
Transfer Learning for Related Reinforcement Learning Tasks via Image-to-Image Translation

Shani Gamrian¹ Yoav Goldberg^{1,2}

Abstract

Despite the remarkable success of Deep RL in learning control policies from raw pixels, the resulting models do not generalize. We demonstrate that a trained agent fails completely when facing small visual changes, and that fine-tuning—the common transfer learning paradigm—fails to adapt to these changes, to the extent that it is faster to re-train the model from scratch. We show that by separating the visual transfer task from the control policy we achieve substantially better sample efficiency and transfer behavior, allowing an agent trained on the source task to transfer well to the target tasks. The visual mapping from the target to the source domain is performed using unaligned GANs, resulting in a control policy that can be further improved using imitation learning from imperfect demonstrations. We demonstrate the approach on synthetic visual variants of the Breakout game, as well as on transfer between subsequent levels of Road Fighter, a Nintendo car-driving game. A visualization of our approach can be seen in <https://youtu.be/4mnkzYyXMn4> and <https://youtu.be/KCGTrQi6Ogo>.

1. Introduction

Deep reinforcement learning is becoming increasingly popular due to various recent success stories, such as the famous achievement of learning to play Atari 2600 video games from pixel-level input (Mnih et al., 2013). However, this approach depends on interacting with the environment a substantial number of times during training. Moreover, it struggles to generalize beyond its experience, the training

process of a new task has to be performed from scratch even for a related one. Recent works have tried to overcome this inefficiency with different approaches such as, learning universal policies that can generalize between related tasks (Schaul et al., 2015), as well as other transfer approaches (Fernando et al., 2017; Rusu et al., 2016).

In this work, we first focus on the Atari game Breakout, in which the main concept is moving the paddle towards the ball in order to maximize the score of the game. We modify the game by introducing visual changes such as adding a rectangle in the middle of the image or diagonals in the background. From a human perspective, it appears that making visual changes that are not significant to the game’s dynamics should not influence the score of the game. We show that the agent fails to transfer. Furthermore, fine-tuning, the main transfer learning method used today in neural networks, also fails to adapt to the small visual change: the information learned in the source task does not benefit the learning process of the very related target task, and can even decelerate it.

Our second focus is attempting to transfer agent behavior across different levels of a video game: can an agent trained on the first level of a game use this knowledge and perform adequately on subsequent levels? We explore the Nintendo game Road Fighter, a car racing game where the goal is to finish the track before the time runs out without crashing. The levels all share the same dynamics, but differ from each other visually and in aspects such as road width. Similar to the Breakout results, an agent trained to play the first level fails to correctly adapt its past experience, causing the learned policy to completely fail on the new levels.

To address the generalization problem, we propose to isolate the visual component and perform zero-shot analogy-based transfer. Concretely, the agent transfers between the tasks by learning to visually map images from the target task back to familiar corresponding images from the source task. Such mapping is naturally achieved using Generative Adversarial Networks (GANs) (Goodfellow et al., 2014). In our setup, it is not realistic to assume paired images in both domains, calling for the use of Unaligned GANs (Liu et al., 2017; Zhu et al., 2017; Kim et al., 2017; Yi et al., 2017).

¹Computer Science Department, Bar-Ilan University, Ramat-Gan, Israel ²Allen Institute for Artificial Intelligence. Correspondence to: Shani Gamrian <gamrianshani@gmail.com>, Yoav Goldberg <yoav.goldberg@gmail.com>.




Figure 1. Various variations of the Breakout game: (a) Standard version, (b) A Constant Rectangle - a rectangle in the same size as the bricks is added to the background in a predefined location, (c) A Moving Square - a square is added to the background and its location changes to one of three predefined locations every 1000 steps, (d) Green Lines - green lines in different sizes are drawn in the background, (e) Diagonals - diagonals are drawn in the left side of the background.

The approach allows the agent to effectively apply its source domain knowledge in the target domain without additional training. For cases where the visual analogy is insufficient for performing optimally, or where the GAN fails to produce sufficiently accurate mappings, we treat the resulting policy as an imperfect demonstration and further improve it using an imitation learning algorithm tailored to the imperfect demonstration scenario. The code is available at <https://github.com/ShaniGam/RL-GAN>.

Contributions This work presents three main contributions. First, in Section 2, we demonstrate how an agent trained with deep reinforcement learning algorithms fails to adapt to small visual changes, and that the common transfer method of fine-tuning fails as well. Second, in Section 3, we propose to separate the visual mapping from the game dynamics, resulting in a new transfer learning approach for related tasks based on visual input mapping. Third, in Section 4, we suggest an imitation learning from imperfect demonstrations algorithm for improving the visual-transfer-based policy in a sample efficient manner. We evaluate this approach on Breakout and Road Fighter in Section 5, and present the results comparing to different baselines. We show that our visual transfer approach is much more sample efficient than the alternatives. Moreover, we use our method as an evaluation setup for unaligned GAN architectures, based on their achieved performance on concrete down-stream tasks.

2. Generalization failures of Deep RL

Many Breakout variations that involve the same dynamics can be constructed. The main idea is to make modifications that are not critical for a human playing the game but are for the algorithm that relies on visual inputs. We demonstrate

the difficulty of deep reinforcement learning to generalize using 4 types of modifications as presented in Figure 1.

