

---

# A Hybrid Stochastic Policy Gradient Algorithm for Reinforcement Learning

---

\*Nhan H. Pham

†Lam M. Nguyen

†Dzung T. Phan

‡Phuong Ha Nguyen

‡§Marten van Dijk

\*Quoc Tran-Dinh

\*UNC Chapel Hill

†IBM Research

‡UConn

§CWI Amsterdam

## Abstract

We propose a novel hybrid stochastic policy gradient estimator by combining an unbiased policy gradient estimator, the REINFORCE estimator, with another biased one, an adapted SARAH estimator for policy optimization. The hybrid policy gradient estimator is shown to be biased, but has variance reduced property. Using this estimator, we develop a new Proximal Hybrid Stochastic Policy Gradient Algorithm (ProxHSPGA) to solve a composite policy optimization problem that allows us to handle constraints or regularizers on the policy parameters. We first propose a single-looped algorithm then introduce a more practical restarting variant. We prove that both algorithms can achieve the best-known trajectory complexity  $\mathcal{O}(\varepsilon^{-3})$  to attain a first-order stationary point for the composite problem which is better than existing REINFORCE/GPOMDP  $\mathcal{O}(\varepsilon^{-4})$  and SVRPG  $\mathcal{O}(\varepsilon^{-10/3})$  in the non-composite setting. We evaluate the performance of our algorithm on several well-known examples in reinforcement learning. Numerical results show that our algorithm outperforms two existing methods on these examples. Moreover, the composite settings indeed have some advantages compared to the non-composite ones on certain problems.

## 1 Introduction

Recently, research on reinforcement learning (RL) (Sutton and Barto, 2018), an area of machine learning

to learn how to make a series of decisions while interacting with the underlying environment, has been immensely active. Unlike supervised learning, reinforcement learning agents often have limited or no knowledge about the environment and the rewards of taking certain actions might not be immediately observed, making these problems more challenging to solve. Over the past decade, there has been a large number of research works developing and using reinforcement learning to solve emerging problems. Notable reinforcement learning agents include, but not limited to, AlphaGo and AlphaZero (Silver et al., 2016, 2018), OpenAIFive (OpenAI, 2018), and AlphaStar (DeepMind, 2019).

In modern RL tasks, the environment is often not known beforehand so the agent has to simultaneously learn the environment while making appropriate decisions. One approach is to estimate the value function or the state-value function, e.g., Q-learning (Watkins and Dayan, 1992) and its variants such as Deep Q-learning (DQN) (Mnih et al., 2013, 2015), Dueling DQN (Wang et al., 2016), and double Q-learning (Hasselt et al., 2016).

It has been observed that learning the state-value function is not efficient when the action space is large or even infinite. In that case, policy gradient methods learn the policy directly with a parameterized function. Silver et al. (2014) presents a framework for deterministic policy gradient algorithms which can be estimated more efficiently than their stochastic counterparts whereas DDPG (Lillicrap et al., 2016) adapts the idea of deep Q-learning into continuous action tasks in RL. TRPO (Schulman et al., 2015) uses a constraint on the KL divergence between the new and old policies to improve the robustness of each update. PPO (Schulman et al., 2017) is an extension of TRPO which uses a clipped surrogate objective resulting a simpler implementation. Other policy gradient methods utilize the actor-critic paradigm including ACER (Wang et al., 2017), A3C (Mnih et al., 2016) and its syn-

---

Proceedings of the 23<sup>rd</sup> International Conference on Artificial Intelligence and Statistics (AISTATS) 2020, Palermo, Italy. PMLR: Volume 108. Copyright 2020 by the author(s).

chronous variant A2C, ACKTR (Wu et al., 2017), and SAC (Haarnoja et al., 2018).

REINFORCE (Williams, 1992) is perhaps one classical method closely related to our work here. It uses an estimator of the policy gradient and applies a gradient ascent step to update the policy. Nevertheless, the REINFORCE estimator is known to have high variance leading to several weaknesses. Other improvements to reduce the variance such as adding baselines (Sutton and Barto, 2018; Zhao et al., 2011), discarding some rewards in the so-called GPOMDP estimator (Baxter and Bartlett, 2001) were proposed. While REINFORCE estimator is an unbiased policy gradient estimator, GPOMDP is shown to be biased (Baxter and Bartlett, 2001) making theoretical analysis harder.

The nature of REINFORCE algorithm appears to be closely related to stochastic gradient descent (SGD) (Robbins and Monro, 1951) in stochastic nonconvex optimization. In particular, the standard SGD estimator is also known to often have fixed variance, which is often high. On the one hand, there are algorithms trying to reduce the oscillation (Tieleman and Hinton, 2012) or introduce momentums or adaptive updates (Allen-Zhu, 2017, 2018; Kingma and Ba, 2014) for SGD methods to accelerate performance. On the other hand, other researchers are searching for new gradient estimators. One approach is the SAGA estimator proposed by Defazio et al. (2014). Another well-known estimator is the SVRG estimator (Johnson and Zhang, 2013) which has been intensively studied in recent works, e.g., in Allen-Zhu and Yuan (2016); Li and Li (2018); Reddi et al. (2016); Zhou et al. (2018). This estimator not only overcomes the storage issue of SAGA but also possesses variance reduced property, i.e., the variance of the estimator decreases over epochs. Methods based on SVRG estimators have recently been developed for reinforcement learning, e.g., SVRPG (Papini et al., 2018). Xu et al. (2019a) refines the analysis of SVRPG to achieve an improved trajectory complexity of  $\mathcal{O}(\varepsilon^{-10/3})$ . Shen et al. (2019) also adopts the SVRG estimator into policy gradient and achieve the trajectory oracle complexity of  $\mathcal{O}(\varepsilon^{-3})$  with the use of a second-order estimator.

