Lecture 8. Adaptive Online Convex Optimization

Advanced Optimization (Fall 2023)

Peng Zhao

zhaop@lamda.nju.edu.cn

Nanjing University

Outline

- Motivation

- Small-Loss Bounds
  - Small-Loss bound for PEA
  - Self-confident Tuning
  - Small-Loss bound for OCO

General Regret Analysis for OMD

OMD update:

$$\mathbf{x}_{t+1} = \arg\min_{\mathbf{x}\in\mathcal{X}} \left\{\eta_t \langle \mathbf{x}, \nabla f_t(\mathbf{x}_t)\rangle + \mathcal{D}_\psi(\mathbf{x}, \mathbf{x}_t)\right\}$$

Lemma 1 (Mirror Descent Lemma). *Let $\mathcal{D}_\psi$ be the Bregman divergence w.r.t. $\psi : \mathcal{X} \to \mathbb{R}$ and assume ψ to be λ-strongly convex with respect to a norm $\|\cdot\|$. Then, $\forall \mathbf{u} \in \mathcal{X}$, the following inequality holds*

$$f_t(\mathbf{x}_t) - f_t(\mathbf{u}) \le \frac{1}{\eta_t}\left(\mathcal{D}_\psi(\mathbf{u}, \mathbf{x}_t) - \mathcal{D}_\psi(\mathbf{u}, \mathbf{x}_{t+1})\right) + \frac{\eta_t}{\lambda}\|\nabla f_t(\mathbf{x}_t)\|_\star^2 - \frac{1}{\eta_t}\mathcal{D}_\psi(\mathbf{x}_{t+1}, \mathbf{x}_t)$$

bias term (range term) — *variance term (stability term)* — *negative term*

Proof of Mirror Descent Lemma

Proof. $f_t(\mathbf{x}_t) - f_t(\mathbf{u}) \le \underbrace{\langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{x}_{t+1} \rangle}_{\text{term (a)}} + \underbrace{\langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_{t+1} - \mathbf{u} \rangle}_{\text{term (b)}}$

Lemma 2 (Stability Lemma).

$$\lambda \|\mathbf{x}_1 - \mathbf{x}_2\| \le \|\mathbf{g}_1 - \mathbf{g}_2\|_\star$$

$\Rightarrow$ term (a) $= \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{x}_{t+1} \rangle \le \frac{\eta_t}{\lambda} \|\nabla f_t(\mathbf{x}_t)\|_\star^2$ (think of two updates: one for $\mathbf{x}_{t+1}$ with $\nabla f_t(\mathbf{x}_t)$ and another one for $\mathbf{x}_t$ with $\mathbf{0}$)

Lemma 3 (Bregman Proximal Inequality).

$$\langle \mathbf{g}_t, \mathbf{x}_{t+1} - \mathbf{u} \rangle \le \mathcal{D}_\psi(\mathbf{u}, \mathbf{x}_t) - \mathcal{D}_\psi(\mathbf{u}, \mathbf{x}_{t+1}) - \mathcal{D}_\psi(\mathbf{x}_{t+1}, \mathbf{x}_t)$$

$\Rightarrow$ term (b) $\le \frac{1}{\eta_t}\Big(\mathcal{D}_\psi(\mathbf{u}, \mathbf{x}_t) - \mathcal{D}_\psi(\mathbf{u}, \mathbf{x}_{t+1}) - \mathcal{D}_\psi(\mathbf{x}_{t+1}, \mathbf{x}_t)\Big)$ (negative term, usually dropped; but sometimes highly useful)

$\Rightarrow$ $f_t(\mathbf{x}_t) - f_t(\mathbf{u}) \le \frac{1}{\eta_t}\big(\mathcal{D}_\psi(\mathbf{u}, \mathbf{x}_t) - \mathcal{D}_\psi(\mathbf{u}, \mathbf{x}_{t+1})\big) + \frac{\eta_t}{\lambda}\|\nabla f_t(\mathbf{x}_t)\|_\star^2 - \frac{1}{\eta_t}\mathcal{D}_\psi(\mathbf{x}_{t+1}, \mathbf{x}_t)$ □

General Analysis Framework for OMD

Lemma 1 (Mirror Descent Lemma). *Let $\mathcal{D}_\psi$ be the Bregman divergence w.r.t. $\psi : \mathcal{X} \to \mathbb{R}$ and assume ψ to be λ-strongly convex with respect to a norm $\|\cdot\|$. Then, $\forall \mathbf{u} \in \mathcal{X}$, the following inequality holds*

$$f_t(\mathbf{x}_t) - f_t(\mathbf{u}) \leq \frac{1}{\eta_t}\big(\mathcal{D}_\psi(\mathbf{u}, \mathbf{x}_t) - \mathcal{D}_\psi(\mathbf{u}, \mathbf{x}_{t+1})\big) + \frac{\eta_t}{\lambda} \|\nabla f_t(\mathbf{x}_t)\|_\star^2$$

Using Lemma 1, we can easily prove the following regret bound for OMD.

Theorem 4 (General Regret Bound for OMD). *Assume ψ is λ-strongly convex w.r.t. $\|\cdot\|$ and $\eta_t = \eta, \forall t \in [T]$. Then, for all $\mathbf{u} \in \mathcal{X}$, the following regret bound holds*

$$\sum_{t=1}^{T} f_t(\mathbf{x}_t) - \sum_{t=1}^{T} f_t(\mathbf{u}) \leq \frac{\mathcal{D}_\psi(\mathbf{u}, \mathbf{x}_1)}{\eta} + \frac{\eta}{\lambda} \sum_{t=1}^{T} \|\nabla f_t(\mathbf{x}_t)\|_\star^2$$

Online Mirror Descent

- Our previous mentioned algorithms can all be covered by OMD.

Algo.	OMD/proximal form	$\psi(\cdot)$	η_t	Regret_T
OGD for convex	$\mathbf{x}_{t+1} = \arg\min_{\mathbf{x}\in\mathcal{X}} \eta_t \langle \mathbf{x}, \nabla f_t(\mathbf{x}_t)\rangle + \frac{1}{2}\\|\mathbf{x}-\mathbf{x}_t\\|_2^2$	$\\|\mathbf{x}\\|_2^2$	$\frac{1}{\sqrt{t}}$	$\mathcal{O}(\sqrt{T})$
OGD for strongly c.	$\mathbf{x}_{t+1} = \arg\min_{\mathbf{x}\in\mathcal{X}} \eta_t \langle \mathbf{x}, \nabla f_t(\mathbf{x}_t)\rangle + \frac{1}{2}\\|\mathbf{x}-\mathbf{x}_t\\|_2^2$	$\\|\mathbf{x}\\|_2^2$	$\frac{1}{\sigma t}$	$\mathcal{O}(\frac{1}{\sigma}\log T)$
ONS for exp-concave	$\mathbf{x}_{t+1} = \arg\min_{\mathbf{x}\in\mathcal{X}} \eta_t \langle \mathbf{x}, \nabla f_t(\mathbf{x}_t)\rangle + \frac{1}{2}\\|\mathbf{x}-\mathbf{x}_t\\|_{A_t}^2$	$\\|\mathbf{x}\\|_{A_t}^2$	$\frac{1}{\gamma}$	$\mathcal{O}(\frac{d}{\gamma}\log T)$
Hedge for PEA	$\mathbf{x}_{t+1} = \arg\min_{\mathbf{x}\in\Delta_N} \eta_t \langle \mathbf{x}, \nabla f_t(\mathbf{x}_t)\rangle + \mathrm{KL}(\mathbf{x}\\|\mathbf{x}_t)$	$\sum_{i=1}^N x_i \log x_i$	$\sqrt{\frac{\ln N}{T}}$	$\mathcal{O}(\sqrt{T\log N})$

