

White-Box Adversarial Policies in Deep Reinforcement Learning

Stephen Casper
MIT CSAIL
scasper@mit.edu

Dylan Hadfield-Menell
MIT CSAIL

Gabriel Kreiman
Boston Children’s Hospital
Center for Brains, Minds, and Machines

Abstract—Adversarial examples against AI systems pose both risks via malicious attacks and opportunities for improving robustness via adversarial training. In multiagent settings, adversarial policies can be developed by training an adversarial agent to minimize a victim agent’s rewards. Prior work has studied black-box attacks where the adversary only sees the state observations and effectively treats the victim as any other part of the environment. In this work, we experiment with white-box adversarial policies to study whether an agent’s internal state can offer useful information for other agents. We make three contributions. First, we introduce white-box adversarial policies in which an attacker can observe a victim’s internal state at each timestep. Second, we demonstrate that white-box access to a victim makes for better attacks in two-agent environments, resulting in both faster initial learning and higher asymptotic performance against the victim. Third, we show that training against white-box adversarial policies can be used to make learners in single-agent environments more robust to domain shifts. Code is available at this [https](https://github.com/stephencasper/white-box-adversarial-policies) url.

Index Terms—adversarial examples, adversarial policies, reinforcement learning, white-box attacks

I. INTRODUCTION

As AI systems become more capable and widely-deployed, it becomes increasingly important to understand and address their vulnerabilities. These include concerns involving *adversarial* attacks that are specifically crafted to make a system fail. Adversarial attacks in the form of subtle perturbations to inputs have been widely studied in supervised learning [15], [43]. However, compared to supervised learning, reinforcement learning (RL) agents can face an expanded set of threats [23], [42], including adversarial *policies* from other agents. The standard approach for developing adversarial policies has been to train an attacker against a black-box victim until the attacker (over)fits a policy that minimizes the victim’s reward. These adversarial policies have been used both to attack victims [12], [14] and to improve a victim’s robustness through adversarial training [36].

This black-box approach often works well, but it fails to utilize any information beyond what the attacker can directly observe, thus treating the victim as any other part of the environment. As a result, this requires cheap query access to the victim, often for many millions of timesteps. Developing a better understanding of threats and opportunities from adversarial policies will be valuable as reinforcement learning systems are showing increasing potential for use in real-world applications. Thus, we set out to expand on the

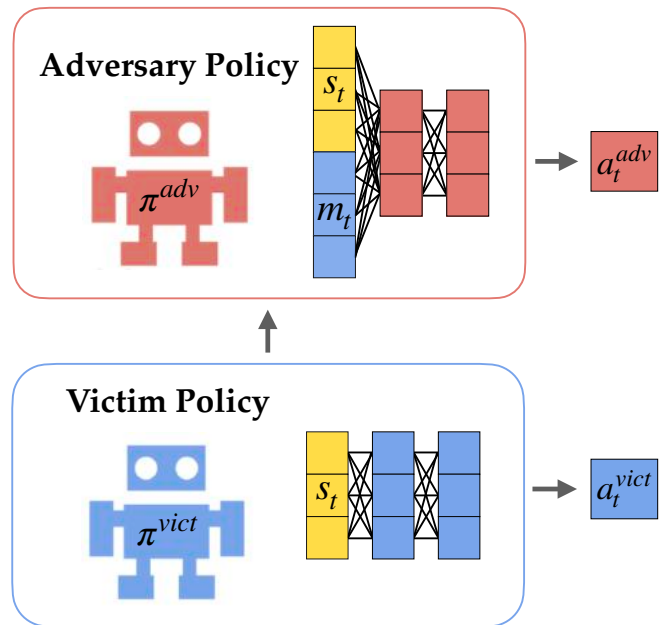


Fig. 1. White-box adversarial policies. At each timestep, both the adversary (adv) and victim (vict) observe the state s_t . The adversary also observes the internal state of the victim and concatenates this extra information, m_t , into its observations. We demonstrate how this type of white-box adversarial policy is more useful than black-box controls for attacks and adversarial training.

conventional threat model with adversarial policies that exploit richer information from the victim.

The analog to training a black-box adversarial policy in supervised learning would be to make a zero-order search through a model’s input space to find examples that make it fail. While black-box attacks like these have been studied in supervised learning [3], they are much less effective and query-efficient than white-box ones which permit access to the model’s internal state. Thus, here we study how using information from the victim can help an attacker learn an adversarial policy more quickly and effectively.

Our version of white-box attacks are adversarial policies that can “read the victim’s mind.” Fig. 1 depicts our general approach. At each timestep, both the adversary and victim observe the state s_t . The adversary, however, is also able to observe internal information, m_t , from the victim. In our experiments, m_t is a vector that consists of the victim’s

