
Hierarchically Decoupled Imitation for Morphological Transfer

Donald J. Hejna III¹ Pieter Abbeel¹ Lerrel Pinto^{1,2}

Abstract

Learning long-range behaviors on complex high-dimensional agents is a fundamental problem in robot learning. For such tasks, we argue that transferring learned information from a morphologically simpler agent can massively improve the sample efficiency of a more complex one. To this end, we propose a hierarchical decoupling of policies into two parts: an independently learned low-level policy and a transferable high-level policy. To remedy poor transfer performance due to mismatch in morphologies, we contribute two key ideas. First, we show that incentivizing a complex agent’s low-level to imitate a simpler agent’s low-level significantly improves zero-shot high-level transfer. Second, we show that KL-regularized training of the high level stabilizes learning and prevents mode-collapse. Finally, on a suite of publicly released navigation and manipulation environments, we demonstrate the applicability of hierarchical transfer on long-range tasks across morphologies. Our code and videos can be found at <https://sites.google.com/berkeley.edu/morphology-transfer>.

1. Introduction

How should one use Reinforcement Learning (RL) to train their new four-legged walking robot? Training a robot from scratch is often infeasible as current RL algorithms typically require millions (Schulman et al., 2017; Haarnoja et al., 2018) to billions (Baker et al., 2019) of samples. Moreover, robots are expensive and slow, which further limits the applicability of learning from scratch. Because of this, tackling the challenge of high sample complexity has received significant interest and prompted a wide array of solutions

¹Department of EECS, University of California, Berkeley

²Computer Science, New York University. Correspondence to: Donald J. Hejna III <jhejna@berkeley.edu>.




Figure 1. In this work, we focus on the problem of transferring long-horizon policies across morphologies. Hence, given an already performant simple agent, we imitate its behaviors on a more complex agent for transfer through hierarchical decoupling.

ranging from off-policy learning (Lillicrap et al., 2015) to temporal abstractions in learning (Kulkarni et al., 2016).

A primary reason behind RL’s sample inefficiency is the lack of prior knowledge used in training new policies. One of the biggest takeaways from advances in computer vision (CV) and natural language processing (NLP) is that priors (both architectural and parametric) are invaluable. Hence this begs the question: why should an agent learn a task from scratch? Why not use strong priors? Recent works have shown the promise of semantic priors such as auxiliary losses (Jaderberg et al., 2016), and architectural priors such as transfer learning networks (Rusu et al., 2016). However unlike passive domains such as object detection, problems in motor control and robotics offer another strong prior: morphology. Instead of learning policies from scratch on a given agent, we can use policies previously learned on morphologically different agents as a prior. For example, instead of training a quadruped walking policy to solve a maze from scratch, requiring millions of expensive samples, we could transfer-learn from a much simpler robot, say a Roomba wheeled robot. This morphological transfer affords two benefits: first, learning a policy first on a simple morphology and then transferring to a harder one induces a natural curriculum for learning (Bengio et al., 2009) which provides richer rewards and learning signal; second, this allows us to use fewer samples from complex robots that are often expensive and time-consuming.

But how do we effectively transfer policies across morpholo-

gies? Direct policy transfer is infeasible, since different morphologies have different state and action spaces (see Figure 1 for examples of morphologies). Another option would be to use morphological latent embeddings (Chen et al., 2018; Pathak et al., 2019), but learning robust embeddings requires training across hundreds of morphologies. Instead, inspired by recent advances in hierarchical learning (Kulkarni et al., 2016; Nachum et al., 2018b), we propose a transfer learning framework using hierarchically decoupled policies. In this framework, the low-level policy is trained specific to a given morphology, while the high-level policy can be re-used across morphologies. For compatible transfer, only the high-level policy, operating on a global agent state is transferred, while the low-level policy is independently learned.

However, as recent work by Nachum et al. (2018a) has shown, high-level policies are intricately tied with low-level policies. Intuitively, if a low-level policy isn’t able to reach a specific sub-goal requested by the high-level policy, the high-level policy will not select that sub-goal. This brings a significant challenge to morphological transfer for agents. Since the low-levels of two agents might be significantly different due to differences in morphology that afford varying low-level capabilities, a zero-shot transfer of the high-level policy might not always be successful. To counter this, we propose a top-down alignment of the low-level policies. This is done by introducing an information theoretic alignment loss that minimizes the mutual information between morphologies and low-level behaviors. This objective is practically optimized using discriminative learning, which allows the low-level policy of a more complex agent to imitate the behavior of the simpler agent’s, making high-level transfer more successful.

Using this low-level alignment significantly improves the transfer of high-level policies. However, even with better alignment, zero-shot transfer of high-level policies will not be able to fully utilize the additional benefits of a complex morphology. One way to improve on this is to finetune (Girshick et al., 2014) the high-level policies after the zero-shot transfer. But, in RL straightforward finetuning suffers from catastrophic forgetting (Rusu et al., 2016) of the simpler agent’s high-level under the changing dynamics of the new agent. To prevent this, we take inspiration from prior work in transfer learning and introduce a KL-regularizer that allows the complex agent to improve performance while staying close to the simpler agent’s high-level. Intuitively, this balances the imitation of the simpler agent’s high-level with its own ability to solve the task. Doing this significantly improves performance on a suite of navigation and manipulation transfer tasks.

In summary, we present the following contributions in this work: (a) we show how hierarchical policies can help morphological transfer. Although recent works in Hierarchical

Reinforcement Learning (Peng et al., 2017; 2019) allude to the potential of morphological transfer, to our knowledge we are the first work that concretely focuses on this problem. (b) We propose two key technical insights for hierarchical imitation, a top-down low-level alignment and a KL-regularized high-level objective to accelerate transfer. (c) Finally, we empirically demonstrate significant improvements in performance of morphological transfer on long-horizon navigation and manipulation tasks.

2. Background

Before we describe our framework, we first discuss relevant background on RL. For an in-depth survey, we refer the reader to (Sutton et al., 1998; Kaelbling et al., 1996).

