Mean-Variance Policy Iteration for Risk-Averse Reinforcement Learning

Shangtong Zhang,[1] **Bo Liu,** [2] **Shimon Whiteson** [1]
[1] University of Oxford
[2] Auburn University
shangtong.zhang@cs.ox.ac.uk, bo.liu@auburn.edu, shimon.whiteson@cs.ox.ac.uk

Abstract

We present a mean-variance policy iteration (MVPI) framework for risk-averse control in a discounted infinite horizon MDP optimizing the variance of a per-step reward random variable. MVPI enjoys *great flexibility* in that any policy evaluation method and risk-neutral control method can be dropped in for risk-averse control off the shelf, in both on- and off-policy settings. This flexibility reduces the gap between risk-neutral control and risk-averse control and is achieved by working on a novel augmented MDP directly. We propose risk-averse TD3 as an example instantiating MVPI, which outperforms vanilla TD3 and many previous risk-averse control methods in challenging Mujoco robot simulation tasks under a risk-aware performance metric. This risk-averse TD3 is the first to introduce deterministic policies and off-policy learning into risk-averse reinforcement learning, both of which are key to the performance boost we show in Mujoco domains.

Introduction

One fundamental task in reinforcement learning (RL, Sutton and Barto 2018) is control, in which we seek a policy that maximizes certain performance metrics. In risk-neutral RL, the performance metric is usually the expectation of some random variable, for example, the expected total (discounted or undiscounted) reward (Puterman 2014; Sutton and Barto 2018). We, however, sometimes want to minimize certain risk measures of that random variable while maximizing its expectation. For example, a portfolio manager usually wants to reduce the risk of a portfolio while maximizing its return. Risk-averse RL is a framework for studying such problems.

Although many real-world applications can potentially benefit from risk-averse RL, e.g., pricing (Wang 2000), healthcare (Parker 2009), portfolio management (Lai et al. 2011), autonomous driving (Maurer et al. 2016), and robotics (Majumdar and Pavone 2020), the development of risk-averse RL largely falls behind risk-neutral RL. Risk-neutral RL methods have enjoyed superhuman performance in many domains, e.g., Go (Silver et al. 2016), protein design (Senior, Jumper, and Hassabis 2018), and StarCraft II (Vinyals et al. 2019), while no human-level performance has been reported for risk-averse RL methods in real-world applications. Risk-neutral RL methods have enjoyed stable off-policy learning (Watkins and Dayan 1992; Maei 2011; Fujimoto, van Hoof, and Meger 2018; Haarnoja et al. 2018), while state-of-the-art risk-averse RL methods, e.g., Xie et al. (2018); Bisi et al. (2020), still require on-policy samples. Risk-neutral RL methods have exploited deep neural network function approximators and distributed training (Mnih et al. 2016; Espeholt et al. 2018), while tabular and linear methods still dominate the experiments of risk-averse RL literature (Tamar, Di Castro, and Mannor 2012; Prashanth and Ghavamzadeh 2013; Xie et al. 2018; Chow et al. 2018). Such a big gap between risk-averse RL and risk-neutral RL gives rise to a natural question: *can we design a meta algorithm that can easily leverage recent advances in risk-neutral RL for risk-averse RL?* In this paper, we give an affirmative answer via the mean-variance policy iteration (MVPI) framework.

Although many risk measures have been used in risk-averse RL, in this paper, we mainly focus on variance (Sobel 1982; Mannor and Tsitsiklis 2011; Tamar, Di Castro, and Mannor 2012; Prashanth and Ghavamzadeh 2013; Xie et al. 2018) given its advantages in interpretability and computation (Markowitz and Todd 2000; Li and Ng 2000). Such an RL paradigm is usually referred to as mean-variance RL, and previous mean-variance RL methods usually consider the variance of the total reward random variable (Tamar, Di Castro, and Mannor 2012; Prashanth and Ghavamzadeh 2013; Xie et al. 2018). Recently, Bisi et al. (2020) propose a reward-volatility risk measure that considers the variance of a per-step reward random variable. Bisi et al. (2020) show that the variance of the per-step reward can better capture the short-term risk than the variance of the total reward and usually leads to smoother trajectories.

In complicated environments with function approximation, pursuing the exact policy that minimizes the variance of the total reward is usually intractable. In practice, all we can hope is to lower the variance of the total reward to a certain level. As the variance of the per-step reward bounds the variance of the total reward from above (Bisi et al. 2020), in this paper, we propose to optimize the variance of the per-step reward as a proxy for optimizing the variance of the total reward. Though the policy minimizing the variance of the per-step reward does not necessarily minimize the variance of the total reward, in this paper, we show, via MVPI, that optimizing the variance of the per-step reward as a proxy is more efficient and scalable than optimizing the variance of the total reward directly.

MVPI enjoys great flexibility in that *any policy evaluation method and risk-neutral control method can be dropped in for risk-averse control off the shelf, in both on- and off-policy settings*. Key to the flexibility of MVPI is that it works on an augmented MDP directly, which we make possible by introducing the Fenchel duality and block cyclic coordinate ascent to solve a policy-dependent reward issue (Bisi et al. 2020). This issue refers to a requirement to solve an MDP whose reward function depends on the policy being followed, i.e., the reward function of this MDP is nonstationary. Consequently, standard tools from the MDP literature are not applicable. We propose risk-averse TD3 as an example instantiating MVPI, which outperforms vanilla TD3 (Fujimoto, van Hoof, and Meger 2018) and many previous mean-variance RL methods (Tamar, Di Castro, and Mannor 2012; Prashanth and Ghavamzadeh 2013; Xie et al. 2018; Bisi et al. 2020) in challenging Mujoco robot simulation tasks in terms of a risk-aware performance metric. To the best of our knowledge, we are the first to benchmark mean-variance RL methods in Mujoco domains, a widely used benchmark for robotic-oriented RL research, and the first to bring off-policy learning and deterministic policies into mean-variance RL.

Mean-Variance RL

We consider an infinite horizon MDP with a state space $\mathcal{S}$, an action space $\mathcal{A}$, a bounded reward function $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, a transition kernel $p : \mathcal{S} \times \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$, an initial distribution $\mu_0 : \mathcal{S} \rightarrow [0, 1]$, and a discount factor $\gamma \in [0, 1]$. The initial state S_0 is sampled from μ_0. At time step t, an agent takes an action A_t according to $\pi(\cdot|S_t)$, where $\pi : \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ is the policy followed by the agent. The agent then gets a reward $R_{t+1} \doteq r(S_t, A_t)$ and proceeds to the next state S_{t+1} according to $p(\cdot|S_t, A_t)$. In this paper, we consider a deterministic reward setting for the ease of presentation, following Chow (2017); Xie et al. (2018). The return at time step t is defined as $G_t \doteq \sum_{i=0}^{\infty} \gamma^i r(S_{t+i}, A_{t+i})$. When $\gamma < 1$, G_t is always well defined. When $\gamma = 1$, to ensure G_t remains well defined, it is usually assumed that all polices are proper (Bertsekas and Tsitsiklis 1996), i.e., for any policy π, the chain induced by π has some absorbing states, one of which the agent will eventually go to with probability 1. Furthermore, the rewards are always 0 thereafter. For any $\gamma \in [0, 1]$, G_0 is the random variable indicating the total reward, and we use its expectation

$$J(\pi) \doteq \mathbb{E}_{\mu_0, p, \pi}[G_0],$$

as our primary performance metric. In particular, when $\gamma = 1$, we can express G_0 as $G_0 = \sum_{t=0}^{T-1} r(S_t, A_t)$, where T is a random variable indicating the first time the agent goes to an absorbing state. For any $\gamma \in [0, 1]$, the state value function and the state-action value function are defined as $v_\pi(s) \doteq \mathbb{E}[G_t|S_t = s]$ and $q_\pi(s, a) \doteq \mathbb{E}[G_t|S_t = s, A_t = a]$ respectively.