Online Mirror Descent

- Our previous mentioned algorithms can all be covere

minimax optimal

Algo.	OMD/proximal form	$\psi(\cdot)$	η_t	Regret_T
OGD for convex	$\mathbf{x}_{t+1} = \arg\min_{\mathbf{x}\in\mathcal{X}} \eta_t \langle \mathbf{x}, \nabla f_t(\mathbf{x}_t)\rangle + \frac{1}{2}\|\mathbf{x}-\mathbf{x}_t\|_2^2$	$\|\mathbf{x}\|_2^2$	$\frac{1}{\sqrt{t}}$	$\mathcal{O}(\sqrt{T})$
OGD for strongly c.	$\mathbf{x}_{t+1} = \arg\min_{\mathbf{x}\in\mathcal{X}} \eta_t \langle \mathbf{x}, \nabla f_t(\mathbf{x}_t)\rangle + \frac{1}{2}\|\mathbf{x}-\mathbf{x}_t\|_2^2$	$\|\mathbf{x}\|_2^2$	$\frac{1}{\sigma t}$	$\mathcal{O}(\frac{1}{\sigma}\log T)$
ONS for exp-concave	$\mathbf{x}_{t+1} = \arg\min_{\mathbf{x}\in\mathcal{X}} \eta_t \langle \mathbf{x}, \nabla f_t(\mathbf{x}_t)\rangle + \frac{1}{2}\|\mathbf{x}-\mathbf{x}_t\|_{A_t}^2$	$\|\mathbf{x}\|_{A_t}^2$	$\frac{1}{\gamma}$	$\mathcal{O}(\frac{d}{\gamma}\log T)$
Hedge for PEA	$\mathbf{x}_{t+1} = \arg\min_{\mathbf{x}\in\Delta_N} \eta_t \langle \mathbf{x}, \nabla f_t(\mathbf{x}_t)\rangle + \text{KL}(\mathbf{x}\|\mathbf{x}_t)$	$\sum_{i=1}^{N} x_i \log x_i$	$\sqrt{\frac{\ln N}{T}}$	$\mathcal{O}(\sqrt{T\log N})$

Beyond the Worst-Case Analysis

- All above regret guarantees hold against the worst case
  - Matching the ***minimax optimality***
  - The environment is ***fully adversarial***

oblivious adversary

examination

adaptive adversary

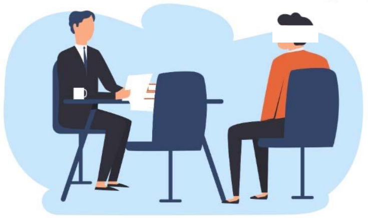

interview

- However, in practice:
  - We are not always interested in the ***worst-case scenario***
  - Environments can exhibit ***specific patterns***: gradual change, periodicity…

$\Longrightarrow$ We are after some more ***problem-dependent*** guarantees.

Beyond the Worst-Case Analysis

- Beyond the worst-case analysis, achieving more adaptive results.

  (1) ***adaptivity***: achieving better guarantees in easy problem instances;

  (2) ***robustness***: maintaining the same worst-case guarantee.

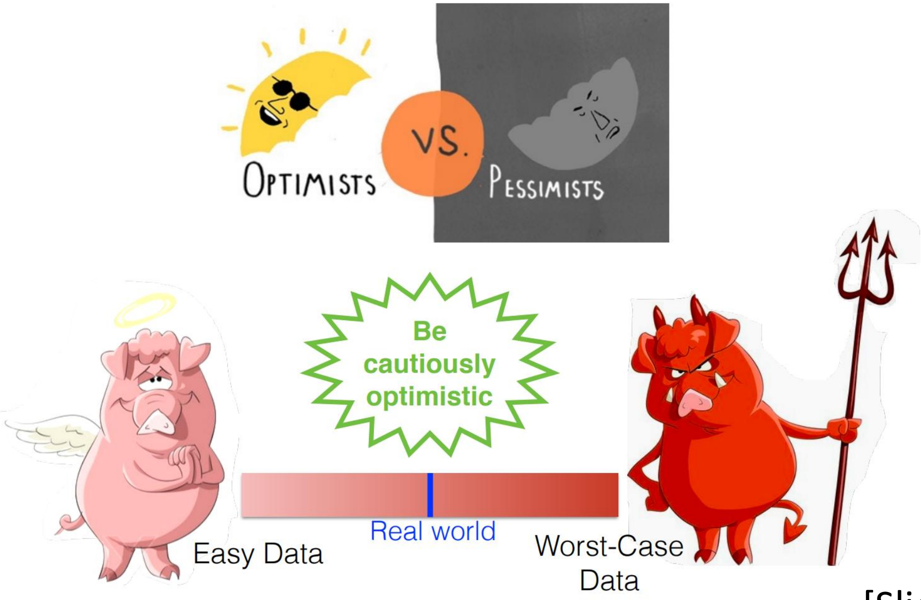

[Slides from Dylan Foster, Adaptive Online Learning @NIPS'15 workshop]

Prediction with Expert Advice

- Recall the PEA setup

At each round $t = 1, 2, \cdots$

(1) the player first picks a weight $\boldsymbol{p}_t$ from a simplex Δ_N;

(2) and simultaneously environments pick a loss vector $\boldsymbol{\ell}_t \in \mathbb{R}^N$;

(3) the player suffers loss $f_t(\boldsymbol{p}_t) \triangleq \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle$, observes $\boldsymbol{\ell}_t$ and updates the model.

- Performance measure: ***regret***

$$\text{Regret}_T \triangleq \sum_{t=1}^{T} \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle - \min_{i \in [N]} \sum_{t=1}^{T} \ell_{t,i}$$

*benchmark the performance with respect to the **best expert***

Prediction with Expert Advice

- Hedge algorithm: taking the "softmax" operation

At each round $t = 1, 2, \cdots$

(1) compute $\boldsymbol{p}_t \in \Delta_N$ such that $p_{t,i} \propto \exp\left(-\eta L_{t-1,i}\right)$ for $i \in [N]$

(2) the player submits $\boldsymbol{p}_t$, suffers loss $\langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle$, and observes loss $\boldsymbol{\ell}_t \in \mathbb{R}^N$

(3) update $\boldsymbol{L}_t = \boldsymbol{L}_{t-1} + \boldsymbol{\ell}_t$

- Another greedy update: $p_{t+1,i} \propto p_{t,i} \exp\left(-\eta \ell_{t,i}\right)$ for $i \in [N]$

  where we set the uniform initialization as $p_{0,i} = 1/N, \forall i \in [N]$.

- The two updates are significantly different when learning rate is changing.

Hedge: Regret Bound

Theorem 1. *Suppose that $\forall t \in [T]$ and $i \in [N], 0 \le \ell_{t,i} \le 1$, then Hedge with learning rate η guarantees*

$$\mathrm{Regret}_T \le \frac{\ln N}{\eta} + \eta T = \mathcal{O}\big(\sqrt{T \log N}\big),$$ **minimax optimal**

where the last equality is by setting η optimally as $\sqrt{(\ln N)/T}$.

- What if there exists an *excellent* expert? i.e., $L_{T,i} \ll T$ holds for some $i \in [N]$.
- Goal: can we achieve a "small-loss" bound? something like $\mathcal{O}(\sqrt{L_{T,i^\star} \log N})$.

