

---

# Robust Reinforcement Learning via Adversarial training with Langevin Dynamics

---

Parameswaran Kamalaruban<sup>1</sup> Yu-Ting Huang<sup>1</sup> Ya-Ping Hsieh<sup>1</sup> Paul Rolland<sup>1</sup> Cheng Shi<sup>1</sup> Volkan Cevher<sup>1</sup>

## Abstract

We introduce a *sampling* perspective to tackle the challenging task of training robust Reinforcement Learning (RL) agents. Leveraging the powerful Stochastic Gradient Langevin Dynamics, we present a novel, scalable two-player RL algorithm, which is a sampling variant of the two-player policy gradient method. Our algorithm consistently outperforms existing baselines, in terms of generalization across different training and testing conditions, on several MuJoCo environments. Our experiments also show that, even for objective functions that entirely ignore potential environmental shifts, our sampling approach remains highly robust in comparison to standard RL algorithms.

## 1. Introduction

Reinforcement learning (RL) promise automated solutions to many real-world tasks with beyond-human performance. Indeed, recent advances in policy gradient methods (Sutton et al., 2000; Silver et al., 2014; Schulman et al., 2015; 2017) and deep reinforcement learning have demonstrated impressive performance in games (Mnih et al., 2015; Silver et al., 2017), continuous control (Lillicrap et al., 2015), and robotics (Levine et al., 2016).

Despite the success of deep RL, the progress is still upset by the fragility in real-life deployments. In particular, the majority of these methods fail to perform well when there is some difference between training and testing scenarios, thereby posing serious safety and security concerns. To this end, learning policies that are *robust* to environmental shifts, mismatched configurations, and even mismatched control actions are becoming increasingly more important.

A powerful framework to learning robust policies is to interpret the changing of the environment as an adversarial perturbation. This notion naturally lends itself to a two-player max-min problem involving a pair of agents, a protagonist

and an adversary, where the protagonist learns to fulfill the original task goals while being robust to the disruptions generated by its adversary. Two prominent examples along this research vein, differing in how they model the adversary, are the Robust Adversarial Reinforcement Learning (RARL) (Pinto et al., 2017) and Noisy Robust Markov Decision Process (NR-MDP) (Tessler et al., 2019).

Despite the impressive empirical progress, the training of the robust RL objectives remains an open and critical challenge. In particular, (Tessler et al., 2019) prove that it is in fact strictly suboptimal to directly apply (deterministic) policy gradient steps to their NR-MDP max-min objectives. Owing to the lack of a better algorithm, the policy gradient is nonetheless still employed in their experiments; similar comments also apply to (Pinto et al., 2017).

The main difficulty originates from the highly non-convex-concave nature of the robust RL objectives, posing significant burdens to all optimization methods. In game-theoretical terms, these methods search for pure Nash Equilibria (pure NE) which might not even exist (Dasgupta & Maskin, 1986). Worse, even when pure NE are well-defined, we show that optimization methods can still get stuck at non-equilibrium stationary points on certain extremely simple non-convex-concave objectives; *cf.* Section 4.

In this paper, we contend that, instead of viewing robust RL as a max-min optimization problem, the *sampling* perspective (Hsieh et al., 2019) from the so-called *mixed Nash Equilibrium* (mixed NE) presents a potential solution to the grand challenge of training robust RL agents. We substantiate our claim by demonstrating the advantages of sampling algorithms over-optimization methods on three fronts:

1. We show in Section 4 that, even in stylized examples that trap the common optimization methods, the sampling algorithms can still make progress towards the optimum in expectation, even tracking the NE points.
2. We conduct extensive experiments to show that sampling algorithms consistently outperform state-of-the-arts in training robust RL agents. Moreover, our experiments on the MuJoCo dataset reveal that sampling algorithms are able to handle previous failure cases of optimization methods, such as the inverted pendulum.

---

<sup>1</sup>LIONS, EPFL, Switzerland. Correspondence to: P. Kamalaruban <kamalaruban.parameswaran@epfl.ch>.

3. Finally, we provide strong empirical evidence that sampling algorithms are inherently more robust than optimization methods for RL. Specifically, we apply sampling algorithms to train an RL agent with *non-robust* objective (*i.e.*, the standard expected cumulative reward maximizing objective in RL), and we compare against the policy learned by *optimizing the robust objective* (*i.e.*, the max-min formulation). Despite the disadvantage, our results show that the sampling algorithms still achieve comparable or better performance than optimization methods (see Figures 14 and 15).

## 2. Background

### 2.1. Stochastic Gradient Langevin Dynamics (SGLD)

For any probability distribution  $p(z) \propto \exp(-g(z))$ , the Stochastic Gradient Langevin Dynamics (SGLD) (Welling & Teh, 2011) iterates as

$$z_{k+1} \leftarrow z_k - \eta \left[ \widehat{\nabla_z g(z)} \right]_{z=z_k} + \sqrt{2\eta\epsilon} \xi_k, \quad (1)$$

where  $\eta$  is the step-size,  $\widehat{\nabla_z g(z)}$  is an unbiased estimator of  $\nabla_z g(z)$ ,  $\epsilon > 0$  is a temperature parameter, and  $\xi_k \sim \mathcal{N}(0, I)$  is a standard normal vector, independently drawn across different iterations. In some cases, the convergence rate of SGLD can be improved by scaling the noise using a positive-definite symmetric matrix  $C$ . We thus define a preconditioned variant of the above update (1) as follows:

$$z_{k+1} \leftarrow z_k - \eta C^{-1} \left[ \widehat{\nabla_z g(z)} \right]_{z=z_k} + \sqrt{2\eta\epsilon} C^{-\frac{1}{2}} \xi_k. \quad (2)$$

In the experiments, we use a RMSProp-preconditioned version of the SGLD (Li et al., 2016).

### 2.2. Saddle Point Problems and Pure NE

Consider the following Saddle Point Problem (SPP):

$$\max_{\theta \in \mathbb{R}^n} \min_{\omega \in \mathbb{R}^m} f(\theta, \omega). \quad (3)$$

Solving (3) equals finding a point  $(\theta^*, \omega^*)$  such that

$$f(\theta, \omega^*) \leq f(\theta^*, \omega^*) \leq f(\theta^*, \omega), \quad \forall \theta \in \mathbb{R}^n, \omega \in \mathbb{R}^m. \quad (4)$$

In the language of game theory, we say that  $(\theta^*, \omega^*)$  is a *pure* Nash Equilibrium (pure NE). If (4) holds only locally, we say that  $(\theta^*, \omega^*)$  is a local pure NE.

A major source of SPPs is the Generative Adversarial Networks (GANs) in deep learning (Goodfellow et al., 2014), which give rise to a variety of algorithms. However, virtually all search for a (local) pure NE; see Section 4.1.

---

### Algorithm 1 MixedNE-LD

---

**Input:** step-size  $\{\eta_t\}_{t=1}^T$ , thermal noise  $\{\epsilon_t\}_{t=1}^T$ , warmup steps  $\{K_t\}_{t=1}^T$ , exponential damping factor  $\beta$ . Initialize (randomly)  $\omega_1, \theta_1$   
**for**  $t = 1, 2, \dots, T - 1$  **do**  
 $\bar{\omega}_t, \omega_t^{(1)} \leftarrow \omega_t$ ;  $\bar{\theta}_t, \theta_t^{(1)} \leftarrow \theta_t$   
**for**  $k = 1, 2, \dots, K_t$  **do**  
 $\xi, \xi' \sim \mathcal{N}(0, I)$   
 $\theta_t^{(k+1)} \leftarrow \theta_t^{(k)} + \eta_t \nabla_{\theta} h(\widehat{\theta}_t^{(k)}, \omega_t) + \sqrt{2\eta_t\epsilon_t} \xi$   
 $\omega_t^{(k+1)} \leftarrow \omega_t^{(k)} - \eta_t \nabla_{\omega} h(\widehat{\theta}_t^{(k)}, \omega_t^{(k)}) + \sqrt{2\eta_t\epsilon_t} \xi'$   
 $\bar{\omega}_t \leftarrow (1 - \beta) \bar{\omega}_t + \beta \omega_t^{(k+1)}$   
 $\bar{\theta}_t \leftarrow (1 - \beta) \bar{\theta}_t + \beta \theta_t^{(k+1)}$   
**end for**  
 $\omega_{t+1} \leftarrow (1 - \beta) \omega_t + \beta \bar{\omega}_t$   
 $\theta_{t+1} \leftarrow (1 - \beta) \theta_t + \beta \bar{\theta}_t$   
**end for**  
**Output:**  $\omega_T, \theta_T$ .

---

### 2.3. Sampling for Mixed NE

In this section, we review some of the key results from (Hsieh et al., 2019). We denote the set of all probability measures on  $\mathcal{Z}$  by  $\mathcal{P}(\mathcal{Z})$ , and the set of all functions on  $\mathcal{Z}$  by  $\mathcal{F}(\mathcal{Z})$ . Given a (sufficiently regular) function  $h : \Theta \times \Omega \rightarrow \mathbb{R}$ , consider the following objective (a two-player game with mixed strategies):

$$\max_{p \in \mathcal{P}(\Theta)} \min_{q \in \mathcal{P}(\Omega)} f(p, q) := \mathbb{E}_{\theta \sim p} \left[ \mathbb{E}_{\omega \sim q} [h(\theta, \omega)] \right]. \quad (5)$$

A pair  $(p^*, q^*)$  achieving the max-min value in (5) is called a *mixed Nash Equilibrium* (mixed NE).

Conceptually, problem (5) can be solved via several infinite-dimensional algorithms, such as the so-called entropic mirror descent or mirror-prox; see (Hsieh et al., 2019). However, these algorithms are infinite-dimensional and require infinite computational power to implement. For practical interest, by leveraging the SGLD sampling techniques and using some practical relaxations, Hsieh et al. (2019) features a simplified variant of these infinite-dimensional algorithms.

For the robust RL formulation (5), it suffices to use the simplest algorithm in (Hsieh et al., 2019). The pseudocode for their resulting algorithm, termed MixedNE-LD (**mixed NE** via Langevin dynamics), can be found in **Algorithm 1**.

## 3. Two-Player Markov Games

**Markov Decision Process:** We consider a Markov Decision Process (MDP) represented by  $\mathcal{M}_1 := (\mathcal{S}, \mathcal{A}, T_1, \gamma, P_0, R_1)$ , where the state and action spaces are

denoted by  $\mathcal{S}$  and  $\mathcal{A}$  respectively. We focus on continuous control tasks, where the actions are real-valued, *i.e.*,  $\mathcal{A} = \mathbb{R}^d$ .  $T_1 : \mathcal{S} \times \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$  captures the state transition dynamics, *i.e.*,  $T_1(s' | s, a)$  denotes the probability of landing in state  $s'$  by taking action  $a$  from state  $s$ . Here  $\gamma$  is the discounting factor,  $P_0 : \mathcal{S} \rightarrow [0, 1]$  is the initial distribution over  $\mathcal{S}$ , and  $R_1 : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$  is the reward.

**Two-Player Zero-Sum Markov Games:** Consider a two-player zero-sum Markov game (Littman, 1994; Perolat et al., 2015), where at each step of the game, both players simultaneously choose an action. The reward each player gets after one step depends on the state and the joint action of both players. Furthermore, the transition kernel of the game is controlled jointly by both players. In this work, we only consider simultaneous games, not the turn-based games.

This game can be described by an MDP  $\mathcal{M}_2 = (\mathcal{S}, \mathcal{A}, \mathcal{A}', T_2, \gamma, R_2, P_0)$ , where  $\mathcal{A}$  and  $\mathcal{A}'$  are the continuous set of actions the players can take,  $T_2 : \mathcal{S} \times \mathcal{A} \times \mathcal{A}' \times \mathcal{S} \rightarrow \mathbb{R}$  is the state transition probability, and  $R_2 : \mathcal{S} \times \mathcal{A} \times \mathcal{A}' \rightarrow \mathbb{R}$  is the reward for both players. Consider an agent executing a policy  $\mu : \mathcal{S} \rightarrow \mathcal{A}$ , and an adversary executing a policy  $\nu : \mathcal{S} \rightarrow \mathcal{A}'$  in the environment  $\mathcal{M}$ . At each time step  $t$ , both players observe the state  $s_t$  and take actions  $a_t = \mu(s_t)$  and  $a'_t = \nu(s_t)$ . In the zero-sum game, the agent gets a reward  $r_t = R_2(s_t, a_t, a'_t)$  while the adversary gets a negative reward  $-r_t$ .

This two-player zero-sum Markov game formulation has been used to model the following robust RL settings:

- Robust Adversarial Reinforcement Learning (RARL) (Pinto et al., 2017), where the power of the adversary is limited by its action space  $\mathcal{A}'$ .
- Noisy Robust Markov Decision Process (NR-MDP) (Tessler et al., 2019), where  $\mathcal{A}' = \mathcal{A}$ ,  $T_2(s_{t+1} | s_t, a_t, a'_t) = T_1(s_{t+1} | s_t, \bar{a}_t)$ , and  $R_2(s_t, a_t, a'_t) = R_1(s_t, \bar{a}_t)$ , with  $\bar{a}_t = (1 - \delta)a_t + \delta a'_t$ , for a chosen  $\delta \in (0, 1)$ , which limits the adversary.

In our adversarial game, we consider the following performance objective:

$$J(\mu, \nu) = \mathbb{E} \left[ \sum_{t=1}^{\infty} \gamma^{t-1} r_t \mid \mu, \nu, \mathcal{M}_2 \right],$$

where  $\sum_{t=1}^{\infty} \gamma^{t-1} r_t$  be the random cumulative return. In particular, we consider the parameterized policies  $\{\mu_\theta : \theta \in \Theta\}$ , and  $\{\nu_\omega : \omega \in \Omega\}$ . By an abuse of notation, we denote  $J(\theta, \omega) = J(\mu_\theta, \nu_\omega)$ . We consider the following objective:

$$\max_{\theta \in \Theta} \min_{\omega \in \Omega} J(\theta, \omega). \quad (6)$$

Note that  $J$  is non-convex-concave in both  $\theta$  and  $\omega$ . Instead of solving (6) directly, we focus on the mixed strategy formulation of (6). In other words, we consider the set of all probability distributions over  $\Theta$  and  $\Omega$ , and we search for the optimal distribution that solves the following program:

$$\max_{p \in \mathcal{P}(\Theta)} \min_{q \in \mathcal{P}(\Omega)} f(p, q) := \mathbb{E}_{\theta \sim p} \left[ \mathbb{E}_{\omega \sim q} [J(\theta, \omega)] \right]. \quad (7)$$

Then, we can use the techniques from Section 2.3 to solve the above problem.

## 4. Simple Non-Convex-Concave SPPs

In Section 3, we have formulated the robust RL problems in either its pure strategy form (6) and mixed strategy form (7).

