent (e.g., the joint angles of a robotic arm). Since we are also concerned with visual observations of agent behavior, we denote these observations as $o \in \mathcal{O}$, i.e., an image of the agent at time t is denoted as o_t . The visual observations of the agent are determined both by the agent’s current proprioceptive state s , and also other factors relating to image formation such as camera position. Importantly, due to phenomena such as occlusion, it is not always possible to infer s from o alone.

In imitation learning (*IL*), agents do not receive reward feedback r . Instead, they have access to expert demonstrations of the task. These demonstrations $\tau_e = \{(s_t, a_t)\}$ are composed of the state and action sequences experienced by the demonstrator. Here, however, we specifically consider the problem of imitation from observation (*IfO*), in which the agent only has access to sequences of visual observations of the demonstrator performing the task, i.e., $\tau_e = \{o_t\}$.

3.2 Adversarial Imitation Learning

Generative adversarial imitation learning (*GAIL*) is a recent imitation learning algorithm developed by Ho and Ermon [2016] that formulates the problem of finding an imitating policy as that of solving the following optimization problem:

$$\min_{\pi \in \prod_{D \in (0,1)^{\mathcal{S} \times \mathcal{A}}} } \max_{D \in (0,1)^{\mathcal{S} \times \mathcal{A}}} -\lambda_H H(\pi) + \mathbb{E}_{\pi}[\log(D(s, a))] + \mathbb{E}_{\pi_E}[\log(1 - D(s, a))], \quad (1)$$

where H is the entropy function, and the discriminator function $D : \mathcal{S} \times \mathcal{A} \rightarrow (0, 1)$ can be thought of as a classifier trained to differentiate between the state-action pairs provided by the demonstrator and those experienced by the imitator. The objective in (1) is similar the one used in generative adversarial networks (*GANs*) [Goodfellow *et al.*, 2014], and the associated algorithm can be thought of as trying to induce an imitator state-action occupancy measure that is similar to that of the demonstrator. Even more recently, there has been research on methods that seek to improve on *GAIL* by, e.g., increasing sample efficiency [Kostrikov *et al.*, 2019; Sasaki *et al.*, 2019] and improving reward representation [Fu *et al.*, 2018; Qureshi *et al.*, 2019].

The method we propose in this paper is most related to generative adversarial imitation from observation [Torabi *et al.*, 2019b], which models the imitating policy using a randomly-initialized convolutional neural network, executes the policy to generate recorded video of the imitator’s behavior, and then trains a discriminator to differentiate between video of the demonstrator and video of the imitator. Next, it uses the discriminator as a reward function for the imitating agent (higher rewards corresponding to behavior the discriminator classifies as coming from the demonstrator), and uses a policy gradient technique (e.g., *TRPO* [Schulman *et al.*, 2015]) to update the policy. The process repeats until convergence. This algorithm differs from what we propose in that *GAIfo* uses visual data *both* in the process of discriminator *and* policy learning. That is, the learned behavior policy maps images

o to actions using a convolutional neural network. The technique we propose, on the other hand, leverages proprioceptive information in the policy learning step, instead learning policies that map proprioceptive states s to actions using a multilayer perceptron architecture.

4 Proposed Method

As presented in Section 3, we are interested in the problem of imitation from observation (*IfO*), where an imitating agent has access to visual demonstrations, $\tau_e = \{o_t\}$, of an expert performing a task, and seeks to learn a behavior that is approximately the same as the expert’s. In many previous approaches to this problem, the imitator selects actions on the basis of visual self-observation alone (i.e., using images of itself). We hypothesize that also leveraging available proprioceptive state information, s , during the learning process will result in better and faster learning.

Inspired by *GAIL*, our algorithm is comprised of two pieces: (1) a generator, which corresponds to the imitation policy, and (2) a discriminator, which serves as the reward function for the imitator. We model the imitation policy as a multilayer perceptron (MLP), π_{θ} . The imitating agent, being aware of its own proprioceptive features s , feeds them into the policy network and receives as output a distribution over actions from which the selected action a can be sampled. The imitator then executes this action and we record a video of the resulting behavior. After several actions have been executed, we have accumulated a collection of visual observations of the imitator’s behavior, $\tau_i = \{o\}$.

