Best-Case Lower Bounds in Online Learning

Cristóbal Guzmán
University of Twente

Nishant Mehta
University of Victoria

Ali Mortazavi
University of Victoria

Problem Setting

Online Convex Optimization

In round t:

(1) Learner plays action w_t in compact, convex set $\mathcal{W} \subset \mathbb{R}^d$

(2) Nature plays convex loss function f_t

(3) Learner suffers loss $f_t(w_t)$

Goal: Obtain low REGRET $\sum_{t=1}^{T} f_t(w_t) - \inf_{w \in \mathcal{W}} \sum_{t=1}^{T} f_t(w)$

Follow the Regularized Leader with adaptive regularizers

Let $(\Phi_t)_{t \geq 1}$ be a sequence of regularizers such that each Φ_t can depend on $(f_s)_{s \leq t}$. In round t, adaptive FTRL plays action:

$$w_{t+1} = \operatorname*{argmin}_{w \in \mathcal{W}} \left\{ \sum_{s=1}^{t} f_s(w) + \Phi_t(w) \right\}$$

THEOREM (McMahan, 2017)
Let each Φ_t be 1-strongly convex with respect to a norm $\| \cdot \|_{(t)}$.
If each Φ_t is nonnegative, the regret of adaptive FTRL is at most

$$\mathcal{R}_T \leq \Phi_T(w^*) - \inf_{w \in \mathcal{W}} \Phi_T(w) + \frac{1}{2} \sum_{t=1}^{T} \|\nabla f_t(w_t)\|_{(t-1),*}^2$$

If nonnegativity might not hold but the regularizers are point wise non-increasing, then the regret of adaptive FTRL is at most

$$\mathcal{R}_T \leq \Phi_T(w^*) + \frac{1}{2} \sum_{t=1}^{T} \|\nabla f_t(w_t)\|_{(t-1),*}^2$$

Adaptivity in Online Learning

Many notions of adaptivity in online learning

Adapting to time horizon
(Decreasing Hedge, DTOL)

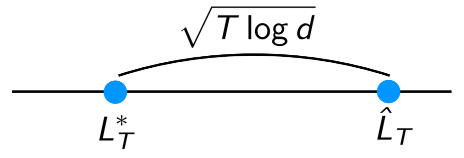

Adapting to small losses
("First-order Hedge", DTOL)

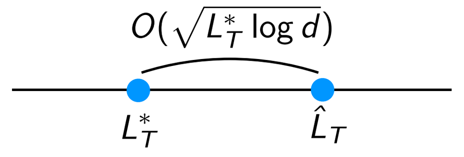

Adapting to geometry
(AdaGrad, OCO)

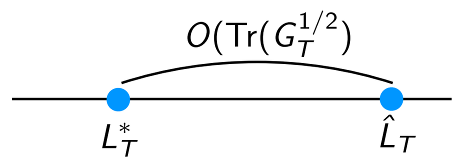

A fundamentally different notion of adaptivity:
Adapting by competing with stronger comparator

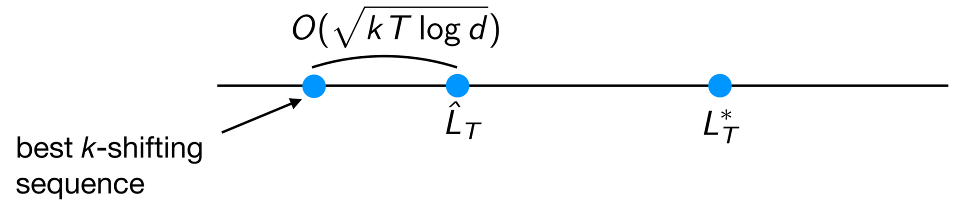

General best-case lower bound

KEY CONCEPT: REGULARIZED REGRET

$$\mathcal{R}_t^{\Phi_t} = \sum_{s=1}^{t} f_s(w_s) - \inf_{w \in \mathcal{W}} \left\{ \sum_{s=1}^{t} f_s(w) + \Phi_t(w) \right\}$$

THEOREM
Consider the adaptive FTRL strategy. Let $(\alpha_t)_{t \in [T]}$ be a sequence such that $\Phi_t(w_t) \leq \Phi_{t-1}(w_t) + \alpha_t$ for all $t \in [T]$. Then

$$\mathcal{R}_T \geq \inf_{w \in \mathcal{W}} \Phi_T(w) - \inf_{w \in \mathcal{W}} \Phi_0(w) - \sum_{t=1}^{T} \alpha_t$$