Total Reward Perspective. Previous mean-variance RL methods (Prashanth and Ghavamzadeh 2013; Tamar, Di Castro, and Mannor 2012; Xie et al. 2018) usually consider the variance of the total reward. Namely, they consider the following problem:

$$\max_\theta \mathbb{E}[G_0] \quad \text{subject to} \quad \mathbb{V}(G_0) \leq \xi, \tag{1}$$

where $\mathbb{V}(\cdot)$ indicates the variance of a random variable, ξ indicates the user's tolerance for variance, and π is parameterized by θ. In particular, Prashanth and Ghavamzadeh (2013) consider the setting $\gamma < 1$ and convert (1) into an unconstrained saddle-point problem: $\max_\lambda \min_\theta L_1(\theta, \lambda) \doteq -\mathbb{E}[G_0] + \lambda(\mathbb{V}(G_0) - \xi)$, where λ is the dual variable. (Prashanth and Ghavamzadeh 2013) use stochastic gradient descent to find the saddle-point of $L_1(\theta, \lambda)$. To estimate $\nabla_{\theta,\lambda} L_1(\theta, \lambda)$, they propose two simultaneous perturbation methods: simultaneous perturbation stochastic approximation and smoothed functional (Bhatnagar, Prasad, and Prashanth 2013), yielding a three-timescale algorithm. Empirical success is observed in a simple traffic control MDP.

Tamar, Di Castro, and Mannor (2012) consider the setting $\gamma = 1$. Instead of using the saddle-point formulation in Prashanth and Ghavamzadeh (2013), they consider the following unconstrained problem: $\max_\theta L_2(\theta) \doteq \mathbb{E}[G_0] - \lambda g(\mathbb{V}(G_0) - \xi)$, where $\lambda > 0$ is a hyperparameter to be tuned and $g(\cdot)$ is a penalty function, which they define as $g(x) \doteq (\max\{0, x\})^2$. The analytical expression of $\nabla_\theta L_2(\theta)$ they provide involves a term $\mathbb{E}[G_0]\nabla_\theta \mathbb{E}[G_0]$. To estimate this term, Tamar, Di Castro, and Mannor (2012) consider a two-timescale algorithm and keep running estimates for $\mathbb{E}[G_0]$ and $\mathbb{V}[G_0]$ in a faster timescale, yielding an episodic algorithm. Empirical success is observed in a simple portfolio management MDP.

Xie et al. (2018) consider the setting $\gamma = 1$ and set the penalty function $g(\cdot)$ in Tamar, Di Castro, and Mannor (2012) to the identity function. With the Fenchel duality $x^2 = \max_y (2xy - y^2)$, they transform the original problem into $\max_{\theta, y} L_3(\theta, y) \doteq 2y(\mathbb{E}[G_0] + \frac{1}{2\lambda}) - y^2 - \mathbb{E}[G_0^2]$, where y is the dual variable. Xie et al. (2018) then propose a solver based on stochastic coordinate ascent, yielding an episodic algorithm.

Per-Step Reward Perspective. Recently Bisi et al. (2020) propose a reward-volatility risk measure, which is the variance of a per-step reward random variable R. In the setting $\gamma < 1$, it is well known that the expected total discounted reward can be expressed as

$$J(\pi) = \frac{1}{1-\gamma} \sum_{s,a} d_\pi(s, a) r(s, a),$$

where $d_\pi(s, a)$ is the normalized discounted state-action distribution:

$$d_\pi(s, a) \doteq (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \Pr(S_t = s, A_t = a|\mu_0, \pi, p).$$

We now define the per-step reward random variable R, a discrete random variable taking values in the image of r, by defining its probability mass function as $p(R = x) = \sum_{s,a} d_\pi(s, a)\mathbb{I}_{r(s,a)=x}$, where $\mathbb{I}$ is the indicator function. It follows that $\mathbb{E}[R] = (1 - \gamma)J(\pi)$. Bisi et al. (2020) argue that $\mathbb{V}(R)$ can better capture short-term risk than $\mathbb{V}(G_0)$ and optimizing $\mathbb{V}(R)$ usually leads to smoother trajectories than optimizing $\mathbb{V}(G_0)$, among other advantages of this risk measure. Bisi et al. (2020), therefore, consider the following objective:

$$J_\lambda(\pi) \doteq \mathbb{E}[R] - \lambda \mathbb{V}(R). \tag{2}$$

Bisi et al. (2020) show that $J_\lambda(\pi) = \mathbb{E}[R - \lambda(R - \mathbb{E}[R])^2]$, i.e., to optimize the risk-aware objective $J_\lambda(\pi)$ is to optimize the canonical risk-neutral objective of a new MDP, which is the same as the original MDP except that the new reward function is

$$r'(s,a) \doteq r(s,a) - \lambda\big(r(s,a) - (1-\gamma)J(\pi)\big)^2.$$

Unfortunately, this new reward function r' depends on the policy π due to the occurrence of $J(\pi)$, implying the reward function is actually nonstationary. By contrast, in canonical RL settings (e.g., Puterman (2014); Sutton and Barto (2018)), the reward function is assumed to be stationary. We refer to this problem as the *policy-dependent-reward* issue. Due to this issue, the rich classical MDP toolbox cannot be applied to this new MDP easily, and the approach of Bisi et al. (2020) *does not and cannot* work on this new MDP directly.

Bisi et al. (2020) instead work on the objective Eq (2) directly *without* resorting to the augmented MDP. They propose to optimize a performance lower bound of $J_\lambda(\pi)$ by extending the performance difference theorem (Theorem 1 in Schulman et al. (2015)) from the risk-neutral objective $J(\pi)$ to the risk-aware objective $J_\lambda(\pi)$, yielding the Trust Region Volatility Optimization (TRVO) algorithm, which is similar to Trust Region Policy Optimization (Schulman et al. 2015).

