Fast Rates in Stochastic Online Convex Optimization by Exploiting the Curvature of Feasible Sets

Taira Tsuchiya [1,2] Shinji Ito [1,2]

[1]The University of Tokyo [2]RIKEN

Online Convex Optimization (OCO)

for $t = 1, 2, \dots, T$ do
- Learner selects x_t from convex body $K \subset \mathbb{R}^d$ (K: **feasible set**)
- Environment reveals **convex loss function** $f_t \colon K \to \mathbb{R}$ (often bounded & Lipschitz)
- Learner incurs loss $f_t(x_t)$ and observes $\nabla f_t(x_t)$ (or f_t)

Learner's Goal: Minimize the **(pseudo-)regret** $\mathsf{R}_T = \max_{x \in K} \mathbb{E}[\sum_{t=1}^T f_t(x_t) - \sum_{t=1}^T f_t(x)]$.

The **optimal decision** $x_\star$ is defined as $x_\star \in \arg\min_{x \in K} \mathbb{E}[\sum_{t=1}^T f_t(x)]$.

When loss function f_t is a linear function, *i.e.*, $f_t(\cdot) = \langle g_t, \cdot \rangle$ for some $g_t \in \mathbb{R}^d$, this problem is called **online linear optimization (OLO)**.

Lower Bound and Fast Rates for Curved Losses

Online Gradient Descent (OGD), $x_{t+1} \leftarrow \Pi_K(x_t - \eta_t \nabla f_t(x_t))$, achieves $\mathsf{R}_T = O(\sqrt{T})$ for Lipschitz continuous f_t (Zinkevich, 2003).

The $O(\sqrt{T})$ bound cannot be improved in general (Hazan et al., 2007).

However, this lower bound can be circumvented when the loss functions are curved! (Hazan et al., 2007)

Definition (strongly convex and exp-concave functions)

A function $f \colon K \to (-\infty, \infty]$ is α-**strongly convex** (w.r.t. a norm $\|\cdot\|$) if for all $x, y \in K$,

$$f(y) \geq f(x) + \langle \nabla f(x), y - x \rangle + \frac{\alpha}{2}\|x - y\|^2 .$$

A function $f \colon K \to (-\infty, \infty]$ is β-**exp-concave** if $\exp(-\beta f(x))$ is concave.

- OGD with $\eta_t = \Theta(1/t) \to \mathsf{R}_T = O(\frac{1}{\alpha} \log T)$ for α-strongly convex losses
- Online Newton Step (ONS) $\to \mathsf{R}_T = O(\frac{d}{\beta} \log T)$ regret β-exp-concave losses

Q. Any other conditions under which we can circumvent the $\Omega(\sqrt{T})$ lower bound?

Exploiting the Curvature of Feasible Sets

Definition (strongly convex sets)

A convex body K is λ-**strongly convex** w.r.t. a norm $\|\cdot\|$ if

$$\forall x, y \in K, \forall \theta \in [0,1] \quad \theta x + (1-\theta) y + \theta(1-\theta)\frac{\lambda}{2}\|x - y\|^2 \cdot \mathbb{B}_{\|\cdot\|} \subseteq K .$$

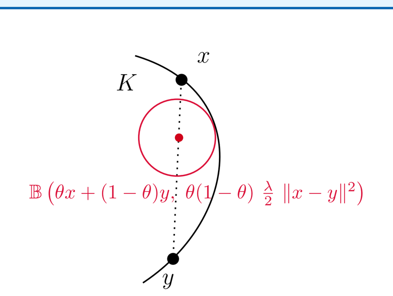

Examples:
- ℓ_p-balls for $p \in (1, 2]$
- Level set $\{x \colon f(x) \leq r\}$ for a strongly convex and smooth function $f \colon \mathbb{R}^d \to \mathbb{R}$

Theorem (Huang, Lattimore, György, and Szepesvári, 2017)

In online **linear** optimization over λ-strongly convex sets, **Follow-the-Leader (FTL)**, $x_t \in \arg\min_{x \in K} \sum_{s=1}^{t-1} \langle g_s, x \rangle$, achieves (for G-Lipschitz losses)

$$\mathsf{R}_T = O\left(\frac{G^2}{\lambda L} \log T \right)$$

if there exists $L > 0$ such that $\|g_1 + \cdots + g_t\|_\star \geq tL$ for all $t \in [T]$ (growth condition).

This upper bound matches their lower bound.

