

Snooping Attacks on Deep Reinforcement Learning

Matthew Inkawhich
Duke University
Durham, North Carolina
matthew.inkawhich@duke.edu

Yiran Chen
Duke University
Durham, North Carolina
yiran.chen@duke.edu

Hai Li
Duke University
Durham, North Carolina
hai.li@duke.edu

ABSTRACT

Adversarial attacks have exposed a significant security vulnerability in state-of-the-art machine learning models. Among these models include deep reinforcement learning agents. The existing methods for attacking reinforcement learning agents assume the adversary either has access to the target agent’s learned parameters or the environment that the agent interacts with. In this work, we propose a new class of threat models, called snooping threat models, that are unique to reinforcement learning. In these snooping threat models, the adversary does not have the ability to interact with the target agent’s environment, and can only eavesdrop on the action and reward signals being exchanged between agent and environment. We show that adversaries operating in these highly constrained threat models can still launch devastating attacks against the target agent by training proxy models on related tasks and leveraging the transferability of adversarial examples.

KEYWORDS

deep reinforcement learning; machine learning; security

ACM Reference Format:

Matthew Inkawhich, Yiran Chen, and Hai Li. 2020. Snooping Attacks on Deep Reinforcement Learning. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), Auckland, New Zealand, May 9–13, 2020*, IFAAMAS, 9 pages.

1 INTRODUCTION

The deep learning revolution has put neural networks at the forefront of the machine learning research and production landscape. Recently, it has been shown that deep neural networks (DNNs) are effective function approximators for solving complex reinforcement learning (RL) problems. Mnih et al. [21] demonstrated that we can leverage the feature extraction capabilities of convolutional neural networks to enable RL agents to learn to play Atari games from raw pixels. Since then, there has been an influx of work dedicated to the field of deep RL (DRL), enabling powerful new solutions for tasks such as game playing [20, 21, 32, 34], continuous control [17, 33, 35], complex robot manipulation [16, 28, 30], and even autonomous vehicle operation [4, 31].

Despite the numerous successes of the applications of deep neural networks, these models have a significant Achilles’ heel. Studies show that these largely uninterpretable models are vulnerable to small adversarial perturbations to the input that cause the models to perform poorly [10, 37]. Recent work has uncovered a variety of adversarial attack algorithms in the domain of image classification

[6, 8, 10, 26]. Alarmingly, results show that adversarial examples effective against one model are also effective against other models trained on the same task [10]. This transferability property has enabled attacks against models in black-box settings [13, 25].

As of now, adversarial attacks against DRL agents primarily consider the white-box threat model, in which the adversary is assumed to have access to the target model’s architecture and parameters. To our knowledge, the study of black-box attacks against DRL has been limited to the acknowledgement of the transferability of adversarial examples across policies and algorithms [12]. Thus, the current method for black-box attacks against DRL agents involves training a surrogate agent, which makes the rather unrealistic assumption that the adversary has unbounded access to the environment. In general, the black-box threat model in RL is more complex than it is in supervised learning. First, in the data generating process for supervised learning the model draws i.i.d. samples from a static dataset. In RL, data generation is in the form of an environment from which new data samples are conditioned on the current environment state and the agent’s actions. This means that an RL adversary cannot trivially approximate the training data generating process the way a supervised learning adversary can. Also, the RL problem consists of an exchange of three signals between the environment and agent (state, action, reward). Even if an RL adversary seeks to approximate the agent or environment using some other method (e.g., inverse RL), they would need access to all signals.

In this work, we investigate threat models for DRL under the assumption that the adversary does not have access to the environment that the target agent interacts with. Specifically, we explore the potential effectiveness of an adversary that only has the ability to eavesdrop on a subset of the RL signals at each time step. We dub these the snooping threat models. For a real-world example of the snooping threat models of a vision-based DRL system, consider a camera-equipped mobile system such as an autonomous vehicle. These systems typically send the camera feed via a wireless network to a server running an instance of the trained DRL agent, which processes the images and decides how to act in the given state. As with all current adversarial threat models, the adversary is assumed to have the ability to alter the input of the model. Action signals are transmitted from the server back to the drone via a separate channel. Finally, reward signals may come from a variety of sources. Given this scenario, it is entirely possible that the adversary cannot access or eavesdrop on all signals due to private channels, encryption, etc.

We show that by training proxy models on tasks *similar to the target agent’s task*, we are able to craft adversarial inputs that significantly reduce the performance of target agents trained on various Atari environments. We feel this to be a more practical threat model for real-world systems. Overall, the contributions of this paper can be summarized as follows:

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand. © 2020 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

- We define the complex threat model taxonomy for snooping attacks on DRL agents.
- We propose the use of proxy models for launching effective adversarial attacks under the different snooping threat models.
- We conduct extensive experiments using the attacks against the state-of-the-art DRL algorithms DQN [21] and PPO [33].
- We empirically show that adversarial examples transfer between models that are trained with very different objectives as long as the tasks are related.

2 RELATED WORK

Adversarial machine learning is a research area that investigates the vulnerabilities of deep neural networks to small adversarial perturbations. In this work, we focus on evasion attacks, which are test-time attacks designed to fool a trained model. This topic was popularized in the image classification domain in the white-box setting, where Szegedy et al. [37] discovered that adding imperceptible noise to an image causes the model to predict the incorrect class with high confidence. Goodfellow et al. [10] followed by suggesting a fast gradient method for crafting adversarial examples in an inexpensive way by assuming local linearity in the target model. Goodfellow et al. [10] also observed that adversarial examples crafted using the gradients of one model transfer to other models trained on the same task. Papernot et al. [25] leveraged this concept of transferability and introduced a framework for black-box attacks that uses the target model as an oracle to train a surrogate model with a similar decision surface, and uses the surrogate to craft adversarial examples that also fool the target.

In the DRL domain, Huang et al. [12] showed that the fast gradient method [10] can be extended to fool agents on the Atari benchmark. Lin et al. [18] introduced a strategically-timed attack, which seeks to make perturbations sparse in time to render them more difficult to detect. Lin et al. [18] further presented an enchanting attack that strives to lure the target agent into a specified state. Behzadan and Munir [1] demonstrated that adversarial examples can be used during training to corrupt the learning process of an agent on the Atari Pong environment. Adversarial examples have also been used for hardening DRL models to environmental parameter variations in continuous control environments [29]. Overall, adversarial attacks on DRL agents have primarily been limited to the white-box setting.

