Lecture 9. Adaptive Online Convex Optimization

Advanced Optimization (Fall 2022)

Peng Zhao

zhaop@lamda.nju.edu.cn

Nanjing University

Review: Hedge Algorithm

Hedge Algorithm

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

(1) compute $\boldsymbol{p}_t \in \Delta_N$ such that $\boldsymbol{p}_t(i) \propto \exp\left(-\eta L_{t-1}(i)\right)$ 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 $L_t = L_{t-1} + \boldsymbol{\ell}_t$

FTL update

$$\boldsymbol{p}_t = \arg\max_{\boldsymbol{p} \in \Delta_N} \langle \boldsymbol{p}, -L_{t-1} \rangle$$

Hedge update

$$\boldsymbol{p}_t(i) \propto \exp\left(-\eta L_{t-1}(i)\right), \forall i \in [N]$$

Review: OMD Framework

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

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*

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

Proof.

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

Review: OMD Analysis

Proof Lemma1. $f_t(\mathbf{x}_t) - f_t(\mathbf{u}) \le \underbrace{\langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{x}_{t+1} \rangle}_{\text{term (a)}} + \underbrace{\langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_{t+1} - \mathbf{u} \rangle}_{\text{term (b)}}$

> **Lemma 2** (Stability Lemma).
> $$\lambda \|\mathbf{x}_1 - \mathbf{x}_2\| \le \|\mathbf{g}_1 - \mathbf{g}_2\|_\star$$

$$\Rightarrow \quad \text{term (a)} = \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{x}_{t+1} \rangle \le \frac{\eta_t}{\lambda} \|\nabla f_t(\mathbf{x}_t)\|_\star^2$$

> **Lemma 3** (Bregman Proximal Inequality).
> $$\langle \mathbf{g}_t, \mathbf{x}_{t+1} - \mathbf{u} \rangle \le \mathcal{D}_\psi(\mathbf{u}, \mathbf{x}_t) - \mathcal{D}_\psi(\mathbf{u}, \mathbf{x}_{t+1}) - \mathcal{D}_\psi(\mathbf{x}_{t+1}, \mathbf{x}_t)$$

$$\Rightarrow \quad \text{term (b)} \le \frac{1}{\eta_t} \left(\mathcal{D}_\psi(\mathbf{u}, \mathbf{x}_t) - \mathcal{D}_\psi(\mathbf{u}, \mathbf{x}_{t+1}) - \mathcal{D}_\psi(\mathbf{x}_{t+1}, \mathbf{x}_t) \right)$$

(negative term, **crucial in this Lec**)

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

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*

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

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*

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

Online Mirror Descent Framework

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

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

Beyond Worst-Case Analysis

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

oblivious adversary

examination

adaptive adversary

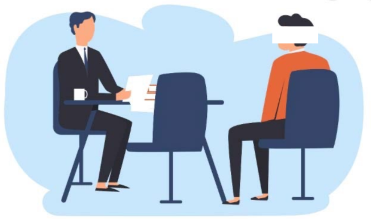

interview

- However, in practice:
  - Not always deal with 'worst-case' scenario
  - Environments can follow specific patterns: gradual change, periodicity…

⇨ We are after some more ***adaptive*** guarantees.

Beyond Worst-Case Analysis

- Beyond the worst-case analysis, achieving more adaptive results.
  - (1) ***adaptivity***: achieving better guarantee in easy problem instance;
  - (2) ***robustness***: maintaining the same worst-case guarantee.

Outline

- Adaptive methods for PEA
  - Small-Loss bound
- A unified framework for OCO: Optimistic OMD
  - Small-Loss bound
  - Gradient-Variance bound
  - Gradient-Variation bounds

Small-loss Bounds for PEA

- Let's first review previous minimax regret result.

Theorem 1. *Suppose that* $\forall t \in [T]$ *and* $i \in [N], 0 \leq \boldsymbol{\ell}_t(i) \leq 1$, *then Hedge with learning rate* η *guarantees*

$$\mathrm{Regret}_T \leq \frac{\ln N}{\eta} + \eta T,$$

which is of order $\mathcal{O}(\sqrt{T \ln N})$ *if* η *is optimally set as* $\sqrt{(\ln N)/T}$.

- Ideally, we want to enjoy a ***smaller*** regret in ***easier*** situations.

Small-loss Bounds for PEA

- Let's first review previous minimax regret result.

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

$$\text{Regret}_T \le \frac{\ln N}{\eta} + \eta T,$$

which is of order $\mathcal{O}(\sqrt{T \ln N})$ *if* η *is optimally set as* $\sqrt{(\ln N)/T}$.

- What if there exists an ***excellent expert***? $\exists i^\star \in [N], \quad L_T(i^\star) \ll T$
- Goal: can we achieve a 'small-loss' bound? i.e.,

$$\mathcal{O}(\sqrt{L_T(i^\star) \ln N})$$

Small-loss Bounds for PEA

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

$$\text{Regret}_T \le \frac{1}{1-\eta}\left(\frac{\ln N}{\eta} + \eta L_T(i^\star)\right),$$

which is of order $\mathcal{O}\left(\sqrt{L_T(i^\star)\ln N} + \ln N\right)$ *if* η *is optimally set as* $\min\left\{\frac{1}{2}, \sqrt{\frac{\ln N}{L_T(i^\star)}}\right\}$.

Small-loss Bounds for PEA

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

$$\mathrm{Regret}_T \le \frac{1}{1-\eta}\left(\frac{\ln N}{\eta} + \eta L_T(i^\star)\right),$$

which is of order $\mathcal{O}\left(\sqrt{L_T(i^\star)\ln N} + \ln N\right)$ *if* η *is optimally set as* $\min\left\{\frac{1}{2}, \sqrt{\frac{\ln N}{L_T(i^\star)}}\right\}$.

Proof. Review the analysis of Hedge.

We present a 'potential-based' proof, where the potential is defined as

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

Review: Potential-based Proof

Proof.

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

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

Review: Potential-based Proof

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

Summing over t, we have

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

$$\begin{aligned}
\sum_{t=1}^T \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle &\le \Phi_0 - \Phi_T + \eta \sum_{t=1}^T \sum_{i=1}^N \boldsymbol{p}_t(i) \boldsymbol{\ell}_t^2(i) \\
&\le \frac{\ln N}{\eta} - \frac{1}{\eta} \ln\left(\exp\left(-\eta L_T(i^\star)\right)\right) + \eta \sum_{t=1}^T \sum_{i=1}^N \boldsymbol{p}_t(i) \boldsymbol{\ell}_t^2(i) \\
&\le \frac{\ln N}{\eta} + L_T(i^\star) + \eta \sum_{t=1}^T \sum_{i=1}^N \boldsymbol{p}_t(i) \boldsymbol{\ell}_t^2(i)
\end{aligned}$$

$$\Longrightarrow \quad \sum_{t=1}^T \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle - L_T(i^\star) \le \frac{\ln N}{\eta} + \eta \sum_{t=1}^T \sum_{i=1}^N \boldsymbol{p}_t(i) \boldsymbol{\ell}_t^2(i)$$

Improved Analysis for Small-loss Bound

Proof. $\Longrightarrow$ $$\sum_{t=1}^T \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle - L_T(i^\star) \leq \frac{\ln N}{\eta} + \eta \sum_{t=1}^T \sum_{i=1}^N \boldsymbol{p}_t(i) \boldsymbol{\ell}_t^2(i)$$

- For previous worst-case analysis, we simply utilize $\boldsymbol{\ell}_t(i) \leq 1$:

$$\eta \sum_{t=1}^T \sum_{i=1}^N \boldsymbol{p}_t(i) \boldsymbol{\ell}_t^2(i) \leq \eta T$$

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

$$\eta \sum_{t=1}^T \sum_{i=1}^N \boldsymbol{p}_t(i) \boldsymbol{\ell}_t^2(i) \leq \eta \sum_{t=1}^T \sum_{i=1}^N \boldsymbol{p}_t(i) \boldsymbol{\ell}_t(i) = \eta \sum_{t=1}^T \langle \boldsymbol{p}_t, \boldsymbol{\ell}_t \rangle$$

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

Improved Analysis for Small-loss Bound

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

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

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

- Therefore, we get the small loss bound $\mathcal{O}\left(\sqrt{L_T(i^\star) \ln N} + \ln N\right)$ if η is optimally set as $\min\left\{\frac{1}{2}, \sqrt{\frac{\ln N}{L_T(i^\star)}}\right\}$.

Improved Analysis for Small-loss Bound

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

- Therefore, we get the small loss bound $\mathcal{O}\left(\sqrt{L_T(i^\star)\ln N} + \ln N\right)$ if η is optimally set as $\min\left\{\frac{1}{2}, \sqrt{\frac{\ln N}{L_T(i^\star)}}\right\}$.

Question: is this algorithm legitimate?

Actually no... as the algorithm requires $L_T(i^\star)$ in advance.

$\Longrightarrow$ Luckily, we can fix it by the ***self-confident tuning*** framework.

Self-confident Tuning

- **Goal**: tuning η without the knowledge of $L_T(i^\star)$
- **Intuition**: two perspective:
  1. The cumulative loss $\widetilde{L}_{t-1} = \sum_{s=1}^{t-1} \langle \boldsymbol{p}_s, \boldsymbol{\ell}_s \rangle$ can be seen as an empirical approximation to the best expert's loss $L_T(i^\star)$;
  2. Use $\widetilde{L}_{t-1} = \sum_{s=1}^{t-1} \langle \boldsymbol{p}_s, \boldsymbol{\ell}_s \rangle$ to approximate $\widetilde{L}_t$ to tune current learning rate η_t.

Hedge update

$$\boldsymbol{p}_t(i) \propto \exp\left(-\eta L_{t-1}(i)\right), \forall i \in [N]$$

Adaptive Hedge update

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

Self-confident Tuning

- **Goal**: tuning η without the knowledge of $L_T(i^\star)$
- **Intuition**: two perspective:
  1. The cumulative loss $\widetilde{L}_{t-1} = \sum_{s=1}^{t-1} \langle \boldsymbol{p}_s, \boldsymbol{\ell}_s \rangle$ can be seen as an empirical approximation to the best expert's loss $L_T(i^\star)$;
  2. Use $\widetilde{L}_{t-1} = \sum_{s=1}^{t-1} \langle \boldsymbol{p}_s, \boldsymbol{\ell}_s \rangle$ to approximate $\widetilde{L}_t$ to tune current learning rate η_t.

