
Escaping from Zero Gradient: Revisiting Action-Constrained Reinforcement Learning via Frank-Wolfe Policy Optimization

Jyun-Li Lin^{1*}

Wei Hung^{1,2*}

Shang-Hsuan Yang^{1*}

Ping-Chun Hsieh¹

Xi Liu³

¹Department of Computer Science, National Yang Ming Chiao Tung University, Hsinchu, Taiwan

²Research Center for Information Technology Innovation, Academia Sinica, Taipei, Taiwan

³Applied Machine Learning, Facebook AI, Menlo Park, CA, USA

*Equal Contribution

Abstract

Action-constrained reinforcement learning (RL) is a widely-used approach in various real-world applications, such as scheduling in networked systems with resource constraints and control of a robot with kinematic constraints. While the existing projection-based approaches ensure zero constraint violation, they could suffer from the zero-gradient problem due to the tight coupling of the policy gradient and the projection, which results in sample-inefficient training and slow convergence. To tackle this issue, we propose a learning algorithm that decouples the action constraints from the policy parameter update by leveraging state-wise Frank-Wolfe and a regression-based policy update scheme. Moreover, we show that the proposed algorithm enjoys convergence and policy improvement properties in the tabular case as well as generalizes the popular DDPG algorithm for action-constrained RL in the general case. Through experiments, we demonstrate that the proposed algorithm significantly outperforms the benchmark methods on a variety of control tasks.

1 INTRODUCTION

Action-constrained reinforcement learning (RL) is a popular approach for sequential decision making in real-world systems. One classic example is maximizing the network-wide utility by optimally allocating the network resource under capacity constraints (Xu et al., 2018; Gu et al., 2019; Zhang et al., 2020a). Another example is robot control under kinematic constraints (Pham et al., 2018; Gu et al., 2017; Jaillet and Porta, 2012; Tsounis et al., 2020), which capture the limitations of the physical components of a robot (e.g., in terms of velocity, torque, or output power). In these examples, the constraints essentially characterize the set of feasible ac-

tions at each state. To ensure the safe and normal operation of these real-world systems, it is required that these action constraints are satisfied throughout the evaluation as well as the training processes (Chow et al., 2018; Liu et al., 2020a; Gu et al., 2017). Therefore, in action-constrained RL, an effective training algorithm is required to achieve the following two tasks simultaneously: (i) iteratively improving the policy and (ii) ensuring zero constraint violation at each training step.

To enable RL with action constraints, one popular generic approach is to include an additional differentiable projection layer at the output of the policy network and follow the standard end-to-end policy gradient approach (Pham et al., 2018; Dalal et al., 2018; Bhatia et al., 2019). While being a general-purpose solution, this projection layer could result in the *zero-gradient issue* during training due to the tight coupling of the policy gradient update and the projection layer. Specifically, zero gradient occurs when the original output of the policy network falls outside of the feasible action set and any small perturbation of the policy parameters does not lead to any change in the final output action due to the projection mechanism. To better understand the zero-gradient issue, let us consider a toy example of a policy network with one hidden layer and a linear output layer used to produce a deterministic scalar action. Suppose the actions are required to be non-negative. To satisfy the non-negativity action constraint, an additional L_2 -projection layer, which is equivalent to a Rectified Linear Unit (ReLU), is added to the output of the policy. It can be seen that the policy network can easily suffer from zero gradient due to the clipping effect of ReLU (Maas et al., 2013). If the zero-gradient issue occurs in a large portion of the state space, the training process could be sample-inefficient as most of the samples are wasted, and therefore the convergence speed could be slow. Notably, the zero-gradient issue can be particularly severe in the early training phase since the pre-projection actions produced by the policy network are likely to be far away from the feasible sets.

The fundamental cause of the zero-gradient issue is the

tight coupling of the policy parameter update and the projection layer under the standard policy gradient framework. Specifically, in the end-to-end policy-gradient-based training process, the update of the policy parameters relies on the gradient of the actual policy output with respect to the policy parameters and thereby involves the gradient of this additional projection layer. To escape from the zero-gradient issue, we take a different approach and propose a learning algorithm that *decouples* the parameter update for policy improvement from constraint satisfaction, without using the policy gradient theorem. The proposed algorithm can be highlighted as follows:

- To accommodate the action constraints, we leverage the Frank-Wolfe method (Frank et al., 1956) to search for feasible action update directions directly *within* the feasible action sets in a *state-wise* manner. Through this procedure, for a collection of states, we obtain the reference actions that are used to guide the update of the policy parameters for improving the current policy.
- To update the parameters of the policy network, we propose to construct a loss function (e.g., mean squared error) that enables the policy network to adjust its outputs toward the reference actions. This update scheme can be viewed as solving a regression problem based on the reference actions by taking one-step gradient descent. In this way, the parameter update is completely decoupled from the action constraints.

Since the proposed framework obviates the need for the gradient of a projection layer, it avoids the zero-gradient issue by nature.

Our Contributions. In this paper, we revisit the action-constrained RL problem and propose a novel learning framework that avoids the zero-gradient issue and achieves zero constraint violation simultaneously:

- We formally identify the important zero gradient issue in the existing projection-based approaches for action-constrained RL. We also pinpoint that the fundamental cause of the zero-gradient issue is the tight coupling of the policy parameter update and the projection layer under the standard policy gradient framework. To the best of our knowledge, this is the first time that the zero-gradient issue is discussed in the context of action-constrained RL.
- To better describe the proposed learning framework, we start from the case of finite state spaces and introduce Frank-Wolfe policy optimization (FWPO) with tabular policy parameterization, which can be viewed as an instance of the generalized policy iteration. By directly searching for update directions within the feasible sets via state-wise Frank-Wolfe, FWPO automatically achieves zero constraint violation and does not require any additional projection. Moreover, we establish the convergence of FWPO as well as its policy improvement property.
- Built on FWPO, we propose Neural FWPO (NFWPO) by extending the idea of FWPO to the general neural policies

via a regression argument. By constructing a loss function and leveraging state-wise Frank-Wolfe, we decouple the policy parameter update from the action constraints. This design automatically prevents the zero-gradient issue. Moreover, we show that the vanilla DDPG is a special case of NFWPO if there is no action constraints.

- Through experiments on various real applications, we empirically show the zero-gradient problem and demonstrate that the proposed algorithms significantly outperform the popular benchmark methods for action-constrained RL.