While SGD, SAGA, and SVRG estimators are unbiased, there have been algorithms developed based on a biased gradient estimator named SARAH (Nguyen et al., 2017b). Such algorithms include SARAH (Nguyen et al., 2017a, 2019), SPIDER (Fang et al., 2018), SpiderBoost (Wang et al., 2018), and Prox-SARAH (Pham et al., 2019b). Similar to SVRG, all these methods can potentially be extended to reinforcement learning. A recent attempt is SARAPO (Yuan et al., 2019) which combines SARAH (Nguyen et al., 2019) with TRPO (Schulman et al., 2015) algo-

rithm but no theoretical guarantee is provided. Yang and Zhang (2019) propose Mirror Policy Optimization (MPO) algorithm which covers the classical policy gradient and the natural policy gradient as special cases. They also introduce a variance reduction variant, called VRMPO, which achieves  $\mathcal{O}(\varepsilon^{-3})$  trajectory complexity. Another notable work is SRVR-PG (Xu et al., 2019b) where the policy gradient estimator is the adapted version of SARAH estimator for reinforcement learning. Note that Yang and Zhang (2019) and Xu et al. (2019b) achieve the same trajectory complexity of  $\mathcal{O}(\varepsilon^{-3})$  as ours. However, our algorithm is essentially different. Xu et al. (2019b) and Yang and Zhang (2019) use two different adaptation of the SARAH estimator for policy gradient. Xu et al. (2019b) uses the importance weight in their estimator to handle distribution shift while Yang and Zhang (2019) remove it as seen in Shen et al. (2019). Meanwhile, we introduce a new policy gradient estimator which can also be calculated recursively. The new estimator is fundamentally different from the other two since it combines the adapted SARAH estimator as in Xu et al. (2019b) with the classical REINFORCE estimator. In addition, our analysis shows that the best-known convergence rate and complexity can be achieved by our single-loop algorithm (Algorithm 1) while SRVR-PG and VRMPO require double loops to achieve the same oracle complexity. Moreover, Xu et al. (2019b); Yang and Zhang (2019) do not consider the composite setting that includes the constraints or regularizers on the policy parameters as we do.

Table 1: A comparison between different methods for the non-composite setting (1) of (2).

Algorithms	Complexity	Composite	Single-loop
REINFORCE (Williams, 1992)	$\mathcal{O}(\varepsilon^{-4})$	✗	✓
GPOMDP (Baxter and Bartlett, 2001)	$\mathcal{O}(\varepsilon^{-4})$	✗	✓
SVRPG (Papini et al., 2018)	$\mathcal{O}(\varepsilon^{-4})$	✗	✗
SVRPG (Xu et al., 2019a)	$\mathcal{O}(\varepsilon^{-10/3})$	✗	✗
HAPG (Shen et al., 2019)	$\mathcal{O}(\varepsilon^{-3})$	✗	✗
VRMPO (Yang and Zhang, 2019)	$\mathcal{O}(\varepsilon^{-3})$	✗	✗
SRVR-PG (Xu et al., 2019b)	$\mathcal{O}(\varepsilon^{-3})$	✗	✗
<b>This work</b>	$\mathcal{O}(\varepsilon^{-3})$	✓	✓

**Our approach:** Our approach lies in the stochastic variance reduction avenue, but using a completely new **hybrid** approach, leading to a novel estimator compared to existing methods in reinforcement learning. We build our estimator by taking a convex combination of the adapted SARAH (Nguyen et al., 2017b) and REINFORCE (Williams, 1992), a classical unbiased policy gradient estimator. This hybrid estimator not only allows us to trade-off the bias and variance between these two estimators but also possesses

useful properties for developing new algorithms. Note that the idea of combining stochastic estimators was first proposed for stochastic optimization in our recent works (Tran-Dinh et al., 2019a,b). Unlike existing policy gradient methods, our algorithm first samples a large batch of trajectories to establish a good search direction. After that, it iteratively updates the policy parameters using our hybrid estimator leading to a single-loop method without any snapshot loop as in SVRG or SARAH variants. In addition, as regularization techniques have shown their effectiveness in deep learning (Neyshabur et al., 2017; Zhang et al., 2017), they possibly have great potential in reinforcement learning algorithms too. A recent study (Liu et al., 2019) shows that regularizations on the policy parameters can greatly improve the performance of policy gradient algorithms. Motivated by these facts, we directly consider a new composite setting (2) as presented in Section 3. For this new composite model, it is not clear if existing algorithms remain convergent by simply adding a projection step on the constraint set, while our method does guarantee convergence.

**Our contribution:** To this end, our contribution in this paper can be summarized as follows:

- (a) We introduce a novel hybrid stochastic policy gradient estimator by combining existing REINFORCE estimator with the adapted SARAH estimator for policy gradient. We investigate some key properties of our estimator that can be used for algorithmic development.
- (b) We propose a new algorithm to solve a composite maximization problem for policy optimization in reinforcement learning. Our model not only covers existing settings but also handles constraints and convex regularizers on policy parameters.
- (c) We provide convergence analysis as the first theoretical result for composite optimization in reinforcement learning and estimate the trajectory complexity of our algorithm and show that our algorithm can achieve the best-known complexity over existing first-order methods (see Table 1).

Our algorithm only has one loop as REINFORCE or GPOMDP, which is fundamentally different from SVRPG, SVRG-adapted, and other SARAH-based algorithms for RL. It can work with single sample or mini-batch and has two steps: proximal gradient step and averaging step with different step-sizes. This makes the algorithm more flexible to use different step-sizes without sacrificing the overall complexity.

**Paper outline:** The rest of this paper is organized as follows. Section 2 describes problem of interest and gives an overview about policy gradient methods. Section 3 introduces our new hybrid estimator for policy gradient and develops the main algorithm. The com-

plexity analysis is presented in Section 4, while Section 5 provides several numerical examples. All technical proofs and experimental details are given in Supplementary Document (Supp. Doc.).

## 2 Model and Problem Statement

**Model:** We consider a Markov Decision Process (MDP) (Sutton and Barto, 2018) equipped with 6 components  $\{\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma, \mathcal{P}_0\}$  where  $\mathcal{S}$ ,  $\mathcal{A}$  are the state and action spaces,  $\mathcal{P}$  denotes the set of transition probabilities when taking certain actions,  $\mathcal{R}$  is the reward function which characterizes the immediate reward earned by taking certain action,  $\gamma$  is a discount factor, and  $\mathcal{P}_0$  is the initial state distribution.

Let  $\pi(\cdot|s)$  be a density function over  $\mathcal{A}$  when current state is  $s$  and  $\pi_\theta(\cdot|s)$  is a policy parameterized by parameter  $\theta$ . A trajectory  $\tau = \{s_0, a_0, s_1, a_1, \dots, s_{H-1}, a_{H-1}\}$  with effective length  $H$  is a collection of states and actions sampled from a stationary policy. Denote  $p_\theta(\cdot)$  as the density induced by policy  $\pi_\theta$  over all possible trajectories and  $p_\theta(\tau)$  is the probability of observing a trajectory  $\tau$ . Also, let  $\mathcal{R}(\tau) = \sum_{j=0}^{H-1} \gamma^j \mathcal{R}(s_j, a_j)$  be the total discounted reward for a trajectory  $\tau$ . Solving an MDP is equivalent to finding the solution that maximizes the expected cumulative discounted rewards.

