

ZERO-SHOT POLICY TRANSFER WITH DISENTANGLED ATTENTION

Anonymous authors

Paper under double-blind review

ABSTRACT

Domain adaptation is an open problem in deep reinforcement learning (RL). Often, agents are asked to perform in environments where data is difficult to obtain. In such settings, agents are trained in similar environments, such as simulators, and are then transferred to the original environment. The gap between visual observations of the source and target environments often causes the agent to fail in the target environment. We present a new RL agent, SADALA (Soft Attention DisentAngled representation Learning Agent). SADALA first learns a compressed state representation. It then jointly learns to ignore distracting features and solve the task presented. SADALA’s separation of important and unimportant visual features leads to robust domain transfer. SADALA outperforms both prior disentangled-representation based RL and domain randomization approaches across RL environments (Visual Cartpole and DeepMind Lab).

1 INTRODUCTION

Reinforcement learning agents learn to maximize rewards within a task by taking actions based on observations. The advent of deep learning has enabled reinforcement learning agents to learn in high dimensional feature spaces and complex action spaces (Krizhevsky et al., 2012; Mnih et al., 2016). Additionally, deep reinforcement learning methods have allowed agents to beat human level performance in a variety of tasks, such as Go (Silver & Hassabis, 2016). However, deep reinforcement learning has two crippling drawbacks: high sample complexity and specificity to training task.

There are many domains where data collection is expensive and time consuming, such as healthcare, autonomous vehicles, and robotics (Gottesman et al., 2019; Barrett et al., 2010). Thus, agents are often trained in simulation and must transfer the resulting knowledge to reality. While this solves the issue of sample complexity, reality and simulated domains are sufficiently different that it is infeasible to train a deep reinforcement learning agent in a simulation and transfer with no additional knowledge. This is known as the reality gap (Sadeghi & Levine, 2016). Jumping the reality gap is difficult for two reasons. The first is that dynamics of a simulation are an approximation to the dynamics of the real world. Previous methods have focused on transfer between domains with different dynamics (Killian et al., 2017; Yao et al., 2018; Doshi-Velez & Konidaris, 2016). The second difficulty in jumping the reality gap is the difference in visual observations of states; due to limitations in photorealistic rendering, simulation and the real world are effectively two different visual domains (Sadeghi & Levine, 2016). Domain randomization is currently the most popular transfer method between different visual domains (Tobin et al., 2017; OpenAI et al., 2018; Sadeghi & Levine, 2016). Though it has shown effectiveness in transfer it requires training on many source environments, leading to high sample complexity costs (Tobin et al., 2017; OpenAI et al., 2018). Work has been done to directly model this transfer by using an image to image translation method which attempts to map images from the target task to the source task (Pan et al., 2017; Liu et al., 2017; Isola et al., 2017). This significantly increases the complexity of the network and heavily relies on the translation network. Therefore, we propose explicit modeling of state representation as a solution to visual transfer learning.

We present a method of robust transfer between visual domains, using attention and a β variational autoencoder to automatically identify task-relevant latent variables. This enables the reinforcement learning agent to pay attention to features relevant to solving the source domain. Additionally, the agent learns to ignore distracting features, leading to efficient transfer to the target domain with a

lower sample complexity than existing methods such as domain randomization (Tobin et al., 2017; Zhang et al., 2017; Bousmalis et al., 2018; Prakash et al., 2019; Chebotar et al., 2019; Tremblay et al., 2018).

2 BACKGROUND

We formalize the concept of transfer between related MDPs. We denote the source domain as D_S and the target domain as D_T . Both domains are MDPs, defined as $D_S = (S_S, A_S, T_S, R_S)$ and $D_T = (S_T, A_T, T_T, R_T)$, where S are states, A are actions, T is the transition function, and R is the reward function. We assume that both D_S and D_T share a discount factor γ . In general transfer, states and actions may be quite different, while transition and reward functions are similar. We denote a family of related MDPs as M . Within M , transitions, rewards, and actions are the same across MDPs, and states are different. These states are generated from a set of latent state factors S^z , each passed through a different MDP-specific renderer G_U for some MDP $U \in M$. In the case of visual cartpole, S^z contains the position and velocity of the cart, the angle and position of the pole, and the colors of the cart, pole, axle, and track. The observed states are then rendered by the renderer. Specifically, $S_U = G_U(S^z)$ and $S_V = G_V(S^z)$, where G_U and G_V are any two MDPs within M .