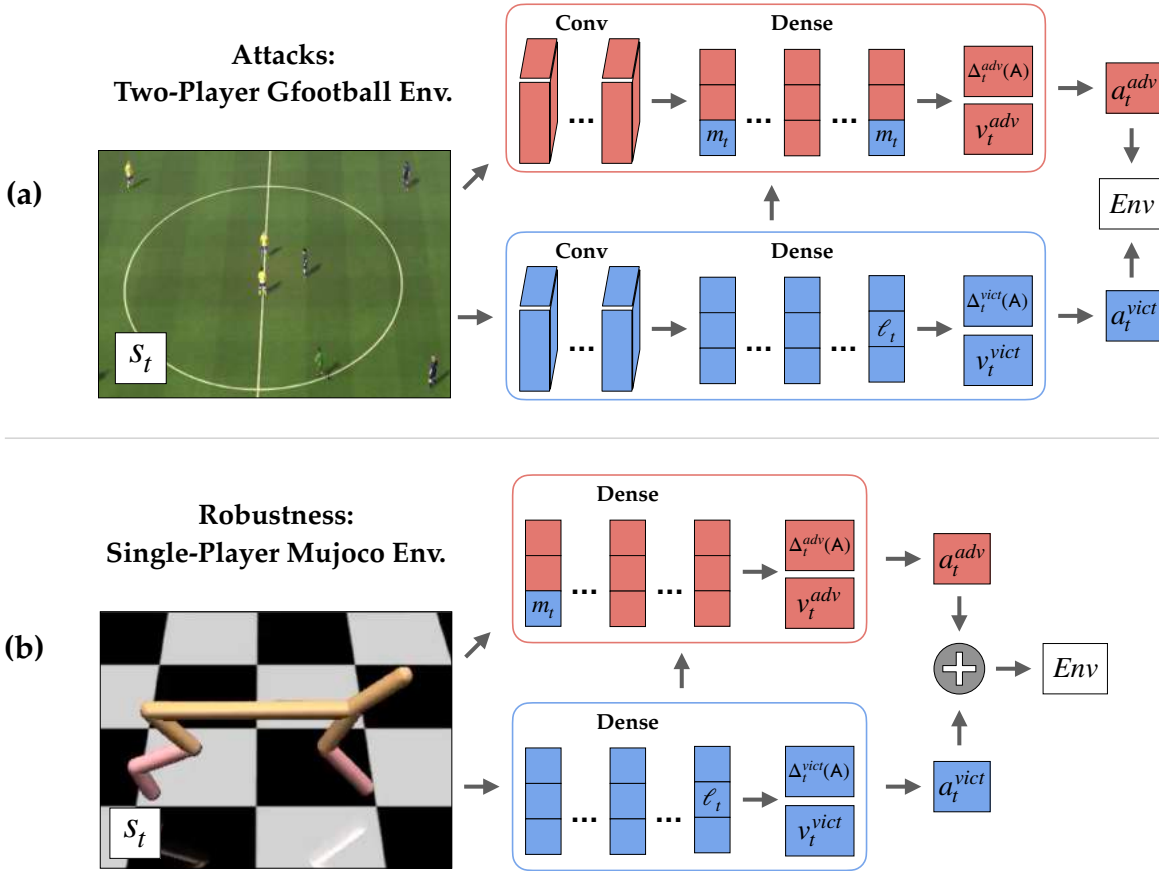


Fig. 2. Our setup for (a) adversarial attacks in the two-player Google Research Football (Gfootball) environment and (b) robust adversarial reinforcement learning (RARL) in single-player Mujoco environments. At each timestep, the state observation s_t is passed to the adversary and victim. The adversary is also given internal information m_t from the victim which is concatenated into its first and last dense layer for Gfootball and its observations for Mujoco environments. The vector m_t can include the victim’s action distribution $\Delta_t^{vict}(\mathcal{A})$, value estimate v_t^{vict} , and/or latent activations ℓ_t . For the two-player Gfootball environment, both actions are passed into the environment’s step function. For single-player Mujoco environment, the adversary’s action is added to the victim’s as a perturbation.

action distribution $\Delta_t^{vict}(\mathcal{A})$, value estimate v_t^{vict} , and/or latent activations ℓ_t .

Specifically, we test this approach in two different settings. First, we test adversarial *attacks* using the two-player Google Research Football (Gfootball) environment [28] and large convolutional policy networks. Both the adversary’s and victim’s actions are passed into the environment’s step function. This setup is illustrated in Fig. 2a. Here, we show that our white-box attackers have both higher initial and asymptotic performance than black-box baselines. Second, we adopt the robust adversarial reinforcement learning (RARL) approach from [36], [46] for experiments in single-player Mujoco environments (HalfCheetah and Hopper) [4] with small fully-connected policy networks. The adversary’s actions are added to the victim’s as a perturbation. This is shown in Fig. 2b. Here, we find that white-box adversaries can be more useful for training robust victims whose policies generalize better to environments with altered transition dynamics.

Given these results, we argue that adversarial policies that exploit inner information pose greater threats from attacks and greater opportunities from adversarial training. More generally,

our results demonstrate that observations from an agent’s internal state can be useful for other agents that interact with it. Following a discussion of related works in section II, Section III details our threat model and methods. Section IV presents results, and Section V a discussion. For a high-level explanation and summary, see the Appendix¹.

II. RELATED WORK

Adversarial Policies: Reinforcement learning agents can be vulnerable to several types of adversarial threats including input perturbations, action perturbations, reward perturbations, environments, and policies from other agents. Both [23] and [42] offer surveys of threats and defenses. Our focus is on adversarial policies. Conventionally, these attacks have been developed by simply training the adversary against the fixed victim policy. The approach has been used by [2], [12], [14], [16], [17], [47] for attacks. These adversaries were even observed unintentionally by [1] and [28] who found that in

¹Code for white-box RARL is available. https://github.com/thestephencasper/white_box_rarl.

competitive multiagent environments, it was key to rotate players in a round-robin fashion to avoid agents overfitting against a particular opponent. Additionally, [37] introduced a approach based on planning, [14] tested the detectability of adversarial policies, [9], [14] explored defense techniques via obfuscating the attacker and using option-based policies respectively, [8], [47] experimented with defense via adversarial training, and [12], [13] offered methods of attacking a victim whose reward is unknown.

Meanwhile, [33], [36], [41], [44], [46], [48] have studied Robust Adversarial Reinforcement Learning (RARL) in which an agent is trained alongside an adversary that perturb’s its body or actions in order for the agent to learn more robust control. [49] studied the stability of this approach. Others [32], [35], [39] have adversarially trained agents under observation or environment perturbations.

Black vs. White-box Attacks: In supervised learning, adversarial attacks are often easy to make with white-box access to the victim’s internal weights. Black-box attacks, however, typically require transfer, zero-order optimization, or gradient estimation, and are usually less successful [3]. Several others including [25], [26], [30], [32], [35] have studied attacks against reinforcement learners analogous to white-box ones in supervised learning. [27] further demonstrated the use of a victim’s internal state by using the value function for scheduling maximally-effective adversarial observation perturbations. These types of attacks perturb the victim’s observations and involve propagating the gradient for an adversarial objective through the policy network. In contrast, our white-box adversarial policies only differ from black-box ones from related work in whether the attacker, a reinforcement learner, can observe the victim’s internal state. Several works [10], [20], [29] have also trained agents with a theory of mind for their opponent in competitive tasks, but only in very simple environments. To our knowledge, we are the first to introduce policies which can exploit inner information from a victim in complex environments.

Open-Source Decision Making: We study victims whose policies are transparent to other agents in the environment. Agents with open source policies pose a number of challenges and pitfalls for decision-making. Several works formalize these challenges in the context of decision theory or game theory [5]–[7], [11], [19]. Our work adds to this by empirically studying one such challenge.