3 PRELIMINARIES

3.1 Adversarial example crafting

The goal of adversarial example crafting is to apply an imperceptible perturbation to a benign input such that the perturbed input fools the target model. In this work, we consider variants of the Fast Gradient Method (FGM) [10] because it is computationally efficient and generalizes to models other than classifiers. FGM is a one-step method that assumes linearity of the decision surface around a given sample. With this assumption, the optimal perturbation for an input x is in the direction to maximize the loss J . We consider constraining the perturbations under the L_∞ and L_2 norm bounds.

Under the L_∞ norm bound $\|x^* - x\|_\infty \leq \epsilon$, an adversarial example x^* is generated as

$$x^* = x + \epsilon * \text{sign}(\nabla_x J(x, y)) \quad (1)$$

where y is the label. Under the L_2 norm bound $\|x^* - x\|_2 \leq \epsilon$, x^* is generated as

$$x^* = x + \epsilon * \frac{\nabla_x J(x, y)}{\|\nabla_x J(x, y)\|_2}. \quad (2)$$

We also consider a momentum iterative variant of FGM (MIFGM) introduced by Dong et al. [8]. This variant disregards the linearity assumption and iteratively perturbs while accumulating a velocity vector in the direction of the gradient to stabilize the perturbation direction and combat overfitting. The velocity vector g is accumulated as

$$g_{t+1} = \mu * g_t + \frac{\nabla_x J(x_t^*, y)}{\|\nabla_x J(x_t^*, y)\|_1}. \quad (3)$$

To craft the example, at each iteration we use Equations 1 and 2, but substitute $\nabla_x J(x, y)$ with g_{t+1} . A visualization of L_∞ and L_2 constrained FGM perturbations on Pong are shown in Figure 1.




Figure 1: Adversarial example comparison between FGM attacks under different norm constraints on the Pong environment. The L_∞ bounded example was crafted with $\epsilon = 0.03$ and the L_2 bounded example used $\epsilon = 2.4$. These examples were crafted using the psychic proxy model (see 5.2).

3.2 Reinforcement learning

Reinforcement Learning is a machine learning paradigm based on sequential interactions between an agent and an environment in which the agent attempts to learn a policy π to maximize a total reward [36]. This interaction is formalized as a Markov Decision Process (MDP), which is the tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma \rangle$. \mathcal{S} is a finite set of states, \mathcal{A} is a finite set of actions, $\mathcal{P}(s, a, s') = P(S_{t+1} = s' | S_t = s, A_t = a)$ is the state transition function, $\mathcal{R}(s, a) = \mathbb{E}[R_{t+1} | S_t = s, A_t = a]$ is the reward function and $\gamma \in [0, 1]$ is a discount factor. The interaction can be described as follows: At each time step t , the environment sends a state s_t to the agent. The agent decides on an action a_t and dispatches it back to the environment. The environment responds by sending a reward r_{t+1} to the agent along with the next state s_{t+1} , and the process repeats. In this work, we consider the state-of-the-art value-based DRL method DQN [21] and policy-based PPO [33].

3.2.1 Deep Q-Networks. Q-learning is a value-based RL algorithm that estimates the cumulative discounted reward of each state-action pair, and chooses the agent’s policy based on these estimated returns. In Q-learning, we approximate the function $Q^*(s, a)$,

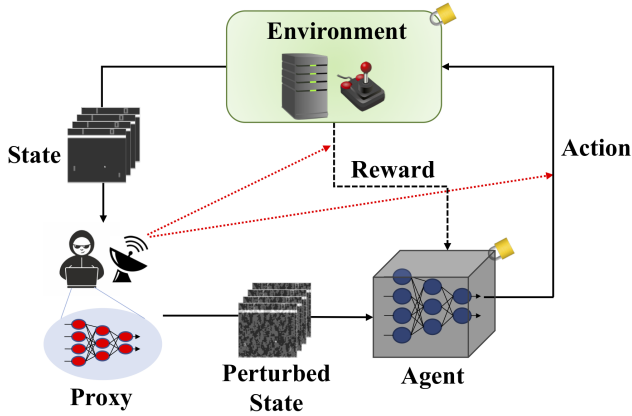


Figure 2: Snooping threat model overview for DRL.

which gives the cumulative discounted reward for taking action a in state s and following an optimal policy π^* thereafter [40]. To train, we iteratively optimize with a temporal difference loss δ based on the Bellman equation [3]:

$$\delta = Q(s, a) - (r + \gamma \max_a Q(s', a)). \quad (4)$$

Once we have an acceptable approximation of Q^* , a common policy is to act ϵ -greedily to aid in exploration. In an effort to apply Q-learning to problems with large state spaces such as Atari games, Mnih et al. [21] combine convolutional neural networks with Q-learning to create deep Q-learning. In practice, this requires the use of a replay memory buffer to enable off-policy learning and a target network for more stable updates.

3.2.2 Proximal policy optimization. PPO [33] is a policy-based RL method, meaning that it directly approximates the policy $\pi_\theta(a|s)$ with a neural network with parameters θ . The algorithm is typically implemented in an actor-critic framework which simultaneously learns to approximate $\pi_\theta(a|s)$ and a variance-reduced advantage estimate to stabilize the policy gradient. At each training step, PPO alternates between (1) collecting experience by running the current policy for a set number of time steps, (2) computing empirical returns and advantages, and (3) using batch learning to optimize a clipped surrogate objective that constrains the amount the updated policy can differ from the old policy:

$$L^{CLIP}(\theta) = \mathbb{E}_t \left[\min(r_t(\theta)\hat{A}_t, \text{clip}(r_t(\theta), 1 - \alpha, 1 + \alpha)\hat{A}_t) \right] \quad (5)$$

where α is a hyperparameter, \hat{A}_t is the empirical advantage, and $r_t(\theta)$ is the probability ratio

$$r_t(\theta) = \frac{\pi_\theta(a_t, s_t)}{\pi_{\theta_{old}}(a_t, s_t)}. \quad (6)$$