In this paper, we target transfers where $T_S \approx T_T$, $R_S = R_T$, $A_S = A_T$, but $S_S = G_S(S^z) \neq S_T = G_T(S^z)$. Namely, when target dynamics are similar to source dynamics, rewards, actions, and true states are identical, but observed states S_S and S_T differ due to differing renderers G_S and G_T .

3 SADALA

3.1 FEATURE EXTRACTION WITH β -VARIATIONAL AUTOENCODERS

Deep RL agents typically learn a mapping from states $s \in S$ to actions $a \in A$, either by directly maximizing reward or through a Q function, using a deep neural network. The deep network implicitly learns a mapping function $F : S_U \rightarrow \hat{S}^z$, for some MDP $U \in M$ and a policy function $\pi : \hat{S}^z \rightarrow A$ that maps from latent state factors to actions. These functions are implicitly learned as a result of the feed-forward nature of deep neural networks. The policy π is a function of the last layer of the network. The output of the last layer of the neural network is a highly processed feature representation of the input that is optimized for use in generating π . While this implicit mapping is sufficient for a deep RL agent to solve a single MDP from M , it is not general enough to transfer to other MDPs within M . By optimizing for π F overfits to the specific S_U from the source domain U .

To circumvent this overfitting, DARLA proposed learning a feature extractor F' that maps from any observed state S_X to the latent state factors S^z (Higgins et al., 2017). DARLA then trains an RL agent on the (approximate) extracted latent state factors \hat{S}^z , separating the “vision” and “action” of the RL agent. DARLA uses a variational autoencoder to learn the mapping F' and an RL network to learn the policy π (Higgins et al., 2017).

The extracted latent state factors \hat{S}^z are (mostly) disentangled in practice, leading to better transfer results than a baseline RL agent. The disentangled factors each correlate with an independent factor in the true state S^z of the task. However, it assumes that all latent state variables $z \in S^z$ are relevant to solving the MDP. This is not usually the case. When attempting to locate an object such as an enemy in a video game, an object for a robot to grasp, or a road for an autonomous vehicle to follow, the position and orientation of the object are the most important factors. Visual distractions such as the color of the enemy, opacity of the object, or construction material of the road are entirely irrelevant to the localization of the object. Dependence on these factors causes the RL agent to fail to robustly transfer to target visual domains.

3.2 ADDING AN ATTENTION MECHANISM

SADALA is similar to DARLA, but has one key difference—it explicitly learns to ignore irrelevant extracted latent state variables.




Figure 1: SADALA Framework



Figure 2: β Variational Autoencoder.

SADALA learns the same mapping function $F' : S_U \rightarrow \hat{S}^z$ as DARLA, also utilizing a β variational autoencoder. It then learns a policy $\pi : \hat{S}^z \rightarrow A$, but utilizes a soft attention mechanism to ignore irrelevant features. Using this mechanism, it implicitly learns two functions: $\mathcal{A} : \hat{S}^z \rightarrow S^a$ that maps from extracted latent state factors to latent state factors relevant to solving the MDP, and $\pi_{\mathcal{A}} : S^a \rightarrow A$ that maps from relevant latent state factors to actions.

Soft attention mechanisms learn an importance weight $w_i \in [0, 1]$ for each feature z_i in part of a neural network. These weights are then multiplied element-wise with the corresponding feature. Thus, the attention layer can be defined as follows:

$$f(Z) = W \odot Z,$$

where W is the vector of importance weights w_i and \odot is element-wise multiplication.

W is defined as

$$W = \sigma(h(Z))$$

where h is a differentiable function parameterized by dense layers of a neural network. Thus, the values $w_i \in W$ are constrained between 0 and 1, inclusive. This allows the RL agent to ignore a feature z_i by setting w_i to 0. It also allows the agent to down-weight the importance of a feature by setting w_i to a fractional value or pay full attention by setting w_i to 1.