III. METHODS

A. Framework

We consider the goal of training an adversary against a victim inside of a two player Markov Decision Process (MDP) defined by a 6-tuple: $(\mathcal{S}, \{\mathcal{A}_{adv}, \mathcal{A}_{vict}\}, T, d_0, \{r_{adv}, r_{vict}\}, \gamma)$ with \mathcal{S} a state set, \mathcal{A}_{adv} and \mathcal{A}_{vict} action sets for the adversary and victim, $T : \mathcal{S} \times \mathcal{A}_{adv} \times \mathcal{A}_{vict} \rightarrow \Delta(\mathcal{S})$ a state transition function which outputs a distribution $\Delta(\mathcal{S})$ over \mathcal{S} , d_0 an initial state distribution, γ a temporal discount factor, and r_{adv} and r_{vict} reward functions for the adversary and victim s.t. $r_{adv}, r_{vict} : \mathcal{S} \times \mathcal{A}_{adv} \times \mathcal{A}_{vict} \times \mathcal{S} \rightarrow \mathcal{R}$. We assume

$r_{adv}(s) \approx -r_{vict}(s) \quad \forall s \in \mathcal{S}$. We only run experiments in which the victim’s policy is fixed, so the two-player MDP reduces to a single-player one. We will use $\pi_{adv} : \mathcal{S} \rightarrow \Delta(\mathcal{A}_{adv})$ and $\pi_{vict} : \mathcal{S} \rightarrow \Delta(\mathcal{A}_{vict})$ to denote the policy of an adversary and victim, and $V_{adv}^{\pi_{adv}}, V_{vict}^{\pi_{vict}} : \mathcal{S} \rightarrow \mathbb{R}$ to refer to their value functions.

B. Threat Model

There are multiple notions that have been used in supervised and reinforcement learning to characterize an adversary. These include being *effective* at making the victim fail, being *subtle* and hard for an observer to detect (e.g., [27]), and being *victim-specific* (e.g., [14]). Here, we use the first criterion and consider any policy that is *effective* at making another fail to be adversarial. For further discussion, see Appendix, A.

Previous works discussed in Section II have assumed a threat model in which the adversary only has black-box access to the victim but can cheaply train against it for many timesteps. We both strengthen and weaken this. First, we make the permissive assumption that the adversary can observe at least some of the victim’s internal state at each timestep and is able to use this information as an observation in the same timestep (see Section III-C for details). This could be a plausible assumption if a malicious attacker could obtain access to a victim agent’s policy parameters – especially if its designers make the victim open-source. Moreover, this will *always* be a realistic assumption for the agent’s designers if they want to test its robustness and/or adversarially train it. Second, we consider the restrictive assumption that the number of timesteps for which the adversary can train against the victim may be limited. Realistically, this could be the case if the victim’s designers limit access to it or if gathering experience is costly.

C. White-Box Adversarial Policies

We train policies using Proximal Policy Optimization (PPO) [40] and Soft Actor Critic (SAC) [18]. Both involve training a value function estimator alongside the policy. We consider attackers that have access to (1) the victim’s action outputs, (2) the victim’s value estimate, and/or (3) the internal activations from the victim’s policy network. Our goal for (1) is to give the adversary a glimpse of the near future so that it can better counter the victim’s behavior. Our goal for (2) is to make it easier for the attacker to quickly learn its own value function because $V_{vict}^{\pi_{vict}}(s_t) \approx -V_{adv}^{\pi_{adv}}(s_t)$. Note this is only possible for victims that have a critic. Finally, our goal for (3) is to give the adversary rich and generally-useful information on how the victim represents the state.

At timestep t , the environment state, s_t , is observed. The victim processes the state and produces an action $a_t^{vict} \sim \pi_{vict}(s_t)$. At the same time, the white-box adversary queries the victim to get its action output $\pi_{vict}(s_t)$, value estimate $V_{vict}(s_t)$, and/or latents $\ell_{vict}(s_t)$ in the form of a vector $m(s_t)$. In a slight abuse of notation, we refer to $\ell_{vict}(s_t)$ as ℓ_t and $m(s_t)$ as m_t . Thus, the adversary’s policy function can be

written as $\pi_{adv}(s_t) = f(s_t, m_t)$, and its value estimate can be written as $V_a^{\pi_a}(s_t) = g(s_t, m_t)$.

We train both adversaries that use large convolutional neural networks (CNNs) and small multilayer perceptrons (MLPs) as policy networks. These architectures are illustrated in Fig. 2. For the large CNNs, we concatenate m_t into the representation of the state twice: once at the first fully-connected layer, and once at the last. We do this so that the adversary can readily learn both complex and simple functions of m_t . In particular, we hypothesized that giving the adversary the victim’s value estimate in its final layer is helpful for learning its own value estimator, which ought to be approximately the negative of the victim’s. For the small MLPs policy networks, we only concatenate m_t with the observation once at the beginning for efficiency.

IV. EXPERIMENTS

A. Stronger Attacks

Environment: We use the two-player Google Research Football environment (Gfootball) [28]. Each agent in the environment controls a team of 11 football (soccer) players. The states are $72 \times 96 \times 4$ pixels with the four channels encoding the left team positions, right team positions, ball position, and active player position. Observations were stacked over four timesteps to give the agents a perception of time, resulting in observations of $72 \times 96 \times 16$ pixels. The agents’ policy networks had a ResNet architecture [21], and the action space was discrete with size 19. We used the same reward shaping as in [28] in which an agent was rewarded 1 for scoring, -1 for being scored on, and 0.1 for advancing the ball one tenth of the way down the field. We trained all Gfootball agents using Proximal Policy Optimization [40] using the Stable Baselines 2 implementation [22].

Victims: First, we trained victims to develop adversarial policies against. For Gfootball, this was done in two stages for a total of 50 million timesteps. First, the victims were trained against a ‘bot’ agent for 25 million timesteps with an entropy reward to encourage exploration. Second, they were trained for another 25 million timesteps against an agent from the first phase with an entropy penalty to encourage more deterministic play. We found this to result in more consistent behavior from adversaries. In Fig. 3 (a) shows the learning curves for these victims.