The goal of the present section is to demonstrate that solving (7) as a *sampling* problem has superior performance over methods that seek pure NE for non-convex-concave SPPs. We do so by providing theoretical and empirical justifications on several simple, yet nontrivial, low-dimensional examples.

Pseudocodes for all algorithms in the section and the omitted proofs can be found in Appendix A.

### 4.1. Existing Algorithms

We will consider three algorithmic frameworks:

1. GAD: Finding pure NE via Gradient ascent-descent.
2. EG: Finding pure NE via Extra-gradient methods.
3. MixedNE-LD: Finding mixed NE via **Algorithm 1**.

Most existing methods to solving SPPs in deep learning can be classified as (adaptive) variants of these frameworks. For instance, Adam, being an adaptive version of GAD, is the predominant algorithm when it comes to learning GANs (Lucic et al., 2018), which was also employed by (Tessler et al., 2019) to train robust RL agents. EG was originally developed by Korpelevich in 1976 to solve variational inequalities for convex problems, and was recently shown to outperform (adaptive) GAD when it comes to training GANs (Gidel et al., 2018). Finally, the MixedNE-LD framework was recently put forth by Hsieh et al. (2019), whose defining feature is to *sample* from the mixed NE.

We now turn to the objectives.

### 4.2. Degree-2 Polynomials: Stationary Points *v.s.* NE

Consider the lowest possible dimension for an SPP:  $\theta, \omega \in \mathbb{R}$ . We aim to construct nontrivial examples where there exist stationary points that are *not* NE. To this end, we

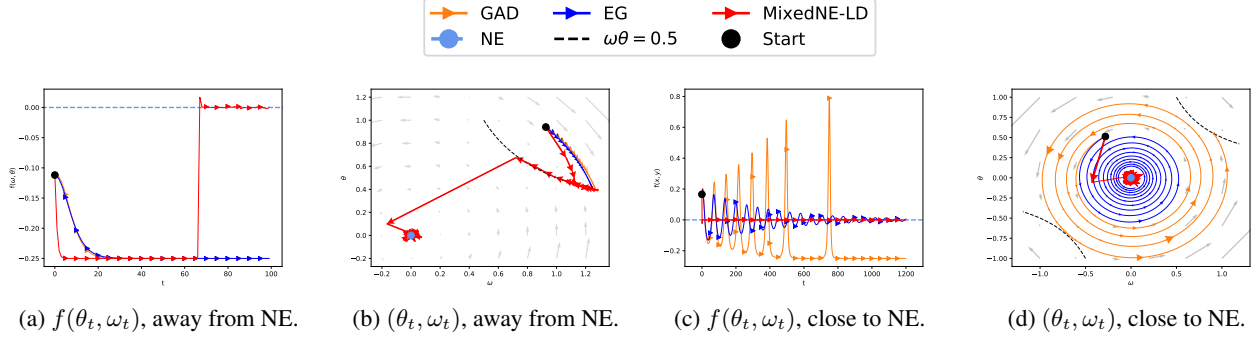


Figure 1.  $f(\theta, \omega) = \theta^2\omega^2 - \theta\omega$ . The NE is  $(0, 0)$  with reward value 0. The dashed curve  $\theta\omega = 0.5$  describe all stationary points that are *not* NE. (a), (b) shows the objective value and the training dynamics when initializing far away from NE. (c), (d) shows the objective value and the training dynamics when  $(\theta_1, \omega_1)$  is initializing close to NE.




Figure 2.  $f(\theta, \omega) = \theta\omega - \theta^2\omega^2$ . The NE is  $(0, 0)$  with reward value 0. The dashed curve  $\theta\omega = 0.5$  are stationary points that are *not* NE. (a), (b) shows the objective value and the training dynamics when initializing far away from NE. (c), (d) shows the objective value and the training dynamics when initializing close to NE.




Figure 3.  $f(\theta, \omega) = \theta^2\omega^2$ . The NE are represented with the line  $\{(\theta, 0) \mid \theta \text{ arbitrary}\}$  with reward value 0. (a), (b) shows the objective value and the training dynamics when initializing far away from NE. (c), (d) shows the objective value and the training dynamics when initializing close to NE.

may simply use the degree-2 polynomials:

$$\max_{\theta \in [-2, 2]} \min_{\omega \in [-2, 2]} f(\theta, \omega) = \theta^2\omega^2 - \theta\omega \quad (8)$$

and

$$\max_{\theta \in [-2, 2]} \min_{\omega \in [-2, 2]} f(\theta, \omega) = \theta\omega - \theta^2\omega^2. \quad (9)$$

The constraint interval  $[-2, 2]$  is included only for ease of presentation; it has no impact on our conclusion. Moreover, the following facts can be readily verified:

- The pure and mixed NE are the same:  $(0, 0)$ .
- The curve  $\{(\theta, \omega) \mid \theta\omega = 0.5\}$  presents stationary points that are *not* NE.

### 4.3. Main Result

We now present the main result in this section.

**Theorem 1.** Consider the (continuous-time) GAD and EG dynamics:

$$\begin{bmatrix} \frac{d\theta}{dt}(t) \\ \frac{d\omega}{dt}(t) \end{bmatrix} = \begin{bmatrix} \nabla_{\theta} f(\theta, \omega) \\ -\nabla_{\omega} f(\theta, \omega) \end{bmatrix} \quad (10)$$

where  $f(\theta, \omega)$  is either (8) or (9). (Note that GAD and EG are different discretizations of the same continuous-time process (Diakonikolas & Orecchia, 2019).) Suppose that the initial point  $(\theta(0), \omega(0))$  is far away from NE:  $\theta(0) \cdot \omega(0) > 0.5$ . Then (10) converges to a non-equilibrium stationary



point on  $\{\theta\omega = 0.5\}$ .

On the other hand, even when initialized at a stationary point such that  $\theta_1 \cdot \omega_1 = 0.5$ , the MixedNE-LD still decreases the distance to NE in expectation:

$$\mathbb{E}\theta_3 \cdot \omega_3 = \theta_1\omega_1 - 4\eta^2 (\eta (\theta_1^2 + \omega_1^2) + 14\eta^2) < \theta_1 \cdot \omega_1 \quad (11)$$

where  $\eta$  is the step-size, and the expectation is over the randomness of the algorithm.

In words, depending on the initialization, the (continuous-time) training dynamics of GAD and EG will either get trapped by non-equilibrium stationary points, or converge to NE. In contrast, the MixedNE-LD is always able to escape non-equilibrium stationary points in expectation.

Figures 1 and 2 demonstrate the empirical behavior of the three algorithms, which is in perfect accordance with the theory. When initialized far away from NE, Figure (1a), (1b), (2a), and (2b) show that GAD and EG get trapped by local stationary points, while MixedNE-LD is able to escape after staying a few iterations near the non-equilibrium states. On the other hand, if initialized sufficiently close to NE, then EG tends to perform better than GAD, as indicated by previous work; see Figure (1c), (1d), (2c), (2d).

Finally, one can ask whether the negative results for GAD and EG are sensitive to the choice of step-size. For instance, we have implemented the vanilla GAD and EG, while in practice one always uses adaptive step-size based on approximate second-order information (Duchi et al., 2011; Kingma & Ba, 2014). However, our next theorem shows that, even with *perfect* second-order information, the training dynamics of GAD and EG still are unable to escape stationary points.

**Theorem 2.** Consider the Newton’s dynamics for solving either (8) or (9):

$$\begin{bmatrix} \frac{d\theta}{dt}(t) \\ \frac{d\omega}{dt}(t) \end{bmatrix} = \begin{bmatrix} \nabla_{\theta}^2 f(\theta, \omega) & 0 \\ 0 & \nabla_{\omega}^2 f(\theta, \omega) \end{bmatrix}^{-1} \begin{bmatrix} \nabla_{\theta} f(\theta, \omega) \\ -\nabla_{\omega} f(\theta, \omega) \end{bmatrix}. \quad (12)$$

Then we have  $\theta(t) \cdot \omega(t) = \theta(0) \cdot \omega(0)$ .

A consequence of of **Theorem 2** is that if we initialize at any point such that  $\theta(0) \cdot \omega(0) \neq 0$ , the training dynamics will remain far away from  $(0, 0)$ , which is the desired NE. Indeed, in Section 5, we shall see that MixedNE-LD outperforms GAD and EG even with adaptivity.

#### 4.4. A Digression: Sampling v.s. Optimization

We would like to demonstrate an additional intriguing behavior of the sampling nature of MixedNE-LD, which we deem as a benefit over deterministic optimization algorithms.

Consider the following SPP:

$$\max_{\theta \in [-2, 2]} \min_{\omega \in [-2, 2]} f(\theta, \omega) = \theta^2 \omega^2. \quad (13)$$

This is a simple SPP where the stationary points  $\{(\theta, 0) \mid \theta \in [-2, 2]\}$  are all NE. Consequently, both GAD and EG succeed in finding an NE, regardless of the initialization; see Figure 3.

The MixedNE-LD, nonetheless, does something slightly more than finding an NE: The MixedNE-LD *explores* among all the NE, inducing a *distribution* on the set of all equilibria; see Figure (3b) and (3d). As exploration is a desirable property in RL, our experiments illustrate yet another advantage of pursuing the mixed NE over pure NE.

## 5. Experiments

In this section, we demonstrate the effectiveness of using the MixedNE-LD framework to solve the robust RL problem.

### 5.1. Off-Policy (DDPG) Experiments

**Two-Player DDPG:** As a case study, we consider NR-MDP setting with  $\delta = 0.1$  (as recommended in Section 6.3 of (Tessler et al., 2019)). We design a two-player variant of DDPG (Lillicrap et al., 2015) algorithm by adapting the Algorithm 1. As opposed to standard DDPG, in two-player DDPG two actor networks output two deterministic policies, the protagonist and adversary policies, denoted by  $\mu_{\theta}$  and  $\nu_{\omega}$ . The critic is trained to estimate the Q-function of the joint-policy. The gradients of the protagonist and adversary parameters are given in Proposition 5 of (Tessler et al., 2019). The resulting algorithm is given in Algorithm 3.

We compare the performance of our algorithm against the baseline algorithm proposed in (Tessler et al., 2019) (see Algorithm 4 with GAD). (Tessler et al., 2019) have suggested a training ratio of 1 : 1 for actors and critic updates. Note that the action noise is injected while collecting transitions for the replay buffer. In (Fujimoto et al., 2018), authors noted that the action noise drawn from the Ornstein-Uhlenbeck (Uhlenbeck & Ornstein, 1930) process offered no performance benefits. Thus we also consider uncorrelated Gaussian noise. In addition to the baseline from (Tessler et al., 2019), we have also considered another baseline, namely Algorithm 4 with Extra-Adam (Gidel et al., 2018).

**Setup:** We evaluate the performance of Algorithm 3 and Algorithm 4 (with GAD and Extra-Adam) on standard continuous control benchmarks available on OpenAI Gym (Brockman et al., 2016) utilizing the MuJoCo environment (Todorov et al., 2012). Specifically, we benchmark on eight tasks: Walker, Hopper, Half-Cheetah, Ant, Swimmer, Reacher, Humanoid, and InvertedPendulum. Details of

