

RL-Based Method for Benchmarking the Adversarial Resilience and Robustness of Deep Reinforcement Learning Policies

Vahid Behzadan and William Hsu
{behzadan, bhsu}@ksu.edu

Kansas State University

Abstract. This paper investigates the resilience and robustness of Deep Reinforcement Learning (DRL) policies to adversarial perturbations in the state space. We first present an approach for the disentanglement of vulnerabilities caused by representation learning of DRL agents from those that stem from the sensitivity of the DRL policies to distributional shifts in state transitions. Building on this approach, we propose two RL-based techniques for quantitative benchmarking of adversarial resilience and robustness in DRL policies against perturbations of state transitions. We demonstrate the feasibility of our proposals through experimental evaluation of resilience and robustness in DQN, A2C, and PPO2 policies trained in the Cartpole environment.

Keywords: Deep Reinforcement Learning · Adversarial Attack · Policy Generalization · Resilience · robustness · benchmarking.

1 Introduction

Since the reports by Behzadan & Munir [1] and Huang et al. [5], the primary emphasis of the state of the art in DRL security [2] has been on the vulnerability of policies to state-space perturbations. In particular, the manipulation of the policy via adversarial examples [4] has remained the main focus of current literature on this issue. However, this bias towards adversarial example attacks gives rise to a critical shortcoming: the analyses of such attacks fail to disentangle the vulnerability caused by the learned representation and that which is due to the sensitivity of the DRL dynamics to distributional shifts in state transitions. Also, the performance of defenses proposed for adversarial example attacks are inherently limited to the considered attack mechanisms. As the most successful technique for mitigation of adversarial examples, adversarial training is known to enhance the robustness of machine learning models to the type of attack used for generating the training adversarial examples, while leaving the model vulnerable to other types of attacks[8]. Furthermore, the current literature fails to provide solutions and approaches which can be used in practice to evaluate and improve the robustness and resilience of DRL policies to attacks that exploit the sensitivity to state transitions. Also, there remains a need for quantitative

approaches to measure and benchmark the resilience and robustness of DRL policies in a reusable and generalizable manner.

In response to these shortcomings, this paper aims to address the problem of quantifying and benchmarking the robustness and resilience of a DRL agent to adversarial perturbations of state transitions at test-time, in a manner that is independent of the attack type. This improves the generalization of current techniques that analyze the model against specific adversarial example attacks. Accordingly, the main contributions of this paper are as follows:

1. We present formulations of the resilience and robustness problems that enable the disentanglement of limitation in representation learning from sensitivity of policies to state transition dynamics.
2. We propose two RL-based techniques and corresponding metrics for the measurement and benchmarking of resilience and robustness of DRL policies to perturbations of state transitions,
3. We demonstrate the feasibility of our proposal through experimental evaluation of their performance on DQN, A2C, and PPO2 agents trained in the Cartpole environment.

The remainder of this paper is organized as follows: Section 2 defines and formulates the problems of adversarial resilience and robustness in DRL. Our proposed methods for benchmarking the test-time resilience and robustness of DRL policies are presented in Sections 3 and 4. Section 5 provides the details of experimental setup for evaluating the performance of our proposals, with the corresponding results presented in Section 6. The paper concludes in Section 7 with a summary of findings and remarks on future directions of research.

2 Problem Formulation

We consider the the generic problem of RL in the settings of a Markov Decision Process (MDP), described by the tuple $MDP := \langle \mathbb{S}, \mathbb{A}, \mathbb{R}, \mathbb{P} \rangle$, where \mathbb{S} is the set of reachable states in the process, \mathbb{A} is the set of available actions, \mathbb{R} is the mapping of transitions to the immediate reward, and \mathbb{P} represents the transition probabilities (i.e., state dynamics), which are initially unknown to RL agents. At any given time-step t , the MDP is at a state $s_t \in \mathbb{S}$. The RL agent’s choice of action at time t , $a_t \in \mathbb{A}$ causes a transition from s_t to a state s_{t+1} according to the transition probability $P(s_{t+1}|s_t, a_t)$. The agent receives a reward $r_{t+1} = R(s_t, a_t, s_{t+1})$ for choosing the action a_t at state s_t . Interactions of the agent with MDP are determined by the policy π . When such interactions are deterministic, the policy $\pi : S \rightarrow \mathbb{A}$ is a mapping between the states and their corresponding actions. A stochastic policy $\pi(s)$ represents the probability distribution of implementing any action $a \in \mathbb{A}$ at state s . The goal of RL is to learn a policy that maximizes the expected discounted return $E[R_t]$, where $R_t = \sum_{k=0}^{\infty} \gamma^k r_{t+k}$; with r_t denoting the instantaneous reward received at time t , and γ is a discount factor $\gamma \in [0, 1]$.