3.3 SADALA FRAMEWORK

The SADALA framework is shown in Figure 1. At inference time, an input state $s \in S$ is fed into a β variational autoencoder. The bottleneck representation of the β -VAE is then used as extracted latent state factors \hat{S}^z . Those features are then fed into the attention mechanism, which outputs the weighted features S^a , according to the formula f above. Finally, the weighted features are used as the inputs to the deep reinforcement learning algorithm.

Though the feature selection and policy learning stages are represented separately in figure 1, they are contained within the same neural network and optimized jointly.

3.4 SADALA TRAINING

Training SADALA requires a few steps. First, the β variational autoencoder, shown in figure 2 must be trained. To do so, we use the loss term of a variational autoencoder, modified to fit the β -VAE (Higgins et al., 2017). The standard VAE loss is

$$\mathcal{L}_{\text{VAE}} = \mathbb{E}_{q(z|x)}[\log p_{\theta}(x|z)] - D_{KL}(q_{\phi}(z|x)||p(z)).$$

The β -VAE loss is

$$\mathcal{L}_{\beta\text{-VAE}} = \mathbb{E}_{q(z|x)}[\log p_{\theta}(x|z)] - \beta D_{KL}(q_{\phi}(z|x)||p(z)),$$

where β is a weighting term that pressures the encoder q to output a latent space similar to an isotropic unit gaussian. Increasing the value of β encourages disentanglement of the latent space while sacrificing minor reconstruction details.

We sample states from the environment, following a random policy. We then train the β -VAE to reconstruct these images and freeze the network weights.

Next, the feature selection and policy learning stages of SADALA must be trained. The weights of the pre-trained β -VAE are frozen and used as weights in the state representation stage.

The feature selection and policy learning stages are then jointly optimized. The SADALA framework can utilize any deep RL algorithm (A3C, REINFORCE, DQN, etc), with minor modification. Specifically, we add a loss term to the loss function \mathcal{L}_{RL} of the original algorithm. Our new loss is

$$\mathcal{L}_{\text{A-RL}} = \mathcal{L}_{\text{RL}} + |W|,$$

where W is the learned attention weights. This added term is an L1 regularization on the attention weights, which encourages sparsity in feature selection (Kolter & Ng, 2009; Ng, 2004).

We additionally test one other modification to the loss term. We make the assumption that the ordering of the extracted latent state variables \hat{S}^z does not change between subsequent frames within an episode. This, the weight w_i for some latent state factor z_i should not change based on the frame. To enforce this loss on the feature selection stage, we use the following loss

$$\mathcal{L}_{\text{PA-RL}} = \mathcal{L}_{\text{RL}} + |W| + \sum_{j \in \text{len}(z)} \text{Var}[W^j],$$

where W^j is vector of weights corresponding to the j th extracted latent state factor vector across the set of training data and Var is the variance. This loss is added to enforce w_i corresponding to the feature z_i to be the same regardless of input. This is similar to learning a pseudo-attention weight vector, rather than an attention weight vector.

See the appendix for additional training details.

4 RESULTS AND DISCUSSION

We test the SADALA framework on two transfer learning tasks, using A3C as the deep RL algorithm. The first task is Visual Cartpole. This domain is the same as Cartpole-v1 in OpenAI Gym with two key differences (Brockman et al., 2016). The observed state is now the pixel rendering of the cartpole as well as the velocities of the cart and pole. Thus, the agent must learn to predict the position of the cart and pole from the rendering. Additionally, we modify this domain to include a transfer task. The agent must learn to transfer its knowledge of the game across different color configurations for the cart, pole, and track. Thus, Visual Cartpole is defined as a family of MDPs M where the true state S^z is the positions and velocities of the cart and pole and the observed state $S_U = G_U(S^z)$ for an MDP $U \in M$ is the pixel observations and velocity values. Optimally, the agent should learn to ignore the factors in extracted latent state factors \hat{S}^z that correspond to color, as they do not aid the agent in balancing the pole. This task tests the agent’s ability to ignore irrelevant latent state factors.