Small-Loss Bounds for PEA

Theorem 2. *Suppose that* $\forall t \in [T]$ *and* $i \in [N], 0 \le \ell_{t,i} \le 1$, *then Hedge with learning rate* $\eta \in (0,1)$ *guarantees*

$$\text{Regret}_T \le \frac{\ln N}{\eta} + \eta \sum_{t=1}^{T} \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle,$$

which can further imply

$$\text{Regret}_T \le \frac{1}{1-\eta}\left(\frac{\ln N}{\eta} + \eta L_{T,i^\star}\right) = \mathcal{O}\left(\sqrt{L_{T,i^\star} \log N} + \log N\right),$$

by setting $\eta = \min\left\{\frac{1}{2}, \sqrt{\frac{\ln N}{L_{T,i^\star}}}\right\}$.

- When $L_{T,i^\star} = \mathcal{O}(T)$, it can recover the ***minimax*** $\mathcal{O}(\sqrt{T \log N})$ guarantee.
- When $L_{T,i^\star} = \mathcal{O}(1)$, the regret bound is $\mathcal{O}(\log N)$, which is independent of T!

Small-Loss Bounds for PEA

Theorem 2. *Suppose that* $\forall t \in [T]$ *and* $i \in [N], 0 \le \ell_{t,i} \le 1$, *then Hedge with learning rate* $\eta \in (0,1)$ *guarantees*

$$\text{Regret}_T \le \frac{\ln N}{\eta} + \eta \sum_{t=1}^{T} \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle,$$

which can further imply

$$\text{Regret}_T \le \frac{1}{1-\eta}\left(\frac{\ln N}{\eta} + \eta L_{T,i^\star}\right) = \mathcal{O}\left(\sqrt{L_{T,i^\star} \log N} + \log N\right),$$

by setting $\eta = \min\left\{\frac{1}{2}, \sqrt{\frac{\ln N}{L_{T,i^\star}}}\right\}$.

- When $L_{T,i^\star} = \mathcal{O}(T)$, it can recover the ***minimax*** $\mathcal{O}(\sqrt{T \log N})$ guarantee.
- When $L_{T,i^\star} = \mathcal{O}(1)$, the regret bound is $\mathcal{O}(\log N)$, which is independent of T!

Review: Potential-based Proof

Proof. Recall the (worst-case regret) analysis of Hedge. We present a 'potential-based' proof, where the potential is defined as

$$\Phi_t \triangleq \frac{1}{\eta} \ln \left(\sum_{i=1}^{N} \exp\left(-\eta L_{t,i}\right) \right).$$

$$\begin{aligned}
\Phi_t - \Phi_{t-1} &= \frac{1}{\eta} \ln \left(\frac{\sum_{i=1}^{N} \exp\left(-\eta L_{t,i}\right)}{\sum_{i=1}^{N} \exp\left(-\eta L_{t-1,i}\right)} \right) \\
&= \frac{1}{\eta} \ln \left(\sum_{i=1}^{N} \left(\frac{\exp\left(-\eta L_{t-1,i}\right)}{\sum_{i=1}^{N} \exp\left(-\eta L_{t-1,i}\right)} \exp(-\eta \ell_{t,i}) \right) \right) \\
&= \frac{1}{\eta} \ln \left(\sum_{i=1}^{N} p_{t,i} \exp\left(-\eta \ell_{t,i}\right) \right) && \text{(update step of } \boldsymbol{p}_t\text{)} \\
&\le \frac{1}{\eta} \ln \left(\sum_{i=1}^{N} p_{t,i} \left(1 - \eta \ell_{t,i} + \eta^2 \ell_{t,i}^2\right) \right) && (\forall x \ge 0, e^{-x} \le 1 - x + x^2)
\end{aligned}$$

Review: Potential-based Proof

Proof. Recall the (worst-case regret) analysis of Hedge. We present a 'potential-based' proof, where the potential is defined as

$$\Phi_t \triangleq \frac{1}{\eta} \ln \left(\sum_{i=1}^{N} \exp\left(-\eta L_{t,i}\right) \right).$$

$$\begin{aligned}
\Phi_t - \Phi_{t-1} &\leq \frac{1}{\eta} \ln \left(\sum_{i=1}^{N} p_{t,i} \left(1 - \eta \ell_{t,i} + \eta^2 \ell_{t,i}^2\right) \right) \\
&= \frac{1}{\eta} \ln \left(1 - \eta \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle + \eta^2 \sum_{i=1}^{N} p_{t,i} \ell_{t,i}^2 \right) \\
&\leq - \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle + \eta \sum_{i=1}^{N} p_{t,i} \ell_{t,i}^2 \qquad (\ln(1+x) \leq x)
\end{aligned}$$

$$\Longrightarrow \qquad \Phi_t - \Phi_{t-1} \leq - \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle + \eta \sum_{i=1}^{N} p_{t,i} \ell_{t,i}^2$$

Review: Potential-based Proof

Proof. $$\Phi_t - \Phi_{t-1} \leq -\langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle + \eta \sum_{i=1}^{N} p_{t,i} \ell_{t,i}^2$$

Summing over t, we have

$$\sum_{t=1}^{T} \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle \leq \Phi_0 - \Phi_T + \eta \sum_{t=1}^{T} \sum_{i=1}^{N} p_{t,i} \ell_{t,i}^2$$

$\Phi_t \triangleq \frac{1}{\eta} \ln \left(\sum_{i=1}^{N} \exp\left(-\eta L_t(i)\right) \right)$

$$\leq \frac{\ln N}{\eta} - \frac{1}{\eta} \ln\left(\exp\left(-\eta L_{T,i^\star}\right)\right) + \eta \sum_{t=1}^{T} \sum_{i=1}^{N} p_{t,i} \ell_{t,i}^2$$

$$\leq \frac{\ln N}{\eta} + L_{T,i^\star} + \eta \sum_{t=1}^{T} \sum_{i=1}^{N} p_{t,i} \ell_{t,i}^2$$

$$\Longrightarrow \quad \sum_{t=1}^{T} \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle - L_{T,i^\star} \leq \frac{\ln N}{\eta} + \eta \sum_{t=1}^{T} \sum_{i=1}^{N} p_{t,i} \ell_{t,i}^2$$

Improved Analysis for Small-Loss Bound

Proof. $$\sum_{t=1}^{T} \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle - L_{T,i^\star} \le \frac{\ln N}{\eta} + \eta \sum_{t=1}^{T} \sum_{i=1}^{N} p_{t,i} \ell_{t,i}^2$$

- For previous worst-case analysis, we simply utilize $\ell_{t,i} \le 1$:

$$\eta \sum_{t=1}^{T} \sum_{i=1}^{N} p_{t,i} \ell_{t,i}^2 \le \eta T$$

- To get a small-loss bound, we improve the analysis to be:

$$\eta \sum_{t=1}^{T} \sum_{i=1}^{N} p_{t,i} \ell_{t,i}^2 \le \eta \sum_{t=1}^{T} \sum_{i=1}^{N} p_{t,i} \ell_{t,i} = \eta \sum_{t=1}^{T} \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle$$

$$\Longrightarrow \quad \sum_{t=1}^{T} \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle - L_{T,i^\star} \le \frac{\ln N}{\eta} + \eta \sum_{t=1}^{T} \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle$$

Improved Analysis for Small-Loss Bound

Proof. $\Longrightarrow \sum_{t=1}^{T}\langle \boldsymbol{p}_t, \boldsymbol{\ell}_t\rangle - L_{T,i^\star} \le \frac{\ln N}{\eta} + \eta \sum_{t=1}^{T}\langle \boldsymbol{p}_t, \boldsymbol{\ell}_t\rangle$

$$(1-\eta)\left(\sum_{t=1}^{T}\langle \boldsymbol{p}_t, \boldsymbol{\ell}_t\rangle - L_{T,i^\star}\right) \le \frac{\ln N}{\eta} + \eta L_{T,i^\star} \qquad \text{(rearrange)}$$

$$\Longrightarrow \sum_{t=1}^{T}\langle \boldsymbol{p}_t, \boldsymbol{\ell}_t\rangle - L_{T,i^\star} \le \frac{1}{1-\eta}\left(\frac{\ln N}{\eta} + \eta L_{T,i^\star}\right) \qquad \square$$

Therefore, we get the small-loss regret bound of order $\mathcal{O}\left(\sqrt{L_{T,i^\star}\log N} + \log N\right)$ when setting the learning rate optimally as $\eta^* = \min\left\{\frac{1}{2}, \sqrt{\frac{\ln N}{L_{T,i^\star}}}\right\}$.

