Minimalistic attacks : how little it takes to fool deep reinforcement learning policies

Qu, Xinghua; Sun, Zhu; Ong, Yew-Soon; Gupta, Abhishek; Wei, Pengfei

2020

Qu, X., Sun, Z., Ong, Y., Gupta, A. & Wei, P. (2020). Minimalistic attacks : how little it takes to fool deep reinforcement learning policies. IEEE Transactions On Cognitive and Developmental Systems. <https://dx.doi.org/10.1109/TCDS.2020.2974509>

<https://hdl.handle.net/10356/153700>

<https://doi.org/10.1109/TCDS.2020.2974509>

© 2020 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The published version is available at: <https://doi.org/10.1109/TCDS.2020.2974509>.

Downloaded on 26 Apr 2023 21:38:29 SGT

Minimalistic Attacks: How Little it Takes to Fool Deep Reinforcement Learning Policies

Xinghua Qu, *Student Member, IEEE*, Zhu Sun, Yew Soon Ong, *Fellow, IEEE*, Abhishek Gupta, Pengfei Wei

Abstract—Recent studies have revealed that neural network based policies can be easily fooled by adversarial examples. However, while most prior works analyze the effects of perturbing every pixel of every frame assuming white-box policy access, in this paper we take a more restrictive view towards adversary generation - with the goal of unveiling the limits of a model's vulnerability. In particular, we explore minimalistic attacks by defining *three key settings*: (1) *black-box policy access*: where the attacker only has access to the input (state) and output (action probability) of an RL policy; (2) *fractional-state adversary*: where only several pixels are perturbed, with the extreme case being a single-pixel adversary; and (3) *tactically-chanced attack*: where only significant frames are tactically chosen to be attacked. We formulate the adversarial attack by accommodating the three key settings, and explore their potency on six Atari games by examining four fully trained state-of-the-art policies. In Breakout, for example, we surprisingly find that: (i) all policies showcase significant performance degradation by merely modifying 0.01% of the input state, and (ii) the policy trained by DQN is totally deceived by perturbing only 1% frames.

Index Terms—Reinforcement Learning, Adversarial Attack.

I. INTRODUCTION

Deep learning [1] has been widely regarded as a promising technique in reinforcement learning (RL), where the goal of an RL agent is to maximize its expected accumulated reward by interacting with a given environment. Although deep neural network (DNN) policies have achieved noteworthy performance on various challenging tasks (e.g., video games, robotics and classical control [2]), recent studies have shown that these policies are easily deceived under adversarial attacks [3]–[5]. Nevertheless, existing works towards adversary generation are found to make some common simplifying assumptions, namely (1) white-box policy access: where the adversarial examples are computed by back-propagating targeted changes in the prescribed action distribution through known neural network architectures and weights, (2) full-state adversary: where the adversary may perturb almost all pixels

Fig. 1: The visualization of the single pixel attack on Breakout.

of an image representing an input state, and (3) fully-chanced attack: where the attacker strikes the policy at every time step (e.g., at every frame of a video game).

Given that most prior works analyze the effects of perturbing every pixel of every frame assuming white-box policy access, in this paper we propose to take a more restrained view towards adversary generation - with the goal of exploring the limits of a DNN model's vulnerability in RL. We thus focus on *minimalistic attacks* by only considering adversarial examples that perturb limited number of pixels in selected frames, and under restricted black-box policy access. In other words, we intend to unveil how little it really takes to successfully fool state-of-the-art RL policies. In what follows, we briefly introduce the three restrictive settings that characterize our study.

Black-box Policy Access (BPA). Most previous studies focus on a white-box setting [5], that allows full access to a policy network for back-propagation. However, most systems do not release their internal configurations (i.e., network structure and weights), only allowing the model to be queried; this makes the white-box assumption too optimistic from an attacker's perspective [6]. In contrast, we use a BPA setting where the attacker treats the policy as an oracle with access only to its input and outputs.

Fractional-State Adversary (FSA). In the FSA setting, the adversary only perturbs a small fraction of the input state. This, in the extreme situation, corresponds to the single-pixel attack shown in Fig. 1, where perturbing a single pixel of the input state is found to change the action prescription from 'RIGHT' to 'LEFT'. In contrast, most previous efforts [5] are mainly based on a full-state adversary (i.e., the number of modified pixels is fairly large, usually spanning the entire frame).

Tactically-Chanced Attack (TCA). In previously studied RL adversarial attacks [3], [7], [8], the adversary strikes the policy

X. Qu, is with Computational Intelligence Lab, School of Computer Science and Engineering, Nanyang Technological University, Singapore, 639798 (email: xinghua001@e.ntu.edu.sg)

Z. Sun is with School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore 639798 (email: sunzhuntu@gmail.com)

Y. S. Ong is Director of the Data Science and Artificial Intelligence Research Centre, School of Computer Science and Engineering, Nanyang Technological University, Singapore 639798 (email: asysong@ntu.edu.sg) and Chief Artificial Intelligence Scientist of Singapore's Agency for Science, Technology and Research (email: Ong_Yew_Soon@hq.a-star.edu.sg)

A. Gupta is with Singapore Institute of Manufacturing Technology (SIMTech), A*STAR, Singapore 138634 (email: abhishek_gupta@simtech.a-star.edu.sg)

P. Wei is with Department of Computer Science, National University of Singapore, Singapore 119077 (email: wpf89928@gmail.com)

on every frame of an episode; this is a setting termed as the fully-chanced attack. Contrarily, we investigate a relatively restrictive case where the attacker only strikes at a few selected frames - a setting we term as *tactically-chanced attack*, where a minimal number of frames is explored to strategically deceive the policy network.

The proposed restrictive settings are deemed to be significant for safety-critical RL applications, such as in the medical treatment of sepsis patients in intensive care units [9], [10] and treatment strategies for HIV [11]. In such domains, there may exist a temporal gap between the acquisition of the input and the action execution; thereby, providing a time window in which a tactical attack could take place. In the future, as we move towards integrating end-to-end vision-based systems for healthcare, the sensitivity of prescribed medical actions (such as drug prescription or other medical interventions) to perturbations in medical images poses a severe threat to the utility of RL in such domains. Moreover, the restrictive nature of the attacks makes them practically imperceptible, greatly reducing the chance of identifying and rectifying them. With this in mind, the present paper explores the vulnerability of deep RL models to restrictive adversarial attacks, with particular emphasis on image-based tasks (i.e., Atari games).