Theorem 3. *Suppose that* $\forall t \in [T]$ *and* $i \in [N], 0 \leq \boldsymbol{\ell}_t(i) \leq 1$, *then Hedge with adaptive learning rate* $\eta_{t-1} = \sqrt{\frac{\ln N}{\tilde{L}_{t-1}+1}}$ *guarantees*

$$\text{Regret}_T \leq 8\sqrt{(L_T(i^\star) + 1)\ln N} + 3\ln N,$$

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

Self-confident Tuning: Proof

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

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

$$\begin{aligned}
\Phi_t(\eta_{t-1}) - \Phi_{t-1}(\eta_{t-1}) &= \frac{1}{\eta_{t-1}} \ln \left(\frac{\sum_{i=1}^{N} \exp\left(-\eta_{t-1} L_t(i)\right)}{\sum_{i=1}^{N} \exp\left(-\eta_{t-1} L_{t-1}(i)\right)} \right) \\
&= \frac{1}{\eta_{t-1}} \ln \left(\sum_{i=1}^{N} \left(\frac{\exp\left(-\eta_{t-1} L_{t-1}(i)\right)}{\sum_{i=1}^{N} \exp\left(-\eta_{t-1} L_{t-1}(i)\right)} \exp(-\eta_{t-1} \boldsymbol{\ell}_t(i)) \right) \right) \\
&= \frac{1}{\eta_{t-1}} \ln \left(\sum_{i=1}^{N} \boldsymbol{p}_t(i) \exp\left(-\eta_{t-1} \boldsymbol{\ell}_t(i)\right) \right) \qquad \text{(update rule of } \boldsymbol{p}_t\text{)}
\end{aligned}$$

Self-confident Tuning: Proof

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

Proof.

$$\Phi_t(\eta_{t-1}) - \Phi_{t-1}(\eta_{t-1}) = \frac{1}{\eta_{t-1}} \ln\left(\sum_{i=1}^N \boldsymbol{p}_t(i) \exp\left(-\eta_{t-1}\boldsymbol{\ell}_t(i)\right)\right)$$

$$\leq \frac{1}{\eta_{t-1}} \ln\left(\sum_{i=1}^N \boldsymbol{p}_t(i)\left(1 - \eta_{t-1}\boldsymbol{\ell}_t(i) + \eta_{t-1}^2 \boldsymbol{\ell}_t^2(i)\right)\right) \quad (\forall x \geq 0, e^{-x} \leq 1 - x + x^2)$$

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

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

Self-confident Tuning: Proof

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

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

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

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

$$\leq \frac{\ln N}{\eta_{T-1}} - \frac{1}{\eta_{T-1}} \ln\left(\exp(-\eta_{T-1} L_T(i^\star))\right) + \sum_{t=1}^T \eta_{t-1} \sum_{i=1}^N \boldsymbol{p}_t(i) \boldsymbol{\ell}_t(i) + \sum_{t=1}^T \left(\Phi_t(\eta_t) - \Phi_t(\eta_{t-1})\right)$$

$(\eta_0 \geq \eta_{T-1})$

$(\boldsymbol{\ell}_t(i) \leq 1)$

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

Self-confident Tuning: 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 non-decreasing w.r.t. η:

$$\eta^2 \nabla \Phi_t(\eta) = \eta^2 \left(-\frac{1}{\eta^2} \ln\left(\frac{1}{N} \sum_{i=1}^{N} \exp(-\eta L_t(i))\right) - \frac{1}{\eta} \frac{\sum_{i=1}^{N} L_t(i) \exp(-\eta L_t(i))}{\sum_{i=1}^{N} \exp(-\eta L_t(i))} \right)$$

$$= \ln N - \sum_{i=1}^{N} p_{t+1}^{\eta}(i) \left(\ln\left(\sum_{j=1}^{N} \exp(-\eta L_t(j))\right) + \eta L_t(i) \right)$$

$$= \ln N - \sum_{i=1}^{N} p_{t+1}^{\eta}(i) \ln \left(\frac{\sum_{j=1}^{N} \exp(-\eta L_t(j))}{\exp(-\eta L_t(i))} \right)$$

$$= \ln N - \sum_{i=1}^{N} p_{t+1}^{\eta}(i) \ln \frac{1}{p_{t+1}^{\eta}(i)} \geq 0 \quad \Longrightarrow \quad \sum_{t=1}^{T} \left(\Phi_t(\eta_t) - \Phi_t(\eta_{t-1})\right) \geq 0$$

Self-confident Tuning: Proof

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

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

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

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

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

How to bound this term?

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.

Self-confident Tuning Lemma: Proof

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

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

Proof.

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

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

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

Self-confident Tuning Lemma: 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:

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

Self-confident Tuning Lemma: 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.

Self-confident Tuning Lemma: 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 $t_0 \triangleq \min\left\{t : \sum_{s=1}^{t-1} x_s \geq B\right\}$:

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

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

Self-confident Tuning Lemma: 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*

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

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

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

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 $x \triangleq \widetilde{L}_T + 1$:

$$x - (\sqrt{\ln N} + 4)\sqrt{x} \leq L_T(i^\star) + 2$$

$$\left(\sqrt{x} - \frac{\sqrt{\ln N} + 4}{2}\right)^2 \leq L_T(i^\star) + 2 + \left(\frac{\sqrt{\ln N} + 4}{2}\right)^2$$

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

$\Longrightarrow$ $$\widetilde{L}_T \leq 3\ln N + L_T(i^\star) + 8\sqrt{(L_T(i^\star) + 1)\ln N}$$ (squaring both sides)

Recall: Small-Loss Bound for PEA

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

Theorem 3. *Suppose that* $\forall t \in [T]$ *and* $i \in [N], 0 \leq \boldsymbol{\ell}_t(i) \leq 1$, *then Hedge with adaptive learning rate* $\eta_t = \sqrt{\frac{\ln N}{\tilde{L}_{t-1}+1}}$ *guarantees*

$$\mathrm{Regret}_T \leq 8\sqrt{(L_T(i^\star) + 1)\ln N} + 3\ln N,$$

where $\widetilde{L}_{t-1} = \sum_{s=1}^{t-1} \langle \boldsymbol{p}_s, \boldsymbol{\ell}_s \rangle$ *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 small-loss bounds for general OCO.

Lemma 3 (Self-bounding Property). *For an L-smooth and non-negative function* $f : \mathcal{X} \mapsto \mathbb{R}_+$, *we have that*

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

Proof. First, we know that if f is L-smooth, then for any $\mathbf{x}, \mathbf{y} \in \mathcal{X}$ we have

$$f(\mathbf{y}) - f(\mathbf{x}) - \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle \leq \frac{L}{2}\|\mathbf{y} - \mathbf{x}\|^2$$

Self-bounding Property

Proof.

$$f(\mathbf{y}) - f(\mathbf{x}) - \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle \le \frac{L}{2}\|\mathbf{y} - \mathbf{x}\|^2$$

Therefore, for any $\mathbf{x}, \mathbf{v} \in \mathbb{R}^d$, we have

$$\langle -\nabla f(\mathbf{x}), \mathbf{v} \rangle - \frac{M}{2}\|\mathbf{v}\|^2 \le f(\mathbf{x}) - f(\mathbf{x} + \mathbf{v}) \le f(\mathbf{x}) - \inf_{\mathbf{y} \in \mathbb{R}^d} f(\mathbf{y}).$$

By the definition of the dual norm:

$$\tfrac{1}{2L}\|\nabla f(\mathbf{x})\|_2^2 = \sup_{\mathbf{v} \in \mathbb{R}^d} \langle -\nabla f(\mathbf{x}), \mathbf{v} \rangle - \tfrac{L}{2}\|\mathbf{v}\|^2 \le f(\mathbf{x}) - \inf_{\mathbf{y} \in \mathbb{R}^d} f(\mathbf{y})$$

Besides, the function f is non-negative in $\mathbb{R}^d$ space, therefore $\inf_{\mathbf{y} \in \mathbb{R}^d} f(\mathbf{y}) \ge 0$,

$$\Longrightarrow \quad \|\nabla f(\mathbf{x})\|_2 \le \sqrt{2Lf(\mathbf{x})}$$

Achieving Small-Loss Bound

Lemma 3 (Self-bounding Property). *For an L-smooth and non-negative function* $f : \mathcal{X} \mapsto \mathbb{R}_+$, *we have that*

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

Online Gradient Descent

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

Achieving Small-Loss Bound

Online Gradient Descent

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

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*

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

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.

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

$$\sum_{t=1}^{T} \eta_t \|\nabla f_t(\mathbf{x}_t)\|_2^2 = D \sum_{t=2}^{T} \frac{\|\nabla f_t(\mathbf{x}_t)\|_2^2}{\sqrt{1 + \widetilde{G}_t}} + G^2 \qquad (\eta_1 \triangleq 1)$$

$$\leq 2D\sqrt{1 + \sum_{t=1}^{T} \|\nabla f_t(\mathbf{x}_t)\|_2^2 + G^2} \qquad \text{(self-confident tuning lemma)}$$

$$\leq 2D\sqrt{1 + 2L\sum_{t=1}^{T} f_t(\mathbf{x}_t) + G^2} \qquad \text{(self-bounding property)}$$

Proof

Proof.

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

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

$$\leq \frac{D^2}{2\eta_T} + \frac{1}{2\eta_1}D^2 \leq \frac{D\sqrt{1+2L\sum_{t=1}^{T} f_t(\mathbf{x}_t)} + D}{2}$$

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

Proof

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

Remember how we solve the similar problem in PEA:

Small-loss bound for PEA: Proof

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

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

$$x - (\sqrt{\ln N} + 4)\sqrt{x} \leq L_T(i^\star) + 2$$

$$\left(\sqrt{x} - \frac{\sqrt{\ln N} + 4}{2}\right)^2 \leq L_T(i^\star) + 2 + \left(\frac{\sqrt{\ln N} + 4}{2}\right)^2$$

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

$$\Longrightarrow \quad \widetilde{L}_T \leq 3\ln N + L_T(i^\star) + 8\sqrt{(L_T(i^\star) + 1)\ln N}$$

(squaring both sides) $\square$

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

Towards a Unified Framework

- Previous small-loss bounds seem to be ad-hoc designed.

- Is there a ***unified framework*** to get all adaptive bounds?

- Review OMD Update:

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

What if we some ***prior knowledge*** of the environments?

Optimistic Online Learning

- **Intuition**: what if we some prior knowledge of the environments?

  ⇨ We can *'guess'* the next move.

