

Policy Optimization as Online Learning with Mediator Feedback

Alberto Maria Metelli*, Matteo Papini*, Pierluca D’Oro, Marcello Restelli

Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano

Piazza Leonardo da Vinci, 32, 20133, Milano, Italy

{albertomaria.metelli, matteo.papini, marcello.restelli}@polimi.it, pierluca.doro@mail.polimi.it

Abstract

Policy Optimization (PO) is a widely used approach to address continuous control tasks. In this paper, we introduce the notion of *mediator feedback* that frames PO as an online learning problem over the policy space. The additional available information, compared to the standard bandit feedback, allows reusing samples generated by one policy to estimate the performance of other policies. Based on this observation, we propose an algorithm, *RANDOMIZED-exploration policy Optimization via Multiple Importance Sampling with Truncation* (RANDOMIST), for regret minimization in PO, that employs a randomized exploration strategy, differently from the existing optimistic approaches. When the policy space is finite, we show that under certain circumstances, it is possible to achieve constant regret, while always enjoying logarithmic regret. We also derive problem-dependent regret lower bounds. Then, we extend RANDOMIST to compact policy spaces. Finally, we provide numerical simulations on finite and compact policy spaces, in comparison with PO and bandit baselines.

1 Introduction

Policy Optimization (PO, Deisenroth, Neumann, and Peters 2013) is a family of Reinforcement Learning (RL, Sutton and Barto 2018) algorithms based on the explicit optimization of the policy parameters. It represents the most promising approach for learning large-scale continuous control tasks and has already achieved marvelous results in video games (e.g., Vinyals et al. 2019) and robotics (e.g., Peng et al. 2020). These achievements, however, rely on massive amounts of simulation rollouts. The efficient use of experience data is essential both to reduce computational costs and to make learning online from real interaction possible. This is still largely an open problem and calls for better theoretical understanding. Any online-learning agent must face the exploration-exploitation dilemma: whether to leverage on its current knowledge to maximize performance or consider new alternatives. Fortunately, the Multi-Armed Bandit (MAB) literature (Bubeck and Cesa-Bianchi 2012; Latimore and Szepesvári 2018) provides a theoretical framework for the problem of efficient exploration under *bandit*

feedback, i.e., observing the effects of the chosen actions. The dilemma is addressed by minimizing the cumulative *regret* of the online performance w.r.t. the optimal one. The most popular exploration strategies are based on the Optimism in the Face of Uncertainty (OFU, Lai and Robbins 1985), of which UCB1 (Auer, Cesa-Bianchi, and Fischer 2002) is the prototypical algorithm, and on Thompson Sampling (TS, Thompson 1933). Both suffer only sublinear regret (Auer, Cesa-Bianchi, and Fischer 2002; Agrawal and Goyal 2012; Kaufmann, Korda, and Munos 2012). TS typically performs better in practice (Chapelle and Li 2011), but it is only computationally efficient in artificial settings (Kveton et al. 2019b). More recent randomized algorithms such as PHE (Perturbed History Exploration) (Kveton et al. 2019a) are able to match the theoretical and practical advantages of TS without the computational burden, and with no assumptions on the payoff distribution.

The OFU principle has been applied to RL (Jaksch, Ortner, and Auer 2010) and recently also to PO (Chowdhury and Gopalan 2019; Efroni et al. 2020), at the level of action selection. These methods are promising but limited to finite actions. A different perspective is proposed by Papini et al. (2019), where the decision problem is not defined over the agent’s actions but over the policy parameters. This change of viewpoint allows exploiting the special structure of the PO problem: for each policy, a sequence of states and actions performed by the agent is collected, constituting, alongside the rewards, a vastly richer signal than the simple bandit feedback. In this paper, we call it *mediator feedback* since this extra information acts as a mediator variable between the policy parameters and the return. OPTIMIST (Papini et al. 2019) is an OFU algorithm that uses Multiple Importance Sampling (MIS, Veach and Guibas 1995) to exploit the mediator feedback, so that the results of one policy provide information on all the others. This allows, in principle, to optimize over an infinite policy space with only finite samples and no regularity assumptions on the underlying process. There are two important limitations in Papini et al. (2019). First, the advantages of the mediator feedback over the bandit feedback are not clear from a theoretical perspective since the regret of OPTIMIST is comparable with that of UCB1 with finite policy space. Second, the policy selection of OPTIMIST requires maximizing a non-convex and non-differentiable index. In the continuous setting, this

*Equal contribution.

is addressed via discretization, with clear scalability issues.

In this work, we provide two major advancements. From the theoretical side, we provide regret lower bounds for the policy optimization problem with finite policy space, and we show that OPTIMIST actually enjoys *constant* regret under the assumptions made in (Papini et al. 2019). In fact, mediator feedback is so special that, under strong-enough assumptions, a greedy algorithm enjoys the same guarantees. We also devise a PHE-inspired randomized algorithm, called RANDOMIST (RANDOMized-exploration policy Optimization via Multiple Importance Sampling with Truncation), with similar regret guarantees as OPTIMIST. From the practical side, this allows replacing the unfeasible index maximization of OPTIMIST with a sampling procedure. Although our regret guarantees apply to the finite setting only, we propose a heuristic version of RANDOMIST for continuous problems, using a Markov Chain Monte Carlo (MCMC, Owen 2013). We show the advantages of this algorithm over continuous OPTIMIST in terms of computational complexity and performance.

The structure of the paper is as follows. We start in Section 2 with the basic background. In Section 3, we formalize the concept of mediator feedback in PO and derive two regret lower bounds. We illustrate, in Section 4, a possible way to exploit mediator feedback, based on importance sampling. Section 5 is devoted to the discussion of deterministic algorithms, providing the improved regret guarantees for OPTIMIST. In Section 6, we present RANDOMIST with its regret guarantees and the heuristic extension to the continuous case. In Section 7, we compare empirically RANDOMIST with relevant baselines on both illustrative examples and continuous-control problems. In Section 8, we discuss relationships with similar approaches from the bandit and RL literature. We conclude in Section 9, discussing the obtained results and proposing future research directions. The proofs of all the results can be found in Appendix D.

2 Preliminaries

In this section, we introduce some notation, the background on Markov decision processes and policy optimization.

Mathematical Background Let $(\mathcal{X}, \mathcal{F})$ be a measurable space, we denote with $\mathcal{P}(\mathcal{X})$ the set of probability measures over \mathcal{X} . Let $P, Q \in \mathcal{P}(\mathcal{X})$ such that $P \ll Q$,¹ for any $\alpha \in [0, \infty]$ the α -Rényi divergence (Rényi 1961) is defined as:²

$$D_\alpha(P\|Q) = \frac{1}{\alpha - 1} \log \int_{\mathcal{X}} \left(\frac{dP}{dQ} \right)^\alpha dQ.$$

We denote with $d_\alpha(P\|Q) = \exp[D_\alpha(P\|Q)]$ the exponentiated Rényi divergence (Cortes, Mansour, and Mohri 2010).

Markov Decision Processes and Policy Optimization A discrete-time Markov Decision Process (MDP, Puterman

¹ P is absolutely continuous w.r.t. Q , i.e., for every measurable set $\mathcal{Y} \subseteq \mathcal{X}$ we have $Q(\mathcal{Y}) = 0 \Rightarrow P(\mathcal{Y}) = 0$.

²In the limit, for $\alpha \rightarrow 1$ we have $D_1(P\|Q) = D_{\text{KL}}(P\|Q)$ and for $\alpha \rightarrow \infty$ we have $D_\infty(P\|Q) = \text{ess sup}_{\mathcal{X}} \frac{dP}{dQ}$.