4 SNOOPING THREAT MODELS

A high level view of the snooping threat model for DRL is portrayed in Figure 2. As in any threat model, we assume the adversary can intercept and manipulate the input game frames (i.e., the state). In the snooping family of threat models, the adversary does not have direct access to the environment, so training a surrogate DRL

model and transferring examples from the surrogate to the target is ruled out. Instead, the adversary can only "snoop", or eavesdrop, on the natural interactions between the agent and environment. We define the **SRA** threat model to represent the scenario in which the adversary can snoop on reward and action signals, **SA** for when the adversary can only snoop on the actions, **SR** for when the adversary can only snoop on the rewards, and **S** for when the adversary only has access to the states. When the reward signal is hidden, the adversary is assumed to have no inclination about the goal of the agent. In the general case of game environments, this means that the adversary cannot explicitly infer the difference between winning and losing. When the action signal is hidden, not only is the adversary unable to determine which actions correspond to certain changes in the state space, but there is also no way of explicitly knowing the scope and magnitude of the action space at all.

5 ATTACK STRATEGIES

5.1 Methodology

Our general attack methodology is to leverage the transferability of adversarial examples. While this concept is not fully understood, recent efforts proclaim that transferability is primarily enabled by decision boundary similarity between models [19, 37, 38]. We argue that this condition, however, generally overestimates the resources necessary to craft effective transferable adversarial examples. The snooping threat models for RL provides one such counterexample, in which the adversary cannot train a surrogate DRL model that closely resembles the target agent. Instead, we posit that if the adversary can train a proxy model that learns a task that is *related* to the target agent's policy to maximize reward, adversarial examples crafted to fool the proxy will also fool the agent. This implies that even proxy models with radically different decision surface dimensionality than the target can reliably produce transferable attacks. We hypothesize that by learning a related task, the proxy will learn to extract and use input features in a similar fashion to the target agent, which is the only necessary condition for crafting transferable attacks. Our explicit goal is to train a proxy model \mathcal{M} under a constraint subset $\kappa \in \{S, SR, SA, SRA\}$ parameterized by $\theta_{\mathcal{M}_\kappa}$ with loss $J_{\mathcal{M}_\kappa}$ that extracts similar input features and learns comparable implications of these features to the target agent \mathcal{T} parameterized by $\theta_{\mathcal{T}}$ trained with loss $J_{\mathcal{T}}$ (e.g., temporal difference error). Using input gradients as an indication of feature saliency, our intention is to train a constrained proxy \mathcal{M}_κ to optimize:

$$\arg \min_{\mathcal{M}_\kappa} \left[\frac{1}{N} \sum_{i=1}^N \left\| \nabla_x J_{\mathcal{M}_\kappa}(\theta_{\mathcal{M}_\kappa}, x_i, y_i) - c \nabla_x J_{\mathcal{T}}(\theta_{\mathcal{T}}, x_i, z_i) \right\| \right] \quad (7)$$

where N is an arbitrary number of inputs, x_i is an input, y_i is the true label for the proxy, z_i is the pseudo-true label [12] for the agent, and c is an unknown constant. We include c as a scaling constant as we desire \mathcal{M}_κ to produce similarly *shaped* input gradients to \mathcal{T} 's, meaning the magnitude of the gradients do not necessarily have to align exactly. Note that this objective is a theoretical representation, and does not represent a tractable optimization objective. The exact choice of task and proxy \mathcal{M}_κ is determined by our discretion.

Also, note that since we are attacking RL agents during test time, we will not have access to ground truth labels required by the FGM

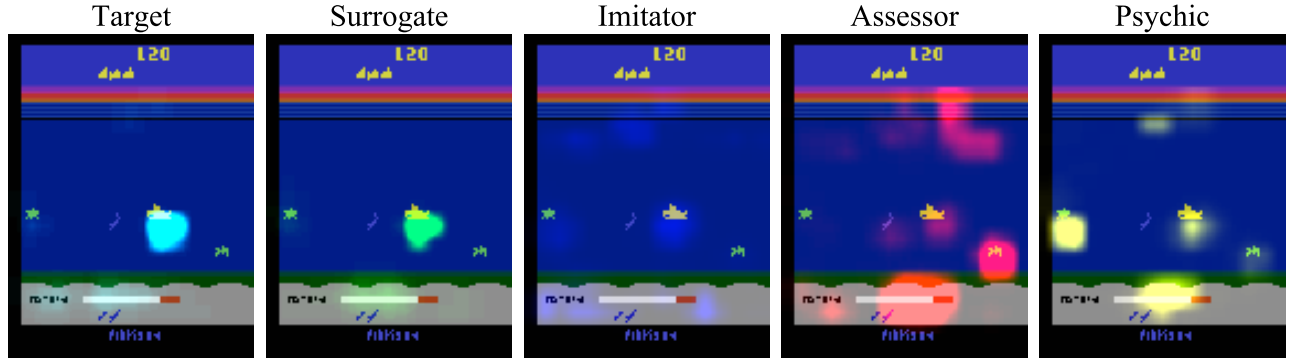


Figure 3: Comparison of features learned by target agent, surrogate agent, and proxy models on Seaquest environment using perturbation-based saliency map visualizations [11].

method (i.e. the y in Equations 1 and 2), as RL data is inherently unlabeled. To work around this, we follow [12] and assume that our surrogates and proxies have converged to a reasonable minima, so we use their outputs as "truth". For surrogate agents or classifier proxies, we use the one-hot output of the model as truth, and for regressor proxies, we use the output (plus a small constant)¹ as truth.