To facilitate the formal statement of adversarial resilience and robustness, we first introduce the following definitions:

- **Adversarial Regret** at time T is the difference between return obtained by the nominal (unperturbed) agent at time T and the return obtained by the perturbed agent at time T . Formally: $\hat{R}_{adv}(T) = R_{nominal}(T) - R_{perturbed}(T)$. The time T may represent either the terminal time step of an episode, or the time-horizon of interest in the analysis.
- **Adversarial Budget** is defined by one or more of the following parameters: the maximum number of features that can be perturbed in the observations ($O_{max} \in [0, \infty]$), the maximum number of observations that can be perturbed ($N_{max} \in [0, \infty]$), and the probability of perturbing each observation ($P(perturb) \in [0, 1]$).

Building on these two concepts, we define the problems of adversarial resilience and robustness as follows:

1. **Test-Time Resilience:** The minimum number of state perturbations required to incur the maximum reduction to the total return at time T (denoted by $\hat{R}_{adv}(T)$) for an agent driven by a policy $\pi(s)$ in an environment with transition dynamics \mathbb{P} .
2. **Test-Time Robustness:** The maximum adversarial regret $\hat{R}_{adv}(T) = \epsilon_{max}$ achievable via a maximum of δ_{max} state perturbations for an agent driven by a policy $\pi(s)$ in an environment with transition dynamics \mathbb{P} .

The following sections provide the details of our proposed solutions to each of the aforementioned problem settings.

3 Benchmarking of Test-Time Resilience

This problem can be modeled as that of finding an optimal adversarial policy $\pi_{adv}(s)$ that minimizes the cost incurred to the adversary C_{adv} in order to impose the maximum adversarial regret $\hat{R}_{adv}(T)$, the worst-case value of which is the highest cumulative reward achieved by the target policy R_{max} . Our proposed approach is through the formulation of this problem in the settings of reinforcement learning. The state space in the corresponding MDP is the set of states in the target MDP, augmented with the action of the target in that state, i.e., $S' = \{\forall s \in \mathbb{S} : (s, \pi(s))\}$. For the purpose of measuring a lower bound for the resilience, we consider the worst-case white-box adversary, which is able to impose targeted state perturbations with a 100% success rate, to induce any action within the permissible action-set of the target \mathbb{A} which has the lowest Q -value at any state s according to the target’s optimal state-action value function Q^* . In this case, the set of permissible adversarial actions at any state s is given by:

$$A_{adv}(s) = \{\text{No Action}\} \cup \mathbb{A} \setminus \pi^*(s) \quad (1)$$

where \mathbb{A} is the action set of the targeted agent, and $\pi : S \rightarrow A$ is the policy of the targeted agent. In the proposed approach, the adversarial reward value is determined via the procedure detailed in Algorithm 1:

Algorithm 1 Reward Assignment of RL Agent for Measuring Adversarial Resilience

Require: Target policy π^* , Perturbation cost function $c_{adv}(\cdot, \cdot)$, Maximum achievable score R_{max} , Optimal state-action value function $Q^*(\cdot, \cdot)$, Current adversarial policy π^{adv} , Current state s_t , Current count of adversarial actions $AdvCount$, Current score R_t

Set $ToPerturb \leftarrow \pi^{adv}(s_t)$

if $ToPerturb$ is False **then**

$a_t \leftarrow \pi^*(s_t)$

$Reward \leftarrow 0$

else

$a'_t \leftarrow \arg \min_a Q^*(s_t, a)$

$Reward \leftarrow -c_{adv}(s_t, a'_t)$

end if

if either s_t or s'_t is terminal **then**

$Reward+ = (R_{max} - R_t)$

end if

where $c(s_t, a'_t)$ is the cost of imposing the state perturbation which induces the adversarial action a'_t at state s_t . It is noteworthy that if the value of $c(s_t, a'_t)$ is invariant with respect to a'_t , the adversarial action set reduces to:

$$A_{adv}(s) = \{\text{No Action}, \underset{a}{\text{Induce } \arg \min Q(s, a)}\} \quad (2)$$

To obtain the test-time resilience of policy π^* to state perturbations, we propose the following procedure:

1. If the state-action value function of the target Q^* is not available (i.e., black-box testing), approximate Q^* via policy imitation [6].
2. Train the adversarial agent against the target following π in its training environment, report the optimal adversarial return $R_{perturbed}^*$ and the maximum adversarial regret $R_{adv}^*(T)$.
3. Apply the adversarial policy against the target in N episodes, record total cost C_{adv} for each episode,
4. Report the average of C_{adv} over N episodes as the mean test-time resilience of π in the given environment.