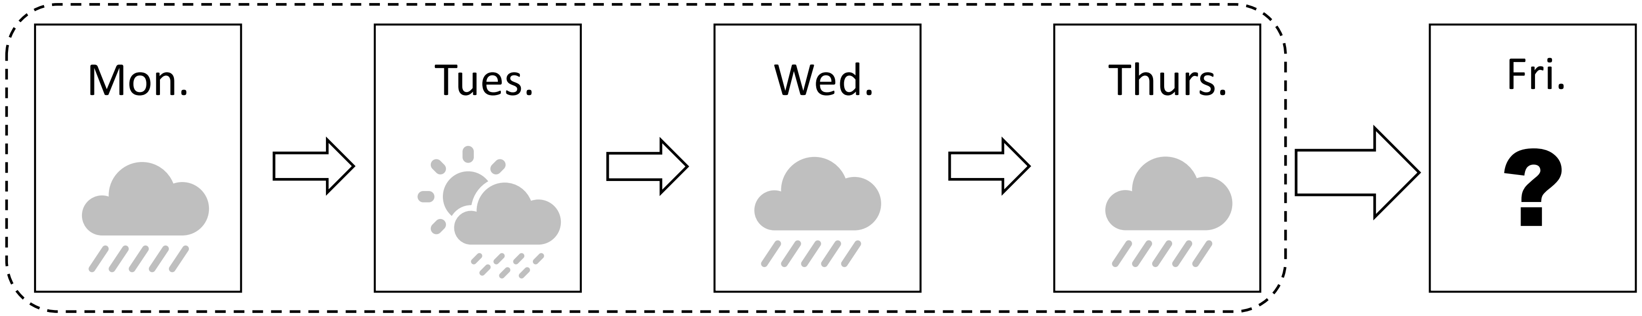

Guess: *It still seems to rain on Friday*

- To formalize the idea, we treat *'guess'* as a hint of future gradients

Optimistic OMD: Algorithmic Framework

- We formulate this intuition as the following two-step update:

Optimistic Online Mirror Descent

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

$$\mathbf{x}_t = \arg\min_{\mathbf{x}\in\mathcal{X}} \eta_t \langle M_t, \mathbf{x} \rangle + \mathcal{D}_\psi(\mathbf{x}, \widehat{\mathbf{x}}_t)$$

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

where $M_t \in \mathbb{R}^d$ is the ***optimistic vector*** at each round.

Optimistic OMD: Generic Analysis

Theorem 4 (Regret for Optimistic OMD). *Assume ψ is 1-strongly convex w.r.t. $\|\cdot\|$, the regret of Optimistic OMD w.r.t. any comparator $\mathbf{u} \in \mathcal{X}$ is bounded as:*

$$\sum_{t=1}^{T} f_t(\mathbf{x}_t) - \sum_{t=1}^{T} f_t(\mathbf{u}) \le \sum_{t=1}^{T} \eta_t \|\nabla f_t(\mathbf{x}_t) - M_t\|_\star^2 \quad \text{(quality of guess)}$$

$$+ \sum_{t=1}^{T} \frac{1}{\eta_t}\Big(\mathcal{D}_\psi(\mathbf{u}, \widehat{\mathbf{x}}_t) - \mathcal{D}_\psi(\mathbf{u}, \widehat{\mathbf{x}}_{t+1})\Big) \quad \text{(telescoping term)}$$

$$- \sum_{t=1}^{T} \frac{1}{\eta_t}\Big(\mathcal{D}_\psi(\widehat{\mathbf{x}}_{t+1}, \mathbf{x}_t) + \mathcal{D}_\psi(\mathbf{x}_t, \widehat{\mathbf{x}}_t)\Big) \quad \text{(negative term)}$$

- Note that for standard OMD, the negative term is usually dropped.
- In Optimistic OMD, the ***negative term*** can be crucial for adaptive regret guarantee.

Optimistic OMD: Generic Analysis

Theorem 4 (Regret for Optimistic OMD). *Assume ψ is 1-strongly convex w.r.t. $\|\cdot\|$, the regret of Optimistic OMD w.r.t. any comparator $\mathbf{u} \in \mathcal{X}$ is bounded as:*

$$\sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}) \le \sum_{t=1}^T \eta_t \|\nabla f_t(\mathbf{x}_t) - M_t\|_\star^2 \quad \text{(quality of guess)}$$

$$+ \sum_{t=1}^T \frac{1}{\eta_t}\Big(\mathcal{D}_\psi(\mathbf{u}, \widehat{\mathbf{x}}_t) - \mathcal{D}_\psi(\mathbf{u}, \widehat{\mathbf{x}}_{t+1})\Big) \quad \text{(telescoping term)}$$

$$- \sum_{t=1}^T \frac{1}{\eta_t}\Big(\mathcal{D}_\psi(\widehat{\mathbf{x}}_{t+1}, \mathbf{x}_t) + \mathcal{D}_\psi(\mathbf{x}_t, \widehat{\mathbf{x}}_t)\Big) \quad \text{(negative term)}$$

Proof. $f_t(\mathbf{x}_t) - f_t(\mathbf{u}) \le \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}_t - \mathbf{u}\rangle$ (convexity)

$$= \underbrace{\langle \nabla f_t(\mathbf{x}_t) - M_t, \mathbf{x}_t - \widehat{\mathbf{x}}_{t+1}\rangle}_{\texttt{term (a)}} + \underbrace{\langle M_t, \mathbf{x}_t - \widehat{\mathbf{x}}_{t+1}\rangle}_{\texttt{term (b)}} + \underbrace{\langle \nabla f_t(\mathbf{x}_t), \widehat{\mathbf{x}}_{t+1} - \mathbf{u}\rangle}_{\texttt{term (c)}}$$

Proof of Optimistic OMD Regret

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

For term (a), we use the stability lemma.

Lemma 2 (Stability Lemma). *Consider the following updates:*

$$\begin{cases} \mathbf{x} = \arg\min_{\mathbf{x} \in \mathcal{X}} \langle \mathbf{g}, \mathbf{x} \rangle + \mathcal{D}_\psi(\mathbf{x}, \mathbf{c}) \\ \mathbf{x}' = \arg\min_{\mathbf{x} \in \mathcal{X}} \langle \mathbf{g}', \mathbf{x} \rangle + \mathcal{D}_\psi(\mathbf{x}, \mathbf{c}) \end{cases}$$

When the regularizer $\psi : \mathcal{X} \mapsto \mathbb{R}$ is a λ-strongly convex function with respect to norm $\|\cdot\|$, we have

$$\lambda \|\mathbf{x} - \mathbf{x}'\| \le \|\mathbf{g} - \mathbf{g}'\|_\star .$$

$$\begin{aligned} \texttt{term (a)} &= \langle \nabla f_t(\mathbf{x}_t) - M_t, \mathbf{x}_t - \widehat{\mathbf{x}}_{t+1} \rangle \\ &\le \|\nabla f_t(\mathbf{x}_t) - M_t\|_\star \|\mathbf{x}_t - \widehat{\mathbf{x}}_{t+1}\| \le \eta_t \|\nabla f_t(\mathbf{x}_t) - M_t\|_\star^2 \end{aligned}$$

Proof of Optimistic OMD Regret

Proof. $f_t(\mathbf{x}_t) - f_t(\mathbf{u}) \leq \underbrace{\langle \nabla f_t(\mathbf{x}_t) - M_t, \mathbf{x}_t - \widehat{\mathbf{x}}_{t+1} \rangle}_{\texttt{term (a)}} + \underbrace{\langle M_t, \mathbf{x}_t - \widehat{\mathbf{x}}_{t+1} \rangle}_{\texttt{term (b)}} + \underbrace{\langle \nabla f_t(\mathbf{x}_t), \widehat{\mathbf{x}}_{t+1} - \mathbf{u} \rangle}_{\texttt{term (c)}}$

For term (b), we adopt the Bregman Proximal lemma.

> **Lemma 3** (Bregman Proximal Inequality). *Consider convex optimization problem with the following update form*
>
> $$\mathbf{x}_{t+1} = \arg\min_{\mathbf{x} \in \mathcal{X}} \left\{ \langle \mathbf{g}_t, \mathbf{x} \rangle + \mathcal{D}_\psi(\mathbf{x}, \mathbf{x}_t) \right\}.$$
>
> *Then, it satisfies the following inequality for any* $\mathbf{u} \in \mathcal{X}$*:*
>
> $$\langle \mathbf{g}_t, \mathbf{x}_{t+1} - \mathbf{u} \rangle \leq \mathcal{D}_\psi(\mathbf{u}, \mathbf{x}_t) - \mathcal{D}_\psi(\mathbf{u}, \mathbf{x}_{t+1}) - \mathcal{D}_\psi(\mathbf{x}_{t+1}, \mathbf{x}_t).$$

Thus, according to update rule: $\mathbf{x}_t = \arg\min_{\mathbf{x} \in \mathcal{X}} \eta_t \langle M_t, \mathbf{x} \rangle + \mathcal{D}_\psi(\mathbf{x}, \widehat{\mathbf{x}}_t)$

$$\texttt{term (b)} = \langle M_t, \mathbf{x}_t - \widehat{\mathbf{x}}_{t+1} \rangle \leq \frac{1}{\eta_t}\Big(\mathcal{D}_\psi(\widehat{\mathbf{x}}_{t+1}, \widehat{\mathbf{x}}_t) - \mathcal{D}_\psi(\widehat{\mathbf{x}}_{t+1}, \mathbf{x}_t) - \mathcal{D}_\psi(\mathbf{x}_t, \widehat{\mathbf{x}}_t) \Big)$$

Proof of Optimistic OMD Regret

Proof. $f_t(\mathbf{x}_t) - f_t(\mathbf{u}) \leq \underbrace{\langle \nabla f_t(\mathbf{x}_t) - M_t, \mathbf{x}_t - \widehat{\mathbf{x}}_{t+1} \rangle}_{\texttt{term (a)}} + \underbrace{\langle M_t, \mathbf{x}_t - \widehat{\mathbf{x}}_{t+1} \rangle}_{\texttt{term (b)}} + \underbrace{\langle \nabla f_t(\mathbf{x}_t), \widehat{\mathbf{x}}_{t+1} - \mathbf{u} \rangle}_{\texttt{term (c)}}$

For term (c), we also adopt the Bregman Proximal lemma.

