

Cross-domain Imitation from Observations

Raychaudhuri, Dripta S.; van Baar, Jeroen; Paul, Sujoy; Roy-Chowdhury, Amit K.

TR2021-074 June 12, 2021

Abstract

Imitation learning seeks to circumvent the difficulty in designing proper reward functions for training agents by utilizing expert behavior. With environments modeled as Markov Decision Processes (MDP), most of the existing imitation algorithms are contingent on the availability of expert demonstrations in the same MDP as the one in which a new imitation policy is to be learned. In this paper, we study how to imitate tasks when discrepancies exist between the expert and agent MDP. These discrepancies across domains could include differing dynamics, viewpoint, or morphology; we present a novel framework to learn correspondences across such domains. Importantly, in contrast to prior works, we use unpaired and unaligned trajectories containing only states in the expert domain, to learn this correspondence. We utilize a cycle-consistency constraint on both the state space and a domain agnostic latent space to accomplish this. In addition, we enforce consistency on the temporal position of states via a normalized position estimator function, to align the trajectories across the two domains. Once this correspondence is found, we can directly transfer the demonstrations on one domain to the other and use it for imitation. Experiments across a wide variety of challenging domains demonstrate the efficacy of our approach.

International Conference on Machine Learning (ICML)

© 2021 MERL. This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Cross-domain Imitation from Observations

Dripta S. Raychaudhuri^{*1} Sujoy Paul^{*2†} Jeroen van Baar³ Amit K. Roy-Chowdhury¹

Abstract

Imitation learning seeks to circumvent the difficulty in designing proper reward functions for training agents by utilizing expert behavior. With environments modeled as Markov Decision Processes (MDP), most of the existing imitation algorithms are contingent on the availability of expert demonstrations in the *same* MDP as the one in which a new imitation policy is to be learned. In this paper, we study how to imitate tasks when discrepancies exist between the expert and agent MDP. These discrepancies across domains could include differing dynamics, viewpoint, or morphology; we present a novel framework to learn correspondences across such domains. Importantly, in contrast to prior works, we use *unpaired* and *unaligned* trajectories containing *only states* in the expert domain, to learn this correspondence. We utilize a cycle-consistency constraint on both the state space and a domain agnostic latent space to accomplish this. In addition, we enforce consistency on the temporal position of states via a normalized position estimator function, to align the trajectories across the two domains. Once this correspondence is found, we can directly transfer the demonstrations on one domain to the other and use it for imitation. Experiments across a wide variety of challenging domains demonstrate the efficacy of our approach.

1. Introduction

Humans possess the innate ability to quickly pick up a new behavior by simply observing others performing the same skill. Not only are we able to learn from demonstrations coming from a third-person point of view, we are also capable of imitation from experts who are morphologically different or have different embodiments - as evidenced by a child imitating an adult with different biomechanics (Jones, 2009). Previous works in neuroscience (Rizzolatti & Craighero, 2004; Marshall & Meltzoff, 2015) have attributed this to the human capacity of learning structure preserving domain correspondences via an invariant feature space (Umiltà et al., 2008), which allows us to re-

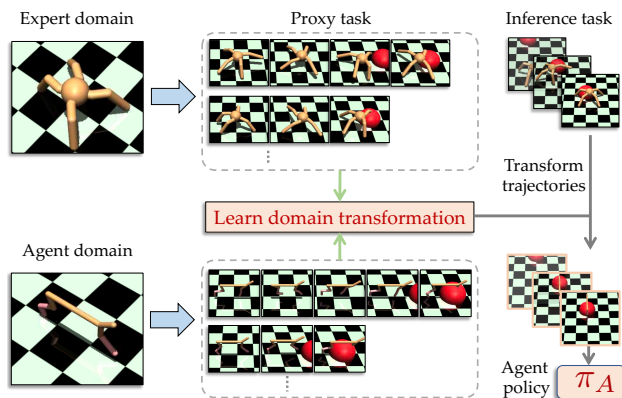


Figure 1. **Problem overview.** Cross-domain Imitation from Observation (\times DIO) entails learning from experts which are different from the agent. Here, the expert is a 4-legged Ant, while the agent is a HalfCheetah. We learn a domain transformation function from unpaired, unaligned, state-only trajectories from a set of proxy tasks and utilize it to imitate the expert on the given inference task.

construct the observed behavior in the self-domain. While imitation learning algorithms (Ho & Ermon, 2016; Ross et al., 2011) are successful, to some extent, in endowing autonomous agents with this ability to imitate expert behavior, they impose the somewhat unrealistic requirement that the demonstrations must come from the same domain, whether that be first-person viewpoint, same morphology or similar dynamics. The question then arises: *can we perform imitation learning which can overcome all such domain discrepancies?*

Prior work on bridging domain disparities in imitation learning have focused on each of these differences in isolation: morphology (Gupta et al., 2017), dynamics (Gangwani & Peng, 2019) and viewpoint mismatch (Stadie et al., 2017; Sharma et al., 2019; Liu et al., 2018). These works (Gupta et al., 2017; Liu et al., 2018; Sharma et al., 2019) utilize *paired, time-aligned* demonstrations from both domains, on a set of *proxy* tasks, to first build a correspondence map across the domains and then perform an extra reinforcement learning (RL) step for learning the final policy on the given task. This limits their applicability since paired demonstrations are rarely available and RL procedures are expensive.

Recently, (Kim et al., 2020) proposed a general framework which can perform imitation across a wide array of such dis-

crepancies from *unpaired, unaligned* demonstrations. However, they require *expert actions*, such as the exact kinematic forces, in order to learn a domain correspondence and assume availability of an expert policy which is utilized in an *interactive learning* setting. This is distinctly different to how humans imitate: we are capable of learning behaviors solely from observations/states, without access to underlying actions. Furthermore, continuously querying the expert might be onerous in several situations. Thus, we require a mechanism for learning policies from observation alone, where the expert demonstrations can originate in a domain which is different from the agent domain and access to the expert is limited. We define this setting as *Cross Domain Imitation from Observation* ($\times\text{DIO}$).

In this work, we propose a novel framework to tackle the $\times\text{DIO}$ problem, encompassing morphological, viewpoint and dynamics mismatch. We follow a two-step approach (see Fig. 1), where we first learn a transformation across the domains using the proxy tasks (Gupta et al., 2017), followed by a transfer process and subsequent learning of the policy. Importantly, in contrast to previous work, we use *unpaired* and *unaligned* trajectories containing *only states* on the expert domain trajectories, to learn this transformation. Additionally, we do not assume any access to the expert policy or the expert domain except for the given demonstrations. To learn the state correspondences, we jointly minimize a divergence between the transition distributions in the state space as well as in the latent space between the expert and the agent proxy task trajectories, while learning to translate between the two domains with the unpaired data via cycle-consistency (Zhu et al., 2017). However, solely learning with such state cycle-consistency may only result in local alignment, and lead to difficulties in optimizing for complex environments. Thus, to impose global alignment, we enforce additional consistency on the temporal position of states across the two domains. This ensures that when a state is mapped from one domain to the other, the degree of completion associated by being in that state remains unchanged. Having learnt this mapping on the proxy tasks, we transfer demonstrations for a new *inference* task from the expert to the agent domain, which are subsequently utilized to learn a policy via imitation.