**Classical policy gradient methods:** Policy gradient methods seek a differentiable parameterized policy  $\pi_\theta$  that maximizes the expected cumulative discounted rewards as

$$\max_{\theta \in \mathbb{R}^q} \left\{ J(\theta) := \mathbb{E}_{\tau \sim p_\theta} [\mathcal{R}(\tau)] \right\}. \quad (1)$$

where  $q$  is the parameter dimension. The policy gradient theorem (Sutton et al., 1999) shows that

$$\nabla J(\theta) = \mathbb{E}_{\tau \sim p_\theta} [\nabla \log p_\theta(\tau) \mathcal{R}(\tau)],$$

where the policy gradient does not depend on the gradient of the state distribution despite the fact that the state distribution depends on the policy parameters (Silver et al., 2014).

This policy gradient can be used in gradient ascent algorithms to update the parameter  $\theta$ . However, we cannot calculate the full gradient at each update as we only get a finite number of samples at each iteration. Consequently, the policy gradient is often estimated by its sample average. At each iteration, a batch of trajectories  $\mathcal{B} = \{\tau_i\}_{i=1, \dots, N}$  will be sampled from the environment to estimate the policy gradient as

$$\tilde{\nabla} J(\theta) := \frac{1}{N} \sum_{i=1}^N g(\tau_i|\theta),$$

where  $g(\tau_i|\theta)$  is a sample estimator of  $\mathbb{E}_{\tau_i \sim p_\theta} [\nabla \log p_\theta(\tau_i) \mathcal{R}(\tau_i)]$ . We call  $\tilde{\nabla} J(\theta)$  a stochastic policy gradient (SPG) estimator. This estimator has been exploited in the two well-known REINFORCE (Williams, 1992) and GPOMDP (Baxter and Bartlett, 2001) methods. The main step of policy gradient ascent methods is to update the parameters as

$$\theta_{t+1} := \theta_t + \eta \nabla J(\theta_t), \quad t = 0, 1, \dots,$$

where  $\eta > 0$  is some appropriate learning rate, which can be fixed or varied over  $t$ . Since the policy changes after each update, the density  $p_\theta(\cdot)$  also changes and creates non-stationarity in the problem which will be handled by importance weight in Section 3.

### 3 A New Hybrid Stochastic Policy Gradient Algorithm

In this section, we first introduce a composite model for policy optimization. Next, we extend the hybrid gradient idea from Tran-Dinh et al. (2019b) to policy gradient estimators. Finally, we develop a new proximal policy gradient algorithm and its restart variant to solve the composite policy optimization problem and analyze their trajectory complexity.

#### 3.1 Composite Policy Optimization Model

While the objective function in (1) is standard in most policy gradient methods, it is natural to have some constraints or regularizers on the policy parameters. In addition, adding constraints can prevent the explosion of parameters in highly nonlinear models as often seen in deep learning (Srivastava et al., 2014). Adopting the idea of composite nonconvex optimization (Pham et al., 2019b), we are interested in the more general optimization problem in reinforcement learning as follow:

$$\max_{\theta \in \mathbb{R}^q} \left\{ J(\theta) - Q(\theta) = \mathbb{E}_{\tau \sim p_\theta} [\mathcal{R}(\tau)] - Q(\theta) \right\}, \quad (2)$$

where  $Q(\theta)$  is a proper, closed, and convex function acting as a regularizer which can be the indicator function of a convex set representing the constraints on the parameters or some standard regularizers such as  $\ell_1$ -norm or  $\ell_2$ -norm. If there is no regularizer  $Q(\theta)$ , the problem (2) reduces to the standard one in (1).

#### 3.2 Assumptions

Let  $F(\theta) := J(\theta) - Q(\theta)$  be the total objective function. We impose the following assumptions for our convergence analysis, which are often used in practice.

**Assumption 3.1.** *The regularizer  $Q : \mathbb{R}^q \rightarrow \mathbb{R} \cup \{+\infty\}$  is a proper, closed, and convex function. We*

*also assume that the domain of  $F$  is nonempty and there exists a finite upper bound*

$$F^* := \sup_{\theta \in \mathbb{R}^q} \{F(\theta) := J(\theta) - Q(\theta)\} < +\infty.$$

**Assumption 3.2.** *The immediate reward function is bounded, i.e., there exists  $R > 0$  such that for all  $a \in \mathcal{A}$ ,  $s \in \mathcal{S}$ ,  $|\mathcal{R}(s, a)| \leq R$ .*

**Assumption 3.3.** *Let  $\pi_\theta(s, a)$  be the policy for a given state-action pair  $(s, a)$ . Then, there exist two positive constants  $G$  and  $M$  such that*

$$\|\nabla \log \pi_\theta(s, a)\| \leq G \text{ and } \|\nabla^2 \log \pi_\theta(s, a)\| \leq M,$$

*for any  $a \in \mathcal{A}$ ,  $s \in \mathcal{S}$  where  $\|\cdot\|$  is the  $\ell_2$ -norm.*

This assumption leads to useful results about the smoothness of  $J(\theta)$  and  $g(\tau|\theta)$  and the upper bound on the variance of the policy gradient estimator.

**Lemma 3.1** ((Papini et al., 2018; Shen et al., 2019; Xu et al., 2019a)). *Under Assumption 3.2 and 3.3, for all  $\theta, \theta_1, \theta_2 \in \mathbb{R}^q$ , we have*

- $\|\nabla J(\theta_1) - \nabla J(\theta_2)\| \leq L \|\theta_1 - \theta_2\|;$
- $\|g(\tau|\theta_1) - g(\tau|\theta_2)\| \leq L_g \|\theta_1 - \theta_2\|;$
- $\|g(\tau, \theta)\| \leq C_g;$  and
- $\|g(\tau|\theta) - \nabla J(\theta)\|^2 \leq \sigma^2,$

*where  $g(\cdot)$  is the REINFORCE estimator and  $L$ ,  $L_g$ ,  $C_g$ , and  $\sigma^2$  are constants depending only on  $R$ ,  $G$ ,  $M$ ,  $H$ ,  $\gamma$ , and the baseline  $b$ .*

For more details about the constants and the proofs of Lemma 3.1 we refer e.g., to Papini et al. (2018); Shen et al. (2019); Xu et al. (2019a).

