CS861: Theoretical Foundations of Machine Learning

Chapter 6: Online Convex Optimization

Kirthevasan Kandasamy

UW-Madison

Outline

1. Set up
2. Follow the regularized leader
3. Follow the perturbed leader
4. Online Shortest Paths

Ch 6.1: Online Convex Optimization

Example 1. *Online linear classification.* We are given a parameter space $\Theta = \{\theta \in \mathbb{R}^d ; \|\theta\|_2 \le 1\}$.

- On each round t, a learner chooses $\theta_t \in \Theta$.
- Simultaneously, the environment picks $(x_t, y_t) \in \mathbb{R}^d \times \{-1, +1\}$.
- The learner incurs the hinge loss $f_t(\theta_t) = \max(0, 1 - y_t \theta_t^\top x_t)$.
- The learner observes (x_t, y_t), and therefore the loss for all $\theta \in \Theta$.

We can define the regret for a policy π as,

$$R_T \left(\pi, \{(x_t, y_t)\}_{t=1}^T\right) = \sum_{t=1}^T f_t(\theta_t) - \min_{\theta \in \Theta} \sum_{t=1}^T f_t(\theta)$$

Online Convex Optimization (cont'd)

Example 2. *The experts problem.* There are K experts.

- On each round t, a learner chooses a probability vector $p_t \in \Delta([K]) = \{p \in \mathbb{R}_+^K; p^\top \mathbf{1}_K = 1\}$.
- Simultaneously, the environment picks a loss vector $\ell_t \in [0,1]^K$.
- Learner incurs loss $\ell_t^\top p_t$ (in expectation, when an expert is sampled from p_t).
- Learner observes ℓ_t, and therefore the losses for all $p_t \in \Delta([K])$.

We can define the regret for a policy π as,

$$R_T\left(\pi, \{\ell_t\}_{t=1}^T\right) = \sum_{t=1}^T p_t^\top \ell_t - \min_{a \in [K]} \sum_{t=1}^T \ell_t(a) = \sum_{t=1}^T p_t^\top \ell_t - \min_{p \in \Delta([K])} \sum_{t=1}^T p^\top \ell_t.$$

Online Convex Optimization

Consider the following framework for online learning:

- Let $\Omega \subset \mathbb{R}^d$.

- On each round t, a learner chooses weight vector $\omega_t \in \Omega$.

- Simultaneously, the environment picks a loss function $f_t : \Omega \to \mathbb{R}$.

- The player incurs loss $f_t(\omega_t)$, but observes f_t (losses for all $\omega \in \Omega$).

Online convex optimization: When Ω is a convex set and $f = (f_1, \ldots, f_T)$ are convex functions, this framework is called online convex optimization.

Regret

Let $f = (f_1, \ldots, f_T)$, be an arbitrary sequence of loss functions. We define regret relative to the best fixed weight vector in hindsight.

$$R_T(\pi, f) = \mathbb{E}\left[\sum_{t=1}^{T} f_t(\omega_t)\right] - \min_{\omega \in \Omega} \sum_{t=1}^{T} f_t(\omega).$$

Here, the expection is with respect to any randomization in the policy.

We wish to design π to achieve small $\sup_f R_T(\pi, f)$.

Examples revisited

Example 1. Online linear classification. Given a parameter space $\Theta = \{\theta \in \mathbb{R}^d; \|\theta\|_2 \leq 1\}$.

- Learner chooses $\theta_t \in \Theta$. Simultaneously, the environment picks $(x_t, y_t) \in \mathbb{R}^d \times \mathbb{R}$.
- Learner incurs hinge loss $f_t(\theta_t) = \max(0, 1 - y_t\theta_t^\top x_t)$.
- The learner observes (x_t, y_t).

Here Θ is convex, and the hinge loss is convex in θ.

Example 2. The experts problem (with minor adjustments). There are K experts.