The key challenge of this research, therefore, lies in assessing how to effectively accommodate the three restrictive settings for adversarial attack generation in deep RL. To this end, we design a mathematical program with a novel objective function to generate the FSA under black-box settings, and propose an entropy-based uncertainty measurement to achieve the TCA. The optimization variables are defined as the discrete 2-D coordinate location(s) and perturbation value(s) of the selectively attacked pixel(s), and the designed mathematical program guarantees a successful deception of the policy as long as a positive objective value is found. The optimization procedure, for a given input frame, is then carried out by a simple genetic algorithm (GA) [12], [13]. Furthermore, the Shannon entropy of the action distribution is specified as the attack uncertainty, and is utilized to only select a few salient frames to be tactically attacked. We demonstrate the sufficiency of the three restrictive settings on four state-of-the-art deep RL policies pre-trained on six Atari games.

To sum up, the contributions of this paper can be encapsulated as follows:

- We unveil how little it takes to deceive an RL policy by considering three restrictive settings, namely, black-box policy access, fractional-state adversary, and tactically-chanced attack;
- We formulate the RL adversarial attack as a black-box optimization problem comprising a novel objective function and discrete FSA optimization variables. We also propose a Shannon entropy-based uncertainty measurement to sparingly select the most vulnerable frames to be attacked;
- We explore the three restrictive settings on the policies trained by four state-of-the-art RL algorithms (i.e., DQN [1], PPO [14], A2C [15], ACKTR [16]) on six representative Atari games (i.e., Pong, Breakout, SpaceInvaders, Seaquest, Qbert, BeamRider);

- Surprisingly, we find that on Breakout: (i) with only *a single pixel* ($\approx 0.01\%$ of the state) attacked, the trained policies show significant performance degradation; and (ii) by merely attacking around 1% frames, the policy trained by DQN is significantly deceived.

The remainder of this paper is organized as follows. In Section II, we provide an overview of related works in RL adversarial attacks, highlighting the novelty of our paper. Section III presents preliminaries on reinforcement learning and adversarial attacks. In section IV, we describe our proposed optimization problem formulation and the associated procedure to generate adversarial attacks for deep RL policies. This is followed by Section V, where we report numerical results demonstrating the effectiveness of our attacks. Finally, Section VI concludes this paper.

II. RELATED WORK

Since Szegedy et al. [17], a number of adversarial attack methods have been investigated to fool deep neural networks (DNNs). Most existing approaches generate adversarial examples through perturbations to images that are imperceptible to the human eye [18]–[20], with the goal of deceiving trained classifiers. As such, we find that in the recent past, attention has largely been focused on adversarial attacks towards supervised learning tasks. In contrast, adversarial attacks in RL have been relatively less explored to date.

In RL, Huang et al. [3] were among the first to demonstrate that neural network policies are vulnerable to adversarial attacks by adding small modifications to the input state of Atari games. A full-state adversary (in which an adversarial example may perturb every pixel in the input state) has previously been generated by white-box policy access [18], where adversarial examples are computed by back-propagating targeted changes to the action distribution prescribed by the policy. Lin et al. [4] proposed strategically-timed attacks and the so-called enchanting attack, but the adversary generation is still based on a white-box policy access assumption and full-state perturbation. Apart from the above, Kos et al. [21] compared the influence of full-state perturbations with random noise, and utilized the value function to guide the adversary injection.

In summary, existing works are largely based on white-box policy access, together with assumptions of a full-state adversary and fully-chanced attack. There is little research studying the potency of input perturbations that may be far less extensive. Therefore, our goal in this paper is to investigate a significantly more restrained view towards analyzing the vulnerability of deep RL models, viz., based on black-box policy access, fractional-state adversary, and tactically-chanced attacks. Additionally, it is contended that studying such restrictive scenarios might give new insights on the geometrical characteristics and overall behavior of deep neural networks in high dimensional spaces [22].

III. PRELIMINARIES

This section first provides some background on RL and several representative approaches to learn RL policies. It then illustrates the basic concepts of adversarial attacks.

Fig. 2: Markov Decision Process in Reinforcement Learning: according to the state s_t , the agent selects an action a_t , and then receives the reward r_{t+1} from the environment.

A. Reinforcement Learning

Reinforcement learning (RL) [23]–[27] can be formulated as a Markov Decision Process (MDP) [28], where the agent interacts with an environment based on the reward and state transition. This decision making process is shown in Fig. 2, where s_t and r_t are the state and reward received from environment at step t , and a_t is the action selected by the agent. Based on a_t , the agent interacts with the environment, transitioning to state s_{t+1} and receiving a new reward r_{t+1} . This procedure continues until the end of the MDP.

The aim of RL [29] is to find an optimal policy $\pi(\theta^*)$ that maximizes the expected accumulated reward,

$$\theta^* = \max_{\theta} \mathbb{E}[\sum_{t=0}^T \gamma^t r_t | \pi_{\theta}], \quad (1)$$

where θ represents the policy parameters (e.g., the weights of a neural network); $\sum_{t=0}^T \gamma^t r_t$ is the sum of discounted rewards over an episode; and $\gamma \in (0, 1)$ denotes the discount factor that balances the long- and short-term rewards.

To find the optimal policy parameters θ^* , many different RL algorithms have been proposed. We select several representative ones for demonstration, including DQN [1], PPO [14], A2C [15], ACKTR [16].

- **Deep Q-Networks (DQN)** [1]: Instead of predicting the probability of each action, DQN computes the Q value for each available action. Such Q value represents the approximated accumulated reward from the current frame. Based on such approximation, DQN is trained via minimizing the temporal difference loss (i.e., squared Bellman error). In this paper, the corresponding policy for DQN is achieved by selecting the action with maximum Q value. To keep the DQN-trained policy consistent with the other three algorithms, a soft-max layer is added at the end of the network.
- **Proximal Policy Optimization (PPO)** [14]: PPO is an off-policy method using policy gradient, and it pursues a balance among ease of implementation, sample complexity, and ease of tuning. PPO is proposed based on trust region policy optimization (TRPO) that solves a constrained optimization problem so as to alleviate performance instability. However, PPO handles this issue in a different manner as compared to TRPO; it involves a penalty term that indicates the Kullback–Leibler divergence between the old policy action prediction and the new one. This operation computes an update at each

step, minimizing the cost function while restricting the updating step to be relatively small.

- **Advantage Actor-Critic (A2C)** [30]: A2C is an actor-critic method that learns both a policy and a state value function (i.e., also called critic). In A2C, the advantage A represents the extra reward if the action a_t is taken, and is formulated as $A = Q(s_t, a_t) - V(s_t)$, where $Q(s_t, a_t)$ is the state-action value and $V(s_t)$ is the state value. This formulation reduces variance of the policy gradient and thus increases stability of the gradient estimation.
- **Actor Critic using Kronecker-Factored Trust Region (ACKTR)** [16]: ACKTR is an extension of the natural policy gradient, which optimizes both the actor and the critic using Kronecker-factored approximate curvature (K-FAC) with trust region. To the best of knowledge, ACKTR is the first scalable trust region natural gradient method for actor-critic methods. The efficacy of ACKTR suggests that Kronecker-factored natural gradient approximations in RL is a promising framework.