**Assumption 3.4.** *There exists a constant  $W \in (0, \infty)$  such that, for each pair of policies encountered in Algorithm 1 the following holds*

$$\text{Var}[\omega(\tau|\theta_1, \theta_2)] \leq W, \quad \theta_1, \theta_2 \in \mathbb{R}^q, \quad \tau \sim p_{\theta_1},$$

*where  $\omega(\tau|\theta_1, \theta_2) = \frac{p_{\theta_2}(\tau)}{p_{\theta_1}(\tau)}$  is the importance weight between  $p_{\theta_2}(\cdot)$  and  $p_{\theta_1}(\cdot)$ .*

Since the importance weight  $\omega$  introduces another source of variance, we require this assumption for our convergence analysis as used in previous works, e.g., in Papini et al. (2018); Xu et al. (2019a).

**Remark 3.1.** *Cortes et al. (2010) shows that if  $\sigma_Q$ ,  $\sigma_P$  are variances of two Gaussian distributions  $P$  and  $Q$ , and  $\sigma_Q > \frac{1}{\sqrt{2}}\sigma_P$  then the variance of the importance weights is bounded, i.e. Assumption 3.4 holds for Gaussian policies which are commonly used to represent the policy in continuous control tasks.*

### 3.3 Optimality Condition

Associated with problem (2), we define

$$\mathcal{G}_\eta(\theta) := \eta^{-1} [\text{prox}_{\eta Q}(\theta + \eta \nabla J(\theta)) - \theta], \quad (3)$$

for some  $\eta > 0$  as the gradient mapping of  $F(\theta)$  (Nesterov, 2014), where  $\text{prox}_Q(\theta) := \text{argmin}_{\theta'} \{Q(\theta') + \frac{1}{2} \|\theta' - \theta\|^2\}$  denotes the proximal operator of  $Q$  (see, e.g., Parikh and Boyd (2014) for more details).

A point  $\theta^*$  is called a stationary point of (2) if

$$\mathbb{E} [\|\mathcal{G}_\eta(\theta^*)\|^2] = 0.$$

Our goal is to design an iterative method to produce an  $\varepsilon$ -approximate stationary point  $\theta_T$  of (2) after at most  $T$  iterations defined as

$$\mathbb{E} [\|\mathcal{G}_\eta(\tilde{\theta}_T)\|^2] \leq \varepsilon^2,$$

where  $\varepsilon > 0$  is a desired tolerance, and the expectation is taken overall the randomness up to  $T$  iterations.

### 3.4 Novel Hybrid SPG Estimator

**Unbiased estimator:** Recall that given a trajectory  $\tau := \{s_0, a_0, \dots, s_{H-1}, a_{H-1}\}$ , the REINFORCE (SPG) estimator is defined as

$$g(\tau|\theta) := \left[ \sum_{t=0}^{H-1} \nabla \log \pi_\theta(a_t|s_t) \right] \mathcal{R}(\tau),$$

where  $\mathcal{R}(\tau) := \sum_{t=0}^{H-1} \gamma^t \mathcal{R}(s_t, a_t)$ .

Note that the REINFORCE estimator is unbiased, i.e.  $\mathbb{E}_{\tau \sim p_\theta} [g(\tau|\theta)] = \nabla J(\theta)$ . In order to reduce the variance of these estimators, a baseline is normally added while maintaining the unbiasedness of the estimators (Sutton and Barto, 2018; Zhao et al., 2011). From now on, we will refer to  $g(\tau|\theta)$  as the baseline-added version defined as

$$g(\tau|\theta) := \sum_{t=0}^{T-1} \nabla \log \pi_\theta(a_t|s_t) A_t,$$

where  $A_t := \mathcal{R}(\tau) - b_t$  with  $b_t$  being a baseline and possibly depending only on  $s_t$ .

**Hybrid SPG estimator:** In order to reduce the number of trajectories sampled, we extend our idea in Tran-Dinh et al. (2019b) for stochastic optimization to develop a new hybrid stochastic policy gradient (HSPG) estimator that helps balance the bias-variance trade-off. The estimator is formed by taking a convex combination of two other estimators: one is an unbiased estimator which can be REINFORCE estimator, and another is the adapted SARAH estimator (Nguyen et al., 2017b) for policy gradient which is biased.

More precisely, if  $\mathcal{B}_t$  and  $\widehat{\mathcal{B}}_t$  are two random batches of trajectories with sizes  $B$  and  $\widehat{B}$ , respectively, sampled from  $p_{\theta_t}(\cdot)$ , the hybrid stochastic policy gradient estimator at  $t$ -th iteration can be expressed as

$$v_t := \beta v_{t-1} + \frac{\beta}{B} \sum_{\tau \in \mathcal{B}_t} \Delta g(\tau|\theta_t) + \frac{(1-\beta)}{\widehat{B}} \sum_{\hat{\tau} \in \widehat{\mathcal{B}}_t} g(\hat{\tau}|\theta_t), \quad (4)$$

where

$$\Delta g(\tau|\theta_t) := g(\tau|\theta_t) - \omega(\tau|\theta_t, \theta_{t-1}) g(\tau|\theta_{t-1}),$$

and

$$v_0 := \frac{1}{N} \sum_{\tau \in \widehat{\mathcal{B}}} g(\tau|\theta_0),$$

with  $\widehat{\mathcal{B}}$  is a batch of trajectories collected at the beginning. Note that  $\omega(\tau|\theta_t, \theta_{t-1})$  is an importance weight added to account for the distribution shift since the trajectories  $\tau \in \mathcal{B}_t$  are sampled from  $p_{\theta_t}(\cdot)$  but not from  $p_{\theta_{t-1}}(\cdot)$ . Note also that  $v_t$  in (4) is also different from the momentum SARAH estimator recently proposed in Cutkosky and Orabona (2019).

### 3.5 The Complete Algorithm

The novel Proximal Hybrid Stochastic Policy Gradient Algorithm (abbreviated by ProxHSPGA) to solve (2) is presented in detail in Algorithm 1.