Meanwhile, we use a convolutional neural network as a discriminator D_{ϕ} . Given visual observations of the demonstrator, τ_e , and observations of the imitator, τ_i , we train the discriminator to differentiate between the data coming from these different sources. Since single video frames lack observability in most cases, we instead stack four frames, $\{o_{t-2}, o_{t-1}, o_t, o_{t+1}\}$, and feed this stack as input to the discriminator.

We train the discriminator to output values closer to zero for the transitions coming from the expert, and values closer to one for those coming from the imitator. Therefore, the discriminator aims to solve the following optimization problem:

$$\max_{\phi} \left(\mathbb{E}_{\tau_i}[\log(D_{\phi}(o_{t-2} : o_{t+1}))] + \mathbb{E}_{\tau_e}[\log(1 - D_{\phi}(o_{t-2} : o_{t+1}))] \right). \quad (2)$$

The lower the value outputted by the discriminator, the higher the chance of the input being from the expert. Recall that the objective for the imitator is to mimic the demonstrator, which can be thought of as fooling the discriminator. Therefore, we use

$$-\left(\mathbb{E}_{\tau_i}[\log(D_{\phi}(o_{t-2} : o_{t+1}))] \right) \quad (3)$$

as the reward to update the imitation policy using *RL*. In particular, we use proximal policy optimization (*PPO*) [Schulman *et al.*, 2017] with gradient steps of

$$\mathbb{E}_{\tau_i}[\nabla_{\theta} \log \pi_{\theta}(a|s)Q(s, a)] - \lambda \nabla_{\theta} H(\pi_{\theta}), \quad (4)$$




Figure 1: A diagrammatic representation of our algorithm. A multilayer perceptron (MLP) is used to model the policy, which takes the proprioceptive features s_t as the input and outputs an action a . The agent then executes the action in its environment. While the agent executes the policy, a video of the resulting behavior is recorded. Stacks of four consecutive grayscale images ($o_{t-2} : o_{t+1}$) from both the demonstrator and the imitator are then prepared as the input for the discriminator, which is trained to discriminate between data coming from these two sources. Finally, the discriminator function is then used as the reward function to train the policy using *PPO* (not shown).

Algorithm 1

- 1: Initialize policy π_θ randomly
- 2: Initialize discriminator D_ϕ randomly
- 3: Obtain visual demonstrations $\tau_e = \{o\}$
- 4: **for** $i \leftarrow 0$ **to** N **do**
- 5: Execute π_θ and record video observation $\tau_i = \{o\}$
- 6: Update the discriminator D_ϕ using loss

$$-\left(\mathbb{E}_{\tau_i}[\log(D_\phi(o_{t-2} : o_{t+1}))] + \mathbb{E}_{\tau_e}[\log(1 - D_\phi(o_{t-2} : o_{t+1}))]\right)$$

- 7: Update π_θ by performing *PPO* updates with gradient steps of

$$\mathbb{E}_{\tau_i}[\nabla_\theta \log \pi_\theta(a|s)Q(s, a)] - \lambda \nabla_\theta H(\pi_\theta),$$

where

$$Q(\hat{s}_t, \hat{a}_t) = -\mathbb{E}_{\tau_i}[\log(D_\phi(o_{t-2} : o_{t+1}))|s_0 = \hat{s}_t, a_0 = \hat{a}_t]$$

- 8: **end for**
-

where $Q(s, a)$ is the state-action value, i.e. the potential reward that the agent receives starting from s and taking action a :

$$Q(\hat{s}_t, \hat{a}_t) = -\mathbb{E}_{\tau_i}[\log(D_\phi(o_{t-2} : o_{t+1}))|s_0 = \hat{s}_t, a_0 = \hat{a}_t]. \quad (5)$$

As presented, our algorithm uses the visual information in order to learn the reward function by comparing visual data generated by the imitator and the demonstrator. It also takes advantage of proprioceptive state features in the process of policy learning by learning a mapping from those features to actions using a reinforcement learning algorithm. Pseudocode and a diagrammatic representation of our proposed algorithm are presented in Algorithm 1 and Figure 1, respectively.

5 Experiments

The algorithm introduced above combines proprioceptive state information with video observations in an adversarial imitation learning paradigm. We hypothesize that using the extra state information in the proposed way will lead to both faster imitation learning and better performance on the imitated task when compared to similar techniques that ignore proprioception. In this section, we describe the experimental procedure by which we evaluated this hypothesis, and discuss the results.