2.1. Reinforcement Learning

In our continuous-control RL setting, an agent receives a state observation $s_t \in \mathcal{S}$ from the environment and applies an action $a_t \in \mathcal{A}$ according to policy π . With stochastic policies, as in our case, we have $a_t \sim \pi(s_t)$. The environment returns a reward for every action r_t . The goal of the agent is to maximize expected cumulative discounted reward $\mathbb{E}_{s_0:T, a_0:T-1, r_0:T-1} \left[\sum_{t=0}^{T-1} \gamma^t r_t \right]$ for discount factor γ and horizon length T . On-policy RL (Schulman et al., 2015; Kakade, 2002; Williams, 1992) optimizes π by iterating between data collection and policy updates. It hence requires new on-policy data every iteration, which is expensive to obtain. On the other hand, off-policy reinforcement learning retains past experiences in a replay buffer and is able to re-use past samples. Thus, in practice, off-policy algorithms have been found to achieve better sample efficiency (Lillicrap et al., 2015). For our experiments we use the off-policy SAC (Haarnoja et al., 2018) algorithm as our base RL optimizer due to its sample efficiency. However, our framework is compatible with any standard RL algorithm.

2.2. Two-layered Hierarchical RL (HRL)

The central idea of HRL is to abstract the policy π into multiple policies that operate at temporally different levels. The most common abstraction is a two-level hierarchy (Nachum et al., 2018a), a low-level policy π^{lo} and a high-level policy π^{hi} . In our formulation, the high-level takes as input a part of the state observation s_t^{hi} and outputs morphology-independent high-level actions a_t^{hi} that serve as subgoals for the low-level policy. These subgoals are morphology-independent and belong to a goal space \mathbb{G} . The low-level takes as input a part of the state observation s_t^{lo} and the commanded subgoal a_t^{hi} , and outputs low-level control actions a_t^{lo} to try and reach the subgoal. Note that we split the state observation s_t into two components, s_t^{hi} a global morphology-independent state and s_t^{lo} a proprioceptive mor-




Figure 2. We transfer Hierarchical policies across morphologies through decoupled imitation. For the low-level policy we use density based discriminative imitation detailed in Section 3.2, which improves zero-shot high-level transfer. For the high-level policy, we use KL-regularized imitation detailed in Section 3.3, which improves finetuning the high-level.

phology dependent state. This particular choice of state, inspired from Marino et al. (2018), allows us to transfer π^{hi} across morphologies that have different state representations.

3. Method

In this work, we focus on the problem of transferring policies from one morphology to another for challenging long-horizon problems. Concretely, we define morphology transfer from morphology A to morphology B on a specific task as transferring a policy trained with A to B , i.e., given access to a trained π_A , what is the fastest way to train π_B (See Figure 2). Practically, A is a simple agent for which training π_A might be a lot easier (and safer) than training on a more complex agent π_B . However, since A and B have different state and action spaces afforded by their morphologies, we build on top of the hierarchical setup described in Section 2.2. The hierarchical demarcation between low-level policies that act on proprioceptive states and high-level policies that act on global states provides a natural way to transfer high-level knowledge. Moreover, hierarchical learning is empirically known to provide massive improvements in sample-efficiency (Nachum et al., 2018a). In the following subsections we detail our proposed technique.

3.1. Zero-shot High-level Transfer

In the hierarchical setting, morphology transfer reduces to transferring $\pi_A \equiv [\pi_A^{lo}, \pi_A^{hi}]$ trained on A to $\pi_B \equiv [\pi_B^{lo}, \pi_B^{hi}]$. One straightforward way to perform transfer is

to set $\pi_A^{hi} \rightarrow \pi_B^{hi}$ as the input s^{hi} and output a^{hi} of the high-level policies are morphology-independent. The low-level policy π_B^{lo} can either be learned with the fixed high-level π_B^{hi} or trained independently. In our experiments, we pre-train the low-level policy π_B^{lo} on uniformly sampled goals from \mathbb{G} , allowing us to learn an effective π_B^{lo} without access to π_A and generalize over tasks without re-training π_B^{lo} .

3.2. Low-level Imitation Through Behavior Alignment

Although directly transferring the high-level policy $\pi_A^{hi} \rightarrow \pi_B^{hi}$ with independently trained low-levels π_B^{lo} allows for zero-shot morphology transfer, it suffers from low-level domain shift especially for tasks requiring precise control. This is because different morphologies afford different low-level behavior. Intuitively, if π_B^{lo} doesn't generate similar behavior to π_A^{lo} , transferring $\pi_A^{hi} \rightarrow \pi_B^{hi}$ may not work since π_B^{lo} does not generate the behavior expected by π_A^{hi} . To solve this problem, we align the low-level π_B^{lo} to π_A^{lo} on the set of goals \mathbb{G} . Note that direct cloning is not possible since the low-level state and action spaces are not the same. Additionally, simply ensuring both agents can reach the same portions of \mathbb{G} used by π_A^{hi} is insufficient for strong alignment since goal-reaching behavior can differ while achieving the same goal (illustrated as low-level behavior in Figure 2). Even in the unlikely scenario where π_A^{lo} and π_B^{lo} can reach the same subset of \mathbb{G} within a single high-level step, different inter-step trajectories can still affect non-static environments.

To incentivize π_B^{lo} (parameterized by θ_B^{lo}) to mimic the behavior of π_A^{lo} , we propose minimizing the following mutual

information objective:

$$\min_{\theta_B^{lo}} I(\text{morphology}; \text{behavior}) \quad (1)$$

Minimizing the mutual information between the morphology type ($M = \{A, B\}$) and the generated low-level behavior will result in a low-level policy π_B^{lo} whose behavior cannot be distinguished from π_A^{lo} . To ensure that the behavior is compatible with both morphologies A and B , we set behavior at time t as $\tau_t = s_{t-k:t}^{hi}$, where k is the horizon of behavior. Let \mathcal{T}_t denote the distribution over behaviors τ_t .

As $I(M; \mathcal{T}_t) = H(M) - H(M|\mathcal{T}_t)$ and $H(M)$ cannot be controlled through π_B^{lo} , our objective from equation 1 reduces to maximizing $H(M|\mathcal{T}_t)$. Here $H(\cdot)$ denotes entropy. Since the probability of agent selection $P(M = m|\mathcal{T}_t = \tau_t)$ cannot be readily estimated, we compute the variational lower bound, which reduces our objective to:

$$\max_{\theta_B^{lo}} \mathbb{E}_{m \sim M, \tau_t \sim \mathcal{T}_t} [-\log q_\phi(m|\tau_t)] \quad (2)$$

Here q parameterized by ϕ is effectively a binary classifier (or discriminator) that outputs the probability of the generated behavior coming from morphology m given input behavior τ_t . Complete derivations of equation 2 can be found in Appendix A. Maximizing this objective through θ_B^{lo} implies generating behaviors that maximally confuse the discriminator. To do this we augment the low-level policy rewards as follows:

$$r_B^{lo} \leftarrow R(\tau_t^B|g) - \lambda_0^{\lambda_1} \log q_\phi(M = B|\tau_t^B) \quad (3)$$

Here λ_0 and λ_1 represent temperature parameters that anneals the rewards from the discriminator over time, τ_t^B represents the trajectory generated by π_B^{lo} while trying to solve the subgoal g , and R represents the sub-goal reward function. The training process is summarized in Algorithm 1. This procedure of fitting a discriminator and then re-optimizing the policy has a similar flavor to recent work in approximate inverse reinforcement learning (Ho & Ermon, 2016), where the discriminator represents the reward density function.