> **Lemma 3** (Bregman Proximal Inequality). *Consider convex optimization problem with the following update form*
>
> $$\mathbf{x}_{t+1} = \arg\min_{\mathbf{x} \in \mathcal{X}} \left\{ \langle \mathbf{g}_t, \mathbf{x} \rangle + \mathcal{D}_\psi(\mathbf{x}, \mathbf{x}_t) \right\}.$$
>
> *Then, it satisfies the following inequality for any $\mathbf{u} \in \mathcal{X}$:*
>
> $$\langle \mathbf{g}_t, \mathbf{x}_{t+1} - \mathbf{u} \rangle \leq \mathcal{D}_\psi(\mathbf{u}, \mathbf{x}_t) - \mathcal{D}_\psi(\mathbf{u}, \mathbf{x}_{t+1}) - \mathcal{D}_\psi(\mathbf{x}_{t+1}, \mathbf{x}_t).$$

Thus, according to update rule: $\widehat{\mathbf{x}}_{t+1} = \arg\min_{\mathbf{x} \in \mathcal{X}} \eta_t \langle \nabla f_t(\mathbf{x}_t), \mathbf{x} \rangle + \mathcal{D}_\psi(\mathbf{x}, \widehat{\mathbf{x}}_t)$

$$\texttt{term (c)} = \langle \nabla f_t(\mathbf{x}_t), \widehat{\mathbf{x}}_{t+1} - \mathbf{u} \rangle \leq \frac{1}{\eta_t} \Big(\mathcal{D}_\psi(\mathbf{u}, \widehat{\mathbf{x}}_t) - \mathcal{D}_\psi(\mathbf{u}, \widehat{\mathbf{x}}_{t+1}) - \mathcal{D}_\psi(\widehat{\mathbf{x}}_{t+1}, \widehat{\mathbf{x}}_t) \Big)$$

Proof of Optimistic OMD Regret

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

Put the three terms together, we can finish the proof.

$$\texttt{term (a)} \le \eta_t \left\| \nabla f_t(\mathbf{x}_t) - M_t \right\|_\star^2$$

$$\texttt{term (b)} \le \frac{1}{\eta_t}\left(\mathcal{D}_\psi(\widehat{\mathbf{x}}_{t+1}, \widehat{\mathbf{x}}_t) - \mathcal{D}_\psi(\widehat{\mathbf{x}}_{t+1}, \mathbf{x}_t) - \mathcal{D}_\psi(\mathbf{x}_t, \widehat{\mathbf{x}}_t) \right)$$

$$\texttt{term (c)} \le \frac{1}{\eta_t}\left(\mathcal{D}_\psi(\mathbf{u}, \widehat{\mathbf{x}}_t) - \mathcal{D}_\psi(\mathbf{u}, \widehat{\mathbf{x}}_{t+1}) - \mathcal{D}_\psi(\widehat{\mathbf{x}}_{t+1}, \widehat{\mathbf{x}}_t) \right)$$

□

Optimistic OMD: Generic Analysis

Theorem 4 (Regret for Optimistic OMD). *Assume ψ is 1-strongly convex w.r.t. $\|\cdot\|$, the regret of Optimistic OMD w.r.t. any comparator $\mathbf{u} \in \mathcal{X}$ is bounded as:*

$$\sum_{t=1}^{T} f_t(\mathbf{x}_t) - \sum_{t=1}^{T} f_t(\mathbf{u}) \leq \sum_{t=1}^{T} \eta_t \|\nabla f_t(\mathbf{x}_t) - M_t\|_\star^2 \quad \text{(quality of guess)}$$

$$+ \sum_{t=1}^{T} \frac{1}{\eta_t}\left(\mathcal{D}_\psi(\mathbf{u}, \widehat{\mathbf{x}}_t) - \mathcal{D}_\psi(\mathbf{u}, \widehat{\mathbf{x}}_{t+1})\right) \quad \text{(telescoping term)}$$

$$- \sum_{t=1}^{T} \frac{1}{\eta_t}\left(\mathcal{D}_\psi(\widehat{\mathbf{x}}_{t+1}, \mathbf{x}_t) + \mathcal{D}_\psi(\mathbf{x}_t, \widehat{\mathbf{x}}_t)\right) \quad \text{(negative term)}$$

Example: Optimistic OGD (with fixed step size)

- Consider the Euclidean regularizer $\mathcal{D}_\psi(\mathbf{x}, \mathbf{y}) = \frac{1}{2}\|\mathbf{x} - \mathbf{y}\|_2^2$, i.e.,:

$$\mathbf{x}_t = \arg\min_{\mathbf{x}\in\mathcal{X}} \eta \langle M_t, \mathbf{x}\rangle + \frac{1}{2}\|\mathbf{x} - \widehat{\mathbf{x}}_t\|_2^2$$

$$\widehat{\mathbf{x}}_{t+1} = \arg\min_{\mathbf{x}\in\mathcal{X}} \eta \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}\rangle + \frac{1}{2}\|\mathbf{x} - \widehat{\mathbf{x}}_t\|_2^2$$

$$\Longrightarrow \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}) \le \eta \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - M_t\|_2^2 \quad \text{(quality of guess)}$$

$$+ \frac{1}{2\eta}\sum_{t=1}^T \left(\|\mathbf{u} - \widehat{\mathbf{x}}_t\|_2^2 - \|\mathbf{u} - \widehat{\mathbf{x}}_{t+1}\|_2^2\right) \quad \text{(telescoping term)}$$

$$- \frac{1}{2\eta}\sum_{t=1}^T \left(\|\widehat{\mathbf{x}}_{t+1} - \mathbf{x}_t\|_2^2 + \|\mathbf{x}_t - \widehat{\mathbf{x}}_t\|_2^2\right) \quad \text{(negative term)}$$

Example: Optimistic OGD (with fixed step size)

- Consider the Euclidean regularizer $\mathcal{D}_\psi(\mathbf{x}, \mathbf{y}) = \frac{1}{2}\|\mathbf{x} - \mathbf{y}\|_2^2$, i.e.,:

$$\mathbf{x}_t = \arg\min_{\mathbf{x}\in\mathcal{X}} \eta \langle M_t, \mathbf{x}\rangle + \frac{1}{2}\|\mathbf{x} - \widehat{\mathbf{x}}_t\|_2^2$$

$$\widehat{\mathbf{x}}_{t+1} = \arg\min_{\mathbf{x}\in\mathcal{X}} \eta \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}\rangle + \frac{1}{2}\|\mathbf{x} - \widehat{\mathbf{x}}_t\|_2^2$$

$$\Longrightarrow \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}) \le \underbrace{\eta \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - M_t\|_2^2}_{\text{(quality of guess)}} + \frac{\|\mathbf{u} - \widehat{\mathbf{x}}_1\|_2^2}{2\eta} \underbrace{- \frac{1}{4\eta}\sum_{t=1}^T \|\mathbf{x}_{t+1} - \mathbf{x}_t\|_2^2}_{\text{(negative term)}}$$

$$\le \eta \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - M_t\|_2^2 + \frac{D^2}{2\eta} \le \mathcal{O}\left(\sqrt{1 + \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - M_t\|_2^2}\right)$$

($\eta = \frac{D}{\sqrt{\sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - M_t\|_2^2}}$, which is not available)

→ ***self-confident tuning***

Optimistic OMD: Applications

- Small-Loss Bound

- Gradient-Variance Bound

- Gradient-Variation Bound

Optimistic OMD: Applications

- Small-Loss Bound

- Gradient-Variance Bound

- Gradient-Variation Bound

Achieving Small-Loss Bound

- Recall the guarantee of optimistic OGD:

$$\mathbf{x}_t = \arg\min_{\mathbf{x}\in\mathcal{X}} \eta \langle M_t, \mathbf{x}\rangle + \frac{1}{2}\|\mathbf{x} - \widehat{\mathbf{x}}_t\|_2^2$$

$$\widehat{\mathbf{x}}_{t+1} = \arg\min_{\mathbf{x}\in\mathcal{X}} \eta \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}\rangle + \frac{1}{2}\|\mathbf{x} - \widehat{\mathbf{x}}_t\|_2^2$$

$$\Longrightarrow \quad \sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}) \leq \mathcal{O}\left(\sqrt{1 + \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - M_t\|_2^2}\right)$$

Setting $M_t = 0$ $\quad \Longrightarrow \quad$ $$\sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{u}) \leq \mathcal{O}\left(\sqrt{1 + \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t)\|_2^2}\right)$$

Achieving Small-Loss Bound

- Employing the self-bounding property of smooth and non-negative functions

Lemma 3 (Self-bounding Property). *For an L-smooth and non-negative function* $f : \mathcal{X} \mapsto \mathbb{R}_+$, *we have that*

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

Setting $M_t = 0$ $\quad \Longrightarrow \quad \sum_{t=1}^{T} f_t(\mathbf{x}_t) - \sum_{t=1}^{T} f_t(\mathbf{u}) \leq \mathcal{O}\left(\sqrt{1 + \sum_{t=1}^{T} \|\nabla f_t(\mathbf{x}_t)\|_2^2}\right)$

$$\leq \mathcal{O}\left(\sqrt{1 + L\sum_{t=1}^{T} f_t(\mathbf{x}_t)}\right)$$

Achieving Small-Loss Bound

$$\Longrightarrow \quad \sum_{t=1}^{T} f_t(\mathbf{x}_t) - \sum_{t=1}^{T} f_t(\mathbf{u}) \leq \mathcal{O}\left(\sqrt{1 + L\sum_{t=1}^{T} f_t(\mathbf{x}_t)}\right)$$

Remember how we solve the similar problem in PEA:

Small-loss bound for PEA: Proof

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

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

$$x - (\sqrt{\ln N} + 4)\sqrt{x} \leq L_T(i^\star) + 2$$

$$\left(\sqrt{x} - \frac{\sqrt{\ln N} + 4}{2}\right)^2 \leq L_T(i^\star) + 2 + \left(\frac{\sqrt{\ln N} + 4}{2}\right)^2$$

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

$$\Longrightarrow \quad \widetilde{L}_T \leq 3\ln N + L_T(i^\star) + 8\sqrt{(L_T(i^\star) + 1)\ln N}$$

(squaring both sides) □

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

Note that this algorithm requires the knowledge of $G_T \triangleq \sum_{t=1}^{T} \|\nabla f_t(\mathbf{x}_t)\|_2^2$

Achieving Small-Loss Bound

Theorem 6 (Small-loss Bound). *Assume that* $\psi(\mathbf{x}) = \frac{1}{2}\|\mathbf{x}\|_2^2$ *and* f_t *is* L*-smooth and non-negative for all* $t \in [T]$*, when setting* $\eta_t = \frac{D}{\sqrt{1+\widetilde{G}_t}}$ *and* $M_t = \mathbf{0}$*, the regret of Optimistic OMD to any comparator* $\mathbf{u} \in \mathcal{X}$ *is bounded as*

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

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*.*

Achieving Small-Loss Bound

Proof.

$$\sum_{t=1}^{T} f_t(\mathbf{x}_t) - \sum_{t=1}^{T} f_t(\mathbf{u}) \leq \sum_{t=1}^{T} \eta_t \left\|\nabla f_t(\mathbf{x}_t) - M_t\right\|_2^2 \qquad \text{(quality of guess, } \texttt{term(a)}\text{)}$$

$$+ \sum_{t=1}^{T} \frac{1}{2\eta_t}\left(\|\mathbf{u} - \widehat{\mathbf{x}}_t\|_2^2 - \|\mathbf{u} - \widehat{\mathbf{x}}_{t+1}\|_2^2\right) \qquad \text{(telescoping term, } \texttt{term(b)}\text{)}$$

$$- \sum_{t=1}^{T} \frac{1}{2\eta_t}\left(\|\widehat{\mathbf{x}}_{t+1} - \mathbf{x}_t\|_2^2 + \|\mathbf{x}_t - \widehat{\mathbf{x}}_t\|_2^2\right) \qquad \text{(negative term, } \texttt{term(c)}\text{)}$$