Transfer-Learning via Fine-Tuning. For all the experiments in this section forward we use the Asynchronous Advantage Actor-Critic (A3C) algorithm (Mnih et al., 2016). The A3C learns the policy and the state-value function using parallel actor-learners exploring different policies for the acceleration and stability of the training. Details of the setup can be seen in the appendix (A.1). This setup successfully trains on Breakout, reaching a score of over 400 points. However, when a network trained on the original game is presented with the game variants, it fails completely, reaching to a maximum score of only 3 points. This shows that the network does not necessarily learn the game’s concepts and heavily relies on the images it receives.

The common approach for transferring knowledge across tasks is fine-tuning. We experiment with common techniques used in deep learning models. In each setting, we have a combination of frozen and fine-tuned layers (Partial/Full) as well as layers that are initialized with the target’s parameters and layers that are initialized with random values (Random). Our settings are inspired by (Yosinski et al., 2014). We train each one of the tasks (before and after the transformation) for 60 million frames, and our evaluation metric is the total reward the agents collect in an episode averaged by the number of episodes, where an episode ends when the game is terminated or when a number of maximum steps is reached. We periodically compute the average during training. Details are available in the appendix (B).

Results. The results presented in Figure 2 show a complete failure of all the fine-tuning approaches to transfer to the target tasks. In the best scenarios the transfer takes just as many epochs as training from scratch, while in other cases starting from a trained network makes it *harder* for the network to learn the target task. As the graphs show, some of the modification interfere more than others. For example, Figure 2a shows that adding a simple rectangle can be destructive for a trained agent: while training from scratch consistently and reliably achieves scores over 300, the settings starting from a trained agent struggle to pass the 200 points mark within the same number of iterations, and have a very high variance. We noticed that during training the agent learns a strategy to maximize the score with a minimum number of actions. None of the experiments we performed showed better results when the layers in the network were fine-tuned, and some showed negative transfer which is a clear indication of an overfitting problem. The A3C model learned the detail and noise in the training data to the extent that it negatively impacted the performance of the model on new data. Our results and conclusions drawn from them are consistent with the results shown when a similar approach was used on Pong (Rusu et al., 2016). In addition to Breakout/A3C,

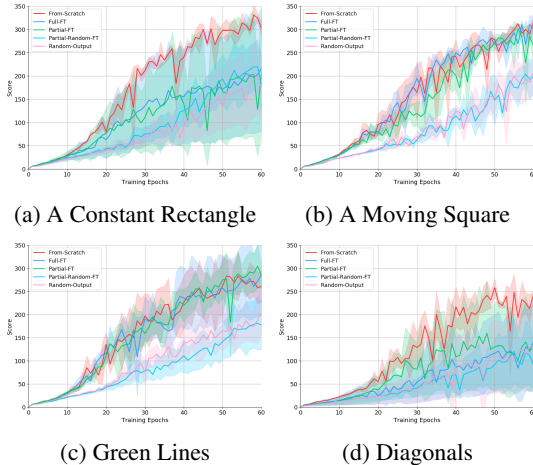


Figure 2. A comparison between the different transfer techniques on Breakout. The y-axis shows the average reward per episode of Breakout during training. The x-axis shows the total number of training epochs where an epoch corresponds to 1 million frames. The plots are averaged on 3 runs with different random seeds. Each curve is the average and its background is the standard deviation.

we also attempted to transfer between a model trained with the synchronous actor-critic variant, A2C, from the first to advanced level of Road Fighter, where the backgrounds change but the dynamics remains the same. This resulted with 0 points on each of the levels, a complete failure of the agent to re-use the driving techniques learned on the first levels on following ones.

3. Analogy-based Zero-Shot Transfer

An agent capable of performing a task in a source domain is now presented with a new domain. Fine-tuning the agent on the target domain fails to transfer knowledge from the source domain. We propose to separate the visual transfer from the dynamics transfer. To perform well, the agent can try and make analogies from the new domain to the old one: after observing a set of states (images) in the new domain, the agent can learn to map them to similar, familiar states from the source domain, and act according to its source domain policy on the mapped state.

More concretely, given a trained policy $\pi(a|s; \theta)$ with trained parameters θ proposing an action a for source domain states $s \in \mathcal{S}$, we wish to learn a mapping function $G : \mathcal{T} \mapsto \mathcal{S}$ from target domain states $t \in \mathcal{T}$ such that interacting with the environment \mathcal{T} by applying the policy $\pi(a|G(t); \theta)$ will result in a good distribution of actions for the states \mathcal{T} , as indicated by high overall scores. In other words, we seek a mapping function G that allows us to re-use the same policy π_θ learned for source environment \mathcal{S} when interacting with the target environment \mathcal{T} .

As both the source and target domain items are images,

we heuristically learn the function G by collecting sets of images from \mathcal{S} and \mathcal{T} and learning to visually map between them using Unaligned GAN (Liu et al., 2017; Zhu et al., 2017; Kim et al., 2017; Yi et al., 2017). We use the scores obtained from interacting with the environment via $\pi(a|G(t); \theta)$ for the GAN model selection and stopping criteria.

Unsupervised Image-to-image Translation In this work, we focus on learning setups that receive only raw image data, without additional domain knowledge about objects or game dynamics. This prohibits us from using supervised paired GANs (Isola et al., 2016) for learning the mapping function G : we cannot collect the needed supervision of corresponding (s, t) pairs. Instead, we use unaligned GANs (Zhu et al., 2017; Liu et al., 2017; Kim et al., 2017; Yi et al., 2017), in which the learner observes two sets of images, one from each domain, with the goal of learning to translate images in one domain to images in another.