1994) is a 6-tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma, \mu)$, where \mathcal{S} is the state space, \mathcal{A} is the action space, \mathcal{P} is the transition model that for each $(s, a) \in \mathcal{S} \times \mathcal{A}$ provides the probability distribution of the next state $\mathcal{P}(\cdot|s, a) \in \mathcal{P}(\mathcal{S})$, $\mathcal{R}(s, a) \in \mathbb{R}$ is the reward function, $\gamma \in [0, 1]$ is the discount factor, and $\mu \in \mathcal{P}(\mathcal{S})$ is the initial-state distribution. In Policy Optimization (PO, Peters and Schaal 2008), we model the agent’s behavior by means of a policy $\pi_\theta(\cdot|s) \in \mathcal{P}(\mathcal{A})$ belonging to a space of parametric policies $\Pi_\Theta = \{\pi_\theta : \theta \in \Theta\}$. The interaction between an agent and an MDP generates a sequence of state-action pairs, named *trajectory*: $\tau = (s_0, a_0, s_1, a_1, \dots, s_{H-1}, a_{H-1})$ where $s_0 \sim \mu$, for all $h \in \{0, \dots, H-1\}$ we have $a_h \sim \pi_\theta(\cdot|s_h)$, $s_{h+1} \sim \mathcal{P}(\cdot|s_h, a_h)$ and $H \in \mathbb{N}$ is the trajectory length. Each parameter $\theta \in \Theta$ determines a policy $\pi_\theta \in \Pi_\Theta$ which, in turn, induces a probability measure $p_\theta \in \mathcal{P}(\mathcal{T})$ over the trajectory space \mathcal{T} . To every trajectory $\tau \in \mathcal{T}$, we associate an index of performance $\mathcal{R}(\tau) = \sum_{h=0}^{H-1} \gamma^h \mathcal{R}(s_h, a_h)$, called *return*. Without loss of generality we assume that $\mathcal{R}(\tau) \in [0, 1]$. Thus, we can evaluate the performance of a policy $\pi_\theta \in \Pi_\Theta$ by means of its *expected return*: $J(\theta) = \mathbb{E}_{\tau \sim p_\theta} [\mathcal{R}(\tau)]$. The goal of the agent consists in finding an optimal parameter, i.e., any θ^* maximizing $J(\theta)$.³

3 Online Policy Optimization and Mediator Feedback

The online PO protocol works as follows. At each round $t \in [n]$, we evaluate a parameter vector $\theta_t \in \Theta$ by running policy π_{θ_t} , collecting one (or more) trajectory $\tau_t \in \mathcal{T}$ and observing the corresponding return $\mathcal{R}(\tau_t)$. Then, based on the history $\mathcal{H}_t = \{(\theta_i, \tau_i, \mathcal{R}(\tau_i))\}_{i=1}^t$, we update θ_t to get θ_{t+1} . From an *online learning* perspective, the goal of the agent consists in maximizing the sum of the expected returns over n rounds or, equivalently, minimizing the cumulative regret $R(n)$:

$$\max_{\theta_1, \dots, \theta_n \in \Theta} \sum_{t=1}^n J(\theta_t) \Leftrightarrow \min_{\theta_1, \dots, \theta_n \in \Theta} R(n) = \sum_{t=1}^n \Delta(\theta_t),$$

where $\Delta(\theta) = J^* - J(\theta)$ is the optimality gap of $\theta \in \Theta$ and $J^* = \sup_{\theta \in \Theta} J(\theta)$. Thus, whenever policy π_{θ_t} is executed the agent receives the trajectory-return pair $(\tau_t, \mathcal{R}(\tau_t))$, that we name *mediator feedback* (MF). The term “mediator” refers to the side information, the trajectory τ_t , that *mediates* between the parameter choice θ_t and the return $\mathcal{R}(\tau_t)$. By naïvely approaching PO as an online-learning problem over policy space, we would only consider *bandit feedback*, in which just the return $\mathcal{R}(\tau_t)$ is observable. In comparison, the MF allows to better exploit the *structure* underlying the PO problem (Figure 1).⁴ Indeed, while the return function \mathcal{R}

³To simplify the presentation, we frame our results for the usual *action-based* PO. Our findings directly extend to *parameter-based* exploration (Sehnke et al. 2008), in which policies are indirectly optimized by learning a hyperpolicy that outputs the policy parameters. Coherently with Papini et al. (2019), the empirical evaluation of Section 7 is carried out in the parameter-based framework.

⁴In this paper, we employ the wording “bandit feedback” with a different meaning compared to some provably efficient approaches to PO (e.g., Efroni et al. 2020). See also Section 8.

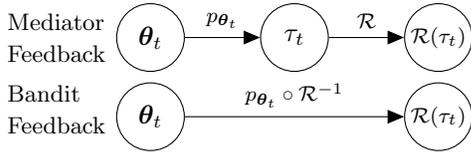


Figure 1: Graphical models comparing mediator and bandit feedbacks.

is unknown, the trajectory distribution p_θ is *partially* known:

$$p_\theta(\tau) = \mu(s_0) \prod_{h=0}^{H-1} \pi_\theta(a_h | s_h) P(s_{h+1} | s_h, a_h). \quad (1)$$

The policy factors π_θ , that depend on θ , are known to the agent, whereas the factors due to the environment (μ and P) are unknown but do not depend on θ . Intuitively, if two policies π_θ and $\pi_{\theta'}$ are sufficiently “similar”, given a trajectory τ from policy π_θ , the return $\mathcal{R}(\tau)$ provides information on the expected return of policy $\pi_{\theta'}$ too.

3.1 Regret Lower Bounds for Finite Policy Space

We focus on the intrinsic complexity of PO with finite policy space, deriving two lower bounds to the regret. The results are phrased, for simplicity, for the case of two policies, i.e., $|\Theta| = 2$, and the proof techniques are inspired to (Bubeck, Perchet, and Rigollet 2013). We start showing that, with enough structure between the policies, i.e., when the KL-divergence between the trajectory distributions is bounded, the best achievable regret is constant.

Theorem 3.1. *There exist an MDP and a parameter space $\Theta = \{\theta_1, \theta_2\}$ with $D_{\text{KL}}(p_{\theta_1} \| p_{\theta_2}) < \infty$, $D_{\text{KL}}(p_{\theta_2} \| p_{\theta_1}) < \infty$ and $J(\theta_1) - J(\theta_2) = \Delta$ such that, for sufficiently large n , all algorithms suffer regret $\mathbb{E} R(n) \geq \frac{1}{32\Delta}$.*

Instead, the presence of policies that are uninformative of one another, i.e., with infinite KL-divergence between the trajectory distributions, leads to a logarithmic regret.

Theorem 3.2. *There exist an MDP and a parameter space $\Theta = \{\theta_1, \theta_2\}$ with $D_{\text{KL}}(p_{\theta_1} \| p_{\theta_2}) = \infty$ or $D_{\text{KL}}(p_{\theta_2} \| p_{\theta_1}) = \infty$, and $J(\theta_1) - J(\theta_2) = \Delta$ such that, for any $n \geq 1$, all algorithms suffer regret $\mathbb{E} R(n) \geq \frac{1}{8\Delta} \log(\Delta^2 n)$.*

4 Exploiting Mediator Feedback with Importance Sampling

In this section, we illustrate how Importance Sampling techniques (IS, Cochran 1977; Owen 2013) can be employed to effectively exploit the mediator feedback in PO.⁵

Monte Carlo Estimation With the bandit feedback at each round $t \in [n]$, the agent has access to the history of parameter-return pairs $\mathcal{H}_t = \{(\theta_i, \mathcal{R}(\tau_i))_{i=1}^{t-1}\}$. Let $T_t(\theta) = \sum_{i=1}^{t-1} \mathbb{1}\{\theta_i = \theta\}$ be the number of trajectories collected with policy $\pi_\theta \in \Pi_\Theta$ up to round $t - 1$. To estimate the

⁵We stress that IS is just *one* method, and not necessarily the best one, to exploit the structure of the PO problem.

expected return $J(\theta)$, if no additional structure is available, we can only use the samples collected when executing π_θ , leading to the Monte Carlo (MC) estimator:

$$\hat{J}_t^{\text{MC}}(\theta) = \frac{1}{T_t(\theta)} \sum_{i=1}^{t-1} \mathcal{R}(\tau_i) \mathbb{1}\{\theta_i = \theta\}. \quad (2)$$

\hat{J}_t^{MC} is unbiased for $J(\theta)$ and its variance scales with $\text{Var}[\hat{J}_t^{\text{MC}}(\theta)] \leq 1/T_t(\theta)$. Clearly, $\hat{J}_t^{\text{MC}}(\theta)$ can be computed only for the policies that have been executed at least once.