The second task is the ‘‘Collect Good Objects’’ task from Deepmind Lab (Beattie et al., 2016). The agent must learn to navigate in first person and pick up ‘‘good’’ objects while avoiding ‘‘bad’’ objects. This task is defined as a family of MDPs M where the true state S^z contains the position of the agent, the position of good and bad objects, the type of good and bad objects, and the color




Figure 3: β -VAE reconstructions of samples from the Visual Cartpole environment. Top row contains original images and the bottom row contains their reconstructions.

of the walls and floor. In a single MDP $U \in M$, all good objects are hats or all good objects are balloons. Similarly, all bad objects are either cans or cakes. The walls and floor can either take a green and orange colorscheme or a red and blue colorscheme. The agent is trained on hats/cans with the green/orange colorscheme and balloons/cakes with both colorschemes. It is then tested on hats/cans with the red/blue colorscheme. Optimally, the agent should learn to ignore the color of the floor and walls. Additionally, it should use the type of object to determine if it is good or bad. This task tests the agent’s ability to ignore distracting latent state factors (the color of the walls and floor) while attending to relevant factors (the positions and types of objects and its own position).

To test the results of the SADALA algorithm, we first test the reconstruction and disentanglement properties of the β -VAE used in the state representation stage. Note that this stage is identical to that of DARLA (Higgins et al., 2017). As such, we expect the disentanglement properties to be similar.

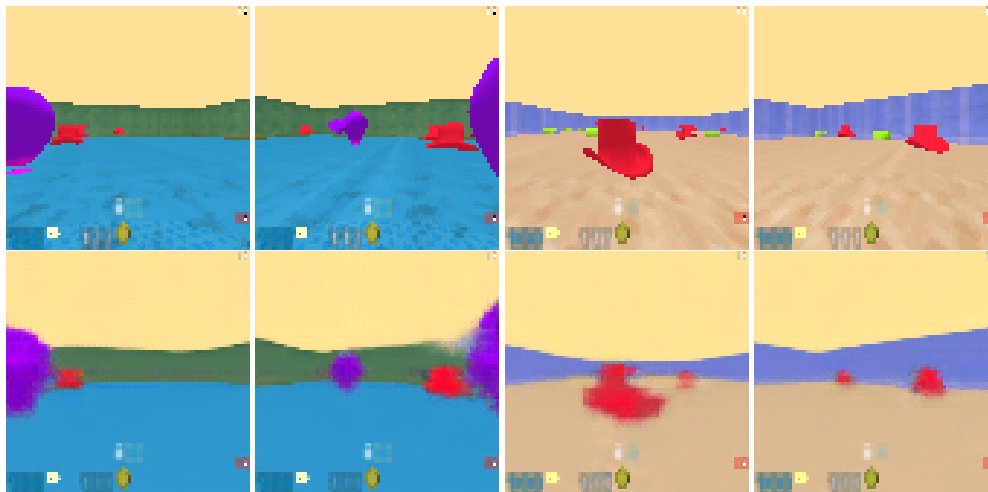


Figure 4: β -VAE reconstructions of samples from the Deepmind Lab environment. Top row contains original images and bottom row contains their reconstructions.

See figure 3 for reconstructions of the cartpole state. Based on the reconstructions, it is apparent that the β -VAE has learned to represent cart position and pole angle. Though the angle of the poles is slightly incorrect in the first set of images, the pole is tilted in the correct direction, yielding sufficiently correct extracted latent state factors.

Additionally, the color of the cart, pole, and background is incorrect in the third pair of images. While this demonstrates that the identification and reconstructions of colors is not infallible, the position of the cart and pole remains correct, yielding a set of extracted latent state parameters that is sufficient to solve the MDP.




Figure 5: β -VAE reconstructions from Visual Cartpole. The image on the left is the original image from the simulation and the image on the right is the reconstruction after multiplying with the learned attention weights.