Adversaries: We trained four types of adversaries, each of which uses observes different information, m_t , from the victim’s internal state:

- 1) **Black-Box Control:** $m_t = \emptyset$. This is the same threat model used by [1], [14] and others mentioned in Section II.
- 2) **Action & Value:** $m_t = V_{vict}(s_t) \oplus \pi_{vict}(s_t)$ where \oplus is the concatenation operator. Here, the adversary sees the scalar value and an $|\mathcal{A}_{vict}|$ -sized observation giving the victim’s distribution over discrete output actions.
- 3) **Latent:** $m_t = \ell_t$ where ℓ_t gives the latent activations from some layer during the forward pass through the

victim’s network from s_t . Here, we use those of the final layer from which both the victim’s actions and value are computed.

- 4) **Full:** $m_t = V_{vict}(s_t) \oplus \pi_{vict}(s_t) \oplus \ell_t$. This combines the Action & Value and Latent threat models.

Results: We train each adversary for 50 million timesteps. Fig. 3b shows the training curves for these attackers. All improve significantly over the black box control, both by having faster initial learning and a higher asymptotic performance. The two types of white-box adversaries that could observe the victim’s latents performed the best. For the action/value, latent, and full attacks, the p values from a one-sided t test for the hypothesis that they were superior to the black box controls were 0.00638, 0.00001, and 0.00002 respectively, demonstrating clear improvements.

B. Improved Robustness

Environment: To evaluate white-box robust adversarial reinforcement learning (RARL), we used HalfCheetah-v3 and Hopper-v3 Mujoco environments from OpenAI Gym. [4]. In both environments, the agent controls a body in a 3D simulated physics environment. Observations are continuous-valued vectors specifying the position of the body, and actions are continuous-valued vectors for controlling it. The agents’ policy networks had a small MLP architecture with two hidden layers of 256 neurons each. We trained all gym agents using SAC [18] with the Stable Baselines 3 implementation [38].

Training: In alternation, we trained a protagonist agent and an ensemble of three adversaries who perturbed the protagonist’s actions. For each training episode for the protagonist, a random adversary from the three was chosen to make the perturbations. We experiment with three methods:

- 1) **RL Control:** An agent is trained with no adversary.
- 2) **RARL:** An agent is trained against an ensemble of black-box adversarial agents. This is the approach used by [46].
- 3) **Latent/Action White-Box RARL (WB-RARL):** An agent is trained against an ensemble of white-box adversaries that each observe its latent activations and action outputs. Thus, $m_t = \pi_{vict}(s_t) \oplus \ell_t$

Results: We trained a total of 40 agents of each type for 2 million timesteps and selected the 20 with the best final performance. Fig. 4a shows the evaluation performance for the HalfCheetah and Hopper agents in an adversary-free environment over the course of training. Performance is comparable between all three conditions with the RL controls seeming to perform the best in HalfCheetah.

To test the robustness of the learned policies, we then test on a set of adversary-free environments with the transition dynamics altered. We selected a range of 8 mass and 8 friction coefficients to modify the environment dynamics by and tested the agents on all 8×8 combinations. The full arrays of results are shown in Fig. 5 in Appendix B. And the mean results over all friction coefficients and mass coefficients are

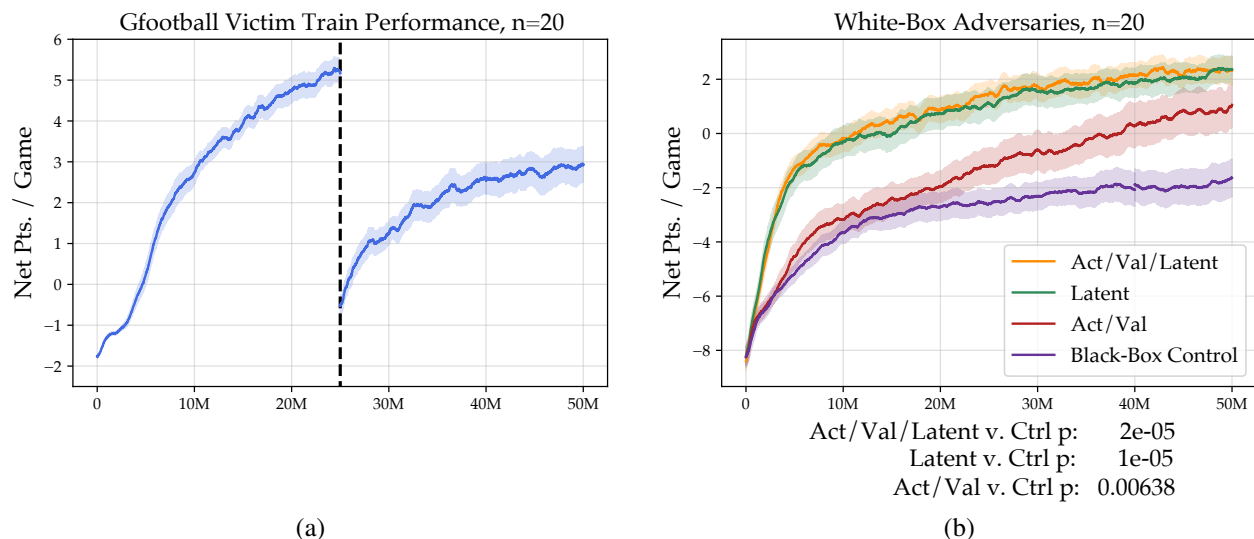


Fig. 3. Results for white-box adversarial attacks. (a) Training curves for Gfootball victims. The curves give the mean and standard error of the mean across $n = 20$ victims. The first 25 million timesteps of training is against a rule-based “bot,” and the action entropy is rewarded while the second 25 million timesteps is against a peer and the action entropy is penalized. (b) Learning curves over 50 million timesteps for various adversarial attackers against the victims from (a) starting from random initialization. The top three curves show the performance of white-box adversaries with access to the victim’s action distribution and value estimate and/or its latent activations. The bottom shows a black-box control. As in (a), the curves give the mean and standard error of the mean across $n = 20$ victims. Three p value are shown below giving the results of a one-sided t test for the hypothesis that each white-box agent beat the black-box control.

plotted in Fig. 4b-c respectively. In Fig. 4b-c, WB-RARL agents generally perform as well or better than the other two. And on average, WB-RARL performs the best over all testing environments. For RL, RARL, and WB-RARL, the HalfCheetah agents achieve mean episode rewards of 902, 914, and 1019, and the Hopper agents achieve 673, 645, and 716 respectively. We performed four one-sided t -tests to test the hypotheses that the WB-RARL agents had superior overall testing performance. For HalfCheetah, the p values were 0.085 and 0.111 for comparing the WB-RARL agents to the RL and RARL ones respectively. For Hopper, the corresponding p values were 0.095 and 0.009. These suggest it is likely that the WB-RARL agents are more robust to these domain shifts.