This procedure introduces three metrics for the quantification of test-time resilience: the optimal adversarial return $R_{perturbed}^*$ achieved in the training process of the adversarial policy, the maximum adversarial regret $R_{adv}^*(T)$ achieved during training, and the mean per-episode of the total cost C_{adv} . These metrics provide the means to benchmark and compare the test-time resilience of different policies trained to optimize the agent’s performance in a given environment.

For the purpose of measuring resilience, we consider convergence to be reached if the average adversarial regret over 200 episodes remains constant. This definition relaxes the instabilities that may arise due to the configuration and architecture of the DRL training process. It is noteworthy that depending on the

training algorithm and design parameters, this procedure is not guaranteed to converge to global optima. However, by reporting the number of iterations and configuration of random number generators with a constant seed, the reported results present a reproducible loose lower bound on the adversarial resilience of the target. Also, the trained adversarial policy can be used to test other policies for comparison of such lower bounds under the same adversarial strategy.

4 Benchmarking of Test-Time Robustness

For this problem, we propose a modified version of the procedure developed for benchmarking the test-time resilience. Accordingly, the reward function is adjusted to account for the lack of a target ϵ , as well as the addition of an adversarial budget constraint δ_{max} . The reward measurement of this process is outlined in Algorithm 2:

Algorithm 2 Reward Assignment of RL Agent for Measuring Adversarial Robustness

Require: Maximum perturbation budget δ_{max} , Perturbation cost function $c_{adv}(\cdot, \cdot)$, Maximum achievable score R_{max} , Optimal state-action value function $Q^*(\cdot, \cdot)$, Current adversarial policy π^{adv} , Current state s , Current count of adversarial actions $AdvCount$, Current score R_t
 Set $AdversarialAction \leftarrow \pi^{adv}(s)$
if $AdversarialAction$ is NoAction **then**
 $Reward \leftarrow 0$
else if $AdvCount \geq \delta_{max}$ **then**
 $Reward \leftarrow -c_{adv}(s, AdversarialAction) \times \delta_{max}$
 $AdvCount+ = 1$
else
 $Reward \leftarrow -c_{adv}(s, AdversarialAction)$
 $AdvCount+ = 1$
end if
if s is terminal **then**
 $Reward+ = 1.0 * (R_{max} - R_t)$
 $AdvCount \leftarrow 0$
end if

The proposed procedure for measuring the test-time robustness of a given DRL policy to adversarial state perturbations is as follows:

1. If the state-action value function of the target Q^* is not available (i.e., black-box testing settings), approximate Q^* from the policy using imitation learning (e.g., [6]),
2. Train the adversarial agent against the target policy π^* in its training environment, report the maximum adversarial regret $R_{adv}^*(T)$ for time T achieved at adversarial optimality,

3. Apply the adversarial policy against the target for N episodes, record the adversarial regret at the end of each episode $R_{adv}(T)$,
4. Report the average of $R_{adv}(T)$ over N episodes as the mean per-episode test-time robustness of π^* in the given environment.

5 Experiment Setup

Environment and Target Policies: To demonstrate the performance of the proposed procedures for benchmarking the test-time robustness and resilience in DRL policies, we present the analysis of the aforementioned measurements for policies trained in the CartPole environment in OpenAI Gym [3]. The considered policies are chosen to represent the commonly-adopted state of the art method from each class of DRL algorithms. From value-iteration approaches, we consider DQN with prioritized replay. From the class of policy gradient approaches, we consider PPO2. As for actor-critic methods, we investigate the A2C method. Table 5 presents the specifications of the CartPole environment, and Tables 5 – 5 provide the parameter settings of each target policy.