Learning Rate Tuning Issue

- We have showed that Hedge with learning rate η enjoys

$$\sum_{t=1}^{T} \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle - L_{T,i^\star} \le \frac{1}{1-\eta}\left(\frac{\ln N}{\eta} + \eta L_{T,i^\star}\right)$$

Therefore, we get the small-loss regret bound of order $\mathcal{O}\left(\sqrt{L_{T,i^\star}\log N} + \log N\right)$

when setting the learning rate optimally as $\eta^* = \min\left\{\frac{1}{2}, \sqrt{\frac{\ln N}{L_{T,i^\star}}}\right\}$.

⇨ However, this online algorithm is not legitimate, due to the requirement of using $L_{T,i^\star}$ (the cummulative loss of the best expert) as the input.

⇨ *Fortunately, we can remedy it by the* ***self-confident tuning*** *framework.*

Self-confident Tuning Framework

- Recall the OGD algorithm for convex function:

$$\mathbf{x}_{t+1} = \Pi_{\mathcal{X}}\left[\mathbf{x}_t - \eta_t \nabla f_t(\mathbf{x}_t)\right]$$

which enjoys the following regret bound

$$\sum_{t=1}^{T} f_t(\mathbf{x}_t) - \min_{\mathbf{x}\in\mathcal{X}} \sum_{t=1}^{T} f_t(\mathbf{x}) \leq \frac{D^2}{\eta} + \eta G^2 T.$$

We can set $\eta = \frac{D}{G\sqrt{T}}$ to obtain an $\mathcal{O}(\sqrt{T})$ regret bound.

> **Question:** can we remove the dependence of ***T*** when tuning the step size?

$\Longrightarrow$ A natural guess is to set $\eta_t = \frac{D}{G\sqrt{t}}$.

Self-confident Tuning Framework

- ***Self-confident* tuning**: utilize the available empirical quantities to approximate the unknown ones.

⇨ use $\eta_t = \frac{D}{G\sqrt{t}}$ to approximate $\eta^\star = \frac{D}{G\sqrt{T}}$, ensuring the same bound (in order).

Theorem 3. *Suppose the diameter of non-empty closed convex set $\mathcal{X}$ is D and $\|\nabla f_t(\mathbf{x})\| \leq G$ for any $\mathbf{x} \in \mathcal{X}$. Then OGD with step size tuning $\eta_t = \frac{D}{G\sqrt{t}}$ ensures the following regret bound:*

$$\sum_{t=1}^{T} f_t(\mathbf{x}_t) - \sum_{t=1}^{T} f_t(\mathbf{u}) \leq \frac{3}{2} GD\sqrt{T}.$$

Self-confident Tuning Framework

Theorem 3. *Suppose the diameter of non-empty closed convex set $\mathcal{X}$ is D and $\|\nabla f_t(\mathbf{x})\| \leq G$ for any $\mathbf{x} \in \mathcal{X}$. Then OGD with step size tuning $\eta_t = \frac{D}{G\sqrt{t}}$ ensures the following regret bound:*

$$\sum_{t=1}^{T} f_t(\mathbf{x}_t) - \sum_{t=1}^{T} f_t(\mathbf{u}) \leq \frac{3}{2} GD\sqrt{T}.$$

Proof.

$$\begin{aligned}
\sum_{t=1}^{T} f_t(\mathbf{x}_t) - f_t(\mathbf{u}) &\leq \frac{1}{2} \sum_{t=1}^{T} \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) \|\mathbf{u} - \mathbf{x}_t\|_2^2 + \sum_{t=1}^{T} \eta_t \|\nabla f_t(\mathbf{x}_t)\|_2^2 \\
&\leq \frac{D^2}{2} \sum_{t=1}^{T} \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) + \frac{G^2}{2} \sum_{t=1}^{T} \eta_t \\
&= \frac{D^2}{2\eta_T} + \frac{GD}{2} \sum_{t=1}^{T} \frac{1}{\sqrt{t}} \qquad \left(\tfrac{1}{\eta_0} = 0\right) \\
&\leq \frac{GD\sqrt{T}}{2} + GD\sqrt{T} = \frac{3}{2} GD\sqrt{T} \qquad \left(\textstyle\sum_{t=1}^{T} \frac{1}{\sqrt{t}} \leq 2\sqrt{T}\right)
\end{aligned}$$

Self-confident Tuning Framework

- Consider the small-loss bound for PEA problem.

Achieving small loss bound $\mathcal{O}\left(\sqrt{L_{T,i^\star}\log N}+\log N\right)$ with $\eta=\min\left\{\frac{1}{2},\sqrt{\frac{\ln N}{L_{T,i^\star}}}\right\}$.

Goal: tuning η without the knowledge of $L_{T,i^\star}$

Deploying self-confident tuning: how can we empirically approximate $L_{T,i^\star}$?

$$L_{T,i}\triangleq\sum_{t=1}^{T}\ell_{t,i},\quad i^\star=\arg\min_{i\in[N]}L_{T,i}$$

$$L_{t,i}\triangleq\sum_{s=1}^{t}\ell_{s,i},\quad i_t^\star=\arg\min_{i\in[N]}L_{t,i}$$

⇨ **Key challenge:** *index $i^\star$ and index sequence $\{i_t^\star\}_{t=1}^{T}$ can be highly different*

Self-confident Tuning Framework

- Consider the small-loss bound for PEA problem.

Achieving small loss bound $\mathcal{O}\left(\sqrt{L_{T,i^\star} \log N} + \log N\right)$ with $\eta = \min\left\{\frac{1}{2}, \sqrt{\frac{\ln N}{L_{T,i^\star}}}\right\}$.

We need to dive into the regret analysis.

$$\sum_{t=1}^{T} \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle - L_{T,i^\star} \leq \frac{\ln N}{\eta} + \eta \sum_{t=1}^{T} \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle$$

$\Longrightarrow$

$$\sum_{t=1}^{T} \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle - L_{T,i^\star} \leq 2\sqrt{(\ln N) \sum_{t=1}^{T} \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle}$$

by setting $\eta = \sqrt{\frac{\ln N}{\sum_{t=1}^{T} \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle}}$.

$\Longleftrightarrow$ Denoted by $\widetilde{L}_T = \sum_{t=1}^{T} \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle$

we obtain $\widetilde{L}_T - L_{T,i^\star} \leq 2\sqrt{(\ln N)\widetilde{L}_T}$

Lemma. *For $x, y, a \in \mathbb{R}_+$ that satisfy $x - y \leq \sqrt{ax}$, it implies $x - y \leq \sqrt{ay} + a$.*

$\Longrightarrow$ $\widetilde{L}_T - L_{T,i^\star} \leq 2\sqrt{(\ln N) L_{T,i^\star}} + 4 \ln N$

by resolving $\widetilde{L}_T$.

Self-confident Tuning Framework

- Consider the small-loss bound for PEA problem.

Achieving small loss bound $\mathcal{O}\left(\sqrt{L_{T,i^\star}\log N} + \log N\right)$ with $\eta = \min\left\{\frac{1}{2}, \sqrt{\frac{\ln N}{L_{T,i^\star}}}\right\}$.