2 PRELIMINARIES

We consider an infinite-horizon discounted Markov decision process (MDP) defined by a tuple $(\mathcal{S}, \mathcal{A}, p, r, \gamma)$, where \mathcal{S} is state space, \mathcal{A} denotes the action space, p is the state transition probability, r is the reward function, and $\gamma \in (0, 1)$ denotes the discount factor. We assume that the action space $\mathcal{A} \subseteq \mathbb{R}^N$ is continuous and the reward function takes value in $[0, 1]$ for all state-action pairs. At each time step $t = 0, 1, \dots$, the learner observes state s_t , takes an action a_t , and receives an immediate reward r_t . In this paper, we consider the *action-constrained* MDPs where for each state $s \in \mathcal{S}$ there is a feasible action set $\mathcal{C}(s) \subseteq \mathcal{A}$ determined by the underlying collection of constraints. We assume that $\mathcal{C}(s)$ is compact and convex. In this paper, we focus on deterministic policies and use $\pi(\cdot; \theta) : \mathcal{S} \rightarrow \mathcal{A}$ to denote a deterministic parametric policy with parameter vector $\theta \in \mathbb{R}^n$. Under a policy π , the value functions are defined as the expected long-term rewards

$$V(s; \pi) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \mid s_0 = s, \pi \right], \quad (1)$$

$$Q(s, a; \pi) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \mid s_0 = s, a_0 = a, \pi \right]. \quad (2)$$

To make a comparison between policies, for any two policies π and π' , we say that $\pi \geq \pi'$ if $V(s; \pi) \geq V(s; \pi')$, for all $s \in \mathcal{S}$. This essentially constructs a partial ordering among policies. To construct a total ordering of all the policies, consider the performance objective defined as a weighted average of the value function

$$J_{\mu}(\pi) := \mathbb{E}_{s \sim \mu} [V(s; \pi)], \quad (3)$$

where μ is called the restarting state distribution (Kakade and Langford, 2002). Note that one common choice of μ is the initial state distribution. It is also convenient to define the discounted state visitation distribution d_{μ}^{π} as $d_{\mu}^{\pi}(s) := (1 - \gamma) \mathbb{E}_{s_0 \sim \mu} [\sum_{t=0}^{\infty} \gamma^t P(s_t = s \mid s_0, \pi)]$, for each $s \in \mathcal{S}$.

Notations. We use the standard notations $\|\cdot\|_p$ and $\|\cdot\|_F$ to denote the L_p -norm of a vector and the Frobenius norm of a matrix, respectively. We use $\langle \cdot, \cdot \rangle$ to denote the inner product of two real vectors. For a set \mathcal{D} , we define the diameter of \mathcal{D} as $\text{diam}_{\|\cdot\|_2}(\mathcal{D}) := \sup_{x_1, x_2 \in \mathcal{D}} \|x_1 - x_2\|_2$. We use $\text{dom } f$ to denote the domain of a function f .

2.1 POLICY GRADIENT

To optimize the objective $J_\mu(\pi)$, the typical approach is to apply gradient ascent based on the policy gradient. Under the standard regularity conditions, the deterministic policy gradient (Silver et al., 2014) can be written as

$$\begin{aligned} & \nabla_\theta J_\mu(\pi(\cdot; \theta)) \\ &= \mathbb{E}_{s \sim d_\mu} \left[\nabla_\theta \pi(s; \theta) \nabla_a Q(s, a; \pi(\cdot; \theta)) \Big|_{a=\pi(s; \theta)} \right]. \end{aligned} \quad (4)$$

As a practical implementation of the deterministic policy gradient approach, DDPG (Lillicrap et al., 2016) extends deep Q -learning (Mnih et al., 2015) to continuous action space in an actor-critic manner. Specifically, DDPG updates the policy parameter θ by applying stochastic gradient ascent according to (4) and obtains an approximated Q -function $Q(s, a; \phi)$ parameterized by ϕ by using a Q -learning-like critic, which updates ϕ by minimizing the loss $\mathbb{E}_{(s, a, s', r) \sim \rho} [(r + \gamma Q(s', \pi(s'; \theta^-); \phi^-) - Q(s, a; \phi))^2]$, where ρ denotes the sampling distribution of the replay buffer, θ^- and ϕ^- are the parameters of the actor and critic target networks, respectively.

2.2 FRANK-WOLFE METHODS

In this section, we provide an overview of the Frank-Wolfe algorithms. Consider an optimization problem in the form

$$\max_{x \in \mathcal{X}} F(x), \quad (5)$$

where $F(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}$ is a differentiable function with a Lipschitz continuous gradient, $\mathcal{X} \subseteq \mathbb{R}^d$ is the feasible set characterized by the underlying constraints on x . One popular approach is to apply the projected gradient ascent method (Bubeck et al., 2015), which combines the standard gradient ascent with a projection step. By contrast, as a projection-free method, the classic Frank-Wolfe algorithms (Frank et al., 1956) and its variants solve the constrained optimization problems in (5) by leveraging a first-order subproblem. We briefly summarize the Frank-Wolfe algorithm for non-convex objective functions in the batch settings as follows (Lacoste-Julien, 2016; Reddi et al., 2016):

- **Initialization.** Let x_k denote the input at the k -th iteration and choose an arbitrary $x_0 \in \mathcal{X}$ to be the initial point.
- **Search for an update direction within the feasible set.** In the k -th iteration, compute $v_k = \operatorname{argmax}_{v \in \mathcal{X}} \langle v, \nabla_x F(x) \Big|_{x=x_k} \rangle$ and update the iterate as $x_{k+1} = x_k + \beta_k (v_k - x_k)$, where $v_k - x_k$ is the update direction and β_k denotes the learning rate.

For unconstrained optimization problems, the convergence properties are typically analyzed in terms of the gradient norm $\|\nabla_x F(x)\|_2$. By contrast, for constrained maximization problems, one widely-used metric of convergence in

the Frank-Wolfe literature is the *Frank-Wolfe gap* defined as $\mathcal{G}(x) := \max_{z \in \mathcal{X}} \langle z - x, \nabla_x F(x) \rangle^1$. It is easy to verify that $\mathcal{G}(x) = 0$ is a necessary and sufficient condition of that x is a stationary point.

3 FRANK-WOLFE POLICY OPTIMIZATION

In this section, we formally present the proposed learning algorithms for action constrained RL. To better describe the proposed learning framework, we start from a stylized setting with tabular policy parameterization for finite state spaces and extend the idea to develop a more practical algorithm for the general parametric policies.

3.1 FRANK-WOLFE POLICY OPTIMIZATION WITH DIRECT POLICY PARAMETERIZATION (FWPO)

For ease of exposition, we first illustrate the proposed algorithm for the case of finite state spaces and tabular policies with direct parameterization, i.e., $\pi(s; \theta) \equiv \theta(s)$, for all $s \in \mathcal{S}$. We consider the performance objective $J_\mu(\pi)$ with some restarting state distribution μ with $\mu(s) > 0$, for all $s \in \mathcal{S}$, and define $\mu_{\min} := \min_{s \in \mathcal{S}} \mu(s)$. For ease of notation, we also define $D_s := \operatorname{diam}_{\|\cdot\|_2}(\mathcal{C}(s))$ for each s and $D_{\max} := \max_{s \in \mathcal{S}} D_s$.