All major approaches to the unaligned image-to-image translation use the Cycle-Consistency principle. We have two mapping (encoding) functions $G_1 : T \mapsto S$ and $G_2 : S \mapsto T$ where $S = \{s_i\}_{i=1}^N$ is a set of images collected from the source task and $T = \{t_j\}_{j=1}^M$ is a set of images collected from the target task. The goal is to generate an image s' , for any given $t \in T$ where $G_1(t) = s'$, that is indistinguishable from $s \in S$. The cycle consistency principle relies on the assumption that the two functions, G_1 and G_2 are inverses of each other. It encourages unsupervised mapping by forcing $G_2(G_1(t)) = t$ and $G_1(G_2(s)) = s$ where s and t are the input images. The second component of the GAN architecture are the discriminators D_1 and D_2 aiming to distinguish between images generated by G_1 and G_2 and the real images from the target and source distributions respectively.

In the following experiments, we use the UNIT framework (Liu et al., 2017), which we found to perform well for the Breakout tasks (in section 5.3 we explicitly compare the UNIT and CycleGAN approaches on both the Breakout and Road Fighter transfer tasks).

GAN Training The Unaligned GAN training dataset requires images from both domains. We collect images from the source domain by running an untrained agent and collecting the observed images, and we do similarly for the target domain. We repeat this procedure for every target task, and create a source-target dataset for each.

For our experiments we use the same architecture and hyperparameters proposed in the UNIT paper. We initialize the weights with Xavier initialization (Glorot & Bengio, 2010), set the batch size to 1 and train the network for a different number of iterations on each task.

Concrete GAN Evaluation Criterion. GAN training, and

unaligned GANs training in particular, are unstable and it is challenging to find a good loss-based stopping criteria for them. A major issue with GANs is the lack of an evaluation metric that works well for all models and architectures, and which can assist in model selection. Fortunately, our setup suggests a natural evaluation criteria: we run the source agent without any further training while using the model to translate each image of the target task back to the source task and collect the rewards the agent receives during the game when presented with the translated image. We use the total accumulated rewards (the score) the agent collects during the game as the criteria for the GAN’s model quality, for model selection. In section 5.3 we use this criteria to compare unaligned GAN variants.

4. Imitation Learning

The visual-analogy transfer method allows the agent to re-use its knowledge from the source domain to act in the target domain, resulting in adequate policies. However, it is limited by the imperfect GAN generation and generalization abilities, which, for more challenging visual mappings, may result in sub-optimal policies.

We propose to use the visual-transfer based policy to create *imperfect demonstrations*, and use imitation learning from these demonstrations to further improve the agent, while maintaining better sample efficiency than learning from scratch.

Accelerating RL with Imitation Learning. We combine RL training with Imitation Learning similarly to (Hester et al., 2017; Gao et al., 2018; Oh et al., 2018), but with special considerations for the imperfect demonstration scenario. More specifically, we focus on combining an actor-critic approach with supervised training on imperfect demonstrations. Our process (Algorithm 1) consists of 3 stages. We start by collecting trajectories of the transferred agent interacting with the environment of the target task. The trajectories are collected while following a stochastic policy, to improve coverage. Each trajectory stochastically follows a visually-transferred agent for an entire game. We find that a small number of 5 trajectories is sufficient for obtaining strong results. For each trajectory, in each time step t , we collect the state of the target task s_t , action a_t , and real values R_t as a triple (s_t, a_t, R_t) and store it in a buffer \mathcal{D} . We use \mathcal{D} to train a new agent on the target task by imitating the action and value of each observation from the demonstrations using supervised learning. After a limited number of iterations we switch to RL training, while combining on-policy A2C updates and off-policy supervised updates from the buffer. To account for the imperfect demonstrations, we stop the off-policy updates and move to exclusive on-policy RL training once the agent perform better than the demonstrations.

Algorithm 1 Imitation Learning

Input: a source trained network $\hat{\theta}$, a generator trained with GANs Gen
 Initialize replay buffer $\mathcal{D} \leftarrow \emptyset$, trajectory buffer $\mathcal{T} \leftarrow \emptyset$
// Collecting trajectories
for $i = 1$ to $Trajectories$ **do**
 Get initial state s_0
 repeat
 Execute an action $a_t, r_t, s_{t+1} \sim \pi_{\hat{\theta}}(a_t | Gen(s_t))$
 Store transition $\mathcal{T} \leftarrow \mathcal{T} \cup (s_t, a_t, r_t)$
 $s_t \leftarrow s_{t+1}$
 until s_t is terminal
 if $r_t > \beta_1 R_{\mathcal{T}}$ **then**
 Compute returns $R_t = \sum_k \gamma^{k-t} r_k$
 $\mathcal{D} \leftarrow \mathcal{D} \cup (s_t, a_t, R_t)$ for all t in \mathcal{T}
 end if
 Clear trajectory buffer $\mathcal{T} \leftarrow \emptyset$
end for
// Supervised Training
 Initialize target network weights θ randomly
for $i = 1$ to $Supervised_Iterations$ **do**
 Train on \mathcal{D} with \mathcal{L}_{IL} using SGD and batch size b
end for
// RL with A2C
for $e = 1$ to $Epochs$ **do**
 for $t = 1$ to $Steps$ **do**
 Execute an action $a_t, r_t, s_{t+1} \sim \pi_{\theta}(a_t | s_t)$
 $s_t \leftarrow s_{t+1}$
 end for
 // On-policy updates
 Update θ according to \mathcal{L}_{a2c} using RMSprop
 // Off-policy updates
 if $t \bmod op_interval = 0$ AND $\hat{R} < \beta_2 R_{\mathcal{T}}$ **then**
 Train on \mathcal{D} with \mathcal{L}_{IL} using SGD and batch size b
 end if
end for