Multiple Importance Sampling Estimation With the mediator feedback, at each round $t \in [n]$ we have access to additional information, i.e., the history of parameter-trajectory-return triples $\mathcal{H}_t = \{(\theta_i, \tau_i, \mathcal{R}(\tau_i))_{i=1}^{t-1}\}$. Thanks to the factorization in Equation (1), we can compute the trajectory distribution ratios without knowing P and μ :

$$\frac{p_\theta(\tau)}{p_{\theta'}(\tau)} = \prod_{h=0}^{H-1} \frac{\pi_\theta(a_h | s_h)}{\pi_{\theta'}(a_h | s_h)}.$$

Thus, we can use *all* the samples to estimate the expected return of *any* policy. Let $\Phi_t = \sum_{j=1}^{t-1} \frac{1}{t-1} p_{\theta_j}$ be the mixture induced by the policies executed up to time $t - 1$: if $p_\theta \ll \Phi_t$, we can employ a Multiple Importance Sampling (MIS, Veach and Guibas 1995) estimator (with balance heuristic).⁶

$$\hat{J}_t(\theta) = \frac{1}{t-1} \sum_{i=1}^{t-1} \omega_{\theta,t}(\tau_i) \mathcal{R}(\tau_i), \quad (3)$$

where $\omega_{\theta,t}(\tau_i) = p_\theta(\tau_i)/\Phi_t(\tau_i)$ is the *importance weight*. Thus, for estimating the expected return $J(\theta)$ of policy π_θ we do not need to execute π_θ , but just require the absolute continuity $p_\theta \ll \Phi_t$ (surely fulfilled if $T_t(\theta) \geq 1$). The statistical properties of the MIS estimator can be phrased in terms of the Rényi divergence. We can prove that $0 \leq \hat{J}_t(\theta) \leq d_\infty(p_\theta \| \Phi_t)$ and the variance can be bounded as $\text{Var}[\hat{J}_t(\theta)] \leq d_2(p_\theta \| \Phi_t)/(t-1)$ (Metelli et al. 2018; Papini et al. 2019; Metelli et al. 2020). Since the variance of $\hat{J}_t(\theta)$ scales with $d_2(p_\theta \| \Phi_t)/(t-1)$ instead of $1/T_t(\theta)$, as for $\hat{J}_t^{\text{MC}}(\theta)$, we refer to $\eta_t(\theta) := (t-1)/d_2(p_\theta \| \Phi_t)$ as the *effective number of trajectories*. It is worth noting that $\eta_t(\theta) \geq T_t(\theta)$ (Lemma C.4); thus, thanks to the structure introduced by the mediator feedback, the MIS estimator variance is always smaller than the MC estimator variance.⁷

Truncated Multiple Importance Sampling Estimation

The main limitation of the MIS estimator is that the importance weight $\omega_{\theta,t}$ displays a *heavy-tail* behavior, preventing exponential concentration, unless $d_\infty(p_\theta \| \Phi_t)$ is finite (Metelli et al. 2018). A common solution consists in

⁶For an extensive discussion of importance sampling and heuristics (e.g., balance heuristic) refer to (Owen 2013).

⁷The effective number of trajectories $\eta_t(\theta)$ is, in fact, the *effective sample size* of $\hat{J}_t(\theta)$ (Martino, Elvira, and Louzada 2017).

truncating the estimator (Ionides 2008) at the cost of introducing a negative bias. Given a (time-variant and policy-dependent) truncation threshold $M_t(\boldsymbol{\theta}) < \infty$, the Truncated MIS (TMIS) was introduced by Papini et al. (2019):

$$\check{J}_t(\boldsymbol{\theta}) = \frac{1}{t-1} \sum_{i=1}^{t-1} \check{\omega}_{\boldsymbol{\theta},t}(\tau_i) \mathcal{R}(\tau_i), \quad (4)$$

where $\check{\omega}_{\boldsymbol{\theta},t}(\tau_i) = \min\{M_t(\boldsymbol{\theta}), \omega_{\boldsymbol{\theta},t}(\tau_i)\}$. TMIS enjoys more desirable theoretical properties than plain MIS. While its variance scales similarly to $\hat{J}_t(\boldsymbol{\theta})$ since $\text{Var}[\check{J}_t(\boldsymbol{\theta})] \leq d_2(p_{\boldsymbol{\theta}}\|\Phi_t)/(t-1)$, the range can be bounded as $0 \leq \check{J}_t(\boldsymbol{\theta}) \leq M_t(\boldsymbol{\theta})$. Thus, the range is controlled by $M_t(\boldsymbol{\theta})$ and no longer by the divergence $d_{\infty}(p_{\boldsymbol{\theta}}\|\Phi_t)$, which may be infinite. Similarly, the bias can be bounded as $J(\boldsymbol{\theta}) - \mathbb{E}_{\tau_i \sim p_{\boldsymbol{\theta}_i}}[\check{J}_t(\boldsymbol{\theta})] \leq d_2(p_{\boldsymbol{\theta}}\|\Phi_t)/M_t(\boldsymbol{\theta})$ (see Papini et al. (2019) and Lemma C.1 for details). If we are interested in minimizing the joint contribution of bias and variance, this suggests to increase $M_t(\boldsymbol{\theta})$ progressively over the rounds.

5 Deterministic Algorithms

In this section, we consider finite policy spaces ($|\Theta| < \infty$) and discuss algorithms for PO that select policies deterministically, i.e., $\boldsymbol{\theta}_t$ is a deterministic function of history \mathcal{H}_{t-1} .

Follow The Leader The simplest algorithm accounting for the mediator feedback is Follow The Leader (FTL). It maintains a TMIS estimator $\check{J}_t(\boldsymbol{\theta})$ and selects the policy with the highest estimated expected return, i.e., $\boldsymbol{\theta}_t \in \arg \max_{\boldsymbol{\theta} \in \Theta} \check{J}_t(\boldsymbol{\theta})$. This is a pure-exploitation algorithm, unsuited for bandit feedback. Surprisingly, under a strong form of mediator feedback, FTL enjoys *constant* regret.

Theorem 5.1. *Let $\Theta = [K]$, $v(\boldsymbol{\theta}) = \max_{\boldsymbol{\theta}' \in \Theta} d_2(p_{\boldsymbol{\theta}}\|p_{\boldsymbol{\theta}'})$ for all $\boldsymbol{\theta} \in \Theta$ and $v^*(\boldsymbol{\theta}) = \max\{v(\boldsymbol{\theta}), v(\boldsymbol{\theta}^*)\}$, where $\pi_{\boldsymbol{\theta}^*}$ is an optimal policy. If $v := \max_{\boldsymbol{\theta} \in \Theta} v(\boldsymbol{\theta}) < \infty$, then, for any $\alpha > 1$, the expected regret of FTL using TMIS with truncation $M_t(\boldsymbol{\theta}) = \sqrt{\frac{td_2(p_{\boldsymbol{\theta}}\|\Phi_t)}{\alpha \log t}}$ is bounded as:*

$$\mathbb{E} R(n) \leq \sum_{\boldsymbol{\theta} \in \Theta: \Delta(\boldsymbol{\theta}) > 0} \frac{48\alpha v^*(\boldsymbol{\theta})}{\Delta(\boldsymbol{\theta})} \log \frac{24\alpha v^*(\boldsymbol{\theta})}{\Delta(\boldsymbol{\theta})^2} + \Delta(\boldsymbol{\theta}_1) + \frac{2K}{\alpha-1} \min\left\{1, \sqrt{2 \log v}\right\}. \quad (5)$$

We refer to the condition when all pairwise Rényi divergences are finite (i.e., $v < \infty$) as *perfect mediator feedback*. In such case, we have the remarkable property that running *any* policy in Π_{Θ} provides information for *all* the others. Indeed, the effective number of trajectories satisfies $\eta_t(\boldsymbol{\theta}) \geq (t-1)/v$ (Lemma C.4). Unfortunately, when $v = \infty$, FTL degenerates to *linear* regret (Fact D.1).