V. DISCUSSION AND BROADER IMPACT

Our goal in this work is to better understand threats and opportunities from adversarial policies in reinforcement learning by studying white-box adversarial attackers. We show that allowing an adversarial policy to observe the internal state of the victim, can result in (1) better initial and asymptotic performance for adversarial attackers and (2) more effective adversarial training for improving the robustness of a learned policy.

More generally, our results show that information about an agent’s internal state offers useful information for other agents interacting with it. This may be the case regardless of whether the setting is adversarial, cooperative, or apathetic. In multiagent settings, it is key to bear in mind that a policy which makes use of white-box information from another agent need not be implemented *by* nor *against* a conventional reinforcement learner. On one hand, policies can be developed

without standard reinforcement learning approaches (e.g., PPO or SAC). For example, human video game players constantly develop strategies to exploit the weaknesses of computer-controlled competitors to great effect. On the other hand, so long as a target agent computes actions via latent information which information could be given to other agents.

Concerning adversarial attacks in particular, one risk of any work that focuses on attack methods is that they could be used for malicious attacks. This is an important concern, but we emphasize that it is better to develop an understanding of adversarial vulnerabilities through exploratory research than from incidents in the real world. We also stress the benefits of adversarial training. Our findings should encourage caution and robustness measures when developing reinforcement learning systems that may vulnerable to these types of attacks. In particular, these should include restricting access to white-box information from agents.

A limitation is that while we show that white-box attacks can be useful, they may be of limited practical relevance. One reason is that for our experiments with RARL, the improvements from granting the adversary white-box access were only modest. Another is that white-box access may often be difficult to obtain in the first place. And even though white-box attacks can help train adversarial policies more quickly, these attacks may still demand many millions of timesteps. Future work on similar black-box attacks that use a model of the victim learned from black-box (and potentially even offline) access may be valuable. Studying ways to more effectively leverage victim information in fewer training timesteps may also be useful. Additional progress like this toward better understanding threats and opportunities from adversaries