The objective function of the off-policy updates before and during the RL training is given by:

$$\mathcal{L}_{IL} = \mathbb{E}_{s,a,R \sim \mathcal{D}} [\mathcal{L}_{IL_{policy}} + \frac{1}{2} \mathcal{L}_{IL_{value}}]$$

$$\mathcal{L}_{IL_{policy}} = \frac{1}{|a|} \sum_{k=0}^{|a|} a_k \log(\hat{a}_k) + (1 - a_k) \log(1 - \hat{a}_k)$$

$$\mathcal{L}_{IL_{value}} = (R - V_{\theta}(s))^2$$

where $\hat{a} = \max_a \pi_{\theta}(a|s)$, $R_t = \sum_k \gamma^{k-t} r_k$ and $\pi_{\theta}(a|s)$, $V_{\theta}(s)$ are the policy and value functions parameterized by θ . For the on-policy updates we use the A2C loss which is given by:

$$\mathcal{L}_{a2c} = \mathbb{E}_{s_t, a_t \sim \pi_{\theta}} [\mathcal{L}_{a2c_{policy}} + \frac{1}{2} \mathcal{L}_{a2c_{value}}]$$

$$\mathcal{L}_{a2c_{policy}} = -\log \pi_{\theta}(a_t | s_t) (V_t^n - V_{\theta}(s_t)) - \alpha \mathcal{H}_t^{\pi_{\theta}}$$

$$\mathcal{L}_{a2c_{value}} = (V_t^n - V_{\theta}(s_t))^2$$

where $V_t^n = \sum_{d=0}^{n-1} \gamma^d r_{t+d} + \gamma^n V_{\theta}(s_{t+n})$ is the n-step bootstrapped value. $\mathcal{H}_t^{\pi_{\theta}} = -\sum_a \pi(a|s_t) \log \pi(a|s_t)$ is the entropy and α is the entropy regularization factor.

Table 1. The score and number of frames needed for it of: the source task (Source), target task when initialized with the source task’ network parameters with no additional training (Target) and the target task when initialized with the source task’ network parameters where every frame is translated to a frame from the source task (Target with GANs).

Source		Target Task	Target		Target with GANs	
Frames	Score		Frames	Score	GAN iterations	Score
43M	302	A Constant Rectangle	0	3	260K	362
43M	302	A Moving Square	0	0	384K	300
43M	302	Green Lines	0	2	288K	300
43M	302	Diagonals	0	0	380K	338




Figure 3. Illustration of a frame taken from the target task (left) and its matching frame of the source task generated with GANs (right) for each one of the Breakout variations. (a)-(d) demonstrate successes, while (e) and (f) show failure modes of the unaligned GAN. In (e) the ball in the input image is not generated in the output and in (f) not all bricks are generated, and some of the generated bricks appear smudged.

Imperfect Demonstrations. While learning from expert demonstrations is a popular technique for learning from good behaviors (Ross et al., 2010; Ho & Ermon, 2016; Torabi et al., 2018), learning from imperfect demonstrations remains a challenging task. We design the algorithm to ensure that the demonstrations benefit and not harm the learning process. First, we collect trajectories by sampling actions rather than following the best ones, for achieving diversity. Stochastic policy leads to more exploration but can also lead to bad trajectories with low scores. We discard trajectories with scores lower than $\beta_1 R_{\mathcal{T}}$ ($\beta_1 = 0.75$), where $R_{\mathcal{T}}$ is the score of a trajectory collected with a deterministic policy (according to Table 2) and β_1 determines how much below the maximum is considered a good trajectory. Second, we combine on-policy and off-policy updates for acceleration and stabilization of the training. Imitating

the demonstration behavior is useful only if it is better than the agent’s current policy. Additionally, to allow substantial improvements over the demonstration the agent may need to deviate from it into other policies that might be worse in order to eventually find a better one. Combining these rationales, we limit the off-policy updates to cases in which the mean reward of all agents \bar{R} is smaller than the $\beta_2 R_{\mathcal{T}}$ ($\beta_2 = 0.6$).¹

5. Experiments

We apply the approach to the Breakout variants (where the visual transfer was sufficient to achieve perfect policy, despite deficiencies in GAN behavior) and to the much more challenging task of transfer between levels of the Road Fighter game, where the visual transfer resulted in an adequate agent, which could then be improved substantially by the imitation learning.

5.1. Breakout

The visual transfer goal in the Breakout variants is removing the visual modifications in each of the target games and mapping back to the unmodified source game. Some tasks turned out to be more challenging to the GAN than others: the Green Lines variation hides parts of the ball in some frames making the training harder, while the Rectangle variation required less training since the number of changed pixels in the image is small.

Overall, the translation tasks—despite their apparent simplicity—proved to be surprisingly challenging for the GANs. While the input and output spaces are highly structured, the network does not have any information about this structure. Instead of learning a “leave game objects intact” policy, it struggles to model and generate them as well. The most common problem was the generator adding or removing blocks from the original image, and in general failing to correctly model block structure (see Fig. 3f and the video). Another issue was with correctly tracking the ball location

¹The off-policies may turn on again if the policy degrades to below the demonstration-based reward.

(Fig. 3e).² Despite these limitations, the learned visual mapping was sufficient for the trained Breakout agent to achieve high scores.