## Robust Reinforcement Learning via Adversarial training with Langevin Dynamics

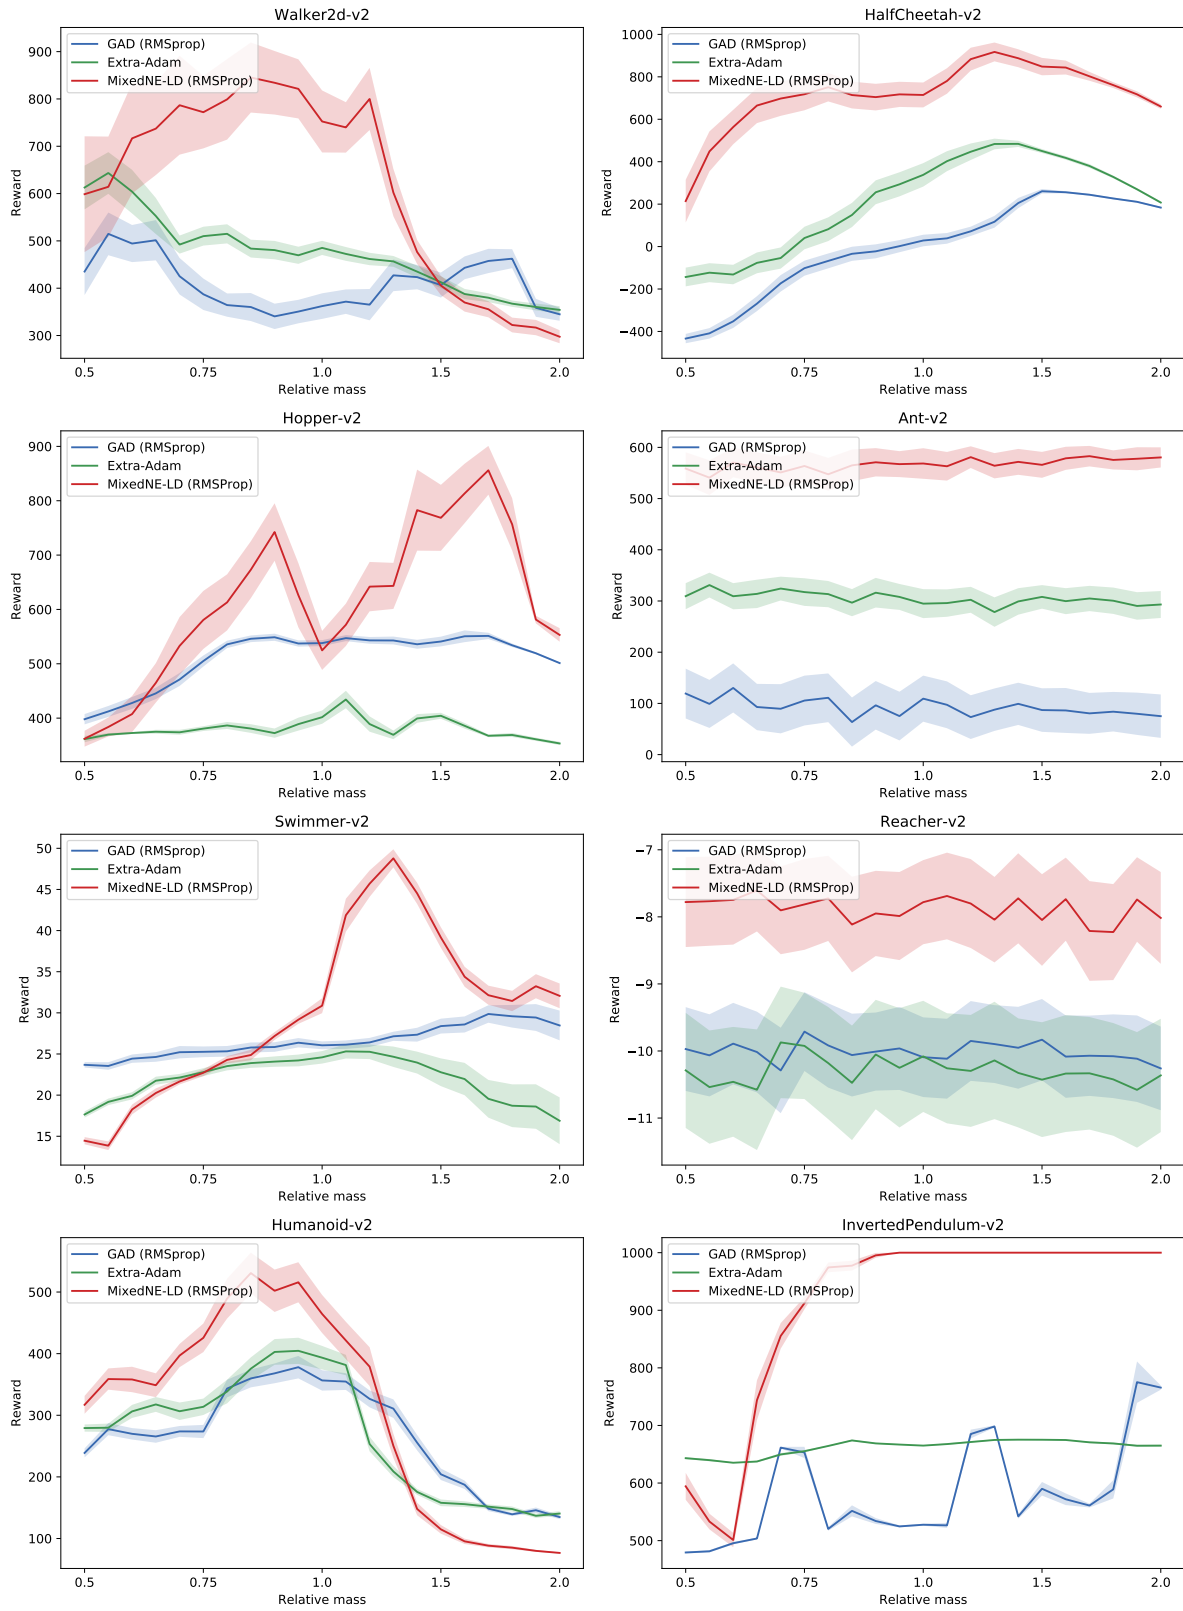


Figure 4. Average performance (over 5 seeds) of Algorithm 3, and Algorithm 4 (with GAD and Extra-Adam), under the NR-MDP setting with  $\delta = 0.1$ . The evaluation is performed without adversarial perturbations, on a range of mass values not encountered during training.

## Robust Reinforcement Learning via Adversarial training with Langevin Dynamics

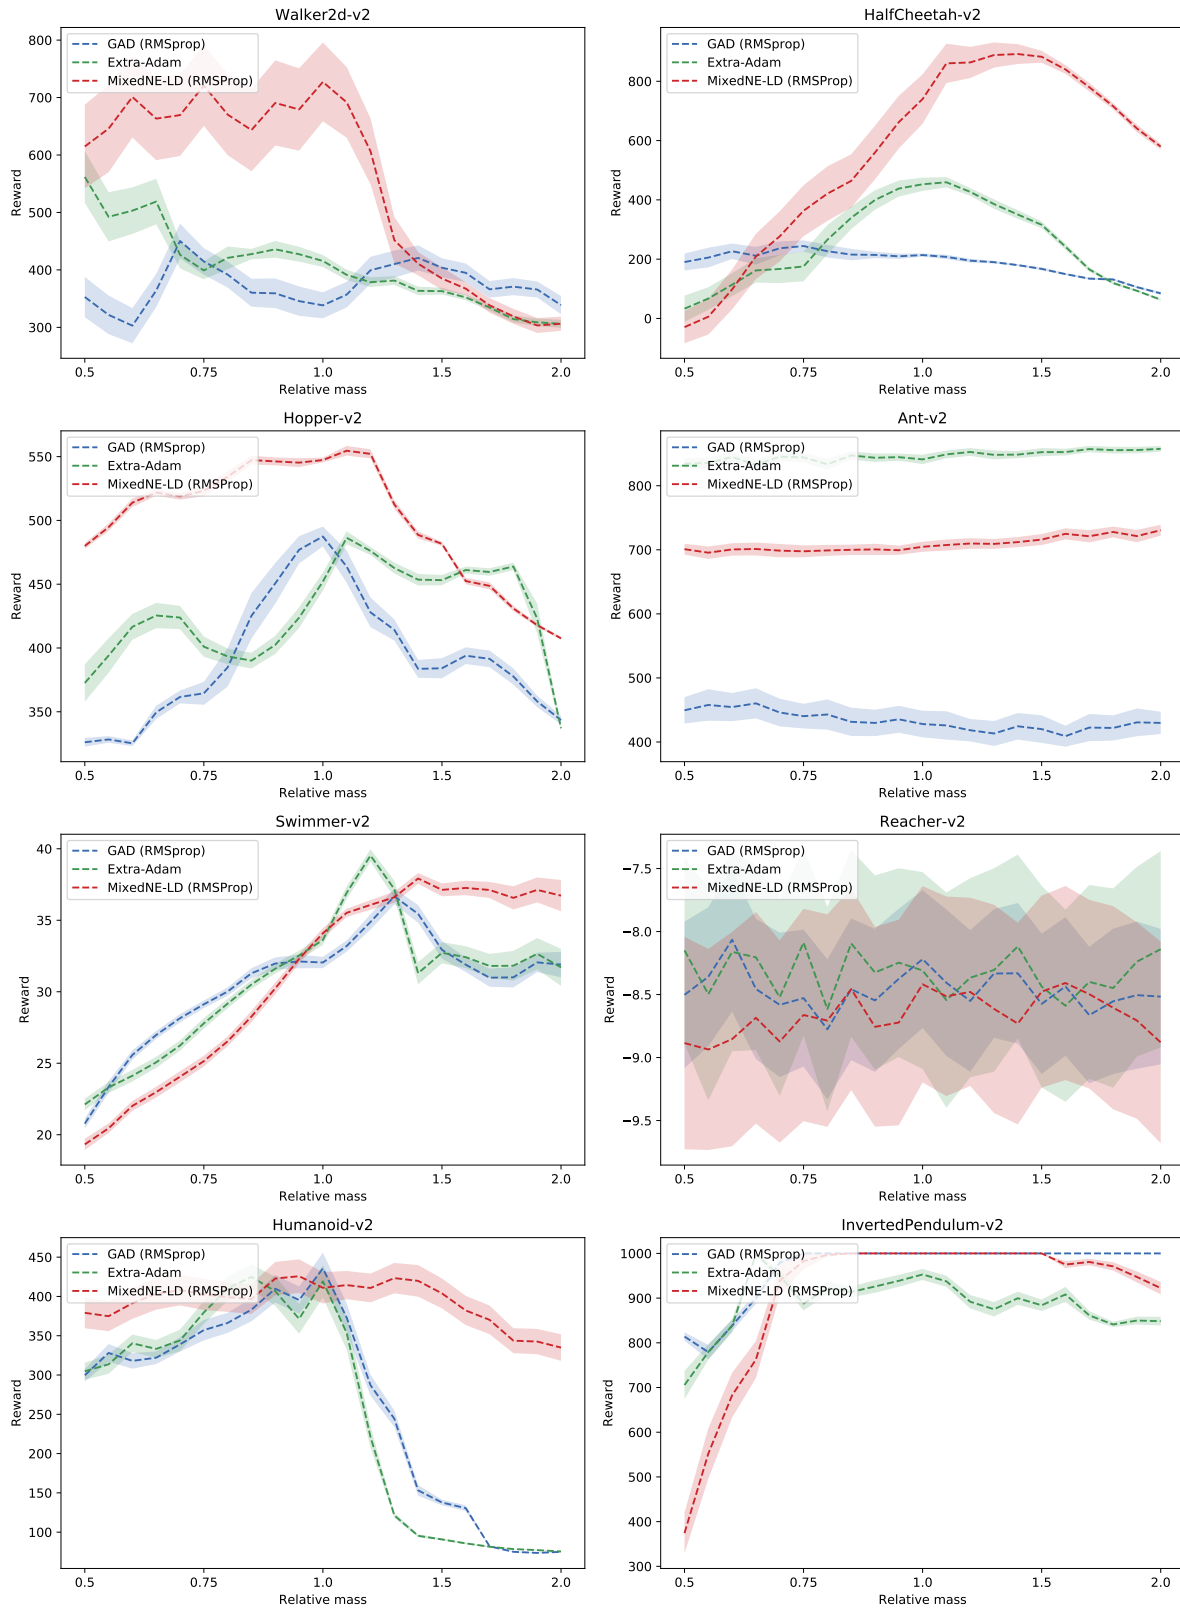


Figure 5. Average performance (over 5 seeds) of Algorithm 3, and Algorithm 4 (with GAD and Extra-Adam), under the NR-MDP setting with  $\delta = 0$ . The evaluation is performed without adversarial perturbations, on a range of mass values not encountered during training.

these environments can be found in (Brockman et al., 2016) and on the GitHub website.

The Algorithm 3 implementation is based on the codebase from (Tessler et al., 2019). For all the algorithms, we use a two-layer feedforward neural network structure of (64, 64, tanh) for both actors (agent and adversary) and critic. The optimizer we use to update the critic is Adam (Kingma & Ba, 2015) with a learning rate of  $10^{-3}$ . The target networks are soft-updated with  $\tau = 0.999$ .

For the GAD baseline, the actors are trained with RMSProp optimizer. For our algorithm (MixedNE-LD), the actors are updated according to Algorithm 1 with warmup steps  $K_t = \min\{15, \lfloor(1 + 10^{-5})^t\rfloor\}$ , and thermal noise  $\sigma_t = \sigma_0 \times (1 - 5 \times 10^{-5})^t$ . The hyperparameters that are not related to exploration (see Table 1) are identical to all the algorithms that are compared.

And we tuned only the exploration-related hyperparameters (for all the algorithms) by grid search: (a) Algorithm 3 with  $(\sigma_0, \sigma) \in \{10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}\} \times \{0, 0.01, 0.1, 0.2, 0.3, 0.4\}$ ; (b) Algorithm 4 with  $\sigma \in \{0, 0.01, 0.1, 0.2, 0.3, 0.4\}$ . For each algorithm-environment pair, we identified the best performing exploration hyperparameter configuration (see Tables 2 and 3).

Each algorithm is trained on 0.5M samples (i.e., 0.5M time steps in the environment). We run our experiments, for each environment, with 5 different seeds. The exploration noise is turned off for evaluation.

**Evaluation:** We evaluate the robustness of all the algorithms under different testing conditions, and in the presence of adversarial disturbances in the testing environment. We train the algorithms with the standard mass variables in OpenAI Gym. At test time, we evaluate the learned policies by changing the mass values (without adversarial perturbations) and estimating the cumulative rewards. As shown in Figure 4, our Algorithm 3 outperforms the baselines Algorithm 4 (with GAD and Extra-Adam) in terms of robustness. Note that we obtain superior performance on the inverted pendulum, which is a failure case for (Tessler et al., 2019). We also evaluate the robustness of the learned policies under both test condition changes, and adversarial disturbances (see Figures 6 and 7 in Appendix B).

**One-Player DDPG:** We evaluate the robustness of one-player variants of Algorithm 3, and Algorithm 4 (with GAD and Extra-Adam), i.e., we consider the NR-MDP setting with  $\delta = 0$ . In this case, we set  $K_t = 1$  for Algorithm 3 (this choice of  $K_t$  makes the computational complexity of both algorithms equal). The results are presented in Figures 5, 8, and 9 (cf. Appendix B).

Here, we remark that Algorithm 3 with  $\delta = 0$ , and  $K_t = 1$

is simply the standard DDPG with actor being updated by preconditioned version of SGLD. Thus we achieve robustness under different testing conditions with just a simple change in the DDPG algorithm and without additional computational cost.

## 5.2. On-Policy (VPG) Experiments

In addition to the off-policy experiments, we test the effectiveness of the MixedNE-LD strategy with the vanilla policy gradient (VPG) method on a toy MDP problem. In particular, we design a two-player variant of VPG (Sutton et al., 2000) algorithm (cf. Algorithm 5) by adapting the Algorithm 1.

**Setup:** We compare the performance of Algorithm 5 and Algorithm 6 (with GAD and Extra-Adam) on a parametrized class of MDPs  $\{\mathcal{M}_\rho = (\mathcal{S}, \mathcal{A}, T_\rho, \gamma, P_0, R) : \rho \in [0, 0.4]\}$ . Here  $\mathcal{S} = [-10, 10]$ ,  $\mathcal{A} = [-1, 1]$ , and  $R(s) = \sin(\sqrt{1.7}s) + \cos(\sqrt{0.3}s) + 3$ . The transition dynamics  $T_\rho$  is defined as follows: given the current state and action  $(s_t, a_t)$ , the next state is  $s_{t+1} = s_t + a_t$  with probability  $1 - \rho$ , and  $s_{t+1} = s_t + a'$  (where  $a' \sim \text{unif}([-1, 1])$ ) with probability  $\rho$ . We also ensure that  $s_{t+1} \in [-10, 10]$ .

For all the algorithms, we use a two-layer feedforward neural network structure of (16, 16, relu) for both actors (agent and adversary). The relevant hyperparameters are given in Tables 4, 5, and 6. Each algorithm is trained for 5000 steps. We run our experiments with 5 different seeds.

**Evaluation:** We train the algorithms with a nominal environment parameter  $\rho = 0.2$ , and evaluate the learned policies on a range of  $\rho \in [0, 0.4]$  values. As shown in Figure 10 (cf. Appendix C), our Algorithm 5 outperforms the baselines Algorithm 6 (with GAD and Extra-Adam) in terms of robustness (in both two-player and one-player settings).

## 5.3. Robustness of One-Player MixedNE-LD

Consider the standard (non-robust) RL objective of maximizing  $J(\theta) = \mathbb{E}[\sum_{t=1}^{\infty} \gamma^{t-1} r_t \mid \mu_\theta, \mathcal{M}_1]$ . We can translate this non-convex problem into an infinite dimensional convex-problem by considering a distribution over deterministic policies as follows (Liu et al., 2017):

$$\max_{p \in \mathcal{P}(\Theta)} \mathbb{E}_{\theta \sim p} [J(\theta)] + \lambda H(p),$$

where  $H(p) = \mathbb{E}_{\theta \sim p} [-\log p(\theta)]$  is the entropy of the distribution  $p$ . The robust behavior of this objective (in the context of loss surface) is discussed in (Chaudhari et al., 2019). The optimal solution to the above problem takes the form:  $p_\lambda^*(\theta) \propto \exp(\frac{1}{\lambda} J(\theta))$ . For a given  $\lambda$ , SGLD can be used to draw samples from  $p_\lambda^*(\theta)$ .