- Learner chooses a probability vector $p_t \in \Delta([K]) = \{p \in \mathbb{R}_+^K; p^\top \mathbf{1}_K = 1\}$. Simultaneously, the environment picks a loss vector $\ell_t \in [0, 1]^K$.
- Learner incurs loss $\ell_t^\top p_t$ and observes ℓ_t.

Here $\Delta([K])$ is convex, and $p_t^\top \ell_t$ is linear (convex).

Follow the (regularized) leader, FTL/FTRL

What is the most straightforward approach?

$$\text{Choose} \qquad \omega_t \in \operatorname*{argmin}_{\omega \in \Omega} \sum_{s=1}^{t-1} f_s(\omega)$$

- This is called *follow the leader* (FTL).

- But this often fails as the chosen weights could fluctuate from round to round (we will see examples shortly).

Follow the regularized leader (FTRL). Stabilize FTL by adding a regularizer $\Lambda(\omega)$,

$$\omega_t \in \operatorname*{argmin}_{\omega \in \Omega} \left(\sum_{s=1}^{t-1} f_s(\omega) + \Lambda(\omega) \right).$$

Follow the (regularized) leader

Follow the regularized leader (FTRL).

$$\omega_t \in \operatorname*{argmin}_{\omega \in \Omega} \left(\sum_{s=1}^{t-1} f_s(\omega) + \Lambda(\omega) \right).$$

Theorem (FTRL). For any $u \in \Omega$, FTRL satisfies

$$\sum_{t=1}^{T} f_t(\omega_t) - \sum_{t=1}^{T} f_t(u) \le \sum_{t=1}^{T} \left(f_t(\omega_t) - f_t(\omega_{t+1})\right) + \Lambda(u) - \min_{\omega \in \Omega} \Lambda(\omega).$$

N.B. Note that FTRL is a deterministic policy. Moreover, this result does not assume convexity of Ω, f_t or Λ.

Proof of FTRL Lemma

FTRL: $\omega_t \in \operatorname{argmin}_{\omega\in\Omega}\left(\sum_{s=1}^{t-1} f_s(\omega) + \Lambda(\omega)\right)$.

Denote $F_t(\omega) = \sum_{s=1}^{t} f_s(\omega) + \Lambda(\omega)$. Therefore, $\omega_{t+1} \in \operatorname{argmin}_\omega F_t(\omega)$.

Let $\Phi_t = \min_{\omega\in\Omega} F_t(\omega) = F_t(\omega_{t+1})$. Now consider,

$$\begin{aligned}
\Phi_{t-1} - \Phi_t &= F_{t-1}(\omega_t) - F_t(\omega_{t+1}) \\
&= F_{t-1}(\omega_t) - (F_{t-1}(\omega_{t+1}) + f_t(\omega_{t+1})) \\
&\leq -f_t(\omega_{t+1}) \qquad \texttt{As } \omega_t \in \operatorname*{argmin}_{\omega\in\Omega} F_{t-1}(\omega)
\end{aligned}$$

Therefore,

$$\Phi_{t-1} - \Phi_t + f_t(\omega_t) \leq f_t(\omega_t) - f_t(\omega_{t+1}).$$

Summing from $t = 1, \ldots, T$ gives,

$$\Phi_0 - \Phi_T + \sum_{t=1}^{T} f_t(\omega_t) \leq \sum_{t=1}^{T} \left(f_t(\omega_t) - f_t(\omega_{t+1})\right).$$

Proof of FTRL Lemma (cont'd)

Recall: *(i)* $\Phi_t = \min_{\omega} F_t(\omega)$, *(ii)* $F_t(\omega) = \left(\sum_{s=1}^{t} f_s(\omega) + \Lambda(\omega)\right)$,

We also have,

$$\Phi_0 = \min_{\omega\in\Omega} F_0(\omega) = \min_{\omega\in\Omega} \Lambda(\omega).$$

$$\Phi_T = \min_{\omega\in\Omega} F_T(\omega) \le F_T(u) = \sum_{t=1}^{T} f_t(u) + \Lambda(u).$$