3.3. High-level Imitation Through KL-regularized Training

In the previous section we discussed aligning the low-levels through density based imitation, which allows for better zero-shot transfer of the high-levels $\pi_A^{hi} \rightarrow \pi_B^{hi}$. However, even with low-level alignment, if the morphologies afford different abilities, direct transfer of the high-level may not reach optimal behavior. One way to remedy this is to finetune the high-level π_B^{hi} after transferring from π_A^{hi} , which already encodes knowledge required to solve the task. However, direct finetuning with RL often leads to catastrophic forgetting of the former policy π_A^{hi} as previously noted in

Algorithm 1 Low-level Alignment for high-level transfer

Input: New agent B , agent A 's behavior \mathcal{T}_A , and a goal distribution \mathbb{G} .

Initialize: Learnable parameters θ for $\pi_{B,\theta}^{lo}$ and ϕ for q_ϕ

for $i=1,2,\dots,N_{\text{iter}}$ **do**

 sample goal $g \sim \mathbb{G}$

for $j=1,2,\dots,T$ **do**

 sample action $a_t^{lo} \sim \pi_{B,\theta}^{lo}$

 collect experience $(s_t^{lo}, a_t^{lo}, s_{t+1}^{lo}, r_t^{lo})$ for \mathcal{T}_B

end for

for $j=1,2,\dots,M_{\text{policy}}$ **do**

 update r_t^{lo} according to eq 3

$\theta \leftarrow \text{policyOptimizer}(\{(s_t^{lo}, a_t^{lo}, s_{t+1}^{lo}, r_t^{lo})\}, \theta)$

end for

for $j=1,2,\dots,M_{\text{discrim}}$ **do**

$\phi \leftarrow \text{discrimOptimizer}(\mathcal{T}_A, \mathcal{T}_B, \phi)$

end for

end for

Return: θ

Rajeswaran et al. (2017); Rusu et al. (2016). To alleviate this, we propose a KL-regularized finetuning that balances staying close to π_A^{hi} and exploiting task driven reward signals.

$$\text{grad}_B^{hi} \leftarrow \text{grad}_{\text{RL}} + \alpha \nabla_{\theta_B^{hi}} \text{KL}(\pi_B^{hi}(a_B^{hi}|s_B^{hi}) || \pi_A^{hi}(a_A^{hi}|s_B^{hi})) \quad (4)$$

Here, under the high-level trajectory $\tau^{hi} \equiv (s^{hi}, a^{hi})$ generated by π_B^{hi} , the gradients for the parameters θ_B^{hi} of π_B^{hi} are represented by grad_B^{hi} . grad_{RL} represents the gradients through the base RL optimizer. While the first term of equation 4 gives the RL gradient, the second term represents the behavior cloning gradient of imitating π_A^{hi} , with $a_A^{hi} \sim \pi_A^{hi}$. This additional imitation objective is a practical form of regularization that penalizes deviations from sub-goal sequences set by the existing high-level π_A^{hi} . Additionally, the imitation loss forces the policy to remain near the portion of the state space in which π_A^{hi} was maximally performant, preventing the aforementioned phenomena of catastrophic forgetting.

4. Experiments

In this section we discuss empirical results on using hierarchically decoupled imitation for transferring policies across morphologies. Note that since there are no standard benchmarks for evaluating morphological transfer, we create an environment suite that will be publicly released. Using this experimental setup, we seek to answer the following questions. First, does hierarchical decoupling provide an effective natural framework for morphology transfer? Second, does discriminative imitation improve zero-shot performance of morphology transfer? Finally, does KL-




Figure 3. Visualization of all agents and tasks used in our environment suite.

regularized fine-tuning improve the performance of transferred policies over baseline methods?

4.1. Experimental Setup

To study morphological transfer for long-horizon tasks, we present a suite of environments with eight agents and four tasks for manipulation and navigation as illustrated in Figure 3. For navigation, we use four agents: *PointMass*, *Ant* (Duan et al., 2016), *3-Legged Ant*, *Quadruped* (Tassa et al., 2018), and two tasks: *Waypoint* navigation and *Maze* navigation. Note that unlike in Nachum et al. (2018a), our maze task offers only a sparse reward. During high-level training we sample goals uniformly across the six segments of the maze, but in evaluation we consider both *Maze Sample*, where goals are sampled as in training and *Maze End*, where the goal is always set to the end of the maze. All morphologies have vastly different action spaces. The simplest agent, the *PointMass*, has an action space of two dimensions while the most complex agent, the *Quadruped*, has an action space of twelve dimensions. Across navigation tasks, the goal space is set to the (x, y) of the agent’s torso. For the manipulation environments, we use four agents with varying degrees of freedom: *PointMass*, *2Link-Arm*, *3Link-Arm*, and *4Link-Arm* adapted from the OpenAI Gym Reacher (Brockman et al., 2016) and two tasks: *BlockPush* (with variants based on goal location) and *PegInsert*. The goal space \mathbb{G} for the manipulation tasks is the (x, y) position of the end-effector. All of these environments are simulated in MuJoCo (Todorov et al., 2012) using the OpenAI Gym interface. Images of the environments can be found in Figure 3. More details of the environments and the agents are provided in Appendix C.

4.2. Training Details

We use SAC as our base RL optimizer, and re-purpose the open-source Stable Baselines (Hill et al., 2018) codebase for our methods. We first train low-level policies by sampling uniformly from the goal space of each task \mathbb{G} . In order to make our hierarchical framework applicable to any task or environment, we use reward functions for the low-level policies that are highly generalizable. For the

navigation tasks, low-level reward is given by the weighted cosine between the vector of agent movement and vector to the goal g . In all manipulation tasks, we use $L2$ distance between the end-effector and goal g with a sparse reward for being within ϵ of g . Across all tasks of the same type, we only train low-level policies for each agent once; i.e. the same low-level policies used for *Waypoint* navigation are used for the *Maze* task. We ran all experiments across five random seeds, except for the high-level and high-level finetunings of the selected navigation tasks, where we ran ten random seeds.