5.2 S threat model

The **S** threat model is the weakest snooping threat model. The adversary only has access to a stream of states, and has no inclination about the motivations or capabilities of the agent. In this case, we propose a proxy \mathcal{M}_S that models the environment dynamics. Previous work shows that we can perform accurate next-frame prediction using the current state, action, and an L_2 reconstruction loss as long as \mathcal{P} is roughly deterministic [15, 23, 39]. This model directly approximates the state transition distribution $\mathcal{P}(s, a, s')$. In **S** however, we do not have access to \mathcal{A} , so the best we can do is approximate an expectation of \mathcal{P} under the target agent’s policy π_T using a *psychic* proxy model:

$$psychic(s_t, \theta_P) \approx \mathbb{E}_{\pi_T} [P(s_{t+1}|s_t)] = \mathbb{E}_{a_t \sim \pi_T} [P(s_{t+1}|s_t, a_t)]. \quad (8)$$

Because the psychic learns the dynamics \mathcal{P} under π_T , it effectively learns an encoding of π_T in the state space. In other words, the displacement of the agent’s representation in the state space (e.g. the pixels of the paddle in Breakout) from s_t to s_{t+1} inherently encodes a representation of the action a_t taken and therefore an instance of the agent’s policy $\pi_T(s_t)$. This approximation will be noisy in general due to displacements in the state space that are not the consequence of the agent’s actions (e.g. the movement of the fish in Seaquest). Despite the noise, the key is that a subset of the features learned by the psychic will be similar to the features learned by the target agent.

5.3 SR threat model

In the **SR** threat model, we can intercept the states and snoop on the rewards. Because we do not have access to the actions, we again cannot make any presumptions about the policy π_T directly. However, by observing the rewards resulting from each state in a

¹We must add a small constant to avoid computing a loss of zero.

rollout we can get insight into the agent’s motivations by training an *assessor* proxy to estimate the value V of a given state under π_T . Since the reward signal is what drives the policy π_T , the features learned to approximate V^{π_T} must correspond to the features used by the agent itself.

$$assessor(s_t, \theta_A) \approx \mathbb{E}_{\pi_T} \left[\sum_{k=0}^{\infty} \gamma^k r_{t+k+1} \right] = V^{\pi_T}(s_t) \quad (9)$$

In our study we use empirical return to estimate value, meaning that we observe the immediate rewards at each step, and when the episode is over we retroactively compute and standardize the discounted sum of rewards at each state. During training, the assessor minimizes a Huber regression objective [9] between predicted value and empirical value. Despite the fact that the $J_{\mathcal{M}_{SR}}$ objective is different than J_T , the task of approximating V^{π_T} is directly related to π_T . For example, value-based agents directly approximate a value estimate $Q^{\pi_T}(s_t, \pi_T(s_t)) = V^{\pi_T}(s_t)$, and for policy-based agents the explicit goal of the resulting policy is to maximize return, so implicitly the agent learns to assess the value of states.

5.4 SA threat model

In the **SA** threat model, we have access to a stream of states and can snoop on the actions taken by the agent at each step. Therefore, approximating the target’s policy π_T can be done directly with supervised learning. We can train an *imitator* model to predict the action that the target will take at a given state as:

$$imitator(s_t, \theta_I) \approx \pi_T(s_t). \quad (10)$$

Since the imitator is a classifier, we train it with cross-entropy loss ($J_{\mathcal{M}_{SA}}$) using the agent’s actions as labels. Although the imitator is trained with a different objective, it is directly related to the target agent’s goal. For value-based agents, the policy π is an ϵ -greedy extension of the learned Q-function and for policy-based agents, the function π_T is directly approximated. This methodology is homologous to the black-box attack strategy in [25], in which a substitute model is trained using the target model as an oracle. Note that if we were to have a continuous action space, the imitator would simply become a regression model.




Figure 4: The effect of uniform random noise on trained agent performance. BASE agents are trained without noise and AUG agents are trained with noise. Rewards are averages over 10 episodes.

5.5 SRA threat model

When we have the ability to eavesdrop on rewards *and* actions, we can attack with an imitator, assessor, or psychic. However, with this additional information we can improve our attacks by strategically timing the perturbations to make them less detectable. Lin et al. [18] show that we can accomplish this in a white-box setting by defining a preference function

$$c(s_t) = \max_{a_t} [\text{Softmax}(Q(s_t, a_t))] - \min_{a_t} [\text{Softmax}(Q(s_t, a_t))] \tag{11}$$

and attacking when $c(s) \geq \beta$ for some threshold β . At first glance, it may seem possible to use a similar preference-based approach with our imitator or assessor models. However, these proxies are only ever exposed to one particular policy π_T during training, so they lack the ability to assign credit to certain states or actions from the delayed reward signal. The RL agent, by contrast, learns this credit assignment through repeated trial and error during training. Nevertheless, we can approximate this ability using a combination of an assessor and an action-conditioned psychic (*AC-psychic*), which predicts s_{t+1} given s_t and a_t [39]. To perform the attack, we use the AC-psychic to generate a hypothetical next-state for every action given the current state, and value each of these with the assessor. We define a preference function whose output increases as the variation of potential future reward increases. In other words, we want to perturb s_t when there is a large difference in the different hypothetical values $V^{\pi_T}(s_{t+1}^H)$. To craft the perturbations, we can choose \mathcal{M}_{SRA} to be the proxy model of our choice (i.e. \mathcal{M}_S , \mathcal{M}_{SR} , or \mathcal{M}_{SA}). For details, see Algorithm 1.

6 IMPLEMENTATION DETAILS

6.1 Agents

We train DQN [21] and PPO [33] agents on Pong, Breakout, Space Invaders, and Seaquest games in the Atari 2600 Arcade Learning Environment [2] via OpenAI Gym [5]. Preprocessing is primarily facilitated using the Atari wrappers in OpenAI’s baselines library [7], which converts frames from RGB to grayscale [0, 1], and resizes frames to 84x84. To make the environment Markovian, a state is created by stacking the last four consecutive frames, making the input volume to the models of shape 4x84x84. We use the PyTorch

Algorithm 1: Strategically-timed snooping attack

Input: Trained assessor, trained AC-psychic, trained proxy \mathcal{M}_κ , trained target agent \mathcal{T} , β

for $t = 1, T$ **do**

- Initialize empty list q ;
- foreach** $a \in \mathcal{A}$ **do**
 - Predict s_{t+1}^H with *AC-psychic*(s_t, a)
 - Estimate V^H with *assessor*(s_{t+1}^H)
 - Append V^H to q ;
- end**
- $c(s_t) = \max [\text{Softmax}(q)] - \min [\text{Softmax}(q)]$
- if** $c(s_t) \geq \beta$ **then**
 - Perturb s_t using $\nabla_x J_{\mathcal{M}_\kappa}$
- end**
- Feed s_t to target \mathcal{T} for action decision;

end

deep learning framework [27], the DQN implementation described in [22], and Ilya Kostrikov’s implementation of PPO [14].