---

#### Algorithm 1 (ProxHSPGA)

---

- 1: **Initialization:** An initial point  $\theta_0 \in \mathbb{R}^q$ , and positive parameters  $m, N, B, \widehat{B}, \beta, \alpha$ , and  $\eta$  (specified later).
- 2: Sample a batch of trajectories  $\widehat{\mathcal{B}}$  of size  $N$  from  $p_{\theta_0}(\cdot)$ .
- 3: Calculate  $v_0 := \frac{1}{N} \sum_{\tau \in \widehat{\mathcal{B}}} g(\tau|\theta_0)$ .
- 4: Update

$$\begin{cases} \widehat{\theta}_1 & := \text{prox}_{\eta Q}(\theta_0 + \eta v_0) \\ \theta_1 & := (1 - \alpha)\theta_0 + \alpha \widehat{\theta}_1. \end{cases}$$

- 5: **For**  $t := 1, \dots, m$  **do**
- 6: Generate 2 independent batches of trajectories  $\mathcal{B}_t$  and  $\widehat{\mathcal{B}}_t$  with size  $B$  and  $\widehat{B}$  from  $p_{\theta_t}(\cdot)$ .
- 7: Evaluate the hybrid estimator  $v_t$  as in (4).
- 8: Update

$$\begin{cases} \widehat{\theta}_{t+1} & := \text{prox}_{\eta Q}(\theta_t + \eta v_t) \\ \theta_{t+1} & := (1 - \alpha)\theta_t + \alpha \widehat{\theta}_{t+1}. \end{cases}$$

- 9: **EndFor**
  - 10: Choose  $\tilde{\theta}_T$  from  $\{\theta_t\}_{t=1}^m$  uniformly randomly.
- 

Unlike SVRPG (Papini et al., 2018; Xu et al., 2019a) and HAPG (Shen et al., 2019), Algorithm 1 only has **one loop** as REINFORCE or GPOMDP. Moreover, Algorithm 1 does not use the estimator for the policy Hessian as in HAPG. At the initial stage, a batch of

trajectories is sampled using  $p_{\theta_0}$  to estimate an initial policy gradient estimator which provides a good initial search direction. At the  $t$ -th iteration, two independent batches of trajectories are sampled from  $p_{\theta_t}$  to evaluate the hybrid stochastic policy gradient estimator. After that, a proximal step followed by an averaging step are performed which are inspired by Pham et al. (2019b). Note that the batches of trajectories at each iteration are sampled from the current distribution which will change after each update. Therefore, the importance weight  $\omega(\tau|\theta_t, \theta_{t-1})$  is introduced to account for the non-stationarity of the sampling distribution. As a result, we still have  $\mathbb{E}_{\tau \sim p_{\theta_t}} [\omega(\tau|\theta_t, \theta_{t-1})g(\tau|\theta_{t-1})] = \nabla J(\theta_{t-1})$ .

### 3.6 Restarting variant

While Algorithm 1 has the best-known theoretical complexity as shown in Section 4, its practical performance may be affected by the constant step-size  $\alpha$  depending on  $m$ . As will be shown later, the step-size  $\alpha \in [0, 1]$  is inversely proportional to the number of iterations  $m$  and it is natural to have  $\alpha$  close to 1 to take advantage of the newly computed information. To increase the practical performance of our algorithm without sacrificing its complexity, we propose to inject a simple restarting strategy by repeatedly running Algorithm 1 for multiple stages as in Algorithm 2.

---

#### Algorithm 2 (Restarting ProxHSPGA)

---

- 1: **Initialization:** Input an initial point  $\theta_0^{(0)}$ .
  - 2: **For**  $s := 0, \dots, S-1$  **do**
  - 3:     Run Algorithm 1 with  $\theta_0 := \theta_0^{(s)}$ .
  - 4:     Output  $\theta_0^{(s+1)} := \theta_{m+1}$ .
  - 5: **EndFor**
  - 6: Choose  $\tilde{\theta}_T$  uniformly randomly from  $\{\theta_t^{(s)}\}_{s=0 \rightarrow S-1, t=0 \rightarrow m}$ .
- 

We emphasize that without this restarting strategy, Algorithm 1 still converges and the restarting loop in Algorithm 2 does not sacrifice the best-known complexity as stated in the next section.

## 4 Convergence Analysis

This section presents key properties of the hybrid stochastic policy gradient estimators as well as the theoretical convergence analysis and complexity estimate.

### 4.1 Properties of the hybrid SPG estimator

Let  $\mathcal{F}_t := \sigma(\tilde{\mathcal{B}}, \mathcal{B}_1, \hat{\mathcal{B}}_1, \dots, \mathcal{B}_{t-1}, \hat{\mathcal{B}}_{t-1})$  be the  $\sigma$ -field generated by all trajectories sampled up to the  $t$ -th iteration. For the sake of simplicity, we assume that  $B = \hat{B}$  but our analysis can be easily extended for the

case  $B \neq \hat{B}$ . Then the hybrid SPG estimator  $v_t$  has the following properties

**Lemma 4.1** (Key properties). *Let  $v_t$  be defined as in (4) and  $\Delta v_t := v_t - \nabla J(\theta_t)$ . Then*

$$\mathbb{E}_{\tau, \hat{\tau} \sim p_{\theta_t}} [v_t] = \nabla J(\theta_t) + \beta \Delta v_{t-1}. \quad (5)$$

*If  $\beta \neq 0$ , then  $v_t$  is an biased estimator. In addition, we have*

$$\mathbb{E}_{\tau, \hat{\tau} \sim p_{\theta_t}} [\|\Delta v_t\|^2] \leq \beta^2 \|\Delta v_{t-1}\|^2 + \frac{(1-\beta)^2 \sigma^2}{B} + \frac{\beta^2 \bar{C}}{B} \|\theta_t - \theta_{t-1}\|^2, \quad (6)$$

*where  $\bar{C} > 0$  is a given constant.*

The proof of Lemma 4.1 and the explicit constants are given in Supp. Doc. A.1 due to space limit.

### 4.2 Complexity Estimates

The following lemma presents a key estimate for our convergence results.

**Lemma 4.2** (One-iteration analysis). *Under Assumption 3.2, 3.3, and 3.4, let  $\{\hat{\theta}_t, \theta_t\}_{t=0}^m$  be the sequence generated by Algorithm 1 and  $\mathcal{G}_\eta$  be the gradient mapping defined in (3). Then*

$$\mathbb{E} [F(\theta_{t+1})] \geq \mathbb{E} [F(\theta_t)] + \frac{\eta^2 \alpha}{2} \mathbb{E} [\|\mathcal{G}_\eta(\theta_t)\|^2] - \frac{\xi}{2} \mathbb{E} [\|v_t - \nabla J(\theta_t)\|^2] + \frac{\zeta}{2} \mathbb{E} [\|\hat{\theta}_{t+1} - \theta_t\|^2], \quad (7)$$

*where  $\xi := \alpha(1 + 2\eta^2)$  and  $\zeta := \alpha(\frac{2}{\eta} - L\alpha - 3) > 0$  provided that  $\alpha \in (0, 1]$  and  $\frac{2}{\eta} - L\alpha - 3 > 0$ .*

The following theorem summarizes the convergence analysis of Algorithm 1.