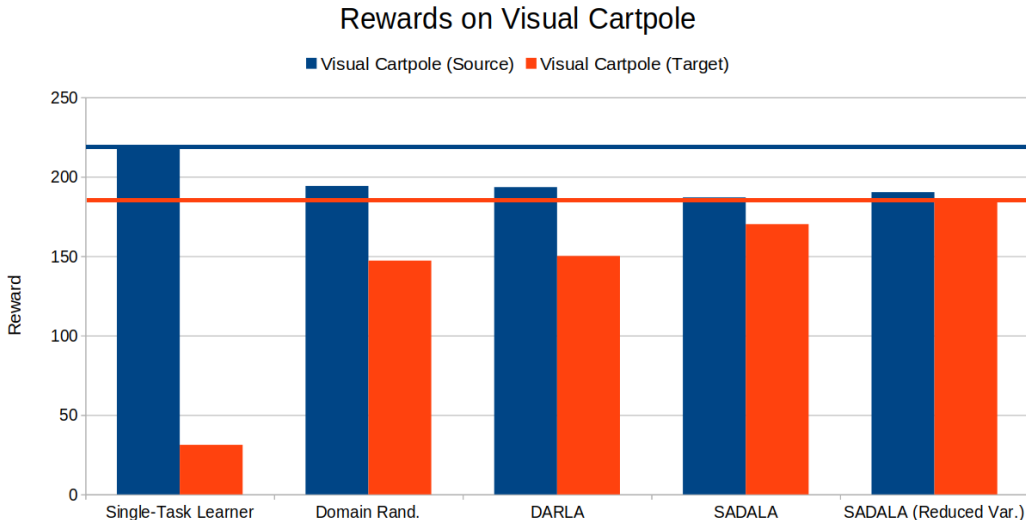


Figure 6: Source and target rewards on Visual Cartpole.

See figure 5 for a visualization of reconstruction with attention. In the original image, the pole is standing straight and the cart is centered. In the reconstruction, the cart is centered, and the pole is almost upright. However, the reconstruction does not include the colors of the cart or pole. Instead it fills the cart and pole with the mean color of the dataset. This shows that the attention weights are properly learning to ignore color and instead pay attention to the position of the cart and pole.

Figures 6 and 7 gives a comparison of the performance of the algorithms across environments. Note that all of the algorithms are sufficient at solving the source task, with the single-task learner performing slightly better. This is due to the fact that the single-task learner can optimize its convolutional filters to the source domain at the expense of failing to transfer to the target domain. For numerical reward values, see table 1 in the appendix.

Domain Randomization was not attempted on the Deepmind Lab task due to limited computational resources. Training the A3C agent on one source domain requires about 24 hours of computational time, making it infeasible to train for tens or hundreds of simulated domains. This reflects the result found by OpenAI et al. (2018), where 100 years of simulation experience was required for transfer.

The single-task learner achieves better rewards on all source tasks than any of the transfer-specific agents. Domain randomization performs less well because of the complexity of the domain randomization task. Rather than optimizing reward for a single domain, the agent must optimize reward across a large set of domains.




Figure 7: Source and target rewards on Deepmind Lab.

DARLA, SADALA, and SADALA with reduced variance also perform less well on the source task than the baseline agent. This is due to imperfections in the β -VAE. Though the β -VAE attempts to reconstruct the input image, it does not do so perfectly, as shown in figure 3. This shows that its extraction of latent state features is not perfect, leading to potentially confusing states given to the RL agent. Additionally, while the β -VAE’s goal is to learn a disentangled representation, it does not do so perfectly. As such, (partially) entangled latent state factors may further confuse the RL agent.

The single-task learner fails to transfer its policy to the target domain. DARLA transfers some of its knowledge to the target domain. Domain randomization also transfers some knowledge. Finally, SADALA transfers more of its knowledge. The single-task agent has no incentive to learn a factored state representation that enables transfer. Its convolutional filters will directly optimize to maximize reward in the source policy. Thus, if the filters are searching for a hat on a blue floor but tested on a hat on a red floor the convolutional filters will fail to transfer.

DARLA learns a state representation that is mostly disentangled. This allows the RL agent to learn to ignore unimportant features such as cart color and utilize important features such as cart position. The factored state representation forces the RL agent to learn a more robust policy. However, the neural network parameterization of the RL policy must implicitly learn to ignore unimportant factors. Therefore, when presented with unseen information, the RL agent may not properly ignore unimportant factors.

Domain randomization forces the neural network to implicitly learn a state representation sufficient to transfer between tasks. This requires large amounts of training data and is less robust than explicit modeling of latent state factors.