6.2 Proxies

For the psychic and AC-psychic, we use a scaled-down implementation of the architecture described in [39], as they work with the raw RGB frames and we predict the preprocessed version of the frames. The imitator architecture that we use is identical to the smaller DQN that was initially introduced in [21], and apply a Softmax operation to the logits for use with the cross-entropy classification loss. Note that we intentionally use a different architecture from the target agents in the interest of strictly adhering to black-box assumptions. The assessor uses the same architecture as the imitator, but we replace the classification layer with a single output node, and train the model with the Huber regression loss [9]. To train these models, we first let the trained agent collect experiences and save these into a buffer of 100,000 state/label pairs. Once we fill the buffer, we train on samples from this buffer of experiences for 30 epochs, using a batch size of 64. We repeat this process until we reach 2.5 million total training iterations. We find that this procedure significantly reduces training time compared to purely online training.

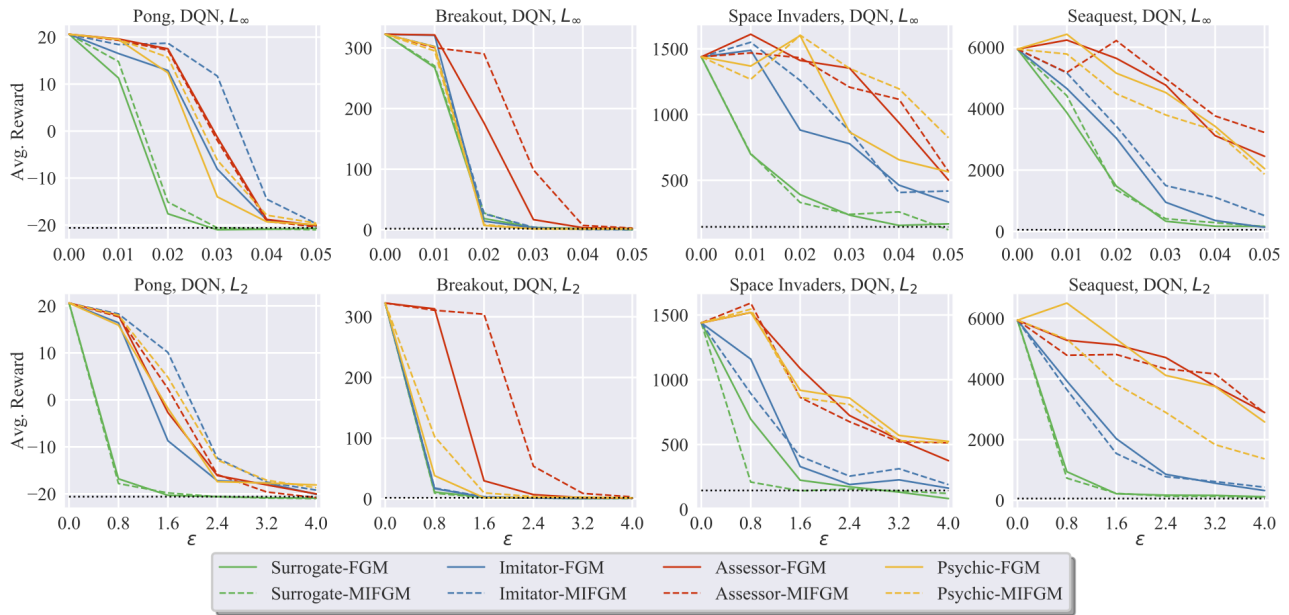


Figure 5: Performance reduction of DQN agents due to L_∞ and L_2 bounded perturbations. The black dotted line represents a random-guess policy.

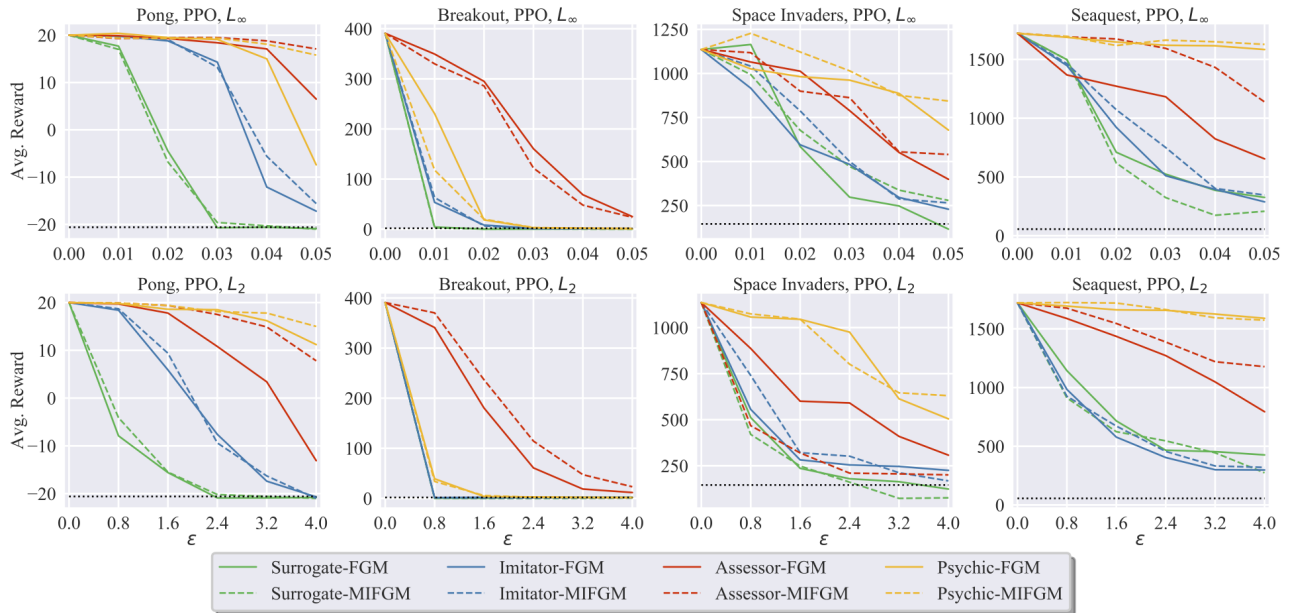


Figure 6: Performance reduction of PPO agents due to L_∞ and L_2 bounded perturbations. The black dotted line represents a random-guess policy.