For term (a),

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

(self-confident tuning lemma)

$$\leq D\sqrt{1 + 2L\sum_{t=1}^{T} f_t(\mathbf{x}_t)} + G^2 \qquad \text{(self-bounding property)}$$

Achieving Small-Loss Bound

Proof. $\mathrm{Regret}_T = \texttt{term (a)} + \texttt{term (b)} - \texttt{term (c)}$

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

$$\texttt{term (b)} \leq \frac{D\sqrt{1 + 2L\sum_{t=1}^{T} f_t(\mathbf{x}_t)} + D}{2}$$

$$\texttt{term (c)} \geq 0$$

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

Achieving Small-Loss Bound

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

Remember how we solve the similar problem in PEA:

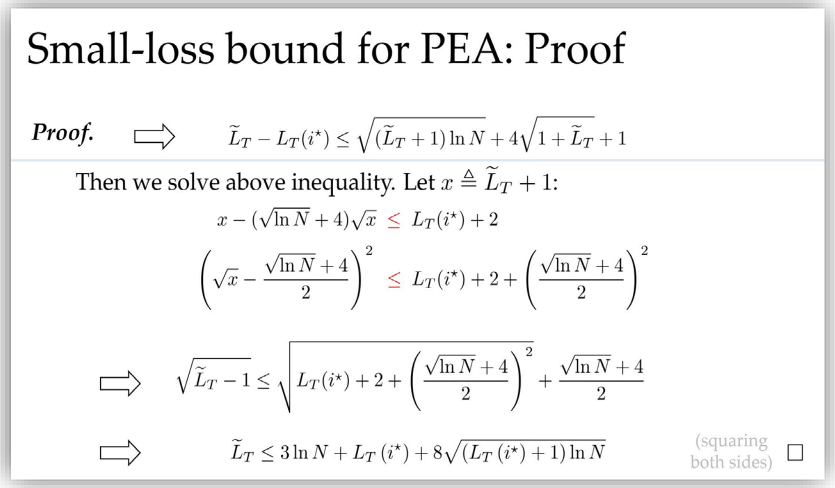
Small-loss bound for PEA: Proof

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

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

$$x - (\sqrt{\ln N} + 4)\sqrt{x} \leq L_T(i^\star) + 2$$

$$\left(\sqrt{x} - \frac{\sqrt{\ln N} + 4}{2}\right)^2 \leq L_T(i^\star) + 2 + \left(\frac{\sqrt{\ln N} + 4}{2}\right)^2$$

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

$$\Longrightarrow \quad \widetilde{L}_T \leq 3\ln N + L_T(i^\star) + 8\sqrt{(L_T(i^\star) + 1)\ln N}$$

(squaring both sides) $\square$

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

Optimistic OMD: Applications

- Small-Loss Bound
- Gradient-Variance Bound
- Gradient-Variation Bound

Achieving Gradient-Variance Bound

Definition 3 (Gradient Variance). For a space $\mathcal{X} \in \mathbb{R}^d$ and a finite T, we say that the gradient variance of a sequence of functions $f_1, \ldots, f_T \in \mathcal{F} : \mathbb{R}^d \to \mathbb{R}$ is defined as

$$\mathrm{Var}_T = \sup_{\{\mathbf{x}_1,\ldots,\mathbf{x}_T\}\in\mathcal{X}} \sum_{t=1}^{T} \|\nabla f_t(\mathbf{x}_t) - \mu_T\|_2^2$$

where $\mu_T = \arg\min_\mu \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - \mu\|_2^2 = \frac{1}{T}\sum_{t=1}^T \nabla f_t(\mathbf{x}_t)$.

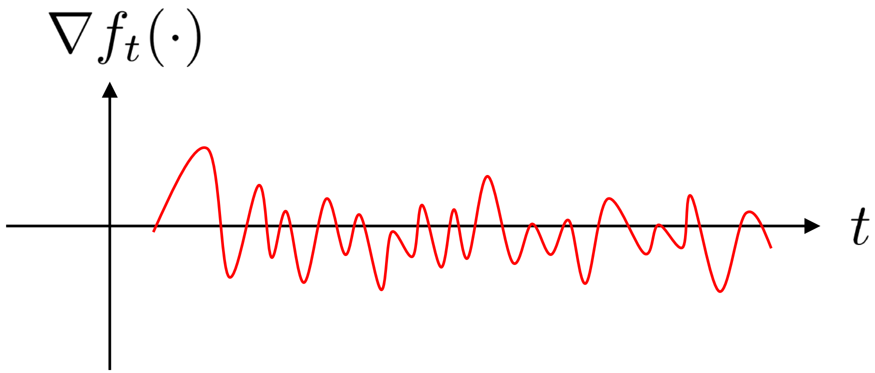

Implicit assumption: there exists a ***latent mean gradient*** $\mathbb{E}_t[\nabla f_t(\mathbf{x}_t)]$.

e.g. SGD (sampled from a set of data)
e.g. Classification (sampled from training set)

Achieving Gradient-Variance Bound

Definition 3 (Gradient Variance). For a space $\mathcal{X} \in \mathbb{R}^d$ and a finite T, we say that the gradient variance of a sequence of functions $f_1, \ldots, f_T \in \mathcal{F} : \mathbb{R}^d \to \mathbb{R}$ is defined as

$$\mathrm{Var}_T = \sup_{\{\mathbf{x}_1,\ldots,\mathbf{x}_T\} \in \mathcal{X}} \sum_{t=1}^{T} \|\nabla f_t(\mathbf{x}_t) - \mu_T\|_2^2$$

where $\mu_T = \arg\min_\mu \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - \mu\|_2^2 = \frac{1}{T} \sum_{t=1}^T \nabla f_t(\mathbf{x}_t)$.

Optimistic Online Mirror Descent

$$\mathbf{x}_t = \arg\min_{\mathbf{x} \in \mathcal{X}} \eta_t \langle M_t, \mathbf{x} \rangle + \frac{1}{2} \|\mathbf{x} - \widehat{\mathbf{x}}_t\|_2^2$$

$$\widehat{\mathbf{x}}_{t+1} = \arg\min_{\mathbf{x} \in \mathcal{X}} \eta_t \langle \nabla f_t(\mathbf{x}_t), \mathbf{x} \rangle + \frac{1}{2} \|\mathbf{x} - \widehat{\mathbf{x}}_t\|_2^2$$

Question:
How to choose M_t?

Achieving Gradient-Variance Bound

Definition 3 (Gradient Variance). For a space $\mathcal{X} \in \mathbb{R}^d$ and a finite T, we say that the gradient variance of a sequence of functions $f_1, \ldots, f_T \in \mathcal{F} : \mathbb{R}^d \to \mathbb{R}$ is defined as

$$\mathrm{Var}_T = \sup_{\{\mathbf{x}_1,\ldots,\mathbf{x}_T\} \in \mathcal{X}} \sum_{t=1}^{T} \|\nabla f_t(\mathbf{x}_t) - \mu_T\|_2^2$$

where $\mu_T = \arg\min_\mu \sum_{t=1}^T \|\nabla f_t(\mathbf{x}_t) - \mu\|_2^2 = \frac{1}{T} \sum_{t=1}^T \nabla f_t(\mathbf{x}_t)$.

Optimistic Online Mirror Descent

$$\mathbf{x}_t = \arg\min_{\mathbf{x} \in \mathcal{X}} \eta_t \langle \widetilde{\mu}_{t-1}, \mathbf{x} \rangle + \frac{1}{2}\|\mathbf{x} - \widehat{\mathbf{x}}_t\|_2^2$$

$$\widehat{\mathbf{x}}_{t+1} = \arg\min_{\mathbf{x} \in \mathcal{X}} \eta_t \langle \nabla f_t(\mathbf{x}_t), \mathbf{x} \rangle + \frac{1}{2}\|\mathbf{x} - \widehat{\mathbf{x}}_t\|_2^2$$

M_t choose to be confident estimation of gradient mean:

$\widetilde{\mu}_t = \frac{1}{t} \sum_{s=1}^t \nabla f_s(\mathbf{x}_s)$

Achieving Gradient-Variance Bound

Theorem 5 (gradient-variance bound). *Assume that* $\psi(\mathbf{x}) = \frac{1}{2}\|\mathbf{x}\|_2^2$, *when setting* $\eta_t = \frac{D}{\sqrt{1+\widetilde{\mathrm{Var}}_{t-1}}}$ *and* $M_t = \widetilde{\mu}_{t-1}$, *the regret of Optimistic OMD to any comparator* $\mathbf{u} \in \mathcal{X}$ *is bounded as*

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

where $\widetilde{\mu}_t = \frac{1}{t}\sum_{s=1}^t \nabla f_s(\mathbf{x}_s)$ *is the empirical estimator of mean, and* $\widetilde{\mathrm{Var}}_{t-1} = \sum_{s=1}^{t-1} \|\nabla f(x_s) - \widetilde{\mu}_s\|_2^2$ *is the confident estimation of variance* Var_T.

Proof.

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

$$- \sum_{t=1}^T \frac{1}{2\eta_t}\left(\|\widehat{\mathbf{x}}_{t+1} - \mathbf{x}_t\|_2^2 + \|\mathbf{x}_t - \widehat{\mathbf{x}}_t\|_2^2\right)$$

(negative term)

Achieving Gradient-Variance Bound

Proof. For term (a),

$$\sum_{t=1}^{T} \eta_t \|\nabla f_t(\mathbf{x}_t) - M_t\|_2^2 = \sum_{t=2}^{T} \eta_t \left\|\nabla f_t(\mathbf{x}_t) - \widetilde{\mu}_{t-1}\right\|_2^2 + G^2 \qquad (\eta_1 \triangleq 1)$$

$$\leq 2\sum_{t=2}^{T} \eta_t \left\|\nabla f_t(\mathbf{x}_t) - \widetilde{\mu}_t\right\|_2^2 + 2\sum_{t=2}^{T} \eta_t \left\|\widetilde{\mu}_t - \widetilde{\mu}_{t-1}\right\|_2^2 + G^2$$

$$\leq 2D\sum_{t=2}^{T} \frac{\|\nabla f_t(\mathbf{x}_t) - \widetilde{\mu}_t\|_2^2}{\sqrt{1 + \sum_{s=1}^{t-1} \|\nabla f_s(\mathbf{x}_s) - \widetilde{\mu}_s\|_2^2}} + 2D\sum_{t=2}^{T} \frac{9G^2}{t^2} + G^2 \qquad \left(\widetilde{\mu}_t = \frac{(t-1)\widetilde{\mu}_{t-1} + \nabla f_t(\mathbf{x}_t)}{t}\right) \; (\|\widetilde{\mu}_t\|_2 \leq G, \forall t \in [T]) \; (\eta_t \leq 1, \forall t \in [T])$$

$$\leq 2D\sum_{t=2}^{T} \frac{\|\nabla f_t(\mathbf{x}_t) - \widetilde{\mu}_t\|_2^2}{\sqrt{1 + \sum_{s=1}^{t-1} \|\nabla f_s(\mathbf{x}_s) - \widetilde{\mu}_s\|_2^2}} + 18DG^2 \cdot \frac{\pi^2}{6} + G^2 \qquad \left(\sum_{x=1}^{\infty} \frac{1}{x^2} = \frac{\pi^2}{6}\right)$$