SADALA builds on DARLA by adding an explicit attention mechanism, allowing it to more effectively ignore unimportant features. Due to the use of the sigmoid activation in the attention mechanism, the attention weights W are bounded between 0 and 1. In addition to providing a direct weight on the importance of a feature, this bound prevents high variance of attention weights across different inputs.

SADALA with the variance reduction term performs worse than both DARLA and SADALA without variance reduction on the Deepmind lab task but better on the other two. In the scenario where the extracted latent state factors from the β -VAE are perfectly disentangled, static attention weights should be sufficient to solve the source task and should transfer better to the target task, as in the

Visual Cartpole and Visual Pendulum tasks. However, the β -VAE does not output perfectly disentangled factors, especially in more complex visual domains such as the Deepmind lab. Thus, the amount of attention paid to each feature from the β -VAE may differ across tasks, violating the assumption that attention weights should be zero variance.

5 CONCLUSION

In this paper we propose SADALA, a three stage method for zero-shot domain transfer. First, SADALA learns a feature extractor that represents input states (images) as disentangled factors. It then filters these latent factors using an attention mechanism to select those most important to solving a source task. Jointly, it learns a policy for the source task that is robust to changes in states, and is able to transfer to related target tasks. We validate the performance of SADALA on both a high-dimensional continuous-control problem (Visual Cartpole) and a 3D naturalistic first-person simulated environments (Deepmind Lab). We show that the attention mechanism introduced is able to differentiate between important and unimportant latent features, enabling robust transfer.

REFERENCES

- Samuel Barrett, Matthew E Taylor, and Peter Stone. Transfer learning for reinforcement learning on a physical robot. In *Ninth International Conference on Autonomous Agents and Multiagent Systems-Adaptive Learning Agents Workshop (AAMAS-ALA)*, volume 1, 2010.
- Charles Beattie, Joel Z Leibo, Denis Teplyashin, Tom Ward, Marcus Wainwright, Heinrich Küttler, Andrew LeFrancq, Simon Green, Víctor Valdés, Amir Sadik, et al. Deepmind lab. *arXiv preprint arXiv:1612.03801*, 2016.
- Konstantinos Bousmalis, Alex Irpan, Paul Wohlhart, Yunfei Bai, Matthew Kelcey, Mrinal Kalakrishnan, Laura Downs, Julian Ibarz, Peter Pastor, Kurt Konolige, et al. Using simulation and domain adaptation to improve efficiency of deep robotic grasping. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4243–4250. IEEE, 2018.
- Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- Yevgen Chebotar, Ankur Handa, Viktor Makoviychuk, Miles Macklin, Jan Issac, Nathan Ratliff, and Dieter Fox. Closing the sim-to-real loop: Adapting simulation randomization with real world experience. In *2019 International Conference on Robotics and Automation (ICRA)*, pp. 8973–8979. IEEE, 2019.
- Finale Doshi-Velez and George Konidaris. Hidden parameter markov decision processes: A semi-parametric regression approach for discovering latent task parametrizations. In *IJCAI: proceedings of the conference*, volume 2016, pp. 1432. NIH Public Access, 2016.
- Omer Gottesman, Fredrik Johansson, Matthieu Komorowski, Aldo Faisal, David Sontag, Finale Doshi-Velez, and Leo Anthony Celi. Guidelines for reinforcement learning in healthcare. *Nat Med*, 25(1):16–18, 2019.
- Irina Higgins, Arka Pal, Andrei A. Rusu, Loïc Matthey, Christopher Burgess, Alexander Pritzel, Matthew Botvinick, Charles Blundell, and Alexander Lerchner. DARLA: improving zero-shot transfer in reinforcement learning. *CoRR*, abs/1707.08475, 2017. URL <http://arxiv.org/abs/1707.08475>.
- Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1125–1134, 2017.
- Taylor W Killian, Samuel Daulton, George Konidaris, and Finale Doshi-Velez. Robust and efficient transfer learning with hidden parameter markov decision processes. In *Advances in Neural Information Processing Systems*, pp. 6250–6261, 2017.