7 EXPERIMENTS AND ANALYSIS

Although previous work has shown that DRL agents are vulnerable to adversarial noise, to our knowledge it is not clear whether it is the *adversarial* nature of the noise or simply the presence of noise that causes the agents to fail. Figure 4 shows that small uniform random noise significantly impacts agent performance. We believe that

this fragility is due to the narrow training data distribution. Atari frames do not naturally contain the noise, occlusion, or lighting variations that natural images do, which leads to a higher risk of overfitting. In the interest of proving that our attacks are effective due to transferability between proxy and agent, for the remainder

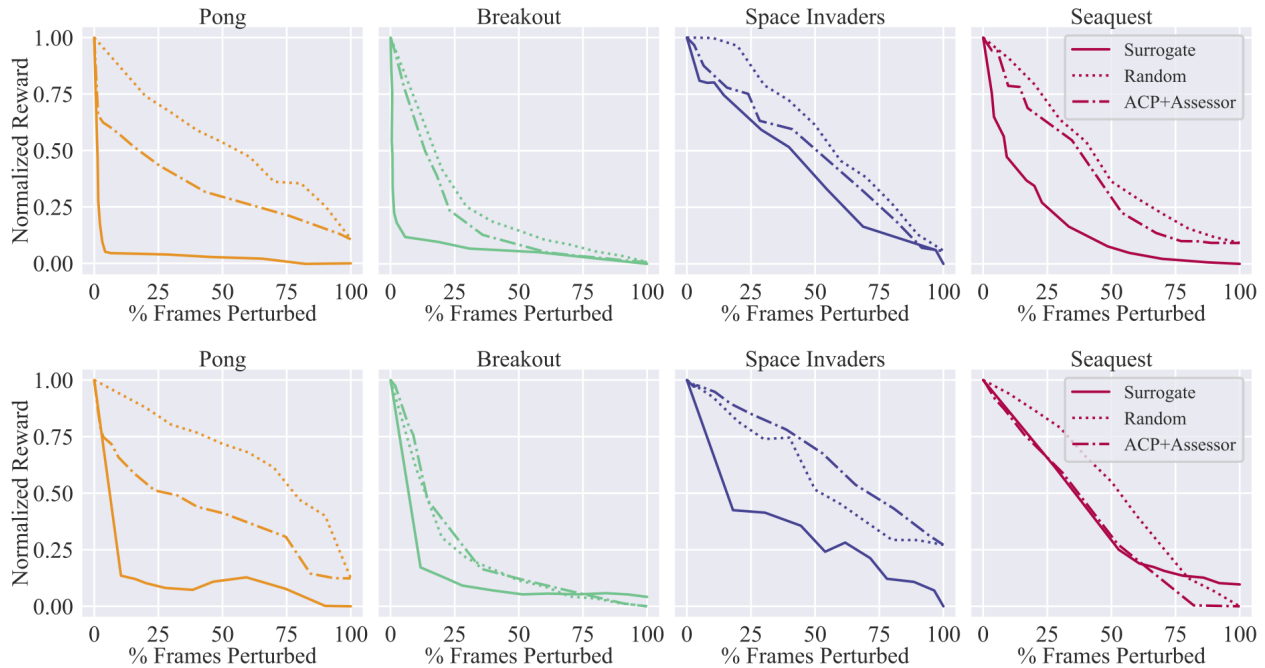


Figure 7: Effectiveness of strategically-timed attacks on agents over 30,000 time steps. Top row corresponds to DQN agents, while bottom corresponds to PPO agents. Our ACP+Assessor method and the random baseline use imitator gradients. We also compare to the technique in [18], which uses the stronger surrogate gradients. The FGM method is used for crafting.

of the work we only consider agents that are hardened with uniform random noise perturbations during training. Our experiments show that this produces more robust models while not significantly reducing clean data performance.

To test the effectiveness of the attacks outlined in Section 5, we average the total reward accumulated by the target agent over 10 episodes of states perturbed by the FGM or MIFGM attacks at every time step (using the proxy gradients $\nabla_x J_{M_k}$). We repeat this on each game for five ϵ values (i.e. attack strengths). For an upper-bound baseline attack, we use the policy-transfer methodology proven to be effective in [12] that assumes the adversary has access to the target agent’s environment and can train an identical model to the target agent. We can then use this model as a surrogate for crafting transferable adversarial examples. Figures 5 and 6 show the effectiveness of the transferred adversarial examples on the DQN and PPO agents, respectively.

The imitators produce the most effective attacks out of the proxy models, performing comparably to the surrogate in all games. We believe this is because these models represent the most direct proxy to the agent’s policy π_T . The assessor attacks generally require a higher ϵ to reach the performance of the baseline and imitator attacks. Work by Tramèr et al. [38] gives the insight that better performing models yield more transferable adversarial examples. We therefore hypothesize that the reason for the lower performance of the assessor attacks is because the task of predicting the exact floating point value of a given state is a difficult objective to approximate precisely given the sparsity and variance of the reward signals.

The attacks from the psychics perform worse than the surrogate and imitator attacks in general. This is not surprising, as the psychics operate under the most constrained threat model. In essence, the psychic model learns a noisy encoding of what the imitator learns directly. However, note that in the Breakout environment the attacks from the psychic model roughly match the effectiveness of the surrogate and imitator attacks. We believe that this is because Breakout has a deterministic state transition function (i.e., the next frame can be exactly predicted by the current frame), unlike the other games which have state elements outside of the agent’s control (e.g., hardcoded paddle in Pong, etc.). Thus, because the psychic model’s learned representation of state dynamics is a direct delineation of the agent’s policy, subsequent attacks are more transferable. This result further supports our hypothesis that proxy models that more directly approximate the target agent’s policy will yield more transferable attacks.