More specifically, setting $\eta = \sqrt{\frac{\ln N}{\widetilde{L}_T}}$, yields

$$\widetilde{L}_T - L_{T,i^\star} \leq \frac{\ln N}{\eta} + \eta \widetilde{L}_T \Longrightarrow \widetilde{L}_T - L_{T,i^\star} \leq 2\sqrt{(\ln N)\widetilde{L}_T} \Longrightarrow \widetilde{L}_T - L_{T,i^\star} \leq \mathcal{O}\left(\sqrt{(\log N)L_{T,i^\star}} + \log N\right)$$

While $\widetilde{L}_T$ cannot be obtained ahead of time, a ***natural*** emperical approxiamtion is:

$$\eta_t = \sqrt{\frac{\ln N}{\widetilde{L}_t}}, \quad \text{where } \widetilde{L}_t = \sum_{s=1}^{t} \langle \boldsymbol{p}_s, \boldsymbol{\ell}_s \rangle \qquad p_{t+1,i} \propto \exp\left(-\eta_t L_{t,i}\right), \forall i \in [N]$$

Self-confident Tuning Framework

Theorem 4. *Suppose that* $\forall t \in [T]$ *and* $i \in [N], 0 \leq \boldsymbol{\ell}_{t,i} \leq 1$, *then Hedge with adaptive learning rate* $\eta_t = \sqrt{\frac{\ln N}{\tilde{L}_t + 1}}$ *guarantees*

$$\begin{aligned} \text{Regret}_T &\leq 8\sqrt{(L_{T,i^\star} + 1)\ln N} + 3\ln N \\ &= \mathcal{O}\Big(\sqrt{L_{T,i^\star} \log N} + \log N\Big), \end{aligned}$$

where $\widetilde{L}_t = \sum_{s=1}^{t} \langle \boldsymbol{p}_s, \boldsymbol{\ell}_s \rangle$ *is cumulative loss the learner suffered at time* t.

Proof

Proof. We again use 'potential-based' proof here, where the potential is defined as

$$\Phi_t(\eta) \triangleq \frac{1}{\eta} \ln \left(\sum_{i=1}^{N} \exp\left(-\eta L_{t,i}\right) \right)$$

$$\Phi_t(\eta_{t-1}) - \Phi_{t-1}(\eta_{t-1}) = \frac{1}{\eta_{t-1}} \ln \left(\frac{\sum_{i=1}^{N} \exp\left(-\eta_{t-1} L_{t,i}\right)}{\sum_{i=1}^{N} \exp\left(-\eta_{t-1} L_{t-1,i}\right)} \right)$$

$$= \frac{1}{\eta_{t-1}} \ln \left(\sum_{i=1}^{N} \left(\frac{\exp\left(-\eta_{t-1} L_{t-1,i}\right)}{\sum_{i=1}^{N} \exp\left(-\eta_{t-1} L_{t-1,i}\right)} \exp(-\eta_{t-1} \ell_{t,i}) \right) \right)$$

$$= \frac{1}{\eta_{t-1}} \ln \left(\sum_{i=1}^{N} p_{t,i} \exp\left(-\eta_{t-1} \ell_{t,i}\right) \right)$$

(update rule of $\boldsymbol{p}_t$)

$(p_{t,i} \propto \exp\left(-\eta_{t-1} L_{t-1,i}\right), \forall i \in [N])$

Proof

Proof.

$$\Phi_t(\eta_{t-1}) - \Phi_{t-1}(\eta_{t-1}) = \frac{1}{\eta_{t-1}} \ln\left(\sum_{i=1}^{N} p_{t,i} \exp\left(-\eta_{t-1}\ell_{t,i}\right)\right)$$

$$\leq \frac{1}{\eta_{t-1}} \ln\left(\sum_{i=1}^{N} p_{t,i}\left(1 - \eta_{t-1}\ell_{t,i} + \eta_{t-1}^2 \ell_{t,i}^2\right)\right) \qquad (\forall x \geq 0, e^{-x} \leq 1 - x + x^2)$$

$$= \frac{1}{\eta_{t-1}} \ln\left(1 - \eta_{t-1}\left\langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \right\rangle + \eta_{t-1}^2 \sum_{i=1}^{N} p_{t,i}\ell_{t,i}^2\right)$$

$$\leq -\left\langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \right\rangle + \eta_{t-1} \sum_{i=1}^{N} p_{t,i}\ell_{t,i}^2 \qquad (\ln(1+x) \leq x)$$

Proof

$$\Phi_t(\eta) \triangleq \frac{1}{\eta} \ln\left(\sum_{i=1}^{N} \exp\left(-\eta L_{t,i}\right)\right)$$

Proof.

$$\Phi_t(\eta_{t-1}) - \Phi_{t-1}(\eta_{t-1}) \leq -\langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle + \eta_{t-1} \sum_{i=1}^{N} p_{t,i} \ell_{t,i}^2$$

$$\Longrightarrow \quad \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle \leq \Phi_{t-1}(\eta_{t-1}) - \Phi_t(\eta_{t-1}) + \eta_{t-1} \sum_{i=1}^{N} p_{t,i} \ell_{t,i}^2$$

$$\sum_{t=1}^{T} \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle \leq \Phi_0(\eta_0) - \Phi_T(\eta_{T-1}) + \sum_{t=1}^{T} \left(\eta_{t-1} \sum_{i=1}^{N} p_{t,i} \ell_{t,i}^2\right) + \sum_{t=1}^{T} \Big(\Phi_t(\eta_t) - \Phi_t(\eta_{t-1})\Big) \qquad \text{(telescoping)}$$

$$\leq \frac{\ln N}{\eta_{T-1}} - \frac{1}{\eta_{T-1}} \ln\left(\exp(-\eta_{T-1} L_{T,i^\star})\right) + \sum_{t=1}^{T} \eta_{t-1} \sum_{i=1}^{N} p_{t,i} \ell_{t,i} + \sum_{t=1}^{T} \left(\Phi_t(\eta_t) - \Phi_t(\eta_{t-1})\right) \qquad (\eta_0 \geq \eta_{T-1}) \quad (\ell_{t,i} \leq 1)$$

$$= \sqrt{\left(\widetilde{L}_{T-1} + 1\right) \ln N} + L_{T,i^\star} + \sum_{t=1}^{T} \eta_{t-1} \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle + \sum_{t=1}^{T} \left(\Phi_t(\eta_t) - \Phi_t(\eta_{t-1})\right)$$

Proof

$$\Phi_t(\eta) \triangleq \frac{1}{\eta}\ln\left(\sum_{i=1}^N \exp\left(-\eta L_{t,i}\right)\right)$$

Proof. $$\sum_{t=1}^T \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t\rangle \leq \sqrt{\left(\widetilde{L}_{T-1}+1\right)\ln N} + L_{T,i^\star} + \sum_{t=1}^T \eta_{t-1}\langle \boldsymbol{p}_t, \boldsymbol{\ell}_t\rangle + \sum_{t=1}^T\left(\Phi_t(\eta_t) - \Phi_t(\eta_{t-1})\right)$$