5.1 Setup

We evaluated our method on a subset of the continuous control tasks available via OpenAI Gym [Brockman *et al.*, 2016] and the MuJoCo simulator [Todorov *et al.*, 2012]: MountainCarContinuous, InvertedPendulum, InvertedDoublePendulum, Hopper, Walker2d, HalfCheetah.

To generate the demonstration data, we first trained an expert agents using pure reinforcement learning (i.e., not from imitation). More specifically, we used proximal policy optimization (*PPO*) [Schulman *et al.*, 2017] and the ground truth reward function provided by OpenAI Gym. After the expert agents were trained, we recorded 64×64 , 30-fps video demonstrations of their behavior.

We compared the proposed method with three other imitation from observation algorithms that do *not* exploit the imitator’s proprioceptive state information: Time Contrastive Networks (TCN) [Sermanet *et al.*, 2018], Behavioral Cloning from Observation (BCO) [Torabi *et al.*, 2018], Generative Adversarial Imitation Learning (GAIfO) [Torabi *et al.*, 2019b; Torabi *et al.*, 2019a]².

5.2 Results

We hypothesized that our method would outperform the baselines with respect to two criteria: (1) the final performance of the trained imitator, i.e., how the imitator performs the task compared to the demonstrator (as measured by the ground truth reward functions), and (2) the speed of the imitation learning process as measured by number of learning iterations. The results shown here were generated using ten inde-

²The considered domains, methods, and implementations are presented in more detail in the longer version of the paper on arXiv [Torabi *et al.*, 2019c]

pendent trials, where each trial used a different random seed to initialize the environments, model parameters, etc.

Figure 2 depicts our experimental results pertaining to the first criterion, i.e., the final task performance of trained imitating agents in each domain. The rectangular bars and error bars represent the mean return and the standard error, respectively, as measured over 1000 trajectories. We report performance using a normalized task score, i.e., scores are scaled in such a way that the demonstrating agent’s performance corresponds to 1.0 and the performance of an agent with random behavior corresponds to 0.0. The x-axis represents the number of demonstration trajectories, i.e., videos, available to the imitator. In general, it can be seen that the proposed method indeed outperforms the baselines in almost all cases, which shows that using the available proprioceptive state information can make a remarkable difference in the final task performance achieved by imitation learning. In the particular case of InvertedPendulum, both *GAIfo* and the proposed method achieve a final task performance equal to that of the demonstrator, likely due to the simplicity of the task. However, for the rest of the tasks, it can be clearly seen that the proposed approach performs better than *GAIfo*³. Further, we can see that increasing the number of demonstrated trajectory results in increased task performance.

To validate our hypothesis with respect to learning speed, we also studied the transient performance of the various learning algorithms. Because only one other method, *GAIfo*, performed as well as the expert in only one domain, InvertedPendulum, Figure 3 only depicts the results for these algorithms in that domain. The x-axis shows the number of iterations, i.e., the number of update cycles for both the policy and the discriminator. Since updating the policy requires interaction with the environment, a smaller number of iterations also corresponds to less overhead during the learning process. As shown in the figure, our method converges to expert-level performance much faster than *GAIfo*, which supports our hypothesis that leveraging proprioception speeds the imitation learning process.

In Figure 2, we can see that two of the baseline methods—*BCO* and *TCN*—do not achieve task performance anywhere near that of the expert.

For InvertedPendulum and InvertedDoublePendulum, we suspect that *TCN* performs poorly due to possible overfitting of the learned state embedding to the specific demonstrations and, therefore, does not generalize well toward supporting the overall goal of keeping the pendulum balanced above the rod. For Hopper, Walker2d, and HalfCheetah, the poor performance of *TCN* may be due to the fact that the tasks are cyclical in nature and therefore not well-suited to the

³Note that the performance of *GAIfo* on Hopper is different from what was presented in the *GAIfo* paper [Torabi *et al.*, 2019b]. We hypothesize that the reason is twofold: (1) different physics engines—MuJoCo is used in this paper, but in the previous work [Torabi *et al.*, 2019b] Pybullet [Coumans and Bai, 2016 2017] was used, and (2) differences in video appearance—in this work we do not alter the default simulator parameters, whereas in the previous work [Torabi *et al.*, 2019b] some of the parameters were modified such as the colors used in the video frames in order to increase the contrast between the agent and the background.