Now we present the proposed FWPO algorithm. We use θ_k to denote the policy parameters in the k -th iteration and choose feasible initial policy parameters θ_0 which satisfy $\theta_0(s) \in \mathcal{C}(s)$, for all $s \in \mathcal{S}$. FWPO adopts the generalized policy iteration framework (Sutton and Barto, 2018) by alternating between two subroutines in each iteration:

- **Policy update via state-wise Frank-Wolfe.** FWPO updates the policy by finding a feasible update direction of each state $s \in \mathcal{S}$ via Frank-Wolfe as

$$c_k(s) = \operatorname{argmax}_{c \in \mathcal{C}(s)} \langle c, \nabla_a Q(s, a; \pi(\cdot; \theta_k)) \Big|_{a=\theta_k(s)} \rangle, \quad (6)$$

$$\theta_{k+1}(s) = \theta_k(s) + \alpha_k(s)(c_k(s) - \theta_k(s)), \quad (7)$$

where $c_k(s) - \theta_k(s)$ is the update direction and $\alpha_k(s)$ denotes the (state-dependent) learning rate. Moreover, it is natural to define the *state-wise Frank-Wolfe gap* of the Q -function at θ_k as

$$g_k(s) := \langle c_k(s) - \theta_k(s), \nabla_a Q(s, a; \pi(\cdot; \theta_k)) \Big|_{a=\theta_k(s)} \rangle. \quad (8)$$

¹In the literature, the Frank-Wolfe gap is typically defined as $\max_{z \in \mathcal{X}} \langle z - x, -\nabla_x F(x) \rangle$ since the goal is to minimize an objective function. By contrast, as the goal of RL is to optimize the policy in terms of rewards, we consider the maximization problem in the form of (5) and make the required changes accordingly.

It is easy to verify that $g_k(s) \geq 0$, for all $k \in \mathbb{N}$ and for all $s \in \mathcal{S}$. As will be shown momentarily, to ensure convergence, the learning rate is configured to be $\alpha_k(s) = \frac{(1-\gamma)\mu_{\min}}{LD_s^2} g_k(s)$.

- **Evaluation of the current policy.** FWPO then evaluates the updated policy and obtain the Q -function (or an approximated version) for the next iteration. This can be done by a standard policy evaluation approach.

The above scheme of FWPO is detailed in Algorithm 1. As suggested by Algorithm 1, FWPO always searches for an update direction within the feasible action sets. Therefore, FWPO automatically achieves zero constraint violation and does not require any additional projection by nature.

Algorithm 1 Frank-Wolfe Policy Optimization (FWPO)

- 1: **Input:** Initialize the policy parameters as θ_0 that satisfies $\theta_0(s) \in \mathcal{C}(s)$ for all $s \in \mathcal{S}$
 - 2: **for** each iteration $k = 0, 1, \dots$ **do**
 - 3: Evaluate $\pi(\cdot; \theta_k)$ and obtain $Q(s, a; \pi(\cdot; \theta_k))$
 - 4: **for** each state $s \in \mathcal{S}$ **do**
 - 5: Compute the Frank-Wolfe update direction by $c_k(s) = \operatorname{argmax}_{c \in \mathcal{C}(s)} \langle c, \nabla_a Q(s, a; \pi(\cdot; \theta_k)) \rangle$
 - 6: $g_k(s) = \langle c_k(s) - \theta_k(s), \nabla_a Q(s, a; \pi(\cdot; \theta_k)) \rangle$
 - 7: $\alpha_k(s) = \frac{(1-\gamma)\mu_{\min}}{LD_s^2} g_k(s)$
 - 8: $\theta_{k+1}(s) = \theta_k(s) + \alpha_k(s)(c_k(s) - \theta_k(s))$
 - 9: **end for**
 - 10: **end for**
-

Remark 1 One salient feature of FWPO is that the policy update in (6)-(7) is done by searching for feasible update directions based on $\nabla_a Q(s, a; \pi)$ on a *per-state* basis with state-dependent learning rates, instead of using the standard policy gradient of the performance objective $J_\mu(\pi)$. As will be seen in Section 3.2, this design plays a critical role in decoupling the policy parameter update from constraint satisfaction. Another advantage of FWPO is that it is agnostic to the discounted state visitation distribution d_μ^π (cf. the deterministic policy gradient in (4)) due to the state-wise nature. This feature allows FWPO to be directly applicable in the off-policy settings in its original form².

Remark 2 As the policy update under FWPO is done on a state-by-state basis instead of directly on $J_\mu(\pi)$, the convergence guarantees of the standard Frank-Wolfe methods do not directly apply to the objective $J_\mu(\pi)$ under FWPO. From this perspective, FWPO is *not* a trivial combination of the Frank-Wolfe methods and policy iteration.

As suggested by Remark 2, we proceed to establish the convergence result of FWPO. For the convergence analysis,

²In the off-policy settings, the deep policy gradient approaches typically require dropping a term in the policy gradient expression to accommodate the behavior policy (Silver et al., 2014).

based on the state-wise Frank-Wolfe gaps defined in (8), we define the *effective Frank-Wolfe gap* of $J_\mu(\pi(\cdot; \theta))$ at θ_k as

$$\mathcal{G}_k := \left(\sum_{s \in \mathcal{S}} g_k(s)^2 \right)^{1/2}. \quad (9)$$

Note that $\mathcal{G}_k = 0$ if and only if the update direction is zero for all the states, i.e., $c_k(s) - \theta_k(s) = 0$. Hence, \mathcal{G}_k indicates whether the $J_\mu(\pi(\cdot; \theta))$ converges to a stationary point. We also define $\bar{\mathcal{G}}_T := \min_{0 \leq k \leq T} \mathcal{G}_k$. To establish the convergence results, we also assume mild regularity conditions on r and p as follows.

Definition 1 A differentiable function $f : \operatorname{dom} f \rightarrow \mathbb{R}$ is said to be L_0 -smooth if there exists $L_0 \geq 0$ such that for any $x, y \in \operatorname{dom} f$, $\|\nabla f(x) - \nabla f(y)\|_2 \leq L_0 \|x - y\|_2$.

Regularity Assumptions:

(A1) The reward function $r(s, a)$ is differentiable and is L_r -smooth in a , for all s, a .

(A2) The transition probability $p(s'|s, a)$ is twice differentiable and L_p -smooth in a , for all s, s', a . Moreover, $p(s'|s, a)$ satisfies $\sup_{s, a, s'} \|\nabla_a p(s'|s, a)\|_2 < C_p$.