**Theorem 4.1.** *Under Assumption 3.1, 3.2, 3.3, and 3.4, let  $\{\theta_t\}_{t=0}^m$  be the sequence generated by Algorithm 1 with*

$$\begin{cases} \beta := 1 - \frac{\sqrt{B}}{\sqrt{N(m+1)}} \\ \alpha := \frac{\hat{c}\sqrt{2}B^{3/4}}{\sqrt{3\bar{C}}N^{1/4}(m+1)^{1/4}} \\ \eta := \frac{2}{4+L\alpha}, \end{cases} \quad (8)$$

*where  $B, \hat{c}, L$ , and  $\bar{C}$  are given constants. If  $\tilde{\theta}_T$  is chosen uniformly at random from  $\{\theta_t\}_{t=0}^m$ , then the following estimate holds*

$$\mathbb{E} [\|\mathcal{G}_\eta(\tilde{\theta}_T)\|^2] \leq \frac{3(4+L)^2 \sigma^2}{4[BN(m+1)]^{1/2}} + \frac{(4+L)^2 \sqrt{3\bar{C}}N^{1/4}}{4\hat{c}\sqrt{2}[B(m+1)]^{3/4}} [F^* - F(\theta_0)]. \quad (9)$$

Consequently, the trajectory complexity is presented in the following corollary.

**Corollary 4.1.** For both Algorithm 1 and Algorithm 2, let us fix  $B \in \mathbb{N}_+$  and set  $N := \tilde{c}\sigma^{8/3}[B(m+1)]^{1/3}$  for some  $\tilde{c} > 0$  in Theorem 4.1. If we also choose  $m$  in Algorithm 1 such that  $m+1 = \frac{\Psi_0^{3/2}\sigma}{B\varepsilon^3}$  and choose  $m, S$  in Algorithm 2 such that  $S(m+1) = \frac{\Psi_0^{3/2}\sigma}{B\varepsilon^3}$  for some constant  $\Psi_0$ , then the number of trajectories  $\mathcal{T}_{traj}$  to achieve  $\tilde{\theta}_T$  such that  $\mathbb{E}[\|\mathcal{G}_\eta(\tilde{\theta}_T)\|^2] \leq \varepsilon^2$  for any  $\varepsilon > 0$  is at most

$$\mathcal{T}_{traj} = \mathcal{O}(\varepsilon^{-3}).$$

where  $\tilde{\theta}_T$  is chosen uniformly at random from  $\{\theta_t^{(s)}\}_{t=0, \dots, m}^{s=0, \dots, S-1}$  if using Algorithm 2.

The proof of Theorem 4.1 and Corollary 4.1 are given in Supp. Doc. A.3, and A.4, respectively.

Comparing our complexity bound with other existing methods in Table 1, we can see that we improve a factor of  $\varepsilon^{-1/3}$  over SVRPG in Xu et al. (2019a) while matching the best-known complexity without the need of using the policy Hessian estimator as HAPG from Shen et al. (2019).

## 5 Numerical Experiments

In this section, we present three examples to provide comparison between the performance of HSPGA and other related policy gradient methods. We also provide an example to illustrate the effect of the regularizer  $Q(\cdot)$  to our model (2). More examples can be found in the Supp. Doc. C. All experiments are run on a Macbook Pro with 2.3 GHz Quad-Core, 8GB RAM.

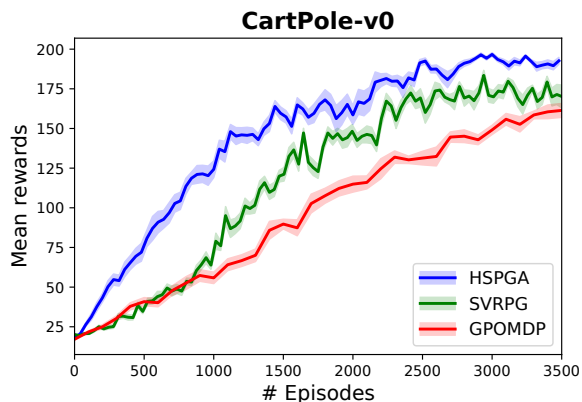


Figure 1: The performance of three algorithms on the Cartpole-v0 environment.

We implement our restarting algorithm, Algorithm 2, on top of the `rllab`<sup>1</sup> library (Duan et al., 2016). The source code is available at <https://github.com/unc-optimization/ProxHSPGA>. We compare our algorithm with two other methods: SVRPG (Papini

<sup>1</sup>Available at <https://github.com/rll/rllab>




Figure 2: The performance of three algorithms on the Acrobot-v1 environment.

et al., 2018; Xu et al., 2019a) and GPOMDP (Baxter and Bartlett, 2001). Although REINFORCE and GPOMDP have the same trajectory complexity, as observed in (Papini et al., 2018), GPOMDP often performs better than REINFORCE, so we only choose to implement GPOMDP in our experiments. Since SVRPG and GPOMDP solves the non-composite problems (1), we set  $Q(\theta) = 0$  in the first three examples and adjust our algorithm, denoted as HSPGA, accordingly. We compare our algorithm with the fixed epoch length variant of SVRPG as reported in Papini et al. (2018); Xu et al. (2019a). For the implementation of SVRPG and GPOMDP, we reuse the implementation of Papini et al.<sup>2</sup> We test these algorithms on three well-studied reinforcement learning tasks: `Cart Pole`, `Acrobot`, and `Mountain Car` which are available in OpenAI `gym` (Brockman et al., 2016), a well-known toolkit for developing and comparing reinforcement learning algorithms. We also test these algorithms on continuous control tasks using other simulators such as `Roboschool` (Klimov and Schulman, 2017) and `Mujoco` (Todorov et al., 2012).

For each environment, we initialize the policy randomly and use it as initial policies for all 10 runs of all algorithms. The performance measure, i.e., mean rewards, is computed by averaging the final rewards of 50 trajectories sampled by the current policy. We then compute the mean and 90% confidence interval across 10 runs of these performance measures at different time point. In all plots, the solid lines represent the mean and the shaded areas are the confidence band of the mean rewards. In addition, detailed configurations of the policy network and parameters can be found in Supp. Doc. B. We note that the architecture of the neural network is denoted as  $[\text{observation space}] \times [\text{hidden layers}] \times [\text{action space}]$ .