Experiments on a wide array of domains that encompass dynamics, morphological and viewpoint mismatch, demonstrate the feasibility of learning domain correspondences from unpaired and unaligned state-only demonstrations. The primary contributions of this work are as follows:

1. We propose an algorithm for cross-domain imitation learning by learning transformations across domains, modeled as Markov Decision Processes (MDP), from *unpaired, unaligned, state-only* demonstrations, thereby ameliorating the need for costly paired, aligned data.

2. Unlike previous work, we do not utilize any costly RL procedure. Neither do we require interactive querying of an expert policy.
3. We adopt multiple tasks in the MuJoCo physics engine (Todorov et al., 2012), and show that our framework can find correspondences and align two domains across *different viewpoints, dynamics and morphologies*.

2. Related Works

Imitation learning. Imitation learning (Schaal, 1999) uses a set of expert demonstrations to learn a policy which successfully mimics the expert. A common approach is behavioral cloning (BC) (Pomerleau, 1989; Bojarski et al., 2016), which amounts to learning to mimic the expert demonstrations via supervised learning. Inverse reinforcement learning (IRL) is another approach, where one seeks to learn a reward function that explains the demonstrated actions (Ho & Ermon, 2016; Abbeel & Ng, 2004; Ziebart et al., 2008). Recent works (Torabi et al., 2018; Yang et al., 2019; Paul et al., 2019) extend imitation learning to state-only demonstrations, where expert actions are not observed - this opens up the possibility of using imitation in robotics and learning from weak-supervision sources such as videos. Unlike these approaches, our work tackles the problem of imitation from state-only demonstrations coming from a *different* domain.

Domain transfer in reinforcement learning. Transfer in the reinforcement learning setting has been attempted by a wide array of works (Taylor & Stone, 2009). (Ammar & Taylor, 2011) manually define a common state space between MDPs and use it to learn a mapping between states. Unsupervised manifold alignment is used in (Ammar et al., 2015) to learn a linear map between states with similar local geometric properties. However, they assume the existence of hand-crafted features along with a distance metric between them, which limits its applicability. Recent works in transfer learning across mismatches in embodiment (Gupta et al., 2017) and viewpoint (Liu et al., 2018; Sharma et al., 2019), obtain state correspondences from an proxy task set comprising paired, time-aligned demonstrations and use them to learn a state map or a state encoder to a domain invariant feature space. (Kim et al., 2020) proposed a framework which can learn a map across domains from unpaired, unaligned demonstrations. However, they require expert actions to train the framework, along with access to an online expert. Furthermore, most of these approaches (Gupta et al., 2017; Liu et al., 2018) utilize an RL step which incurs additional computational cost. In contrast to these methods, our approach learns an MDP structure preserving state map from unpaired, unaligned demonstrations without requiring access to expert actions, additional RL or online experts. Viewpoint agnostic imitation has also been tackled in (Stadie et al., 2017), where a combination of adversarial learning (Ho & Ermon, 2016) and domain confusion (Tzeng et al.,

Table 1. Comparison to prior work using attributes demonstrated in the paper. \times DIO satisfies all the criteria desired in a holistic domain adaptive imitation framework.

METHOD	UNPAIRED TRAJECTO- RIES	ONLY STATES	NO ONLINE EXPERT	NO RL
IF (Gupta et al., 2017)	✗	✓	✓	✗
DAIL (Kim et al., 2020)	✓	✗	✗	✓
Ours	✓	✓	✓	✓

2014) is used to learn a policy without an proxy set. However, it fails to account for large variations in viewpoint, in addition to sub-optimal trajectories from the expert domain. From a theoretical perspective, our approach aligns with the objective of MDP homomorphisms (Ravindran, 2004). Similar ideas are explored in learning MDP similarity metric via bisimulation (Ferns et al., 2011) and Boltzmann machine reconstruction error (Ammar et al., 2014). However, these works find homomorphisms within an MDP and do not provide ways to discover homomorphisms across MDPs.

Cycle-consistency. Our work draws inspiration from the literature on cycle-consistency (Zhu et al., 2017; Hoffman et al., 2018; Smith et al., 2019). CycleGAN (Zhu et al., 2017) introduced cycle-consistency to learn bidirectional transformations between domains via Generative Adversarial Networks (Goodfellow et al., 2014) for unpaired image-to-image translation. This was extended to domain adaptation in (Hoffman et al., 2018). Similar techniques are applied in sim-to-real transfer (Ho et al., 2020; Gamrian & Goldberg, 2019). Recently, (Rao et al., 2020) propose RL-CycleGAN to perform sim-to-real transfer by adding extra supervision from the Q-value function. Unlike these works, which are restricted on visual alignments, we propose to learn alignments across differing dynamics/morphology.

3. Problem Setting

Before formally defining the \times DIO problem, we first lay the groundwork in terms of notation. Following (Kim et al., 2020), we define a domain as a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{P}_0)$, where \mathcal{S} denotes the state space, \mathcal{A} is the action space, \mathcal{P} is the dynamics or transition function, and \mathcal{P}_0 is the initial distribution over the states. Given an action $a \in \mathcal{A}$, the next state is governed by the transition dynamics as $s' \sim \mathcal{P}(s'|s, a)$. An infinite horizon Markov Decision Process (MDP) is defined subsequently by including a reward function $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, and a discount factor $\gamma \in [0, 1]$ to the domain tuple. Thus, while the domain typifies only the agent morphology and the dynamics, augmenting the domain with a reward and discount factor describes an MDP for a particular task. We define an MDP in some domain x for a task \mathcal{T} as $\mathcal{M}_x^{\mathcal{T}} = (\mathcal{S}_x, \mathcal{A}_x, \mathcal{P}_x, r_x^{\mathcal{T}}, \gamma_x^{\mathcal{T}}, \mathcal{P}_{0x})$. A policy is a map $\pi_x^{\mathcal{T}} : \mathcal{S}_x \rightarrow \mathcal{B}(\mathcal{A}_x)$, where \mathcal{B} is the set of all probability measures on \mathcal{A}_x . A trajectory corresponding to the task \mathcal{T} in

domain x is a sequence of states $\eta_{\mathcal{M}_x^{\mathcal{T}}} = \{s_x^0, s_x^1, \dots, s_x^{H_\eta}\}$, where H_η denotes the length of the trajectory. We denote $\mathcal{D}_{\mathcal{M}_x^{\mathcal{T}}} = \{\eta_{\mathcal{M}_x^{\mathcal{T}}}^i\}_{i=1}^N$ to be a set of such trajectories. In our work, we consider two domains: expert and agent, indicated by $\mathcal{M}_E^{\mathcal{T}}$ and $\mathcal{M}_A^{\mathcal{T}}$ respectively.

The objective of \times DIO is to learn an optimal policy $\pi_A^{\mathcal{T}}$ in the agent domain, given state-only demonstrations $\mathcal{D}_{\mathcal{M}_E^{\mathcal{T}}}$ in the expert domain. In this paper, we propose to first learn a transformation $\psi : \mathcal{S}_E \rightarrow \mathcal{S}_A$ between the domains and then leverage ψ to imitate from the expert demonstrations. Following prior work (Gupta et al., 2017; Liu et al., 2018; Kim et al., 2020), we assume access to a dataset consisting of expert-agent trajectories for M different proxy tasks: $\mathcal{D} = \{(\mathcal{D}_{\mathcal{M}_E^{\mathcal{T}_j}}, \mathcal{D}_{\mathcal{M}_A^{\mathcal{T}_j}})\}_{j=1}^M$. Proxy tasks encompass simple primitive skills in both domains and are different from the inference task \mathcal{T} , for which we wish to learn the policy.