Then, the resulting algorithm is equivalent to Algorithm 3 with  $\delta = 0$ . Note that in our one-player DDPG experiments, we obtained significant improvement over both robust and non-robust baselines (see Figure 14) even with single inner loop iteration ( $K_t = 1$ ). Since the Algorithm 3 is computationally demanding even though it uses a mean approximation in its inner loop, this new approximation by setting  $\delta = 0$  and  $K_t = 1$  is preferred in practice.

## 6. Conclusion

In this work, we study the robust reinforcement learning problem. By adapting the approximate infinite-dimensional entropic mirror descent from (Hsieh et al., 2019), we design a robust variant of DDPG algorithm, under the NR-MDP setting. To the best of our knowledge, this is the first work to apply SGLD for the robust RL problem. In our experiments, we evaluate the robustness of our algorithm on several continuous control tasks, and found that our algorithm clearly outperforms the robust and non-robust baselines while tackling the failure case (i.e., inverted pendulum) for the earlier literature. Intriguingly, even the simple version of the algorithm with a single Langevin step results in competitive results with a desirable computational complexity.

## Acknowledgements

This work has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement n 725594 - time-data), the Swiss National Science Foundation (SNSF) under grant number 407540.167319, and the Army Research Office under grant number W911NF-19-1-0404.

## References

- Abernethy, J., Lai, K. A., and Wibisono, A. Last-iterate convergence rates for min-max optimization. *arXiv preprint arXiv:1906.02027*, 2019.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- Bubeck, S., Cohen, M. B., Lee, Y. T., Lee, J. R., and Madry, A. K-server via multiscale entropic regularization. In *Proceedings of the 50th annual ACM SIGACT symposium on theory of computing*, pp. 3–16, 2018.
- Chaudhari, P., Choromanska, A., Soatto, S., LeCun, Y., Baldassi, C., Borgs, C., Chayes, J., Sagun, L., and Zecchina, R. Entropy-sgd: Biasing gradient descent into wide valleys. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(12):124018, 2019.
- Dasgupta, P. and Maskin, E. The existence of equilibrium in discontinuous economic games, i: Theory. *The Review of economic studies*, 53(1):1–26, 1986.
- Daskalakis, C. and Panageas, I. Last-iterate convergence: Zero-sum games and constrained min-max optimization. *Innovations in Theoretical Computer Science*, 2019.
- Dhariwal, P., Hesse, C., Klimov, O., Nichol, A., Plappert, M., Radford, A., Schulman, J., Sidor, S., Wu, Y., and Zhokhov, P. Openai baselines. <https://github.com/openai/baselines>, 2017.
- Diakonikolas, J. and Orecchia, L. The approximate duality gap technique: A unified theory of first-order methods. *SIAM Journal on Optimization*, 29(1):660–689, 2019.
- Duchi, J., Hazan, E., and Singer, Y. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(Jul):2121–2159, 2011.
- Fujimoto, S., van Hoof, H., and Meger, D. Addressing function approximation error in actor-critic methods. *arXiv preprint arXiv:1802.09477*, 2018.
- Gidel, G., Berard, H., Vignoud, G., Vincent, P., and Lacoste-Julien, S. A variational inequality perspective on generative adversarial networks. *arXiv preprint arXiv:1802.10551*, 2018.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in neural information processing systems*, pp. 2672–2680, 2014.

- Hsieh, Y.-P., Liu, C., and Cevher, V. Finding mixed nash equilibria of generative adversarial networks. In *International Conference on Machine Learning*, pp. 2810–2819, 2019.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kingma, D. P. and Ba, J. A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, volume 5, 2015.
- Levine, S., Finn, C., Darrell, T., and Abbeel, P. End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research*, 17(1):1334–1373, 2016.
- Li, C., Chen, C., Carlson, D. E., and Carin, L. Preconditioned stochastic gradient langevin dynamics for deep neural networks. In *AAAI*, volume 2, pp. 4, 2016.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- Littman, M. L. Markov games as a framework for multi-agent reinforcement learning. In *Machine Learning Proceedings*. Elsevier, 1994.
- Liu, Y., Ramachandran, P., Liu, Q., and Peng, J. Stein variational policy gradient. In *Proceedings of the 33rd Conference on Uncertainty in Artificial Intelligence*, 2017.
- Lucic, M., Kurach, K., Michalski, M., Gelly, S., and Bousquet, O. Are gans created equal? a large-scale study. In *Advances in neural information processing systems*, pp. 700–709, 2018.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529, 2015.
- Perolat, J., Scherrer, B., Piot, B., and Pietquin, O. Approximate dynamic programming for two-player zero-sum Markov games. In *International Conference on Machine Learning*, 2015.
- Pinto, L., Davidson, J., Sukthankar, R., and Gupta, A. Robust adversarial reinforcement learning. In *International Conference on Machine Learning*, 2017.
- Schulman, J., Levine, S., Abbeel, P., Jordan, M., and Moritz, P. Trust region policy optimization. In *International Conference on Machine Learning*, pp. 1889–1897, 2015.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Silver, D., Lever, G., Heess, N., Degris, T., Wierstra, D., and Riedmiller, M. Deterministic policy gradient algorithms. In *ICML*, 2014.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354, 2017.
- Sutton, R. S., McAllester, D. A., Singh, S. P., and Mansour, Y. Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, pp. 1057–1063, 2000.
- Tessler, C., Efroni, Y., and Mannor, S. Action robust reinforcement learning and applications in continuous control. *arXiv preprint arXiv:1901.09184*, 2019.
- Todorov, E., Erez, T., and Tassa, Y. Mujoco: A physics engine for model-based control. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pp. 5026–5033. IEEE, 2012.
- Uhlenbeck, G. E. and Ornstein, L. S. On the theory of the brownian motion. *Physical review*, 36(5):823, 1930.
- Welling, M. and Teh, Y. W. Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pp. 681–688, 2011.

**Algorithm 2** Algorithms in Section 4 (MixedNE-LD / GAD / EG)

---

**Input:** step-size  $\{\eta_t\}_{t=1}^T$ , thermal-noise  $\{\epsilon_t\}_{t=1}^T$ , warmup steps  $\{K_t\}_{t=1}^T$ , exponential damping factor  $\beta$ .  
**for**  $t = 1, 2, \dots, T - 1$  **do**

**MixedNE-LD:**

$$\bar{\omega}_t, \omega_t^{(1)} \leftarrow \omega_t; \bar{\theta}_t, \theta_t^{(1)} \leftarrow \theta_t$$

**for**  $k = 1, 2, \dots, K_t$  **do**

$$\xi, \xi' \sim \mathcal{N}(0, I)$$

$$\theta_t^{(k+1)} \leftarrow \Pi_{\Theta} \left( \theta_t^{(k)} + \eta_t \nabla_{\theta} f(\theta_t^{(k)}, \omega_t) + \epsilon_t \sqrt{2\eta_t} \xi' \right)$$

$$\omega_t^{(k+1)} \leftarrow \Pi_{\Omega} \left( \omega_t^{(k)} - \eta_t \nabla_{\omega} f(\theta_t, \omega_t^{(k)}) + \epsilon_t \sqrt{2\eta_t} \xi \right)$$

$$\bar{\omega}_t \leftarrow (1 - \beta) \bar{\omega}_t + \beta \omega_t^{(k+1)}$$

$$\bar{\theta}_t \leftarrow (1 - \beta) \bar{\theta}_t + \beta \theta_t^{(k+1)}$$

**end for**

$$\theta_{t+1} \leftarrow (1 - \beta) \theta_t + \beta \bar{\theta}_t$$

$$\omega_{t+1} \leftarrow (1 - \beta) \omega_t + \beta \bar{\omega}_t$$

**GAD (Gradient Ascent Descent):**

$$\theta_{t+1} \leftarrow \Pi_{\Theta} (\theta_t + \eta_t \nabla_{\theta} f(\theta_t, \omega_t))$$

$$\omega_{t+1} \leftarrow \Pi_{\Omega} (\omega_t - \eta_t \nabla_{\omega} f(\theta_{t+1}, \omega_t))$$

**EG (Extra-Gradient):**

$$\theta_{t+\frac{1}{2}} \leftarrow \Pi_{\Theta} (\theta_t + \eta_t \nabla_{\theta} f(\theta_t, \omega_t))$$

$$\omega_{t+\frac{1}{2}} \leftarrow \Pi_{\Omega} (\omega_t - \eta_t \nabla_{\omega} f(\theta_t, \omega_t))$$

$$\theta_{t+1} \leftarrow \Pi_{\Theta} \left( \theta_t + \eta_t \nabla_{\theta} f(\theta_{t+\frac{1}{2}}, \omega_{t+\frac{1}{2}}) \right)$$

$$\omega_{t+1} \leftarrow \Pi_{\Omega} \left( \omega_t - \eta_t \nabla_{\omega} f(\theta_{t+\frac{1}{2}}, \omega_{t+\frac{1}{2}}) \right)$$

**end for**

**Output:**  $\omega_T, \theta_T$ .

---

## A. Algorithms and Omitted Proofs for Section 4

### A.1. Algorithms and Hyperparameters

The pseudocode of the algorithms can be found in **Algorithm 2** (the symbol  $\Pi$  denotes the projection). The hyperparameter setting for experiments in Section 4 is:

- Algorithm 2 with GDA, and  $\eta_t = 0.1$
- Algorithm 2 with EG, and  $\eta_t = 0.1$
- Algorithm 2 with MixedNE-LD,  $\eta_t = 0.1$ ,  $\epsilon_t = 0.01$ ,  $K_t = 50$ , and  $\beta = 0.5$ .

We also note that we focus on the “last iterate” convergence for EG (Abernethy et al., 2019; Daskalakis & Panageas, 2019), instead of the usual ergodic average in convex optimization literature. This is because, in practice, people almost exclusively use the last iterate.

### A.2. Proof of Theorem 1

We will focus on the case  $f(\theta, \omega) = \theta^2 \omega^2 - \theta \omega$ . Without loss of generality, we may also assume that  $\omega(0) > \theta(0) > 0$ ; the proof of the other cases follows the same argument.

Let  $(\theta(t), \omega(t))$  follow the dynamics (10) with  $\theta(0) \cdot \omega(0) > 0.5$ . Assume, for the moment, that both  $\theta$  and  $\omega$  are without constraint. Then we have

$$\begin{aligned} \frac{1}{2} \frac{d}{dt} (\theta(t)^2 + \omega(t)^2) &= \theta \frac{d\theta}{dt} + \omega \frac{d\omega}{dt} \\ &= 2\theta^2\omega^2 - \theta\omega + (-2\theta^2\omega^2 + \theta\omega) \\ &= 0 \end{aligned}$$

implying that  $\theta^2(t) + \omega^2(t) = \theta^2(0) + \omega^2(0)$  for all  $t$ . Therefore  $(r \cos(t + \phi_1), r \sin(t + \phi_2))$ , where  $(r \cos \phi_1, r \sin \phi_2) = (\theta(0), \omega(0))$ , is a solution for dynamics for small enough  $t$ .

On the other hand, we have

$$\begin{aligned} \frac{d}{dt} (\theta(t)\omega(t)) &= \frac{d\theta}{dt}(t) \cdot \omega(t) + \theta(t) \cdot \frac{d\omega}{dt}(t) \\ &= 2\theta(t)\omega^3(t) - \omega^2(t) + (-2\theta^3(t)\omega(t) + \theta^2(t)) \\ &= (\theta^2(t) - \omega^2(t)) (1 - 2\theta(t)\omega(t)) \\ &= (\theta^2(t) - \omega^2(t)) (1 - 2r^2 \cos(t + \phi_1) \sin(t + \phi_2)). \end{aligned}$$

When  $t = 0$ , we have  $1 - 2r^2 \cos(t + \phi_1) \sin(t + \phi_2) = 1 - 2\theta(0)\omega(0) < 0$ . When  $t = \frac{\pi}{t}$ , we have

$$\begin{aligned} 2r^2 \cos(t + \phi_1) \sin(t + \phi_2) &= 2r^2 \left( \frac{\sqrt{2}}{2} \cos \phi_1 - \frac{\sqrt{2}}{2} \sin \phi_1 \right) \left( \frac{\sqrt{2}}{2} \cos \phi_2 + \frac{\sqrt{2}}{2} \sin \phi_2 \right) \\ &= (\theta(0) - \sqrt{r^2 - \theta(0)^2}) (\sqrt{r^2 - \omega(0)^2} + \omega(0)) \\ &= (\theta^2(0) - \omega^2(0)) < 0 \end{aligned}$$

whence  $1 - 2r^2 \cos(t + \phi_1) \sin(t + \phi_2) > 0$ . The intermediate value theorem then implies that there exists a  $\tilde{t}$  such that  $1 - 2\theta(\tilde{t})\omega(\tilde{t}) = 0$ . But since  $\{(\theta, \omega) \mid 2\theta\omega = 1\}$  are the stationary points of the dynamics (10), we conclude that  $\frac{d}{dt} (\theta(t)\omega(t)) = 0$  whenever  $t \geq \tilde{t}$ ; that is,  $(\theta(t), \omega(t))$  gets trapped at the stationary point  $(\theta(\tilde{t}), \omega(\tilde{t}))$ . This concludes the first part of the theorem when there is no boundary.

If the boundary is present, the dynamics (10) should be modified to the *projected dynamics* (Bubeck et al., 2018) and the proof remains the same, except that when  $(\theta(t), \omega(t))$  hits the boundary, the curve needs to traverse along the boundary to decrease the norm.