Although notable performances have been achieved by these algorithms on many challenging tasks (e.g., video games [1] and board games [31], [32]), recent studies have revealed that the policies trained by these algorithms are easily fooled by adversarial perturbations [3]–[5], as introduced next.

B. Adversarial Attack

Recent studies have shown that deep learning is vulnerable against well-designed input perturbations (i.e., adversaries) [6]. These adversaries can easily fool even seemingly high performing deep learning models through small input perturbations that are often imperceptible to the human eye. Such vulnerabilities have been well studied in supervised learning, and also to a lesser extent in RL [3], [4], [20].

In RL, the aim of an adversarial attack is to find the optimal adversary δ_t that minimizes the accumulated reward. Let T represent the length of an episode, and $r(\pi, s_t)$ indicate a function that returns the reward given state s_t and DNN based policy π . Accordingly, the problem setting for generating adversarial attacks in RL can be formulated as,

$$\min_{\delta_t} \sum_{t=1}^T r(\pi, s_t + \delta_t) : \forall t \|\delta_t\| \leq L \quad (2)$$

where the adversary generated at time-step t is represented by δ_t , and its norm (i.e., $\|\delta_t\|$) is bounded by L . The basic assumption of Eq. (2) is that the misguided action selection of policy π will result in a reward degradation. In other words, an action prediction a_t^p obtained from the perturbed state $s_t + \delta_t$ may differ from the originally unperturbed action a_t^o , thus threatening the reward value r_t . This corresponds to the definition of untargeted attacks [6], which is stated as $a_t^o \neq a_t^p$.

To solve the optimization problem formulated in Eq. (2), many white-box approaches (e.g., fast gradient sign method [33]) have been applied in previous works [3], [4]. These white-box approaches are essentially derivative-based, as they generate adversarial examples by back-propagating through the policy network to calculate the gradient of a cost function with respect to the input state, i.e., $\nabla_{s_t} J(\pi, \theta, \Delta a_t, s_t)$. Here, θ represents the weights of the neural network; Δa_t is the

change in action space; J indicates the loss function (e.g., cross-entropy loss).

However, an essential precondition of white-box approaches is complete knowledge of the policy, including the model structure and the model parameters θ . Contrarily, a more restrictive setting is black-box policy access [6], that only allows an attacker to present the input state to the policy and observe the output.

In addition to the above, most prior studies have only analyzed the effects of perturbing every pixel of every frame in RL, which is deemed too intensive to be of much practical relevance. Therefore, this paper aims to provide a far more restrained view towards adversarial attack generation, with the goal of unveiling how little it takes to fool deep RL policies.

IV. THE PROPOSED METHODOLOGY

This section provides details of the three key ingredients of the proposed restrictive attack setting, namely *black-box policy access* (BPA), *fractional-state adversary* (FSA), and *tactically-chanced attack* (TCA).

A. Black-box Policy Access

We adopt the commonly used black-box definition [6], [35] from supervised learning, where the attacker can only query the policy. In other words, the attacker is unable to compute the gradient $\nabla_{s_t} \pi(\cdot|s_t, \delta_t)$, but only has the privilege to query a targeted policy so as to obtain useful information for crafting adversarial examples. We note that such a setting has been rarely investigated in RL adversarial attacks.

We realize the proposed BPA setting by formulating the adversarial attack as a black-box optimization problem. To this end, we first define a measure $\mathcal{D}(\cdot)$ to quantify the discrepancy between the original action distribution $\pi(\cdot|s_t)$ (produced by an RL policy without input perturbation) and the corresponding perturbed distribution $\pi(\cdot|s_t + \delta_t)$. Assuming a finite set of m available actions $a_t^1, a_t^2, \dots, a_t^m$, the probability distribution over actions is represented as $\pi(\cdot|s_t) = [p(a_t^1), p(a_t^2), \dots, p(a_t^m)]$, where $\sum_{j=1}^m p(a_t^j) = 1$. Typically, a deterministic policy selects the action $o = \arg \max_j p(a_t^j)$. With this, the black-box attacker considers the discrepancy measure as the optimization objective function, where the state s_t is perturbed by adversary δ_t such that the measure $\mathcal{D}(\cdot)$ is maximized.

The overall problem formulation can thus be stated as,

$$\max_{\delta_t} \mathcal{D}(\pi(\cdot|s_t), \pi(\cdot|s_t + \delta_t)) : \forall t \|\delta_t\| \leq L. \quad (3)$$

The above mathematical program changes the action distribution to $\pi(\cdot|s_t + \delta_t)$, from the original (optimal) action distribution $\pi(\cdot|s_t)$. *If the change in action distribution leads to a change in action selection at any point in the trajectory of an RL agent, then, based on Bellman's principle of optimality [30], the resulting sub-trajectory shall lead to a sub-optimal cumulative reward as a consequence of the attack.* It is based on this principle that we propose to substitute the reward minimization problem of Eq. (2) with the discrepancy maximization formulation in Eq. (3).

The exact choice of the discrepancy measure $\mathcal{D}(\cdot)$ is expected to have a significant impact on the attack performance, as different measures shall capture different variations in the action distributions. Previous works [3], [4], [20] apply the Euclidean norm (e.g., L_1 , L_2 , L_∞) between $\pi(\cdot|s_t + \delta_t)$ and $\pi(\cdot|s_t)$. However, such L_p norm cannot guarantee a successful untargeted attack, since maximizing the L_p norm may not guarantee that the action selection has been altered by the perturbed state. Thus, a successful untargeted attack under deterministic action selection must ensure that $\arg \max_j \pi(\cdot|s_t + \delta_t)_j \neq \arg \max_j \pi(\cdot|s_t)_j$, where $\pi(\cdot|s_t)_j = p(a_t^j)$.

To enable a more consistent discrepancy measure for untargeted attacks, we design the following function $\tilde{\mathcal{D}}$ based on the query feedback of the policy,

$$\tilde{\mathcal{D}} = \max_{j \neq o} \pi(\cdot|s_t + \delta_t)_j - \pi(\cdot|s_t + \delta_t)_o \quad (4)$$

where: $o = \arg \max_j \pi(\cdot|s_t)_j$.

This formulation is different from the Euclidean norm, guaranteeing a successful untargeted attack if $\tilde{\mathcal{D}}$ is positive. To support this claim, we refer to the theorem and proof below.