Proof sketch
Step 1: Regret is not much smaller than regularized regret

$$\mathcal{R}_t \geq \mathcal{R}_t^{\Phi_t} + \inf_{w \in \mathcal{W}} \Phi_t(w)$$

Step 2: Regularized regret is approximately non-decreasing

$$\begin{aligned} \mathcal{R}_{t+1}^{\Phi_{t+1}} &\geq \sum_{s=1}^{t+1} f_s(w_s) - \left(\sum_{s=1}^{t+1} f_s(w_{t+1}) + \Phi_{t+1}(w_{t+1}) \right) \\ &\geq \sum_{s=1}^{t} f_s(w_s) - \left(\sum_{s=1}^{t} f_s(w_{t+1}) + \Phi_t(w_{t+1}) \right) - \alpha_t = \mathcal{R}_t^{\Phi_t} - \alpha_t \end{aligned}$$

Therefore:

$$\begin{aligned} \mathcal{R}_T \geq \inf_{w \in \mathcal{W}} \Phi_T(w) + \mathcal{R}_T^{\Phi_T} &\geq \inf_{w \in \mathcal{W}} \Phi_T(w) + \mathcal{R}_0^{\Phi_0} - \sum_{t=1}^{T} \alpha_t \\ &= \inf_{w \in \mathcal{W}} \Phi_T(w) - \inf_{w \in \mathcal{W}} \Phi_0(w) - \sum_{t=1}^{T} \alpha_t \end{aligned}$$

Applications

Constant regularizer ($\Phi_t = \Phi$ for all t): Nonnegative regret!
Known for OLO from Gofer and Mansour, 2016)
Includes Follow the Leader as special case (no regularization)

Decreasing Hedge for DTOL: $-\frac{1}{2}\sqrt{T \log d} \leq \mathcal{R}_T \leq \sqrt{T \log d}$

Timeless Hedge for DTOL

$$\min \left\{ 0, -\sqrt{\frac{L_T^* \log d}{2}} + 4 \log d \right\} \leq \mathcal{R}_T = O\left(\sqrt{L_T^* \log d} + \log d \right)$$

Adaptive Gradient FTRL $w_{t+1} = \operatorname*{argmin}_{w \in \mathcal{W}} \left\{ \sum_{s=1}^{t} f_s(w) + \frac{1}{2\eta} \|w\|_{\delta \mathrm{I} + G_t^{1/2}}^2 \right\}$

$$-\frac{D_2}{2} \mathrm{Tr}(G_T^{1/2}) \leq \mathcal{R}_T \leq \frac{D_2 M_2}{2} + D_2 \mathrm{Tr}(G_T^{1/2})$$

Group Fair Online Learning

GROUP FAIR DECISION-THEORETIC ONLINE LEARNING

OLO on simplex, assume experts are group fair:

Goal: Obtain low regret while being group fair

Previous results (Blum, Gunasekar, Lykouris, Srebro)
Fixed horizon group fairness
"Interleaved Hedge"

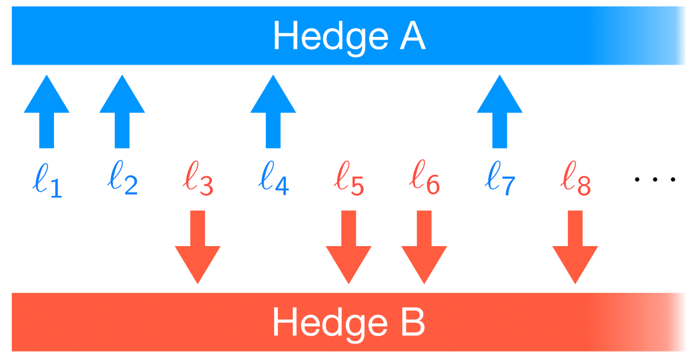

(I) Experts are group fair
$\Rightarrow$ best error rate for group A $\approx$ best error rate for group B

(II) Hedge with constant learning rate has low worst-case regret
AND nonnegative best-case regret!