To bound $\sum_{t=1}^T \left(\Phi_t(\eta_t) - \Phi_t(\eta_{t-1})\right)$, we prove that $\Phi_t(\eta)$ is increasing w.r.t. η:

$$\begin{aligned}
\eta^2\Phi_t'(\eta) &= \eta^2\left(-\frac{1}{\eta^2}\ln(\frac{1}{N}\sum_{i=1}^N \exp(-\eta L_{t,i})) - \frac{1}{\eta}\frac{\sum_{i=1}^N L_{t,i}\exp(-\eta L_{t,i})}{\sum_{i=1}^N \exp(-\eta L_{t,i})}\right)\\
&= \ln N - \sum_{i=1}^N p_{t+1,i}^\eta\left(\ln\left(\sum_{j=1}^N \exp(-\eta L_{t,j})\right) + \eta L_{t,i}\right) \qquad (p_{t+1,i}^\eta \propto \exp(-\eta L_{t,i}))\\
&= \ln N - \sum_{i=1}^N p_{t+1,i}^\eta \ln\left(\frac{\sum_{j=1}^N \exp(-\eta L_{t,j})}{\exp(-\eta L_{t,i})}\right)\\
&= \ln N - \sum_{i=1}^N p_{t+1,i}^\eta \ln\frac{1}{p_{t+1,i}^\eta} \geq 0
\end{aligned}$$

$$\Longrightarrow \sum_{t=1}^T\left(\Phi_t(\eta_t) - \Phi_t(\eta_{t-1})\right) \leq 0 \qquad (\eta_t \leq \eta_{t-1})$$

Proof

Proof. From the potential-based proof, we already know that

$$\sum_{t=1}^{T} \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle - L_{T,i^\star} \leq \sqrt{(\widetilde{L}_{T-1} + 1) \ln N} + \sum_{t=1}^{T} \eta_{t-1} \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle$$

$$\leq \sqrt{(\widetilde{L}_{T-1} + 1) \ln N} + \sum_{t=1}^{T} \frac{\langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle}{\sqrt{\sum_{s=1}^{t-1} \langle \boldsymbol{p}_s, \boldsymbol{\ell}_s \rangle + 1}}$$

$(\eta_{t-1} = \sqrt{\frac{\ln N}{\widetilde{L}_{t-1}+1}})$

$(\widetilde{L}_{t-1} = \sum_{s=1}^{t-1} \langle \boldsymbol{p}_s, \boldsymbol{\ell}_s \rangle)$

How to bound this term?

⇨ *This is actually a common structure to handle.*

Self-confident Tuning Lemma

Lemma 1. *Let $a_1, a_2, \ldots, a_T$ be non-negative real numbers. Then*

$$\sum_{t=1}^{T} \frac{a_t}{\sqrt{1 + \sum_{s=1}^{t} a_s}} \leq 2\sqrt{1 + \sum_{t=1}^{T} a_t}$$

Lemma 2. *Let $a_1, a_2, \ldots, a_T$ be non-negative real numbers. Then*

$$\sum_{t=1}^{T} \frac{a_t}{\sqrt{1 + \sum_{s=1}^{t-1} a_s}} \leq 4\sqrt{1 + \sum_{t=1}^{T} a_t + \max_{t \in [T]} a_t}$$

The two lemmas are useful for analyzing algorithms with self-confident tuning.

Proof

Lemma 1. *Let $a_1, a_2, \ldots, a_T$ be non-negative real numbers. Then*

$$\sum_{t=1}^{T} \frac{a_t}{\sqrt{1+\sum_{s=1}^{t} a_s}} \leq 2\sqrt{1+\sum_{t=1}^{T} a_t}$$

Proof.

$$\frac{1}{2}x \leq 1-\sqrt{1-x}, \forall x \in [0,1]$$

Let $a_0 \triangleq 1$, by set $x = a_t / \sum_{s=0}^{t} a_s$:

$$\frac{a_t}{2\sum_{s=0}^{t} a_s} \leq 1-\sqrt{1-\frac{a_t}{\sum_{s=0}^{t} a_s}}$$

Proof

Proof.

$$\frac{a_t}{2\sum_{s=0}^{t} a_s} \leq 1 - \sqrt{1 - \frac{a_t}{\sum_{s=0}^{t} a_s}}$$

$$\Longrightarrow \quad \frac{a_t}{2\sqrt{\sum_{s=0}^{t} a_s}} \leq \sqrt{\sum_{s=0}^{t} a_s} - \sqrt{\sum_{s=0}^{t} a_s - \sum_{s=0}^{t-1} a_s}$$

By telescopling from $t = 1$ to T:

$$\sum_{t=1}^{T} \left(\frac{a_t}{2\sqrt{1 + \sum_{s=1}^{t} a_s}} \right) \leq \sqrt{\sum_{s=0}^{T} a_s} - \sqrt{\sum_{s=0}^{1} a_s - \sum_{s=0}^{0} a_s} \leq \sqrt{1 + \sum_{t=1}^{T} a_t} \qquad \square$$

Proof

Lemma 2. *Let $a_1, a_2, \ldots, a_T$ be non-negative real numbers. Then*

$$\sum_{t=1}^{T} \frac{a_t}{\sqrt{1 + \sum_{s=1}^{t-1} a_s}} \leq 4\sqrt{1 + \sum_{t=1}^{T} a_t + \max_{t \in [T]} a_t}$$

Proof. We define that $\max_{t \in [T]} a_t = B$.

- **Case 1.** If $\sum_{t=1}^{T} a_t \leq B$:

$$\sum_{t=1}^{T} \frac{a_t}{\sqrt{1 + \sum_{s=1}^{t-1} a_s}} \leq \sum_{t=1}^{T} a_t \leq B$$, Lemma 2 is obviously satisfied.

Proof

Lemma 2. *Let $a_1, a_2, \ldots, a_T$ be non-negative real numbers. Then*

$$\sum_{t=1}^{T} \frac{a_t}{\sqrt{1+\sum_{s=1}^{t-1} a_s}} \leq 4\sqrt{1+\sum_{t=1}^{T} a_t + \max_{t\in[T]} a_t}$$

Proof. We define that $\max_{t\in[T]} a_t = B$.

- **Case 2.** If $\sum_{t=1}^{T} a_t \geq B$, we define $t_0 \triangleq \min\left\{t : \sum_{s=1}^{t-1} x_s \geq B\right\}$:

$$\sum_{t=1}^{T} \frac{a_t}{\sqrt{1+\sum_{s=1}^{t-1} a_s}} \leq B + \sum_{t=t_0}^{T} \frac{a_t}{\sqrt{1+\sum_{s=1}^{t-1} a_s}} \leq B + \sum_{t=t_0}^{T} \frac{a_t}{\sqrt{1+\frac{\sum_{s=1}^{t-1} a_s + a_t}{2}}}$$

$\left(\frac{x+y}{2} \leq x \text{ for } x \geq y\right)$

Proof

Lemma 2. *Let $a_1, a_2, \ldots, a_T$ be non-negative real numbers. Then*

$$\sum_{t=1}^{T} \frac{a_t}{\sqrt{1 + \sum_{s=1}^{t-1} a_s}} \leq 4\sqrt{1 + \sum_{t=1}^{T} a_t + \max_{t\in[T]} a_t}$$

Proof. We define that $\max_{t\in[T]} a_t = B$.