Theorem 1. *Suppose the discrepancy measure $\tilde{\mathcal{D}}$ in Eq. (4) is positive and policy π is a deterministic, i.e., action a_t^o is chosen such that $o = \arg \max_j \pi(\cdot|s_t)_j$. Then, the adversarial example δ_t for policy π at state s_t is guaranteed to be a successful untargeted attack, i.e.,*

$$\arg \max_j [\pi(\cdot|s_t + \delta_t)]_j \neq \arg \max_j [\pi(\cdot|s_t)]_j.$$

Proof. We use the symbol o and p to represent the index of the selection action from state s_t and perturbed state $s_t + \delta_t$, respectively. As the policy π is deterministic, o and p are given as,

$$o = \arg \max_j [\pi(\cdot|s_t)]_j,$$

$$p = \arg \max_j [\pi(\cdot|s_t + \delta_t)]_j.$$

If the discrepancy measure $\tilde{\mathcal{D}}$ is positive, then we have,

$$\begin{aligned} \tilde{\mathcal{D}} > 0 &\Rightarrow \max_{j \neq o} \pi(\cdot|s_t + \delta_t)_j - \pi(\cdot|s_t + \delta_t)_o > 0 \\ &\Rightarrow \max_{j \neq o} \pi(\cdot|s_t + \delta_t)_j > \pi(\cdot|s_t + \delta_t)_o \\ &\Rightarrow \max_j \pi(\cdot|s_t + \delta_t)_j = \max_{j \neq o} \pi(\cdot|s_t + \delta_t)_j. \end{aligned}$$

Therefore, the perturbed action index p is given by,

$$\begin{aligned} p &= \arg \max_j \pi(\cdot|s_t + \delta_t)_j = \arg \max_{j \neq o} \pi(\cdot|s_t + \delta_t)_j \neq o \\ &\Rightarrow \arg \max_j [\pi(\cdot|s_t + \delta_t)]_j \neq \arg \max_j [\pi(\cdot|s_t)]_j. \quad \blacksquare \end{aligned}$$

With the aforementioned discrepancy measure $\tilde{\mathcal{D}}$, the mathematical program for generating attacks is given by,

$$\begin{aligned} \max_{\delta_t} \max_{j \neq o} \pi(\cdot|s_t + \delta_t)_j - \pi(\cdot|s_t + \delta_t)_o \\ \text{where: } \forall t \|\delta_t\| \leq L, o = \arg \max_j \pi(\cdot|s_t)_j \end{aligned} \quad (5)$$

When resolving the optimization problem in Eq. (5), δ_t is determined according to its parameterisation. In this paper, δ_t is limited to perturbing only a small fraction of the input state (i.e., fractional-state adversary), as described next.

(a) Full-state adversary [33]

(b) Fractional-state adversary [34]

Fig. 3: Full-state adversary versus the FSA

B. Fractional-State Adversary

To explore adversarial attacks limited to a few pixels in RL, we investigate the fractional-state adversary (FSA) setting. In comparison to the full-state adversary depicted in Fig. 3(a), FSA only perturbs a fraction of the state; shown in Fig. 3(b). The extreme scenario for FSA is merely a single-pixel attack, which is deemed to be more physically realizable for the attacker than a full-state one. For instance, pasting a sticker or a simple fading of color in the “STOP” sign could easily lead to a FSA in Fig. 3(b).

To achieve the FSA setting, we parameterize the adversary δ_t by its corresponding pixel coordinates (i.e., $\mathbf{x}_t, \mathbf{y}_t$) and perturbation value $\mathbf{p}_t$ as follows,

$$\begin{aligned} \delta_t &\leftarrow \mathcal{P}(\mathbf{x}_t, \mathbf{y}_t, \mathbf{p}_t) \\ &\leftarrow \mathcal{P}(x_t^1, y_t^1, p_t^1, \dots, x_t^n, y_t^n, p_t^n) \end{aligned} \quad (6)$$

where n is the number of pixels that are attacked in a frame, x_t^i, y_t^i are coordinate values of the i^{th} perturbed pixel, and p_t^i is the adversarial perturbation value of the i^{th} pixel. With Eq. (6), the final black-box optimization problem is stated as,

$$\max_{\mathbf{x}_t, \mathbf{y}_t, \mathbf{p}_t} \max_{j \neq o} [\pi(\cdot | s_t + \mathcal{P}(\mathbf{x}_t, \mathbf{y}_t, \mathbf{p}_t))]_j - [\pi(\cdot | s_t + \mathcal{P}(\mathbf{x}_t, \mathbf{y}_t, \mathbf{p}_t))]_o$$

where: $\forall t \ o = \arg \max_j [\pi(\cdot | s_t)]_j$;

$$\begin{aligned} \forall t \ 0 \leq x_t^i \leq \mathcal{I}_x, x_t^i \in \mathbb{N}, i &= [1, 2, \dots, n]; \\ \forall t \ 0 \leq y_t^i \leq \mathcal{I}_y, y_t^i \in \mathbb{N}, i &= [1, 2, \dots, n]; \\ \forall t \ 0 \leq u_t^i + p_t^i \leq \mathcal{I}_p, p_t^i \in \mathbb{N}, i &= [1, 2, \dots, n]; \end{aligned} \quad (7)$$

where $\mathcal{I}_x, \mathcal{I}_y$ are the integral upper bounds of $\mathbf{x}_t$ and $\mathbf{y}_t$ respectively, u_t^i is the unperturbed value of the i^{th} pixel in s_t , $\mathcal{I}_p$ is the maximum allowable perturbed pixel value, and n is the FSA size that controls the number of pixels to be attacked given any input state s_t . The extreme case is $n = 1$, implying that there is only *one pixel* attacked in an input frame. To the best of our knowledge, such one-pixel attack has never been explored in the context of RL adversarial attacks.

To obtain the optimal $\mathbf{x}_t, \mathbf{y}_t, \mathbf{p}_t$, we utilize a simple genetic algorithm (GA) [12], [13], [36] (i.e., a *derivative-free* evolutionary computation [37]–[44] method for black-box optimization). The pseudo code of the GA is illustrated in Algorithm 1 (lines 7-17). A population pop containing λ individuals is evolved, where each individual $\delta_t^i, i \in [1, \lambda]$ is a candidate adversarial example that is evaluated by the objective function in Eq. (7). The top- $\lambda\beta$ individuals with respect to the objective value are selected as elites to survive in the new generation (line 10). Based on standard crossover (line 12) and mutation (line 13) operations, another population pop' is reproduced to replace the old pop . This process repeats until a specified maximum number of function evaluations M^{eval} is reached.