We now turn to the statement for MixedNE-LD. Let  $(\theta_1, \omega_1)$  be initialized at any stationary point:  $\theta_1\omega_1 = 0.5$ . Consider the two-step evolution of MixedNE-LD:

$$\begin{aligned} \theta_2 &= \theta_1 + \sqrt{2\eta}\xi, \\ \omega_2 &= \omega_1 + \sqrt{2\eta}\xi', \\ \theta_3 &= \theta_2 + \eta(2\theta_2\omega_2^2 - \omega_2) + \sqrt{2\eta}\xi'', \\ \omega_3 &= \omega_2 - \eta(2\theta_2^2\omega_2 - \theta_2) + \sqrt{2\eta}\xi''' \end{aligned}$$

where  $\xi, \xi', \xi''$ , and  $\xi'''$  are independent standard Gaussian. Since we initialize at a stationary point  $\theta_1\omega_1 = 0.5$ , we have

$$\begin{aligned} 2\theta_2\omega_2 - 1 &= 2\theta_1\omega_1 + \sqrt{2\eta}\omega_1\xi + \sqrt{2\eta}\theta_1\xi' + 2\eta\xi\xi' - 1 \\ &= \sqrt{2\eta}\omega_1\xi + \sqrt{2\eta}\theta_1\xi' + 2\eta\xi\xi'. \end{aligned} \tag{14}$$

Using the towering property of the expectation, (14), and the fact that  $\xi, \xi', \xi''$ , and  $\xi'''$  are independent standard Gaussian, we compute

$$\begin{aligned} \mathbb{E}\theta_3\omega_3 &= \mathbb{E}[\mathbb{E}[\theta_3\omega_3 \mid \theta_2, \omega_2]] \\ &= \mathbb{E}\left[\mathbb{E}\left[\left(\theta_2 + \eta(2\theta_2\omega_2^2 - \omega_2) + \sqrt{2\eta}\xi''\right) \left(\omega_2 - \eta(2\theta_2^2\omega_2 - \theta_2) + \sqrt{2\eta}\xi'''\right) \mid \theta_2, \omega_2\right]\right] \end{aligned}$$

$$\begin{aligned}
 &= \mathbb{E} \left[ \mathbb{E} \left[ (\theta_2 + \eta (2\theta_2\omega_2^2 - \omega_2)) (\omega_2 - \eta (2\theta_2^2\omega_2 - \theta_2)) \mid \theta_2, \omega_2 \right] \right] \\
 &= \mathbb{E} \left[ (\theta_2 + \eta\omega_2 (2\theta_2\omega_2 - 1)) (\omega_2 - \eta\theta_2 (2\theta_2\omega_2 - 1)) \right] \\
 &= \mathbb{E} \left[ \theta_2\omega_2 - \eta\theta_2^2 (2\theta_2\omega_2 - 1) + \eta\omega_2^2 (2\theta_2\omega_2 - 1) - \eta^2\theta_2\omega_2 (2\theta_2\omega_2 - 1)^2 \right] \\
 &= \mathbb{E} \left[ \theta_1\omega_1 - \eta \left( \theta_1^2 + 2\eta\xi^2 + 2\sqrt{2\eta}\theta_1\xi - \omega_1^2 - 2\eta\xi'^2 - 2\sqrt{2\eta}\omega_1\xi' \right) \left( \sqrt{2\eta}\omega_1\xi + \sqrt{2\eta}\theta_1\xi' + 2\eta\xi\xi' \right) \right. \\
 &\quad \left. - 4\eta^2 \left( \sqrt{2\eta}\omega_1\xi + \sqrt{2\eta}\theta_1\xi' + 2\eta\xi\xi' \right) \left( 2\eta\omega_1^2\xi^2 + 2\eta\theta_1^2\xi'^2 + 4\eta^2\xi^2\xi'^2 + 2\eta\xi\xi' + 4\sqrt{2\eta}\theta_1\xi\xi'^2 + 4\sqrt{2\eta}\omega_2\xi^2\xi' \right) \right] \\
 &= \theta_1\omega_1 - 0 - 4\eta^2 (\eta\omega_2^2 + \eta\theta_1^2 + 2\eta^2 + 4\eta^2 + 4\eta^2 + 4\eta^2) \\
 &= \theta_1\omega_1 - 4\eta^2 (\eta (\theta_1^2 + \omega_1^2) + 14\eta^2)
 \end{aligned}$$

which is (11).

### A.3. Proof of Theorem 2

Spelling out the Newton dynamics (12), we get

$$\begin{aligned}
 \frac{d\theta}{dt}(t) &= \frac{1}{2\omega^2(t)} (2\theta(t)\omega^2(t) - \omega(t)) \\
 &= \theta(t) - \frac{1}{2\omega(t)}
 \end{aligned}$$

and similarly  $\frac{d\omega}{dt}(t) = -\omega(t) + \frac{1}{2\theta(t)}$ . As a result, we have

$$\begin{aligned}
 \frac{d}{dt} (\theta(t)\omega(t)) &= \frac{d\theta}{dt}(t) \cdot \omega(t) + \theta(t) \cdot \frac{d\omega}{dt}(t) \\
 &= \theta(t)\omega(t) - \frac{1}{2} - \theta(t)\omega(t) + \frac{1}{2} \\
 &= 0
 \end{aligned}$$

which concludes the proof.



## B. Off-Policy (DDPG) Experiments: Algorithms, Hyperparameters, and Results

- Algorithms:
  1. MixedNE-LD: Algorithm 3
  2. Baselines: Algorithm 4 (with GAD and Extra-Adam)
- Hyperparameters:
  1. Common hyperparameters for Algorithm 3 and Algorithm 4: Table 1
  2. Exploration-related hyperparameters for Algorithm 3 and Algorithm 4 (the best performing values for every environment are presented): Tables 2 and 3
- Results:
  1. Heat maps for NR-MDP setting with  $\delta = 0.1$  (Figures 6 and 7)
  2. Heat maps for NR-MDP setting with  $\delta = 0$  (Figures 8 and 9)

## C. On-Policy (VPG) Experiments: Algorithms, and Hyperparameters, and Results

- Algorithms:
  1. MixedNE-LD: Algorithm 5
  2. Baselines: Algorithm 6 (with GAD and Extra-Adam)
- Hyperparameters:
  1. Common hyperparameters for Algorithm 5 and Algorithm 6: Table 4
  2. Additional hyperparameters for Algorithm 5 and Algorithm 6 (the best performing values are presented): Tables 5 and 6
- Results:
  1. NR-MDP setting with  $\delta = 0.1$  (Figure 10a)
  2. NR-MDP setting with  $\delta = 0$  (Figure 10b)

## D. Ablation Study

- Ablation on  $(\beta, K_t)$ : see Figures 11, 12, and 13.
- Ablation on  $\delta$ : see Figures 14 and 15. If the  $\delta$  value is way larger (overly conservative) than the requirement (range of environmental changes), it could negatively impact the generalization ability. Choosing the appropriate value of delta is problem dependent.

## E. Code

The code repository (for all the experiments): <https://github.com/DaDaCheng/LIONS-RL/tree/master/Robust-Reinforcement-Learning-via-Adversarial-training-with-Langevin-Dynamics>.

Table 1. Common hyperparameters for Algorithm 3 and Algorithm 4, where most of the values are chosen from (Dhariwal et al., 2017).

Hyperparameter	Value
critic optimizer	Adam
critic learning rate	$10^{-3}$
target update rate $\tau$	0.999
mini-batch size $N$	128
discount factor $\gamma$	0.99
damping factor $\beta$	0.9
replay buffer size	$10^6$
action noise parameter $\sigma$	$\{0, 0.01, 0.1, 0.2, 0.3, 0.4\}$
RMSProp parameter $\alpha$	0.999
RMSProp parameter $\epsilon$	$10^{-8}$
RMSProp parameter $\eta$	$10^{-4}$
thermal noise $\sigma_t$ (Algorithm 3)	$\sigma_0 \times (1 - 5 \times 10^{-5})^t$ , where $\sigma_0 \in \{10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}\}$
warmup steps $K_t$ (Algorithm 3)	$\min \{15, \lfloor (1 + 10^{-5})^t \rfloor\}$

Table 2. Exploration-related hyperparameters for Algorithm 3 and Algorithm 4 chosen via grid search (for NR-MDP setting with  $\delta = 0.1$ ).

	Algorithm 3: $(\sigma_0, \sigma)$	Algorithm 4 (with GAD): $\sigma$	Algorithm 4 (with Extra-Adam): $\sigma$
Walker-v2	$(10^{-2}, 0.01)$	0	0.3
HalfCheetah-v2	$(10^{-2}, 0)$	0.2	0.01
Hopper-v2	$(10^{-3}, 0.2)$	0.2	0.3
Ant-v2	$(10^{-4}, 0.2)$	0.4	0.01
Swimmer-v2	$(10^{-5}, 0.4)$	0.4	0.4
Reacher-v2	$(10^{-3}, 0.2)$	0.4	0.2
Humanoid-v2	$(10^{-4}, 0.01)$	0	0.01
InvertedPendulum-v2	$(10^{-3}, 0.01)$	0.1	0.01

Table 3. Exploration-related hyperparameters for Algorithm 3 and Algorithm 4 chosen via grid search (for NR-MDP setting with  $\delta = 0$ ).

	Algorithm 3: $(\sigma_0, \sigma)$	Algorithm 4 (with GAD): $\sigma$	Algorithm 4 (with Extra-Adam): $\sigma$
Walker-v2	$(10^{-2}, 0.1)$	0.01	0.2
HalfCheetah-v2	$(10^{-2}, 0.01)$	0.4	0.01
Hopper-v2	$(10^{-5}, 0.3)$	0.4	0.1
Ant-v2	$(10^{-2}, 0.4)$	0.4	0.01
Swimmer-v2	$(10^{-2}, 0.2)$	0.3	0.3
Reacher-v2	$(10^{-3}, 0.2)$	0.3	0.2
Humanoid-v2	$(10^{-2}, 0.1)$	0	0.01
InvertedPendulum-v2	$(10^{-3}, 0)$	0.01	0.01

**Algorithm 3** DDPG with MixedNE-LD (pre-conditioner = RMSProp)

**Hyperparameters:** see Table 1

Initialize (randomly) policy parameters  $\omega_1, \theta_1$ , and Q-function parameter  $\phi$ .

Initialize the target network parameters  $\omega_{\text{targ}} \leftarrow \omega_1, \theta_{\text{targ}} \leftarrow \theta_1$ , and  $\phi_{\text{targ}} \leftarrow \phi$ .

Initialize replay buffer  $\mathcal{D}$ .

Initialize  $m \leftarrow \mathbf{0}$ ;  $m' \leftarrow \mathbf{0}$ .

$t \leftarrow 1$ .

**repeat**

Observe state  $s$ , and select actions  $a = \mu_{\theta_t}(s) + \xi$ ;  $a' = \nu_{\omega_t}(s) + \xi'$ , where  $\xi, \xi' \sim \mathcal{N}(0, \sigma I)$

Execute the action  $\bar{a} = (1 - \delta)a + \delta a'$  in the environment.

Observe reward  $r$ , next state  $s'$ , and done signal  $d$  to indicate whether  $s'$  is terminal.

Store  $(s, \bar{a}, r, s', d)$  in replay buffer  $\mathcal{D}$ .

If  $s'$  is terminal, reset the environment state.

**if** it's time to update **then**

**for** however many updates **do**

$\bar{\omega}_t, \omega_t^{(1)} \leftarrow \omega_t$ ;  $\bar{\theta}_t, \theta_t^{(1)} \leftarrow \theta_t$

**for**  $k = 1, 2, \dots, K_t$  **do**

Sample a random minibatch of  $N$  transitions  $B = \{(s, \bar{a}, r, s', d)\}$  from  $\mathcal{D}$ .

Compute targets  $y(r, s', d) = r + \gamma(1 - d)Q_{\phi_{\text{targ}}}(s', (1 - \delta)\mu_{\theta_{\text{targ}}}(s') + \delta\nu_{\omega_{\text{targ}}}(s'))$ .

Update critic by one step of (preconditioned) gradient descent using  $\nabla_{\phi} L(\phi)$ , where

$$L(\phi) = \frac{1}{N} \sum_{(s, \bar{a}, r, s', d) \in B} (y(r, s', d) - Q_{\phi}(s, \bar{a}))^2.$$

Compute the (agent and adversary) policy gradient estimates:

$$\nabla_{\theta} \widehat{J}(\theta, \omega_t) = \frac{1 - \delta}{N} \sum_{s \in \mathcal{D}} \nabla_{\theta} \mu_{\theta}(s) \nabla_{\bar{a}} Q_{\phi}(s, \bar{a}) \Big|_{\bar{a} = (1 - \delta)\mu_{\theta}(s) + \delta\nu_{\omega_t}(s)}$$

$$\nabla_{\omega} \widehat{J}(\theta_t, \omega) = \frac{\delta}{N} \sum_{s \in \mathcal{D}} \nabla_{\omega} \nu_{\omega}(s) \nabla_{\bar{a}} Q_{\phi}(s, \bar{a}) \Big|_{\bar{a} = (1 - \delta)\mu_{\theta_t}(s) + \delta\nu_{\omega}(s)}.$$

$g \leftarrow \left[ \nabla_{\theta} \widehat{J}(\theta, \omega_t) \right]_{\theta = \theta_t^{(k)}}; m \leftarrow \alpha m + (1 - \alpha)g \odot g$ ;  $C \leftarrow \text{diag}(\sqrt{m} + \epsilon)$

$\theta_t^{(k+1)} \leftarrow \theta_t^{(k)} + \eta C^{-1}g + \sqrt{2\eta}\sigma_t C^{-\frac{1}{2}}\xi$ , where  $\xi \sim \mathcal{N}(0, I)$

$g' \leftarrow \left[ \nabla_{\omega} \widehat{J}(\theta_t, \omega) \right]_{\omega = \omega_t^{(k)}}; m' \leftarrow \alpha m' + (1 - \alpha)g' \odot g'$ ;  $D \leftarrow \text{diag}(\sqrt{m'} + \epsilon)$

$\omega_t^{(k+1)} \leftarrow \omega_t^{(k)} - \eta D^{-1}g' + \sqrt{2\eta}\sigma_t D^{-\frac{1}{2}}\xi'$ , where  $\xi' \sim \mathcal{N}(0, I)$

$\bar{\omega}_t \leftarrow (1 - \beta)\bar{\omega}_t + \beta\omega_t^{(k+1)}$ ;  $\bar{\theta}_t \leftarrow (1 - \beta)\bar{\theta}_t + \beta\theta_t^{(k+1)}$

Update the target networks:

$$\phi_{\text{targ}} \leftarrow \tau\phi_{\text{targ}} + (1 - \tau)\phi$$

$$\theta_{\text{targ}} \leftarrow \tau\theta_{\text{targ}} + (1 - \tau)\theta_t^{(k+1)}$$

$$\omega_{\text{targ}} \leftarrow \tau\omega_{\text{targ}} + (1 - \tau)\omega_t^{(k+1)}$$

**end for**

$\omega_{t+1} \leftarrow (1 - \beta)\omega_t + \beta\bar{\omega}_t$ ;  $\theta_{t+1} \leftarrow (1 - \beta)\theta_t + \beta\bar{\theta}_t$

$t \leftarrow t + 1$ .

**end for**

**end if**

**until** convergence

**Output:**  $\omega_T, \theta_T$ .

**Algorithm 4** DDPG with GAD (pre-conditioner = RMSProp) / Extra-Adam

**Hyperparameters:** see Table 1

Initialize (randomly) policy parameters  $\omega_1, \theta_1$ , and Q-function parameter  $\phi$ .

Initialize the target network parameters  $\omega_{\text{targ}} \leftarrow \omega_1, \theta_{\text{targ}} \leftarrow \theta_1$ , and  $\phi_{\text{targ}} \leftarrow \phi$ .

Initialize replay buffer  $\mathcal{D}$ .

Initialize  $m \leftarrow \mathbf{0}$ ;  $m' \leftarrow \mathbf{0}$ .

$t \leftarrow 1$ .