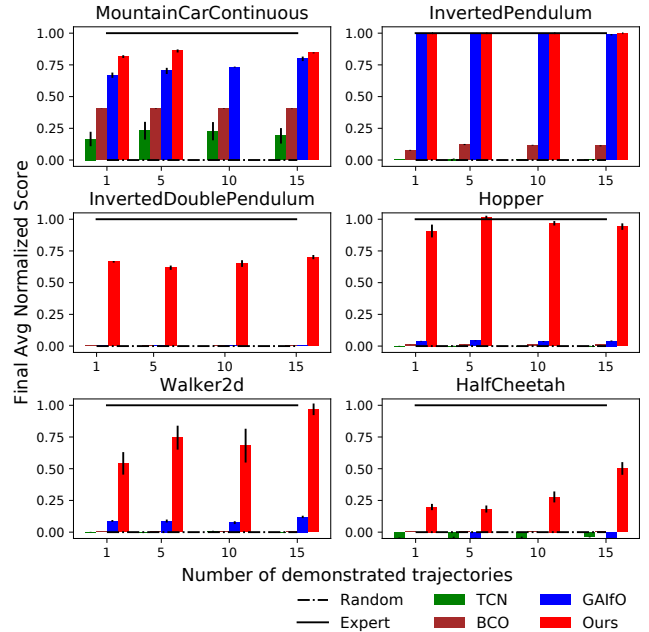


Figure 2: The rectangular bars and error bars represent the mean normalized return and the standard error, respectively, as measured over 1000 trials. The normalized values have been scaled in such a way that expert and random performance are 1.0 and 0.0, respectively. The x-axis represents the number of available video demonstration trajectories.

time-dependent learned state embedding. *TCN* performs relatively better in MountainCarContinuous, compared to other domains because this domain does have the properties required by *TCN*. As for *BCO*, we posit that the low performance is due to the well-known compounding-error issue present in behavioral cloning.

One interesting thing to note is that Walker2d results in larger error bars for our technique than those seen for any of the other domains. We hypothesize that the reason for this is that the video frames provide very poor information regarding the state of the demonstrator—here, the agent has two legs, which sometimes results in occlusion and, therefore, uncertainty regarding which action the agent should take.

Finally, we can see that the proposed technique performs the most poorly in the HalfCheetah domain. We hypothesize that this is due to the speed at which the demonstrator acts: frame-to-frame differences are large, e.g., three to four consecutive frames cover a complete cycle of the agent jumping forwards. This rate of change may make it difficult for our discriminator to extract a pattern of behavior, which, consequently, would make it much more difficult for the agent to move its behavior closer to that of the demonstrator. Therefore, one way that performance might be improved is to increase the frame rate at which the demonstrations are sampled. Another way, as suggested by Figure 2, would be to increase the number of demonstration trajectories beyond what is shown here.




Figure 3: Performance of imitation agents with respect to the number of iterations (N). Solid colored lines represent the mean return and shaded areas represent standard errors. The returns are scaled so that the performance of the expert and random policies be zero and one, respectively.

6 Conclusion and Future Work

In this paper, we hypothesized that including proprioception would be beneficial to the learning process in the *IfO* paradigm. To test this hypothesis, we presented a new imitation from observation algorithm that leverages both available visual and proprioceptive information. It uses visual information to compare the imitator’s behavior to that of the demonstrator, and uses this comparison as a reward function for training a policy over proprioceptive states. We showed that leveraging this state information can significantly improve both the performance and the efficiency of the learning process.

However, to achieve the end-goal of true imitation from observation, several challenges remain. For example, *IfO* algorithms should be able to overcome embodiment mismatch (the imitator and the demonstrator have different embodiments), and viewpoint mismatch (the visual demonstrations are recorded from different viewpoints.). Resolving these limitations is a natural next step for extending this research. Another way to improve upon the proposed method is to attempt to make the training more reliable by incorporating techniques developed to improve the stability of *GANs*, such as the work of Arjovsky *et al.* [2017]. Further, to the best of our knowledge, nobody has been able to deploy *GAN*-like methods on real robots due to high sample complexity. Therefore, techniques that seek to improve the learning process with respect to this metric should also be investigated further.