When training policies for low-level imitation we collect goal-conditioned transition data in the form $(g - s_t^{hi}, g - s_{t+1}^{hi})$ from both agents in an off-policy manner. Data from agent A is collected by running π_A^{lo} on uniformly sampled goals from \mathbb{G} , however one could imagine re-using previous data generated while learning π_A^{lo} . Data from agent B is reused from the RL algorithm. For every step of agent B during optimization, we add its transition along with a transition from agent A to a circular buffer that maintains class balance. While training, we only update the discriminator periodically by randomly sampling data from the buffer. In addition to annealing the weight of discriminator rewards as per equation 3, we also anneal the learning rate of the discriminator to zero, preventing over-fitting and allowing agent B to train against an increasingly stationary target.

For KL-regularized finetuning, we directly incorporate a term for the KL-divergence between π_A^{hi} and π_B^{hi} into the policy optimizer loss. This is easily calculated as all of our policies are parameterized by diagonal gaussians. Just as with the discriminator, the KL-loss coefficient is annealed to zero during training as it is no longer needed once the transferred policy is stable. Additional training details including hyper-parameters set are included in Appendix D.

4.3. How well does zero-shot hierarchical transfer work?

We perform straightforward high-level transfer described in Section 3.1 across our task-morphology environment suite. In Table 1, we present results from combining the high-level of an agent (column-wise) with a specific agent (row-wise).

Hierarchically Decoupled Imitation for Morphological Transfer

Waypoint	PM HL	Ant HL	Ant3 HL	Quad HL	Maze E	PM HL	Ant HL	Maze S	PM HL	Ant HL
PM LL	1e3 ± 43	603 ± 34	716 ± 58	577 ± 70	PM LL	.80 ± .18	.19 ± .17	PM LL	.96 ± .04	.56 ± .04
Ant LL	483 ± 39	476 ± 19	473 ± 51	407 ± 36	Ant LL	.30 ± .08	.16 ± .14	Ant LL	.62 ± .04	.50 ± .08
Ant3 LL	489 ± 74	484 ± 72	499 ± 65	432 ± 65	Ant3 LL	.58 ± .13	.12 ± .10	Ant3 LL	.84 ± .04	.56 ± .05
Quad LL	169 ± 33	182 ± 22	220 ± 27	257 ± 19	Quad LL	.40 ± .21	.13 ± .15	Quad LL	.78 ± .08	.58 ± .04
Block Push 1	PM HL	2Link HL	3Link HL	4Link HL	Insert	PM HL	3Link HL	4Link HL		
PM LL	.99 ± .01	.21 ± .15	.33 ± .18	.26 ± .17	PM LL	1.0 ± .00	.60 ± .22	.60 ± .22		
3Link LL	.17 ± .08	.20 ± .09	.96 ± .02	.49 ± .14	3Link LL	1.0 ± .00	1.0 ± .00	.80 ± .18		
4Link LL	.24 ± .11	.07 ± .04	.10 ± .05	.89 ± .04	4Link LL	1.0 ± .00	.80 ± .18	.90 ± .06		

Table 1. Selected Zero-Shot performance results averaged across a hundred episodes per run with navigation tasks on top and manipulation tasks below. For all tasks except waypoint navigation, values reported indicate the fraction of successes. “Maze E” refers to the *Maze End* evaluation and “Maze S” refers to the *Maze Sample* evaluation.

On relatively easy tasks like *Waypoint*, zero-shot transfer works well. For instance, on the *Ant* morphology, using a high-level learned on *PointMass* achieves the same performance within confidence bounds. On the sparse *Maze* task, transferring from *PointMass* to *Ant* does better than learning the *Ant* policy from scratch. Though initially surprising, this result can be explained by the *PointMass*’s superior exploration—the *PointMass* is speedy allowing it to easily discover sparse reward signals, while the *Ant* is slow and topples over when it runs into walls. This is a concrete example of how a simple agent can provide a valuable learning curriculum to a complex one. However, the 30% *MazeEnd* success rate after zero-shot transfer still leaves much to be desired in comparison to the *PointMass*’s 80% success rate. Interestingly, even on the simple *Waypoint* task the nearly ideal *PointMass* experiences a significant performance degradation when using a different high-level policy, indicating that high-level policies are indeed overfit to the morphology they are trained on. This is manifested further in the *BlockPush* task where zero-shot performance deteriorates significantly. For example, when transferring the high-level form the *PointMass* to the *4Link Arm*, performance drops by around 74%. The poor high-level transfer on harder environments and morphologies motivates the need for better transfer algorithms. Additional results are presented in Appendix E.

4.4. Does discriminative imitation improve transfer?

To improve zero shot high-level transfer, we perform discriminative imitation on the low-level policies as described in Section 3.2. Results are presented in table 2 across a wide set of poorly performing plain high-level transfers. We measure the effectiveness of the discriminator by comparing zero-shot performance of the plain low-level and the low-level trained with a discriminator for a given high-level. Across nearly all tasks, discriminative imitation of the low-level improves zero-shot performance, especially in manipulation tasks. For example, transferring the *PointMass* high-level to the *3Link Arm* is twice as successful across




Figure 4. Learning curves for low-level policies with and without the discriminator. For the Ant Agent (left) the discriminator policy achieves slightly lower though comparable performance to the regular low-level, while the 3 and 4 Link arms (right) learn slightly faster with discriminator supervision.

both block push tasks using our method. The *Quadruped*’s performance in the *Waypoint* task with discriminative imitation was better than training the high-level from scratch for the limited number of low-level steps we used in training. A substantial benefit of discriminative imitation is that the aligned low-level policy only needs to be trained once for transfer across any number of high-level policies, making it particularly useful when training the high-level policy is expensive or an increase in performance is desired across a large set of tasks. Besides the *2Link* to *3Link* transfer, the same imitated low level boosts performance on all tasks.

To verify that our agents are indeed learning to discriminate rather than just learning objectively better policies, we examine the low-level learning curves in Figure 4. Though for the arm agents the discriminator provides an extra form of supervision that speeds up learning, final rewards for agents with the discriminator are near those of agents without, indicating that the performance benefits in Table 2 can indeed be attributed to imitation.