Algorithm 1: Attack Generation via GA

Input: s_t – state at time-step t ; n – number of pixels for attack; λ – population size; β – selection rate; γ – rate of mutation; M^{eval} – maximum number of function evaluations; ζ^* – TCA threshold value

```

1 Load the policy  $\pi$  pre-trained by a particular RL algorithm;
2 Obtain the RL input state  $s$  in a particular Atari game;
3 Calculate the attack uncertainty  $\zeta_t$  by Eq. (8);
4 if  $\zeta_t \geq \zeta^*$  then
5   // No adversarial attack
6    $\delta_t = \text{none}$ ;
7 else
8   // Explore the  $n$  pixel attack
9   Initialize the population  $pop$  with randomly generating  $\lambda$ 
10  individuals;
11  repeat
12    Evaluate the discrepancy measure  $\tilde{\mathcal{D}}$  of each individual;
13    Save the top- $\lambda\beta$  individuals in  $pop'$ ;
14    for  $j=1: \frac{\lambda(1-\beta)}{2}$  do
15      Crossover: generate  $\delta_c^j$  with one-point crossover [12];
16      Mutation:  $\delta_m^j = \delta^j + \gamma \cdot \epsilon_j, \epsilon_j \in \mathcal{N}(0, 1)$ ;
17      Append  $\delta_c$  and  $\delta_m$  to  $pop'$ ;
18     $pop = pop'$ ;
19  until the maximum objective evaluations  $M^{eval}$  reached;
20  Return the optimal FSA  $\delta^* = [\mathbf{x}^*, \mathbf{y}^*, \mathbf{p}^*]$ ;

```

C. Tactically-Chanced Attack

To explore adversarial attacks on a restricted number of frames, we design tactically-chanced attacks (TCA) where the attacker strategically strikes *only* salient frames that are likely to be most contributing to the accumulated reward. In contrast, most existing approaches [3] apply adversarial examples on every frame, which is referred to as fully-chanced attack. Our proposed TCA is more restrictive, realizable, and threatening, as can further be highlighted from three different perspectives: (1) due to the communication budget restriction, the attacker may not be able to strike the policy in a fully-chanced fashion; (2) a tactically-chanced attack is less likely to be detected by the defender; and (3) only striking the salient frames improves the attack efficiency, as many frames contribute trivially to the accumulated reward.

To this end, we define the normalized Shannon entropy of the action distribution $\pi(\cdot | s_t) = [p(a_t^1), p(a_t^2), \dots, p(a_t^m)]$ as a measure of attack uncertainty. Specifically, the attack uncertainty (ζ_t) of each frame is given by,

$$\zeta_t = - \sum_{i=0}^m \frac{p(a_t^i) \cdot \log p(a_t^i)}{\log m} \quad (8)$$

where $p(a_i)$ is the probability of action a_i ; m is the dimensionality of the action space. Note that ζ_t must lie within $(0, 1]$. Under this definition, the frame with relatively low ζ_t indicates that the policy has high confidence in its prescribed action. Hence, we assume that attacking such frames will effectively disrupt the policy.

Fig. 4: The six representative Atari games [45] considered in the experiments

Fig. 5: The attack uncertainty ζ_t on Pong.

We demonstrate the attack uncertainty on Pong in Fig. 5; similar trends can be observed on the other games as well. In the figure, the frames with smaller ζ_t , marked by brown rectangles, depict that the ball is close to the paddle. On the other hand, the blue rectangle marked ones with larger ζ_t (close to 1) indicate that the ball is distant from the paddle. Attacking the brown marked frames are intuitively more effective in fooling the policy, as they are likely to lead to a more considerable reward loss. In contrast, an attack may be relatively inconsequential for the blue marked frames. As the ball is distant from the paddle, attacking such frames will have little impact on the cumulative reward.

Therefore, we formulate TCA by defining a TCA threshold (ζ^*), as shown by the red dashed line in Fig. (5); this threshold controls the proportion of attacked frames, such that only those frames with an uncertainty value below ζ^* are perturbed. In the experimental section, we analyze the sensitivity of our approach to the choice of ζ^* by varying its values; this also helps us to explore how little it actually takes to deceive a policy (from the perspective of the number of frames attacked). The adversary δ_t under TCA setting is thus given by,

$$\delta_t = \begin{cases} \mathcal{P}(\mathbf{x}_t^*, \mathbf{y}_t^*, \mathbf{p}_t^*), & \text{if } \zeta_t < \zeta^*, \\ \text{none}, & \text{if } \zeta_t \geq \zeta^*. \end{cases} \quad (9)$$

Eq. (9) implies that if the attack uncertainty (ζ_t) is smaller than ζ^* , an adversary shall be generated by solving the optimization problem in Eq. (7). Otherwise, the attacker will tactically hide without wasting compute resources on trivial frames.

Fig. 6: The neural network structure for playing Atari games

V. EXPERIMENTS AND ANALYSIS

We evaluate the three restrictive settings on six Atari games in OpenAI Gym [45] with various difficulty levels, including Pong, Breakout, SpaceInvaders, Seaquest, Qbert, and BeamRider¹. These games are shown in Fig. 4. For each game, the policies are trained by four start-of-the-art RL algorithms, including DQN [1], PPO2 [14], A2C [15], ACKTR [16]. The network structure is taken from [46], and is kept the same for all the four RL algorithms.

A. Experiment Setup

We utilize the fully-trained policies from the RL baseline zoo [47], where each policy is trained for 4×10^7 time-steps. The other parameter settings (e.g., learning rate, discount factor and max gradient norm) for each training algorithm follow the default settings in the RL baseline zoo [47]. Each policy follows the same pre-processing steps and neural network architecture (i.e., shown as Fig. 6) as in [1]. The input state s_t of the neural network is the concatenation of the last four screen images, where each image is resized to 84×84 . The pixel value of the grey scale image is in the range $[0, 255]$ stepped by 1. The output of the policy is a distribution over possible actions for PPO2, A2C, ACKTR, and an estimation of Q values for DQN.

To calculate the attack uncertainty for DQN, the soft-max operation is applied to normalize the output. Moreover, given the image size and pixel value, the constraints in Eq. (7) are set as $\mathcal{I}_x = 84$, $\mathcal{I}_y = 84$, $\mathcal{I}_p = 255$. In the GA, the population size λ , maximum number of objective function evaluations M^{eval} , selection rate β and mutation rate γ are set to 10, 400, 0.2 and 0.1 respectively. However, following Theorem 1, the optimization process is terminated when a positive function value is found. To investigate the impact of FSA size n , we alter its value from 1 to 10 in steps of 1. In other experiments, we analyze the effect of the TCA threshold ζ^* by setting

¹Our code is available at <https://github.com/RLMA2019/RLMA>

Fig. 7: The results of adversarial attack with different FSA size n (i.e. number of attacked pixels), where the line and shaded area illustrate the mean and standard deviation of 30 independent runs respectively.

Fig. 8: The results of adversarial attack with different TCA threshold ζ^* , where the line and shaded area illustrate the mean and standard deviation of 30 independent runs respectively.

different values in $[0, 1]$. This leads to different proportions of attacked frames.