- J Zico Kolter and Andrew Y Ng. Regularization and feature selection in least-squares temporal difference learning. In *Proceedings of the 26th annual international conference on machine learning*, pp. 521–528. ACM, 2009.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (eds.), *Advances in Neural Information Processing Systems 25*, pp. 1097–1105. Curran Associates, Inc., 2012. URL <http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf>.
- Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *Advances in neural information processing systems*, pp. 700–708, 2017.
- Volodymyr Mnih, Adrià Puigdomènech Badia, Mehdi Mirza, Alex Graves, Timothy P. Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. *CoRR*, abs/1602.01783, 2016. URL <http://arxiv.org/abs/1602.01783>.
- Andrew Y Ng. Feature selection, l1 vs. l2 regularization, and rotational invariance. In *Proceedings of the twenty-first international conference on Machine learning*, pp. 78. ACM, 2004.
- OpenAI, :, M. Andrychowicz, B. Baker, M. Chociej, R. Jozefowicz, B. McGrew, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray, J. Schneider, S. Sidor, J. Tobin, P. Welinder, L. Weng, and W. Zaremba. Learning Dexterous In-Hand Manipulation. *ArXiv e-prints*, August 2018.
- Xinlei Pan, Yurong You, Ziyang Wang, and Cewu Lu. Virtual to real reinforcement learning for autonomous driving. *arXiv preprint arXiv:1704.03952*, 2017.
- Aayush Prakash, Shaad Boochoon, Mark Brophy, David Acuna, Eric Cameracci, Gavriel State, Omer Shapira, and Stan Birchfield. Structured domain randomization: Bridging the reality gap by context-aware synthetic data. In *2019 International Conference on Robotics and Automation (ICRA)*, pp. 7249–7255. IEEE, 2019.
- Fereshteh Sadeghi and Sergey Levine. Cad2rl: Real single-image flight without a single real image. *arXiv preprint arXiv:1611.04201*, 2016.
- David Silver and Demis Hassabis. Alphago: Mastering the ancient game of go with machine learning. *Research Blog*, 2016.
- Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 23–30. IEEE, 2017.
- Jonathan Tremblay, Aayush Prakash, David Acuna, Mark Brophy, Varun Jampani, Cem Anil, Thang To, Eric Cameracci, Shaad Boochoon, and Stan Birchfield. Training deep networks with synthetic data: Bridging the reality gap by domain randomization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 969–977, 2018.
- Jiayu Yao, Taylor Killian, George Konidaris, and Finale Doshi-Velez. Direct policy transfer via hidden parameter markov decision processes. 2018.
- Fangyi Zhang, Jürgen Leitner, Michael Milford, and Peter Corke. Sim-to-real transfer of visuomotor policies for reaching in clutter: Domain randomization and adaptation with modular networks. *world*, 7(8), 2017.

A APPENDIX

A.1 TASKS

We test the SADALA framework on two transfer learning tasks, using A3C as the deep RL algorithm. These tasks are described in detail below.

A.1.1 VISUAL CARTPOLE

As in the original cartpole domain, an agent must learn to keep a pole balanced on a cart upright for as long as possible. The states in the original domain are the position and velocity of the cart and pole, and the states in the modified domain are image renderings of the cart and pole. The agent acts in both domains by applying a force of $+1$ or -1 to the cart, and the environment transitions based on a deterministic physics simulation of the cart and pole. The reward is $+1$ at every timestep that the pole has not fallen over, and the episode terminates when the pole falls over or the cart moves too far to the left or right. Similar to OpenAI Gym’s Cartpole V1, Visual Cartpole ends after a maximum of 500 steps, yielding a maximum reward of 500 per episode (Brockman et al., 2016).

When framing this MDP as a transfer learning task, the agent is trained on some number of source domains, where the colors of the cart, pole, axle, and track are varied. The agent is then tested on a target domain with held-out colors, which are sufficiently different from those of the source environments. Specifically the colors of the source environments are constrained to be outside of a hypersphere of the test environment.

Note that visual cartpole as described above is a Partially Observable Markov Decision Process (POMDP), in which there is a hidden state that is not directly observable. Specifically, it is not possible for an agent to determine the velocity of the cart and pole solely based on one image. In order to simplify the visual cartpole task to an MDP, velocity information is included in the state. This enforces the Markov property on the environment, making it better suited to test visual transfer.