We relax certain assumptions made in previous work, which are critical for real-world applications. Firstly, the trajectories derived from proxy tasks are not paired, i.e., time-aligned trajectories do not exist in \mathcal{D} . This is crucial in real-world cases, as the tasks may not be executed at the same rate in different domains. Secondly, expert actions are not observed: such actions are difficult to obtain in various scenarios such as videos of humans performing some task. Finally, we train in an offline fashion and do not require any expert policy for interactive querying, to guide the learning process, beyond the provided demonstrations. Table 1 explicitly details how our setting differs from the ones tackled in the literature.

Once the domain transformation function ψ is learnt, we use it to translate the expert domain trajectories $\mathcal{D}_{\mathcal{M}_E^{\mathcal{T}}}$, for the inference task \mathcal{T} , to the agent domain to obtain $\mathcal{D}_{\mathcal{M}_A^{\mathcal{T}}}$. An inverse dynamics model $\mathcal{I}_A : \mathcal{S}_A \times \mathcal{S}_A \rightarrow \mathcal{A}_A$ is then learnt to augment these translated trajectories with actions, similar to (Torabi et al., 2018). These are subsequently used to learn the policy $\pi_A^{\mathcal{T}}$ via behavioral cloning.

4. Method

A crucial characteristic of a good domain transformation ψ lies in MDP dynamics preservation. In our framework, we enforce this from both the local and global perspectives. For local alignment, we aim to ensure that optimal state transitions in $\mathcal{M}_E^{\mathcal{T}}$ map to optimal transitions in $\mathcal{M}_A^{\mathcal{T}}$. Our proposed method achieves this local alignment by matching the state-transition distributions defined for the true and transferred trajectories on the proxy tasks in an adversarial manner, while maintaining cycle-consistency. A latent space is learned via a mutual information objective to only preserve task-specific information. On the other hand, a learned temporal position function aims to enforce consistency on the temporal position of the states across the two

domains to ensure global alignment. In the following parts, we define each of these components in more detail.

4.1. Local alignment via distribution matching

State cycle-consistency. We seek to map optimal transitions in the expert domain to the agent domain, and propose to learn domain transformation ψ such that the state transition distribution is matched over the trajectories derived from the proxy tasks. We utilize adversarial training to accomplish this. Given unpaired samples $\{(s_E^t, s_E^{t+1})\} \in \mathcal{D}_{\mathcal{M}_E^{\mathcal{T}_j}}$ and $\{(s_A^t, s_A^{t+1})\} \in \mathcal{D}_{\mathcal{M}_A^{\mathcal{T}_j}}$ drawn from the j^{th} proxy task, the function ψ is learned in an adversarial manner with a discriminator D_A^j , where ψ tries to map (s_E^t, s_E^{t+1}) onto the distribution of (s_A^t, s_A^{t+1}) , while D_A^j tries to distinguish translated samples $(\psi(s_E^t), \psi(s_E^{t+1}))$ against real samples (s_A^t, s_A^{t+1}) :

$$\begin{aligned} \min_{\psi} \max_{D_A^j} \mathcal{L}_{adv}^j &= \mathbb{E}_{(s_A^t, s_A^{t+1}) \sim \mathcal{D}_{\mathcal{M}_A^{\mathcal{T}_j}}} \left[\log D_A^j(s_A^t, s_A^{t+1}) \right] \\ &+ \mathbb{E}_{(s_E^t, s_E^{t+1}) \sim \mathcal{D}_{\mathcal{M}_E^{\mathcal{T}_j}}} \left[\log(1 - D_A^j(\psi(s_E^t), \psi(s_E^{t+1}))) \right] \quad (1) \end{aligned}$$

Solely optimizing this adversarial loss can lead to the model mapping the same set of states to any random permutation of states in the agent domain, where any of the learned mappings can induce an output distribution that matches the agent state transition distribution. Following (Zhu et al., 2017), we introduce cycle consistency as a means to control this undesired effect. We learn another state map in the opposite direction $\phi : \mathcal{S}_A \rightarrow \mathcal{S}_E$ by optimizing an adversarial loss, $\min_{\phi} \max_{D_E^j} \mathcal{L}_{adv}^j$, with a discriminator D_E^j . Cycle consistency is then imposed as:

$$\begin{aligned} \min_{\psi, \phi} \mathcal{L}_{cyc}^j &= \mathbb{E}_{s_E \sim \mathcal{D}_{\mathcal{M}_E^{\mathcal{T}_j}}} [\|\phi \circ \psi(s_E) - s_E\|_2^2] + \\ &\mathbb{E}_{s_A \sim \mathcal{D}_{\mathcal{M}_A^{\mathcal{T}_j}}} [\|\psi \circ \phi(s_A) - s_A\|_2^2] \quad (2) \end{aligned}$$

Domain invariant latent space. To incentivize ψ, ϕ to generalize beyond proxy tasks, we use an encoder-decoder structure for the transformation function ψ . Concretely, $\psi = D_E \circ E_E$, where $E_E : \mathcal{S}_E \rightarrow \mathcal{Z}$ represents an encoder which maps a state in the expert domain to a domain agnostic latent space \mathcal{Z} , while $D_E : \mathcal{Z} \rightarrow \mathcal{S}_A$ represents the decoding function. $\phi = D_A \circ E_A$ is defined similarly via the same latent space \mathcal{Z} . Prior work (Gupta et al., 2017) has explored learning such invariant spaces, but use paired data from both domains, which is a very strong and often unrealistic assumption, as explained above. Inspired from work based on information theoretic objectives (Eysenbach et al., 2018; Wan et al., 2020), we learn the latent space by minimizing the mutual information between the domain and the latent transitions:

$$\min_{E_E, E_A} I(d; (z^t, z^{t+1})) \quad (3)$$

where (z^t, z^{t+1}) denotes an encoded transition from either of the domains. Minimizing the mutual information between the domain $(\Delta = \{E, A\})$ and the encoded latent transition for the *same* proxy task will result in a latent space which encodes the *task-specific* information and filters out the *domain-specific* nuances.

Note that we can decompose the mutual information term as $I(\Delta; (z^t, z^{t+1})) = H(\Delta) - H(\Delta|(z^t, z^{t+1}))$, where $H(\cdot)$ denotes the entropy. Thus, our objective in Equation 3 reduces to just maximizing the conditional entropy $H(\Delta|(z^t, z^{t+1}))$. Due to intractability of this expression (Aleml et al., 2016; Poole et al., 2019), we optimize a variational lower bound $\mathbb{E}_{d \sim \Delta, (s_d^t, s_d^{t+1}) \sim \mathcal{D}_{\mathcal{M}_d^{\mathcal{T}_j}}} [-\log q^j(d|(z^t, z^{t+1}))]$ instead, where q^j denotes a variational distribution which approximates the true posterior.