- **Case 2.** If $\sum_{t=1}^{T} a_t \geq B$, we define $t_0 \triangleq \min\left\{t : \sum_{s=1}^{t-1} x_s \geq B\right\}$:

$$B + \sum_{t=t_0}^{T} \frac{a_t}{\sqrt{1 + \frac{\sum_{s=1}^{t-1} a_s + a_t}{2}}} \leq B + \sum_{t=t_0}^{T} \frac{2a_t}{\sqrt{1 + \sum_{s=1}^{t} a_s}} \overset{\text{(Lemma 1)}}{\leq} B + 4\sqrt{1 + \sum_{t=1}^{T} a_t} \qquad \square$$

Small-Loss bound for PEA: Proof

Proof. From previous potential-based proof, we already known that

$$\sum_{t=1}^{T} \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle - L_{T,i^\star} \leq \sqrt{(\widetilde{L}_{T-1} + 1) \ln N} + \sum_{t=1}^{T} \frac{\langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle}{\sqrt{\sum_{s=1}^{t-1} \langle \boldsymbol{p}_s, \boldsymbol{\ell}_s \rangle + 1}}$$

Lemma 2. *Let $a_1, a_2, \ldots, a_T$ be non-negative real numbers. Then*

$$\sum_{t=1}^{T} \frac{a_t}{\sqrt{1 + \sum_{s=1}^{t-1} a_s}} \leq 4\sqrt{1 + \sum_{t=1}^{T} a_t + \max_{t \in [T]} a_t}$$

$$\Longrightarrow \quad \widetilde{L}_T - L_{T,i^\star} \leq \sqrt{(\widetilde{L}_{T-1} + 1) \ln N} + 4\sqrt{1 + \widetilde{L}_T + 1} \qquad (\ell_i \leq 1, \forall i \in [N])$$

$$\leq \sqrt{(\widetilde{L}_T + 1) \ln N} + 4\sqrt{1 + \widetilde{L}_T + 1}$$

Small-Loss bound for PEA: Proof

Proof. $\Longrightarrow$ $$\widetilde{L}_T - L_{T,i^\star} \leq \sqrt{(\widetilde{L}_T + 1)\ln N} + 4\sqrt{1 + \widetilde{L}_T} + 1$$

Then we solve above inequality. Let $x \triangleq \widetilde{L}_T + 1$:

$$x - (\sqrt{\ln N} + 4)\sqrt{x} \leq L_{T,i^\star} + 2 \quad \Longrightarrow \quad \left(\sqrt{x} - \frac{\sqrt{\ln N} + 4}{2}\right)^2 \leq L_{T,i^\star} + 2 + \left(\frac{\sqrt{\ln N} + 4}{2}\right)^2$$

This implies that

$$\sqrt{\widetilde{L}_T - 1} \leq \sqrt{L_{T,i^\star} + 2 + \left(\frac{\sqrt{\ln N} + 4}{2}\right)^2} + \frac{\sqrt{\ln N} + 4}{2}$$

$$\Longrightarrow \quad \widetilde{L}_T \leq 3\ln N + L_{T,i^\star} + 8\sqrt{(L_{T,i^\star} + 1)\ln N} = \mathcal{O}\left(\sqrt{L_{T,i^\star}\log N} + \log N\right).$$ (squaring both sides) $\square$

Recall: Small-Loss Bound for PEA

- So far, we have obtained a PEA algorithm with small-loss bound.

Theorem 4. *Suppose that* $\forall t \in [T]$ *and* $i \in [N], 0 \le \boldsymbol{\ell}_{t,i} \le 1$, *then Hedge with adaptive learning rate* $\eta_t = \sqrt{\frac{\ln N}{\tilde{L}_t + 1}}$ *guarantees*

$$\begin{aligned} \text{Regret}_T &\le 8\sqrt{(L_{T,i^\star} + 1)\ln N} + 3\ln N \\ &= \mathcal{O}\Big(\sqrt{L_{T,i^\star}\log N} + \log N\Big), \end{aligned}$$

where $\widetilde{L}_t = \sum_{s=1}^t \langle \boldsymbol{p}_s, \boldsymbol{\ell}_s \rangle$ *is cumulative loss the learner suffered at time* t.

- Can we further extend the result to more ***general OCO*** setting?

Small Loss in General OCO Setting

Definition 4 (Small Loss)**.** The small-loss quantity of the OCO problem (online function $f_t : \mathcal{X} \mapsto \mathbb{R}$) is defined as

$$F_T = \min_{\mathbf{x} \in \mathcal{X}} \sum_{t=1}^{T} f_t(\mathbf{x})$$

- By taking $f_t(\mathbf{x}) = \langle \mathbf{x}, \boldsymbol{\ell}_t \rangle$ and $\mathcal{X} = \Delta_N$, we recover the definition of the small-loss quantity of PEA problem:

$$F_T = \min_{\mathbf{x} \in \Delta_N} \sum_{t=1}^{T} \langle \mathbf{x}, \boldsymbol{\ell}_t \rangle = \sum_{t=1}^{T} \boldsymbol{\ell}_{t,i^*} = L_{T,i^*}$$

Self-bounding Property

- We require the following ***self-bounding property*** to ensure the small-loss bound for general OCO.

Lemma 3 (Self-bounding Property). *For an L-smooth function $f : \mathbb{R}^d \mapsto \mathbb{R}$ with $\mathbf{x}^\star \in \arg\min_{\mathbf{v}\in\mathbb{R}^d} f(\mathbf{v})$, we have that*

$$\|\nabla f(\mathbf{x})\|_2 \leq \sqrt{2L(f(\mathbf{x}) - f(\mathbf{x}^\star))}, \quad \forall \mathbf{x} \in \mathcal{X}.$$

Proof. By the smoothness of f, for any $\mathbf{x}, \boldsymbol{\delta} \in \mathbb{R}^d$ we have

$$f(\mathbf{x}+\boldsymbol{\delta}) \leq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \boldsymbol{\delta} \rangle + \frac{L}{2}\|\boldsymbol{\delta}\|_2^2.$$

Self-bounding Property

Proof.

$$f(\mathbf{x}+\boldsymbol{\delta}) \leq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \boldsymbol{\delta} \rangle + \frac{L}{2} \|\boldsymbol{\delta}\|_2^2$$

Choosing $\boldsymbol{\delta} = -\frac{\nabla f(\mathbf{x})}{L}$ gives

$$f\left(\mathbf{x} - \frac{\nabla f(\mathbf{x})}{L}\right) \leq f(\mathbf{x}) - \frac{\|\nabla f(\mathbf{x})\|_2^2}{2L}.$$

(actually one-step improvement lemma)

Notice that $f(\mathbf{x}^\star) \leq f(\mathbf{x} - \frac{\nabla f(\mathbf{x})}{L})$ by definition, which implies

$$f(\mathbf{x}^\star) \leq f\left(\mathbf{x} - \frac{\nabla f(\mathbf{x})}{L}\right) \leq f(\mathbf{x}) - \frac{\|\nabla f(\mathbf{x})\|_2^2}{2L}.$$

Rearranging the above terms finishes the proof. □

Self-bounding Property

Lemma 3 (Self-bounding Property). *For an L-smooth function* $f : \mathbb{R}^d \mapsto \mathbb{R}$ *with* $\mathbf{x}^\star \in \arg\min_{\mathbf{v} \in \mathbb{R}^d} f(\mathbf{v})$, *we have that*

$$\|\nabla f(\mathbf{x})\|_2 \leq \sqrt{2L(f(\mathbf{x}) - f(\mathbf{x}^\star))}, \quad \forall \mathbf{x} \in \mathcal{X}.$$

Corollary 1. *For an L-smooth and non-negative function* $f : \mathbb{R}^d \mapsto \mathbb{R}$ *, we have that*

$$\|\nabla f(\mathbf{x})\|_2 \leq \sqrt{2Lf(\mathbf{x})}, \quad \forall \mathbf{x} \in \mathcal{X}.$$

Achieving Small-Loss Bound

- We show that under the ***self-bounding condition***, OGD can yield the desired small-loss regret bound.

$$\mathbf{x}_{t+1} = \Pi_{\mathbf{x}\in\mathcal{X}}[\mathbf{x}_t - \eta_t \nabla f_t(\mathbf{x}_t)]$$

Theorem 6 (Small-loss Bound). *Assume that f_t is L-smooth and non-negative for all $t \in [T]$, when setting $\eta_t = \frac{D}{\sqrt{1+\widetilde{G}_t}}$, the regret of OGD to any comparator $\mathbf{u} \in \mathcal{X}$ is bounded as*

$$\text{Regret}_T = \sum_{t=1}^{T} f_t(\mathbf{x}_t) - \sum_{t=1}^{T} f_t(\mathbf{u}) \le \mathcal{O}\left(\sqrt{1+F_T}\right)$$

where $\widetilde{G}_t = \sum_{s=1}^{t} \|\nabla f_s(\mathbf{x}_s)\|_2^2$ is the empirical estimator of cumulative gradient G_T.