4.5. Does finetuning improve transfer?

After transferring the high-level policy from a simpler agent to a more complex one, we finetune it by retraining to improve performance. Results for this are visualized in

Hierarchically Decoupled Imitation for Morphological Transfer

Task	Waypoint			Maze End		Maze Sampled		Insert
Transfer	PM→Ant	PM→Ant3	PM→Quad	PM→Ant	PM→Quad	PM→Ant	PM→Quad	3Link→4Link
Zero-Shot	483 ± 39	489 ± 74	169 ± 33	.30 ± .08	.40 ± .22	.62 ± .04	.78 ± .08	.80 ± .18
Discrim (ours)	546 ± 33	495 ± 38	338 ± 39	.55 ± .13	.40 ± .21	.72 ± .03	.89 ± .06	.93 ± 0.05

Task	Block Push 1				Block Push 2			
Transfer	PM→3Link	2Link→3Link	2Link→4Link	3Link→4Link	PM→3Link	2Link→3Link	2Link→4Link	3Link→4Link
Zero-Shot	.17 ± .08	.20 ± .09	0.07 ± .04	.10 ± .05	.24 ± .12	.61 ± .16	.19 ± .12	.39 ± .16
Discrim (ours)	.34 ± .10	.43 ± .15	.17 ± .08	.23 ± .13	.43 ± .12	.42 ± .11	.46 ± .30	.44 ± .14

Table 2. Selected discriminator results. We train low-level policies with a discriminator for the transferred-to agents, then assess zero-shot performance across a hundred episodes for each run. We find that the same low-level aligned policy improves transfer across both tasks, commonly leading to a doubling of performance in the *Block Push* tasks.

Figure 6 and Figure 7 where we examine returns when using extra samples collected from agent *B*. In most cases, finetuning works better than learning the high-level from scratch (in green) and reaches substantially higher performance compared to the zero-shot high-level transfer (dotted purple line). However, in some cases regular finetuning suffers from unstable training. For instance in the transfer finetuning of *2Link Arm* from *PointMass* on *BlockPush*, we see poor confidence bounds. Empirically, we notice massive fluctuations in training performance across seeds that cause this behavior, most likely explained by high-level policies shifting far out of distribution during morphology transfer, causing catastrophic forgetting of task-solving knowledge. This motivates the need for better transfer learning methods for finetuning.

4.6. How much does KL-regularized training help?

To improve performance during finetuning, we use KL-regularized imitation as described in Section 3.3. Empirically, we find that the addition of an imitation loss during high-level transfer substantially improves performance as seen in Figure 6 and Figure 7. On the *BlockPush* we notice at least a 2X speedup compared to already strong direct finetuning method. On the *Maze* environments, we again notice improvements in performance; however the gains are modest compared to the *BlockPush* environment. We additionally compare our KL-regularized finetuning method to training the high-level policy from scratch with a pre-trained low-level and learning the task without hierarchy, denoted “Full”. Note that we make our comparisons with respect to the number of samples used in training. Learning without hierarchy was unsuccessful for the complex agents on navigation tasks, particularly the sparse-reward maze. Though it is cut off in the graphs, learning without hierarchy was eventually successful for all the manipulation tasks.

4.7. In what cases does hierarchical transfer fail?

The aforementioned results demonstrate the significant promise of hierarchically decoupled transfer in scenarios where all agents can similarly cover the task state space.




Figure 5. Trajectories for the *Ant* agent on a “steps” environment for training the high-level from scratch, finetuning from *Point Mass*, and zero-shot transfer from *Point Mass* from left to right respectively. The dark blue (taller) and light blue (shorter) steps are reduced in height to allow the ant agent to climb over.

Zero-Shot	PM HL	Ant HL	Finetune	.53 ± .09
PM LL	.60 ± .22	.00 ± .00	Finetune KL	.52 ± .13
Ant LL	.32 ± .16	.59 ± .05	HL Scratch	.59 ± .05

Table 3. Comparison of transfer performance on the step task. We evaluate only on the goal at the end of the maze as depicted in Figure 5

However, what happens if a more complex agent’s morphology endows it with additional abilities? In this specific context, hierarchical decoupling may not lead to a perfectly optimal transfer as the more complex agent may be able to reach new states other agents could not. To examine such scenarios, we design a simple “step maze” task, with low barriers that an *Ant* could scale but a *Point Mass* could not. We then examine the trajectories of the *Ant* when its high-level is trained from scratch, finetuned from *Point Mass*, and zero-shot transferred from *Point Mass* in Figure 5. The *Ant* high-level trained from scratch learns to climb over the highest barrier while the zero-shot trajectory mostly avoids the barriers just as the *Point Mass* would. By finetuning, the *Ant* agent moves closer to the policy learned from scratch climbing the lower barrier, but remains somewhat tied to its prior and less consistently climbs the high barrier. Empirical results are included in Table 3, where we see that for a limited number of samples from *Ant*, training from scratch achieves the highest performance unlike in previous environments. By training from scratch, the *Ant* reaches the




Figure 6. Comparison of performance finetuning from *PointMass* for *Ant Waypoint*, *Ant Maze End*, *Ant Maze Sample*, and *Quadrupe Maze Sample* from left to right. For *Waypoint*, the agent receives a reward of 100 per waypoint reached and for *Maze*, the agent receives a reward of 1000 for reaching the set goal. “Full” denotes that the policy was trained without hierarchy. For *Maze End*, we finetune with the goal always set to the end of the maze.




Figure 7. Learning curves of finetuning for 2 *Link* from *PointMass*, 3 *Link* from *PointMass*, 4 *Link* from *PointMass*, and 4 *Link* from 3 *Link* from left to right on the block push 1 task. The agent receives a reward of 200 for completing the task successfully.

final goal 59% of the time while zero-shot and finetuning are only successful 32% and 53% of the time respectively. Furthermore, KL-regularized finetuning does not seem to help here, indicating that high-level imitation is only useful when the optimal solution for a task is similar between agents.

5. Related Work

Our work is inspired by and builds on top of a broad range of topics, including transfer learning, morphological transfer, imitation learning, and information theoretic RL. In this section, we overview the most relevant ones.

5.1. Multi-task and Transfer Learning

Learning models that can share information across tasks has been concretely studied in the context of multi-task learning (Caruana, 1997), where models for multiple tasks are simultaneously learned. More recently, Kokkinos (2017); Doersch & Zisserman (2017) looks at shared learning across visual tasks, Pinto & Gupta (2017) looks at shared learning across robotic tasks, and Pathak et al. (2019) looks at message passing for low-level control of articulated agents. In contrast to multi-task learning where knowledge is simultaneously learned, we focus on the disparate setting in which knowledge from one task (a simpler agent) is transferred to another (a more complex agent). Murali et al. (2018) learn

manipulation policies via curriculum over joints, but focus only on a single agent. More relevant to our work, Chen et al. (2018) focuses on transfer across different hardware by training hardware conditioned embeddings across a large number of morphologies. However, access to more than two or three morphologies is unrealistic in practice. Our method is best suited for these scenarios.