Here, q^j is parameterized as a discriminator which outputs the probability that the generated transition comes from domain d for the j th proxy task. Maximizing this objective over the encoder parameters ensures that the discriminator is maximally confused and the latent transitions for the task, coming from both domains, are well aligned. The overall objective is as follows:

$$\min_{q^j} \max_{E_E, E_A} \mathcal{L}_{MI} = \mathbb{E}_{d \sim \Delta, (s_d^t, s_d^{t+1}) \sim \mathcal{D}_{\mathcal{M}_d^{\mathcal{T}_j}}} [-\log q^j(d|(z^t, z^{t+1}))] \quad (4)$$

Additionally, we enforce consistency in the latent embedding to further constrain the learnt mapping:

$$\begin{aligned} \min_{\psi, \phi} \mathcal{L}_z^j &= \mathbb{E}_{s_E \sim \mathcal{D}_{\mathcal{M}_E^{\mathcal{T}_j}}} [\|E_A \circ \psi(s_E) - E_E(s_E)\|_2^2] \\ &+ \mathbb{E}_{s_A \sim \mathcal{D}_{\mathcal{M}_A^{\mathcal{T}_j}}} [\|E_E \circ \phi(s_A) - E_A(s_A)\|_2^2] \quad (5) \end{aligned}$$

4.2. Global alignment via temporal position preservation

Solely learning with state cycle-consistency may result only in local alignment: an optimal state pair in the expert domain may get mapped to an optimal transition in the agent domain while violating task semantics (transitions from beginning of a task get mapped to terminal ones), and then back without breaking cycle-consistency. In order to constrain the mapping to maintain temporal semantics for a task, we enforce additional consistency on the temporal position of states across the two domains.

We encode the temporal position of a state by computing a normalized score of proximity to the terminal state in the trajectory. Each state is assigned a value of 1 if they are terminating goal states and 0 otherwise. These discrete values are then exponentially weighted by a discount factor $\gamma \in (0, 1)$ to obtain a continuous estimate of the state

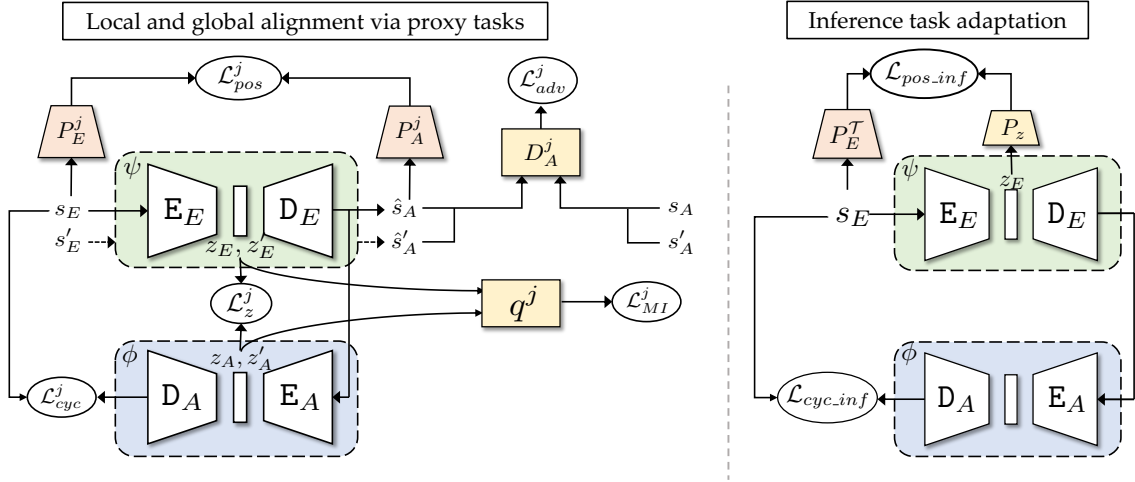


Figure 2. Framework overview. An illustration of our MDP correspondence learning framework. We perform local alignment via state-transition distribution matching and cycle-consistency in the state space using \mathcal{L}_{adv}^j and \mathcal{L}_{cyc}^j , as well as in a learnt latent space using \mathcal{L}_z^j and \mathcal{L}_{MI}^j (only proxy task is j shown here). The inverse cycle from agent to expert is omitted here for clarity. Global alignment is performed via consistency on the temporal position of states across the two domains, using the pre-trained position estimators P_A^j, P_E^j in \mathcal{L}_{pos}^j . Further improvement is obtained via inference task adaptation using $\mathcal{L}_{pos.inf}^j$ and $\mathcal{L}_{cyc.inf}^j$ - this prevents overfitting to the proxy tasks and makes the learnt transformation more robust and well-conditioned to the target data.

temporal position. Using these temporal encodings, we pre-train temporal position estimators P_E^j, P_A^j in a supervised fashion by optimizing a squared error loss as follows:

$$\min_{P_E^j} \mathbb{E}_{\eta \sim \mathcal{D}} \sum_{t=1}^{H_\eta} \left(P_E^j(s_E^t) - \gamma^{H_\eta - t} \right)^2 \quad (6)$$

P_A^j is learnt in a similar fashion by optimizing Equation 6 with respect to the agent trajectories. These estimators are subsequently used to enforce temporal preservation as:

$$\begin{aligned} \min_{\psi, \phi} \mathcal{L}_{pos}^j &= \mathbb{E}_{s_E \sim \mathcal{D}_{\mathcal{M}_E^j}} \left[\|P_A^j \circ \psi(s_E) - P_E^j(s_E)\|_2^2 \right] \\ &+ \mathbb{E}_{s_A \sim \mathcal{D}_{\mathcal{M}_A^j}} \left[\|P_E^j \circ \phi(s_A) - P_A^j(s_A)\|_2^2 \right]. \quad (7) \end{aligned}$$

Our temporal position estimators may be interpreted as state value functions: trajectories are from a greedy optimal policy with reward 1 for terminal states, and 0 otherwise.

4.3. Inference task adaptation

As discussed in Section 3, we are provided with the state-only trajectories $\mathcal{D}_{\mathcal{M}_E^T}$ on solely the expert domain for the inference task \mathcal{T} . We propose to use these trajectories during the learning process as additional regularization, referred to as inference task adaptation. First, we enforce cycle consistency on the states:

$$\min_{\psi, \phi} \mathcal{L}_{cyc.inf} = \mathbb{E}_{s_E \sim \mathcal{D}_{\mathcal{M}_E^T}} \left[\|\phi \circ \psi(s_E) - s_E\|_2^2 \right]. \quad (8)$$

In addition, we also enforce temporal preservation in the latent space. Concretely, we first train a position estimator

P_E^T by optimizing Equation 6 on the given trajectories as discussed in Section 4.2. We use the trained position estimator, along with a latent space position predictor P_z to enforce temporal preservation by:

$$\min_{E_E, P_z} \mathcal{L}_{pos.inf} = \mathbb{E}_{s_E \sim \mathcal{D}_{\mathcal{M}_E^T}} \left[\|P_z \circ E_E(s_E) - P_E^T(s_E)\|_2^2 \right]. \quad (9)$$