Table 1. Specifications of the CartPole Environment

Observation Space	Cart Position [-4.8, +4.8] Cart Velocity [-inf, +inf] Pole Angle [-24 deg, +24 deg] Pole Velocity at Tip [-inf, +inf]
Action Space	0 : Push cart to the left 1 : Push cart to the right
Reward	+1 for every step taken
Termination	Pole Angle is more than 12 degrees Cart Position is more than 2.4 Episode length is greater than 500

Table 2. Parameters of DQN Policy

No. Timesteps	10^5
γ	0.99
Learning Rate	10^{-3}
Replay Buffer Size	50000
First Learning Step	1000
Target Network Update Freq.	500
Prioritized Replay	True
Exploration	Parameter-Space Noise
Exploration Fraction	0.1
Final Exploration Prob.	0.02
Max. Total Reward	500

Table 3. Parameters of A2C Policy

No. Timesteps	5×10^5
γ	0.99
Learning Rate	7×10^{-4}
Entropy Coefficient	0.0
Value Function Coefficient	0.25
Max. Total Reward	500

Table 4. Parameters of A2C Policy

No. Environments	8
No. Timesteps	10^6
No. Runs per Environment per Update	2048
No. Minibatches per update	32
Bias-Variance Trade-Off Factor	0.95
No. Surrogate Epochs	10
γ	0.99
Learning Rate	3×10^{-4}
Entropy Coefficient	0.0
Value Function Coefficient	0.5
Max. Total Reward	500

Adversarial Agent: In these experiments, the adversarial agent is a DQN agent with the hyperparameters provided in Table 5. We consider a homogeneous perturbation cost function for all state perturbations, that is, $\forall s, a' : c_{adv}(s, a') = c_{adv}$. For both the resilience and robustness measurements, we set $c_{adv} = 1$ (i.e., each perturbation incurs a cost of 1 to the adversary). The training process is terminated when the adversarial regret is maximized and the 100-episode average of the number of adversarial perturbations is quasi-stable for 200 episodes.

Table 5. Parameters of DQN Policy

Max. Timesteps	10^5
γ	0.99
Learning Rate	10^{-3}
Replay Buffer Size	50000
First Learning Step	1000
Target Network Update Freq.	500
Experience Selection	Prioritized Replay
Exploration	Parameter-Space Noise
Exploration Fraction	0.1
Final Exploration Prob.	0.02

6 Results

6.1 Resilience Benchmarks

We consider the white-box settings in the training of adversarial agents for resilience measurement. For the DQN target, the optimal state-action value function Q^* of the target is directly utilized. As for the A2C and PPO2 targets, the state-action value function is calculated from the internally-available state value estimations $V^*(s)$ according to the following transformation:

$$Q^*(s_t, a) = r(s_t, a) + \gamma V^*(s_{t+1}) \quad (3)$$

where s_{t+1} is the state resulting from a transition out of state s_t by implementing action a .

Training Results: The training progress plots of adversarial DQN policy on the three target policies are presented in Fig.1–3. It can be seen that all three policies converge to the same optima. However, for the adversary targeting the DQN policies, the convergence is achieved at a higher number of training steps.

It is noteworthy that for all three policies, the mean-per-100 episodes of the minimum number of perturbations at convergence is almost similar (as reported in Table 6.2), with A2C having the largest value of 7.69 perturbations, PPO2 at a value of 7.49 perturbations, and DQN having the lowest value of 7.13. Also, the test-time performance of these trained policies indicate similar results, with DQN requiring 6.95 perturbations to incur an adversarial regret of 491.15, PPO2 requiring 7.72 perturbations for an adversarial regret of 490.47, and A2C requiring 8.71 perturbations for an adversarial regret of 488.16. Accordingly, we can interpret these results as follows: the DQN policy has the lowest adversarial resilience among the three, followed by the PPO2 policy. Within the context of this comparison, the A2C policy is found to be the most resilient to state-space perturbation attacks.




Fig. 1. Adversarial Training Progress for Resilience Benchmarking of the DQN Policy




Fig. 2. Adversarial Training Progress for Resilience Benchmarking of the A2C Policy

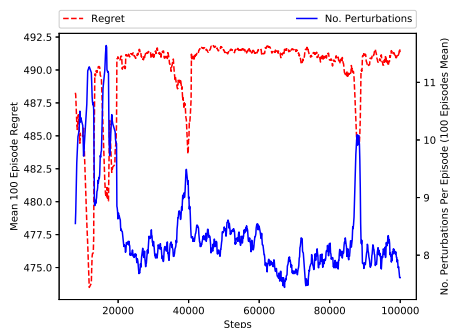


Fig. 3. Adversarial Training Progress for Resilience Benchmarking of the PPO2 Policy

6.2 Test-Time Step-Perturbation Distribution

To investigate the state-transition vulnerability of each policy, we also study the frequency of perturbing states at each time step of an episode for the three adversarial policies. The results, presented in Fig. 4–6, illustrate that in all three policies, the initial time steps have been the subject of most perturbations. This result is noteworthy, as it contradicts the assumption of Lin et al.[7] that the most effective adversarial perturbations are those that are mounted towards the terminal state of the environment.