Importantly, Bisi et al. (2020) show that $\mathbb{V}(G_0) \leq \frac{\mathbb{V}(R)}{(1-\gamma)^2}$, indicating that minimizing the variance of R implicitly minimizes the variance of G_0. We, therefore, can optimize $\mathbb{V}(R)$ as a proxy (upper bound) for optimizing $\mathbb{V}(G_0)$. In this paper, we argue that $\mathbb{V}(R)$ is easier to optimize than $\mathbb{V}(G_0)$. The methods of Tamar, Di Castro, and Mannor (2012); Xie et al. (2018) optimizing $\mathbb{V}(G_0)$ involve terms like $(\mathbb{E}[G_0])^2$ and $\mathbb{E}[G_0^2]$, which lead to terms like $G_0^2 \sum_{t=0}^{T-1} \nabla_\theta \log \pi(A_t|S_t)$ in their update rules, yielding large variance. In particular, it is computationally prohibitive to further expand G_0^2 explicitly to apply variance reduction techniques like baselines (Williams 1992). By contrast, we show in the next section that by considering $\mathbb{V}(R)$, MVPI involves only $r(s,a)^2$, which is easier to deal with than G_0^2.

Mean-Variance Policy Iteration

Although in many problems our goal is to maximize the expected total undiscounted reward, practitioners often find that optimizing the discounted objective ($\gamma < 1$) as a proxy for the undiscounted objective ($\gamma = 1$) is better than optimizing the undiscounted objective directly, especially when deep neural networks are used as function approximators (Mnih et al. 2015; Lillicrap et al. 2015; Espeholt et al. 2018; Xu, van Hasselt, and Silver 2018; Van Seijen, Fatemi, and Tavakoli 2019). We, therefore, focus on the discounted setting in the paper, which allows us to consider optimizing the variance of the per-step reward as a proxy (upper bound) for optimizing the variance of the total reward.

To address the policy-dependent reward issue, we use the Fenchel duality to rewrite $J_\lambda(\pi)$ as

$$\begin{aligned} J_\lambda(\pi) &= \mathbb{E}[R] - \lambda \mathbb{E}[R^2] + \lambda(\mathbb{E}[R])^2 \\ &= \mathbb{E}[R] - \lambda \mathbb{E}[R^2] + \lambda \max_y \big(2\mathbb{E}[R]y - y^2\big), \end{aligned} \tag{3}$$

yielding the following problem:

$$\begin{aligned} &\max_{\pi, y} J_\lambda(\pi, y) \\ \doteq &\textstyle\sum_{s,a} d_\pi(s,a)\big(r(s,a) - \lambda r(s,a)^2 + 2\lambda r(s,a)y\big) - \lambda y^2. \end{aligned} \tag{4}$$

We then propose a *block cyclic coordinate ascent* (BCCA, Luenberger and Ye 1984; Tseng 2001; Saha and Tewari 2010, 2013; Wright 2015) framework to solve (4), which updates y and π alternatively as shown in Algorithm 1. At

Algorithm 1: Mean-Variance Policy Iteration (MVPI)

for k = *1, …* **do**
  Step 1: $y_{k+1} \doteq (1-\gamma)J(\pi_k)$ `// The exact solution for` $\arg\max_y J_\lambda(\pi_k, y)$
  Step 2:
  $\pi_{k+1} \doteq \arg\max_\pi \Big(\sum_{s,a} d_\pi(s,a)\big(r(s,a) - \lambda r(s,a)^2 + 2\lambda r(s,a) y_{k+1}\big) - \lambda y_{k+1}^2\Big)$
end

the k-th iteration, we first fix π_k and update y_{k+1} (Step 1). As $J_\lambda(\pi_k, y)$ is quadratic in y, y_{k+1} can be computed analytically as $y_{k+1} = \sum_{s,a} d_{\pi_k}(s,a) r(s,a) = (1-\gamma)J(\pi_k)$, i.e., all we need in this step is $J(\pi_k)$, which is exactly the performance metric of the policy π_k. We, therefore, refer to Step 1 as *policy evaluation*. We then fix y_{k+1} and update π_{k+1} (Step 2). Remarkably, Step 2 can be reduced to the following problem:

$$\pi_{k+1} = \arg\max_\pi \textstyle\sum_{s,a} d_\pi(s,a) \hat{r}(s,a; y_{k+1}),$$

where $\hat{r}(s,a;y) \doteq r(s,a) - \lambda r(s,a)^2 + 2\lambda r(s,a)y$. In other words, to compute π_{k+1}, we need to solve a new MDP, which is the same as the original MDP except that the reward function is $\hat{r}$ instead of r. This new reward function $\hat{r}$ does not depend on the policy π, avoiding the policy-dependent-reward issue of Bisi et al. (2020). In this step, a new policy π_{k+1} is computed. An intuitive conjecture is that this step is a *policy improvement* step, and we confirm this with the following proposition:

Proposition 1. *(Monotonic Policy Improvement)* $\forall k, J_\lambda(\pi_{k+1}) \geq J_\lambda(\pi_k)$.

Though the monotonic improvement w.r.t. the objective $J_\lambda(\pi, y)$ in Eq (4) follows directly from standard BCCA theories, Theorem 1 provides the monotonic improvement w.r.t. the objective $J_\lambda(\pi)$ in Eq (3). The proof is provided in the appendix. Given Theorem 1, we can now consider the whole BCCA framework in Algorithm 1 as a policy iteration framework, which we call *mean-variance policy iteration* (MVPI). Let $\{\pi_\theta : \theta \in \Theta\}$ be the function class for policy optimization, we have

Assumption 1. $\{\theta \in \Theta, y \in \mathbb{R} \mid J_\lambda(\theta, y) \geq J_\lambda(\theta_0)\}$ *is compact, where* θ_0 *is the initial parameters.*

Assumption 2. $\sup_{\theta\in\Theta} ||\frac{\partial \log \pi_\theta(a|s)}{\partial\theta_i \partial\theta_j}|| < \infty$, $\sup_{\theta\in\Theta} ||\nabla_\theta \log \pi_\theta(a|s)|| < \infty$, Θ *is open and bounded.*

Proposition 2. *(Convergence of MVPI with function approximation) Under Assumptions 1 & 2, let*

$$y_{k+1} \doteq \arg\max_y J_\lambda(\theta_k, y),$$
$$\theta_{k+1} \doteq \arg\max_{\theta\in\Theta} J_\lambda(\theta, y_{k+1}), \quad k = 0, 1, \dots$$

then $J_\lambda(\theta_{k+1}) \geq J_\lambda(\theta_k)$, $\{J_\lambda(\theta_k)\}_{k=1,\dots}$ *converges, and* $\liminf_k ||\nabla_\theta J_\lambda(\theta_k)|| = 0$.

Remark 1. *Assumption 1 is standard in BCCA literature (e.g., Theorem 4.1 in Tseng (2001)). Assumption 2 is standard in policy optimization literature (e.g., Assumption 4.1 in Papini et al. (2018)). Convergence in the form of* lim inf *also appears in other literature (e.g., Luenberger and Ye (1984); Tseng (2001); Konda (2002); Zhang et al. (2020b)).*

The proof is provided in the appendix. MVPI enjoys great flexibility in that any policy evaluation method and risk-neutral control method can be dropped in off the shelf, which makes it possible to leverage all the advances in risk-neutral RL. MVPI differs from the standard policy iteration (PI, e.g., see Bertsekas and Tsitsiklis (1996); Puterman (2014); Sutton and Barto (2018)) in two key ways: **(1)** policy evaluation in MVPI requires only a scalar performance metric, while standard policy evaluation involves computing the value of all states. **(2)** policy improvement in MVPI considers an augmented reward $\hat{r}$, which is different at each iteration, while standard policy improvement always considers the original reward. Standard PI can be used to solve the policy improvement step in MVPI.