This yields,

$$\min_{\omega\in\Omega} \Lambda(\omega) \; - \; \sum_{t=1}^{T} f_t(u) - \Lambda(u) \; + \; \sum_{t=1}^{T} f_t(\omega_t) \; \le \; \sum_{t=1}^{T} (f_t(\omega_t) - f_t(\omega_{t+1}))$$

The theorem follows by rearranging the terms,

$$\sum_{t=1}^{T} f_t(\omega_t) \; - \; \sum_{t=1}^{T} f_t(u) \le \sum_{t=1}^{T} (f_t(\omega_t) - f_t(\omega_{t+1})) \; + \; \Lambda(u) \; - \; \min_{\omega\in\Omega} \Lambda(\omega).$$

FTRL vs FTL

Theorem (FTRL). For all $u \in \Omega$, FTRL satisfies

$$\sum_{t=1}^{T} f_t(\omega_t) - \sum_{t=1}^{T} f_t(u) \leq \sum_{t=1}^{T} \left(f_t(\omega_t) - f_t(\omega_{t+1})\right) + \Lambda(u) - \min_{\omega \in \Omega} \Lambda(\omega).$$

This means, for FTL the regret satisfies

$$R_T(\pi^{\mathrm{FTL}}, f) \leq \sum_{t=1}^{T} \left(f_t(\omega_t) - f_t(\omega_{t+1})\right).$$

- This bound suggests that if ω_t fluctuates frequently, the regret will be bad.
- The purpose behind the regularizer is to stabilize ω_t.

To motivate how a regularizer will be chosen, we will consider 3 examples for FTL with $\Omega = [0, 1]$ and when the losses are bounded, $f_t : [0, 1] \to [0, 1]$.

Example 1: FTL with linear losses

$$\text{Example 1:} \qquad \Omega = [0,1], \qquad f_t(\omega) = \begin{cases} \frac{1}{2}\omega & \text{if } t = 1, \\ \omega & \text{if } t \text{ is odd, } t > 1, \\ 1 - \omega & \text{if } t \text{ is even.} \end{cases}$$

We have

$$F_t(\omega) = \sum_{t=1}^{T} f_s(\omega) = \begin{cases} \frac{1}{2}\omega + \frac{t-1}{2} & \text{if } t \text{ is odd,} \\ -\frac{1}{2}\omega + \frac{t}{2} & \text{if } t \text{ is even,} \end{cases}$$

Therefore,

$$\omega_t = \operatorname*{argmin}_{\omega \in [0,1]} F_{t-1}(\omega) = \begin{cases} 1 & \text{if } t \text{ is odd,} \\ 0 & \text{if } t \text{ is even} \end{cases}$$

Total loss of FTL is at least $T - 1$.
Best action $\omega \in [0,1]$ in hindsight will have regret at most $T/2$.
Therefore, regret is at least $T/2 - 1$.
Bound from the theorem: $\sum_{t=1}^{T} (f_t(\omega_t) - f_t(\omega_{t+1})) \approx \sum_t 1 = T$.

Example 2: FTL with quadratic losses

$$\text{Example 2:} \qquad \Omega = [0,1], \qquad f_t(\omega) = \begin{cases} \omega^2 & \text{if } t \text{ is odd,} \\ (1-\omega)^2 & \text{if } t \text{ is even.} \end{cases}$$

We have

$$F_t(\omega) = \sum_{t=1}^{T} f_s(\omega) = \begin{cases} \frac{t+1}{2}\omega^2 + \frac{t-1}{2}(1-\omega)^2 & \text{if } t \text{ is odd,} \\ \frac{t}{2}\left(\omega^2 + (1-\omega)^2\right) & \text{if } t \text{ is even,} \end{cases}$$

Therefore,

$$\omega_t = \operatorname*{argmin}_{\omega \in [0,1]} F_{t-1}(\omega) = \begin{cases} \frac{1}{2} & \text{if } t \text{ is odd,} \\ \frac{1}{2} - \frac{1}{2t} & \text{if } t \text{ is even,} \end{cases}$$

In this example, there is not much fluctuation.