Achieving Gradient-Variance Bound

Proof. $$\sum_{t=1}^{T} \eta_t \|\nabla f_t(\mathbf{x}_t) - M_t\|_2^2 \le 2D \sum_{t=2}^{T} \frac{\|\nabla f_t(\mathbf{x}_t) - \widetilde{\mu}_t\|_2^2}{\sqrt{1 + \sum_{s=1}^{t-1} \|\nabla f_s(\mathbf{x}_s) - \widetilde{\mu}_s\|_2^2}} + 18DG^2 \cdot \frac{\pi^2}{6} + G^2$$

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

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

$\Longrightarrow$

$$\le 8D\sqrt{1 + \sum_{t=1}^{T} \|\nabla f_t(\mathbf{x}_t) - \widetilde{\mu}_t\|_2^2 + 8DG^2} + 18DG^2 \cdot \frac{\pi^2}{6} + G^2$$

$$\le 8D\sqrt{1 + \mathrm{Var}_T} + (39D + 1)G^2$$

$(\mathrm{Var}_T = \sup_{\{\mathbf{x}_1,\ldots,\mathbf{x}_T\} \in \mathcal{X}} \sum_{t=1}^{T} \|\nabla f_t(\mathbf{x}_t) - \mu_T\|_2^2)$

Achieving Gradient-Variance Bound

Proof. We then analyze term (b) in the same way as before:

$$\texttt{term (b)} = \sum_{t=1}^{T} \frac{1}{2\eta_t}\left(\|\mathbf{u}-\widehat{\mathbf{x}}_t\|_2^2 - \|\mathbf{u}-\widehat{\mathbf{x}}_{t+1}\|_2^2\right)$$

$$= \sum_{t=2}^{T}\left(\frac{1}{2\eta_t}-\frac{1}{2\eta_{t-1}}\right)\|\mathbf{u}-\widehat{\mathbf{x}}_t\|_2^2 + \frac{1}{2\eta_1}\|\mathbf{u}-\widehat{\mathbf{x}}_1\|_2^2$$

$$\leq \sum_{t=2}^{T}\left(\frac{1}{2\eta_t}-\frac{1}{2\eta_{t-1}}\right)D^2 + \frac{1}{2\eta_1}D^2 \quad (\eta_t \leq \eta_{t-1} \text{ and } \|\mathbf{x}-\mathbf{y}\|_2 \leq D, \forall \mathbf{x},\mathbf{y}\in\mathcal{X})$$

$$\leq \frac{D^2}{2\eta_T} + \frac{1}{2\eta_1}D^2 \leq \frac{D\sqrt{1+\mathrm{Var}_T}+D}{2} \quad \left(\frac{1}{\eta_T} = \frac{\sqrt{1+\widetilde{\mathrm{Var}}_{T-1}}}{D} \leq \frac{\sqrt{1+\mathrm{Var}_T}}{D}\right)$$

Achieving Gradient-Variance Bound

Proof. Finally, put three terms together:

$$\mathrm{Regret}_T = \texttt{term (a)} + \texttt{term (b)} - \texttt{term (c)}$$

$$\texttt{term (a)} \leq 8D\sqrt{1 + \mathrm{Var}_T} + (39D + 1)G^2$$

$$\texttt{term (b)} \leq \frac{D^2}{2\eta_T} + \frac{1}{2\eta_1}D^2 \leq \frac{D\sqrt{1 + \mathrm{Var}_T} + D}{2}$$

$$\texttt{term (c)} \geq 0$$

$$\Longrightarrow \mathrm{Regret}_T \leq 9D\sqrt{1 + \mathrm{Var}_T} + 39DG^2 + G^2$$

Optimistic OMD: Applications

- Small-Loss Bound
- Gradient-Variance Bound
- Gradient-Variation Bound

Recall: L-smooth Function

Definition 12 (Smoothness)**.** A function f is L-smooth if, for any $\mathbf{x}, \mathbf{y} \in \operatorname{dom} f$,

$$f(\mathbf{y}) \leq f(\mathbf{x}) + \nabla f(\mathbf{x})^\top (\mathbf{y} - \mathbf{x}) + \frac{L}{2}\|\mathbf{y} - \mathbf{x}\|^2,$$

or equivalently, $\nabla^2 f(\mathbf{x}) \preceq LI$, or

$$\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\| \leq L\|\mathbf{x} - \mathbf{y}\|.$$

- For linear functions, $L = 0$.
- Smoothness is in fact the ***Lipshitzness of gradients***.

Gradient-Variation Bound

Definition 2 (Gradient Variation). For a space $\mathcal{X} \in \mathbb{R}^d$ and a finite T, we say that the gradient variation of a sequence of functions $f_1, \ldots, f_T \in \mathcal{F} : \mathbb{R}^d \to \mathbb{R}$ is defined as

$$V_T = \sum_{t=2}^{T} \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla f_t(\mathbf{x}) - \nabla f_{t-1}(\mathbf{x})\|_2^2$$

Gradient variation characterizes online functions' shifting rate.

- ***Adaptivity***: it can be small in slowly changing environments.
- ***Robustness***: $V_T \leq 4G^2T$ in the worst case. ($\|\nabla f_t(\mathbf{x})\| \leq G, \forall \mathbf{x} \in \mathcal{X}$ and $t \in [T]$)

Gradient-Variation Bound

Definition 2 (Gradient Variation)**.** For a space $\mathcal{X} \in \mathbb{R}^d$ and a finite T, we say that the gradient variation of a sequence of functions $f_1, \ldots, f_T \in \mathcal{F} : \mathbb{R}^d \to \mathbb{R}$ is defined as

$$V_T = \sum_{t=2}^{T} \sup_{\mathbf{x} \in \mathcal{X}} \|\nabla f_t(\mathbf{x}) - \nabla f_{t-1}(\mathbf{x})\|_2^2$$

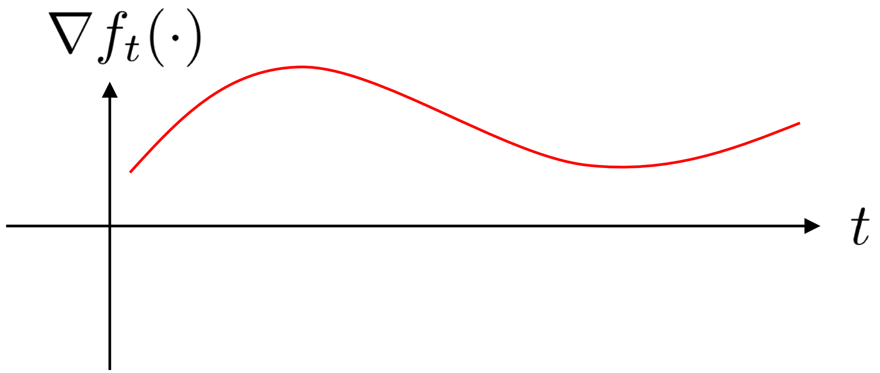

Implicit assumption:
Gradient (online function) shifts slowly

e.g., age forecasting by portraits

Optimistic OMD for Gradient-Variation Bound

Optimistic Online Mirror Descent

$$\mathbf{x}_t = \arg\min_{\mathbf{x}\in\mathcal{X}} \eta_t \langle M_t, \mathbf{x}\rangle + \frac{1}{2}\|\mathbf{x} - \widehat{\mathbf{x}}_t\|_2^2$$

$$\widehat{\mathbf{x}}_{t+1} = \arg\min_{\mathbf{x}\in\mathcal{X}} \eta_t \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}\rangle + \frac{1}{2}\|\mathbf{x} - \widehat{\mathbf{x}}_t\|_2^2$$

Question: How to choose M_t?

⇨ Imposing a prior on the change of the online functions

*setting M_t as the **last-round gradient*** $M_t = \nabla f_{t-1}(\mathbf{x}_{t-1})$

Optimistic OMD for Gradient-Variation Bound

Optimistic Online Mirror Descent

$$\mathbf{x}_t = \arg\min_{\mathbf{x}\in\mathcal{X}} \eta_t \langle M_t, \mathbf{x}\rangle + \frac{1}{2}\|\mathbf{x} - \widehat{\mathbf{x}}_t\|_2^2$$

$$\widehat{\mathbf{x}}_{t+1} = \arg\min_{\mathbf{x}\in\mathcal{X}} \eta_t \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}\rangle + \frac{1}{2}\|\mathbf{x} - \widehat{\mathbf{x}}_t\|_2^2$$

Optimistic OMD for Gradient-Variation Bound

$$\mathbf{x}_t = \arg\min_{\mathbf{x}\in\mathcal{X}} \eta_t \langle \nabla f_{t-1}(\mathbf{x}_{t-1}), \mathbf{x}\rangle + \frac{1}{2}\|\mathbf{x} - \widehat{\mathbf{x}}_t\|_2^2$$

$$\widehat{\mathbf{x}}_{t+1} = \arg\min_{\mathbf{x}\in\mathcal{X}} \eta_t \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}\rangle + \frac{1}{2}\|\mathbf{x} - \widehat{\mathbf{x}}_t\|_2^2$$

Gradient-Variation Bound

Theorem 4 (Gradient Variation Regret Bound). *Assume that* $\psi(\mathbf{x}) = \frac{1}{2}\|\mathbf{x}\|_2^2$ *and* f_t *is* *L-smooth* *for all* $t \in [T]$, *when setting* $\eta_t = \min\{\frac{1}{4L}, \frac{D}{\sqrt{1+\widetilde{V}_{t-1}}}\}$ *and* $M_t = \nabla f_{t-1}(\mathbf{x}_{t-1})$, *the regret of Optimistic OMD to any comparator* $\mathbf{u} \in \mathcal{X}$ *is*

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

where $\widetilde{V}_{t-1} = \sum_{s=2}^{t-1} \|\nabla f_s(\mathbf{x}_{s-1}) - \nabla f_{s-1}(\mathbf{x}_{s-1})\|_2^2$ *is the empirical estimates of* V_t.

Proof.