As the first step, we introduce the following proposition on the smoothness of the performance objective $J_\mu(\pi(\cdot; \theta))$. Notably, given the regularity assumptions of r and p in action, it remains non-trivial to establish the smoothness of $J_\mu(\pi(\cdot; \theta))$ in θ due to the multi-step compound effect of the changes in policy parameters on the value functions.

Proposition 1 Under the regularity assumptions (A1)-(A2), there exists some constant $L > 0$ such that for any restarting state distribution μ , $J_\mu(\pi(\cdot; \theta))$ is L -smooth in θ .

The proof of Proposition 1 is provided in Appendix A.1. Now we are ready to present the convergence result.

Proposition 2 Under the FWPO algorithm with $\alpha_k(s) = \frac{(1-\gamma)\mu_{\min}}{LD_s^2} g_k(s)$, $\{\pi(\cdot; \theta_k)\}$ form a non-decreasing sequence of policies in the sense that $\pi(\cdot; \theta_{k+1}) \geq \pi(\cdot; \theta_k)$, for all k . Moreover, the effective Frank-Wolfe gap of FWPO converges to zero as $k \rightarrow \infty$, and the convergence rate can be quantified as

$$\sum_{k=0}^{\infty} \mathcal{G}_k^2 \leq \frac{2LD_{\max}^2}{(1-\gamma)^3 \mu_{\min}^2}, \quad (10)$$

which implies that $\bar{\mathcal{G}}_T = O(T^{-1/2})$.

Proof Due to space limitation, we provide a sketch of proof: (i) To show the non-decreasing property, we leverage the policy difference lemma (Kakade and Langford, 2002) and verify a sufficient condition of strict policy improvement; (ii) To show the convergence result, we leverage the smoothness of the value functions as well as the objective and use the technique for convergence of non-convex optimization similar to that in (Reddi et al., 2016; Lacoste-Julien, 2016);

(iii) A proper learning rate can be selected by taking the smoothness conditions as well as the restarting state distribution into account. For completeness, the detailed proof is provided in Appendix A.2.

Remark 3 The style of the convergence guarantee in Proposition 2 is common in the analysis of gradient descent methods for non-convex smooth functions (Bottou et al., 2018). Moreover, the result (i.e., convergence to a stationary point) in Proposition 2 resembles those of the policy gradient algorithms (Sutton, 2000; Silver et al., 2014), but for the action-constrained RL settings. On the other hand, in (6), the search of the update direction requires the gradient of the Q -function. In practice, it may not be feasible to obtain the whole true Q -function, and a value function approximator can be included. In practice, it can be expected that a sufficiently accurate critic shall provide a sufficiently good update direction.

3.2 NEURAL FRANK-WOLFE POLICY OPTIMIZATION (NFWPO)

In this section, we formally present the proposed NFWPO algorithm for general parametric policies for action-constrained RL. As highlighted in Section 1, we propose to decouple constraint satisfaction from the policy parameter update. Specifically, to accommodate the action constraints, we extend the state-wise Frank-Wolfe subroutine to the general parametric policies. One inherent challenge of such extension is that the Frank-Wolfe method searches for an update direction within the feasible set by nature. However, under neural parameterization, an action produced by the neural network is not guaranteed to stay in the feasible action set. To address this, we propose to incorporate a projection step into the state-wise Frank-Wolfe subroutine. Define a projection operator as

$$\Pi_{C(s)}(z) = \operatorname{argmin}_{y \in C(s)} \|y - z\|_2. \quad (11)$$

For ease of exposition, in the sequel we call the input z a *pre-projection action* and $\Pi_{C(s)}(z)$ a *post-projection action*.

NFWPO adopts the actor-critic architecture. Let $\bar{\theta}$ and $\bar{\phi}$ be the current parameters of the actor and the critic, respectively. The main features of NFWPO are captured by the actor part as below.

- **Derive reference actions via state-wise Frank-Wolfe.** For each s in the mini-batch \mathcal{B} , NFWPO uses Frank-Wolfe to compute the reference action at each state s as

$$\tilde{a}_s = \Pi_{C(s)}(\pi(s; \bar{\theta})) + \alpha(\bar{c}(s) - \Pi_{C(s)}(\pi(s; \bar{\theta}))), \quad (12)$$

where α is the learning rate of Frank-Wolfe and

$$\bar{c}(s) = \operatorname{argmax}_{c \in C(s)} \langle c, \nabla_a Q(s, a; \bar{\phi}) \rangle|_{a=\Pi_{C(s)}(\pi(s; \bar{\theta}))}. \quad (13)$$

(Note that the projection $\Pi_{C(s)}(\cdot)$ is only for generating feasible actions and does not require backpropagation.)

- **Construct an MSE loss function.** NFWPO constructs a loss function $\mathcal{L}_{\text{NFWPO}}(\theta; \bar{\theta})$ as the MSE between the actions of the current policy and the reference actions, i.e.,

$$\mathcal{L}_{\text{NFWPO}}(\theta; \bar{\theta}) = \sum_{s \in \mathcal{B}} (\pi(s; \theta) - \tilde{a}_s)^2. \quad (14)$$

- **Update policy by gradient descent.** NFWPO updates the policy parameter by minimizing the MSE loss in (14) by using gradient descent for one step, i.e.,

$$\theta \leftarrow \theta - \beta \nabla_{\theta} \mathcal{L}_{\text{NFWPO}}(\theta; \bar{\theta}). \quad (15)$$

On the other hand, the critic of NFWPO can be based on any standard policy evaluation technique. For ease of exposition, for NFWPO, we use the same critic as the vanilla DDPG (as described in Section 2.1). The detailed pseudo code of NFWPO is provided in the supplementary material.

Notably, similar to (6), NFWPO only uses $\nabla_a Q(s, a; \bar{\phi})$ for deriving reference actions, *without* using the deterministic policy gradient in (4). This design allows NFWPO to decouple constraint satisfaction in (12)-(13) from the parameter update in (14)-(15). As highlighted in Section 1, this decoupling obviates the need for the gradient of a projection layer and hence automatically avoids the zero-gradient issue. Moreover, below we show that DDPG is actually a special case of NFWPO when there is no action constraints. The proof is provided in Appendix B

Proposition 3 *If there is no action constraints, then the policy update scheme of NFWPO in (12)-(15) is equivalent to the vanilla DDPG by (Lillicrap et al., 2016).*

Remark 4 While NFWPO leverages a projection step in (12), this projection step is only for deriving reference actions and does not take part in the policy parameter update. As a result, NFWPO does not require backpropagation of the projection step (as shown in (12)-(15)) and therefore automatically avoids the zero-gradient issue. Hence, NFWPO is essentially different from the existing solutions that combine DDPG with a projection layer for end-to-end training (Pham et al., 2018; Dalal et al., 2018).