A.1.2 DEEPMIND LAB

The second task is the Collect Good Objects task from Deepmind Lab (Beattie et al., 2016). This environment can be defined as a family of MDPs that share transition function, reward function, and action space. The goal of the agent is to pick up as many “good” objects ($+1$ reward) and avoid picking up “bad” objects (-1 reward) as it can within one minute.

These objects can be hats, cans, cakes, and balloons. Additionally the color combinations of the walls and floors can vary between green walls and orange floors or red walls and blue floors. These variations in state space correspond to the different MDPs within the family N of Deepmind Lab MDPs. The β -VAE is trained on a set of MDPs covering all latent state factor combinations, the RL network and attention mechanism is trained on a subset of N and tested on a held-out MDP.

A.2 ADDITIONAL RESULTS

Table 1 compares the numerical reward values displayed in figures 6 and 7. The source domains VC (Source) and DM (Source) are listed for completion. It is better to sacrifice some reward in the source domain for additional reward in target domains.

| Task | Single-Task Learner | Domain Rand. | DARLA | SADALA | SADALA (Var. Reduction) |
|-------------|---------------------|--------------|---------|-------------|-------------------------|
| VC (Source) | 219.53 | 194.2 | 193.48 | 187.0 | 190.3 |
| VC (Target) | 31.2 | 147.2 | 150.136 | 170.16 | 185.87 |
| DM (Source) | 30.8 | — | 29.09 | 28.45 | 28.06 |
| DM (Target) | 6.2 | — | 19.7 | 20.7 | 15.3 |

Table 1: Mean reward of A3C, Domain Randomization, DARLA, and SADALA agents on Visual Cartpole and Deepmind Lab. Higher is better. Each mean reward is taken over a series of 5000 test episodes according to 5 trials.

B NEURAL NETWORK DETAILS

The neural network for the Visual Cartpole task consists of two parts: the β -VAE network and the policy network. First, the encoder for the β -VAE network is composed of three convolutional layers, each with kernel size 3 and stride 1. In order, they have 32, 32, and 64 filters. The output of the last

convolutional layer is then frozen and passed through a dense layer with 256 neurons. This is then passed through another dense layer with 64 neurons, which outputs the latent space. The latent space is composed of 32 gaussians, each parameterized by a mean and variance. The decoder is an inverse of the encoder, utilizing deconvolutional layers. The encoder from the β -VAE is then frozen and used as a feature extractor for the policy network. The policy network takes the encoders outputs and passes them through one dense layer with 64 neurons, which outputs the attention weights. These weights are multiplied elementwise with the encoder’s outputs and passed through two dense layers, each with 128 neurons. Finally, the output is passed through a dense layer with two neurons, outputting the policy given the input state. The output is also passed through a dense layer with one neuron, which outputs the value of the input state.

The networks used for the Domain Randomization agent and Single-Task learner have a similar structure. They first pass the input through three convolutional filters, identical in architecture to the encoder. They then pass the outputs of the convolutional filters through two dense layers, each with 128 neurons. Finally, this output is passed through the dense layer with two neurons which outputs the policy, and the dense layer with one neuron, which outputs the value.

The neural networks used for the Deepmind Lab task are slight modifications to the neural networks described above. They have four convolutional layers instead of three, with 32, 32, 64, and 64 filters, respectively. Additionally, the hidden dense layers in the policy network have 256 rather than 128 neurons. As in DARLA, the reconstruction loss of the β -VAE for Deepmind Lab takes place in the latent space of a denoising autoencoder with 100 latent factors Higgins et al. (2017).

Relu activations are used throughout. All Visual Cartpole networks are trained with Adam with a learning rate of $5e - 4$ and all Deepmind Lab networks are trained Adam with a learning rate of $1e - 4$. The β -VAE networks are trained on a set of 1,000,000 images sampled from the observed state space of source policies. The RL networks are trained on 1,000,000 and 16,000,000 states, actions, and rewards for the Visual Cartpole and Deepmind Lab tasks, respectively. The domain randomization agent is trained on 3,000,000 Visual Cartpole states, actions, and rewards.