Average Reward Setting

So far we have considered the total reward as the primary performance metric for mean-variance RL. We now show that MVPI can also be used when we consider the average reward as the primary performance metric. Assuming the chain induced by π is ergodic and letting $\bar{d}_\pi(s)$ be its stationary distribution, Filar, Kallenberg, and Lee (1989); Prashanth and Ghavamzadeh (2013) consider the *long-run variance* risk measure $\Lambda(\pi) \doteq \sum_{s,a} \bar{d}_\pi(s,a)\big(r(s,a) - \bar{J}(\pi)\big)^2$ for the average reward setting, where $\bar{d}_\pi(s,a) \doteq \bar{d}_\pi(s)\pi(a|s)$ and $\bar{J}(\pi) = \sum_{s,a} \bar{d}_\pi(s,a) r(s,a)$ is the average reward. We now define a risk-aware objective

$$\begin{aligned}\bar{J}_\lambda(\pi) &\doteq \bar{J}(\pi) - \lambda\Lambda(\pi) \quad (5)\\ &= \max_y \sum_{s,a} \bar{d}_\pi(s,a)\hat{r}(s,a;y) - \lambda y^2,\end{aligned}$$

where we have used the Fenchel duality and BCCA can take over to derive MVPI for the average reward setting as Algorithm 1. It is not a coincidence that the only difference between (4) and (5) is the difference between d_π and $\bar{d}_\pi$. The root cause is that the total discounted reward of an MDP is always equivalent to the average reward of an artificial MDP (up to a constant multiplier), whose transition kernel is $\tilde{p}(s'|s,a) = \gamma p(s'|s,a) + (1-\gamma)\mu_0(s')$ (e.g., see Section 2.4 in Konda (2002) for details).

Off-Policy Learning

Off-policy learning has played a key role in improving data efficiency (Lin 1992; Mnih et al. 2015) and exploration (Osband et al. 2016; Osband, Aslanides, and Cassirer 2018) in risk-neutral control algorithms. Previous mean-variance RL methods, however, consider only the on-policy setting and cannot be easily made off-policy. For example, it is not clear whether perturbation methods for estimating gradients (Prashanth and Ghavamzadeh 2013) can be used off-policy. To reweight terms like $G_0^2 \sum_{t=0}^{T-1} \nabla_\theta \log \pi(A_t|S_t)$ from Tamar, Di Castro, and Mannor (2012); Xie et al. (2018) in the off-policy setting, we would need to compute the product of importance sampling ratios $\Pi_{i=0}^{T-1} \frac{\pi(a_i|s_i)}{\mu(a_i|s_i)}$, where μ is the behavior policy. This product usually suffers from high variance (Precup, Sutton, and Dasgupta 2001; Liu et al. 2018) and requires knowing the behavior policy μ, both of which are practical obstacles in real applications. By contrast, as MVPI works on an augmented MDP directly, any risk-neutral off-policy learning technique can be used for risk-averse off-policy control directly. In this paper, we consider MVPI in both on-line and off-line off-policy settings.

On-line setting. In the on-line off-policy setting, an agent interacts with the environment following a behavior policy μ to collect transitions, which are stored into a replay buffer (Lin 1992) for future reuse. Mujoco robot simulation tasks (Brockman et al. 2016) are common benchmarks for this paradigm (Lillicrap et al. 2015; Haarnoja et al. 2018), and TD3 is a leading algorithm in Mujoco tasks. TD3 is a risk-neutral control algorithm, reducing the over-estimation bias (Hasselt 2010) of DDPG (Lillicrap et al. 2015), which is a neural network implementation of the deterministic policy gradient theorem (Silver et al. 2014). Given the empirical success of TD3, we propose MVPI-TD3 for risk-averse control in this setting. In the policy evaluation step of MVPI-TD3, we set y_{k+1} to the average of the recent K rewards, where K is a hyperparameter to be tuned and we have assumed the policy changes slowly. Theoretically, we should use a weighted average as $d_\pi(s,a)$ is a discounted distribution. Though implementing this weighted average is straightforward, practitioners usually ignore discounting for state visitation in policy gradient methods to improve sample efficiency (Mnih et al. 2016; Schulman et al. 2015, 2017; Bacon, Harb, and Precup 2017). Hence, we do not use the weighted average in MVPI-TD3. In the policy improvement step of MVPI-TD3, we sample a mini-batch of transitions from the replay buffer and perform one TD3 gradient update. The pseudocode of MVPI-TD3 is provided in the appendix.

Off-line setting In the off-line off-policy setting, we are presented with a batch of transitions $\{s_i, a_i, r_i, s_i'\}_{i=1,\dots,K}$ and want to learn a good target policy π for control solely from this batch of transitions. Sometimes those transitions are generated by following a known behavior policy μ. But more commonly, those transitions are generated from multiple unknown behavior policies, which we refer to as the behavior-agnostic off-policy setting (Nachum et al. 2019a). Namely, the state-action pairs (s_i, a_i) are distributed according to some unknown distribution d, which may result from multiple unknown behavior policies. The successor state s_i' is distributed according to $p(\cdot|s_i, a_i)$ and $r_i = r(s, a)$. The degree of off-policyness in this setting is usually larger than the on-line off-policy setting.

In the off-line off-policy setting, the policy evaluation step in MVPI becomes the standard *off-policy evaluation* problem (OPE, Thomas, Theocharous, and Ghavamzadeh (2015); Thomas and Brunskill (2016); Jiang and Li (2016); Liu et al. (2018)), where we want to estimate a scalar performance metric of a policy with off-line samples. One promising approach to OPE is *density ratio learning*, where we use function approximation to learn the density ratio $\frac{d_\pi(s,a)}{d(s,a)}$ directly, which we then use to reweight $r(s, a)$. All off-policy evaluation algorithms can be integrated into MVPI in a plug-and-play manner. In the off-line off-policy setting, the policy improvement step in MVPI becomes the standard *off-policy policy optimization* problem, where we can reweight the canonical on-policy actor-critic (Sutton et al. 2000; Konda 2002) with the density ratio as in Liu et al. (2019) to achieve off-policy policy optimization. Algorithm 2 provides an example of Off-line MVPI.

In the on-line off-policy learning setting, the behavior policy and the target policy are usually closely correlated (e.g., in MVPI-TD3), we, therefore, do not need to learn the density ratio. In the off-line off-policy learning setting, the dataset may come from behavior policies that are arbitrarily different from the target policy. We, therefore, resort to density ratios to account for this discrepancy. Density ratio learning itself is an active research area and is out of scope of this paper. See Hallak and Mannor (2017); Liu et al. (2018); Gelada and Bellemare (2019); Nachum et al. (2019a); Zhang et al. (2020a); Zhang, Liu, and Whiteson (2020); Mousavi et al. (2020) for more details about density ratio learning.