**repeat**

Observe state  $s$ , and select actions  $a = \mu_{\theta_t}(s) + \xi$ ;  $a' = \nu_{\omega_t}(s) + \xi'$ , where  $\xi, \xi' \sim \mathcal{N}(0, \sigma I)$

Execute the action  $\bar{a} = (1 - \delta)a + \delta a'$  in the environment.

Observe reward  $r$ , next state  $s'$ , and done signal  $d$  to indicate whether  $s'$  is terminal.

Store  $(s, \bar{a}, r, s', d)$  in replay buffer  $\mathcal{D}$ .

If  $s'$  is terminal, reset the environment state.

**if** it's time to update **then**

**for** however many updates **do**

Sample a random minibatch of  $N$  transitions  $B = \{(s, \bar{a}, r, s', d)\}$  from  $\mathcal{D}$ .

Compute targets  $y(r, s', d) = r + \gamma(1 - d)Q_{\phi_{\text{targ}}}(s', (1 - \delta)\mu_{\theta_{\text{targ}}}(s') + \delta\nu_{\omega_{\text{targ}}}(s'))$ .

Update critic by one step of (preconditioned) gradient descent using  $\nabla_{\phi} L(\phi)$ , where

$$L(\phi) = \frac{1}{N} \sum_{(s, \bar{a}, r, s', d) \in B} (y(r, s', d) - Q_{\phi}(s, \bar{a}))^2.$$

Compute the (agent and adversary) policy gradient estimates:

$$\nabla_{\theta} \widehat{J}(\theta, \omega_t) = \frac{1 - \delta}{N} \sum_{s \in \mathcal{D}} \nabla_{\theta} \mu_{\theta}(s) \nabla_{\bar{a}} Q_{\phi}(s, \bar{a}) \Big|_{\bar{a}=(1-\delta)\mu_{\theta}(s)+\delta\nu_{\omega_t}(s)}$$

$$\nabla_{\omega} \widehat{J}(\theta_t, \omega) = \frac{\delta}{N} \sum_{s \in \mathcal{D}} \nabla_{\omega} \nu_{\omega}(s) \nabla_{\bar{a}} Q_{\phi}(s, \bar{a}) \Big|_{\bar{a}=(1-\delta)\mu_{\theta_t}(s)+\delta\nu_{\omega}(s)}.$$

**GAD (pre-conditioner = RMSProp):**

$g \leftarrow \left[ \nabla_{\theta} \widehat{J}(\theta, \omega_t) \right]_{\theta=\theta_t}$ ;  $m \leftarrow \alpha m + (1 - \alpha)g \odot g$ ;  $C \leftarrow \text{diag}(\sqrt{m} + \epsilon)$

$\theta_{t+1} \leftarrow \theta_t + \eta C^{-1}g$

$g' \leftarrow \left[ \nabla_{\omega} \widehat{J}(\theta_t, \omega) \right]_{\omega=\omega_t}$ ;  $m' \leftarrow \alpha m' + (1 - \alpha)g' \odot g'$ ;  $D \leftarrow \text{diag}(\sqrt{m'} + \epsilon)$

$\omega_{t+1} \leftarrow \omega_t - \eta D^{-1}g'$

**Extra-Adam:** use Algorithm 4 from (Gidel et al., 2018).

Update the target networks:

$$\phi_{\text{targ}} \leftarrow \tau \phi_{\text{targ}} + (1 - \tau)\phi$$

$$\theta_{\text{targ}} \leftarrow \tau \theta_{\text{targ}} + (1 - \tau)\theta_{t+1}$$

$$\omega_{\text{targ}} \leftarrow \tau \omega_{\text{targ}} + (1 - \tau)\omega_{t+1}$$

$t \leftarrow t + 1$ .

**end for**

**end if**

**until** convergence

**Output:**  $\omega_T, \theta_T$ .

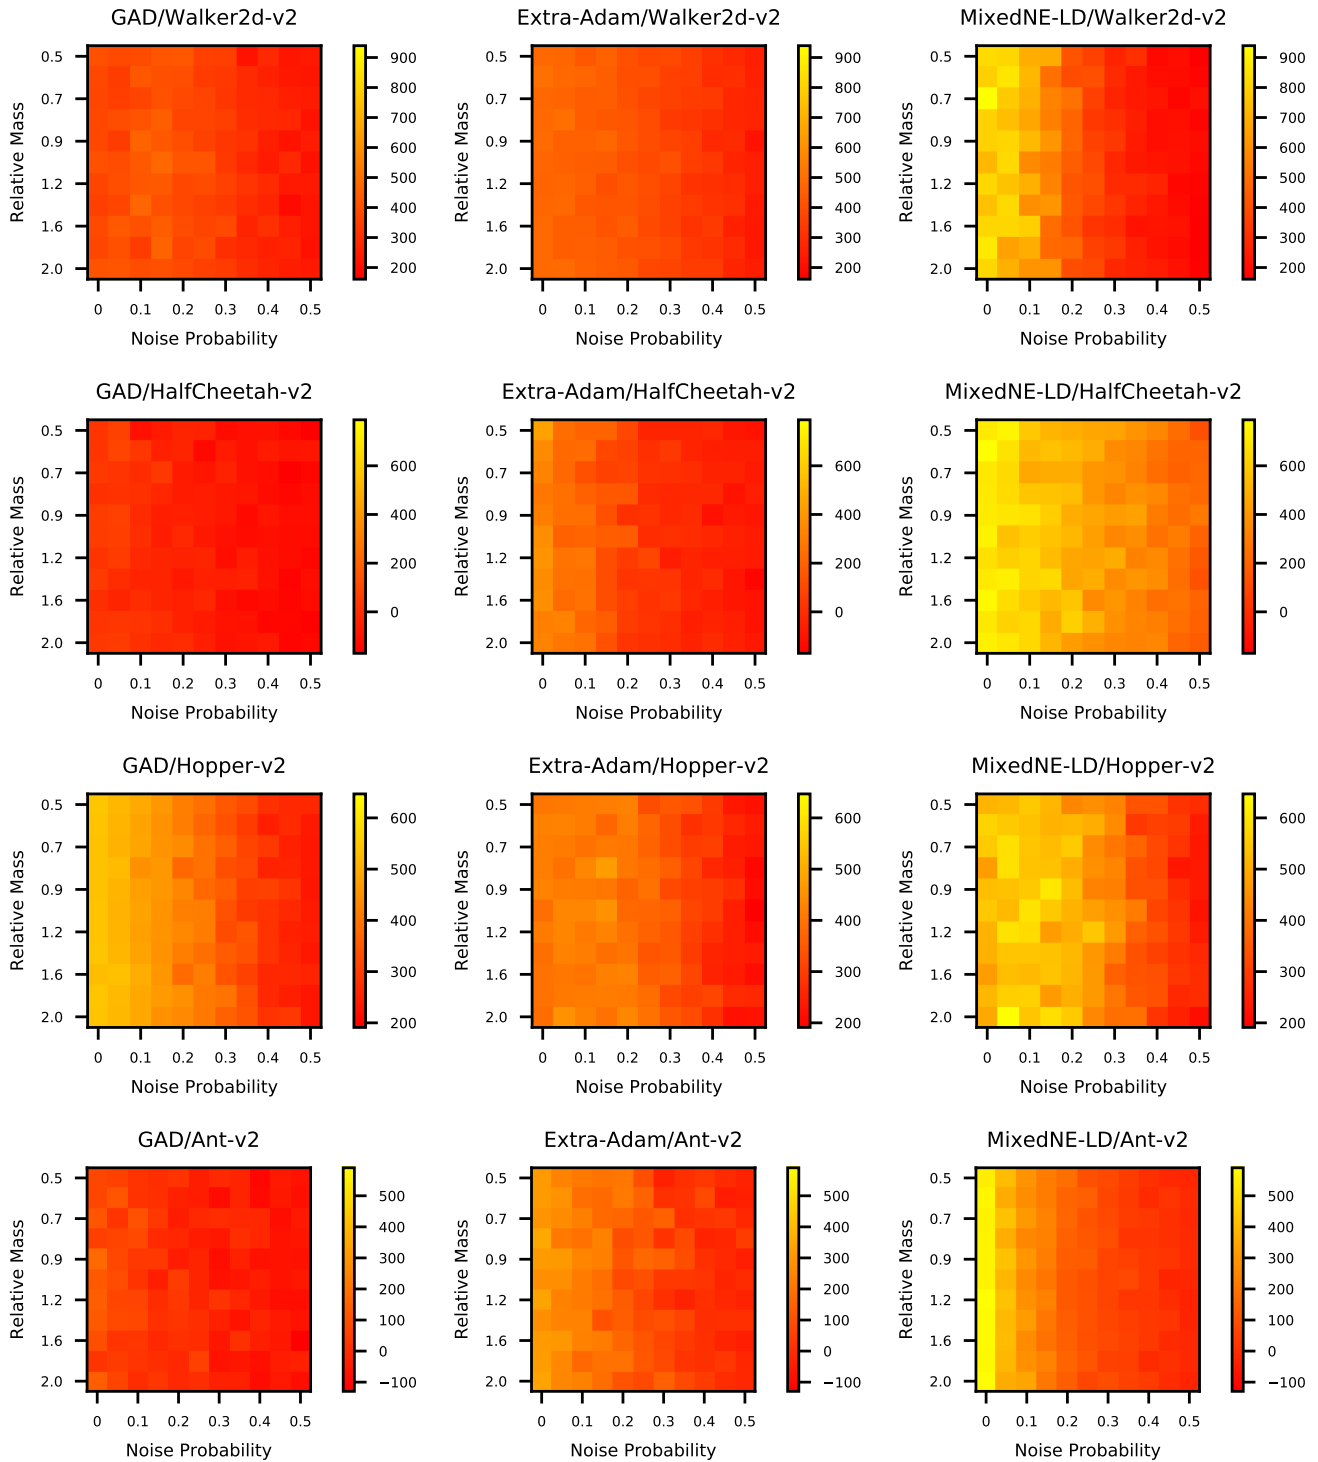


Figure 6. Average performance (over 5 seeds) of Algorithm 3, and Algorithm 4 (with GAD and Extra-Adam), under the NR-MDP setting with  $\delta = 0.1$ . The evaluation is performed on a range of noise probability and mass values not encountered during training. Environments: Walker, HalfCheetah, Hopper, and Ant.






Figure 7. Average performance (over 5 seeds) of Algorithm 3, and Algorithm 4 (with GAD and Extra-Adam), under the NR-MDP setting with  $\delta = 0.1$ . The evaluation is performed on a range of noise probability and mass values not encountered during training. Environments: Swimmer, Reacher, Humanoid, and InvertedPendulum.

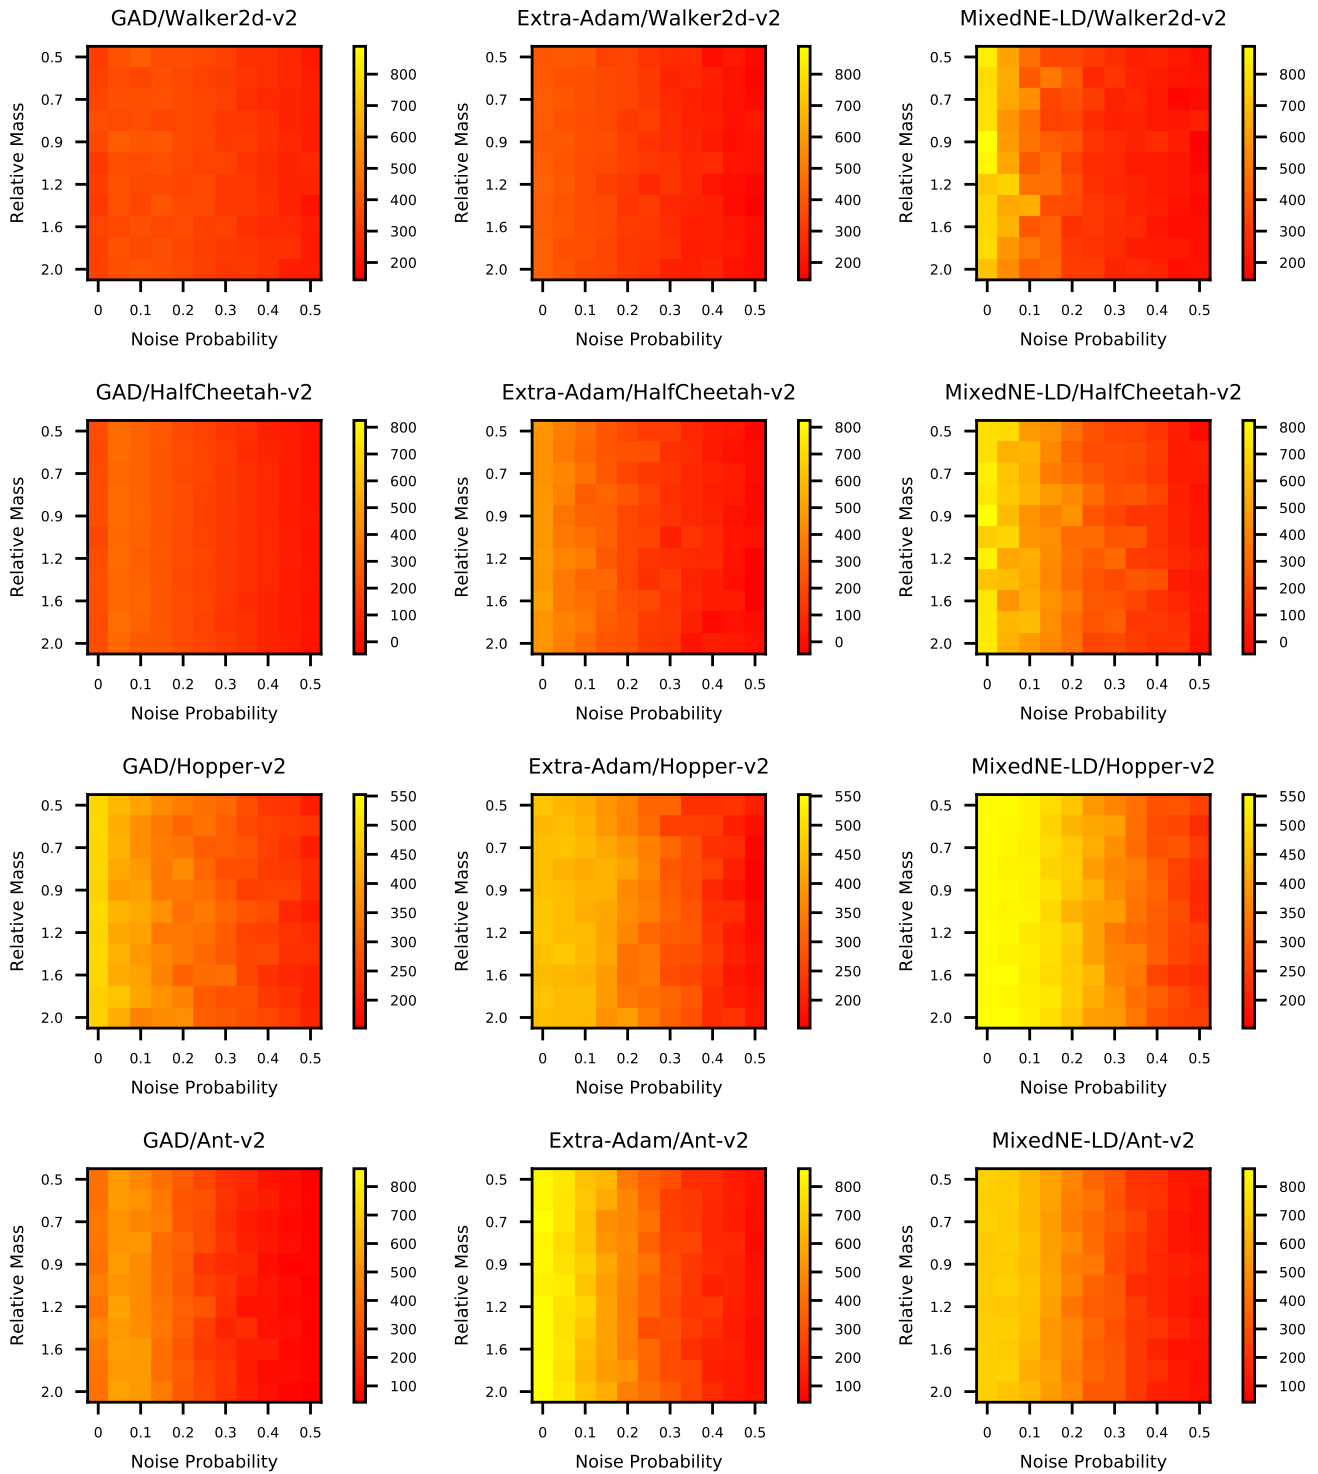


Figure 8. Average performance (over 5 seeds) of Algorithm 3, and Algorithm 4 (with GAD and Extra-Adam), under the NR-MDP setting with  $\delta = 0$ . The evaluation is performed on a range of noise probability and mass values not encountered during training. Environments: Walker, HalfCheetah, Hopper, and Ant.