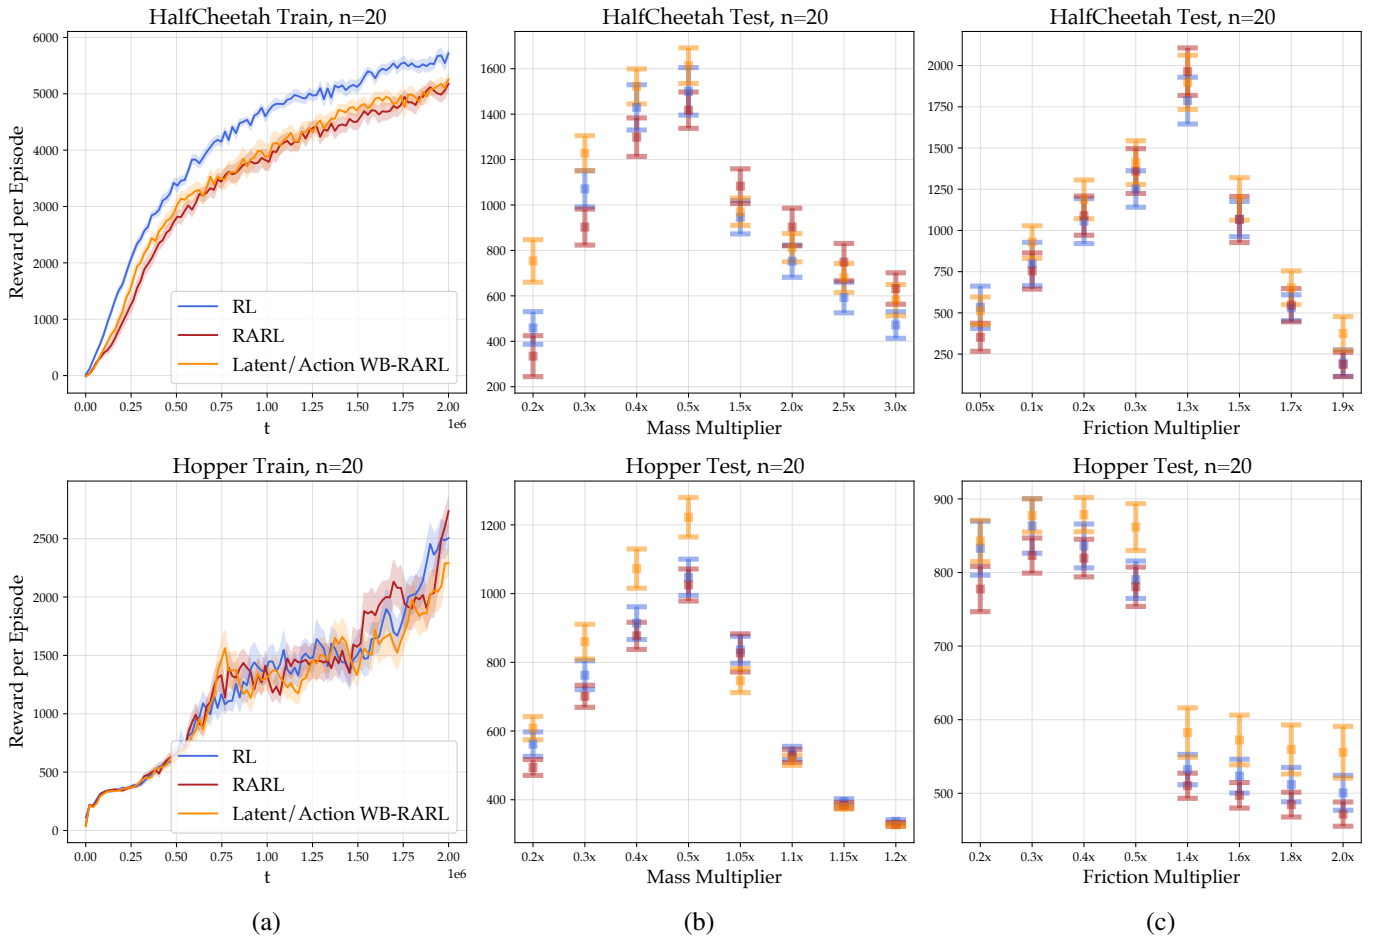


Fig. 4. Results for white-box adversarial training. Training and testing performance for (top) HalfCheetah and (bottom) Hopper agents. (a) Performance over training for robust adversarial reinforcement learning (RARL) experiments. Results are obtained from adversary-free testing environments. The curves show the mean and standard error of the mean across $n = 20$ agents. We then tested the final agents across a range of environments with perturbed mass and friction coefficients. The full results are shown in Fig. 5 in Appendix B. Here, (b-c) show the mean and standard error of the mean for testing results averaged across the friction and mass coefficients respectively. Again, all errorbars show standard error of the mean across $n = 20$ agents. In general, agents trained with white-box adversarial training perform as well or better than controls.

in reinforcement learning will be a promising direction for expanding the toolbox for more trustworthy AI.

REFERENCES

- [1] Trapit Bansal, Jakub Pachocki, Szymon Sidor, Ilya Sutskever, and Igor Mordatch. Emergent complexity via multi-agent competition. *arXiv preprint arXiv:1710.03748*, 2017.
- [2] Vahid Behzadan and William Hsu. Adversarial exploitation of policy imitation. *arXiv preprint arXiv:1906.01121*, 2019.
- [3] Siddhant Bhambrani, Sumanyu Muku, Avinash Tulasi, and Arun Balaji Buduru. A survey of black-box adversarial attacks on computer vision models. *arXiv preprint arXiv:1912.01667*, 2019.
- [4] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- [5] Stephen Casper. Achilles heels for agi/asi via decision theoretic adversaries. *arXiv preprint arXiv:2010.05418*, 2020.
- [6] Andrew Critch. A parametric, resource-bounded generalization of l b’s theorem, and a robust cooperation criterion for open-source game theory. *The Journal of Symbolic Logic*, 84(4):1368–1381, 2019.
- [7] Andrew Critch, Michael Dennis, and Stuart Russell. Cooperative and uncooperative institution designs: Surprises and problems in open-source game theory. *arXiv preprint arXiv:2208.07006*, 2022.
- [8] Pavel Czempin and Adam Gleave. Reducing exploitability with population based training. *arXiv preprint arXiv:2208.05083*, 2022.
- [9] Prithviraj Dasgupta. Using options to improve robustness of imitation learning against adversarial attacks. In *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications III*, volume 11746, page 1174610. International Society for Optics and Photonics, 2021.
- [10] Aaron Davidson. Using artificial neural networks to model opponents in texas hold’em. *Unpublished manuscript*, 1999.
- [11] Abram Demski and Scott Garrabrant. Embedded agency. *arXiv preprint arXiv:1902.09469*, 2019.
- [12] Ted Fujimoto, Timothy Doster, Adam Attarian, Jill Brandenberger, and Nathan Hodas. The effect of antagonistic behavior in reinforcement learning. 2021.
- [13] Ted Fujimoto, Timothy Doster, Adam Attarian, Jill Brandenberger, and Nathan Hodas. Reward-free attacks in multi-agent reinforcement learning. *arXiv preprint arXiv:2112.00940*, 2021.
- [14] Adam Gleave, Michael Dennis, Neel Kant, Cody Wild, Sergey Levine, and Stuart Russell. Adversarial policies: Attacking deep reinforcement learning. *arXiv preprint arXiv:1905.10615*, 2019.
- [15] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- [16] Jun Guo, Yonghong Chen, Yihang Hao, Zixin Yin, Yin Yu, and Simin Li. Towards comprehensive testing on the robustness of cooperative multi-agent reinforcement learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 115–