Algorithm 2: Off-line MVPI

Input: A batch of transitions $\{s_i, a_i, r_i, s_i'\}_{i=1,\dots,K}$ and a learning rate α
while *True* **do**
  Learn the density ratio $\rho_\pi(s, a) \doteq \frac{d_\pi(s,a)}{d(s,a)}$ with $\{s_i, a_i, r_i, s_i'\}_{i=1,\dots,K}$
  `// For example, use GradientDICE (Zhang, Liu, and Whiteson 2020)`
  $y \leftarrow \frac{1}{K}\sum_{i=1}^{K} \rho_\pi(s_i, a_i) r_i$
  for $i = 1, \dots, K$ **do**
    $\hat{r}_i \leftarrow r_i - \lambda r_i^2 + 2\lambda r_i y$
    $a_i' \sim \pi(\cdot|s_i')$
  end
  Learn $q_\pi(s, a)$ with $\{s_i, a_i, \hat{r}_i, s_i', a_i'\}_{i=1,\dots,K}$
  `// For example, use TD(0) (Sutton 1988) in` $\mathcal{S} \times \mathcal{A}$
  $\theta \leftarrow \theta + \alpha \rho_\pi(s_i, a_i) \nabla_\theta \log \pi(a_i|s_i) q_\pi(s_i, a_i)$, where i is randomly selected
end

Experiments

All curves in this section are averaged over 10 independent runs with shaded regions indicate standard errors. All implementations are publicly available.[1]

On-line learning setting. In many real-world robot applications, e.g., in a warehouse, it is crucial that the robots' performance be consistent. In such cases, risk-averse RL is an appealing option to train robots. Motivated by this, we benchmark MVPI-TD3 on eight Mujoco robot manipulation tasks from OpenAI gym. Though Mujoco tasks have a stochastic initial state distribution, they are usually equipped with a deterministic transition kernel. To make them more suitable for investigating risk-averse control algorithms, we add a Gaussian noise $\mathcal{N}(0, 0.1^2)$ to every action. As we are not aware of any other off-policy mean-variance RL method, we use several recent on-policy mean-variance RL method as baselines, namely, the methods of Tamar, Di Castro, and Mannor (2012); Prashanth and Ghavamzadeh (2013), MVP (Xie et al. 2018), and TRVO (Bisi et al. 2020). The methods of Tamar, Di Castro, and Mannor (2012); Prashanth and Ghavamzadeh (2013) and MVP are not designed for deep RL settings. To make the comparison fair, we improve those baselines with parallelized actors to stabilize the training of neural networks as in Mnih et al. (2016).[2] TRVO is essentially MVPI with TRPO for the policy improvement. We, therefore, implement TRVO as MVPI with Proximal Policy Optimization (PPO, Schulman et al. 2017) to improve its performance. We also use the vanilla risk-neutral TD3 as a baseline. We use two-hidden-layer neural networks for function approximation.

We run each algorithm for 10^6 steps and evaluate the algorithm every 10^4 steps for 20 episodes. We report the mean of those 20 episodic returns against the training steps in Figure 1. The curves are generated by setting $\lambda = 1$. More details are provided in the appendix. The results show that MVPI-TD3 outperforms all risk-averse baselines in all tested domains (in terms of both final episodic return and learning speed). Moreover, the curves of the methods from the total reward perspective are always flat in all domains with only one exception that MVP achieves a reasonable performance in `Reacher`, though exhaustive hyperparameter tuning is conducted, including λ and ξ. This means they fail to achieve the risk-performance trade-off in our tested domains, indicating that those methods are not able to scale up to Mujoco domains with neural network function

[1] https://github.com/ShangtongZhang/DeepRL

[2] They are on-policy algorithms so we cannot use experience replay.

	$\Delta_{\mathrm{J}}^{\mathrm{TRVO}}$	$\Delta_{\mathrm{mean}}^{\mathrm{TRVO}}$	$\Delta_{\mathrm{variance}}^{\mathrm{TRVO}}$	$\Delta_{\mathrm{SR}}^{\mathrm{TRVO}}$	$\Delta_{\mathrm{J}}^{\mathrm{MVPI}}$	$\Delta_{\mathrm{mean}}^{\mathrm{MVPI}}$	$\Delta_{\mathrm{variance}}^{\mathrm{MVPI}}$	$\Delta_{\mathrm{SR}}^{\mathrm{MVPI}}$
InvertedP.	-321%	-1%	$3 \cdot 10^7$%	-100%	0%	0%	0%	0%
InvertedD.P.	-367%	-17%	365%	-62%	42%	1%	-41%	32%
HalfCheetah	92%	-87%	-92%	-54%	99%	-44%	-98%	285%
Walker2d	-1%	-61%	0%	-61%	88%	-55%	-88%	30%
Swimmer	-10%	-7%	18%	-14%	4%	-2%	-53%	42%
Hopper	57%	-25%	-57%	13%	72%	-3%	-71%	81%
Reacher	-48%	-12%	102%	21%	-1%	-1%	0%	-1%
Ant	97%	-85%	-97%	-18%	96%	-53%	-96%	130%

Table 1: Normalized statistics of TRVO and MVPI-TD3. MVPI is shorthand for MVPI-TD3 in this table. For algo $\in \{\text{MVPI-TD3}, \text{TRVO}, \text{TD3}\}$, we compute the risk-aware performance metric as $\mathrm{J}_{\mathrm{algo}} \doteq \mathrm{mean}_{\mathrm{algo}} - \lambda \mathrm{variance}_{\mathrm{algo}}$ with $\lambda = 1$, where $\mathrm{mean}_{\mathrm{algo}}$ and $\mathrm{variance}_{\mathrm{algo}}$ are mean and variance of the 100 evaluation episodic returns. Then we compute the normalized statistics as $\Delta_{\mathrm{J}}^{\mathrm{algo}} \doteq \frac{\mathrm{J}_{\mathrm{algo}} - \mathrm{J}_{\mathrm{TD3}}}{|\mathrm{J}_{\mathrm{TD3}}|}$, $\Delta_{\mathrm{mean}}^{\mathrm{algo}} \doteq \frac{\mathrm{mean}_{\mathrm{algo}} - \mathrm{mean}_{\mathrm{TD3}}}{|\mathrm{mean}_{\mathrm{TD3}}|}$, $\Delta_{\mathrm{variance}}^{\mathrm{algo}} \doteq \frac{\mathrm{variance}_{\mathrm{algo}} - \mathrm{variance}_{\mathrm{TD3}}}{|\mathrm{variance}_{\mathrm{TD3}}|}$, $\mathrm{SR}_{\mathrm{algo}} \doteq \frac{\mathrm{mean}_{\mathrm{algo}}}{\sqrt{\mathrm{variance}_{\mathrm{algo}}}}$, $\Delta_{\mathrm{SR}}^{\mathrm{algo}} \doteq \frac{\mathrm{SR}_{\mathrm{algo}} - \mathrm{SR}_{\mathrm{TD3}}}{|\mathrm{SR}_{\mathrm{TD3}}|}$ Both MVPI-TD3 and TRVO are trained with $\lambda = 1$. All J_{algo} are averaged over 10 independent runs.

approximation. Those flat curves suggest that perturbation-based gradient estimation in Prashanth and Ghavamzadeh (2013) may not work well with neural networks, and the $G_0^2 \sum_{t=0}^{T-1} \nabla_\theta \log \pi(a_t|s_t)$ term in Tamar, Di Castro, and Mannor (2012) and MVP may suffer from high variance, yielding instability. By contrast, the two algorithms from the per-step reward perspective (MVPI-TD3 and TRVO) do learn a reasonable policy, which experimentally supports our argument that optimizing the variance of the per-step reward is more efficient and scalable than optimizing the variance of the total reward.