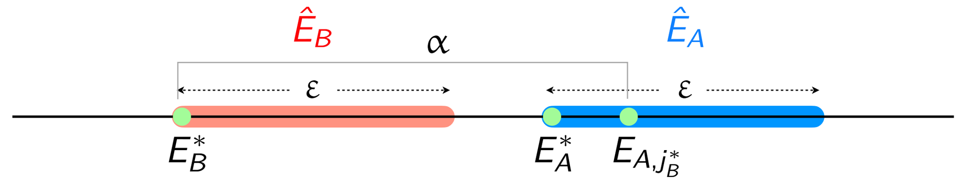

Problem with fixed horizon fairness

Learner's loss

Group A	1		1		1		1		0		0		0		0	
Group B		0		0		0		0		1		1		1		1

An unfair stopping point

New result: Decreasing Hedge "enjoys" best-case lower bound
Interleaved Decreasing Hedge can satisfy anytime group fairness

Anytime equalized error rate:

For all time horizons $t = 1, 2, \ldots, T$:
Learner's error rate for group A up until time t $\approx$ Learner's error rate for group B up until time t

If all experts are anytime group fair, Interleaved Decreasing Hedge is anytime group fair up to BLAH and has regret BLAH

Adaptivity of Linearized FTRL

Our results do not hold for Linearized FTRL, and there is a simple example where Linearized FTRL has negative linear regret!

Brier game: f_t(w_t) = (w_t - y_t)^2

THEOREM
Exponentiated gradient method has negative linear regret

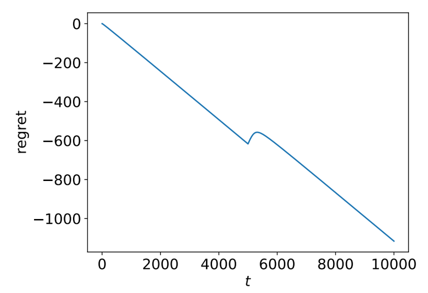

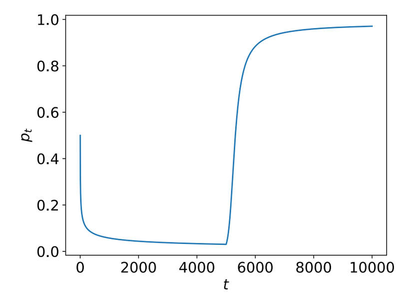

Best-case sequence for binary DTOL

DECREASING HEDGE FOR DTOL WITH 2 EXPERTS & BINARY LOSSES

THEOREM
Any loss sequence like $\binom{1}{0}\binom{0}{1}\binom{0}{0}\cdots\binom{0}{1}$ can be converted to

canonical best-case sequence $\binom{0}{1}^{k-1}\binom{0}{0}^{2}\binom{1}{0}^{k-1}$

without increasing the regret of Decreasing Hedge

Op 1: Remove leader switches

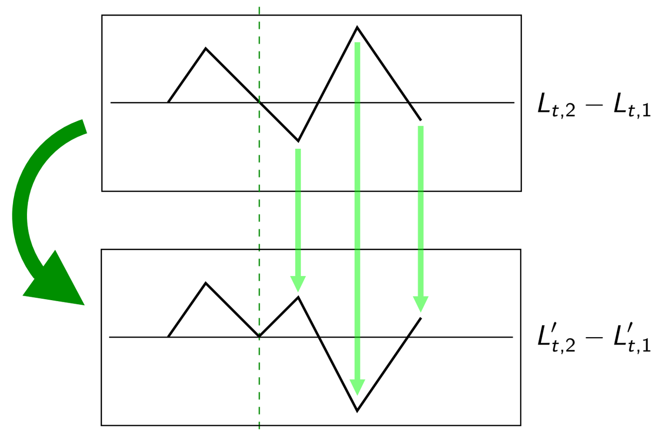

Op 2: Increase leader's lead

Since expert 1 is leader,
Push leader's loss to later
Pull non-leaders loss to earlier

$$\ell_1 \ldots \ell_T = \binom{\cdots}{\cdots}\binom{1}{0}\binom{0}{1}\binom{\cdots}{\cdots}$$

$$\ell'_1 \ldots \ell'_T = \binom{\cdots}{\cdots}\binom{0}{1}\binom{1}{0}\binom{\cdots}{\cdots}$$

Op 3: Convert pairs of zero vectors

$$\binom{0}{1}\cdots\binom{0}{1}\binom{0}{0}\binom{0}{0}\cdots\binom{0}{0}\binom{0}{0}\binom{1}{0}\cdots\binom{1}{0}$$

$$\binom{0}{1}\cdots\binom{0}{1}\binom{0}{1}\binom{0}{0}\cdots\binom{0}{0}\binom{1}{0}\binom{1}{0}\cdots\binom{1}{0}$$

Decreasing Hedge's regret for best-case sequence

$$-12\sqrt{T} - \frac{1}{2} \leq \mathcal{R}_T \leq -0.2\sqrt{T} + 8$$