4.4. Optimization

Given the alignment dataset D containing trajectories from the M proxy tasks, we first pre-train the temporal position estimators $\{(P_E^j, P_A^j)\}_{j=1}^M$ using Equation 6. This is followed by adversarial training of the state maps ψ, ϕ , where we use separate discriminators on the state space and latent space for each proxy task. The full objective is then:

$$\begin{aligned} \min_{\psi, \phi} \max_{\{D_E^j\}, \{D_A^j\}, \{q^j\}} \mathcal{L} &= \sum_{j=1}^M \left[\lambda_1 \left(\mathcal{L}_{adv}^j(D_A^j) + \mathcal{L}_{adv}^j(D_E^j) \right) \right. \\ &+ \lambda_2 \left(\mathcal{L}_{cyc}^j + \mathcal{L}_z^j \right) + \lambda_3 \mathcal{L}_{pos}^j - \lambda_4 \mathcal{L}_{MI}^j \left. \right] \\ &+ \lambda_5 \left(\mathcal{L}_{cyc.inf} + \mathcal{L}_{pos.inf} \right), \quad (10) \end{aligned}$$

where $\{\lambda_i\}_{i=1}^5$ denote hyper-parameters which control the contribution of each loss term. A pictorial description of the overall framework is shown in Figure 2.

4.5. Imitation from observation

We use the learned ψ to map the states in the inference task expert demonstrations $\mathcal{D}_{\mathcal{M}_E^T}$ to the agent domain. Given

Table 2. Cross-domain imitation performance of the policy learnt on transferred trajectories for inference tasks. All rewards are normalized by expert performance on corresponding task.

METHOD	V-R2R	V-R2W	D-R2R	M-R2R	M-A2A	M-A2C
IF	0.32 ± 0.10	0.57 ± 0.20	0.48 ± 0.30	0.61 ± 0.23	0.09 ± 0.08	0.00 ± 0.00
CCA	0.16 ± 0.27	0.86 ± 0.30	0.47 ± 0.20	0.16 ± 0.13	0.30 ± 0.30	0.75 ± 0.50
CYCLEGAN	0.17 ± 0.10	0.72 ± 0.16	0.13 ± 0.02	0.12 ± 0.06	0.22 ± 0.20	0.80 ± 0.28
OURS	0.95 ± 0.03	0.93 ± 0.01	0.99 ± 0.02	0.96 ± 0.07	0.78 ± 0.08	1.00 ± 0.00

the set of transferred state-only demonstrations $\hat{\mathcal{D}}_{\mathcal{M}_A^T}$, we can use any imitation from observation algorithm to learn the final policy. In this work, we follow the Behavioral Cloning from Observation (BCO) approach proposed in (Torabi et al., 2018). BCO entails learning an inverse dynamics model $\mathcal{I}_A : \mathcal{S}_A \times \mathcal{S}_A \rightarrow \mathcal{A}_A$ to infer missing action information. First, we collect a dataset of state-action triplets $\mathcal{P} = \{(s_A^t, a_A^t, s_A^{t+1})\}$ by random exploration. The inverse model is subsequently estimated by Maximum Likelihood Estimation (MLE) of the observed transitions in \mathcal{P} . Assuming a Gaussian distribution over actions, this reduces to minimizing an ℓ_2 loss as follows,

$$\min_{\mathcal{I}_A} \sum_{(s_A^t, a_A^t, s_A^{t+1}) \in \mathcal{P}} \|a_A^t - \mathcal{I}_A(s_A^t, s_A^{t+1})\|_2^2 \quad (11)$$

Next, the learnt inverse model is used to augment $\hat{\mathcal{D}}_{\mathcal{M}_A^T}$ with agent specific actions. Finally, these action augmented trajectories are used to learn the final policy π_A^T via behavioral cloning. Note that our correspondence learning framework is agnostic to the imitation from observation algorithm used for learning the agent policy.

5. Experiments

In this section, we analyze the efficacy of our proposed method on the xDIO task. We adopt MuJoCo (Todorov et al., 2012) as the experimental test-bed and evaluate on several cross-domain tasks, along with a thorough ablation study of different modules in our overall framework. Implementation details are presented in the supplementary materials. Code and videos are available at: <https://driptarc.github.io/xdio.html>.

5.1. Tasks

We use a total of 7 environments derived from the OpenAI Gym (Brockman et al., 2016): 2-link Reacher, 3-link Reacher, Friction-modified 2-link Reacher, Third-person 2-link Reacher, 4-legged Ant, 6-legged Ant and HalfCheetah. We use the joint level state-action space for all environments. These are used to construct six cross-domain tasks:

Dynamics-Reacher2Reacher (D-R2R): Agent domain is the 2-link Reacher and expert domain is the Friction-modified 2-link Reacher, created by doubling the friction

Algorithm 1 Learn domain transformation ψ

```

Input: Proxy task set  $\{(\mathcal{D}_{\mathcal{M}_E^T}, \mathcal{D}_{\mathcal{M}_A^T})\}_{j=1}^M$ , inference task trajectories  $\mathcal{D}_{\mathcal{M}_E^T}$ 
while not done do
  for  $j = 1, \dots, M$  do //Global and local alignment
    Sample  $(s_E, s'_E) \sim \mathcal{D}_{\mathcal{M}_E^T}, (s_A, s'_A) \sim \mathcal{D}_{\mathcal{M}_A^T}$  and store in buffers  $B_E^j, B_A^j$ 
    for  $i = 1, \dots, N$  do
      Sample mini-batch  $i$  from  $B_E^j, B_A^j$ 
      Update  $D_E^j, D_A^j$  by maximizing  $\mathcal{L}_{adv}^i(D_E^j)$  and  $\mathcal{L}_{adv}^j(D_A^j)$  respectively
      Update  $q^j$  by minimizing  $\mathcal{L}_{MI}^j$ 
      Update  $\psi, \phi$  by minimizing  $\lambda_1 (\mathcal{L}_{adv}^j(D_A^j) + \mathcal{L}_{adv}^i(D_E^j)) + \lambda_2 (\mathcal{L}_{cyc}^j + \mathcal{L}_z^j) + \lambda_3 \mathcal{L}_{pos}^j - \lambda_4 \mathcal{L}_{MI}^j$ 
    end for
  end for
  Sample  $(s_E, s'_E) \sim \mathcal{D}_{\mathcal{M}_E^T}$  and store in buffers  $B_E^{M+1}$  //Inference task adaptation
  for  $i = 1, \dots, N$  do
    Sample mini-batch  $i$  from  $B_E^{M+1}$ 
    Update  $V_z$  by minimizing  $\mathcal{L}_{pos.inf}$ 
    Update  $\psi, \phi$  by minimizing  $\mathcal{L}_{cyc.inf} + \mathcal{L}_{pos.inf}$ 
  end for
end while