Table 1 shows the results of a test game played by the agent with and without the GAN transfer. The source agent was trained until it reached 300 points, which we consider to be a high score. This required 43M frame interactions. When applied to the target tasks, the agent fails with scores ≤ 3 . As discussed in Section 2, training from scratch for these tasks will require a similar number of frames. With the GAN based transfer the agent achieves scores ≥ 300 , while observing only $100k$ target task frames and performing hundreds of thousands of GAN iterations, a 100x fold increase in sample efficiency.

5.2. Road Fighter



Figure 4. Road Fighter levels from left to right: Level 1, Level 2, Level 3 and Level 4.

While the Breakout variants work well to demonstrate the RL transfer failure cases, we consider them as “toy examples”. We proceed to demonstrate the effectiveness of our transfer method on a “real” task: learning to transfer between different levels of the Road Fighter game. Road Fighter contains 4 different levels (Fig. 4), each with a different background where some are more difficult than others. The levels mostly differ visually and all have the same rules and require the same driving techniques. We train an RL agent to play the first level of the game. To master the game, the agent has to acquire 3 main capabilities: driving fast, avoiding collision with obstacles, and if a car collision occurs reacting fast to avoid crashing. We use the A2C algorithm, the synchronous version of the Advantage Actor-Critic which performs better on this game than A3C, reaching over 10,000 game points on level 1. Training an RL agent to play this level requires observing over 100M frames. The trained agent fails completely on more advanced levels, reaching a score of 0.

For the visual-transfer training, we collect $100k$ frames from each of levels 2, 3 and 4 by running an untrained agent repeatedly until we have collected sufficient samples. Using the collected images we train a mapping function from each

²We believe that this highly structured setting exposes an inherent deficiency in unaligned GAN models, suggesting an interesting avenue for future GAN research.

new level (target task) to the first one (source task). We use the same GAN architecture used for Breakout, but initialize the weights with Orthogonal initialization. Compared to Breakout, these tasks introduce new challenges: rather than removing a mostly static element, the GAN has to be able to change the background and road size while keeping the cars in the right position on the road. On the other hand, this setup may be closer to the one unaligned GANs are usually applied on. We restrict ourselves to collecting images from the beginning of the game, before the agent had any training. This restricts the phenomena the GAN can observe, leading to some target tasks’ images without a clear corresponding situation in the first level, potentially causing unpredictable behaviors. For example, the generator matches the diagonal shaped roads to one of the first rare and unique images of level 1 (Fig. 5e).

Data Efficiency. We measure the number of frames of game-interaction needed for the analogy-transfer method. We collect 100k frames, and then train the GAN for up to $500k$ iterations, evaluating it every 10,000 iterations by running the game and observing the score, and pick the best scoring model. This amounts to $100k + 50 * F$ frames, where $F = 3000$ is roughly the average number of frames in a game. This amounts to about $250k$ frames of game interaction for each transferred level, an order of magnitude fewer interaction frames than training an RL agent to achieve a comparable score.

Results. The results in table 2 show that the visual transfer manages to achieve scores of 5350, 5350 and 2050 on level 2, 3 and 4 after $320k$, $450k$ and $270k$ GAN iterations respectively, while performing only a fraction of the game interactions required to train an RL agent from scratch to achieve these scores on these levels.

Qualitatively, the agent applies many of the abilities it gained when training on the first level, most notably driving fast, staying on the road, avoiding some cars, and, most importantly, recovering from car crashes.

Limitations of purely-visual transfer. While the GAN works well in generating objects it has seen before, such as the agent’s car, it does have limitations on objects it has not seen. As a result, it ends up generating differently colored cars all in red, or not generating them at all, as shown in Fig. 5a, 5d and 5f. Colorful cars can be “collected” by the agent and are worth 1000 points each. Generating them in red makes the agent avoid them, losing these extra points and achieving overall lower scores even if finishing the track. When cars are not fully generated, the agent is less likely to avoid them, and eventually crashes.

In general, as the game progresses, the more obstacles are presented making it harder to play. On level 3, the challenge increases as the road is very narrow, making it harder for

Table 2. The scores of the agent on every level of Road Fighter with and without analogy-transfer and imitation learning, as well as the number of game-interaction frames needed for the analogy transfer and for achieving a similar score with RL training, and number of frames needed for an RL agent to achieve over 10,000 points with and without imitation learning.

	Score (no transfer)	Score (analogy transfer)	# Frames (analogy)	# Frames (from scratch)	Score (+imitation)	#Frames (imitation)	# Frames (from scratch)
Level 2	0	5350	250K	12.4M	10230	38.6M	159M
Level 3	0	5350	250K	31M	10300	21M	54.4M
Level 4	0	2050	250K	13.6M	10460	13.4M	111M

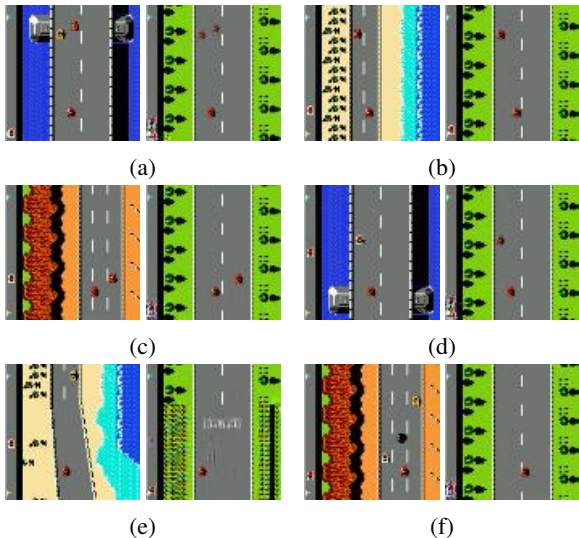


Figure 5. Left: the original frame. Right: GAN generated. The first three examples show the success cases of the GAN while the last three show representative failures: in (d) and (f) the only object generated on the road is the player’s car and in (e) the diagonal shaped road of level 2 in matched to the starting point of level 1.