Example 2: FTL with quadratic losses (cont'd)

Example 2: $f_t(\omega) = \begin{cases} \omega^2 & \text{if } t \text{ is odd,} \\ (1-\omega)^2 & \text{if } t \text{ is even.} \end{cases}$ $\quad \omega_t = \underset{\omega \in [0,1]}{\operatorname{argmin}} F_{t-1}(\omega) = \begin{cases} \frac{1}{2} & \text{if } t \text{ is odd,} \\ \frac{1}{2} - \frac{1}{2t} & \text{if } t \text{ is even,} \end{cases}$

We have the following bound on the regret,

$$
\begin{aligned}
R_T &\leq \sum_{t=1}^{T} \left(f_t(\omega_t) - f_t(\omega_{t+1})\right) \\
&= \sum_{\text{odd } t} \left[\left(\frac{1}{2}\right)^2 - \left(\frac{1}{2} - \frac{1}{2(t+1)}\right)^2 \right] + \sum_{\text{even } t} \left[\left(\frac{1}{2} + \frac{1}{2t}\right)^2 - \left(\frac{1}{2}\right)^2 \right] \\
&= \sum_{\text{odd } t} \left[\frac{1}{2(t+1)} - \frac{1}{4(t+1)^2} \right] + \sum_{\text{even } t} \left[\frac{1}{2t} + \frac{1}{4t^2} \right] \\
&\leq \sum_{t=1}^{T} \left(\mathcal{O}\left(\frac{1}{t}\right) + \mathcal{O}\left(\frac{1}{t^2}\right) \right) \in \mathcal{O}(\log(T)).
\end{aligned}
$$

Example 2: FTL with quadratic losses (cont'd)

$$\text{Example 1:} \qquad \Omega = [0,1], \qquad f_t(\omega) = \begin{cases} \frac{1}{2}\omega & \text{if } t = 1, \\ \omega & \text{if } t \text{ is odd, } t > 1, \\ 1 - \omega & \text{if } t \text{ is even.} \end{cases}$$

$$\text{Example 2:} \qquad \Omega = [0,1], \qquad f_t(\omega) = \begin{cases} \omega^2 & \text{if } t \text{ is odd,} \\ (1 - \omega)^2 & \text{if } t \text{ is even.} \end{cases}$$

Like in Example 1, *the best action for a given round* i.e $\operatorname{argmin}_\omega f_t(\omega)$ fluctuates from 0 to 1. However, the regret is not large since $\operatorname{argmin}_\omega F_t(\omega)$ does not fluctuate.

Question: Let us consider linear losses again, but with FTRL. What type of regularizer should we use?

Example 3: FTRL with linear losses and a quadratic regularizer

$$\text{FTRL:} \qquad \omega_t \in \operatorname*{argmin}_{\omega \in \Omega} \left(\sum_{s=1}^{t-1} f_s(\omega) + \Lambda(\omega) \right).$$

Let us revisit Example 1,

$$\Omega = [0,1], \qquad f_t(\omega) = \begin{cases} \frac{1}{2}\omega & \text{if } t = 1, \\ \omega & \text{if } t \text{ is odd, } t > 1, \\ 1 - \omega & \text{if } t \text{ is even.} \end{cases},$$

but use FTRL with a quadratic regularizer $\Lambda(\omega) = \frac{1}{\eta}(\omega - 1/2)^2$. Will specify η later.