```

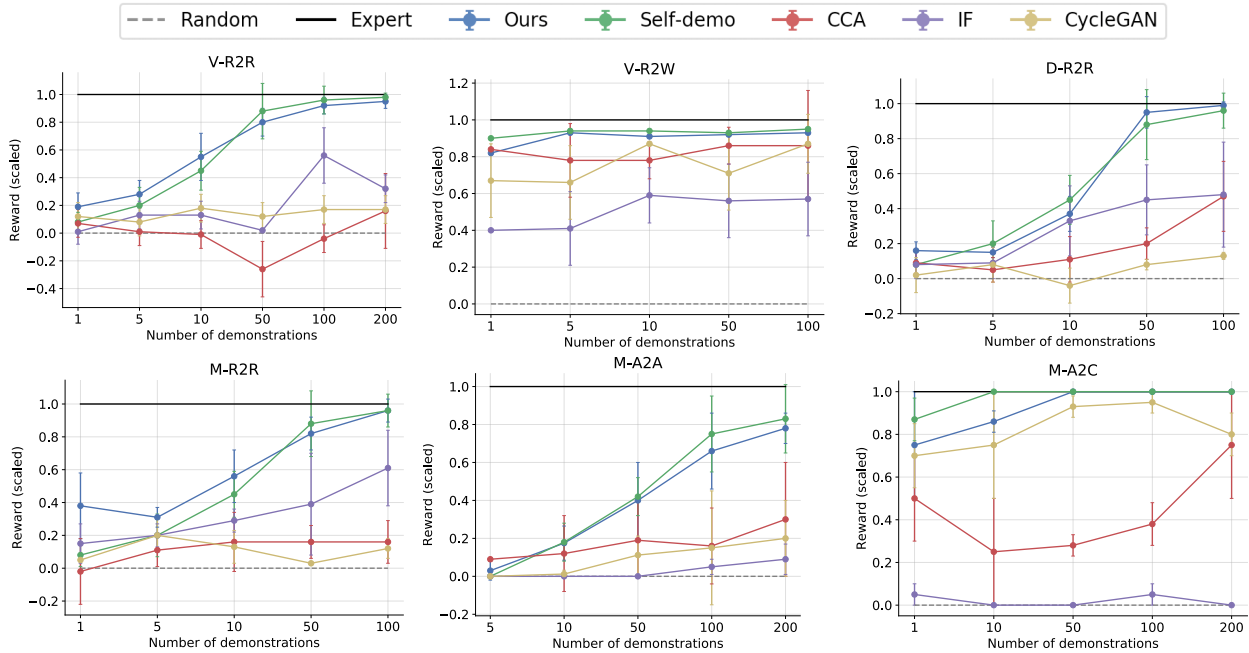


Figure 3. **Adaptation complexity.** Performance of learned policy as the number of cross-domain demonstrations is varied. Our framework consistently performs better than baselines and achieves results close to Self-demo.

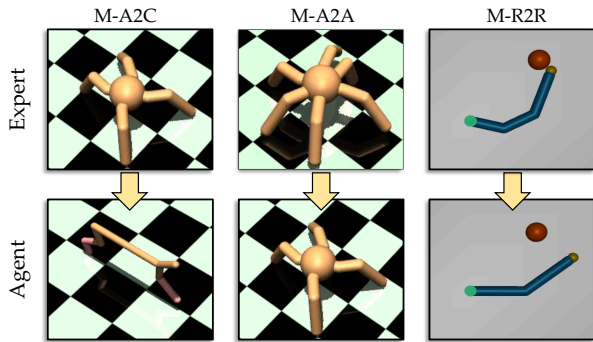


Figure 4. **Cross-domain tasks.** Different morphologically mismatched tasks used in our experiments.

co-efficient of the former. The proxy tasks are reaching for M goals and the inference tasks are reaching for 4 new goals, placed maximally far away from the proxy goals. See the supplementary for more details on goal placement.

Viewpoint-Reacher2Reacher (V-R2R): Agent domain is the 2-link Reacher and expert domain is Third-person 2-link Reacher that has a “third person” view state space with a 180° planar offset. Tasks are the same as D-R2R.

Viewpoint-Reacher2Writer (V-R2W): Agent domain is the 2-link Reacher and expert domain is Third-person 2-link Reacher. The proxy tasks are the same as D-R2R and the inference task is tracing a letter on a plane as fast as possible (Kim et al., 2020). The inference task differs from the proxy tasks in two key aspects: the end effector must draw a straight line from the letter’s vertex to vertex and not

slow down at the vertices.

Morphology-Reacher2Reacher (M-R2R): Agent domain is the 2-link Reacher, while expert domain is the 3-link Reacher. Otherwise same as D-R2R.

Morphology-Ant2Ant (M-A2A): Agent domain is the 4-legged Ant, while expert domain is the 6-legged Ant. Otherwise same as D-R2R.

Morphology-Ant2Cheetah (M-A2C): Agent domain is the HalfCheetah, while expert domain is the 4-legged Ant. Otherwise same as D-R2R.

5.2. Baselines

We compare our framework to other methods which are able to learn state correspondences from unpaired and unaligned demonstrations without access to expert actions - Canonical Correlation Analysis (Hotelling, 1992), Invariant Features (Gupta et al., 2017) and CycleGAN (Zhu et al., 2017). Canonical Correlation Analysis (CCA) (Hotelling, 1992) finds invertible linear transformations to a space where domain data are maximally correlated when given unpaired, unaligned demonstrations. Invariant Features (IF) learns state maps via a domain agnostic space from paired and aligned demonstrations - we use Dynamic Time Warping (Müller, 2007) on the learned latent space to compute the pairings from the unpaired data. CycleGAN learns the state correspondence via adversarial learning with an additional cycle-consistency on state reconstruction. For all the base-

Table 3. Ablation study on each module’s contribution to final policy performance.

METHOD	V-R2R	V-R2W	D-R2R	M-R2R	M-A2A	M-A2C
OURS	0.95 ± 0.05	0.93 ± 0.00	0.99 ± 0.02	0.96 ± 0.07	0.78 ± 0.08	1.00 ± 0.00
- W/O INFERENCE ADAPTATION	0.81 ± 0.11	0.88 ± 0.03	0.74 ± 0.22	0.78 ± 0.11	0.46 ± 0.12	0.78 ± 0.23
- W/O \mathcal{L}_{MI}	0.60 ± 0.30	0.92 ± 0.03	0.76 ± 0.30	0.67 ± 0.34	0.28 ± 0.20	0.80 ± 0.21
- W/O TEMPORAL PRESERVATION	0.64 ± 0.31	0.84 ± 0.00	0.70 ± 0.32	0.72 ± 0.32	0.36 ± 0.50	0.43 ± 0.50