To qualitatively evaluate the features that these models learn, we create perturbation-based saliency maps [11] designed to be human interpretable visualizations of $\nabla_x J(\theta, x_i, y_i)$ on Atari games. Figure 3 shows a representative example of what the saliency maps tell us about these models. As expected, the surrogate model’s map is very similar to the target’s. The imitator’s salient regions tend to be the most similar to the target’s out of the proxy models, which further supports our beliefs regarding attack effectiveness. The assessor tends to focus less on the agent itself (i.e. the submarine), and more on the objects around it which correspond to potential future reward. This intuition also helps to explain the inferior performance




Figure 8: Noise visualization produced FGM attack by a surrogate and each proxy model for a DQN agent trained on Atari Pong environment. On the left is the current unperturbed game state. The top and bottom rows of noise are the L_∞ and L_2 constrained perturbations, respectively. We note that the saliency patterns in the gradients of the different proxies show considerable visual similarities.

of the assessor’s attacks. Finally, because the psychic model predicts the entire next-frame, it must attend to any object in the state space that moves. Because the psychic learns an approximation of the state transition dynamics, a subset of the features that it learns (e.g. the submarine, fish, divers) do correspond to features that are important to the agent’s policy. Figure 8 gives a direct visualization of the gradient-based perturbations of each proxy on a Pong state. Note that the L_2 perturbations from the imitator and assessor proxies share considerable visual similarities to the surrogate agent’s “upper bound” perturbation. The psychic model’s perturbations show notable saliency in regions of the ball and paddles, but the gradients are not as strongly localized.

The results of the strategically-timed snooping attack on DQN and PPO agents can be seen in Figure 7. Note that we use the FGM crafting method with imitator gradients in this experiment. We compare this performance to (1) the black-box variant of the original strategically-timed attack described in [18] using FGM with surrogate gradients, and (2) an attack which perturbs at random time steps with FGM and imitator gradients. The strategically-timed attacks will typically be more effective on the paddle-based games such as Pong, as there is a clear distinction between time steps that are critical for good performance (i.e. as the ball approaches the paddle) and time steps when the actions of the agent do not matter. As expected, our method does not perform as well as the surrogate that was trained to assign credit to certain states and actions. Our method does show promise, however, as it performs markedly better than the random strategy without requiring any interaction with the environment. We believe that our method’s effectiveness is bottlenecked by the high-variance assessor predictions, and reducing the variance of the reward signal learned would lead to better performance.

An unexpected result of our experiments is the unimpressive performance of MIFGM compared to vanilla FGM. To understand why, we consider that iterative attacks like MIFGM create a more custom perturbation for the model that it is crafted by. This introduces the possibility that the resulting adversarial examples are overfit to a local maxima of the white-box model. In the case of the the proxy models, we posit that because of the significant differences in the

loss functions used by the agents and the proxies, adversarial examples that are more tailored to a specific objective will not perform as well. In the case of the baseline surrogate agents that *do* use the same objectives, the high variance of the RL data generation process is more likely to yield models with differently-shaped decision surfaces compared to supervised classifiers that draw i.i.d. samples. Again, this renders these more tailored perturbations less effective for transfer-based attacks.

Finally, we propose that our work offers a unique prospective on adversarial example transferability. Previous work has shown that adversarial examples transfer between different classifiers, etc. This phenomena has been justified by decision surface similarity between the classifiers [19, 24, 38]. However, to our knowledge, we are among the first to demonstrate that adversarial examples transfer across models that learn with very different optimization objectives as long as the tasks are related such that the gradient of the loss with respect to the input has a similar shape. For example, although the loss, output dimensionality, and decision surface of the psychic proxy is radically different from the target model, adversarial examples generated by the psychic also fool the agent. This implies that an adversary can effectively attack agents *even if they have no knowledge of what the target system is trying to do*. These findings alone force us to reconsider what makes adversarial examples transfer between models. We believe that our snooping threat models and proxies offer a unique and effective test environment for further work on the study of transferability of adversarial examples between models with fundamentally different but *related* objectives.

8 CONCLUSIONS

It is often infeasible for an adversary to have access to a target agent’s parameters or train a DRL surrogate agent on a proprietary environment. However, we show that an adversary only needs the ability to intercept the states and optionally snoop on the action and reward signals to launch powerful adversarial attacks on a black-box agent. Further, we show that adversaries can use trained proxy models to make their attacks sparse in time, rendering them more difficult to detect. We believe these to be important security vulnerabilities to consider and address before deploying these models to systems that we trust.

REFERENCES

- [1] Vahid Behzadan and Arslan Munir. 2017. Vulnerability of Deep Reinforcement Learning to Policy Induction Attacks. In *MLDM*.
- [2] Marc G. Bellemare, Yavar Naddaf, Joel Veness, and Michael H. Bowling. 2013. The Arcade Learning Environment: An Evaluation Platform for General Agents. *J. Artif. Intell. Res.* 47 (2013), 253–279.
- [3] Richard Bellman. 1966. Dynamic programming. *Science* 153 3731 (1966), 34–7.
- [4] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Prasoon Goyal, Lawrence D. Jackel, Miguel Pozuelo Monfort, Urs Muller, Jiakai Zhang, Xin Zhang, Junbo Jake Zhao, and Karol Zieba. 2016. End to End Learning for Self-Driving Cars. *CoRR abs/1604.07316* (2016).
- [5] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. 2016. OpenAI Gym. (2016). arXiv:arXiv:1606.01540
- [6] Nicholas Carlini and David A. Wagner. 2017. Towards Evaluating the Robustness of Neural Networks. *2017 IEEE Symposium on Security and Privacy (SP)* (2017), 39–57.
- [7] Prafulla Dhariwal, Christopher Hesse, Oleg Klimov, Alex Nichol, Matthias Plappert, Alec Radford, John Schulman, Szymon Sidor, Yuhuai Wu, and Peter Zhokhov. 2017. OpenAI Baselines. <https://github.com/openai/baselines>. (2017).