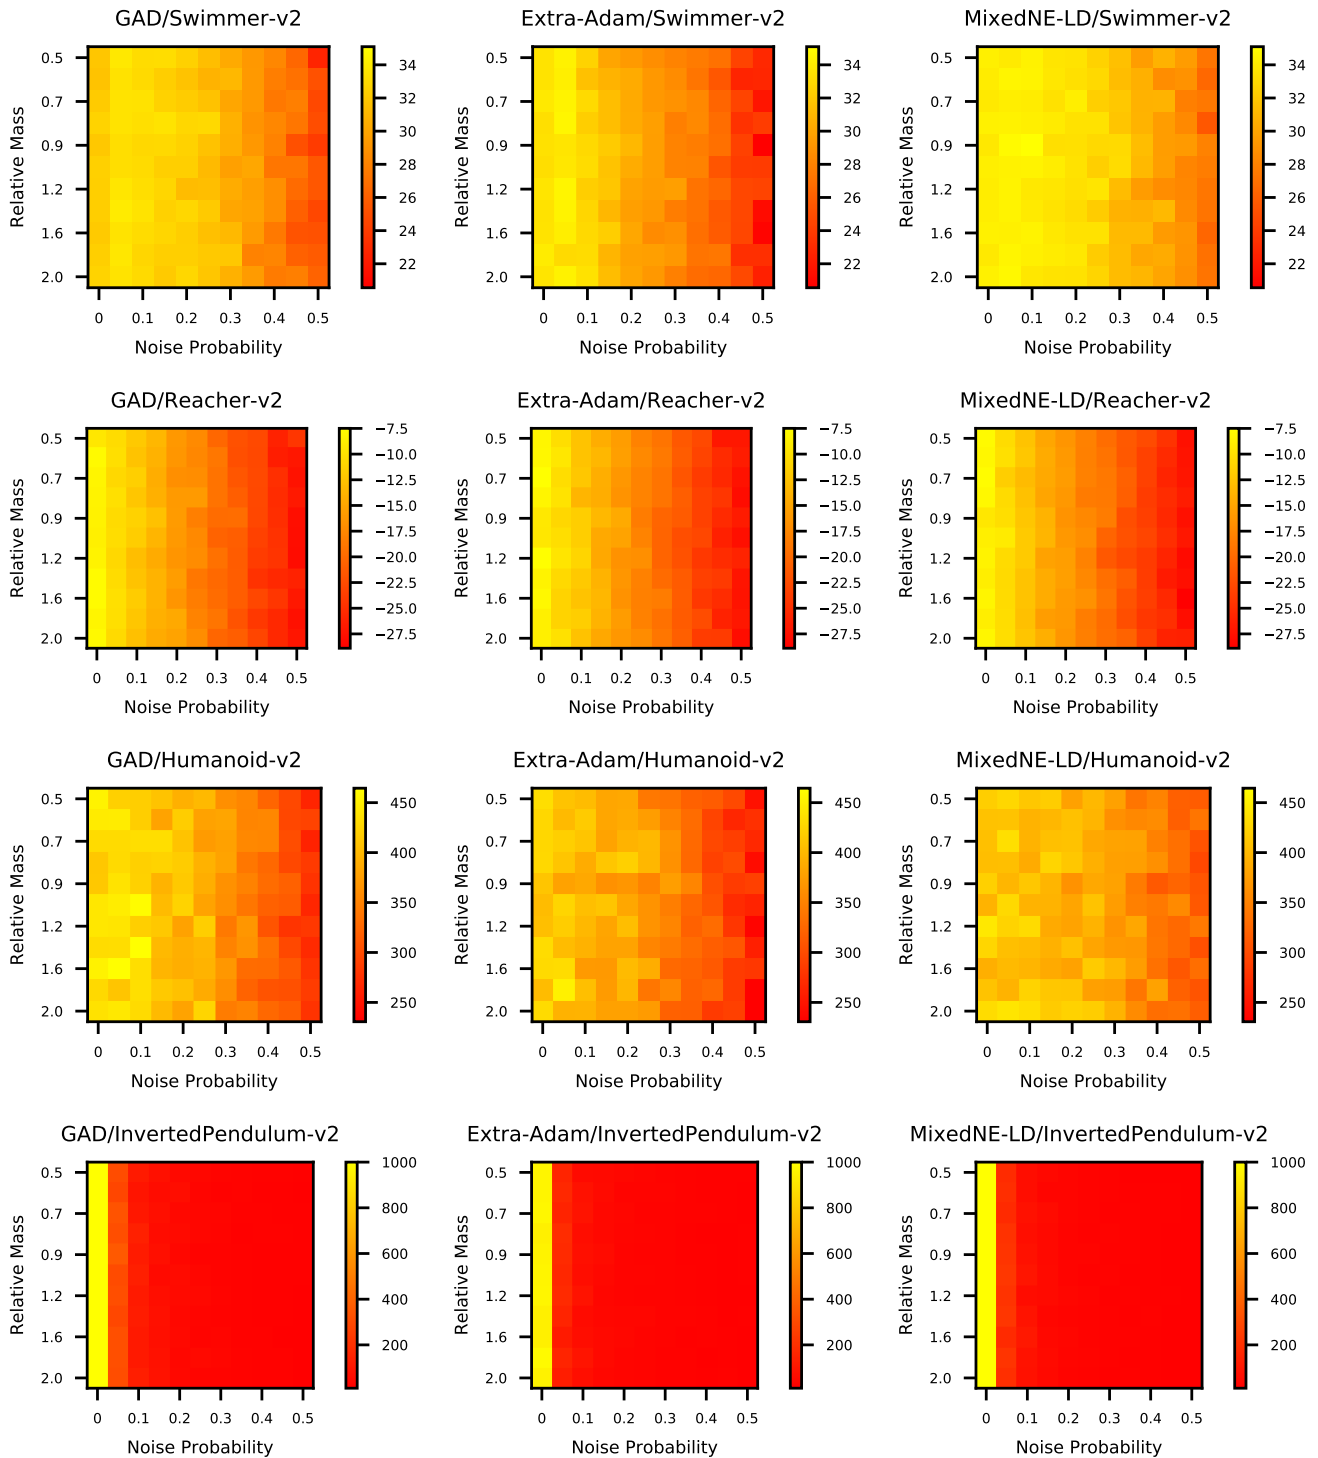


Figure 9. Average performance (over 5 seeds) of Algorithm 3, and Algorithm 4 (with GAD and Extra-Adam), under the NR-MDP setting with  $\delta = 0$ . The evaluation is performed on a range of noise probability and mass values not encountered during training. Environments: Swimmer, Reacher, Humanoid, and InvertedPendulum.

Table 4. Common hyperparameters for Algorithm 5 and Algorithm 6.

Hyperparameter	Value
discount factor $\gamma$	0.99
trajectory length $H$	500
number of trajectories per step $ \mathcal{D}_k $	1
RMSProp parameter $\alpha$	0.99
RMSProp parameter $\epsilon$	$10^{-8}$
learning rate $\eta$	$\{10^{-3}, 10^{-4}, 10^{-5}\}$
damping factor $\beta$	0.9

 Table 5. Additional hyperparameters for Algorithm 5 and Algorithm 6 chosen via grid search (for NR-MDP setting with  $\delta = 0.1$ )

	Algorithm 5: $(\sigma_0, \eta, N_k)$	Algorithm 6 (with GAD): $\eta$	Algorithm 6 (with Extra-Adam): $\eta$
$\rho = 0.2$	$(10^{-5}, 10^{-3}, 1)$	$10^{-4}$	$10^{-4}$

 Table 6. Additional hyperparameters for Algorithm 5 and Algorithm 6 chosen via grid search (for NR-MDP setting with  $\delta = 0$ )

	Algorithm 5: $(\sigma_0, \eta, N_k)$	Algorithm 6 (with GAD): $\eta$	Algorithm 6 (with Extra-Adam): $\eta$
$\rho = 0.2$	$(10^{-4}, 10^{-4}, 10)$	$10^{-4}$	$10^{-3}$

**Algorithm 5** VPG with MixedNE-LD (pre-conditioner = RMSProp)

**Hyperparameters:** see Table 4

Initialize (randomly) policy parameters  $\theta_0, w_0$

**for**  $k = 0, 1, 2, \dots$  **do**

$\bar{\theta}_k, \theta_k^{(0)} \leftarrow \theta_k; \bar{w}_k, w_k^{(0)} \leftarrow w_k$

**for**  $n = 0, 1, \dots, N_k$  **do**

Collect set of trajectories  $\mathcal{D}_k^{(n)} = \{(\dots, s_t^{(\tau)}, \bar{a}_t^{(\tau)}, r_t^{(\tau)}, \dots)\}_\tau$  by running  $\pi_{\theta_k^{(n)}}$ , and  $\pi'_{w_k^{(n)}}$  in  $\mathcal{M}$ , i.e.,  $a_t \sim \pi_{\theta_k^{(n)}}(s_t)$ ,  $a'_t \sim \pi'_{w_k^{(n)}}(s_t)$ ,  $\bar{a}_t = (1 - \delta)a_t + \delta a'_t$ , and  $s_{t+1} \sim T_\rho(\cdot | s_t, \bar{a}_t)$ .

Estimate the policy gradient (where  $G_t = \sum_{s=0}^T \gamma^s r_{t+s}$ )

$$g = \frac{1 - \delta}{|\mathcal{D}_k^{(n)}|} \sum_{\tau \in \mathcal{D}_k^{(n)}} \sum_t \gamma^t G_t^{(\tau)} \left[ \nabla_{\theta} \log \pi_{\theta}(a_t^{(\tau)} | s_t^{(\tau)}) \right]_{\theta = \theta_k^{(n)}}$$

$$g' = \frac{\delta}{|\mathcal{D}_k^{(n)}|} \sum_{\tau \in \mathcal{D}_k^{(n)}} \sum_t \gamma^t G_t^{(\tau)} \left[ \nabla_w \log \pi_w(a'_t^{(\tau)} | s_t^{(\tau)}) \right]_{w = w_k^{(n)}}$$

$m \leftarrow \alpha m + (1 - \alpha) g \odot g$ ;  $C \leftarrow \text{diag}(\sqrt{m} + \epsilon)$

$\theta_k^{(n+1)} \leftarrow \theta_k^{(n)} + \eta C^{-1} g + \sqrt{2\eta\sigma_k} C^{-\frac{1}{2}} \xi$ , where  $\xi \sim \mathcal{N}(0, I)$

$\bar{\theta}_k \leftarrow (1 - \beta) \bar{\theta}_k + \beta \theta_k^{(n+1)}$

$m' \leftarrow \alpha m' + (1 - \alpha) g' \odot g'$ ;  $D \leftarrow \text{diag}(\sqrt{m'} + \epsilon)$

$w_k^{(n+1)} \leftarrow w_k^{(n)} - \eta D^{-1} g + \sqrt{2\eta\sigma_k} D^{-\frac{1}{2}} \xi'$ , where  $\xi' \sim \mathcal{N}(0, I)$

$\bar{w}_k \leftarrow (1 - \beta) \bar{w}_k + \beta w_k^{(n+1)}$

**end for**

$\theta_{k+1} \leftarrow (1 - \beta) \theta_k + \beta \bar{\theta}_k$

$w_{k+1} \leftarrow (1 - \beta) w_k + \beta \bar{w}_k$

**end for**

**Algorithm 6** VPG with GAD (pre-conditioner = RMSProp) / Extra-Adam

**Hyperparameters:** see Table 4

 Initialize (randomly) policy parameters  $\theta_0, w_0$ 
**for**  $k = 0, 1, 2, \dots$  **do**

 Collect set of trajectories  $\mathcal{D}_k = \{(\dots, s_t^{(\tau)}, \bar{a}_t^{(\tau)}, r_t^{(\tau)}, \dots)\}_\tau$  by running  $\pi_{\theta_k}$ , and  $\pi'_{w_k}$  in  $\mathcal{M}$ , i.e.,  $a_t \sim \pi_{\theta_k}(s_t)$ ,  $a'_t \sim \pi'_{w_k}(s_t)$ ,  $\bar{a}_t = (1 - \delta)a_t + \delta a'_t$ , and  $s_{t+1} \sim T_\rho(\cdot | s_t, \bar{a}_t)$ .