a player to avoid crashing. However, the GAN manages to generate the road in the right shape in most frames and position the cars in the matching ratio. Level 4 gets the lowest score. In this level the main obstacles are the sudden turns in the road causing the car to be very close to the sideways and the increasing dangers a player has to avoid. These difficulties make this level much harder than level 1 and might require more training even from a human player.

Improving with Imitation Learning. Using zero-shot visual transfer, the agent manages to apply 2 out of 3 of the main capabilities it gained when training on level 1. It is missing the third capability: avoiding collisions and collecting bonus cars, mainly because of bad generation. It is also challenged by winding roads (levels 3 and 4), which are not available on level 1. We now train an RL agent to play the target levels by imitating the imperfect demonstration of the visual transfer policy (see Section 4).

The graphs in Figure 6 demonstrates that the combination of visual transfer and imitation learning works well, surpassing training from scratch on the target tasks in terms of both

sample efficiency *and* accuracy. The final agents clear levels 2, 3 and 4 with scores of over 10,000 points on each.

5.3. Towards Task-oriented GAN evaluation

Evaluating GAN models and comparing them to each other is a major challenge. Our setup introduces a natural, measurable objective for unaligned GAN evaluation: using them for visual transfer of an RL agent, and measuring the agent’s performance on the translated environment. We use the approach to compare Cycle-GAN and UNIT-GAN with somewhat mixed results: UNIT works better for breakout, while the methods are mostly tied for Road Fighter, with some advantage to Cycle-GAN (exact numbers are available in the appendix, Table 6). The main difference between the two methods is the weight-sharing constraint applied in UNIT, making the domains dependent on each other by sharing and updating the weights of one or several decoders and encoders layers. We hypothesize this constraint is an advantage in tasks where the representation of the images in the different domains are similar, such as the Breakout variants. In contrast, the more distinct Road Fighter levels could benefit from the independence in Cycle-GAN.

6. Related Work

Transfer Learning (TL) is a machine learning technique used to improve the training speed of a target task with knowledge learned in a source task. Pretraining and fine-tuning was proposed in (Hinton & Salakhutdinov, 2006) and applied to TL in (Bengio, 2012) and (Dauphin et al., 2012). In this procedure, the approach is to train the base network and then copy its first n layers to the first n layers of a target network. One can choose to update the feature layers transferred to the new task with the error backpropagated from its output, or they can be left frozen, meaning that they do not change during training on the new task. Unfortunately, as we have shown, while fine-tuning might have the ability to accelerate the training process in some cases, it can also have a damaging impact on others.

Generalization is a key element in training deep learning models with time or data size constraints. Recent discussions on overfitting in Deep RL algorithms (Zhang et al.,




Figure 6. Learning curves of the agent trained to play levels 2 (a), 3 (b) and 4 (c) of Road Fighter from scratch using A2C, and using our A2C with Imitation Learning method. Each point represents the average reward of 84 agents every 100 RL updates.

2018) encouraged better evaluation (e.g. OpenAI Retro Contest³) and generalization methods. In Atari, there are many similarities between the goals and the mechanism of the games. For this reason, there have been many works attempting to transfer between games or between different variations of the same game, one approach trying to do both is the *progressive networks* (Rusu et al., 2016). A progressive network is constructed by successively training copies of A3C on each task of interest. In that work, they transferred between different games as well as from different variations of the game Pong. The drawback of this approach is the number of parameters growing quadratically with the number of tasks.

Zero-shot generalization is a popular research topic. In the context of games, (Kansky et al., 2017) performs zero-shot transfer between modified versions of the same game using schema networks. Like us, they also demonstrate their method on the game Breakout, using Object Oriented Markov Decision Process. In contrast, we do not use the representation of the objects in the game, and wish to preserve the accomplishments of DQN and transfer using only raw data. Other attempted to achieve robust policies using learned disentangled representation of the image (Higgins et al., 2017), analogies between sets of instructions (Oh et al., 2017), *interactive replay* (Bruce et al., 2017) while training and learn general policies by training on multiple tasks in parallel (Espeholt et al., 2018; Sohn et al., 2018).

Finally, the idea of using GANs for transfer learning and domain adaptation was explored for supervised image classification and robotics applications by several authors (Bousmalis et al., 2016; Hoffman et al., 2017; Liu & Tuzel, 2016; Bousmalis et al., 2017). Our work is the first to combine it with imitation learning and successfully demonstrate it in an RL setup.

Imitation learning is the paradigm of learning from demonstrations by imitating their behaviours. Combining it with RL seems natural, while RL provides more exploration of new states, imitation improves RL by providing prior knowledge to exploit. Recent works has shown a success of this

combination in a difficult setting of learning from imperfect demonstrations. DQfD (Hester et al., 2017) and SIL (Oh et al., 2018) merge temporal difference and imitation losses for training and prioritize better demonstrations by choosing the ones that are most likely to improve the current behaviour. In contrast, NAC (Gao et al., 2018) is an RL algorithm that uses a unified actor-critic objective function that is capable of resisting poor behaviours. Our approach is similar to these works. However, we prioritize good behaviours from the start, by selecting trajectories with the highest scores. Moreover, we use a separate supervised loss function for imitation and for RL and train our agent with both as long as it benefits the learning.