- [8] Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. 2018. Boosting Adversarial Attacks with Momentum. In *CVPR*.
- [9] Ross B. Girshick. 2015. Fast R-CNN. *2015 IEEE International Conference on Computer Vision (ICCV)* (2015), 1440–1448.
- [10] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and Harnessing Adversarial Examples. *CoRR abs/1412.6572* (2015).
- [11] Sam Greysdanus, Anurag Koul, Jonathan Dodge, and Alan Fern. 2018. Visualizing and Understanding Atari Agents. In *ICML*.
- [12] Sandy H. Huang, Nicolas Papernot, Ian J. Goodfellow, Yan Duan, and Pieter Abbeel. 2017. Adversarial Attacks on Neural Network Policies. *CoRR abs/1702.02284* (2017).
- [13] Nathan Inkawhich, Wei Wen, Hai Li, and Yiran Chen. 2019. Feature Space Perturbations Yield More Transferable Adversarial Examples. In *CVPR*.
- [14] Ilya Kostrikov. 2018. PyTorch Implementations of Reinforcement Learning Algorithms. <https://github.com/ikostrikov/pytorch-a2c-ppo-acktr-gail>. (2018).
- [15] Felix Leibfried, Nate Kushman, and Katja Hofmann. 2017. A Deep Learning Approach for Joint Video Frame and Reward Prediction in Atari Games. *CoRR abs/1611.07078* (2017).
- [16] Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. 2016. End-to-End Training of Deep Visuomotor Policies. *Journal of Machine Learning Research* 17 (2016), 39:1–39:40.
- [17] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. 2016. Continuous control with deep reinforcement learning. *CoRR abs/1509.02971* (2016).
- [18] Yen-Chen Lin, Zhang-Wei Hong, Yuan-Hong Liao, Meng-Li Shih, Ming-Yu Liu, and Min Sun. 2017. Tactics of Adversarial Attack on Deep Reinforcement Learning Agents. In *IJCAI*.
- [19] Yanpei Liu, Xinyun Chen, Chang Liu, and Dawn Xiaodong Song. 2017. Delving into Transferable Adversarial Examples and Black-box Attacks. *CoRR abs/1611.02770* (2017).
- [20] Volodymyr Mnih, Adria Puigdomènech Badia, Mehdi Mirza, Alex Graves, Timothy P. Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. 2016. Asynchronous Methods for Deep Reinforcement Learning. In *ICML*.
- [21] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin A. Riedmiller. 2013. Playing Atari with Deep Reinforcement Learning. *CoRR abs/1312.5602* (2013).
- [22] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin A. Riedmiller, Andreas Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharmarshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. 2015. Human-level control through deep reinforcement learning. *Nature* 518 (2015), 529–533.
- [23] Junhyuk Oh, Xiaoxiao Guo, Honglak Lee, Richard L. Lewis, and Satinder P. Singh. 2015. Action-Conditional Video Prediction using Deep Networks in Atari Games. In *NIPS*.
- [24] Nicolas Papernot, Patrick D. McDaniel, and Ian J. Goodfellow. 2016. Transferability in Machine Learning: from Phenomena to Black-Box Attacks using Adversarial Samples. *CoRR abs/1605.07277* (2016).
- [25] Nicolas Papernot, Patrick D. McDaniel, Ian J. Goodfellow, Somesh Jha, Z. Berkay Celik, and Ananthram Swami. 2017. Practical Black-Box Attacks against Machine Learning. In *AsiaCCS*.
- [26] Nicolas Papernot, Patrick D. McDaniel, Somesh Jha, Matt Fredrikson, Z. Berkay Celik, and Ananthram Swami. 2016. The Limitations of Deep Learning in Adversarial Settings. *2016 IEEE European Symposium on Security and Privacy (EuroS&P)* (2016), 372–387.
- [27] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in PyTorch. (2017).
- [28] Deepak Pathak, Parsa Mahmoudieh, Guanghao Luo, Pulkit Agrawal, Dian Chen, Yide Shentu, Evan Shelhamer, Jitendra Malik, Alexei A. Efros, and Trevor Darrell. 2018. Zero-Shot Visual Imitation. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)* (2018), 2131–21313.
- [29] Anay Pattanaik, Zhenyi Tang, Shuijing Liu, Gautham Bommanan, and Girish Chowdhary. 2018. Robust Deep Reinforcement Learning with Adversarial Attacks. In *AAMAS*.
- [30] Lerrel Pinto, Marcin Andrychowicz, Peter Welinder, Wojciech Zaremba, and Pieter Abbeel. 2018. Asymmetric Actor Critic for Image-Based Robot Learning. *CoRR abs/1710.06542* (2018).
- [31] Ahmad El Sallab, Mohammed Abdou, Etienne Perot, and Senthil Yogamani. 2017. Deep Reinforcement Learning framework for Autonomous Driving. *CoRR abs/1704.02532* (2017).
- [32] John Schulman, Sergey Levine, Philipp Moritz, Michael I. Jordan, and Pieter Abbeel. 2015. Trust Region Policy Optimization. In *ICML*.
- [33] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal Policy Optimization Algorithms. *CoRR abs/1707.06347* (2017).
- [34] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Vedavyas Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy P. Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. 2016. Mastering the game of Go with deep neural networks and tree search. *Nature* 529 (2016), 484–489.
- [35] David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin A. Riedmiller. 2014. Deterministic Policy Gradient Algorithms. In *ICML*.
- [36] Richard S. Sutton and Andrew G. Barto. 1988. Reinforcement Learning: An Introduction. *IEEE Transactions on Neural Networks* 16 (1988), 285–286.
- [37] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. 2014. Intriguing properties of neural networks. *CoRR abs/1312.6199* (2014).
- [38] Florian Tramèr, Nicolas Papernot, Ian J. Goodfellow, Dan Boneh, and Patrick D. McDaniel. 2017. The Space of Transferable Adversarial Examples. *CoRR abs/1704.03453* (2017).
- [39] Elias Wang. 2017. Deep Action Conditional Neural Network for Frame Prediction in Atari Games.
- [40] Christopher J. C. H. Watkins and Peter Dayan. 1992. Technical Note: Q-Learning. *Machine Learning* 8 (1992), 279–292.