Table 6. Comparison of Test-Time and Training-Time Resilience Measurements for DQN, A2C, and PPO2 Policies

Target Policy	Max. Regret	Avg. Regret (Training)	Avg. No. Perturbations (Training)	Avg. Regret	Avg. No. Perturbations
DQN	492	491.24	7.13	491.15	6.95
A2C	492	491.44	7.69	488.16	8.71
PPO2	492	491.72	7.49	490.47	7.72




Fig. 4. Perturbation Count Per Episodic Time Step in 100 Runs Targeting DQN Policy




Fig. 5. Perturbation Count Per Episodic Time Step in 100 Runs Targeting A2C Policy

6.3 Robustness Benchmarks

To demonstrate the performance of our proposed technique for benchmarking the robustness of DRL policies, we provide the training-time results for two cases of $\delta_{max} = 10$ and $\delta_{max} = 5$ for DQN, A2C, and PPO2 Policies. As illustrated in Fig.7a – 7c, all three adversarial policies converge with similar minimum perturbation counts as those obtained in resilience analysis. This is expected, as the resilience analysis established that the minimum number of actions required for maximum regret is 7.5, which is less than the available budget of $\delta_{max} = 10$. As for the case of $\delta_{max} = 5$, Fig.8a – 8c demonstrate significant differences between the three policies. In Fig.8a, it can be seen that at 5 actions, the convergence occurs with an adversarial regret of 462.5, while for A2C, the best 5-action indication of convergence occurs at an adversarial regret of 224. As for PPO2, this value is at 268.2. These results indicate a similar ranking of the robustness in these policies, with DQN being the least-robust to maximum of 5 perturbations, and the A2C prevailing as the most robust policy to maximum of 5 perturbations.




Fig. 6. Perturbation Count Per Episodic Time Step in 100 Runs Targeting PPO Policy




Fig. 7. Adversarial Training Progress for Robustness Benchmarking, $\delta_{max} = 10$

7 Conclusion

We presented two RL-based techniques for benchmarking the resilience and robustness of DRL policies to adversarial perturbations of state transition dynamics. Experimental evaluation of our proposals demonstrates the feasibility of these techniques for quantitative analysis of policies with regards to their sensitivity to state transition dynamics. A promising venue of further exploration is to study and extend the proposed methodologies for evaluation of generalization in DRL policies.

References

- Behzadan, V., Munir, A.: Vulnerability of deep reinforcement learning to policy induction attacks. In: International Conference on Machine Learning and Data Mining in Pattern Recognition. pp. 262–275. Springer (2017)
- Behzadan, V., Munir, A.: The faults in our pi stars: Security issues and open challenges in deep reinforcement learning. arXiv preprint arXiv:1810.10369 (2018)




Fig. 8. Adversarial Training Progress for Robustness Benchmarking, $\delta_{max} = 5$

3. Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., Zaremba, W.: Openai gym. arXiv preprint arXiv:1606.01540 (2016)
4. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples (2014). arXiv preprint arXiv:1412.6572 (2014)
5. Huang, S., Papernot, N., Goodfellow, I., Duan, Y., Abbeel, P.: Adversarial attacks on neural network policies. arXiv preprint arXiv:1702.02284 (2017)
6. Hussein, A., Gaber, M.M., Elyan, E., Jayne, C.: Imitation learning: A survey of learning methods. ACM Computing Surveys (CSUR) **50**(2), 21 (2017)
7. Lin, Y.C., Hong, Z.W., Liao, Y.H., Shih, M.L., Liu, M.Y., Sun, M.: Tactics of adversarial attack on deep reinforcement learning agents. arXiv preprint arXiv:1703.06748 (2017)
8. Tramèr, F., Kurakin, A., Papernot, N., Goodfellow, I., Boneh, D., McDaniel, P.: Ensemble adversarial training: Attacks and defenses. arXiv preprint arXiv:1705.07204 (2017)