7. Conclusions

We demonstrated the lack of generalization by looking at artificially constructed visual variants of a game (Breakout), and different levels of a game (Road Fighter). We further show that transfer learning by fine-tuning fails. The policies learned using model-free RL algorithms on the original game are not directly transferred to the modified games even when the changes are irrelevant to the game’s dynamics.

We present a new approach for transfer learning between related RL environments using GANs without the need for any additional training of the RL agent, and while requiring orders of magnitude less interactions with the environment. We further suggest this setup as a way to evaluate GAN architectures by observing their behavior on concrete tasks, revealing differences between the Cycle-GAN and UNIT-GAN architectures. While we report a success in analogy transfer using Unaligned GANs, we also encountered limitations in the generation process that made it difficult for the agent to maximize the results on the Road Fighter’s tasks. We overcome these difficulties by using the imperfect behaviours as demonstrations to improve and accelerate the RL training on each one of the target tasks. We believe that these tasks and results demonstrate a success of the analogy transfer method across different levels of a video game. They also suggest a potential of performing well on additional real world tasks in which visual analogies can be made.

³<https://contest.openai.com/>

Acknowledgements

We thank Hal Daum III for the helpful discussions on the Imitation Learning algorithm during the development of the work.

References

- Bengio, Y. Deep learning of representations for unsupervised and transfer learning. In Guyon, I., Dror, G., Lemaire, V., Taylor, G., and Silver, D. (eds.), *Proceedings of ICML Workshop on Unsupervised and Transfer Learning*, volume 27 of *Proceedings of Machine Learning Research*, pp. 17–36, Bellevue, Washington, USA, 02 Jul 2012. PMLR. URL <http://proceedings.mlr.press/v27/bengio12a.html>.
- Bousmalis, K., Silberman, N., Dohan, D., Erhan, D., and Krishnan, D. Unsupervised pixel-level domain adaptation with generative adversarial networks. *CoRR*, abs/1612.05424, 2016. URL <http://arxiv.org/abs/1612.05424>.
- Bousmalis, K., Irpan, A., Wohlhart, P., Bai, Y., Kelcey, M., Kalakrishnan, M., Downs, L., Ibarz, J., Pastor, P., Konolige, K., Levine, S., and Vanhoucke, V. Using simulation and domain adaptation to improve efficiency of deep robotic grasping. *CoRR*, abs/1709.07857, 2017. URL <http://arxiv.org/abs/1709.07857>.
- Bruce, J., Sünderhauf, N., Mirowski, P., Hadsell, R., and Milford, M. One-shot reinforcement learning for robot navigation with interactive replay. *CoRR*, abs/1711.10137, 2017. URL <http://arxiv.org/abs/1711.10137>.
- Dauphin, G. M. Y., Glorot, X., Rifai, S., Bengio, Y., Goodfellow, I., Lavoie, E., Muller, X., Desjardins, G., Warde-Farley, D., Vincent, P., Courville, A., and Bergstra, J. Unsupervised and transfer learning challenge: a deep learning approach. In Guyon, I., Dror, G., Lemaire, V., Taylor, G., and Silver, D. (eds.), *Proceedings of ICML Workshop on Unsupervised and Transfer Learning*, volume 27 of *Proceedings of Machine Learning Research*, pp. 97–110, Bellevue, Washington, USA, 02 Jul 2012. PMLR. URL <http://proceedings.mlr.press/v27/mesnil12a.html>.
- Espeholt, L., Soyer, H., Munos, R., Simonyan, K., Mnih, V., Ward, T., Doron, Y., Firoiu, V., Harley, T., Dunning, I., Legg, S., and Kavukcuoglu, K. IMPALA: scalable distributed deep-rl with importance weighted actor-learner architectures. *CoRR*, abs/1802.01561, 2018. URL <http://arxiv.org/abs/1802.01561>.
- Fernando, C., Banarse, D., Blundell, C., Zwols, Y., Ha, D., Rusu, A. A., Pritzel, A., and Wierstra, D. Pathnet: Evolution channels gradient descent in super neural networks. *CoRR*, abs/1701.08734, 2017. URL <http://arxiv.org/abs/1701.08734>.
- Gao, Y., Xu, H., Lin, J., Yu, F., Levine, S., and Darrell, T. Reinforcement learning from imperfect demonstrations. *CoRR*, abs/1802.05313, 2018. URL <http://arxiv.org/abs/1802.05313>.
- Glorot, X. and Bengio, Y. Understanding the difficulty of training deep feedforward neural networks. In *In Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS10)*. Society for Artificial Intelligence and Statistics, 2010.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N. D., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems 27*, pp. 2672–2680. Curran Associates, Inc., 2014. URL <http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>.
- Hester, T., Vecerik, M., Pietquin, O., Lanctot, M., Schaul, T., Piot, B., Sendonaris, A., Dulac-Arnold, G., Osband, I., Agapiou, J., Leibo, J. Z., and Gruslys, A. Learning from demonstrations for real world reinforcement learning. *CoRR*, abs/1704.03732, 2017. URL <http://arxiv.org/abs/1704.03732>.
- Higgins, I., Pal, A., Rusu, A. A., Matthey, L., Burgess, C., Pritzel, A., Botvinick, M. M., Blundell, C., and Lerchner, A. Darla: Improving zero-shot transfer in reinforcement learning. In *ICML*, 2017.
- Hinton, G. E. and Salakhutdinov, R. R. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006. ISSN 0036-8075. doi: 10.1126/science.1127647. URL <http://science.sciencemag.org/content/313/5786/504>.
- Ho, J. and Ermon, S. Generative adversarial imitation learning. *CoRR*, abs/1606.03476, 2016. URL <http://arxiv.org/abs/1606.03476>.
- Hoffman, J., Tzeng, E., Park, T., Zhu, J., Isola, P., Saenko, K., Efros, A. A., and Darrell, T. Cycada: Cycle-consistent adversarial domain adaptation. *CoRR*, abs/1711.03213, 2017. URL <http://arxiv.org/abs/1711.03213>.
- Isola, P., Zhu, J., Zhou, T., and Efros, A. A. Image-to-image translation with conditional adversarial networks. *CoRR*, abs/1611.07004, 2016. URL <http://arxiv.org/abs/1611.07004>.