4 EXPERIMENTAL RESULTS

In this section, we empirically evaluate FWPO and NFWPO in various real-world applications, including bike sharing systems, communication networks, and continuous control in MuJoCo. We compare the proposed algorithms against the following popular benchmark methods:

- **DDPG+Projection:** The training procedure is identical to the vanilla DDPG (Lillicrap et al., 2016) except that the action is post-processed by the L_2 -projection operator $\Pi_{C(s)}(\cdot)$ before being applied to the environment.

- **DDPG+RewardShaping**: Built on DDPG+Projection, this algorithm adds the L_2 -norm between the pre-projection and post-projection actions as a penalty to the intrinsic reward.
- **DDPG+OptLayer**: This design uses a differentiable projection layer, namely the OptLayer, that supports end-to-end training via gradient descent (Pham et al., 2018).

Moreover, for the projection step (without the need of backpropagation) required by DDPG+Projection, DDPG+RewardShaping, and NFWPO, we implement this functionality on the Gurobi optimization solver (Gurobi Optimization, 2021). Therefore, the post-projection actions are guaranteed to satisfy the action constraints for all the algorithms. For each task, each algorithm is trained under the common set of 5 random seeds. Each evaluation consists of 10 episodes, and we report the average performance along with the standard deviation in Figures 1-5. We also summarize the average return over the final 10 evaluations in Table 1 and Table 2. The detailed training setup can be found in Appendix D. The code of our experiments is available ³.

4.1 BIKE SHARING SYSTEMS

We use the open-source BSS simulator⁴, which was originally proposed by (Ghosh and Varakantham, 2017) and later used for evaluating action-constrained RL by (Bhatia et al., 2019). In a bike-sharing problem, there are m bikes and n stations, each of which has a pre-determined bike storage capacity C . An action is to allocate m bikes to n stations under random demands. The reward signal consists of three parts: (i) Moving cost: the cost of moving one or multiple bikes from one station to another; (ii) Lost-demand cost: the cost of unserved demand due to bike outage. (iii) Overflow cost: the cost incurred when the number of bikes in one station exceeds its capacity.

Evaluating FWPO. Since the bike sharing environment has a finite state space, we first use it to evaluate FWPO against the baseline methods, all with tabular policy parameterization. For the action value function, we use the same Q -learning-like critic as the vanilla DDPG for all the algorithms. A medium-sized system with $n = 3$, $m = 90$, and $C = 35$ is chosen that allows to analytically find the optimal policy. There are two types of constraints: (i) Global constraint: all the action entries shall sum to 90; (ii) Local constraints: each entry of the action shall be between 0 and 35. Figure 1(a) shows the average return of the three algorithms. We observe that FWPO performs the best, while DDPG+Projection and DDPG+RewardShaping both suffer from slow learning. This is also reflected by Figure 1(b), which shows that FWPO converges to a near-optimal policy much faster than the baselines. The above phenomenon is mainly due to the inaccurate policy gradient of DDPG under

action constraints. Specifically, the critics of the two baselines are trained with samples with feasible actions while the gradients $\nabla_a Q(s, a; \phi)$ are mostly evaluated at those actions outside the feasible sets. By contrast, FWPO always stays in the feasible action sets and hence naturally avoids the issue of inaccurate gradients.

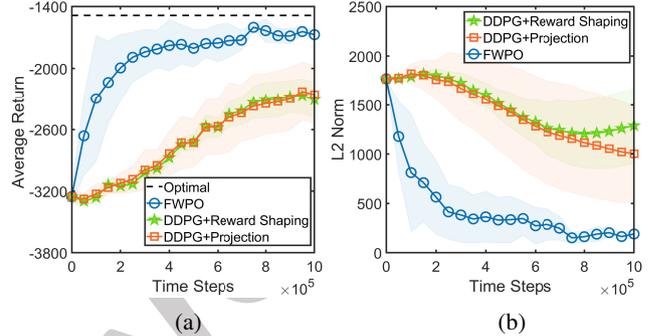


Figure 1: Bike sharing problem with $n = 3$ (BSS-3) under tabular policies: (a) Average return over 5 random seeds; (b) L_2 -norm between the learned policies and the optimal policy at each training step.

Evaluating NFWPO. We proceed to compare NFWPO with the other three baselines in solving a larger-scale bike-sharing problem with $m = 150$, $n = 5$, and $C = 35$. As shown by Figure 2(a), NFWPO converges faster and achieves a larger return than the other baselines. To better understand its behavior, Figure 2(b) shows the cumulative constraint violations of the *pre-projection* actions. Interestingly, the pre-projection actions of NFWPO can largely avoid constraint violation, and thus requires less help from the projection during training. By contrast, all the baselines rely heavily on the projection step to stay feasible, because most of their pre-projection actions fail to satisfy the constraints. We also observe that DDPG+Projection and DDPG+RewardShaping attain similar average return and frequency of violation. This is because they both produce pre-projection actions far from the feasible sets and thereby obtain similar post-projection actions. Meanwhile, DDPG+OptLayer suffers from nearly zero learning progress due to the zero-gradient issue. Figure 2(c) compares the sample-based gradient with respect to the pre and post-OptLayer actions. Since the gradients of the pre-OptLayer actions (green line in Figure 2(c)) are mostly close to zero, the sample-based policy gradients $\nabla_{\theta} \hat{J}_{\mu}(\pi(\cdot; \theta))$ are therefore close to zero for most of the training steps. As the gradients with respect to the post-OptLayer actions are always *non-zero* (blue dotted line in Figure 2(c)), we know that the zero-gradient issue of $\nabla_{\theta} \hat{J}_{\mu}(\pi(\cdot; \theta))$ indeed results from the projection layer. This confirms that the additional OptLayer could easily lead to the zero-gradient issue and sample-inefficient training.

³https://github.com/upupsheep/NFWPO_Final_Code

⁴BSS: <https://github.com/bhatiaabhinav/gym-BSS>

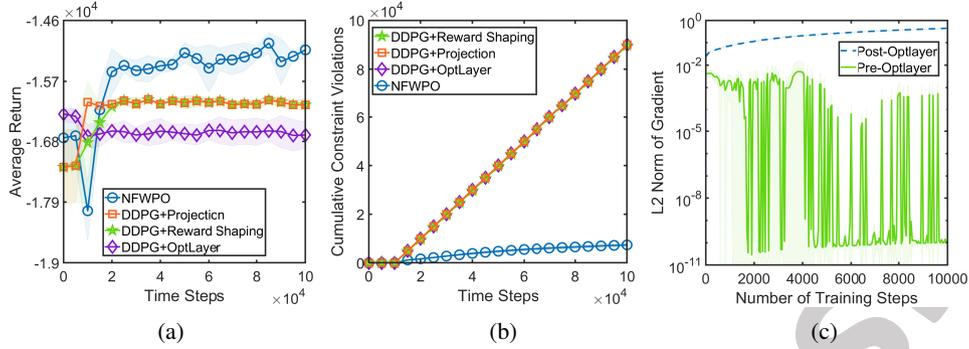


Figure 2: Bike sharing problem with $n = 5$ (BSS-5) under neural policies: (a) Average return over 5 random seeds; (b) Cumulative number of constraint violations of the pre-projection actions during training; (c) L_2 -norm of the sample-based gradients with respect to the pre- and post-OptLayer actions of DDPG+OptLayer.

