Open Problem: Parameter-Free and Scale-Free Online Algorithms

Francesco Orabona **Dávid Pál**

Yahoo Research, New York

June 25, 2016

COLT 2016

Online Linear Optimization

Given a convex set $K \subseteq \mathbb{R}^N$

For $t = 1, 2, \ldots$

- predict $w_t \in K$
- receive loss vector $\ell_t \in \mathbb{R}^N$
- suffer loss $\langle \ell_t, w_t \rangle$

$$\text{Regret}_T(u) = \underbrace{\sum_{t=1}^{T} \langle \ell_t, w_t \rangle}_{\text{algorithm's loss}} - \underbrace{\sum_{t=1}^{T} \langle \ell_t, u \rangle}_{\text{competitor's loss}}$$

We focus on:

1. $K = \mathbb{R}^N$
2. $K = \Delta_N = \{x \in \mathbb{R}^N \ : \ x \geq 0, \|x\|_1 = 1\}$

Two Types of Adaptivity

1. Adaptivity to competitor u (parameter-free, quantile bounds, ...)
2. Adaptivity to scale of $\ell_1, \ell_2, \ldots, \ell_T$ (scale-free, second-order bounds, ...)

Open Problem (Informal)

Design efficient **doubly adaptive** algorithms.

FTRL Bound

Theorem (CBL'06, SS'11)

If $R : K \to \mathbb{R}$ *is a non-negative* 1*-strongly convex function w.r.t.* $\|\cdot\|$*, then FTRL with regularizer* R *and learning rate* $\eta > 0$ *satisfies*

$$\forall u \in K \qquad \mathrm{Regret}_T(u) \le \frac{R(u)}{\eta} + \eta \sum_{t=1}^{T} \|\ell_t\|_*^2$$

With learning rate $\eta = \sqrt{R(u) / \sum_{t=1}^{T} \|\ell_t\|_*^2}$

$$\mathrm{Regret}_T(u) \le \sqrt{R(u) \sum_{t=1}^{T} \|\ell_t\|_*^2}$$

Two cheats

1. Bound holds only for **fixed** u
2. Need to know $\sum_{t=1}^{T} \|\ell_t\|_*^2$

Existing Results for $K = \Delta_N$

Regularizer $R(u) = D(u\|\pi)$ $\sup_{u \in \Delta_N} R(u) = \max_i \ln\left(\frac{1}{\pi_i}\right)$

1. For **any** $\ell_1, \ell_2, \ldots, \ell_T \in \mathbb{R}^N$ [deREGK'11, OP'15]

$$\text{Regret}_T(u) \le \sqrt{\max_i \ln\left(\frac{1}{\pi_i}\right) \sum_{t=1}^{T} \|\ell_t\|_\infty^2}$$

2. Assuming that $\|\ell_t\|_\infty \le 1$ [CFH'09, CV'10, LS'14, LS'15, KE'15, FRS'15, OP'16]

$$\text{Regret}_T(u) \le \sqrt{T\left(1 + D(u\|\pi)\right)}$$

3. For **any** $\ell_1, \ell_2, \ldots, \ell_T \in \mathbb{R}^N$ [FRS'15+OP'15]

$$\text{Regret}_T(u) \le \sqrt{\left(1 + D(u\|\pi)\right) \sum_{t=1}^{T} \|\ell_t\|_\infty^2}$$

$O(N \max_i \log\log \frac{1}{\pi_i})$ *memory and time per round*

Existing Results for $K = \mathbb{R}^N$

Regularizer $R(u) = \frac{1}{2}\|u\|_2^2$ $\qquad \sup_{u \in \mathbb{R}^N} R(u) = +\infty$

1. For **any** $\ell_1, \ell_2, \ldots, \ell_T \in \mathbb{R}^N$ [OP'15]

$$\text{Regret}_T(u) \le \left(1 + \|u\|_2^2\right) \sqrt{\sum_{t=1}^{T} \|\ell_t\|_2^2} + \sqrt{T} \max_{1 \le t \le T} \|\ell_t\|_2$$

2. Assuming that $\|\ell_t\|_2 \le 1$ [SM'12, O'13, MA'13, MO'14, O'14, OP'16]

$$\text{Regret}_T(u) \le \left(1 + \|u\|_2^1\right) \sqrt{T \log(1 + \|u\|_2)}$$

Open Problems

Open Problem #1

Find an algorithm for $K = \Delta_N$ with $O(N)$ per-round time complexity such that for any $\pi \in \Delta_N$ and any $\ell_1, \ell_2, \ldots, \ell_T$

$$\forall u \in \Delta_N \qquad \text{Regret}_T(u) \leq \sqrt{(1 + D(u \| \pi)) \sum_{t=1}^{T} \|\ell_t\|_\infty^2}$$

Open Problem #2 — Reward $100 for positive solution

Find an algorithm for $K = \mathbb{R}^N$ with $O(N)$ per-round time complexity such that for any $\ell_1, \ell_2, \ldots, \ell_T$

$$\forall u \in \mathbb{R}^N \quad \text{Regret}_T(u) \leq (1 + \|u\|_2 \cdot \text{polylog}(1 + \|u\|_2)) \sqrt{\sum_{t=1}^{T} \|\ell_t\|_2^2}$$