UCB1 We can always apply an algorithm for standard bandit feedback, like UCB1 (Lai and Robbins 1985; Auer, Cesa-Bianchi, and Fischer 2002), to PO with finite policy space, ignoring the mediator feedback. UCB1 maintains the sample mean $\hat{J}_t^{\text{MC}}(\boldsymbol{\theta})$ of the observed returns for

Algorithm 1 OPTIMIST

Input: initial parameter $\boldsymbol{\theta}_1, \alpha > 1$

Execute $\pi_{\boldsymbol{\theta}_1}$, observe $\tau_1 \sim p_{\boldsymbol{\theta}_1}$ and $\mathcal{R}(\tau_1)$

for $t = 2, \dots, n$ **do**

 Compute expected return estimate $\check{J}_t(\boldsymbol{\theta})$

 Compute index:

$$B_t(\boldsymbol{\theta}) = \check{J}_t(\boldsymbol{\theta}) + (1 + \sqrt{2}) \sqrt{\frac{\alpha \log t}{\eta_t(\boldsymbol{\theta})}}$$

 Select $\boldsymbol{\theta}_t \in \arg \max_{\boldsymbol{\theta} \in \Theta} B_t(\boldsymbol{\theta})$

 Execute $\pi_{\boldsymbol{\theta}_t}$, observe $\tau_t \sim p_{\boldsymbol{\theta}_t}$ and $\mathcal{R}(\tau_t)$

end for

each $\boldsymbol{\theta} \in \Theta$ and selects the one that maximizes $\hat{J}_t^{\text{MC}}(\boldsymbol{\theta}) + \sqrt{(\alpha \log t)/T_t(\boldsymbol{\theta})}$. The optimistic bonus favors policies that have been selected less often, in accordance with the OFU principle. Being designed for bandit feedback, UCB1 guarantees $\mathcal{O}(\Delta^{-1} \log n)$ regret (Auer, Cesa-Bianchi, and Fischer 2002) even if $v = \infty$, but it cannot exploit mediator feedback when actually present.

In principle, we could employ FTL or UCB1 based on whether v is finite or infinite. There are two reasons why this approach might be inappropriate. First, we would disregard the possibility to share information among pairs of policies with finite divergence, losing possible practical benefits (not captured by the current regret analysis). Second, even when $v < \infty$, the regret of FTL is $\mathcal{O}(v\Delta^{-1} \log(v\Delta^{-2}))$ that, at finite time, might be worse than $\mathcal{O}(\Delta^{-1} \log n)$, especially for large v . Note that deriving the conditions on v so that the regret of UCB1 is smaller than that of FTL is not practical since it would require the knowledge of the gap Δ .

OPTIMIST The difficulty in combining the advantages of FTL and UCB1 is overcome by OPTIMIST (Algorithm 1), an OFU-based algorithm introduced by Papini et al. (2019).⁸ It selects policies as to maximize an *optimistic* TMIS expected return estimate that favors policies with a lower effective number of trajectories. In the original paper (Papini et al. 2019), OPTIMIST is only shown to enjoy sublinear regret in high probability under perfect mediator feedback ($v < \infty$). We show here that OPTIMIST actually enjoys constant regret under perfect mediator feedback (like FTL) without ever degenerating into linear regret (like UCB1).

Theorem 5.2. *Let $\Theta = [K]$ and $v(\boldsymbol{\theta}) = \max_{\boldsymbol{\theta}' \in \Theta} d_2(p_{\boldsymbol{\theta}}\|p_{\boldsymbol{\theta}'})$ for all $\boldsymbol{\theta} \in \Theta$ ($v(\boldsymbol{\theta})$ can be infinite). For any $\alpha > 1$, the expected regret of OPTIMIST with*

truncation $M_t(\boldsymbol{\theta}) = \sqrt{\frac{td_2(p_{\boldsymbol{\theta}}\|\Phi_t)}{\alpha \log t}}$ is bounded as:

(a) *if $v := \max_{\boldsymbol{\theta} \in \Theta} v(\boldsymbol{\theta}) < \infty$:*

$$\mathbb{E} R(n) \leq \sum_{\boldsymbol{\theta} \in \Theta: \Delta(\boldsymbol{\theta}) > 0} \frac{48\alpha v(\boldsymbol{\theta})}{\Delta(\boldsymbol{\theta})} \log \frac{24\alpha v(\boldsymbol{\theta})}{\Delta(\boldsymbol{\theta})^2}$$

⁸We consider here a slight variant of OPTIMIST with an explicit exploration parameter α in place of the original confidence parameter δ from (Papini et al. 2019), since we focus on expected regret rather than high-probability regret.

Algorithm 2 RANDOMIST

Input: initial parameter θ_1 , scale $a \geq 0$, translation $b \geq 0$, $\alpha > 1$

Execute π_{θ_1} , observe $\tau_1 \sim p_{\theta_1}$ and $\mathcal{R}(\tau_1)$

for $t = 2, \dots, n$ **do**

 Compute expected return estimate $\check{J}_t(\theta)$

 Generate perturbation:

$$U_t(\theta) = \frac{1}{\eta_t(\theta)} \sum_{l=1}^{a\eta_t(\theta)} \tau_l + b, \text{ with } \tau_l \sim \text{Ber}(1/2)$$

 Select $\theta_t \in \arg \max_{\theta \in \Theta} \check{J}_t(\theta) + U_t(\theta)$

 Execute π_{θ_t} , observe $\tau_t \sim p_{\theta_t}$ and $\mathcal{R}(\tau_t)$

end for

$$+ \Delta(\theta_1) + \frac{2K}{\alpha - 1} \min \left\{ 1, \sqrt{2 \log v} \right\};$$

(b) in any case:

$$\mathbb{E} R(n) \leq \sum_{\theta \in \Theta: \Delta(\theta) > 0} \frac{24\alpha}{\Delta(\theta)} \log n + \frac{\alpha + 1}{\alpha - 1} K,$$

with an instance-independent expected regret of $\mathbb{E} R(n) \leq 4\sqrt{6\alpha K n \log n} + (\alpha + 1)K/(\alpha - 1)$.

Note also that the regret correctly goes to zero with the divergence (when $v = 1$, all the policies are equivalent). It is an interesting open problem whether better regret guarantees can be provided for the intermediate case, i.e., when some (but not all) the Rényi divergences are finite.

6 Randomized Algorithms

In this section, we propose a novel algorithm for regret minimization in PO that selects the policies with a randomized strategy. RANDOMIST (RANDOMized-exploration policy Optimization via Multiple Importance Sampling with Truncation, Algorithm 2) is based on PHE (Kveton et al. 2019a) and employs additional samples to *perturb* the TMIS expected return estimate $\check{J}_t(\theta)$, enforcing exploration.⁹ Clearly, RANDOMIST shares the randomized nature of exploration with the Bayesian approaches for bandits (e.g., Thompson Sampling (Thompson 1933)) although no prior-posterior mechanism is explicitly implemented and no assumption (apart for boundedness) on the return distribution is needed. At each round $t = 2, \dots, n$, we update the TMIS expected return estimate for each policy $\check{J}_t(\theta)$ and we generate the perturbation $U_t(\theta)$ that is obtained through $a\eta_t(\theta)$ *pseudo-rewards* sampled from a Bernoulli distribution $\text{Ber}(1/2)$. Then, we play the policy maximizing the *perturbed estimated expected return*, i.e., the sum of the estimated expected return $\check{J}_t(\theta)$ and the perturbation $U_t(\theta)$. The two hyperparameters are the *perturbation scale* $a > 0$ and the *perturbation translation* $b > 0$. Informally, a and b are responsible for the amount of exploration: a governs the variance of the perturbation, while b (which is absent in PHE) accounts for the negative bias introduced by the TMIS estimator. We now present the properties of RANDOMIST with finite parameter space and propose an extension to deal with compact parameter spaces.