4.2 UTILITY MAXIMIZATION OF COMMUNICATION NETWORKS

In this section, we evaluate the proposed methods over the task of utility maximization in communication networks. We simulate the network with the open-source network simulator from PCC-RL⁵ (Jay et al., 2019). For the network topology, we consider the classic T3 NSFNET Backbone and set the bandwidth of each link to be 50 packets per second throughout the experiments. We generate three network flows, each of which has three candidate paths from its source to the destination. The action is to determine the rate allocation of each flow along each candidate path. The reward consists of three parts: (i) Throughput: the number of received packets per second; (ii) Drop rate: the number of dropped packets per second; (iii) Latency: the average latency of the packets in the last second. For each flow i , its immediate reward is $\log\left(\frac{\text{throughput}_i}{\text{drop rate}_i^{0.5} \times \text{latency}_i^{1.5}}\right)$, which corresponds to the widely-used proportional fairness criteria (Kelly, 1997). One salient feature of a communication network is that when the total packet arrival rate of a link approaches its bandwidth, the latency will grow rapidly, and accordingly most of the packets would be dropped. Therefore, in this environment, the action constraints correspond to the link bandwidth constraints, i.e., the total assigned packet arrival rate of each link should be bounded by 50.

Figure 3(a) shows the training curves and indicates that NFWPO still converges fast (in about 10^5 steps) and achieves much larger return than the baselines. Moreover, similar to the bike-sharing problems, we see from Figure 3(b) that most of the pre-projection actions of NFWPO already satisfy the constraints. In this task, we find that reward shaping does help in guiding the pre-projection actions towards the feasible action sets, but only under some random seeds (and therefore the large variance in Figures 3(a)-(b)). Regarding DDPG+OptLayer, in the initial training phase,

we observe that it mostly produces pre-OptLayer actions with small flow rates, which lead to a smaller number of constraint violations and moderate returns. To achieve a higher return, DDPG+OptLayer then gradually increases the flow rates but accidentally causes more constraint violations of pre-OptLayer actions and suffers from the inaccurate gradient issue described in Section 4.1. Ultimately, DDPG+OptLayer can only achieve a fairly low return.

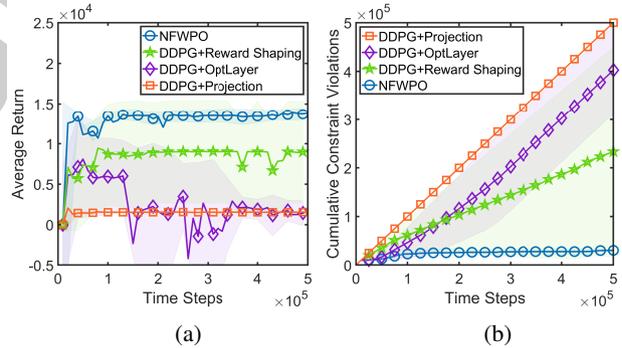


Figure 3: Utility maximization in NSFNET: (a) Average return over 5 random seeds; (b) Cumulative constraint violations of the pre-projection actions during training.

4.3 MUJOCO CONTINUOUS CONTROL TASKS

To further validate NFWPO, we consider popular continuous control tasks in MuJoCo (Todorov et al., 2012) with non-linear and state-dependent action constraints. We further compare the proposed algorithm with important benchmarks RL algorithms for MuJoCo control tasks, including PPO (Schulman et al., 2017), TRPO (Schulman et al., 2015), and SAC (Haarnoja et al., 2018). To make the comparison even more comprehensive, we also evaluate FOCOPS (Zhang et al., 2020b), which is a recent approach designed to address long-term total discounted cost constraints. As

⁵PCC-RL: <https://github.com/PCCproject/PCC-RL>

Table 1: Average return over the final 10 evaluations.

Methods	BSS-3	BSS-5	NSFNET
NFWPO	-1673.04	-15132.21	13770.67
DDPG+Projection	-2254.52	-16123.48	1514.44
DDPG+Reward Shaping	-2308.00	-16123.48	9010.46
DDPG+OptLayer	-	-16686.04	1667.59

Table 2: Average return over the final 10 evaluations in the MuJoCo environments.

Methods	Reacher	Halfcheetah
NFWPO	-4.76	6513.26
DDPG+Projection	-11.15	2746.72
DDPG+Reward Shaping	-8.66	3065.37
DDPG+OptLayer	-7.25	1399.37
SAC+Projection	-10.50	4874.45
TRPO+Projection	-11.04	2247.82
PPO+Projection	-10.68	1459.04
FOCOPS+Projection	-12.23	1916.46

FOCOPS is not designed for handling state-wise action constraints, we relax the action constraints into the long-term total discounted constraints required by FOCOPS. To ensure constraint satisfaction, we use the same technique as DDPG+Projection, i.e., the actions of PPO, TRPO, SAC, and FOCOPS are post-processed by L_2 -projection before being applied to the environment.

Reacher with non-linear constraints. In this task, the action space is 2-dimensional (denoted by u_1, u_2), and each action entry corresponds to the torque of a joint of a 2-DoF robot. To validate the applicability of NFWPO, we impose nonlinear action constraints as: $u_1^2 + u_2^2 \leq 0.05$. From Figure 4(a), we observe a similar trend that NFWPO still converges faster and achieves a larger return than the other baseline methods. In this task, DDPG+OptLayer and DDPG+RewardShaping can achieve a return closer to NFWPO as it has fewer pre-OptLayer (pre-projection) violations as shown in Figure 4(b). The other algorithms still perform poorly as they always produce actions far from the feasible sets and relies heavily on the projection step. This is reflected by the fact that SAC, TRPO, and PPO violate the constraint at almost every training step, as shown in Figure 4(b). Regarding FOCOPS, the violation of the constraint in the second half of training is less than that in the first half. This indicates that FOCOPS needs much more training steps to find a policy that violate lesser, Despite this, the return of FOCOPS remains fairly low as other benchmarks.