 Estimate the policy gradient (where  $G_t = \sum_{s=0}^T \gamma^s r_{t+s}$ )

$$g = \frac{1 - \delta}{|\mathcal{D}_k|} \sum_{\tau \in \mathcal{D}_k} \sum_t \gamma^t G_t^{(\tau)} \left[ \nabla_{\theta} \log \pi_{\theta}(a_t^{(\tau)} | s_t^{(\tau)}) \right]_{\theta=\theta_k}$$

$$g' = \frac{\delta}{|\mathcal{D}_k|} \sum_{\tau \in \mathcal{D}_k} \sum_t \gamma^t G_t^{(\tau)} \left[ \nabla_w \log \pi'_w(a_t^{(\tau)} | s_t^{(\tau)}) \right]_{w=w_k}$$

**GAD (pre-conditioner = RMSProp):**
 $m \leftarrow \alpha m + (1 - \alpha) g \odot g$ ;  $C \leftarrow \text{diag}(\sqrt{m} + \epsilon)$ 
 $\theta_{k+1} \leftarrow \theta_k + \eta C^{-1} g$ 
 $m' \leftarrow \alpha m' + (1 - \alpha) g' \odot g'$ ;  $D \leftarrow \text{diag}(\sqrt{m'} + \epsilon)$ 
 $w_{k+1} \leftarrow w_k - \eta D^{-1} g'$ 
**Extra-Adam:** use Algorithm 4 from (Gidel et al., 2018).

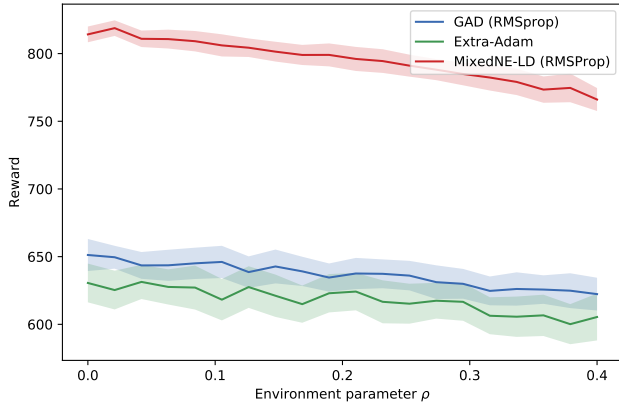
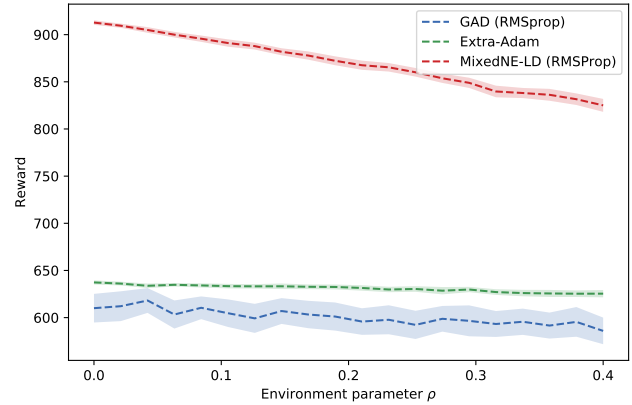
**end for**

 (a)  $\delta = 0.1$ 

 (b)  $\delta = 0$ 

Figure 10. Average performance (over 5 seeds) of Algorithm 5, and Algorithm 6 (with GAD and Extra-Adam), under the NR-MDP setting with  $\delta = 0.1$  and 0 (training on nominal environment  $\rho_0 = 0.2$ ). The evaluation is performed without adversarial perturbations, on a range of environment parameters not encountered during training.






Figure 11. Ablation study: Average performance (over 5 seeds) of MixedNE-LD (with different  $\beta, K_t$ ) under the NR-MDP setting with  $\delta = 0.1$  (training on Half-cheetah with relative mass 1). The evaluation is performed on a range of noise probability and mass values not encountered during training.

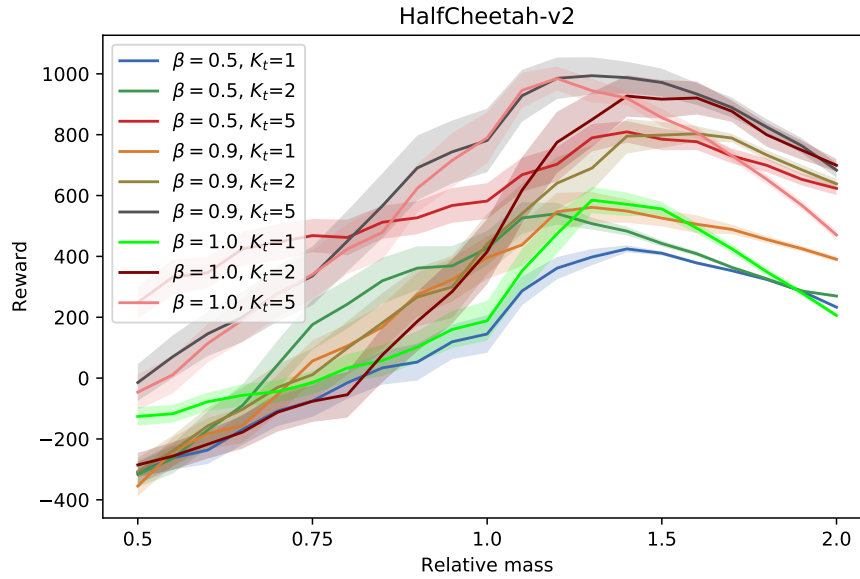


Figure 12. Ablation study: Average performance (over 5 seeds) of MixedNE-LD (with different  $\beta, K_t$ ) under the NR-MDP setting with  $\delta = 0.1$  (training on Half-cheetah with relative mass 1). The evaluation is performed without adversarial perturbations, on a range of mass values not encountered during training.

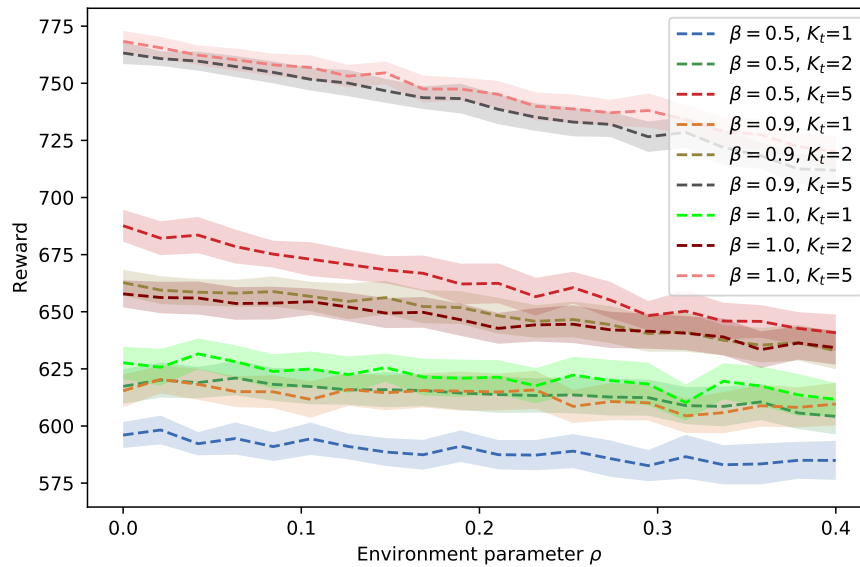


Figure 13. Ablation study: Average performance (over 5 seeds) of MixedNE-LD (with different  $\beta, K_t$ ) under the NR-MDP setting with  $\delta = 0$  (training on nominal environment  $\rho_0 = 0.2$ ). The evaluation is performed without adversarial perturbations, on a range of environment parameters not encountered during training.

## Robust Reinforcement Learning via Adversarial training with Langevin Dynamics

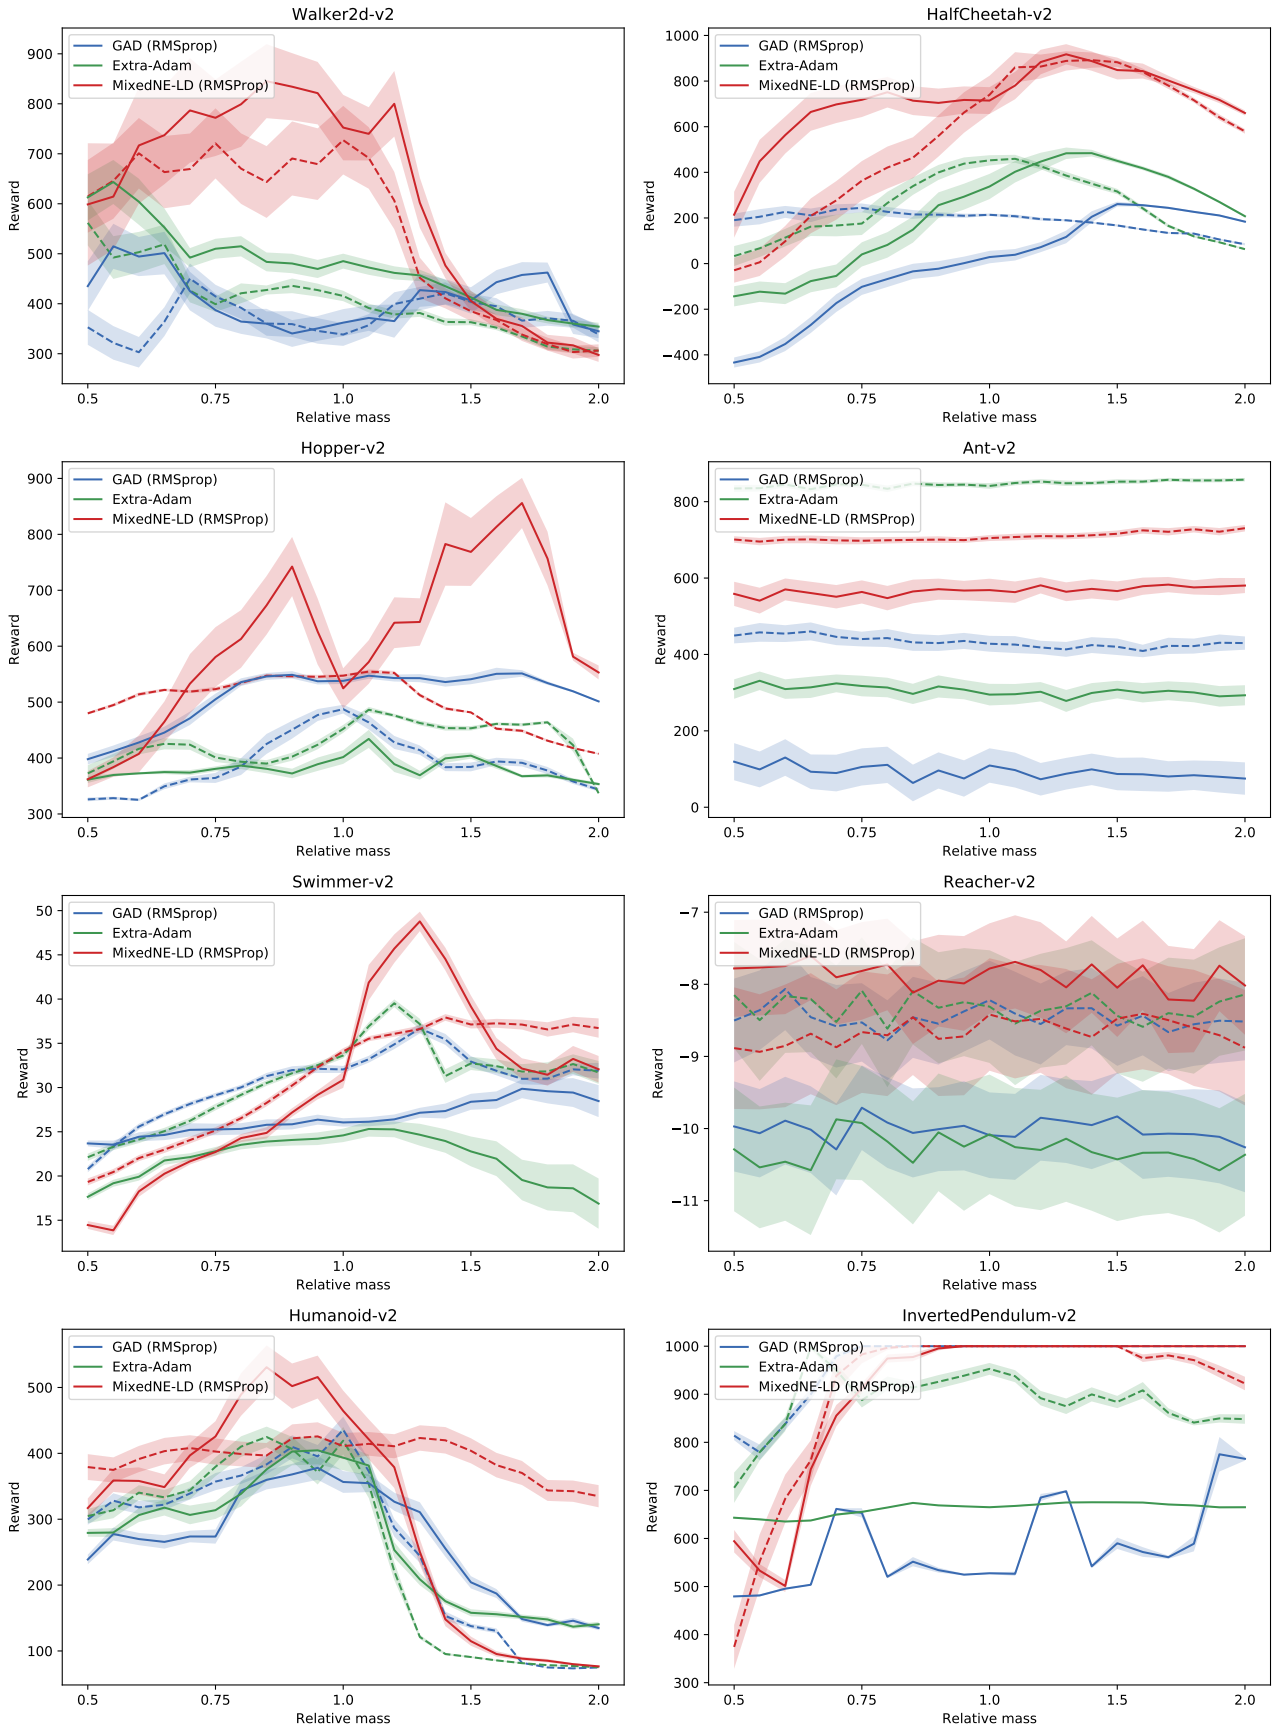


Figure 14. Ablation study: Average performance (over 5 seeds) of Algorithm 3, and Algorithm 4 (with GAD and Extra-Adam), under the NR-MDP setting with  $\delta = 0.1$  (solid lines) and  $\delta = 0$  (dashed lines). The evaluation is performed without adversarial perturbations, on a range of mass values not encountered during training.




Figure 15. Ablation study: Average performance (over 5 seeds) of different algorithms under the NR-MDP setting with  $\delta = 0.1$  (solid lines) and  $\delta = 0$  (dashed lines). The evaluation (after training on the nominal environment  $\rho_0 = 0.2$ ) is performed without adversarial perturbations, on a range of environment parameters not encountered during training.