As shown in Figure 1, the vanilla risk-neutral TD3 outperforms all risk-averse algorithms (in terms of episodic return). This is expected as it is in general hard for a risk-averse algorithm to outperform its risk-neutral counterpart in terms of a risk-neutral performance metric. We now compare TD3, MVPI-TD3 and TRVO in terms of a risk-aware performance metric. To this end, we test the agent at the end of training for an extra 100 episodes to compute a risk-aware performance metric. We report the normalized statistics in Table 1. The results show that MVPI-TD3 outperforms TD3 in 6 out of 8 tasks in terms of the risk-aware performance metric. Moreover, MVPI-TD3 outperforms TRVO in 7 out of 8 tasks. We also compare the algorithms in terms of the sharp ratio (SR, Sharpe 1966). Although none of the algorithms optimizes SR directly, MVPI-TD3 outperforms TD3 and TRVO in 6 and 7 tasks respectively in terms of SR. This performance boost of MVPI-TD3 over TRVO indeed results from the performance boost of TD3 over PPO, and it is the flexibility of MVPI that makes this off-the-shelf application TD3 in risk-averse RL possible. We also provide versions of Figure 1 and Table 1 with $\lambda = 0.5$ and $\lambda = 2$ in the appendix. The relative performance is the same as $\lambda = 1$. Though we propose to optimize the variance of the per-step reward as a proxy for optimizing the variance of the total reward, we also provide in the appendix curves comparing MVPI-TD3 and TRVO w.r.t. risk measures consisting of both mean and variance of the per-step reward. With those risk measures, MVPI-TD3 generally outperforms TRVO as well.

Off-line learning setting. We consider an infinite horizon MDP (Figure 2). Two actions a_0 and a_1 are available at s_0, and we have $p(s_3|s_0, a_1) = 1, p(s_1|s_0, a_0) = p(s_2|s_0, a_0) = 0.5$. The discount factor is $\gamma = 0.7$ and the agent is initialized at s_0. We consider the objective $J_\lambda(\pi)$ in Eq (3). If $\lambda = 0$, the optimal policy is to choose a_0. If λ is large enough, the optimal policy is to choose a_1. We consider the behavior-agnostic off-policy setting, where the sampling distribution d satisfies $d(s_0, a_0) = d(s_0, a_1) = d(s_1) = d(s_2) = d(s_3) = 0.2$. This sampling distribution may result from multiple unknown behavior policies. Although the representation is tabular, we use a softmax policy. So the problem we consider is *nonlinear* and *nonconvex*. As we are not aware of any other behavior-agnostic off-policy risk-averse RL method, we benchmark only Off-line MVPI (Algorithm 2). Details are provided in the appendix. We report the probability of selecting a_0 against training iterations. As shown in Figure 3, $\pi(a_0|s_0)$ decreases as λ increases, indicating Off-line MVPI copes well with different risk levels. The main challenge in Off-line MVPI rests on learning the density ratio. Scaling up density ratio learning algorithms reliably to more challenging domains like Mujoco is out of the scope of this paper.

Related Work

Both MVPI and Bisi et al. (2020) consider the per-step reward perspective for mean-variance RL. In this work, we mainly use the variance of the per-step reward as a proxy (upper bound) for optimizing the variance of the total reward. Though TRVO in Bisi et al. (2020) is the same as instantiating MVPI with TRPO, the derivation is dramatically different. In particular, it is not clear whether the performance-lower-bound-based derivation for TRVO can be adopted to deterministic policies, off-policy learning, or other policy optimization paradigms, and this is not explored in Bisi et al. (2020). By contrast, MVPI is compatible with *any* existing risk-neural policy optimization technique. Furthermore, MVPI works for both the total discounted reward setting and the average reward setting. It is not clear how the performance lower bound in Bisi et al. (2020), which plays a

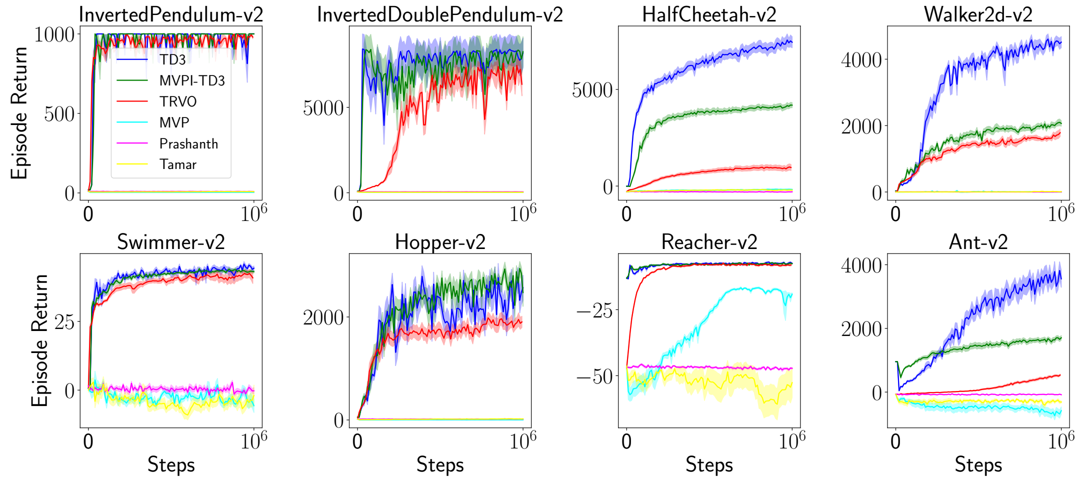

Figure 1: Training progress of MVPI-TD3 and baseline algorithms. Curves are averaged over 10 independent runs with shaded regions indicating standard errors.

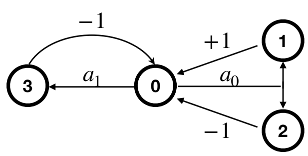

Figure 2: A tabular MDP

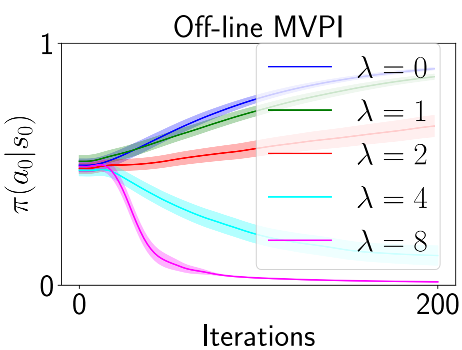

Figure 3: The training progress of Off-line MVPI. Curves are averaged over 30 independent runs with shaded regions indicating standard errors.

central role in TRVO, can be adapted to the average reward setting. All the advantages of MVPI over TRVO result from addressing the policy-dependent-reward issue in Bisi et al. (2020).