Halfcheetah with state-dependent constraints. In this task, an action is 6-dimensional and is denoted by (v_1, \dots, v_6) . We consider a challenging scenario where the constraint is state-dependent. Specifically, the imposed constraint is $\sum_{i=1}^6 |v_i w_i| \leq 20$, where w_i denotes the angular velocity of the i -th joint and is part of the state. This constraint is meant to capture the limitation of total output

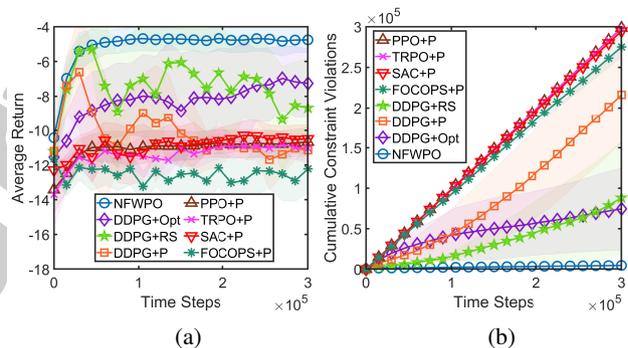


Figure 4: Reacher-v2 with a non-linear action constraint: (a) Average return over 5 random seeds; (b) Cumulative constraint violations of the pre-projection actions during training (“+P” stands for “+Projection”, “+RS” stands for “Reward Shaping”, and “+Opt” stands for “+OptLayer”).

power. Similar to the other environments, from Figure 5(a)-(b), we still observe that NFWPO achieves better sample efficiency than the other baselines. Moreover, NFWPO violates the constraint for only 3% of the steps while getting the highest return. On the other hand, PPO, TRPO, and SAC all violate the constraint for more than 75% of the steps. FOCOPS violates the constraint for about 15% of the time but only achieves a fairly low return. Again, we see that NFWPO outperforms all the baseline methods with much less constraint violation.

5 RELATED WORK

The constrained RL problems have been extensively studied from two main perspectives. The first category encodes the constraints via cost signals, which are incurred at each step

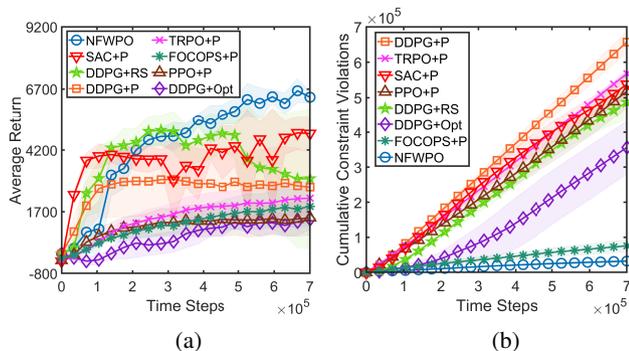


Figure 5: Halfcheetah-v2 with a state-dependent action constraint: (a) Average return over 5 random seeds; (b) Cumulative constraint violations of the pre-projection actions during training (“+P” stands for “+Projection”, “+RS” stands for “Reward Shaping”, and “+Opt” stands for “+OptLayer”).

along with the reward signals, and accordingly focuses on the average-cost constraints. This line of research works borrows a variety of ideas from the optimization literature. For example, (Chow et al., 2015) addressed chance constraints by using a primal-dual approach to achieve a trade-off between return and risk. Similarly, (Tessler et al., 2018) proposed Reward Constrained Policy Optimization which applied Lagrangian relaxation and converted the constraints into penalty terms in the objective. (Achiam et al., 2017) proposed Constrained Policy Optimization to achieve strict policy improvement under the average cost constraints by using the trust-region approach. In (Chow et al., 2018), Lyapunov-based safe reinforcement learning was proposed to address the constraints by solving a linear program. (Yang et al., 2019) proposed Projection-Based Constrained Policy Optimization to achieve no constraint violation by taking a projection step after the local reward improvement update. In (Liu et al., 2020b), Interior-point Policy Optimization was proposed to handle the average cost constraints by augmenting the objective with logarithmic barrier functions. (Satija et al., 2020) took a different approach by converting the trajectory-level constraints into per-step state-wise constraints and accordingly defining a safe policy improvement step. (Zhang et al., 2020b) proposed an approach to solving the constrained policy optimization problem by first finding a target policy directly in the policy space and thereafter converting the target policy to a parameterized one through a projection step onto the parameter space. Different from all the above prior works, this paper considers the RL settings with state-wise action constraints.

The second category is on the state-wise constraints that need to be satisfied on a step-by-step basis. (Pham et al., 2018) studied the state-wise action constraints of robotic systems and proposed a projection-based OptLayer to enforce the constraints. (Dalal et al., 2018) also considered

state-wise safety constraints under linearization and proposed a projection-based safety layer to handle the constraints. Similarly, (Bhatia et al., 2019) considered resource constraints and proposed variants of OptLayer to improve the computational efficiency. (Shah et al., 2020) proposed a more efficient projection scheme for enforcing linear action constraints. Despite the similarity in the problem setting, we take a different approach and propose a decoupling framework by leveraging Frank-Wolfe to address the action constraints and completely avoid the zero-gradient issue.

6 CONCLUSION

This paper revisits action-constrained RL to tackle the zero-gradient issue and ensure zero constraint violation simultaneously. We achieve this goal by developing a learning framework that decouples the policy parameter update from constraint satisfaction by leveraging state-wise Frank-Wolfe and a regression argument. Our theoretical and experimental results demonstrate that the proposed learning algorithm is indeed a promising approach for action-constrained RL.

Acknowledgements

This material is based upon work partially supported by the Ministry of Science and Technology of Taiwan under Contract No. MOST 108-2636-E-009-014 and Contract No. MOST 109-2636-E-009-012.

References

- Joshua Achiam. Spinning Up in Deep Reinforcement Learning. 2018.
- Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization. In *International Conference on Machine Learning*, pages 22–31. PMLR, 2017.
- Alekh Agarwal, Sham M Kakade, Jason D Lee, and Gaurav Mahajan. On the theory of policy gradient methods: Optimality, approximation, and distribution shift. *arXiv:1908.00261*, 2019.
- Alekh Agarwal, Sham M Kakade, Jason D Lee, and Gaurav Mahajan. Optimality and approximation with policy gradient methods in Markov decision processes. In *Conference on Learning Theory*, pages 64–66, 2020.
- Akshay Agrawal, Brandon Amos, Shane Barratt, Stephen Boyd, Steven Diamond, and Zico Kolter. Differentiable convex optimization layers. In *Advances in Neural Information Processing Systems*, 2019.
- Brandon Amos and J Zico Kolter. Optnet: Differentiable optimization as a layer in neural networks. In *International Conference on Machine Learning*, pages 136–145. PMLR, 2017.