Transfer learning (Pan & Yang, 2009; Torrey & Shavlik, 2010) focuses on transferring knowledge from one domain to another. One of the simplest forms of transfer is finetuning (Girshick et al., 2014) using models initialized on different tasks. Several other works look at more complex forms of transfer (Yang et al., 2007; Hoffman et al., 2014; Aytar & Zisserman, 2011; Saenko et al., 2010; Kulis et al., 2011; Fernando et al., 2013; Gopalan et al., 2011; Jhuo et al., 2012; Ammar et al., 2015). The most relevant to our work is Tzeng et al. (2017), where a discriminator is used to align intermediate features across domains. Similar in spirit, in our proposed low-level alignment through density imitation, we use a discriminator to align the behavior of low-level policies across morphologies.

In the context of RL, transfer learning (Taylor & Stone, 2009) research has focused on learning transferable features across tasks (Parisotto et al., 2015; Barreto et al., 2017; Omidshafiei et al., 2017). Another line of work by (Rusu et al., 2016; Kinsky et al., 2017; Devin et al., 2017) has focused on network architectures that improves transfer of

RL policies. Since the focus of our work is on transfer across morphologies, we note that the aforementioned works are orthogonal and complementary to ours.

5.2. Hierarchical Reinforcement Learning

HRL based techniques (Barto & Mahadevan, 2003; Bacon et al., 2017; Li et al., 2019a) have been able to solve complex or long-horizon tasks through temporal abstraction across hierarchies as seen in Levy et al. (2017) and Nachum et al. (2018a). In a similar vein, several works have focused on the discovery of primitives (Eysenbach et al., 2018; Sharma et al., 2019a; Shankar et al., 2020), which are useful for hierarchical RL. These ideas have already been combined, as in Stochastic Neural Networks by Florensa et al. (2017), where skills are learned in pretraining and then used to solve diverse complex tasks. Similarly, Andreas et al. (2016) learn modular sub-policies to solve a temporally extended task.

These prior works inform our choice to use HRL as the backbone of our framework. Moreover, hierarchical decoupling allows for a natural delineation of control (Wolpert & Kawato, 1998). However, we note that unlike standard hierarchical RL techniques, we train our low-level policies independent of the high-level policies. Although a few recent works have alluded to the potential of hierarchical policies for morphology transfer (Peng et al., 2017; Tirumala et al., 2019; Li et al., 2019b; Hu & Montana, 2019), to our knowledge we are the first to focus on this problem.

5.3. Imitation Learning

The central contribution of this work is decoupled imitation learning for the low-level and high-level policies in a hierarchical setup. In the context of control, the field of learning from demonstrations (LfD) (Nicolescu & Mataric, 2003; Argall et al., 2009) learns to reproduce a set of demonstrated expert behavior. A popular technique called behavior cloning (Esmaili et al., 1995) focuses on fitting parametric models to expert demonstrations (Kober & Peters, 2009; Peters et al., 2013). Several works (Niekum et al., 2012; Krishnan et al., 2018; Murali et al., 2016; Meier et al., 2011) focus on segmenting demonstrations followed by fitting models to each of the segments. Rajeswaran et al. (2017) focus on regularizing the behavior cloning objective for dexterous hand manipulation. Inspired by this idea, we use a KL-regularized objective in the context of regularizing the high-level policy. More recently, Goyal et al. (2019) use KL-regularization in the context of goal-conditioned RL and Galashov et al. (2019) use KL-regularization in conjunction with a divided state space for learning re-usable behaviors. Additionally, Sharma et al. (2019b) address the third person imitation problem, where the demonstrator is different than the acting agent, by using a hierarchical setup. Instead of imitation through cloning, inverse RL (Ng et al.,

2000; Abbeel & Ng, 2004) focuses on recovering the underlying reward function from expert demonstrations. Ho & Ermon (2016) extends inverse RL to higher dimensional state-action demonstrations by learning a parametric expert density model through discriminative learning between expert demonstration and learned behavior. Following this technique, several works have extended this idea to third person demonstrations (Stadie et al., 2017) and stochastic demonstrations (Li et al., 2017). Solving the information theoretic formulation for low-level state alignment reduces to a similar discriminative learning approach. However, instead of differentiating between expert demonstrations and learned behavior, our discriminator differentiates between the simpler agent’s low-level behavior and the more complex agent’s low-level.

6. Conclusions

In this work, we have presented one of the first steps towards morphology transfer by using hierarchically decoupled imitation. This technique allows transferring complex long horizon behavior from morphologically simple agents to more complex ones in a fraction of sample complexity compared to standard RL techniques. Although this work focuses on simulated environments, we believe that this opens the door to research in morphological transfer on real robots. Moreover, although our technique for decoupled imitation is presented in the context of morphological transfer, we believe that the technique is flexible enough to be applied to general purpose imitation learning.

Acknowledgements

We thank AWS for computing resources. We also gratefully acknowledge the support from Berkeley DeepDrive, NSF, and the ONR Pecase award. Finally, we thank Alex Li, Laura Smith, and the rest of the Robot Learning Lab community for their insightful comments and suggestions.

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.
- Ammar, H. B., Eaton, E., Ruvolo, P., and Taylor, M. E. Unsupervised cross-domain transfer in policy gradient reinforcement learning via manifold alignment. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- Andreas, J., Klein, D., and Levine, S. Modular multitask reinforcement learning with policy sketches, 2016.
- Argall, B. D., Chernova, S., Veloso, M., and Browning, B.