We have

$$F_t(\omega) = \sum_{t=1}^{T} f_s(\omega) = \begin{cases} \frac{1}{\eta}(\omega - 1/2)^2 + \frac{1}{2}\omega + \frac{t-1}{2} & \text{if } t \text{ is odd,} \\ \frac{1}{\eta}(\omega - 1/2)^2 - \frac{1}{2}\omega + \frac{t}{2} & \text{if } t \text{ is even,} \end{cases}$$

Example 3: FTRL with linear losses and a quadratic regularizer (cont'd)

$$F_t(\omega) = \sum_{t=1}^{t} f_s(\omega) = \begin{cases} \frac{1}{\eta}(\omega - 1/2)^2 + \frac{1}{2}\omega + \frac{t-1}{2} & \text{if } t \text{ is odd,} \\ \frac{1}{\eta}(\omega - 1/2)^2 - \frac{1}{2}\omega + \frac{t}{2} & \text{if } t \text{ is even,} \end{cases}$$

Therefore, for odd t, we have,

$$\omega_t = \operatorname*{argmin}_{\omega \in [0,1]} F_{t-1}(\omega) = \operatorname*{argmin}_{\omega \in [0,1]} \left(\frac{1}{\eta}(\omega - 1/2)^2 - \frac{1}{2}\omega + \frac{t-1}{2} \right) = \frac{1}{2} + \frac{\eta}{4}.$$

A similar calculation for even t reveals,

$$\omega_t = \operatorname*{argmin}_{\omega \in [0,1]} F_{t-1}(\omega) = \begin{cases} \frac{1}{2} + \frac{\eta}{4} & \text{if } t \text{ is odd,} \\ \frac{1}{2} - \frac{\eta}{4} & \text{if } t \text{ is even} \end{cases}$$

Example 3: FTRL with linear losses and a quadratic regularizer (cont'd)

$$f_t(\omega) = \begin{cases} \omega & \text{if } t \text{ is odd, } t > 1, \\ 1 - \omega & \text{if } t \text{ is even.} \end{cases}, \qquad \omega_t = \begin{cases} \frac{1}{2} + \frac{\eta}{4} & \text{if } t \text{ is odd,} \\ \frac{1}{2} - \frac{\eta}{4} & \text{if } t \text{ is even} \end{cases}$$

We have the following bound on the regret,

$$R_T \leq \sum_{t=1}^{T} (f_t(\omega_t) - f_t(\omega_{t+1})) \; + \; \Lambda(u) - \min_{\omega \in \Omega} \Lambda(\omega). \qquad \texttt{FTRL lemma.}$$

$$\leq \sum_{\text{odd } t} \left(\frac{1}{2} + \frac{\eta}{4} - \left(\frac{1}{2} - \frac{\eta}{4} \right) \right) + \sum_{\text{even } t} \left(\left(1 - \frac{1}{2} + \frac{\eta}{4} \right) - \left(1 - \frac{1}{2} - \frac{\eta}{4} \right) \right)$$

$$+ \frac{1}{\eta} \left(\frac{1}{4} - 0 \right)$$

$$= \sum_{t} \frac{\eta}{2} + \frac{1}{4\eta} = \frac{\eta T}{2} + \frac{1}{4\eta}.$$

If we choose $\eta = 1/\sqrt{T}$, then $R_T \in \mathcal{O}(\sqrt{T})$.

Take-aways from the three examples

- Linear losses have bad behavior in FTL due to the instability of ω_t selected.
- We should add a “nice” regularizer to stabilize the fluctuations.