- 122, 2022.
- [17] Wenbo Guo, Xian Wu, Sui Huang, and Xinyu Xing. Adversarial policy learning in two-player competitive games. In *International Conference on Machine Learning*, pages 3910–3919. PMLR, 2021.
- [18] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*, 2018.
- [19] Joseph Y Halpern and Rafael Pass. Game theory with translucent players. *International Journal of Game Theory*, 47(3):949–976, 2018.
- [20] He He, Jordan Boyd-Graber, Kevin Kwok, and Hal Daumé III. Opponent modeling in deep reinforcement learning. In *International conference on machine learning*, pages 1804–1813. PMLR, 2016.
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [22] Ashley Hill, Antonin Raffin, Maximilian Ernestus, Adam Gleave, Anssi Kanervisto, Rene Traore, Prafulla Dhariwal, Christopher Hesse, Oleg Klimov, Alex Nichol, Matthias Plappert, Alec Radford, John Schulman, Szymon Sidor, and Yuhuai Wu. Stable baselines. <https://github.com/hill-a/stable-baselines>, 2018.
- [23] Inaam Ilahi, Muhammad Usama, Junaid Qadir, Muhammad Umar Janjua, Ala Al-Fuqaha, Dinh Thai Huang, and Dusit Niyato. Challenges and countermeasures for adversarial attacks on deep reinforcement learning. *IEEE Transactions on Artificial Intelligence*, 2021.
- [24] Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Logan Engstrom, Brandon Tran, and Aleksander Madry. Adversarial examples are not bugs, they are features. *arXiv preprint arXiv:1905.02175*, 2019.
- [25] Ezgi Korkmaz. Adversarially trained neural policies in the fourier domain. In *ICML 2021 Workshop on Adversarial Machine Learning*, 2021.
- [26] Ezgi Korkmaz. Investigating vulnerabilities of deep neural policies. In *Uncertainty in Artificial Intelligence*, pages 1661–1670. PMLR, 2021.
- [27] Jernej Kos and Dawn Song. Delving into adversarial attacks on deep policies. *arXiv preprint arXiv:1705.06452*, 2017.
- [28] Karol Kurach, Anton Raichuk, Piotr Stańczyk, Michał Zajac, Olivier Bachem, Lasse Espeholt, Carlos Riquelme, Damien Vincent, Marcin Michalski, Olivier Bousquet, et al. Google research football: A novel reinforcement learning environment. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 4501–4510, 2020.
- [29] Alan J Lockett, Charles L Chen, and Risto Miikkulainen. Evolving explicit opponent models in game playing. In *Proceedings of the 9th annual conference on Genetic and evolutionary computation*, pages 2106–2113, 2007.
- [30] Björn Lütjens, Michael Everett, and Jonathan P How. Certified adversarial robustness for deep reinforcement learning. In *Conference on Robot Learning*, pages 1328–1337. PMLR, 2020.
- [31] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- [32] Tuomas Oikarinen, Wang Zhang, Alexandre Megretski, Luca Daniel, and Tsui-Wei Weng. Robust deep reinforcement learning through adversarial loss. *Advances in Neural Information Processing Systems*, 34, 2021.
- [33] Xinlei Pan, Daniel Seita, Yang Gao, and John Canny. Risk averse robust adversarial reinforcement learning. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 8522–8528. IEEE, 2019.
- [34] Nicolas Papernot, Patrick McDaniel, and Ian Goodfellow. Transferability in machine learning: from phenomena to black-box attacks using adversarial samples. *arXiv preprint arXiv:1605.07277*, 2016.
- [35] Anay Pattanaik, Zhenyi Tang, Shuijing Liu, Gautham Bommanan, and Girish Chowdhary. Robust deep reinforcement learning with adversarial attacks. *arXiv preprint arXiv:1712.03632*, 2017.
- [36] Lerrel Pinto, James Davidson, Rahul Sukthankar, and Abhinav Gupta. Robust adversarial reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2817–2826. JMLR. org, 2017.
- [37] Alberto Pozanco, Susana Fernández, Daniel Borrajo, et al. Anticipatory counterplanning. *arXiv preprint arXiv:2203.16171*, 2022.
- [38] Antonin Raffin, Ashley Hill, Adam Gleave, Anssi Kanervisto, Maximilian Ernestus, and Noah Dormann. Stable-baselines3: Reliable reinforcement learning implementations. *Journal of Machine Learning Research*, 22(268):1–8, 2021.
- [39] Lucas Schott, Manon Césaire, Hatem Hajri, and Sylvain Lamprier. Improving robustness of deep reinforcement learning agents: Environment attacks based on critic networks. *arXiv preprint arXiv:2104.03154*, 2021.
- [40] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [41] Hiroaki Shioya, Yusuke Iwasawa, and Yutaka Matsuo. Extending robust adversarial reinforcement learning considering adaptation and diversity. 2018.
- [42] Samuel Henrique Silva and Peyman Najafirad. Opportunities and challenges in deep learning adversarial robustness: A survey. *arXiv preprint arXiv:2007.00753*, 2020.
- [43] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- [44] Kai Liang Tan, Yasaman Esfandiari, Xian Yeow Lee, Soumik Sarkar, et al. Robustifying reinforcement learning agents via action space adversarial training. In *2020 American control conference (ACC)*, pages 3959–3964. IEEE, 2020.
- [45] Florian Tramèr, Nicolas Papernot, Ian Goodfellow, Dan Boneh, and Patrick McDaniel. The space of transferable adversarial examples. *arXiv preprint arXiv:1704.03453*, 2017.
- [46] Eugene Vinitzky, Yuqing Du, Kanaad Parvate, Kathy Jang, Pieter Abbeel, and Alexandre Bayen. Robust reinforcement learning using adversarial populations. *arXiv preprint arXiv:2008.01825*, 2020.
- [47] Xian Wu, Wenbo Guo, Hua Wei, and Xinyu Xing. Adversarial policy training against deep reinforcement learning. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 1883–1900, 2021.
- [48] Peng Zhai, Jie Luo, Zhiyan Dong, Lihua Zhang, Shunli Wang, and Dingkan Yang. Robust adversarial reinforcement learning with dissipation inequation constraint. 2022.
- [49] Kaiqing Zhang, Bin Hu, and Tamer Basar. On the stability and convergence of robust adversarial reinforcement learning: A case study on linear quadratic systems. *Advances in Neural Information Processing Systems*, 33:22056–22068, 2020.

VI. ACKNOWLEDGMENTS

We thank Lucas Janson for valuable ideas and feedback throughout the course of this work.

APPENDIX

A. Understanding Adversarial Policies

The notion of an *adversary* for a deep learning system was popularized by [15], [43] and subsequent research. These works developed adversarial images that are both *effective*, meaning that they fool an image classifier, and *subtle*, meaning that they only differ from a benign image by a very small-norm perturbation. While they often transfer to other models [24], [31], [34], [45], these adversaries are also typically *victim-specific* in the sense that they are created specifically to fool a particular model.

As in supervised learning, “effectiveness” is used as part of the definition for adversarial policies across the literature. “Victim-specificity” sometimes is, but many RL works (e.g., [2]) including ours do not require an adversary to be victim-specific. Finally, “subtlety” has not been adopted as a standard for adversaries research in RL. A notion of subtlety for adversaries in RL that would be analogous to supervised learning would be that the adversary produces distributions over actions or trajectories that are very similar to a benign agent. However, in this and all related work in RL of which we know, no notion of subtlety is part of the definition of an adversarial policy. So ultimately, we use “adversarial” here to simply refer to a policy which is good at beating a victim.