While the application of Fenchel duality in RL is not new, previously it has been used only to address double sampling issues (e.g., Liu et al. (2015); Dai et al. (2017); Xie et al. (2018); Nachum et al. (2019a)). By contrast, we use Fenchel duality together with BCAA to address the policy-dependent-reward issue in Bisi et al. (2020) and derive a policy iteration framework that appears to be novel to the RL community.

Besides variance, value at risk (VaR, Chow et al. 2018), conditional value at risk (CVaR, Chow and Ghavamzadeh 2014; Tamar, Glassner, and Mannor 2015; Chow et al. 2018), sharp ratio (Tamar, Di Castro, and Mannor 2012), and exponential utility (Howard and Matheson 1972; Borkar 2002) are also used for risk-averse RL. In particular, it is straightforward to consider exponential utility for the per-step reward, which, however, suffers from the same problems as the exponential utility for the total reward, e.g., it overflows easily (Gosavi, Das, and Murray 2014).

Conclusion

In this paper, we propose MVPI for risk-averse RL. MVPI enjoys great flexibility such that any policy evaluation method and risk-neutral control method can be dropped in for risk-averse control off the shelf, in both on- and off-policy settings. This flexibility dramatically reduces the gap between risk-neutral control and risk-averse control. To the best of our knowledge, MVPI is the first empirical success of risk-averse RL in Mujoco robot simulation domains, and is also the first success of off-policy risk-averse RL and risk-averse RL with deterministic polices. Deterministic policies play an important role in reducing the variance of a policy (Silver et al. 2014). Off-policy learning is important for improving data efficiency (Mnih et al. 2015) and exploration (Osband, Aslanides, and Cassirer 2018). Incorporating those two elements in risk-averse RL appears novel and is key to the observed performance improvement.

Possibilities for future work include considering other risk measures (e.g., VaR and CVaR) of the per-step reward random variable, integrating more advanced off-policy policy optimization techniques (e.g., Nachum et al. 2019b) in off-policy MVPI, optimizing λ with meta-gradients (Xu, van Hasselt, and Silver 2018), analyzing the sample complexity of MVPI, and developing theory for approximate MVPI.

References

Bacon, P.-L.; Harb, J.; and Precup, D. 2017. The option-critic architecture. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*.

Bertsekas, D. P.; and Tsitsiklis, J. N. 1996. *Neuro-Dynamic Programming*. Athena Scientific Belmont, MA.

Bhatnagar, S.; Prasad, H.; and Prashanth, L. 2013. *Stochastic Recursive Algorithms for Optimization*. Springer London.

Bisi, L.; Sabbioni, L.; Vittori, E.; Papini, M.; and Restelli, M. 2020. Risk-Averse Trust Region Optimization for Reward-Volatility Reduction. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence*. Special Track on AI in FinTech.

Borkar, V. S. 2002. Q-learning for risk-sensitive control. *Mathematics of operations research* .

Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; and Zaremba, W. 2016. Openai gym. *arXiv preprint arXiv:1606.01540* .

Chow, Y. 2017. *Risk-Sensitive and Data-Driven Sequential Decision Making*. Ph.D. thesis, Stanford University.

Chow, Y.; and Ghavamzadeh, M. 2014. Algorithms for CVaR optimization in MDPs. In *Advances in Neural Information Processing Systems*.

Chow, Y.; Ghavamzadeh, M.; Janson, L.; and Pavone, M. 2018. Risk-constrained reinforcement learning with percentile risk criteria. *The Journal of Machine Learning Research* .

Dai, B.; Shaw, A.; Li, L.; Xiao, L.; He, N.; Liu, Z.; Chen, J.; and Song, L. 2017. SBEED: Convergent reinforcement learning with nonlinear function approximation. *arXiv preprint arXiv:1712.10285* .

Espeholt, L.; Soyer, H.; Munos, R.; Simonyan, K.; Mnih, V.; Ward, T.; Doron, Y.; Firoiu, V.; Harley, T.; Dunning, I.; et al. 2018. Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures. *arXiv preprint arXiv:1802.01561* .

Filar, J. A.; Kallenberg, L. C.; and Lee, H.-M. 1989. Variance-penalized Markov decision processes. *Mathematics of Operations Research* .

Fujimoto, S.; van Hoof, H.; and Meger, D. 2018. Addressing function approximation error in actor-critic methods. *arXiv preprint arXiv:1802.09477* .

Gelada, C.; and Bellemare, M. G. 2019. Off-Policy Deep Reinforcement Learning by Bootstrapping the Covariate Shift. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*.

Gosavi, A. A.; Das, S. K.; and Murray, S. L. 2014. Beyond exponential utility functions: A variance-adjusted approach for risk-averse reinforcement learning. In *2014 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning*.

Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290* .

Hallak, A.; and Mannor, S. 2017. Consistent on-line off-policy evaluation. In *Proceedings of the 34th International Conference on Machine Learning*.

Hasselt, H. V. 2010. Double Q-learning. In *Advances in neural information processing systems*.

Howard, R. A.; and Matheson, J. E. 1972. Risk-sensitive Markov decision processes. *Management Science* .

Jiang, N.; and Li, L. 2016. Doubly robust off-policy value evaluation for reinforcement learning. In *International Conference on Machine Learning*.

Konda, V. R. 2002. *Actor-critic algorithms*. Ph.D. thesis, Massachusetts Institute of Technology.

Lai, T. L.; Xing, H.; Chen, Z.; et al. 2011. Mean-variance portfolio optimization when means and covariances are unknown. *The Annals of Applied Statistics* .

Li, D.; and Ng, W.-L. 2000. Optimal dynamic portfolio selection: Multiperiod mean-variance formulation. *Mathematical finance* .

Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2015. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971* .

Lin, L.-J. 1992. Self-improving reactive agents based on reinforcement learning, planning and teaching. *Machine Learning* .

Liu, B.; Liu, J.; Ghavamzadeh, M.; Mahadevan, S.; and Petrik, M. 2015. Finite-Sample Analysis of Proximal Gradient TD Algorithms. In *Proceedings of the 31st Conference on Uncertainty in Artificial Intelligence*.

Liu, Q.; Li, L.; Tang, Z.; and Zhou, D. 2018. Breaking the curse of horizon: Infinite-horizon off-policy estimation. In *Advances in Neural Information Processing Systems*.

Liu, Y.; Swaminathan, A.; Agarwal, A.; and Brunskill, E. 2019. Off-Policy Policy Gradient with State Distribution Correction. *arXiv preprint arXiv:1904.08473* .

Luenberger, D. G.; and Ye, Y. 1984. *Linear and nonlinear programming (3rd Edition)*. Springer.