B. Results for Fractional-State Adversary

We investigate the impact of FSA size n , where $n = 1$ corresponds to the extreme case that only one pixel is perturbed. For a fair comparison with respect to different algorithms and games, in the experiments reported in Fig. 7 we have set the TCA threshold ζ^* as the mean of all frames' attack uncertainty values. Fig. 7 reveals the performance (i.e. accumulated reward) drop of the policies trained by the four RL algorithms on the six Atari games. Several interesting observations can be noted.

(1) Overall, the FSA size n is positively related to the performance drop. That is, the FSA with larger n is able to deceive the policy to a larger extent. (2) On most of the games, the policies are almost completely fooled with $n \leq 4$.

This indicates a small FSA size is often sufficient to achieve a successful adversarial attack, indicating the efficiency of FSA. (3) The results indicate that high performance in the absence of attacks is not a guarantee of robustness against adversarial attacks. For instance, in SpaceInvaders, the original performance of A2C is higher than DQN, but the performance of A2C drops faster than that of DQN as shown in Fig. 7(c). (4) We *surprisingly* find that: on Breakout, with only a single pixel attacked, the policies trained by all the four RL algorithms are deceived, with the accumulated reward dropping by more than 50%; the visualisation for the single-pixel attack on Breakout is shown in Fig. 1. Such single pixel attack corresponds to an approximate perturbation proportion of 0.01% (i.e., $1/(84 \times 84)$) of the total number of pixels in a frame. On Qbert, the accumulated reward of the policy trained by DQN drops from 10000 to 0 with two pixels attacked.

Fig. 9: The relationship between attacked frames and TCA threshold ζ^* , where the line and shaded area illustrate the mean and standard deviation of 30 independent runs respectively.

Fig. 10: The relationship between attacked frames and total frames, where the line and shaded area illustrate the mean and standard deviation of 30 independent runs respectively.

C. Results for Tactically-Chanced Attack

We study the tactically-chanced attack (TCA) by altering the TCA threshold ζ^* . Different ζ^* values lead to different proportions of attacked frames. In our exploration, we set the ζ^* in the range of $(0, 1]$, where $\zeta^* = 0$ and $\zeta^* = 1$ correspond to 0% and 100% proportion of attacked frames, respectively. Recall that according to the results in Fig. 7, most of the policies are sufficiently fooled with no more than 4 pixels attacked. Hence, for a fair comparison, we set the FSA size $n = 4$ in subsequent experiments. The results are displayed in Fig. 8, where we observe that in general the performance (i.e., accumulated reward) decreases for higher TCA threshold ζ^* values. This demonstrates that the policies are more easily fooled with more frames attacked. In addition, we make some other instructive observations.

(1) In Fig. 8 (a) on Pong and Fig. 8 (d) on Seaquest, the policies trained by PPO2 are more sensitive to the increasing of TCA threshold values. In contrast, the performance of policies trained by DQN appears to be more stable with the increasing values of TCA threshold, until it is close to 1. We infer that this observation is because DQN provides Q values instead of directly predicting action probabilities; the resultant increased variance of the derived action distribution leads to much fewer attacked frames under a given threshold value. (2) In Fig. 8 (b) on Breakout, the accumulated rewards of the policies trained by PPO2, A2C and ACKTR decrease to 0 when the TCA threshold value is higher than 0.3. This indicate that a small TCA threshold value is enough to completely deceive those policies. (3) In Fig. 8 (f) on Qbert, when the TCA threshold value is higher than 0.6, all policies can only

Fig. 11: The impact of attacked frames for accumulated reward on Breakout

obtain a reward less than 2000. Especially for the policy trained by PPO2, the accumulated reward drops from 14000 to 2000.

We further analyze the exact number of attacked frames with different TCA threshold values. As shown in Fig. 9, the number of attacked frames naturally increases with the prescribed threshold value. In Fig. 9 (b), on Breakout, when the TCA threshold value is smaller than 0.2, less than 50 frames are attacked. A similar finding is also observed on Beamrider, where less than 100 frames are attacked when the TCA threshold value is smaller than 0.2.

We also analyze the relationship between the number of attacked frames and the number of total frames in Fig. 10. We find similar trends on the six games, where the total number of frames in the game decreases when more frames are attacked. This observation is particularly apparent for Breakout in Fig. 10 (b), where attacking less than 50 frames considerably reduces the total number of frames. This results from the fact that the attacked frames may cause (i) termination of the game, or (ii) a loss of life. Both outcomes greatly shorten the episode, resulting in a smaller number of total frames. For instance, if the paddle in Breakout is deceived to miss the ball, the agent would lose 1/5 life.

To explore a more restrictive setting, we additionally examine the single-pixel ($n = 1$) attack on limited frames. We illustrate the results on Breakout in Fig. 11, and similar trends can be observed on other games as well. From this figure, we find that by only attacking six frames on average, the policy trained by DQN is totally fooled with reward decreasing from 224 to 13.7. Here, six frames correspond to an attack ratio of $6/535 \approx 1\%$, where 535 is the averaged number of total frames. Similar findings are also observed on policies trained by PPO2, A2C and ACKTR. For instance, the policy trained by PPO2 shows a significant reward decline with only 27 frames attacked under the same TCA setting as other policies.

D. Attack Initialization Analysis

The GA population initialization plays an important role in efficiently generating successful adversarial examples, and thus highlights the practicality of the approach. To this end, we study the effects of random initialization versus warm starting based on prior data. In RL, the simplest form of prior data originates from the correlation between states in a Markov Decision Process (MDP); this also distinguishes RL from other supervised learning tasks (e.g., classification), where the predictions on different data points are independent. Thus, with the existence of correlations between adjacent frames in mind, we set up different types of population initializations in the

GA: (1) **GA-RI**: random initialization (RI) of pop_0 ; (2) **GA-WSI**: warm start initialization (WSI) of pop_0 by inserting the optimal adversary x_p^* found in the previous attacked frame.

We randomly attack four frames for all the six games. On each attacked frame, the GA optimization is executed 10 times independently on both types of initialization. The comparison results are shown in Fig. 12. In general, we clearly see that GA-WSI significantly speeds up the optimization convergence in finding successful adversarial examples (i.e., corresponding to positive objective function values). In particular, on many frames (e.g., Pong: ξ_1, ξ_2, ξ_4 ; Breakout: ξ_2 ; Space Invaders: ξ_3 ; Qbert: ξ_3, ξ_4), the GA-RI takes significantly more time (or is even unable) to obtain a positive objective value; whereas GA-WSI achieves it easily.