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

$$- \sum_{t=1}^T \frac{1}{2\eta_t}\left(\|\widehat{\mathbf{x}}_{t+1} - \mathbf{x}_t\|_2^2 + \|\mathbf{x}_t - \widehat{\mathbf{x}}_t\|_2^2\right)$$

(negative term)

Proof of Gradient-Variation Bound

Proof. For term 1,

$$\sum_{t=1}^{T} \eta_t \|\nabla f_t(\mathbf{x}_t) - M_t\|_2^2 \leq \sum_{t=2}^{T} \eta_t \left\|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\right\|_2^2 + G^2 \qquad (\eta_1 \triangleq 1)$$

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

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

(L-smooth)

Proof of Gradient-Variation Bound

Proof. For term 1,

$$\sum_{t=1}^{T} \eta_t \|\nabla f_t(\mathbf{x}_t) - M_t\|_2^2 \leq \sum_{t=2}^{T} \eta_t \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2 + G^2 \qquad (\eta_1 \triangleq 1)$$

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

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

(L-smooth)

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 of Gradient-Variation Bound

Proof. $\texttt{term (a)} \leq 2\sum_{t=2}^{T}\eta_t L^2 \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2 + 4D\sqrt{1 + \sum_{t=2}^{T}\|\nabla f_t(\mathbf{x}_{t-1}) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2} + (4D+1)G^2$

$$\leq 2\sum_{t=2}^{T}\eta_t L^2 \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2 + 4D\sqrt{1 + V_T} + (4D+1)G^2$$

$\left(V_T = \sum_{t=2}^{T}\sup_{\mathbf{x}\in\mathcal{X}}\|\nabla f_t(\mathbf{x}) - \nabla f_{t-1}(\mathbf{x})\|_2^2\right)$

This term depends on our algorithm,
how to deal with it?

Proof of Gradient-Variation Bound

Proof. We then analysis term (b),

$$\texttt{term (b)} = \sum_{t=1}^{T} \frac{1}{2\eta_t}\left(\|\mathbf{u} - \widehat{\mathbf{x}}_t\|_2^2 - \|\mathbf{u} - \widehat{\mathbf{x}}_{t+1}\|_2^2\right)$$

$$= \sum_{t=2}^{T} \left(\frac{1}{2\eta_t} - \frac{1}{2\eta_{t-1}}\right)\|\mathbf{u} - \widehat{\mathbf{x}}_t\|_2^2 + \frac{1}{2\eta_1}\|\mathbf{u} - \widehat{\mathbf{x}}_1\|_2^2$$

$$\leq \sum_{t=2}^{T} \left(\frac{1}{2\eta_t} - \frac{1}{2\eta_{t-1}}\right) D^2 + \frac{1}{2\eta_1} D^2 \quad (\eta_t \leq \eta_{t-1} \text{ and } \|\mathbf{x} - \mathbf{y}\|_2 \leq D, \forall \mathbf{x}, \mathbf{y} \in \mathcal{X})$$

$$\leq \frac{D^2}{2\eta_T} + \frac{1}{2\eta_1} D^2 \leq \frac{\max\{4(L+1)D, D\sqrt{1+V_T} + D\}}{2}$$

$$(\eta_T = \min\{\tfrac{1}{4L}, \tfrac{D}{\sqrt{1+\widetilde{V}_{T-1}}}\} \geq \min\{\tfrac{1}{4L}, \tfrac{D}{\sqrt{1+V_T}}\})$$

Proof of Gradient-Variation Bound

Proof. For the term (c), we have

$$\texttt{term (c)} = \sum_{t=1}^{T} \frac{1}{2\eta_t}\left(\|\widehat{\mathbf{x}}_{t+1} - \mathbf{x}_t\|_2^2 + \|\mathbf{x}_t - \widehat{\mathbf{x}}_t\|_2^2\right)$$

$$\geq \sum_{t=2}^{T} \frac{1}{2\eta_t}\left(\|\widehat{\mathbf{x}}_t - \mathbf{x}_{t-1}\|_2^2 + \|\widehat{\mathbf{x}}_t - \mathbf{x}_t\|_2^2\right) \qquad \left(\frac{1}{\eta_t} \geq \frac{1}{\eta_{t-1}}\right)$$

$$\geq \sum_{t=2}^{T} \frac{1}{4\eta_t}\|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2 \qquad (a^2 + b^2 \geq (a+b)^2/2)$$

Proof of Gradient-Variation Bound

Proof. For the term (c), we have

$$\texttt{term (c)} = \sum_{t=1}^{T} \frac{1}{2\eta_t}\left(\|\widehat{\mathbf{x}}_{t+1} - \mathbf{x}_t\|_2^2 + \|\mathbf{x}_t - \widehat{\mathbf{x}}_t\|_2^2\right)$$

$$\geq \sum_{t=2}^{T} \frac{1}{2\eta_t}\left(\|\widehat{\mathbf{x}}_t - \mathbf{x}_{t-1}\|_2^2 + \|\widehat{\mathbf{x}}_t - \mathbf{x}_t\|_2^2\right) \qquad \left(\frac{1}{\eta_t} \geq \frac{1}{\eta_{t-1}}\right)$$

$$\geq \sum_{t=2}^{T} \frac{1}{4\eta_t}\|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2 \qquad (a^2 + b^2 \geq (a+b)^2/2)$$

Does this term look familiar?

Proof of Gradient-Variation Bound

Proof. Finally, put three terms together:

$$\mathrm{Regret}_T = \texttt{term (a)} + \texttt{term (b)} - \texttt{term (c)}$$

$$\texttt{term (a)} \leq 2\sum_{t=2}^{T} \eta_t L^2 \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2 + 4D\sqrt{1+V_T} + (4D+1)G^2$$

$$\texttt{term (b)} \leq \frac{\max\{4(L+1)D, D\sqrt{1+V_T}+D\}}{2}$$

$$\texttt{term (c)} \geq \sum_{t=2}^{T} \frac{1}{4\eta_t} \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2$$

Proof of Gradient-Variation Bound

Proof. Finally, put three terms together:

$$\mathrm{Regret}_T = \texttt{term (a)} + \texttt{term (b)} - \texttt{term (c)}$$

$$\texttt{term (a)} \leq 2\sum_{t=2}^{T} \eta_t L^2 \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2 + 4D\sqrt{1+V_T} + (4D+1)G^2$$

$$\texttt{term (b)} \leq \frac{\max\{4(L+1)D, D\sqrt{1+V_T} + D\}}{2}$$

$$\texttt{term (c)} \geq \sum_{t=2}^{T} \frac{1}{4\eta_t} \|\mathbf{x}_t - \mathbf{x}_{t-1}\|_2^2 \qquad \left(\eta_t = \min\{\tfrac{1}{4L}, \tfrac{D}{\sqrt{1+\widetilde{V}_{t-1}}}\}\right)$$

$$\Longrightarrow \mathrm{Regret}_T \leq 5D\sqrt{1+V_T} + (4D+1)G^2 + 2LD = \mathcal{O}\left(\sqrt{1+V_T}\right) \qquad \square$$

Comparison

$$\mathbf{x}_t = \arg\min_{\mathbf{x}\in\mathcal{X}} \eta \langle M_t, \mathbf{x}\rangle + \frac{1}{2}\|\mathbf{x} - \widehat{\mathbf{x}}_t\|_2^2$$

$$\widehat{\mathbf{x}}_{t+1} = \arg\min_{\mathbf{x}\in\mathcal{X}} \eta \langle \nabla f_t(\mathbf{x}_t), \mathbf{x}\rangle + \frac{1}{2}\|\mathbf{x} - \widehat{\mathbf{x}}_t\|_2^2$$

Different priors are imposed by designing suitable M_t for specific environments.

	Assumption(s)	**Setting of Optimism**	**Setting of η_t**	**Adaptive Regret Bound**
Small-loss Bound	L-Smooth + Non-negative	$M_t = \mathbf{0}$	$\approx \frac{D}{\sqrt{1+\widetilde{G}_t}}$	$\mathcal{O}\left(\sqrt{1+F_T}\right)$
Variance Bound	–	$M_t = \widetilde{\mu}_{t-1}$	$\approx \frac{D}{\sqrt{1+\mathrm{Var}_{t-1}}}$	$\mathcal{O}\left(\sqrt{1+\mathrm{Var}_T}\right)$
Variation Bound	L-Smooth	$M_t = \nabla f_{t-1}(\mathbf{x}_{t-1})$	$\approx \frac{D}{\sqrt{1+\widetilde{V}_{t-1}}}$	$\mathcal{O}\left(\sqrt{1+V_T}\right)$

Recovering Gradient-Variance Bound

By using algorithm for gradient-variation Bound (OMD with $M_t = \nabla f_{t-1}(\mathbf{x}_{t-1})$):

$$\sum_{t=1}^{T} \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2 \leq 3 \sum_{t=1}^{T} \|f_t(\mathbf{x}_t) - \widetilde{\mu}_t\|_2^2 \qquad (\leq 3\,\mathrm{Var}_T)$$

$$(\approx V_T)$$

$$+ 3 \sum_{t=1}^{T} \|f_{t-1}(\mathbf{x}_{t-1}) - \widetilde{\mu}_{t-1}\|_2^2 \qquad (\leq 3\,\mathrm{Var}_T)$$

$$+ 3 \sum_{t=1}^{T} \|\widetilde{\mu}_t - \widetilde{\mu}_{t-1}\|_2^2 \qquad (\leq 3 \cdot \tfrac{\pi^2}{6})$$

$\Longrightarrow$ algorithm with V_T bound can also attain Var_T bound

Recovering Small-Loss Bound

By using algorithm for gradient-variation Bound (OMD with $M_t = \nabla f_{t-1}(\mathbf{x}_{t-1})$):

$$\sum_{t=1}^{T} \|\nabla f_t(\mathbf{x}_t) - \nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2 \leq 2\sum_{t=1}^{T} \|\nabla f_t(\mathbf{x}_t)\|_2^2 + 2\sum_{t=2}^{T} \|\nabla f_{t-1}(\mathbf{x}_{t-1})\|_2^2 \quad ((a+b)^2 <= 2(a^2+b^2))$$

$(\approx V_T)$

$$\leq 4L\sum_{t=1}^{T} f_t(\mathbf{x}_t) + 4L\sum_{t=2}^{T} f_{t-1}(\mathbf{x}_{t-1}) \quad \text{(self-bounding property)}$$

$$\leq 8LF_T$$

$\Longrightarrow$ algorithm with V_T bound can also attain F_T bound

Variation-type Bounds: History Bits

Extracting Certainty from Uncertainty: Regret Bounded by Variation in Costs

Elad Hazan
IBM Almaden
650 Harry Rd, San Jose, CA 95120
hazan@us.ibm.com

Satyen Kale
Microsoft Research
1 Microsoft Way, Redmond, WA 98052
sakale@microsoft.com

Abstract

Prediction from expert advice is a fundamental problem in machine learning. A major pillar of the field is the existence of learning algorithms whose average loss approaches that of the best expert in hindsight (in other words, whose average regret approaches zero). Traditionally the regret of online algorithms was bounded in terms of the number of prediction rounds.