Acknowledgments

This work has taken place in the Learning Agents Research Group (LARG) at the Artificial Intelligence Laboratory, The University of Texas at Austin. LARG research is supported in part by grants from the National Science Foundation (IIS-1637736, IIS-1651089, IIS-1724157), the Office of Naval Research (N00014-18-2243), Future of Life Institute (RFP2-

000), Army Research Lab, DARPA, Intel, Raytheon, and Lockheed Martin. Peter Stone serves on the Board of Directors of Cogitai, Inc. The terms of this arrangement have been reviewed and approved by the University of Texas at Austin in accordance with its policy on objectivity in research.

References

- [Abbeel and Ng, 2004] Pieter Abbeel and Andrew Y Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first International Conference on Machine learning*, page 1. ACM, 2004.
- [Argall *et al.*, 2009] Brenna D Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57(5):469–483, 2009.
- [Arjovsky *et al.*, 2017] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *ICML*, pages 214–223, 2017.
- [Aytar *et al.*, 2018] Yusuf Aytar, Tobias Pfaff, David Budden, Thomas Paine, Ziyu Wang, and Nando de Freitas. Playing hard exploration games by watching youtube. In *NeurIPS*, pages 2935–2945, 2018.
- [Bagnell *et al.*, 2007] JA Bagnell, Joel Chestnutt, David M Bradley, and Nathan D Ratliff. Boosting structured prediction for imitation learning. In *Advances in Neural Information Processing Systems*, pages 1153–1160, 2007.
- [Bain and Sammut, 1995] Michael Bain and Claude Sammut. A framework for behavioral cloning. *Machine Intelligence 14*, 1995.
- [Bain and Sommut, 1999] Michael Bain and Claude Sommut. A framework for behavioural cloning. *Machine Intelligence 15*, 15:103, 1999.
- [Baker *et al.*, 2009] Chris L Baker, Rebecca Saxe, and Joshua B Tenenbaum. Action understanding as inverse planning. *Cognition*, 113(3):329–349, 2009.
- [Brockman *et al.*, 2016] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI Gym, 2016.
- [Coumans and Bai, 2016 2017] Erwin Coumans and Yunfei Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning. 2016-2017.
- [Daftry *et al.*, 2016] Shreyansh Daftry, J Andrew Bagnell, and Martial Hebert. Learning transferable policies for monocular reactive mav control. In *International Symposium on Experimental Robotics*, pages 3–11. Springer, 2016.
- [Edwards *et al.*, 2019] Ashley D Edwards, Himanshu Sahni, Yannick Schroeker, and Charles L Isbell. Imitating latent policies from observation. *ICML*, 2019.
- [Fu *et al.*, 2018] Justin Fu, Katie Luo, and Sergey Levine. Learning robust rewards with adversarial inverse reinforcement learning. In *ICLR*, 2018.
- [Ganin *et al.*, 2016] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François

- Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *JMLR*, 17(1):2096–2030, 2016.
- [Goo and Niekum, 2019] Wonjoon Goo and Scott Niekum. One-shot learning of multi-step tasks from observation via activity localization in auxiliary video. *ICRA*, 2019.
- [Goodfellow *et al.*, 2014] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, 2014.
- [Gupta *et al.*, 2017] Abhishek Gupta, Coline Devin, YuXuan Liu, Pieter Abbeel, and Sergey Levine. Learning invariant feature spaces to transfer skills with reinforcement learning. In *International Conference on Learning Representations*, 2017.
- [Ho and Ermon, 2016] Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. In *Advances in Neural Information Processing Systems*, pages 4565–4573, 2016.
- [Kostrikov *et al.*, 2019] Ilya Kostrikov, Kumar Krishna Agrawal, Debidatta Dwibedi, Sergey Levine, and Jonathan Tompson. Discriminator-actor-critic: Addressing sample inefficiency and reward bias in adversarial imitation learning. In *ICLR*, 2019.
- [Liu *et al.*, 2018] YuXuan Liu, Abhishek Gupta, Pieter Abbeel, and Sergey Levine. Imitation from observation: Learning to imitate behaviors from raw video via context translation. 2018.
- [Merel *et al.*, 2017] Josh Merel, Yuval Tassa, Sriram Srinivasan, Jay Lemmon, Ziyu Wang, Greg Wayne, and Nicolas Heess. Learning human behaviors from motion capture by adversarial imitation. *arXiv preprint arXiv:1707.02201*, 2017.
- [Misra *et al.*, 2016] Ishan Misra, C Lawrence Zitnick, and Martial Hebert. Shuffle and learn: unsupervised learning using temporal order verification. In *ECCV*, pages 527–544. Springer, 2016.
- [Osa *et al.*, 2018] Takayuki Osa, Joni Pajarinen, Gerhard Neumann, J Andrew Bagnell, Pieter Abbeel, Jan Peters, et al. An algorithmic perspective on imitation learning. *Foundations and Trends® in Robotics*, 7(1-2):1–179, 2018.
- [Pavse *et al.*, 2019] Brahma Pavse, Faraz Torabi, Josiah Hanna, Garrett Warnell, and Peter Stone. Ridm: Reinforced inverse dynamics modeling for learning from a single observed demonstration. In *International Conference on Machine Learning Workshop on Imitation, Intent, and Interaction (I3)*, 2019.
- [Qureshi *et al.*, 2019] Ahmed H. Qureshi, Byron Boots, and Michael C. Yip. Adversarial imitation via variational inverse reinforcement learning. In *ICLR*, 2019.
- [Ross *et al.*, 2011] Stéphane Ross, Geoffrey J Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *International Conference on Artificial Intelligence and Statistics*, pages 627–635, 2011.
- [Sasaki *et al.*, 2019] Fumihiko Sasaki, Tetsuya Yohira, and Atsuo Kawaguchi. Sample efficient imitation learning for continuous control. In *ICLR*, 2019.
- [Schaal, 1997] Stefan Schaal. Learning from demonstration. In *Advances in Neural Information Processing Systems*, pages 1040–1046, 1997.
- [Schulman *et al.*, 2015] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *ICML*, pages 1889–1897, 2015.
- [Schulman *et al.*, 2017] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [Sermanet *et al.*, 2018] Pierre Sermanet, Corey Lynch, Jasmine Hsu, and Sergey Levine. Time-contrastive networks: Self-supervised learning from multi-view observation. In *International Conference in Robotics and Automation (ICRA)*, 2018.
- [Stadie *et al.*, 2017] Bradly C Stadie, Pieter Abbeel, and Ilya Sutskever. Third-person imitation learning. In *ICLR*, 2017.
- [Sutton and Barto, 1998] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge, 1998.
- [Todorov *et al.*, 2012] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *IROS*, pages 5026–5033. IEEE, 2012.
- [Torabi *et al.*, 2018] Faraz Torabi, Garrett Warnell, and Peter Stone. Behavioral cloning from observation. In *IJCAI*, pages 4950–4957, 2018.
- [Torabi *et al.*, 2019a] Faraz Torabi, Garrett Warnell, and Peter Stone. Adversarial imitation learning from state-only demonstrations. In *AAMAS*, 2019.
- [Torabi *et al.*, 2019b] Faraz Torabi, Garrett Warnell, and Peter Stone. Generative adversarial imitation from observation. In *International Conference on Machine Learning Workshop on Imitation, Intent, and Interaction (I3)*, 2019.
- [Torabi *et al.*, 2019c] Faraz Torabi, Garrett Warnell, and Peter Stone. Imitation learning from video by leveraging proprioception. *arXiv preprint arXiv:1905.09335*, 2019.
- [Torabi *et al.*, 2019d] Faraz Torabi, Garrett Warnell, and Peter Stone. Recent advances in imitation learning from observation. In *IJCAI*. AAAI Press, 2019.
- [Vered *et al.*, 2016] Mor Vered, Gal A Kaminka, and Sivan Biham. Online goal recognition through mirroring: Humans and agents. In *The Fourth Annual Conference on Advances in Cognitive Systems*, 2016.
- [Vered *et al.*, 2018] Mor Vered, Ramon Fraga Pereira, Maurício C Magnaguagno, Gal A Kaminka, and Felipe Meneguzzi. Towards online goal recognition combining goal mirroring and landmarks. In *AAMAS*, pages 2112–2114, 2018.