- Abhinav Bhatia, Pradeep Varakantham, and Akshat Kumar. Resource constrained deep reinforcement learning. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 29, pages 610–620, 2019.
- Léon Bottou, Frank E Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *Siam Review*, 60(2):223–311, 2018.
- Sébastien Bubeck et al. Convex Optimization: Algorithms and Complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.
- Yinlam Chow, Aviv Tamar, Shie Mannor, and Marco Pavone. Risk-sensitive and robust decision-making: a CVaR optimization approach. In *Advances in Neural Information Processing Systems*, pages 1522–1530, 2015.
- Yinlam Chow, Ofir Nachum, Edgar Dueñez-Guzman, and Mohammad Ghavamzadeh. A Lyapunov-based approach to safe reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 8103–8112, 2018.
- Gal Dalal, Krishnamurthy Dvijotham, Matej Vecerik, Todd Hester, Cosmin Paduraru, and Yuval Tassa. Safe exploration in continuous action spaces. *arXiv:1801.08757*, 2018.
- Marguerite Frank, Philip Wolfe, et al. An algorithm for quadratic programming. *Naval research logistics quarterly*, 3(1-2):95–110, 1956.
- Scott Fujimoto, Herke Hoof, and David Meger. Addressing function approximation error in actor-critic methods. In *International Conference on Machine Learning*, pages 1587–1596, 2018.
- Yasuhiro Fujita and Shin-ichi Maeda. Clipped action policy gradient. In *International Conference on Machine Learning*, pages 1597–1606, 2018.
- Yasuhiro Fujita, Prabhat Nagarajan, Toshiki Kataoka, and Takahiro Ishikawa. Chainerrl: A deep reinforcement learning library. *Journal of Machine Learning Research*, 22(77):1–14, 2021. URL <http://jmlr.org/papers/v22/20-376.html>.
- Supriyo Ghosh and Pradeep Varakantham. Incentivizing the use of bike trailers for dynamic repositioning in bike sharing systems. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 27, 2017.
- Lin Gu, Deze Zeng, Wei Li, Song Guo, Albert Y Zomaya, and Hai Jin. Intelligent VNF orchestration and flow scheduling via model-assisted deep reinforcement learning. *IEEE Journal on Selected Areas in Communications*, 38(2):279–291, 2019.
- Shixiang Gu, Ethan Holly, Timothy Lillicrap, and Sergey Levine. Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 3389–3396, 2017.
- LLC Gurobi Optimization. Gurobi optimizer reference manual, 2021. URL <http://www.gurobi.com>.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International Conference on Machine Learning*, pages 1861–1870. PMLR, 2018.
- Léonard Jaillet and Josep M Porta. Path planning under kinematic constraints by rapidly exploring manifolds. *IEEE Transactions on Robotics*, 29(1):105–117, 2012.
- Nathan Jay, Noga H. Rotman, P. Brighten Godfrey, Michael Schapira, and Aviv Tamar. Internet congestion control via deep reinforcement learning, 2019.
- Sham Kakade and John Langford. Approximately optimal approximate reinforcement learning. In *International Conference on Machine Learning*, 2002.
- Frank Kelly. Charging and rate control for elastic traffic. *European transactions on Telecommunications*, 8(1):33–37, 1997.
- Simon Lacoste-Julien. Convergence rate of Frank-Wolfe for non-convex objectives. *arXiv:1607.00345*, 2016.
- Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. In *International Conference on Learning Representations*, 2016.
- Anqi Liu, Guanya Shi, Soon-Jo Chung, Anima Anandkumar, and Yisong Yue. Robust regression for safe exploration in control. In *Learning for Dynamics and Control*, pages 608–619. PMLR, 2020a.
- Yongshuai Liu, Jiabin Ding, and Xin Liu. IPO: Interior-point policy optimization under constraints. In *AAAI Conference on Artificial Intelligence*, volume 34, pages 4940–4947, 2020b.
- Andrew L Maas, Awni Y Hannun, and Andrew Y Ng. Rectifier nonlinearities improve neural network acoustic models. In *International Conference on Machine Learning*, volume 30, page 3, 2013.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.

- Tu-Hoa Pham, Giovanni De Magistris, and Ryuki Tachibana. Oplayer - practical constrained optimization for deep reinforcement learning in the real world. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 6236–6243, 2018.
- Sashank J Reddi, Suvrit Sra, Barnabás Póczos, and Alex Smola. Stochastic frank-wolfe methods for nonconvex optimization. In *Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1244–1251, 2016.
- Harsh Satija, Philip Amortila, and Joelle Pineau. Constrained markov decision processes via backward value functions. In *International Conference on Machine Learning*, pages 8502–8511. PMLR, 2020.
- John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897. PMLR, 2015.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms, 2017.
- Sanket Shah, Sinha Arunesh, Varakantham Pradeep, Perrault Andrew, and Tambe Milind. Solving online threat screening games using constrained action space reinforcement learning. In *AAAI Conference on Artificial Intelligence*, volume 34, pages 2226–2235, 2020.
- David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic policy gradient algorithms. In *International conference on machine learning*, pages 387–395, 2014.
- Richard S Sutton. Policy gradient methods for reinforcement learning with function approximation. *Advances in Neural Information Processing Systems*, 12:1057–1063, 2000.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- Chen Tessler, Daniel J Mankowitz, and Shie Mannor. Reward constrained policy optimization. In *International Conference on Learning Representations*, 2018.
- E. Todorov, T. Erez, and Y. Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033, 2012. doi: 10.1109/IROS.2012.6386109.
- Vassilios Tsounis, Mitja Alge, Joonho Lee, Farbod Farshidian, and Marco Hutter. Deepgait: Planning and control of quadrupedal gaits using deep reinforcement learning. *IEEE Robotics and Automation Letters*, 5(2):3699–3706, 2020.
- Zhiyuan Xu, Jian Tang, Jingsong Meng, Weiyi Zhang, Yanzhi Wang, Chi Harold Liu, and Dejun Yang. Experience-driven networking: A deep reinforcement learning based approach. In *IEEE Conference on Computer Communications (INFOCOM)*, pages 1871–1879, 2018.
- Tsung-Yen Yang, Justinian Rosca, Karthik Narasimhan, and Peter J Ramadge. Projection-based constrained policy optimization. In *International Conference on Learning Representations*, 2019.
- Junjie Zhang, Minghao Ye, Zehua Guo, Chen-Yu Yen, and H Jonathan Chao. CFR-RL: Traffic engineering with reinforcement learning in SDN. *IEEE Journal on Selected Areas in Communications*, 38(10):2249–2259, 2020a.
- Yiming Zhang, Quan Vuong, and Keith Ross. First order constrained optimization in policy space. *Advances in Neural Information Processing Systems*, 2020b.