Maei, H. R. 2011. *Gradient temporal-difference learning algorithms*. Ph.D. thesis, University of Alberta.

Majumdar, A.; and Pavone, M. 2020. How should a robot assess risk? Towards an axiomatic theory of risk in robotics. In *Robotics Research*. Springer.

Mannor, S.; and Tsitsiklis, J. 2011. Mean-variance optimization in Markov decision processes. *arXiv preprint arXiv:1104.5601* .

Markowitz, H. M.; and Todd, G. P. 2000. *Mean-variance analysis in portfolio choice and capital markets*. John Wiley & Sons.

Maurer, M.; Christian Gerdes, J.; Lenz, B.; and Winner, H. 2016. *Autonomous driving: technical, legal and social aspects*. Springer Nature.

Mnih, V.; Badia, A. P.; Mirza, M.; Graves, A.; Lillicrap, T.; Harley, T.; Silver, D.; and Kavukcuoglu, K. 2016. Asynchronous methods for deep reinforcement learning. In *Proceedings of the 33rd International Conference on Machine Learning*.

Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *Nature* .

Mousavi, A.; Li, L.; Liu, Q.; and Zhou, D. 2020. Black-box Off-policy Estimation for Infinite-Horizon Reinforcement Learning. In *International Conference on Learning Representations*.

Nachum, O.; Chow, Y.; Dai, B.; and Li, L. 2019a. DualDICE: Behavior-Agnostic Estimation of Discounted Stationary Distribution Corrections. *arXiv preprint arXiv:1906.04733* .

Nachum, O.; Dai, B.; Kostrikov, I.; Chow, Y.; Li, L.; and Schuurmans, D. 2019b. AlgaeDICE: Policy Gradient from Arbitrary Experience. *arXiv preprint arXiv:1912.02074* .

Osband, I.; Aslanides, J.; and Cassirer, A. 2018. Randomized prior functions for deep reinforcement learning. In *Advances in Neural Information Processing Systems*.

Osband, I.; Blundell, C.; Pritzel, A.; and Van Roy, B. 2016. Deep exploration via bootstrapped DQN. In *Advances in Neural Information Processing Systems*.

Papini, M.; Binaghi, D.; Canonaco, G.; Pirotta, M.; and Restelli, M. 2018. Stochastic variance-reduced policy gradient. *arXiv preprint arXiv:1806.05618* .

Parker, D. 2009. Managing risk in healthcare: understanding your safety culture using the Manchester Patient Safety Framework (MaPSaF). *Journal of nursing management* .

Prashanth, L.; and Ghavamzadeh, M. 2013. Actor-critic algorithms for risk-sensitive MDPs. In *Advances in neural information processing systems*.

Precup, D.; Sutton, R. S.; and Dasgupta, S. 2001. Off-policy temporal-difference learning with function approximation. In *Proceedings of the 18th International Conference on Machine Learning*.

Puterman, M. L. 2014. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.

Saha, A.; and Tewari, A. 2010. On the finite time convergence of cyclic coordinate descent methods. *arXiv preprint arXiv:1005.2146* .

Saha, A.; and Tewari, A. 2013. On the nonasymptotic convergence of cyclic coordinate descent methods. *SIAM Journal on Optimization* 23(1): 576–601.

Schulman, J.; Levine, S.; Abbeel, P.; Jordan, M.; and Moritz, P. 2015. Trust region policy optimization. In *Proceedings of the 32nd International Conference on Machine Learning*.

Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* .

Senior, A.; Jumper, J.; and Hassabis, D. 2018. AlphaFold: Using AI for scientific discovery. *DeepMind. Recuperado de: https://deepmind. com/blog/alphafold* .

Sharpe, W. F. 1966. Mutual fund performance. *The Journal of business* .

Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. 2016. Mastering the game of Go with deep neural networks and tree search. *Nature* .

Silver, D.; Lever, G.; Heess, N.; Degris, T.; Wierstra, D.; and Riedmiller, M. 2014. Deterministic policy gradient algorithms. In *Proceedings of the 31st International Conference on Machine Learning*.

Sobel, M. J. 1982. The variance of discounted Markov decision processes. *Journal of Applied Probability* .

Sutton, R. S. 1988. Learning to predict by the methods of temporal differences. *Machine Learning* .

Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement Learning: An Introduction (2nd Edition)*. MIT press.

Sutton, R. S.; McAllester, D. A.; Singh, S. P.; and Mansour, Y. 2000. Policy gradient methods for reinforcement learning with function approximation. In *Advances in Neural Information Processing Systems*.

Tamar, A.; Di Castro, D.; and Mannor, S. 2012. Policy gradients with variance related risk criteria. *arXiv preprint arXiv:1206.6404* .

Tamar, A.; Glassner, Y.; and Mannor, S. 2015. Optimizing the CVaR via sampling. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*.

Thomas, P.; and Brunskill, E. 2016. Data-efficient off-policy policy evaluation for reinforcement learning. In *International Conference on Machine Learning*, 2139–2148.

Thomas, P. S.; Theocharous, G.; and Ghavamzadeh, M. 2015. High-confidence off-policy evaluation. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*.

Tseng, P. 2001. Convergence of a block coordinate descent method for nondifferentiable minimization. *Journal of optimization theory and applications* 109(3): 475–494.

Van Seijen, H.; Fatemi, M.; and Tavakoli, A. 2019. Using a Logarithmic Mapping to Enable Lower Discount Factors in Reinforcement Learning. In *Advances in Neural Information Processing Systems*.

Vinyals, O.; Babuschkin, I.; Czarnecki, W. M.; Mathieu, M.; Dudzik, A.; Chung, J.; Choi, D. H.; Powell, R.; Ewalds, T.; Georgiev, P.; et al. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature* .

Wang, S. S. 2000. A class of distortion operators for pricing financial and insurance risks. *Journal of risk and insurance* .

Watkins, C. J.; and Dayan, P. 1992. Q-learning. *Machine Learning* .

Williams, R. J. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* .

Wright, S. J. 2015. Coordinate descent algorithms. *Mathematical Programming* 151(1): 3–34.

Xie, T.; Liu, B.; Xu, Y.; Ghavamzadeh, M.; Chow, Y.; Lyu, D.; and Yoon, D. 2018. A block coordinate ascent algorithm for mean-variance optimization. In *Advances in Neural Information Processing Systems*.

Xu, Z.; van Hasselt, H. P.; and Silver, D. 2018. Meta-gradient reinforcement learning. In *Advances in neural information processing systems*.

Zhang, R.; Dai, B.; Li, L.; and Schuurmans, D. 2020a. GenDICE: Generalized Offline Estimation of Stationary Values. In *International Conference on Learning Representations*.

Zhang, S.; Liu, B.; and Whiteson, S. 2020. GradientDICE: Rethinking Generalized Offline Estimation of Stationary Values. In *Proceedings of the 37th International Conference on Machine Learning*.

Zhang, S.; Liu, B.; Yao, H.; and Whiteson, S. 2020b. Provably Convergent Two-Timescale Off-Policy Actor-Critic with Function Approximation. In *Proceedings of the 37th International Conference on Machine Learning*.