Research Questions

Some of the existing algorithms
1. are only applicable to online **linear** optimization (→ cannot leverage the curvature of loss functions)
2. can suffer a large regret when some ideal conditions (*e.g.*, the growth condition) are not satisfied
3. requires curvature over the entire boundary of the feasible set

Research Questions
1. Can we resolve these three limitations?
2. Are there any other characterizations of feasible sets for which we can achieve fast rates?

Sphere-enclosed Sets: A New Characterization of Feasible Sets

Definition (sphere-enclosed sets)

Let $K \subset \mathbb{R}^d$ be a convex body, $u \in \mathrm{bd}(K)$, and $f \colon K \to \mathbb{R}$. Then, convex body K is (ρ, u, f)-**sphere-enclosed** if there exists a ball $\mathbb{B}(c, \rho)$ with $c \in \mathbb{R}^d$ and $\rho > 0$ satisfying
1. $u \in \mathrm{bd}(\mathbb{B}(c, \rho))$
2. $K \subseteq \mathbb{B}(c, \rho)$
3. there exists $k > 0$ such that $u + k \nabla f(u) = c$

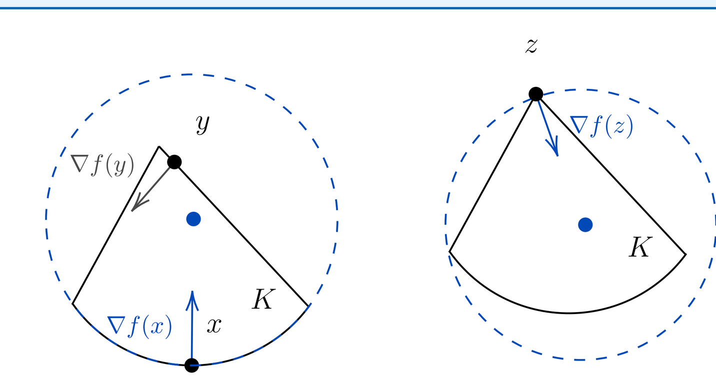

Figure 1. Examples of sphere-enclosed sets

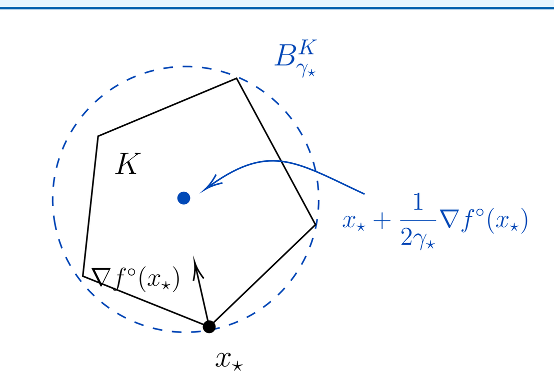

Figure 2. Fig for proof

Main Result (1): Fast Rate over Sphere-enclosed Sets

Stochastic Environment: $f_1, f_2, \dots \sim \mathcal{D}$, $f^\circ = \mathbb{E}_{f \sim \mathcal{D}}[f]$, and $x_\star = \arg\min_{x \in K} f^\circ(x)$

Adversarial Environment: $f_1, f_2, \dots$ are fully adversarial

Theorem

Consider online **convex** optimization. Suppose that K is $(\rho, x_\star, f^\circ)$-sphere-enclosed and that $\nabla f^\circ(x_\star) \neq 0$. Then, there exists an algorithm (MetaGrad or universal online learning algorithm by van Erven and Koolen 2016) such that

$$\mathsf{R}_T = O\left(\frac{G^2 \rho}{\|\nabla f^\circ(x_\star)\|_2} \log T \right) \quad \text{in stochastic environments}$$

and $\mathsf{R}_T = O(GD\sqrt{T})$ in adversarial environments. (D: diam of K, G: Lipschitzness of f_t)

Matches the lower bound in Huang et al. (2017)

Benefits of our bound:
1. Can achieve the $O(\log T)$ regret if the boundary of K is curved around the optimal decision $x_\star$ or $x_\star$ is on corners
2. Can handle convex loss functions and thus the curvature of loss functions (*e.g.*, strong convexity or exp-concavity) can be simultaneously exploited
3. Can achieve $O(\sqrt{T})$ regret even in the worst-case scenarios

Limitation: Achieve fast rates only in stochastic environments
→ Our bounds can be extended to **corrupted stochastic environments** with optimal guarantees! (omitted)