<sup>2</sup>Available at <https://github.com/Dam930/rllab>

**Cart Pole-v0 environment:** For the `Cart pole` environment, we use a deep soft-max policy network (Bridle, 1990; Levine, 2017; Sutton and Barto, 2018) with one hidden layer of 8 neurons. Figure 1 depicts the results where we run each algorithm for 10 times and compute the mean and 90% confidence intervals.

From Figure 1, we can see that HSPGA outperforms the other 2 algorithms while SVRPG works better than GPOMDP as expected. HSPGA is able to reach the maximum reward of 200 in less than 4000 episodes.




Figure 3: The performance of three algorithms on the `Mountain Car-v0` environment.

**Acrobot environment:** Next, we evaluate three algorithms on the `Acrobot-v1` environment. Here, we use a deep soft-max policy with one hidden layer of 16 neurons. The performance of these 3 algorithms are illustrated in Figure 2.

We observe similar results as in the previous example where HSPGA has the best performance over three candidates. SVRPG is still better than GPOMDP in this example.

**Mountain Car environment:** For the `MountainCar-v0` environment, we use a deep Gaussian policy (Sutton and Barto, 2018) where the mean is the output of a neural network containing one hidden layer of 8 neurons and the standard deviation is fixed at 1. The results of three algorithms are presented in Figure 3.

Figure 3 shows that HSPGA highly outperforms the other two algorithms. Again, SVRPG remains better than GPOMDP as expected.

**The effect of regularizers:** We test the effect of the regularizer  $Q(\cdot)$  by adding a Tikhonov one as

$$\max_{\theta \in \mathbb{R}^q} \left\{ J(\theta) - \lambda \|\theta\|_2^2 \right\}.$$

This model was intensively studied in Liu et al. (2019).

We also compare all non-composite algorithms with ProxHSPGA in the `Roboschool Inverted`

`Pendulum-v1` environment. In this experiment, we set the penalty parameter  $\lambda = 0.001$  for ProxHSPGA. The results are depicted in Figure 4 and more information about the configuration of each algorithm is in Supp. Document B.




Figure 4: The performance of composite vs. non-composite algorithms on the `Roboschool Inverted Pendulum-v1` environment.

From Figure 4, in terms of non-composite algorithms, HSPGA is the best followed by SVRPG and then by GPOMDP. Furthermore, ProxHSPGA shows its advantage by reaching the maximum reward of 1000 faster than HSPGA.

## 6 Conclusion

We have presented a novel policy gradient algorithm to solve regularized reinforcement learning models. Our algorithm uses a novel policy gradient estimator which is a combination of an unbiased estimator, i.e. REINFORCE estimator, and a biased estimator adapted from SARAH estimator for policy gradient. Theoretical results show that our algorithm achieves the best-known trajectory complexity to attain an  $\varepsilon$ -approximate first-order solution for the problem under standard assumptions. In addition, our numerical experiments not only help confirm the benefit of our algorithm compared to other closely related policy gradient methods but also verify the effectiveness of regularization in policy gradient methods.

## Acknowledgments

Q. Tran-Dinh has partly been supported by the National Science Foundation (NSF), grant no. DMS-1619884 and the Office of Naval Research (ONR), grant no. N00014-20-1-2088 (2020-2023). Q. Tran-Dinh and N. H. Pham are partly supported by The Statistical and Applied Mathematical Sciences Institute (SAMSI).



## References

- Z. Allen-Zhu. Katyusha: The first direct acceleration of stochastic gradient methods. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1200–1205, New York, NY, USA, 2017.
- Z. Allen-Zhu. Natasha 2: Faster non-convex optimization than sgd. In *Advances in Neural Information Processing Systems 31*, pages 2675–2686, 2018.
- Z. Allen-Zhu and Y. Yuan. Improved svrg for non-strongly-convex or sum-of-non-convex objectives. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, pages 1080–1089, 2016.
- J. Baxter and P. L. Bartlett. Infinite-horizon policy-gradient estimation. *J. Artif. Int. Res.*, 15(1):319–350, November 2001.
- J. S. Bridle. Training stochastic model recognition algorithms as networks can lead to maximum mutual information estimation of parameters. In *Advances in Neural Information Processing Systems 2*, pages 211–217. Morgan-Kaufmann, 1990.
- G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba. Openai gym, 2016.
- C. Cortes, Y. Mansour, and M. Mohri. Learning bounds for importance weighting. In *Advances in Neural Information Processing Systems 23*, pages 442–450. 2010.
- A. Cutkosky and F. Orabona. Momentum-based variance reduction in non-convex sgd. In *Advances in Neural Information Processing Systems*, pages 15210–15219, 2019.
- DeepMind. AlphaStar: Mastering the Real-Time Strategy Game StarCraft II. <https://deepmind.com/blog>, 2019.
- A. Defazio, F. Bach, and S. Lacoste-Julien. Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 1*, pages 1646–1654, Cambridge, MA, USA, 2014.
- Y. Duan, X. Chen, R. Houthoof, J. Schulman, and P. Abbeel. Benchmarking deep reinforcement learning for continuous control. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, pages 1329–1338, 2016.
- C. Fang, C. J. Li, Z. Lin, and T. Zhang. Spider: Near-optimal non-convex optimization via stochastic path-integrated differential estimator. In *NeurIPS*, 2018.
- T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1861–1870, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR.
- H. van Hasselt, A. Guez, and D. Silver. Deep reinforcement learning with double q-learning. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pages 2094–2100, 2016.
- R. Johnson and T. Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in Neural Information Processing Systems (NIPS)*, pages 315–323, 2013.
- D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014.
- O. Klimov and J. Schulman. Roboschool. <https://openai.com/blog/roboschool/>, 2017.
- S. Levine. CS 294-112: Deep reinforcement learning lecture notes, 2017.
- Z. Li and J. Li. A simple proximal stochastic gradient method for nonsmooth nonconvex optimization. In *Proceedings of the 32Nd International Conference on Neural Information Processing Systems*, pages 5569–5579, USA, 2018.
- T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra. Continuous control with deep reinforcement learning. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.
- Z. Liu, X. Li, B. Kang, and T. Darrell. Regularization matters in policy optimization. *arXiv preprint arXiv:1910.09191*, 2019.
- V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. A. Riedmiller. Playing atari with deep reinforcement learning. *ArXiv*, abs/1312.5602, 2013.
- V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. A. Riedmiller, A. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518:529–533, 2015.
- V. Mnih, A. P. Badia, M. Mirza, A. Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48, pages 1928–1937, 20–22 Jun 2016.

- Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Springer Publishing Company, Incorporated, 1 edition, 2014.
- B. Neyshabur, S. Bhojanapalli, D. Mcallester, and N. Srebro. Exploring generalization in deep learning. In *Advances in Neural Information Processing Systems 30*, pages 5947–5956. 2017.
- L. M. Nguyen, J. Liu, K. Scheinberg, and M. Takáč. Stochastic recursive gradient algorithm for nonconvex optimization. *CoRR*, abs/1705.07261, 2017a.
- L. M. Nguyen, J. Liu, K. Scheinberg, and M. Takáč. SARAH: A novel method for machine learning problems using stochastic recursive gradient. In *Proceedings of the 34th International Conference on Machine Learning*, pages 2613–2621, 2017b.
- L. M. Nguyen, M. van Dijk, D. T. Phan, P. H. Nguyen, T.-W. Weng, and J. R. Kalagnanam. Finite-sum smooth optimization with sarah. *arXiv preprint arXiv:1901.07648*, 2019.
- OpenAI. OpenAI Five. <https://blog.openai.com/openai-sfive/>, 2018.
- M. Papini, D. Binaghi, G. Canonaco, M. Pirota, and M. Restelli. Stochastic variance-reduced policy gradient. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80, pages 4026–4035, 10–15 Jul 2018.
- N. Parikh and S. Boyd. Proximal algorithms. *Found. Trends Optim.*, 1(3):127–239, January 2014.
- N. H. Pham, L. M. Nguyen, D. T. Phan, and Q. Tran-Dinh. Proxsarah: An efficient algorithmic framework for stochastic composite nonconvex optimization. *ArXiv*, abs/1902.05679, 2019b.
- S. J. Reddi, A. Hefny, S. Sra, B. Póczós, and A. Smola. Stochastic variance reduction for nonconvex optimization. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, pages 314–323, 2016.
- H. Robbins and S. Monro. A stochastic approximation method. *Ann. Math. Statist.*, 22(3):400–407, 09 1951.
- J. Schulman, S. Levine, P. Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37, pages 1889–1897, Lille, France, 07–09 Jul 2015.
- J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *ArXiv*, abs/1707.06347, 2017.
- Z. Shen, A. Ribeiro, H. Hassani, H. Qian, and C. Mi. Hessian aided policy gradient. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 5729–5738, Long Beach, California, USA, 09–15 Jun 2019.
- D. Silver, G. Lever, N. Heess, T. Degris, D. Wierstra, and M. Riedmiller. Deterministic policy gradient algorithms. In *Proceedings of the 31st International Conference on International Conference on Machine Learning - Volume 32*, pages I–387–I–395, 2014.
- D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. v. d. Drisssche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis. Mastering the game of go with deep neural networks and tree search. *Nature*, 529:484–503, 2016.
- D. Silver, T. Hubert, J. Schrittwieser, I. Antonoglou, M. Lai, A. Guez, M. Lanctot, L. Sifre, D. Kumaran, T. Graepel, T. Lillicrap, K. Simonyan, and D. Hassabis. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419):1140–1144, 2018.
- N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- R. S. Sutton and A. G. Barto. *Introduction to Reinforcement Learning, 2nd Edition*. MIT Press, 2018.
- R. S. Sutton, D. McAllester, S. Singh, and Y. Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Proceedings of the 12th International Conference on Neural Information Processing Systems*, pages 1057–1063, 1999.
- T. Tieleman and G. Hinton. Lecture 6.5—RmsProp: Divide the gradient by a running average of its recent magnitude. COURSERA: Neural Networks for Machine Learning, 2012.
- E. Todorov, T. Erez, and Y. Tassa. Mujoco: A physics engine for model-based control. *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033, 2012.
- Q. Tran-Dinh, N. H. Pham, D. T. Phan, and L. M. Nguyen. Hybrid stochastic gradient descent algorithms for stochastic nonconvex optimization. *arXiv preprint arXiv:1905.05920*, 2019a.
- Q. Tran-Dinh, N. H. Pham, D. T. Phan, and L. M. Nguyen. A hybrid stochastic optimization framework for stochastic composite nonconvex optimization. *arXiv preprint arXiv:1907.03793*, 2019b.
- Z. Wang, T. Schaul, M. Hessel, H. Van Hasselt, M. Lanctot, and N. De Freitas. Dueling network architectures for deep reinforcement learning. In *Pro-*

*ceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, pages 1995–2003, 2016.

- Z. Wang, V. Bapst, N. Heess, V. Mnih, R. Munos, K. Kavukcuoglu, and N. d. Freitas. Sample efficient actor-critic with experience replay. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*, 2017.
- Z. Wang, K. Ji, Y. Zhou, Y. Liang, and V. Tarokh. Spiderboost: A class of faster variance-reduced algorithms for nonconvex optimization. *ArXiv*, abs/1810.10690, 2018.
- C. J. C. H. Watkins and P. Dayan. Q-learning. *Machine Learning*, 8(3):279–292, May 1992.
- R. J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3):229–256, May 1992.
- Y. Wu, E. Mansimov, S. Liao, R. Grosse, and J. Ba. Scalable trust-region method for deep reinforcement learning using kronecker-factored approximation. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, pages 5285–5294, USA, 2017.
- P. Xu, F. Gao, and Q. Gu. An improved convergence analysis of stochastic variance-reduced policy gradient. *Conference on Uncertainty in Artificial Intelligence*, 2019a.
- P. Xu, F. Gao, and Q. Gu. Sample efficient policy gradient methods with recursive variance reduction. *ArXiv*, abs/1909.08610, 2019b.
- L. Yang and Y. Zhang. Policy optimization with stochastic mirror descent. *CoRR*, abs/1906.10462, 2019.
- H. Yuan, C. J. Li, Y. Tang, and Y. Zhou. Policy optimization via stochastic recursive gradient algorithm, 2019. URL <https://openreview.net/forum?id=rJ13S2A9t7>.
- C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals. Understanding deep learning requires rethinking generalization. 2017. URL <https://arxiv.org/abs/1611.03530>.
- T. Zhao, H. Hachiya, G. Niu, and M. Sugiyama. Analysis and improvement of policy gradient estimation. In *Advances in Neural Information Processing Systems 24*, pages 262–270. 2011.
- D. Zhou, P. Xu, and Q. Gu. Stochastic nested variance reduction for nonconvex optimization. *ArXiv*, abs/1806.07811, 2018.