⁹In this sense, RANDOMIST, as well as PHE, resembles the Follow the Perturbed Leader (Hannan 1957) strategy.

Finite Parameter Space If the policy space is finite, we can show that RANDOMIST enjoys guarantees similar to those of OPTIMIST on the expected regret.

Theorem 6.1. *Let $\Theta = [K]$, $v(\theta) = \max_{x' \in \Theta} d_2(p_\theta \| p_{\theta'})$ for all $\theta \in \Theta$ ($v(\theta)$ can be infinite) and $v^*(\theta) = \max\{v(\theta), v(\theta^*)\}$ where π_{θ^*} is an optimal policy. For any $\alpha > 1$, the expected regret of RANDOMIST with truncation $M_t(\theta) = \sqrt{\frac{td_2(p_\theta \| \Phi_t)}{\alpha \log t}}$ is bounded as follows:*

(a) if $v := \max_{\theta \in \Theta} v(\theta) < \infty$, $b \leq \sqrt{(\alpha \log t)/\eta_t(\theta)}$ and $a \geq 0$:

$$\mathbb{E} R(n) \leq \sum_{\theta \in \Theta: \Delta(\theta) > 0} \frac{(188 + 32a)\alpha v^*(\theta)}{\Delta(\theta)} \log \frac{(94 + 16a)\alpha v^*(\theta)}{\Delta(\theta)^2} + \Delta(\theta_1) + \frac{\alpha + 3}{\alpha - 1} \min \left\{ 1, \sqrt{2 \log v} \right\} K;$$

(b) no matter the value of v , if $a > 8$ and $J(\theta) - \mathbb{E}[\check{J}_t(\theta)] \leq b \leq \sqrt{(\alpha \log t)/\eta_t(\theta)}$:

$$\mathbb{E} R(n) \leq \sum_{\theta \in \Theta: \Delta(\theta) > 0} \frac{(52 + 110a)c\alpha}{\Delta(\theta)} \log n + 2 \frac{\alpha + 1}{\alpha - 1} K,$$

where $c = 2 + \frac{e^2 \sqrt{a}}{\sqrt{2\pi}} \exp\left[\frac{16}{a-8}\right] \left(1 + \sqrt{\frac{\pi a}{a-8}}\right)$, with an instance-independent regret of $\mathbb{E} R(n) \leq 2\sqrt{(52 + 110a)c\alpha K n \log n} + 2 \frac{\alpha + 1}{\alpha - 1} K$.

Under perfect mediator feedback RANDOMIST enjoys constant regret, like OPTIMIST, although with a dependence on $v^*(\theta)$, which involves the divergence w.r.t. an optimal policy. Moreover, in such case, since exploration is not needed, we could even set $a = b = 0$ reducing RANDOMIST to FTL. Similarly to OPTIMIST, when we allow $v = \infty$, the regret becomes logarithmic and the hyperparameters a and b must be carefully set to enforce exploration.

Compact Parameter Space When the parameter space is a compact set, i.e., $\Theta = [-M, M]^d$, the $\arg \max$ in Algorithm 2 cannot be explicitly computed. However, the random variable $\theta \in \arg \max_{\theta' \in \Theta} \check{J}_t(\theta) + U_t(\theta)$ can be seen as sampled from the distribution for θ of being the parameter in Θ with the largest perturbed estimated expected return, whose p.d.f. is given by (D'Eramo et al. 2017):

$$\begin{aligned} \mathbf{g}_t^*(\theta) &= g\left(\check{J}_t(\theta) + U_t(\theta) = \sup_{\theta' \in \Theta} \check{J}_t(\theta') + U_t(\theta') \mid \mathcal{H}_{t-1}\right) \\ &= \int_{\mathbb{R}} \frac{g_\theta(y)}{G_\theta(y)} \prod_{\Theta} G_{\theta'}(y) d\theta' dy, \end{aligned} \quad (6)$$

where $\prod_{\Theta} G_\theta(y) d\theta = \exp\left(\int_{\Theta} \log G_\theta(y) d\theta\right)$ is the *product integral* (Davis and Chatfield 1970), g_θ and G_θ are the p.d.f. and the c.d.f. of the random variable $\check{J}_t(\theta) + U_t(\theta)$ conditioned to the history \mathcal{H}_{t-1} . The *computation* of \mathbf{g}_t^* (even up to a constant) is challenging as the product integral requires a numerical integration over the parameter space Θ . Provided that an approximation (up to a constant) \mathbf{g}_t^\dagger of \mathbf{g}_t^* is available, we can use a Monte Carlo Markov Chain method (Owen

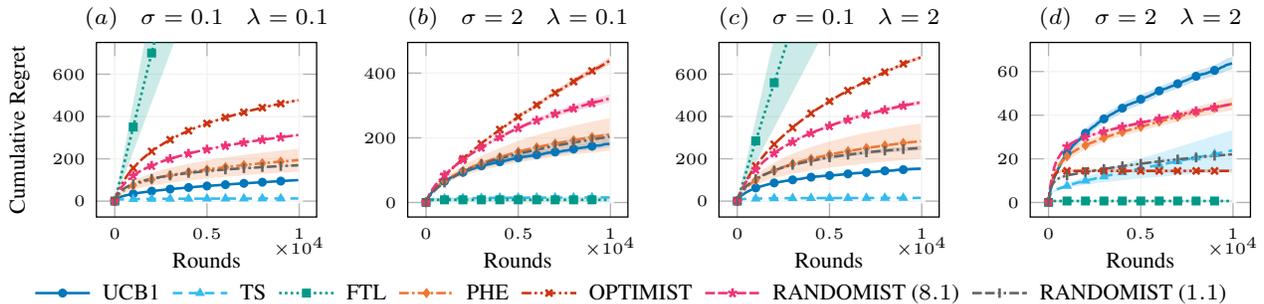


Figure 2: Cumulative regret on the illustrative PO for four values of σ and λ . 20 runs, 95% c.i.

2013) to generate a sample $\theta_t \sim \mathfrak{g}_t^\dagger$. As a practical approximation, we consider the p.d.f. for θ of having a perturbed estimated expected return larger than that of the previously executed policies:¹⁰ $\mathfrak{g}_t^\dagger(\theta) \propto \int_{\mathbb{R}} g_\theta(y) \prod_{i=1}^{t-1} G_{\theta_i}(y) dy$.

Since $\mathcal{O}(d)$ iterations of MCMC are sufficient to generate a sample (Beskos and Stuart 2009), where d is the dimensionality of Θ , and one evaluation of \mathfrak{g}_t^\dagger can be performed in time $\mathcal{O}(t^3)$, the per-round complexity of RANDOMIST is $\mathcal{O}(dt^3)$. This can be further reduced to $\mathcal{O}(dt^2)$ via clever caching (see Appendix F). OPTIMIST (Papini et al. 2019) can also be applied to continuous parameter spaces, with an $\tilde{\mathcal{O}}(\sqrt{vdn})$ high-probability regret bound. However, it is not clear how to perform the maximization step of OPTIMIST efficiently in this setting, since the optimistic index is non-differentiable and non-convex in the parameter variable. Discretization is adopted in (Papini et al. 2019), leading to $\mathcal{O}(t^{1+d/2})$ time complexity, that is exponential in d . The RANDOMIST variant proposed here, although heuristic, has only polynomial dependence on d , thus scaling more favorably to high-dimensional problems.