- Kansky, K., Silver, T., Mély, D. A., Eldawy, M., Lázaro-Gredilla, M., Lou, X., Dorfman, N., Sidor, S., Phoenix, D. S., and George, D. Schema networks: Zero-shot transfer with a generative causal model of intuitive physics. *CoRR*, abs/1706.04317, 2017. URL <http://arxiv.org/abs/1706.04317>.
- Kim, T., Cha, M., Kim, H., Lee, J. K., and Kim, J. Learning to discover cross-domain relations with generative adversarial networks. *CoRR*, abs/1703.05192, 2017. URL <http://arxiv.org/abs/1703.05192>.
- Kostrikov, I. Pytorch implementations of asynchronous advantage actor critic. <https://github.com/ikostrikov/pytorch-a3c>, 2018a.
- Kostrikov, I. Pytorch implementations of reinforcement learning algorithms. <https://github.com/ikostrikov/pytorch-a2c-ppo-acktr>, 2018b.
- Liu, M. and Tuzel, O. Coupled generative adversarial networks. *CoRR*, abs/1606.07536, 2016. URL <http://arxiv.org/abs/1606.07536>.
- Liu, M., Breuel, T., and Kautz, J. Unsupervised image-to-image translation networks. *CoRR*, abs/1703.00848, 2017. URL <http://arxiv.org/abs/1703.00848>.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. A. Playing atari with deep reinforcement learning. *CoRR*, abs/1312.5602, 2013. URL <http://arxiv.org/abs/1312.5602>.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T. P., Harley, T., Silver, D., and Kavukcuoglu, K. Asynchronous methods for deep reinforcement learning. *CoRR*, abs/1602.01783, 2016. URL <http://arxiv.org/abs/1602.01783>.
- Oh, J., Singh, S. P., Lee, H., and Kohli, P. Zero-shot task generalization with multi-task deep reinforcement learning. *CoRR*, abs/1706.05064, 2017. URL <http://arxiv.org/abs/1706.05064>.
- Oh, J., Guo, Y., Singh, S., and Lee, H. Self-imitation learning. *CoRR*, abs/1806.05635, 2018. URL <http://arxiv.org/abs/1806.05635>.
- Ross, S., Gordon, G. J., and Bagnell, J. A. No-regret reductions for imitation learning and structured prediction. *CoRR*, abs/1011.0686, 2010. URL <http://arxiv.org/abs/1011.0686>.
- Rusu, A. A., Rabinowitz, N. C., Desjardins, G., Soyer, H., Kirkpatrick, J., Kavukcuoglu, K., Pascanu, R., and Hadsell, R. Progressive neural networks. *CoRR*, abs/1606.04671, 2016. URL <http://arxiv.org/abs/1606.04671>.
- Schaul, T., Quan, J., Antonoglou, I., and Silver, D. Prioritized experience replay. *CoRR*, abs/1511.05952, 2015. URL <http://arxiv.org/abs/1511.05952>.
- Sohn, S., Oh, J., and Lee, H. Multitask reinforcement learning for zero-shot generalization with subtask dependencies. *CoRR*, abs/1807.07665, 2018.
- Torabi, F., Warnell, G., and Stone, P. Behavioral cloning from observation. *CoRR*, abs/1805.01954, 2018. URL <http://arxiv.org/abs/1805.01954>.
- Yi, Z., Zhang, H., Tan, P., and Gong, M. Dualgan: Unsupervised dual learning for image-to-image translation. *CoRR*, abs/1704.02510, 2017. URL <http://arxiv.org/abs/1704.02510>.
- Yosinski, J., Clune, J., Bengio, Y., and Lipson, H. How transferable are features in deep neural networks? In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2, NIPS'14*, pp. 3320–3328, Cambridge, MA, USA, 2014. MIT Press. URL <http://dl.acm.org/citation.cfm?id=2969033.2969197>.
- Zhang, C., Vinyals, O., Munos, R., and Bengio, S. A study on overfitting in deep reinforcement learning. *CoRR*, abs/1804.06893, 2018. URL <http://arxiv.org/abs/1804.06893>.
- Zhu, J., Park, T., Isola, P., and Efros, A. A. Unpaired image-to-image translation using cycle-consistent adversarial networks. *CoRR*, abs/1703.10593, 2017. URL <http://arxiv.org/abs/1703.10593>.