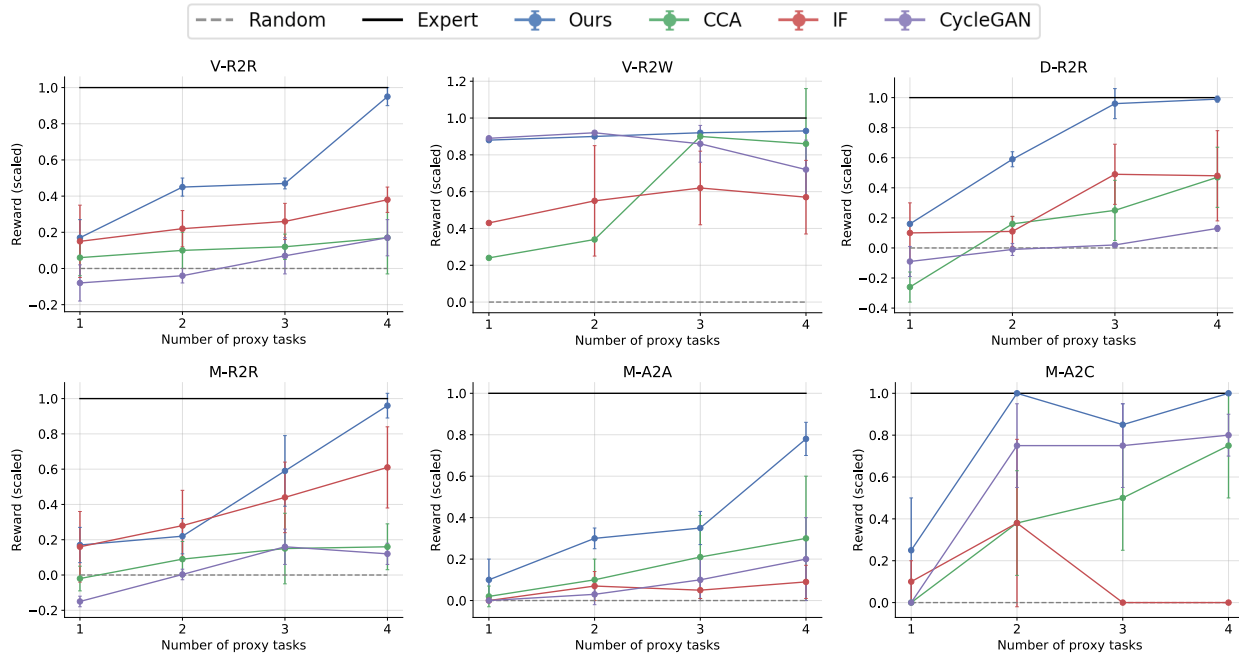


Figure 5. **Alignment Complexity.** Performance of learned policy as as the number of proxy tasks is varied. Notably, even with a reduced number of proxy tasks, our method outperforms the baselines in most cases.

lines, we follow a similar procedure towards learning the final policy - the correspondence is learnt through the proxy tasks and then is used to transfer trajectories for policy training via BCO. Reported results are averaged across 10 runs. Experts on Reacher tasks are trained using PPO (Schulman et al., 2017), while those for Ant/Cheetah are trained using A3C (Mnih et al., 2016).

5.3. Cross-domain imitation performance

We compare imitation policies learnt by our framework against those learnt using baselines in Table 2. As may be observed, the proposed method achieves near expert performance across all the cross-domain tasks encompassing viewpoint, dynamics and morphological mismatch. On the other hand, baselines consistently fail to generalize across the same tasks. There are two key reasons which can be hypothesized for this poor performance. Firstly, IF requires time-aligned trajectories, and the alignment when done by algorithms like DTW, rather than human intervention, may not be good enough given that our experiments involve diverse starting states, up to 1.5x differences in demonstra-

tion lengths (shown in Figure 6), and varying task execution rates. Secondly, baselines which learn from unpaired data (CCA and CycleGAN), also fail due to the lack of a mechanism to preserve MDP task characteristics, which is taken care of in our method via temporal order preservation and domain alignment. Figure 7 illustrates the learnt state-maps for some of the cross-domain tasks. The proposed framework translates the expert states in a manner that preserves task semantics.

Varying the number of demonstrations. Given an adequate set of proxy tasks, we experiment by varying the number of cross-domain demonstrations required for training the policy on the inference task. To serve as an upper-bound on performance, we imitate on agent domain demonstrations, drawn from an expert, on the inference task and denote this as the Self-demo baseline. As shown in Figure 3, our framework produces transferred demonstrations of equal effectiveness to the self-demonstrations. This clearly demonstrates the effectiveness of our framework.

Varying the number of proxy tasks. The number of proxy tasks play a vital role in learning the correspondence across

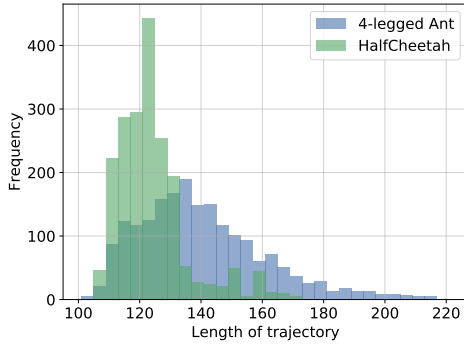


Figure 6. **Trajectory length distributions.** Distributions of trajectory lengths of the proxy tasks used for M-A2C.

the domains. We perform experiments by varying the number of proxy tasks in the alignment set needed to learn the state-map for imitation, given sufficient cross-domain demonstrations for the inference tasks. The results are shown in Figure 5. In general, more proxy tasks equate to better domain alignment as the solution space over possible state maps is constrained, and the learnt mapping generalizes better to the inference tasks.

5.4. Ablation study

We perform a set of ablation studies by removing each piece of the framework, demonstrating the importance of including each component. The results are shown in Table 3. We begin by excluding inference task adaptation. This leads to a small drop in performance across all tasks, reinforcing the need for adapting on the inference task to incorporate the new state distribution introduced by the inference task. Notably, even without adaptation, the performance in almost all the tasks exceeds those of the baselines. Removing the mutual information objective leads to a similar drop in performance across all tasks. Excluding temporal position preservation also reduces performance – demonstrating the significance of preserving task semantics via global alignment, which cycle-consistency alone fails to ensure.

6. Conclusion

In this paper, we present a novel framework to tackle the xDIO task by learning a state-map across domains using both local and global alignment. Local alignment is performed via transition distribution matching and cycle-consistency in both the state and latent space, while global alignment is enforced via the idea of temporal position preservation. While previous approaches rely on paired data and expert actions, we provide a general framework that can learn the mapping from unpaired, unaligned demonstrations without expert actions. We demonstrate the efficacy of our approach on multiple cross-domain tasks encompassing dynamics, viewpoint and morphological mismatch. Our future work will concentrate on extending our method for learning

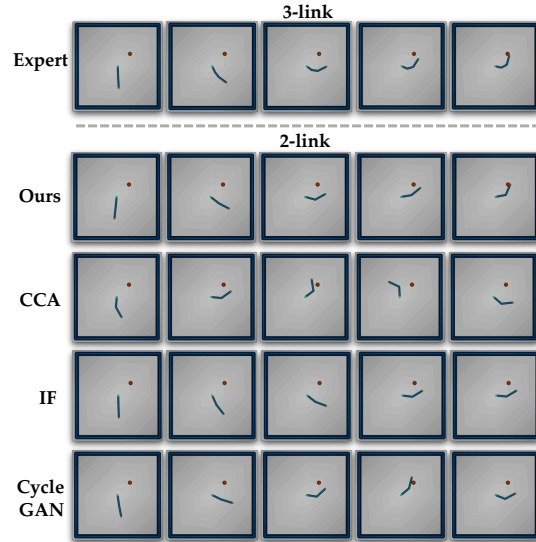


Figure 7. **Visualization of domain transformations.** State maps learned by our framework and the baselines on the M-R2R task. Our framework is able to map the end effector in a manner which preserves task semantics.

correspondence using random trajectories, thus mitigating the need for proxy tasks.

Acknowledgements

This work was partially supported by Mitsubishi Electric Research Labs and US National Institute of Food and Agriculture Award No: 2021-67022-33453 through the National Robotics Initiative.