Cesa-Bianchi, Mansour and Stoltz [4] posed the question whether it is be possible to bound the regret of an online algorithm by the *variation* of the observed costs. In this paper we resolve this question, and prove such bounds in the fully adversarial setting, in two important online learning scenarios: prediction from expert advice, and online linear optimization.

1 Introduction

A cornerstone of modern machine learning are algorithms for prediction from expert advice. The seminal work of Littlestone and Warmuth [12], Vovk [13] and Freund and Schapire [6] gave algorithms which, under fully adversarial cost sequences, attain average cost approaching that of the best expert in hindsight.

To be more precise, consider a prediction setting in which an online learner has access to n experts. Iteratively, the learner may chose the advice of any expert deterministically or randomly. After choosing a course of action, an adversary reveals the cost of following the advice of the different experts, from which the expected cost of the online learner is derived. The classic results mentioned above give algorithms which sequentially produce randomized decisions, such that the difference between the (expected) cost of the algorithm and the best expert in hindsight grows like $O(\sqrt{T \log n})$, where T is the number of prediction iterations. This extra additive cost is known as the regret of the online learning algorithm.

However, *a priori* it is not clear why online learning algorithms should have high regret (growing with the number of iterations) in an unchanging environment. As an extreme example, consider a setting in which there are only two experts. Suppose that the first expert always incurs cost 1, whereas the second expert always incurs cost $\frac{1}{2}$. One would expect to "figure out" this pattern quickly, and focus on the second expert, thus incurring a total cost that is at most $\frac{T}{2}$ plus at most a constant extra cost (irrespective of the number of rounds T), thus having only constant regret. However, any straightforward application of previously known analyses of expert learning algorithms only gives a regret bound of $\Theta(\sqrt{T})$ in this simple case (or very simple variations of it).

More generally, the natural bound on the regret of a "good" learning algorithm should depend on *variation* in the sequence of costs, rather than purely on the number of iterations. If the cost sequence has low variation, we expect our algorithm to be able to perform better.

This intuition has a direct analog in the stochastic setting: here, the sequence of experts' costs are independently sampled from a distribution. In this situation, a natural bound on the rate of convergence to the optimal expert is controlled by the variance of the distribution (low variance should imply faster convergence). This was formalized by Cesa-Bianchi, Mansour and Stoltz [4], who assert that *"proving such a rate in the fully adversarial setting would be a fundamental result"*.

In this paper we prove the first such regret bounds on online learning algorithms in two important scenarios: prediction from expert advice, and the more general framework of online linear optimization. Our algorithms have regret bounded by the variation of the cost sequence, in a manner that is made precise in the following sections. Thus, our bounds are tighter than *all* previous bounds, and hence yield better bounds on the applications of previous bounds (see, for example, the applications in [4]).

1.1 Online linear optimization

Online linear optimization [10] is a general framework for online learning which has received much attention recently. In this framework the decision set is an arbitrary bounded, closed, convex set in Euclidean space $K \subseteq \mathbb{R}^n$ rather than a fixed set of experts, and the costs are determined by adversarially constructed vectors, $f_1, f_2, \ldots \in \mathbb{R}^n$, such that the cost of point $x \in K$ is given by $f_t \cdot x$. The online learner iteratively chooses a point in the convex set $x_t \in K$, and then the cost vector f_t is revealed and the cost $f_t \cdot x_t$ is occurred. The performance of online learning algorithms is measured by the regret, which is defined as the difference in the total cost of the sequence of points chosen by the algorithm, viz.

JMLR: Workshop and Conference Proceedings vol 23 (2012) 6.1–6.20 25th Annual Conference on Learning Theory

Online Optimization with Gradual Variations

Chao-Kai Chiang[1,2] CHAOKAI@IIS.SINICA.EDU.TW
Tianbao Yang[3] YANGTIA1@MSU.EDU
Chia-Jung Lee[1] LEECJ@IIS.SINICA.EDU.TW
Mehrdad Mahdavi[3] MAHDAVIM@CSE.MSU.EDU
Chi-Jen Lu[1] CJLU@IIS.SINICA.EDU.TW
Rong Jin[3] RONGJIN@CSE.MSU.EDU
Shenghuo Zhu[4] ZSH@SV.NEC-LABS.COM
[1] *Institute of Information Science, Academia Sinica, Taipei, Taiwan.*
[2] *Department of Computer Science and Information Engineering, National Taiwan University, Taipei, Taiwan.*
[3] *Department of Computer Science and Engineering Michigan State University, East Lansing, MI, 48824, USA*
[4] *NEC Laboratories America Cupertino, CA, 95014, USA*

Editor: Shie Mannor, Nathan Srebro, Robert C. Williamson

Abstract

We study the online convex optimization problem, in which an online algorithm has to make repeated decisions with convex loss functions and hopes to achieve a small regret. We consider a natural restriction of this problem in which the loss functions have a small deviation, measured by the sum of the distances between every two consecutive loss functions, according to some distance metrics. We show that for the linear and general smooth convex loss functions, an online algorithm modified from the gradient descend algorithm can achieve a regret which only scales as the square root of the deviation. For the closely related problem of prediction with expert advice, we show that an online algorithm modified from the multiplicative update algorithm can also achieve a similar regret bound for a different measure of deviation. Finally, for loss functions which are strictly convex, we show that an online algorithm modified from the online Newton step algorithm can achieve a regret which is only logarithmic in terms of the deviation, and as an application, we can also have such a logarithmic regret for the portfolio management problem.
Keywords: Online Learning, Regret, Convex Optimization, Deviation.

1. Introduction

We study the online convex optimization problem in which a player has to make decisions iteratively for a number of rounds in the following way. In round t, the player has to choose a point x_t from some convex feasible set $\mathcal{X} \subseteq \mathbb{R}^N$, and after that the player receives a convex loss function f_t and suffers the corresponding loss $f_t(x_t) \in [0, 1]$. The player would like to have an online algorithm that can minimize its regret, which is the difference between the total loss it suffers and that of the best fixed point in hindsight. It is known

COLT 2012
best student paper award

Extracting Certainty from Uncertainty: Regret Bounded by Variation in Costs. COLT 2008.

Online Optimization with Gradual Variations. COLT 2012.

Summary

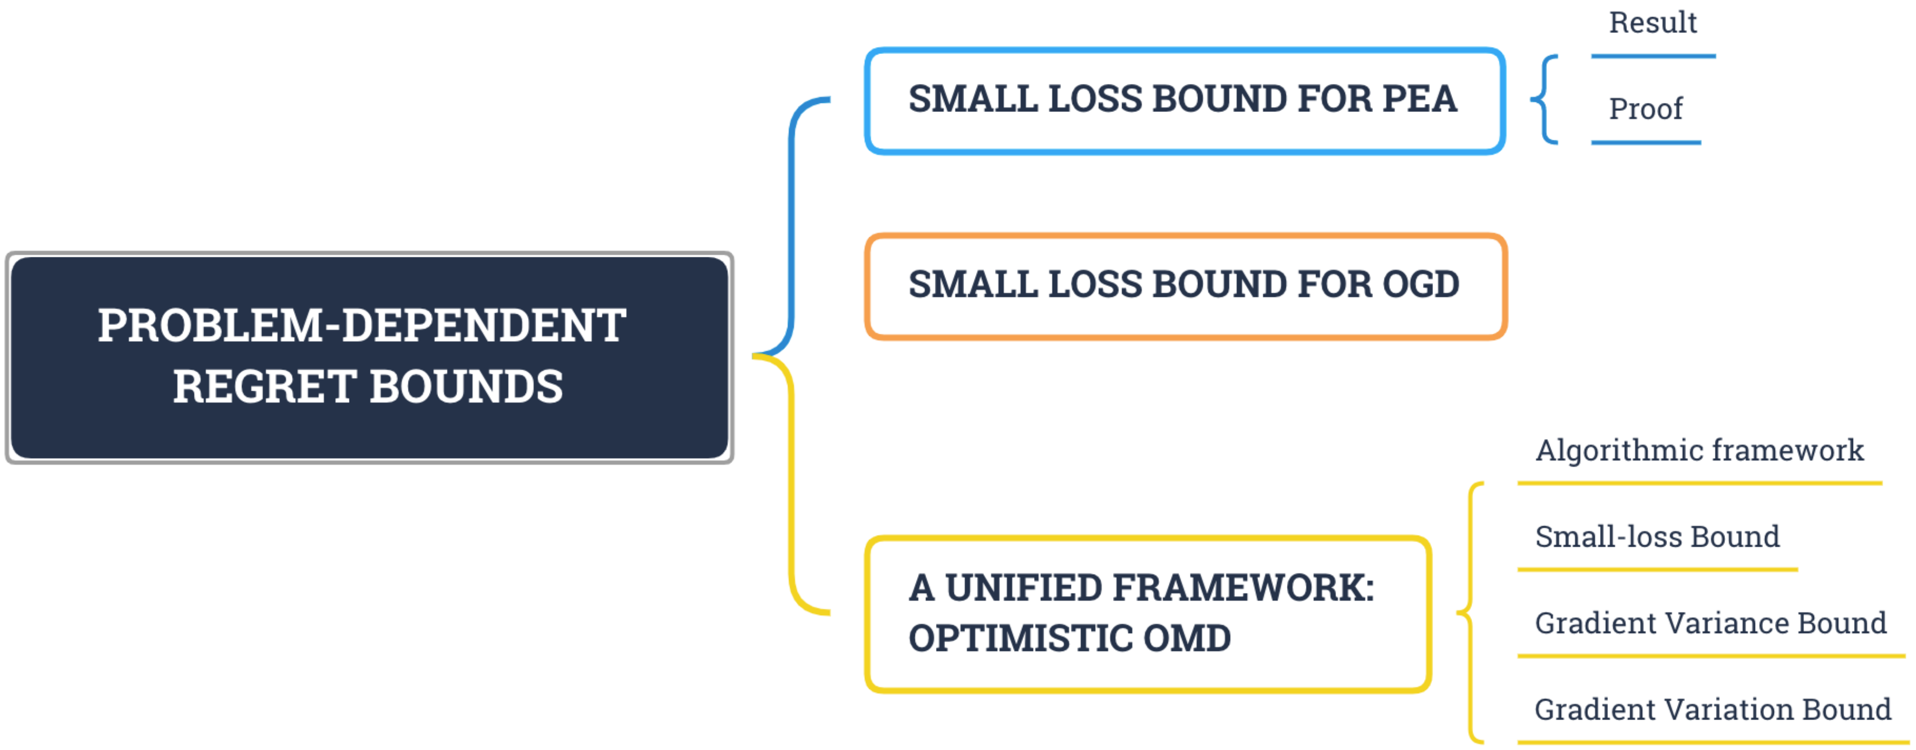

Q & A

Thanks!