We interestingly find that, in some cases, GA-WSI provides a successful adversary without the need of any optimization. This observation underlines the real possibility of launching successful adversarial attacks on deep RL policies in **real-time** at little/no computational cost, a threat that raises severe concerns about their deployment especially in safety critical applications. As such, it points to two critical directions for future research. On the one hand, there is a need for more comprehensive exploration of procedures of learning from experiential frames, that could lead to the generation of potent adversaries on-the-fly. On the other hand, a significant void in the trustworthiness of existing RL policies is revealed, one that could possibly be addressed through more robust training algorithms.

CONCLUSION

This paper explores *minimalistic scenarios* for adversarial attack in deep reinforcement learning (RL), comprising (1) black-box policy access where the attacker only has access to the input (state) and output (action probability) of an RL policy; (2) fractional-state adversary where only a small number of pixels are perturbed, with the extreme case of one-pixel attack; and (3) tactically-chanced attack where only some significant frames are selected to be attacked. We verify these settings on policies that are trained by four state-of-the-art RL algorithms on six Atari games. We surprisingly find that: (i) with only **a single pixel** ($\approx 0.01\%$ of the state) attacked, the trained policies can be significantly deceived; and (ii) by merely attacking around 1% frames, the policy trained by DQN is totally fooled on certain games. Our immediate future work will focus on further investigating different initialization schemes that could better exploit the correlations across different attacked frames. Our goal is to thoroughly analyze the plausibility of generating adversarial attacks in real-time for time-sensitive RL applications.

Fig. 12: The comparison between random initialization (GA-RI) and warm start initialization (GA-WSI), where the line and shaded area illustrate the mean and standard deviation of 10 independent runs respectively.

ACKNOWLEDGEMENT

This work is funded by the National Research Foundation, Singapore under its AI Singapore programme [Award No.: AISG-RP-2018-004] and the Data Science and Artificial Intelligence Research Center (DSAIR) at Nanyang Technological University. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not reflect the views of National Research Foundation, Singapore.

REFERENCES

- [1] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, "Playing atari with deep reinforcement learning," *arXiv preprint arXiv:1312.5602*, 2013.
- [2] Z. Li, J. Liu, Z. Huang, Y. Peng, H. Pu, and L. Ding, "Adaptive impedance control of human-robot cooperation using reinforcement learning," *IEEE Transactions on Industrial Electronics*, vol. 64, no. 10, pp. 8013–8022, 2017.
- [3] S. Huang, N. Papernot, I. Goodfellow, Y. Duan, and P. Abbeel, "Adversarial attacks on neural network policies," *arXiv preprint arXiv:1702.02284*, 2017.
- [4] Y.-C. Lin, Z.-W. Hong, Y.-H. Liao, M.-L. Shih, M.-Y. Liu, and M. Sun, "Tactics of adversarial attack on deep reinforcement learning agents," *arXiv preprint arXiv:1703.06748*, 2017.
- [5] X. Yuan, P. He, Q. Zhu, and X. Li, "Adversarial examples: Attacks and defenses for deep learning," *IEEE transactions on neural networks and learning systems*, 2019.
- [6] P.-Y. Chen, H. Zhang, Y. Sharma, J. Yi, and C.-J. Hsieh, "Zoo: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models," in *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*. ACM, 2017, pp. 15–26.
- [7] A. Pattanaik, Z. Tang, S. Liu, G. Bommannan, and G. Chowdhary, "Robust deep reinforcement learning with adversarial attacks," in *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2018, pp. 2040–2042.
- [8] L. Pinto, J. Davidson, R. Sukthankar, and A. Gupta, "Robust adversarial reinforcement learning," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2017, pp. 2817–2826.
- [9] A. Raghu, "Reinforcement learning for sepsis treatment: Baselines and analysis," *Reinforcement Learning for Real Life*, 2019.
- [10] A. Raghu, M. Komorowski, L. A. Celi, P. Szolovits, and M. Ghassemi, "Continuous state-space models for optimal sepsis treatment-a deep reinforcement learning approach," *arXiv preprint arXiv:1705.08422*, 2017.
- [11] S. Parbhoo, J. Bogojeska, M. Zazzi, V. Roth, and F. Doshi-Velez, "Combining kernel and model based learning for hiv therapy selection," *AMIA Summits on Translational Science Proceedings*, vol. 2017, p. 239, 2017.
- [12] J. H. Holland, "Genetic algorithms," *Scientific american*, vol. 267, no. 1, pp. 66–73, 1992.
- [13] K. Deep, K. P. Singh, M. L. Kansal, and C. Mohan, "A real coded genetic algorithm for solving integer and mixed integer optimization problems," *Applied Mathematics and Computation*, vol. 212, no. 2, pp. 505–518, 2009.
- [14] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [15] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu, "Asynchronous methods for deep reinforcement learning," in *International conference on machine learning*, 2016, pp. 1928–1937.
- [16] Y. Wu, E. Mansimov, R. B. Grosse, S. Liao, and J. Ba, "Scalable trust-region method for deep reinforcement learning using kronecker-factored approximation," in *Advances in neural information processing systems*, 2017, pp. 5279–5288.
- [17] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," *arXiv preprint arXiv:1312.6199*, 2013.
- [18] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," *arXiv preprint arXiv:1412.6572*, 2014.
- [19] A. Kurakin, I. Goodfellow, and S. Bengio, "Adversarial machine learning at scale," *arXiv preprint arXiv:1611.01236*, 2016.
- [20] C. Xiao, X. Pan, W. He, J. Peng, M. Sun, J. Yi, B. Li, and D. Song, "Characterizing attacks on deep reinforcement learning," *arXiv preprint arXiv:1907.09470*, 2019.
- [21] J. Kos and D. Song, "Delving into adversarial attacks on deep policies," *arXiv preprint arXiv:1705.06452*, 2017.
- [22] A. Fawzi, S.-M. Moosavi-Dezfooli, and P. Frossard, "The robustness of deep networks: A geometrical perspective," *IEEE Signal Processing Magazine*, vol. 34, no. 6, pp. 50–62, 2017.
- [23] X. Qu, Y.-S. Ong, Y. Hou, and X. Shen, "Memetic evolution strategy for reinforcement learning," in *2019 IEEE Congress on Evolutionary Computation (CEC)*. IEEE, 2019, pp. 1922–1928.
- [24] Z. Xie and Y. Jin, "An extended reinforcement learning framework to model cognitive development with enactive pattern representation," *IEEE Transactions on Cognitive and Developmental Systems*, vol. 10, no. 3, pp. 738–750, 2018.
- [25] F. de La Bourdonnaye, C. Teulière, J. Triesch, and T. Chateau, "Learning to touch objects through stage-wise deep reinforcement learning," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 1–9.
- [26] Y. Hou, Y.-S. Ong, L. Feng, and J. M. Zurada, "An evolutionary transfer reinforcement learning framework for multiagent systems," *IEEE Transactions on Evolutionary Computation*, vol. 21, no. 4, pp. 601–615, 2017.
- [27] Z. Zhang, Y.-S. Ong, D. Wang, and B. Xue, "A collaborative multiagent reinforcement learning method based on policy gradient potential," *IEEE transactions on cybernetics*, 2019.
- [28] C. H. Papadimitriou and J. N. Tsitsiklis, "The complexity of markov decision processes," *Mathematics of operations research*, vol. 12, no. 3, pp. 441–450, 1987.
- [29] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, "Trust region policy optimization," in *International Conference on Machine Learning*, 2015, pp. 1889–1897.
- [30] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [31] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot *et al.*, "Mastering the game of go with deep neural networks and tree search," *nature*, vol. 529, no. 7587, p. 484, 2016.
- [32] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton *et al.*, "Mastering the game of go without human knowledge," *Nature*, vol. 550, no. 7676, p. 354, 2017.
- [33] I. Goodfellow, N. Papernot, S. Huang, Y. Duan, P. Abbeel, and J. Clark, "Attacking machine learning with adversarial examples," *OpenAI*. <https://blog.openai.com/adversarial-example-research>, 2017.
- [34] T. B. Brown, D. Mané, A. Roy, M. Abadi, and J. Gilmer, "Adversarial patch," *arXiv preprint arXiv:1712.09665*, 2017.
- [35] J. Su, D. V. Vargas, and K. Sakurai, "One pixel attack for fooling deep neural networks," *IEEE Transactions on Evolutionary Computation*, 2019.
- [36] D. P. Bertsekas, "Nonlinear programming," *Journal of the Operational Research Society*, vol. 48, no. 3, pp. 334–334, 1997.
- [37] Y. Jin, J. Branke *et al.*, "Evolutionary optimization in uncertain environments-a survey," *IEEE Transactions on evolutionary computation*, vol. 9, no. 3, pp. 303–317, 2005.
- [38] Y. Jin, "A comprehensive survey of fitness approximation in evolutionary computation," *Soft computing*, vol. 9, no. 1, pp. 3–12, 2005.
- [39] Y. Jin, M. Olhofer, and B. Sendhoff, "A framework for evolutionary optimization with approximate fitness functions," *IEEE Transactions on evolutionary computation*, vol. 6, no. 5, pp. 481–494, 2002.
- [40] Y. Jin, H. Wang, T. Chugh, D. Guo, and K. Miettinen, "Data-driven evolutionary optimization: an overview and case studies," *IEEE Transactions on Evolutionary Computation*, vol. 23, no. 3, pp. 442–458, 2018.
- [41] B. Liu, L. Wang, Y.-H. Jin, F. Tang, and D.-X. Huang, "Improved particle swarm optimization combined with chaos," *Chaos, Solitons & Fractals*, vol. 25, no. 5, pp. 1261–1271, 2005.
- [42] B. Liu, L. Wang, and Y.-H. Jin, "An effective pso-based memetic algorithm for flow shop scheduling," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 37, no. 1, pp. 18–27, 2007.
- [43] X. Ma, Q. Zhang, G. Tian, J. Yang, and Z. Zhu, "On tchebycheff decomposition approaches for multiobjective evolutionary optimization,"