Proof

Proof. $$\sum_{t=1}^{T} f_t(\mathbf{x}_t) - \sum_{t=1}^{T} f_t(\mathbf{u}) \le \sum_{t=1}^{T} \eta_t \|\nabla f_t(\mathbf{x}_t)\|_2^2 + \sum_{t=1}^{T} \frac{1}{2\eta_t}\left(\|\mathbf{u}-\mathbf{x}_t\|_2^2 - \|\mathbf{u}-\mathbf{x}_{t+1}\|_2^2\right)$$

$$\sum_{t=1}^{T} \eta_t \|\nabla f_t(\mathbf{x}_t)\|_2^2 = D\sum_{t=2}^{T} \frac{\|\nabla f_t(\mathbf{x}_t)\|_2^2}{\sqrt{1+\widetilde{G}_t}} + G^2 \le 2D\sqrt{1+\sum_{t=1}^{T}\|\nabla f_t(\mathbf{x}_t)\|_2^2} + G^2$$

$(\eta_1 \triangleq 1)$ $(\widetilde{G}_t = \sum_{s=1}^{t} \|\nabla f_s(\mathbf{x}_s)\|_2^2)$

Lemma 1. *Let $a_1, a_2, \ldots, a_T$ be non-negative real numbers. Then*

$$\sum_{t=1}^{T} \frac{a_t}{\sqrt{1+\sum_{s=1}^{t} a_s}} \le 2\sqrt{1+\sum_{t=1}^{T} a_t}$$

Proof

Proof. $$\sum_{t=1}^{T} f_t(\mathbf{x}_t) - \sum_{t=1}^{T} f_t(\mathbf{u}) \le \sum_{t=1}^{T} \eta_t \|\nabla f_t(\mathbf{x}_t)\|_2^2 + \sum_{t=1}^{T} \frac{1}{2\eta_t}\left(\|\mathbf{u}-\mathbf{x}_t\|_2^2 - \|\mathbf{u}-\mathbf{x}_{t+1}\|_2^2\right)$$

$$\sum_{t=1}^{T} \eta_t \|\nabla f_t(\mathbf{x}_t)\|_2^2 = D\sum_{t=2}^{T} \frac{\|\nabla f_t(\mathbf{x}_t)\|_2^2}{\sqrt{1+\widetilde{G}_t}} + G^2 \le 2D\sqrt{1+\sum_{t=1}^{T}\|\nabla f_t(\mathbf{x}_t)\|_2^2} + G^2$$

$(\eta_1 \triangleq 1)$ $\quad$ $(\widetilde{G}_t = \sum_{s=1}^{t} \|\nabla f_s(\mathbf{x}_s)\|_2^2)$

$$\le 2D\sqrt{1+2L\sum_{t=1}^{T} f_t(\mathbf{x}_t)} + G^2$$

(self-bounding property)

Proof

Proof. $$\sum_{t=1}^{T} f_t(\mathbf{x}_t) - \sum_{t=1}^{T} f_t(\mathbf{u}) \leq \sum_{t=1}^{T} \eta_t \left\|\nabla f_t(\mathbf{x}_t)\right\|_2^2 + \sum_{t=1}^{T} \frac{1}{2\eta_t}\left(\|\mathbf{u}-\mathbf{x}_t\|_2^2 - \|\mathbf{u}-\mathbf{x}_{t+1}\|_2^2\right)$$

$$\sum_{t=1}^{T} \eta_t \|\nabla f_t(\mathbf{x}_t)\|_2^2 \leq 2D\sqrt{1 + 2L\sum_{t=1}^{T} f_t(\mathbf{x}_t) + G^2}$$

$$\sum_{t=1}^{T} \frac{1}{2\eta_t}\left(\|\mathbf{u}-\mathbf{x}_t\|_2^2 - \|\mathbf{u}-\mathbf{x}_{t+1}\|_2^2\right) \leq \frac{D}{2}\sqrt{1 + 2L\sum_{t=1}^{T} f_t(\mathbf{x}_t)} + \frac{D}{2}$$

$$\Longrightarrow \quad \text{Regret}_T = \sum_{t=1}^{T} f_t(\mathbf{x}_t) - \sum_{t=1}^{T} f_t(\mathbf{u}) \leq 3D\sqrt{1 + 2L\sum_{t=1}^{T} f_t(\mathbf{x}_t) + G^2}$$

Proof

Proof. $\text{Regret}_T = \sum_{t=1}^{T} f_t(\mathbf{x}_t) - \sum_{t=1}^{T} f_t(\mathbf{u}) \leq 3D\sqrt{1 + 2L\sum_{t=1}^{T} f_t(\mathbf{x}_t) + G^2}$

Remember how we solve a similar problem in PEA:

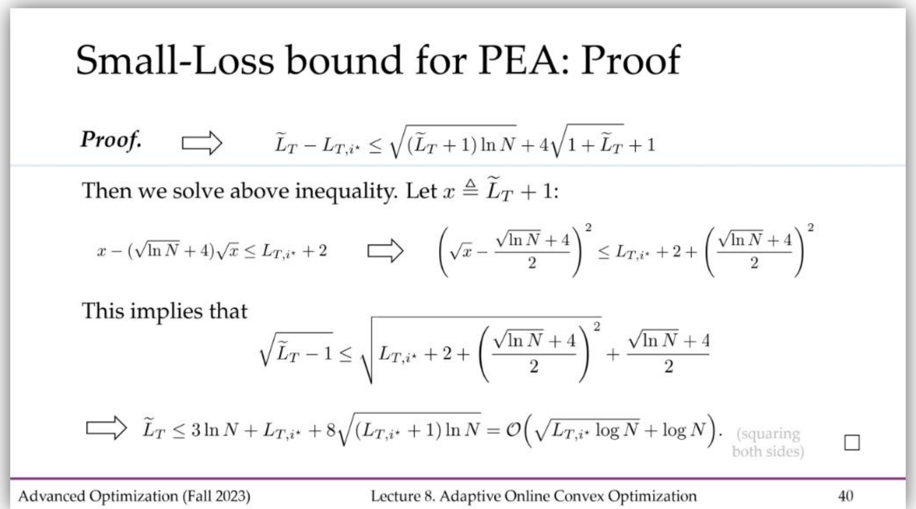

$$\Longrightarrow \quad \text{Regret}_T = \sum_{t=1}^{T} f_t(\mathbf{x}_t) - \sum_{t=1}^{T} f_t(\mathbf{u}) = \mathcal{O}\left(D\sqrt{L\sum_{t=1}^{T} f_t(\mathbf{u}) + 1 + G^2}\right). \qquad \square$$

Several Remarks

- Remark 1: about the non-negative assumption

  When the online functions are non-negative, it is possible to redefine the small-loss quantity by incorporating each-round minimal function value.

- Remark 2: about the smoothness assumption

  Smoothness is necessary to obtain small-loss regret bound by the first-order method (can be proved by the online-to-batch conversion and existing lower bounds for deterministic optimization).

- Remark 3: take care of the way dealing with variance term

  In OGD here we use Lemma 1, while in PEA Hedge for PEA we use Lemma 2.

Summary

Q & A

Thanks!