- A survey of robot learning from demonstration. *Robotics and autonomous systems*, 2009.
- Aytar, Y. and Zisserman, A. Tabula rasa: Model transfer for object category detection. In *2011 international conference on computer vision*, pp. 2252–2259. IEEE, 2011.
- Bacon, P.-L., Harb, J., and Precup, D. The option-critic architecture. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- Baker, B., Kanitscheider, I., Markov, T., Wu, Y., Powell, G., McGrew, B., and Mordatch, I. Emergent tool use from multi-agent autotutorials. *arXiv preprint arXiv:1909.07528*, 2019.
- Barreto, A., Dabney, W., Munos, R., Hunt, J. J., Schaul, T., van Hasselt, H. P., and Silver, D. Successor features for transfer in reinforcement learning. In *Advances in neural information processing systems*, pp. 4055–4065, 2017.
- Barto, A. G. and Mahadevan, S. Recent advances in hierarchical reinforcement learning. *Discrete event dynamic systems*, 13(1-2):41–77, 2003.
- Bengio, Y., Louradour, J., Collobert, R., and Weston, J. Curriculum learning. In *ICML*, 2009.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- Caruana, R. Multitask learning. *Machine learning*, 28(1): 41–75, 1997.
- Chen, T., Murali, A., and Gupta, A. Hardware conditioned policies for multi-robot transfer learning. In *Advances in Neural Information Processing Systems*, pp. 9333–9344, 2018.
- Devin, C., Gupta, A., Darrell, T., Abbeel, P., and Levine, S. Learning modular neural network policies for multi-task and multi-robot transfer. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2169–2176. IEEE, 2017.
- Doersch, C. and Zisserman, A. Multi-task self-supervised visual learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2051–2060, 2017.
- Duan, Y., Chen, X., Houthoofd, R., Schulman, J., and Abbeel, P. Benchmarking deep reinforcement learning for continuous control. In *International Conference on Machine Learning*, pp. 1329–1338, 2016.
- Esmaili, N., Sammut, C., and Shirazi, G. Behavioural cloning in control of a dynamic system. IEEE, 1995.
- Eysenbach, B., Gupta, A., Ibarz, J., and Levine, S. Diversity is all you need: Learning skills without a reward function. *CoRR*, abs/1802.06070, 2018. URL <http://arxiv.org/abs/1802.06070>.
- Fernando, B., Habrard, A., Sebban, M., and Tuytelaars, T. Unsupervised visual domain adaptation using subspace alignment. In *Proceedings of the IEEE international conference on computer vision*, pp. 2960–2967, 2013.
- Florensa, C., Duan, Y., and Abbeel, P. Stochastic neural networks for hierarchical reinforcement learning, 2017.
- Galashov, A., Jayakumar, S. M., Hasenclever, L., Tirumala, D., Schwarz, J., Desjardins, G., Czarnecki, W. M., Teh, Y. W., Pascanu, R., and Heess, N. Information asymmetry in kl-regularized rl. *arXiv preprint arXiv:1905.01240*, 2019.
- Girshick, R., Donahue, J., Darrell, T., and Malik, J. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 580–587, 2014.
- Gopalan, R., Li, R., and Chellappa, R. Domain adaptation for object recognition: An unsupervised approach. In *2011 international conference on computer vision*, pp. 999–1006. IEEE, 2011.
- Goyal, A., Islam, R., Strouse, D., Ahmed, Z., Botvinick, M., Larochelle, H., Bengio, Y., and Levine, S. Infobot: Transfer and exploration via the information bottleneck. *arXiv preprint arXiv:1901.10902*, 2019.
- Haarnoja, T., Zhou, A., Hartikainen, K., Tucker, G., Ha, S., Tan, J., Kumar, V., Zhu, H., Gupta, A., Abbeel, P., et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018.
- Hill, A., Raffin, A., Ernestus, M., Gleave, A., Kanervisto, A., Traore, R., Dhariwal, P., Hesse, C., Klimov, O., Nichol, A., Plappert, M., Radford, A., Schulman, J., Sidor, S., and Wu, Y. Stable baselines. <https://github.com/hill-a/stable-baselines>, 2018.
- Ho, J. and Ermon, S. Generative adversarial imitation learning. In *Advances in neural information processing systems*, pp. 4565–4573, 2016.
- Hoffman, J., Darrell, T., and Saenko, K. Continuous manifold based adaptation for evolving visual domains. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 867–874, 2014.
- Hu, Y. and Montana, G. Skill transfer in deep reinforcement learning under morphological heterogeneity. *arXiv preprint arXiv:1908.05265*, 2019.

- Jaderberg, M., Mnih, V., Czarnecki, W. M., Schaul, T., Leibo, J. Z., Silver, D., and Kavukcuoglu, K. Reinforcement learning with unsupervised auxiliary tasks. *arXiv preprint arXiv:1611.05397*, 2016.
- Jhuo, I.-H., Liu, D., Lee, D., and Chang, S.-F. Robust visual domain adaptation with low-rank reconstruction. In *2012 IEEE conference on computer vision and pattern recognition*, pp. 2168–2175. IEEE, 2012.
- Kaelbling, L. P., Littman, M. L., and Moore, A. W. Reinforcement learning: A survey. *Journal of artificial intelligence research*, 4:237–285, 1996.
- Kakade, S. M. A natural policy gradient. In *Advances in neural information processing systems*, pp. 1531–1538, 2002.
- Kansky, K., Silver, T., Mély, D. A., Eldawy, M., Lázaro-Gredilla, M., Lou, X., Dorfman, N., Sidor, S., Phoenix, S., and George, D. Schema networks: Zero-shot transfer with a generative causal model of intuitive physics. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1809–1818. JMLR. org, 2017.
- Kober, J. and Peters, J. Learning motor primitives for robotics. In *ICRA*, 2009.
- Kokkinos, I. Ubernet: Training a universal convolutional neural network for low-, mid-, and high-level vision using diverse datasets and limited memory. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6129–6138, 2017.
- Krishnan, S., Garg, A., Patil, S., Lea, C., Hager, G., Abbeel, P., and Goldberg, K. Transition state clustering: Unsupervised surgical trajectory segmentation for robot learning. In *RR*. 2018.
- Kulis, B., Saenko, K., and Darrell, T. What you saw is not what you get: Domain adaptation using asymmetric kernel transforms. In *CVPR 2011*, pp. 1785–1792. IEEE, 2011.
- Kulkarni, T. D., Narasimhan, K., Saeedi, A., and Tenenbaum, J. Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation. In *Advances in neural information processing systems*, pp. 3675–3683, 2016.
- Levy, A., Konidaris, G., Platt, R., and Saenko, K. Learning multi-level hierarchies with hindsight, 2017.
- Li, A. C., Florensa, C., Clavera, I., and Abbeel, P. Sub-policy adaptation for hierarchical reinforcement learning. *arXiv preprint arXiv:1906.05862*, 2019a.
- Li, S., Wang, R., Tang, M., and Zhang, C. Hierarchical reinforcement learning with advantage-based auxiliary rewards. In *Advances in Neural Information Processing Systems*, pp. 1407–1417, 2019b.
- Li, Y., Song, J., and Ermon, S. Infogail: Interpretable imitation learning from visual demonstrations. In *Advances in Neural Information Processing Systems*, pp. 3812–3822, 2017.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- Marino, K., Gupta, A., Fergus, R., and Szlam, A. Hierarchical rl using an ensemble of proprioceptive periodic policies. 2018.
- Meier, F., Theodorou, E., Stulp, F., and Schaal, S. Movement segmentation using a primitive library. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3407–3412. IEEE, 2011.
- Murali, A., Garg, A., Krishnan, S., Pokorny, F. T., Abbeel, P., Darrell, T., and Goldberg, K. Tsc-dl: Unsupervised trajectory segmentation of multi-modal surgical demonstrations with deep learning. In *ICRA*, 2016.
- Murali, A., Pinto, L., Gandhi, D., and Gupta, A. Cassl: Curriculum accelerated self-supervised learning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6453–6460. IEEE, 2018.
- Nachum, O., Gu, S., Lee, H., and Levine, S. Data-efficient hierarchical reinforcement learning. *CoRR*, abs/1805.08296, 2018a. URL <http://arxiv.org/abs/1805.08296>.
- Nachum, O., Gu, S., Lee, H., and Levine, S. Near-optimal representation learning for hierarchical reinforcement learning. *arXiv preprint arXiv:1810.01257*, 2018b.
- Ng, A. Y., Russell, S. J., et al. Algorithms for inverse reinforcement learning. In *Icml*, volume 1, pp. 663–670, 2000.
- Nicolescu, M. N. and Mataric, M. J. Natural methods for robot task learning: Instructive demonstrations, generalization and practice. In *AAMAS*, 2003.
- Niekum, S., Osentoski, S., Konidaris, G., and Barto, A. G. Learning and generalization of complex tasks from unstructured demonstrations. IEEE, 2012.
- Omidshafiei, S., Pazis, J., Amato, C., How, J. P., and Vian, J. Deep decentralized multi-task multi-agent reinforcement learning under partial observability. In *Proceedings of*