IEEE Transactions on Evolutionary Computation, vol. 22, no. 2, pp. 226–244, 2017.

- [44] L. Feng, Y.-S. Ong, S. Jiang, and A. Gupta, “Autoencoding evolutionary search with learning across heterogeneous problems,” *IEEE Transactions on Evolutionary Computation*, vol. 21, no. 5, pp. 760–772, 2017.
- [45] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, “Openai gym,” *arXiv preprint arXiv:1606.01540*, 2016.
- [46] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski *et al.*, “Human-level control through deep reinforcement learning,” *Nature*, vol. 518, no. 7540, p. 529, 2015.
- [47] A. RAFFIN, “Reinforcement learning baseline zoo,” <https://github.com/araffin/rl-baselines-zoo>, 2018.

Xinghua Qu received the B.S. degree in Aircraft Design and Engineering from Northwestern Polytechnic University, China, in 2014, and the M.S. degree in Aerospace Engineering, School of Astronautics, Beihang University, China, in 2017. He is currently the Phd student from school of computer science and engineering, Nanyang Technological University, Singapore. His current research focus includes reinforcement learning and adversarial machine learning.

Zhu Sun received her Ph.D. degree from School of Computer Science and Engineering, Nanyang Technological University (NTU), Singapore, in 2018. During her Ph.D. study, she focused on designing efficient recommendation algorithms by considering side information. Her research has been published in leading conferences and journals in related domains (e.g., IJCAI, AAAI, CIKM, ACM RecSys and ACM UMAP). Currently, she is a research fellow at School of Electrical and Electronic Engineering, NTU, Singapore.

Yew-Soon Ong (M'99-SM'12-F'18) received a Ph.D. degree on artificial intelligence in complex design from the University of Southampton, U.K., in 2003. He is a President's Chair Professor in Computer Science at the Nanyang Technological University (NTU), and holds the position of Chief Artificial Intelligence Scientist at the Agency for Science, Technology and Research Singapore. At NTU, he currently serves as Director of the Data Science and Artificial Intelligence Research Center and Director of the Singtel-NTU Cognitive & Artificial Intelligence Joint Lab. His current research focus is in artificial and computational intelligence. He is founding Editor-in-Chief of the *IEEE Transactions on Emerging Topics in Computational Intelligence* and Associate Editor of several *IEEE Transactions*. He has received several IEEE outstanding paper awards, listed as a Thomson Reuters highly cited researcher and among the World's Most Influential Scientific Minds.

Abhishek Gupta received his PhD in Engineering Science from the University of Auckland, New Zealand, in the year 2014. He currently serves as a Scientist in the Singapore Institute of Manufacturing Technology (SIMTech), at the Agency for Science, Technology and Research (A*STAR), Singapore. He has diverse research experience in the field of computational science, ranging from numerical methods in engineering physics, to topics in computational intelligence. His current research focus is in the development of memetic computation and probabilistic model-based algorithms for automated knowledge extraction and transfer across optimization problems, with diverse applications spanning cyber physical production systems and engineering design.

Pengfei Wei received his Ph.D. Degree from School of Computer Science and Engineering, Nanyang Technological University (NTU), Singapore, in 2018. During his Ph.D. study, he focused on domain adaptation and gaussian processes. His research has been published in leading conferences and journals in related domains (e.g., ICML, IJCAI, ICDM, IEEE TNNLS, IEEE TKDE). Currently, he is a research fellow at School of Computing, National University of Singapore.