7 Numerical Simulations

We present the numerical simulations, starting with an illustrative example and then moving to RL benchmarks. For the RL experiments, similarly to Papini et al. (2019), the evaluation is carried out in the parameter-based PO setting (Sehnke et al. 2008), where the policy parameters θ are sampled from a *hyperpolicy* ν_ξ and the optimization is performed in the space of *hyperparameters* Ξ (Appendix A). This setting is particularly convenient since the Rényi divergence between hyperpolicies can be computed exactly (at least for Gaussians). Details and an additional experiment on the Cartpole domain are reported in Appendix F.

Illustrative Problems The goal of this experiment is to show the advantages of the additional structure offered by the mediator feedback over the bandit feedback. We design a class of 5-policy PO problems, isomorphic to bandit problems, in which trajectories are collapsed to a single real action $\mathcal{T} = \mathbb{R}$ and $\mathcal{R}(\tau) = \max\{0, \min\{1, \tau/4\}\}$.

¹⁰ \mathfrak{g}_t^\dagger can be seen as obtained from \mathfrak{g}_t^* applying a quadrature with $\{\theta_1, \dots, \theta_{t-1}\}$ as nodes for the inner integral.

The policies are Gaussians ($\mathcal{N}(0, \sigma^2)$, $\mathcal{N}(1, \sigma^2)$, $\mathcal{N}(2, \sigma^2)$, $\mathcal{N}(2.95, \lambda^2)$, $\mathcal{N}(3, \sigma^2)$) defined in terms of the two values $\sigma, \lambda > 0$. The optimal policy is the fifth one and we have a near-optimal parameter, the fourth, with a different variance. Intuitively, we can tune the parameters σ and λ to vary the Rényi divergences. We compare RANDOMIST with $a = 8.1$ (as prescribed in Theorem 6.1) and $a = 1.1$, and $b = \sqrt{(\alpha \log t)/\eta_t(\theta)}$ for both cases, with OPTIMIST (Papini et al. 2019), FTL, UCB1 (Auer, Cesa-Bianchi, and Fischer 2002), PHE (Kveton et al. 2019a), and TS with Gaussian prior (Agrawal and Goyal 2013a). The cumulative regret is shown in Figure 2 for four combinations of σ and λ . In (a) and (d) we are in a perfect mediator feedback, but in (a) $\log v \simeq 2.25$ and (d) $\log v \simeq 900$. Instead, in (b) or (c), we have $v = \infty$. We notice that FTL displays a (near-)linear regret in (a) as expected since $v = \infty$ but also in (c) where v is finite but very large. RANDOMIST with theoretical value of $a = 8.1$ always displays a good behavior and better than OPTIMIST, except in (d) where the latter shows a remarkable constant regret. We also note that when the amount of information shared among parameters is small, UCB1 performs better than OPTIMIST as well as PHE over RANDOMIST. Furthermore, TS with Gaussian prior performs very well across the tasks, although it considers the bandit feedback. This can be explained since TS assumes the correct return distribution. It also suggests that RANDOMIST could be improved when coped with other perturbation distributions (e.g., Gaussian). Finally, we observe that RANDOMIST with $a = 1.1$, although violating the conditions of Theorem 6.1, keeps showing a sublinear regret even in (b) and (c) when $v = \infty$.

Linear Quadratic Gaussian Regulator The Linear Quadratic Gaussian Regulator (LQG, Curtain 1997) is a benchmark for continuous control. We consider the monodimensional case and a Gaussian hyperpolicy $\nu_\xi = \mathcal{N}(\xi, 0.15^2)$ where ξ is the learned parameter. From ν_ξ , we sample the gain θ of a deterministic linear policy: $a_h = \theta s_h$. This experiment aims at comparing RANDOMIST with UCB1 (Auer, Cesa-Bianchi, and Fischer 2002), GPUCB (Srinivas et al. 2010), and OPTIMIST (Papini et al. 2019) in a finite policy space by discretizing $[-1, 1]$ in $K = 100$ parameters. In Figure 3, we notice that OPTIMIST and RANDOMIST outperform UCB1. While RAN-

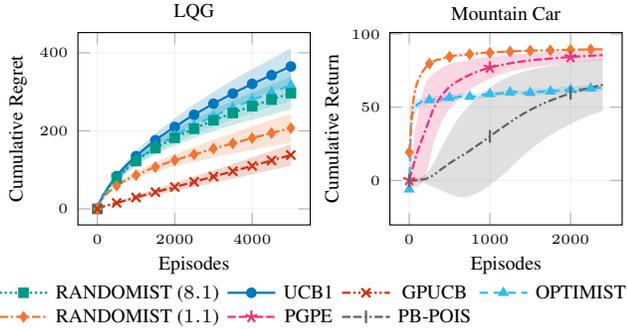


Figure 3: Cumulative regret in the LQG (30 runs, 95% c.i.) and cumulative return in the Mountain Car (5 runs, 95% c.i.).

DOMIST with $a = 8.1$ and OPTIMIST have similar performance, RANDOMIST improves significantly when setting a to 1.1. As in (Papini et al. 2019), the good performance of GPUCB is paired with a lack of theoretical guarantees due to the arbitrary choice of the GP kernel.

Mountain Car To test RANDOMIST in a continuous parameter space, we employ the approximation described above in the Mountain Car environment (Sutton and Barto 2018). We consider the setting of (Papini et al. 2019), employing PGPE (Sehnke et al. 2008) and PB-POIS (Metelli et al. 2018) as baselines. We use a Gaussian hyperpolicy $\nu_{\xi} = \mathcal{N}(\xi, \text{diag}(0.15, 3)^2)$ with learned mean ξ , from which we sample the parameters of a deterministic policy, linear in position and velocity. The exploration phase is performed by sampling from the approximate density g_t^{\dagger} , taking 10 steps of the Metropolis-Hastings algorithm (Owen 2013) with Gaussian proposal $q_m = \mathcal{N}(\theta_m, \text{diag}(0.15, 3)^2)$. Figure 3 shows that RANDOMIST outperforms both policy gradient baselines and OPTIMIST, in terms of learning speed and final performance.

8 Related Works

In this section, we revise the related literature, with attention to bandits with expert advice and to provably efficient PO. Additional comparisons are reported in Appendix B.

Mediator Feedback and Expert Advice A related formulation are the *Bandits with Expert Advice* (BEA, Bubeck and Cesa-Bianchi 2012, Section 4.2), introduced as an approach to adversarial contextual bandits. To draw a parallelism with PO, let \mathcal{T} be the set of arms and $\Theta = [K]$ the finite set of experts. At each step t , the agent receives *advice* $p_{\theta}^t \in \mathcal{P}(\mathcal{T})$ from each expert $\theta \in \Theta$, selects one expert θ_t , and pulls arm $\tau_t \sim p_{\theta_t}^t$. The goal is to minimize the *in-class* regret, competing with the best expert in hindsight. Differently from the trajectory distributions of PO, expert advice can change with time. A major concern of BEA, also relevant to PO, is the dependency of the regret on the number K of experts (resp. policies). A naïve application of Exp3 (Auer et al. 2002)

yields $\mathcal{O}(\sqrt{nK \log K})$ regret. Like our PO algorithms, this is impractical when the experts are exponentially many. Exp4 (Auer et al. 2002) achieves $\mathcal{O}(\sqrt{n|\mathcal{T}| \log K})$ regret, which scales well with K , but is vacuous in the case of infinite arms. McMahan and Streeter (2009) replace $|\mathcal{T}|$ with the *degree of agreement* of the experts, which has interesting similarities with our distributional-divergence approach. *Meta-bandit* approaches (Agarwal et al. 2017; Pacchiano et al. 2020) are so general that could be applied both to continuous-arm BEA and PO, but also exhibit a superlogarithmic dependence on K . Beygelzimer et al. (2011) obtain $\tilde{\mathcal{O}}(\sqrt{dn})$ regret competing with an infinite set of experts of VC-dimension d , mirrored in PO by OPTIMIST on compact spaces of dimension d (Papini et al. 2019, Theorem 3).