- the 34th International Conference on Machine Learning- Volume 70*, pp. 2681–2690. JMLR. org, 2017.
- Pan, S. J. and Yang, Q. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10): 1345–1359, 2009.
- Parisotto, E., Ba, J. L., and Salakhutdinov, R. Actor-mimic: Deep multitask and transfer reinforcement learning. *arXiv preprint arXiv:1511.06342*, 2015.
- Pathak, D., Lu, C., Darrell, T., Isola, P., and Efros, A. A. Learning to control self-assembling morphologies: a study of generalization via modularity. In *Advances in Neural Information Processing Systems*, pp. 2292–2302, 2019.
- Peng, X. B., Berseth, G., Yin, K., and Van De Panne, M. Deeploco: Dynamic locomotion skills using hierarchical deep reinforcement learning. *ACM Transactions on Graphics (TOG)*, 36(4):1–13, 2017.
- Peng, X. B., Chang, M., Zhang, G., Abbeel, P., and Levine, S. Mpc: Learning composable hierarchical control with multiplicative compositional policies. In *Advances in Neural Information Processing Systems*, pp. 3681–3692, 2019.
- Peters, J., Kober, J., Mülling, K., Krämer, O., and Neumann, G. Towards robot skill learning: From simple skills to table tennis. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 627–631. Springer, 2013.
- Pinto, L. and Gupta, A. Learning to push by grasping: Using multiple tasks for effective learning. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2161–2168. IEEE, 2017.
- Rajeswaran, A., Kumar, V., Gupta, A., Vezzani, G., Schulman, J., Todorov, E., and Levine, S. Learning complex dexterous manipulation with deep reinforcement learning and demonstrations, 2017.
- Rusu, A. A., Rabinowitz, N. C., Desjardins, G., Soyer, H., Kirkpatrick, J., Kavukcuoglu, K., Pascanu, R., and Hadsell, R. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016.
- Saenko, K., Kulis, B., Fritz, M., and Darrell, T. Adapting visual category models to new domains. In *European conference on computer vision*, pp. 213–226. Springer, 2010.
- Schulman, J., Levine, S., Abbeel, P., Jordan, M. I., and Moritz, P. Trust region policy optimization. In *ICML*, pp. 1889–1897, 2015.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Shankar, T., Tulsiani, S., Pinto, L., and Gupta, A. Discovering motor programs by recomposing demonstrations. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=rkgHY0NYwr>.
- Sharma, A., Gu, S., Levine, S., Kumar, V., and Hausman, K. Dynamics-aware unsupervised discovery of skills. *CoRR*, abs/1907.01657, 2019a. URL <http://arxiv.org/abs/1907.01657>.
- Sharma, P., Pathak, D., and Gupta, A. Third-person visual imitation learning via decoupled hierarchical controller. In *Advances in Neural Information Processing Systems*, pp. 2593–2603, 2019b.
- Stadie, B. C., Abbeel, P., and Sutskever, I. Third-person imitation learning. *arXiv preprint arXiv:1703.01703*, 2017.
- Sutton, R. S., Barto, A. G., et al. *Introduction to reinforcement learning*, volume 2. MIT press Cambridge, 1998.
- Tassa, Y., Doron, Y., Muldal, A., Erez, T., Li, Y., Casas, D. d. L., Budden, D., Abdolmaleki, A., Merel, J., Lefrancq, A., et al. Deepmind control suite. *arXiv preprint arXiv:1801.00690*, 2018.
- Taylor, M. E. and Stone, P. Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research*, 10(Jul):1633–1685, 2009.
- Tirumala, D., Noh, H., Galashov, A., Hasenclever, L., Ahuja, A., Wayne, G., Pascanu, R., Teh, Y. W., and Heess, N. Exploiting hierarchy for learning and transfer in kl-regularized rl. *arXiv preprint arXiv:1903.07438*, 2019.
- Todorov, E., Erez, T., and Tassa, Y. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–5033. IEEE, 2012.
- Torrey, L. and Shavlik, J. Transfer learning. In *Handbook of research on machine learning applications and trends: algorithms, methods, and techniques*, pp. 242–264. IGI Global, 2010.
- Tzeng, E., Hoffman, J., Saenko, K., and Darrell, T. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7167–7176, 2017.
- Williams, R. J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.

Wolpert, D. M. and Kawato, M. Multiple paired forward and inverse models for motor control. *Neural networks*, 11(7-8):1317–1329, 1998.

Yang, J., Yan, R., and Hauptmann, A. G. Adapting svm classifiers to data with shifted distributions. In *Seventh IEEE International Conference on Data Mining Workshops (ICDMW 2007)*, pp. 69–76. IEEE, 2007.