Q. Any other condition for which we can achieve fast rates? →uniformly convex sets!

Proof Sketch

In stochastic environments, the regret is bounded from below by

$$\mathsf{R}_T = \mathbb{E}\left[\sum_{t=1}^T (f^\circ(x_t) - f^\circ(x_\star)) \right] \geq \mathbb{E}\left[\sum_{t=1}^T \langle \nabla f^\circ(x_\star), x_t - x_\star \rangle \right] \quad \text{(convexity of } f^\circ)$$

$$\geq \mathbb{E}\left[\sum_{t=1}^T \gamma_\star \|x_t - x_\star\|_2^2 \right] \quad \text{for some } \gamma_\star > 0 \quad \text{(sphere-enclosedness of } K)$$

The universal online learning algorithms achieve

$$\mathsf{R}_T \lesssim \mathbb{E}\left[\sqrt{ \sum_{t=1}^T \|x_t - x_\star\|_2^2 \log T } \right] .$$

Combining upper and lower bounds of regret and Jensen's inequality,

$$\mathsf{R}_T \lesssim \sqrt{ \mathbb{E}\left[\sum_{t=1}^T \|x_t - x_\star\|_2^2 \right] \log T } - \gamma_\star \mathbb{E}\left[\sum_{t=1}^T \|x_t - x_\star\|_2^2 \right] \underset{ax - bx^2 \leq a^2/(4b)}{\lesssim} \frac{\log T}{\gamma_\star} .$$

[Check $\geq$] Consider a ball facing at $x_\star$ (see the left figure):

$$B_\gamma^K = \mathbb{B}\left(x_\star + \tfrac{1}{2\gamma} \nabla f^\circ(x_\star) , \tfrac{1}{2\gamma} \|\nabla f^\circ(x_\star)\|_2 \right) \subseteq \mathbb{R}^d$$

Observation: $z \in B_\gamma^K$ is equivalent to $\langle \nabla f^\circ(x_\star), z - x_\star \rangle \geq \gamma \|z - x_\star\|_2^2$.

From the $(\rho, x_\star, f^\circ)$-sphere-enclosedness of K, there exists $\gamma > 0$ so that $K \subseteq B_\gamma^K$, and thus

$$\langle \nabla f^\circ(x_\star), x_t - x_\star \rangle \geq \gamma \|x_t - x_\star\|_2^2 .$$

One can set $\gamma_\star$ to $\gamma_\star = \sup\{\gamma \geq 0 \colon K \subseteq B_\gamma^K\}$. □

(Since K is $(\rho, x_\star, f^\circ)$-sphere-enclosing, $\gamma_\star$ satisfies $\gamma_\star < \infty$ and $\frac{1}{2\gamma_\star}\|\nabla f^\circ(x_\star)\| = \rho$.)

Main Result (2): Faster Rates over Uniformly Convex Sets

Definition (uniformly convex sets)

A convex body K is (κ, q)-**uniformly convex w.r.t. a norm** $\|\cdot\|$ (or q-uniformly convex) if

$$\forall x, y \in K, \forall \theta \in [0,1] \quad \theta x + (1-\theta) y + \theta(1-\theta)\frac{\kappa}{2}\|x - y\|^q \cdot \mathbb{B}_{\|\cdot\|} \subseteq K .$$

Examples: ℓ_p-balls for $p \in (1, \infty)$, $(\kappa, 2)$-uniformly convex set is κ-strongly convex

Theorem

Consider online **convex** optimization. Suppose that K is (κ, q)-uniformly convex and that $\nabla f^\circ(x_\star) \neq 0$. Then, there exists an algorithm such that

$$\mathsf{R}_T = O\left(\frac{G^{\frac{q}{q-1}}}{(\kappa \|\nabla f^\circ(x_\star)\|_\star)^{\frac{1}{q-1}}} T^{\frac{q-2}{2(q-1)}} (\log T)^{\frac{q}{2(q-1)}} \right) \quad \text{in stochastic environments}$$

and $\mathsf{R}_T = O(GD\sqrt{T})$ in adversarial environments. (D: diam of K, G: Lipschitzness of f_t)

- Becomes $O(\log T)$ when $q = 2$ and $\widetilde{O}(\sqrt{T})$ s, thus interpolating between the bound over the strongly convex sets and non-curved feasible sets
- Strictly better than the $O\left(T^{\frac{q-2}{q-1}}\right)$ bound, which can be achieved by FTL and becomes smaller than $O(\sqrt{T})$ only when $q \in (2, 3)$, in Kerdreux, d'Aspremont, and Pokutta (2021a).