Provably Efficient PO Recently, a surge of approaches to deal with PO in a theoretically sound way, with both stochastic or adversarial environments, has emerged. These works consider either *full-information*, i.e., the agent observes the whole reward function $\{\mathcal{R}(s_h, a)\}_{a \in \mathcal{A}}$ regardless the played action (e.g., Rosenberg and Mansour 2019; Cai et al. 2019), or the *bandit feedback* (with a different meaning compared to the use we have made in this paper), in which only the reward of the chosen action is observed $\mathcal{R}(s_h, a_h)$ (e.g., Jin et al. 2019; Efroni et al. 2020). These methods are not directly comparable with the mediator feedback, although both settings exploit the structure of the PO problem. While with MF we explicitly model the policy space Π_{Θ} , these methods search in the space of all Markovian stationary policies. Furthermore, they are limited to tabular MDPs, while MF can deal natively with continuous state-action spaces.

9 Discussion and Conclusions

We have deepened the understanding of policy optimization as an online learning problem with additional feedback. We believe that mediator feedback has potential applications even beyond PO. Indeed, the problem of optimizing over probability distributions also encompasses GANs and variational inference (Chu, Blanchet, and Glynn 2019) and, more generally, MF emerges in any Bayesian network in which we control the conditional distributions on some vertices, via parameters θ , while the other are fixed and independent from θ . Furthermore, we have introduced a novel randomized algorithm, RANDOMIST, and we have shown its advantages both in terms of computational complexity and performance. The algorithm could be improved by adopting a different perturbation, e.g., Gaussian, as already hinted in (Kveton et al. 2019b). Further work is needed to match the theoretical regret lower bounds. Currently, a major discrepancy is the use of the KL-divergence in the lower bounds instead of the larger Rényi divergence required by algorithms based on IS. Moreover, the algorithm employs the ratio importance weight and, thus, it might suffer from the curse of horizon (Liu et al. 2018). Finally, the case of non-perfect mediator feedback could be related to graphical bandits (Alon et al. 2017), where finite Rényi divergences are the edges of a directed feedback graph, in order to capture the actual difficulty of this intermediate case.

Acknowledgments

This work has been partially supported by the Italian MIUR PRIN 2017 Project ALGADIMAR "Algorithms, Games, and Digital Markets".

References

- Abbasi-Yadkori, Y.; Pál, D.; and Szepesvári, C. 2011. Improved algorithms for linear stochastic bandits. In *NeurIPS*.
- Abeille, M.; and Lazaric, A. 2017a. Linear thompson sampling revisited. *Electronic Journal of Statistics* 11(2): 5165–5197.
- Abeille, M.; and Lazaric, A. 2017b. Thompson Sampling for Linear-Quadratic Control Problems. In *AISTATS*.
- Abeille, M.; and Lazaric, A. 2018. Improved regret bounds for thompson sampling in linear quadratic control problems. In *ICML*.
- Agarwal, A.; Luo, H.; Neyshabur, B.; and Schapire, R. E. 2017. Corraling a Band of Bandit Algorithms. In *COLT*.
- Agrawal, S.; and Goyal, N. 2012. Analysis of Thompson Sampling for the Multi-armed Bandit Problem. In *COLT*.
- Agrawal, S.; and Goyal, N. 2013a. Further Optimal Regret Bounds for Thompson Sampling. In *AISTATS*.
- Agrawal, S.; and Goyal, N. 2013b. Thompson sampling for contextual bandits with linear payoffs. In *ICML*.
- Alon, N.; Cesa-Bianchi, N.; Dekel, O.; and Koren, T. 2015. Online learning with feedback graphs: Beyond bandits. In *COLT*.
- Alon, N.; Cesa-Bianchi, N.; Gentile, C.; Mannor, S.; Mansour, Y.; and Shamir, O. 2017. Nonstochastic Multi-Armed Bandits with Graph-Structured Feedback. *SIAM J. Comput.* 46(6): 1785–1826.
- Auer, P.; Cesa-Bianchi, N.; and Fischer, P. 2002. Finite-time Analysis of the Multiarmed Bandit Problem. *Mach. Learn.* 47(2-3): 235–256.
- Auer, P.; Cesa-Bianchi, N.; Freund, Y.; and Schapire, R. E. 2002. The Nonstochastic Multiarmed Bandit Problem. *SIAM J. Comput.* 32(1): 48–77.
- Beskos, A.; and Stuart, A. 2009. Computational complexity of Metropolis-Hastings methods in high dimensions. In *Monte Carlo and Quasi-Monte Carlo Methods 2008*, 61–71.
- Beygelzimer, A.; Langford, J.; Li, L.; Reyzin, L.; and Schapire, R. E. 2011. Contextual Bandit Algorithms with Supervised Learning Guarantees. In *AISTATS*.
- Bubeck, S.; and Cesa-Bianchi, N. 2012. Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems. *Foundations and Trends in Machine Learning* 5(1): 1–122.
- Bubeck, S.; Eldan, R.; and Lehec, J. 2018. Sampling from a log-concave distribution with projected langevin monte carlo. *Discrete & Computational Geometry* 59(4): 757–783.
- Bubeck, S.; Perchet, V.; and Rigollet, P. 2013. Bounded regret in stochastic multi-armed bandits. In *COLT*.
- Cai, Q.; Yang, Z.; Jin, C.; and Wang, Z. 2019. Provably Efficient Exploration in Policy Optimization. *arXiv preprint arXiv:1912.05830*.
- Caron, S.; Kveton, B.; Lelarge, M.; and Bhagat, S. 2012. Leveraging Side Observations in Stochastic Bandits. In *UAI*.
- Casella, G.; and George, E. I. 1992. Explaining the Gibbs sampler. *The American Statistician* 46(3): 167–174.
- Chapelle, O.; and Li, L. 2011. An Empirical Evaluation of Thompson Sampling. In *NeurIPS*.
- Chen, W.; Wang, Y.; Yuan, Y.; and Wang, Q. 2016. Combinatorial multi-armed bandit and its extension to probabilistically triggered arms. *JMLR* 17(1): 1746–1778.
- Cheng, X.; and Bartlett, P. 2018. Convergence of Langevin MCMC in KL-divergence. In *ALT*.
- Chowdhury, S. R.; and Gopalan, A. 2017. On kernelized multi-armed bandits. In *ICML*.
- Chowdhury, S. R.; and Gopalan, A. 2019. Online Learning in Kernelized Markov Decision Processes. In *AISTATS*.
- Chu, C.; Blanchet, J. H.; and Glynn, P. W. 2019. Probability Functional Descent: A Unifying Perspective on GANs, Variational Inference, and Reinforcement Learning. In Chaudhuri, K.; and Salakhutdinov, R., eds., *ICML*.
- Chung, F.; and Lu, L. 2006. Old and new concentration inequalities. *Complex Graphs and Networks* 107: 23–56.
- Cochran, W. G. 1977. *Sampling Techniques, 3rd Edition*. John Wiley. ISBN 0-471-16240-X.
- Combes, R.; Magureanu, S.; and Proutiere, A. 2017. Minimal exploration in structured stochastic bandits. In *NeurIPS*.
- Corless, R. M.; Gonnet, G. H.; Hare, D. E.; Jeffrey, D. J.; and Knuth, D. E. 1996. On the LambertW function. *Advances in Computational mathematics* 5(1): 329–359.
- Cortes, C.; Mansour, Y.; and Mohri, M. 2010. Learning Bounds for Importance Weighting. In *NeurIPS*.
- Curtain, R. F. 1997. Linear-quadratic control: An introduction. *Autom.* 33(5): 1004.
- Dani, V.; Hayes, T. P.; and Kakade, S. M. 2008. Stochastic Linear Optimization under Bandit Feedback. In *COLT*.
- Davis, W.; and Chatfield, J. 1970. Concerning product integrals and exponentials. *AMS*.
- Dean, S.; Mania, H.; Matni, N.; Recht, B.; and Tu, S. 2018. Regret bounds for robust adaptive control of the linear quadratic regulator. In *NeurIPS*.
- Deisenroth, M. P.; Neumann, G.; and Peters, J. 2013. A Survey on Policy Search for Robotics. *Foundations and Trends in Robotics* 2(1-2): 1–142.
- D’Eramo, C.; Nuara, A.; Pirota, M.; and Restelli, M. 2017. Estimating the maximum expected value in continuous reinforcement learning problems. In *AAAI*.
- Eckles, D.; and Kaptein, M. 2014. Thompson sampling with the online bootstrap. *arXiv preprint arXiv:1410.4009*.
- Efroni, Y.; Shani, L.; Rosenberg, A.; and Mannor, S. 2020. Optimistic Policy Optimization with Bandit Feedback. *arXiv preprint arXiv:2002.08243*.
- Gil, M.; Alajaji, F.; and Linder, T. 2013. Rényi divergence measures for commonly used univariate continuous distributions. *Information Sciences* 249: 124–131.
- Grant, J. A.; and Leslie, D. S. 2020. On Thompson Sampling for Smoother-than-Lipschitz Bandits. *arXiv preprint arXiv:2001.02323*.
- Hannan, J. 1957. Approximation to Bayes risk in repeated play. *Contributions to the Theory of Games* 3: 97–139.
- Ionides, E. L. 2008. Truncated importance sampling. *JCGS* 17(2): 295–311.
- Jaksch, T.; Ortner, R.; and Auer, P. 2010. Near-optimal Regret Bounds for Reinforcement Learning. *J. Mach. Learn. Res.* 11: 1563–1600.