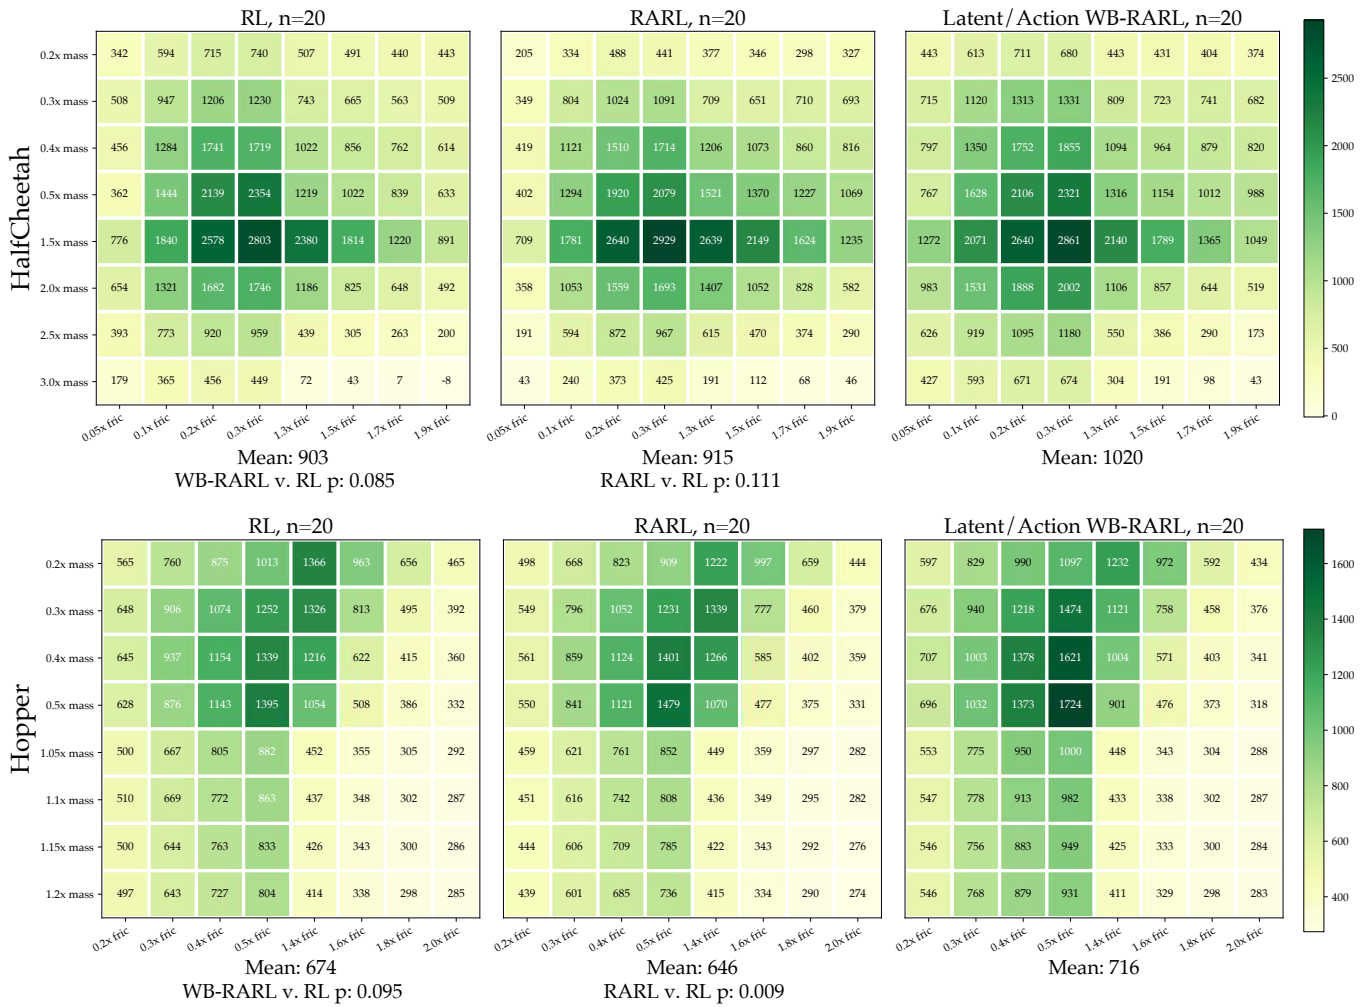


Fig. 5. Evaluations for Robust Adversarial Reinforcement Learning Experiments for $n = 20$ agents with (top) HalfCheetah and (bottom) Hopper agents. Each grid shows mean episode reward for adversary-free environments with the mass and friction coefficients altered. Under each grid in the grid is displayed. Under the RL and RARL grids cols 1 and 2), the one-sided p value for the hypothesis that WB-RARL is superior to RL and RARL is shown.

B. Full Robust Adversarial Reinforcement Learning Results

As discussed in Section IV-B, we tested agents on environments with altered mass and friction parameters. For both the HalfCheetah and Hopper environments, we used a set of 8×8 different mass and friction values. Testing results across all testing environments for control, RARL, and WB-RARL agents are shown here in Fig. 5. Under each grid, the mean for all results in the grid is displayed. Under the RL and RARL grids (columns 1 and 2), the p value from a one-sided t-test for the hypothesis that WB-RARL is superior to RL and RARL is shown.

C. High-Level Summary

Here, we provide a summary of this work which does not assume that the reader has a technical background.

“Reinforcement Learning” (RL) is the process by which an agent learns via some formalized process of trial and error to accomplish a goal. Humans are reinforcement learners.

And so are some algorithms that are commonly studied in machine learning research today. For example, is common to use reinforcement learning algorithms to train AI systems to play video games. Using experience, they can infer what types of actions lead to higher scores and adjust their behavior accordingly.

Multiagent RL describes settings in which there is more than one agent acting in some setting. Past research has shown that in multiagent settings, training “adversarial” reinforcement learners to make other reinforcement learners fail can be useful. One one hand, an adversarial agent can often learn to act in a way that renders the “victim” agent unable to accomplish its goals. For example, an adversary can sometimes act in ways that make a victim in a two player video game seem to take actions that are as bad as – or even worse than – random ones. On the other hand, training a victim against an adversarial agent can make it much more robust to some failures. For example, this might make the victim particularly

effective at avoiding failures due to changes to its environment.

In this work, we study a new approach to adversarial attacks and adversarial training in RL. We experiment with “white-box” attacks in which the adversary can observe the internal state of the victim. For humans, this would be analogous to one person playing a game against someone else while being able to view scans of their brain. We show that these white-box adversarial agents are more effective than controls for both attacks and adversarial training. We argue that this helps us to better understand threats and opportunities from adversarial RL. And based on these results, we call for increased caution and more effective robustness measures when deploying RL systems in the real world.