References

Abbeel, P. and Ng, A. Y. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first International Conference on Machine learning*, pp. 1, 2004.

Alemi, A. A., Fischer, I., Dillon, J. V., and Murphy, K. Deep variational information bottleneck. *arXiv preprint arXiv:1612.00410*, 2016.

Ammar, H. B. and Taylor, M. E. Reinforcement learning transfer via common subspaces. In *International Workshop on Adaptive and Learning Agents*, pp. 21–36. Springer, 2011.

Ammar, H. B., Eaton, E., Taylor, M. E., Mocanu, D. C., Driessens, K., Weiss, G., and Tuyls, K. An automated measure of mdp similarity for transfer in reinforcement learning. 2014.

Ammar, H. B., Eaton, E., Ruvolo, P., and Taylor, M. Un-supervised cross-domain transfer in policy gradient rein-

- forcement learning via manifold alignment. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29, 2015.
- Bojarski, M., Del Testa, D., Dworakowski, D., Firner, B., Flepp, B., Goyal, P., Jackel, L. D., Monfort, M., Muller, U., Zhang, J., et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- Eysenbach, B., Gupta, A., Ibarz, J., and Levine, S. Diversity is all you need: Learning skills without a reward function. *arXiv preprint arXiv:1802.06070*, 2018.
- Ferns, N., Panangaden, P., and Precup, D. Bisimulation metrics for continuous markov decision processes. *SIAM Journal on Computing*, 40(6):1662–1714, 2011.
- Gamrian, S. and Goldberg, Y. Transfer learning for related reinforcement learning tasks via image-to-image translation. In *International Conference on Machine Learning*, pp. 2063–2072. PMLR, 2019.
- Gangwani, T. and Peng, J. State-only imitation with transition dynamics mismatch. In *International Conference on Learning Representations*, 2019.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. *Advances in neural information processing systems*, 27:2672–2680, 2014.
- Gupta, A., Devin, C., Liu, Y., Abbeel, P., and Levine, S. Learning invariant feature spaces to transfer skills with reinforcement learning. In *International Conference on Learning Representations*. OpenReview.net, 2017.
- Ho, D., Rao, K., Xu, Z., Jang, E., Khansari, M., and Bai, Y. Retinagan: An object-aware approach to sim-to-real transfer. *arXiv preprint arXiv:2011.03148*, 2020.
- Ho, J. and Ermon, S. Generative adversarial imitation learning. In *Advances in Neural Information Processing Systems*, pp. 4565–4573, 2016.
- Hoffman, J., Tzeng, E., Park, T., Zhu, J.-Y., Isola, P., Saenko, K., Efros, A., and Darrell, T. Cycada: Cycle-consistent adversarial domain adaptation. In *International conference on machine learning*, pp. 1989–1998. PMLR, 2018.
- Hotelling, H. Relations between two sets of variates. In *Breakthroughs in statistics*, pp. 162–190. Springer, 1992.
- Jones, S. S. The development of imitation in infancy. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1528):2325–2335, 2009.
- Kim, K., Gu, Y., Song, J., Zhao, S., and Ermon, S. Domain adaptive imitation learning. In *International Conference on Machine Learning*, pp. 5286–5295. PMLR, 2020.
- Liu, Y., Gupta, A., Abbeel, P., and Levine, S. Imitation from observation: Learning to imitate behaviors from raw video via context translation. In *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1118–1125. IEEE, 2018.
- Marshall, P. J. and Meltzoff, A. N. Body maps in the infant brain. *Trends in Cognitive Sciences*, 19(9):499–505, 2015.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., and Kavukcuoglu, K. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pp. 1928–1937. PMLR, 2016.
- Müller, M. Dynamic time warping. *Information retrieval for music and motion*, pp. 69–84, 2007.
- Paul, S., Vanbaars, J., and Roy-Chowdhury, A. Learning from trajectories via subgoal discovery. In *Advances in Neural Information Processing Systems*, pp. 8411–8421, 2019.
- Pomerleau, D. A. Alvin: An autonomous land vehicle in a neural network. In *Advances in Neural Information Processing Systems*, pp. 305–313, 1989.
- Poole, B., Ozair, S., Van Den Oord, A., Alemi, A., and Tucker, G. On variational bounds of mutual information. In *International Conference on Machine Learning*, pp. 5171–5180, 2019.
- Rao, K., Harris, C., Irpan, A., Levine, S., Ibarz, J., and Khansari, M. RI-cyclelegan: Reinforcement learning aware simulation-to-real. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11157–11166, 2020.
- Ravindran, B. *An algebraic approach to abstraction in reinforcement learning*. PhD thesis, University of Massachusetts at Amherst, 2004.
- Rizzolatti, G. and Craighero, L. The mirror-neuron system. *Annu. Rev. Neurosci.*, 27:169–192, 2004.
- Ross, S., Gordon, G., and Bagnell, D. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pp. 627–635, 2011.
- Schaal, S. Is imitation learning the route to humanoid robots? *Trends in cognitive sciences*, 3(6):233–242, 1999.

- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Sharma, P., Pathak, D., and Gupta, A. Third-person visual imitation learning via decoupled hierarchical controller. In *Advances in Neural Information Processing Systems*, pp. 2597–2607, 2019.
- Smith, L., Dhawan, N., Zhang, M., Abbeel, P., and Levine, S. Avid: Learning multi-stage tasks via pixel-level translation of human videos. *arXiv preprint arXiv:1912.04443*, 2019.
- Stadie, B. C., Abbeel, P., and Sutskever, I. Third person imitation learning. In *International Conference on Learning Representations*. OpenReview.net, 2017.
- Taylor, M. E. and Stone, P. Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research*, 10(7), 2009.
- Todorov, E., Erez, T., and Tassa, Y. Mujoco: A physics engine for model-based control. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–5033. IEEE, 2012.
- Torabi, F., Warnell, G., and Stone, P. Behavioral cloning from observation. *arXiv preprint arXiv:1805.01954*, 2018.
- Tzeng, E., Hoffman, J., Zhang, N., Saenko, K., and Darrell, T. Deep domain confusion: Maximizing for domain invariance. *arXiv preprint arXiv:1412.3474*, 2014.
- Umiltà, M., Intskirveli, I., Grammont, F., Rochat, M., Caruana, F., Jezzini, A., Gallese, V., Rizzolatti, G., et al. When pliers become fingers in the monkey motor system. *Proceedings of the National Academy of Sciences*, 105(6): 2209–2213, 2008.
- Wan, M., Gangwani, T., and Peng, J. Mutual information based knowledge transfer under state-action dimension mismatch. *arXiv preprint arXiv:2006.07041*, 2020.
- Yang, C., Ma, X., Huang, W., Sun, F., Liu, H., Huang, J., and Gan, C. Imitation learning from observations by minimizing inverse dynamics disagreement. In *Advances in Neural Information Processing Systems*, pp. 239–249, 2019.
- Zhu, J.-Y., Park, T., Isola, P., and Efros, A. A. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pp. 2223–2232, 2017.
- Ziebart, B. D., Maas, A. L., Bagnell, J. A., and Dey, A. K. Maximum entropy inverse reinforcement learning. In *AAAI*, volume 8, pp. 1433–1438. Chicago, IL, USA, 2008.