- Jin, C.; Allen-Zhu, Z.; Bubeck, S.; and Jordan, M. I. 2018. Is q-learning provably efficient? In *NeurIPS*.
- Jin, C.; Jin, T.; Luo, H.; Sra, S.; and Yu, T. 2019. Learning adversarial markov decision processes with bandit feedback and unknown transition. *arXiv preprint arXiv:1912.01192*.
- Kallus, N. 2018. Instrument-Armed Bandits. In *Algorithmic Learning Theory*, 529–546.
- Kaufmann, E.; Korda, N.; and Munos, R. 2012. Thompson Sampling: An Asymptotically Optimal Finite-Time Analysis. In *ALT*.
- Kveton, B.; Szepesvári, C.; Ghavamzadeh, M.; and Boutilier, C. 2019a. Perturbed-History Exploration in Stochastic Multi-Armed Bandits. In *IJCAI*.
- Kveton, B.; Szepesvári, C.; Vaswani, S.; Wen, Z.; Lattimore, T.; and Ghavamzadeh, M. 2019b. Garbage In, Reward Out: Bootstrapping Exploration in Multi-Armed Bandits. In *ICML*.
- Lai, T. L.; and Robbins, H. 1985. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics* 6(1): 4–22.
- Lattimore, T.; and Munos, R. 2014. Bounded Regret for Finite-Armed Structured Bandits. In *NeurIPS*.
- Lattimore, T.; and Szepesvári, C. 2018. Bandit algorithms .
- Liu, Q.; Li, L.; Tang, Z.; and Zhou, D. 2018. Breaking the Curse of Horizon: Infinite-Horizon Off-Policy Estimation. In *NeurIPS*, 5361–5371.
- Lu, X.; and Van Roy, B. 2017. Ensemble sampling. In *NeurIPS*.
- Magureanu, S.; Combes, R.; and Proutiere, A. 2014. Lipschitz bandits: Regret lower bounds and optimal algorithms. In *COLT*.
- Martino, L.; Elvira, V.; and Louzada, F. 2017. Effective sample size for importance sampling based on discrepancy measures. *Signal Processing* 131: 386–401.
- McMahan, H. B.; and Streeter, M. J. 2009. Tighter Bounds for Multi-Armed Bandits with Expert Advice. In *COLT*.
- Metelli, A. M.; Papini, M.; Faccio, F.; and Restelli, M. 2018. Policy Optimization via Importance Sampling. In *NeurIPS*.
- Metelli, A. M.; Papini, M.; Montali, N.; and Restelli, M. 2020. Importance Sampling Techniques for Policy Optimization. *JMLR* 21(141): 1–75.
- Osband, I.; Russo, D.; and Van Roy, B. 2013. (More) efficient reinforcement learning via posterior sampling. In *NeurIPS*.
- Osband, I.; Russo, D.; Wen, Z.; and Van Roy, B. 2017. Deep exploration via randomized value functions. *JMLR*.
- Owen, A. B. 2013. Monte Carlo theory, methods and examples. *Monte Carlo Theory, Methods and Examples*.
- Pacchiano, A.; Phan, M.; Abbasi-Yadkori, Y.; Rao, A.; Zimmert, J.; Lattimore, T.; and Szepesvári, C. 2020. Model Selection in Contextual Stochastic Bandit Problems. *CoRR* abs/2003.01704.
- Papini, M.; Metelli, A. M.; Lupo, L.; and Restelli, M. 2019. Optimistic Policy Optimization via Multiple Importance Sampling. In *ICML*.
- Peng, X. B.; Coumans, E.; Zhang, T.; Lee, T.-W.; Tan, J.; and Levine, S. 2020. Learning Agile Robotic Locomotion Skills by Imitating Animals. *arXiv preprint arXiv:2004.00784*.
- Peters, J.; and Schaal, S. 2008. Reinforcement learning of motor skills with policy gradients. *Neural Networks* 21(4): 682–697.
- Phan, M.; Yadkori, Y. A.; and Domke, J. 2019. Thompson Sampling and Approximate Inference. In *NeurIPS*.
- Puterman, M. L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley.
- Rényi, A. 1961. On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*.
- Rosenberg, A.; and Mansour, Y. 2019. Online Convex Optimization in Adversarial Markov Decision Processes. In *ICML*.
- Rusmevichientong, P.; and Tsitsiklis, J. N. 2010. Linearly parameterized bandits. *Mathematics of Operations Research* 35(2): 395–411.
- Russo, D. J.; Van Roy, B.; Kazerouni, A.; Osband, I.; and Wen, Z. 2018. A Tutorial on Thompson Sampling. *Foundations and Trends® in Machine Learning* 11(1): 1–96.
- Sehnke, F.; Osendorfer, C.; Rückstieß, T.; Graves, A.; Peters, J.; and Schmidhuber, J. 2008. Policy Gradients with Parameter-Based Exploration for Control. In *ICANN*.
- Srinivas, N.; Krause, A.; Kakade, S. M.; and Seeger, M. W. 2010. Gaussian Process Optimization in the Bandit Setting: No Regret and Experimental Design. In *ICML*.
- Strehl, A. L.; and Littman, M. L. 2008. An analysis of model-based interval estimation for Markov decision processes. *JCSS* 74(8): 1309–1331.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement learning: An introduction*. MIT press.
- Tange, O. 2011. GNU Parallel - The Command-Line Power Tool. *login: The USENIX Magazine* 36(1): 42–47.
- Thompson, W. R. 1933. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika* 25(3/4): 285–294.
- Veach, E.; and Guibas, L. J. 1995. Optimally combining sampling techniques for Monte Carlo rendering. In Mair, S. G.; and Cook, R., eds., *SIGGRAPH*, 419–428. ACM.
- Vinyals, O.; Babuschkin, I.; Czarnecki, W.; Mathieu, M.; Dudzik, A.; Chung, J.; Choi, D.; Powell, R.; Ewalds, T.; Georgiev, P.; Oh, J.; Horgan, D.; Kroiss, M.; Danihelka, I.; Huang, A.; Sifre, L.; Cai, T.; Agapiou, J.; Jaderberg, M.; and Silver, D. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature* 575(7782): 350–354.