Data-Dependent Regret Bounds for Online Portfolio Selection

Chung-En Tsai, Ying-Ting Lin, and Yen-Huan Li
National Taiwan University
chungentsai@ntu.edu.tw

Contributions

- *First data-dependent bounds for OPS and for non-Lipschitz, non-smooth losses.*
- *Novel smoothness characterizations of log-loss.*
- A general analysis of optimistic FTRL with self-concordant regularizers, which are not necessarily barriers.
- *Current theoretically fastest stochastic method for minimizing expected log-loss [1].*

Poster | This work (NeurIPS 2023) | [1] AISTATS 2024

Online Portfolio Selection (OPS)

"The single most iconic online learning problem."

At the t-th round,
1. INVESTOR chooses a portfolio $x_t \in \Delta$;
2. MARKET announces a price relative $a_t \in [0, \infty)^d$;
3. INVESTOR suffers a loss $f_t(x_t) := -\log \langle a_t, x_t \rangle$.

- Goal: minimize $\mathrm{Regret}_T := \sum_{t=1}^T f_t(x_t) - \min_{x \in \Delta} \sum_{t=1}^T f_t(x)$.
- Assumption (does not affect Regret_T): $\|a_t\|_\infty = 1$ for all t.

Challenges

- *Lack of Lipschitz continuity and smoothness.*
- Lipschitz continuity and smoothness are standard assumptions to obtain sub-linear worst-case regret and data-dependent bounds, respectively.

Existing Algorithms

Algorithms	Regret_T Bound ($\tilde{O}$)		Per-round time ($\tilde{O}$)
	Best case	Worst case	
$\widetilde{\mathrm{EG}}$	$d^{1/3} T^{2/3}$		d
BSM, Soft-Bayes, LB-OMD	$\sqrt{dT}$		d
This work	$d \log^2 T$	$\sqrt{dT}$	d
This work	$d \log T$	$\sqrt{dT}$	d
BISONS	$d^2 \log^2 T$		d^3
PAE+DONS	$d^2 \log^5 T$		d^3
VB-FTRL	$d \log T$		$d^2 T$
LB-FTRL without linearized losses	$d \log^{d+1} T$		$d^2 T$
ADA-BARRONS	$d^2 \log^4 T$		$d^{2.5} T$
UPS	$d \log T$		$d^4 T^{14}$

- Minimax regret: $\min_{\text{all algorithms}} \max_{a_1, \dots, a_T} \mathrm{Regret}_T = \Theta(d \log T)$.

Properties of the Log-Loss

Lemma ("Lipschitz continuity" and "Smoothness"). For any $x, y \in \mathrm{ri}\, \Delta$, we have

$$\|\nabla f_t(x)\|_{x,*} \leq 1,$$

$$\|x \odot \nabla f_t(x) - y \odot \nabla f_t(y)\|_2 \leq 4 \|x - y\|_x,$$

$$\min_{\alpha \in \mathbb{R}} \|\nabla f_t(x) + \alpha \mathbf{1}\|_{x,*}^2 \leq 4 \left(f_t(x) - \min_{x \in \Delta} f_t(x) \right),$$

where $\|v\|_{x,*} := \langle v, \nabla^{-2} h(x) v \rangle^{1/2}$ and $h(x) := \sum_{i=1}^d -\log x(i)$.

Tools. Relative smoothness of f_t and self-concordance of h.

The third inequality can be derived from geodesic smoothness of f_t w.r.t. the Poincaré metric $\langle u, v \rangle_x := \langle u, \nabla^2 h(x) v \rangle$ on Δ.

First Data-Dependent Bounds

Theorem. There exist two algorithms that satisfy

$$\mathrm{Regret}_T \leq O\left(d \log^2 T + \sqrt{d L_T^\star \log T} \right) \quad (1)$$

$$\mathrm{Regret}_T \leq O\left(d \log T + \sqrt{d V_T \log T} \right), \quad (2)$$

respectively, where

$$L_T^\star := \min_{x \in \Delta} \sum_{t=1}^T f_t(x), \quad V_T := \sum_{t=2}^T \|\nabla f_t(x_{t-1}) - \nabla f_{t-1}(x_{t-1})\|_{x_{t-1},*}^2.$$

Implicitly Defined Optimistic LB-FTRL

At the t-th round, pick $p_{t+1} \in -\Delta$ and solve

$$\begin{cases} x_{t+1} \odot \hat{g}_{t+1} = p_{t+1}, \\ x_{t+1} \in \mathrm{argmin}_{x \in \Delta} \eta_t \left\langle \sum_{\tau=1}^t \nabla f_\tau(x_\tau) + \hat{g}_{t+1}, x \right\rangle + h(x). \end{cases}$$

Theorem. $(x_{t+1}, \hat{g}_{t+1})$ can be solved in $\tilde{O}(d)$ time.

Examples.

- Eq. (1): $p_{t+1} = 0$ and $\eta_t = O\left(\frac{\sqrt{d}}{\sqrt{\sum_{\tau=1}^t \min_{\alpha \in \mathbb{R}} \|\nabla f_\tau(x_\tau) + \alpha \mathbf{1}\|_{x_\tau,*}^2}} \right)$.
- Eq. (2): $p_{t+1} = x_t \odot \nabla f_t(x_t)$ and $\eta_t = O\left(\sqrt{d} / \sqrt{V_t} \right)$.

Implication for Minimizing Log-Loss

Current theoretically fastest stochastic method [1].

Consider $\min_{x \in \Delta} \{ F(x) := \mathbf{E}_a[-\log \langle a, x \rangle] \}$. Assume the stochastic first-order oracle $\mathcal{O}$ satisfies $\mathbf{E}_\xi \|\mathcal{O}(x; \xi) - \nabla F(x)\|_{x,*}^2 \leq \sigma^2$.

Theorem. There exists a stochastic algorithm satisfying

$$\mathbf{E}\left[F(\bar{x}_T) - \min_{x \in \Delta} F(x) \right] \leq O\left(\frac{d \log^3 T}{T} + \frac{\sigma \sqrt{d} \log T}{\sqrt{T}} \right),$$

with $\tilde{O}(d)$ per-iteration time.

This matches convergence rate of SGD for minimizing smooth functions, regardless of the non-smoothness of the log-loss.