Ch 6.2: FTRL with convex losses and strongly convex regularizers

$$\textbf{FTRL: } \omega_t \in \operatorname*{argmin}_{\omega\in\Omega} \left(\sum_{s=1}^{t-1} f_s(\omega) + \Lambda(\omega) \right).$$

Theorem. Suppose f_t is convex for all t and let $\Lambda(\omega) = \frac{1}{\eta}\lambda(\omega)$ where λ is 1–strongly convex with respect to some norm $\|\cdot\|$. Let $\|\cdot\|_\star$ be the dual norm and let $\omega_\star \in \operatorname{argmin}_{\omega\in\Omega} \sum_{t=1}^{T} f_t(\omega)$. Then, the following bound holds for any sequence $(g_1, \ldots, g_T)$ where $g_t \in \partial f_t(\omega_t)$ for all t,

$$R_T(\pi, f) \leq \frac{1}{\eta}\left(\lambda(\omega_\star) - \min_{\omega\in\Omega} \lambda(\omega) \right) + \eta \sum_{t=1}^{T} \|g_t\|_\star^2.$$

Proof: FTRL with a strongly convex regularizer

Recall the following bound for FTRL. For all $u \in \Omega$,

$$\sum_{t=1}^{T} f_t(\omega_t) - \sum_{t=1}^{T} f_t(u) \le \Lambda(u) - \min_{\omega \in \Omega} \Lambda(\omega) \; + \; \sum_{t=1}^{T} \left(f_t(\omega_t) - f_t(\omega_{t+1})\right).$$

We will apply this theorem with $u \leftarrow \omega_\star \in \operatorname{argmin}_{\omega \in \Omega} \sum_{t=1}^{T} f_t(\omega)$. We have,

$$\begin{aligned} R_T(\pi, f) &\triangleq \sum_{t=1}^{T} f_t(\omega_t) - \sum_{t=1}^{T} f_t(\omega_\star) \\ &\le \frac{1}{\eta}\left(\lambda(\omega_\star) - \min_{\omega \in \Omega} \lambda(\omega)\right) + \sum_{t=1}^{T} \left(f_t(\omega_t) - f_t(\omega_{t+1})\right) \end{aligned}$$

It is sufficient to show $(f_t(\omega_t) - f_t(\omega_{t+1})) \le \eta \|g_t\|_\star^2$.

Proof: FTRL with a strongly convex regularizer (cont'd)

$$R_T(\pi, f) \leq \frac{1}{\eta}\left(\lambda(\omega_\star) - \min_{\omega\in\Omega}\lambda(\omega)\right) + \sum_{t=1}^{T}(f_t(\omega_t) - f_t(\omega_{t+1}))$$

By convexity, as $g_t \in \partial f_t(\omega_t)$, we have $f_t(\omega_{t+1}) \geq f_t(\omega_t) + g_t^\top(\omega_{t+1} - \omega_t)$.

Hence, by Hölder's inequality, we have

$$f_t(\omega_t) - f_t(\omega_{t+1}) \leq g_t^\top(\omega_t - \omega_{t+1}) \leq \|\omega_{t+1} - \omega_t\|\|g_t\|_\star.$$

Now, denote $F_t(\omega) = \sum_{s=1}^{t} f_t(\omega) + \frac{1}{\eta}\lambda(\omega)$. We have that F_t is $\frac{1}{\eta}$-strongly convex, as λ is 1–strongly convex and f_t's are convex.

Proof: FTRL with a strongly convex regularizer (cont'd)

(i) F_t is $\frac{1}{\eta}$–strongly convex.
(ii) If $\omega_\star = \mathrm{argmin}_{\omega\in\Omega} f(\omega)$, where f is α–strongly convex, then $f(\omega) \geq f(\omega_\star) + \frac{\alpha}{2}\|\omega - \omega_\star\|_2^2$.

Recall, in FTRL, we have $\omega_t = \mathrm{argmin}_\omega F_{t-1}(\omega)$.
Therefore, ω_{t+1} minimizes F_t and ω_t minimizes F_{t-1}. Using *(i)*, *(ii)* we have,

$$F_{t-1}(\omega_{t+1}) - F_{t-1}(\omega_t) \geq \frac{1}{2\eta}\|\omega_t - \omega_{t+1}\|^2,$$
$$F_t(\omega_t) - F_t(\omega_{t+1}) \geq \frac{1}{2\eta}\|\omega_t - \omega_{t+1}\|^2.$$

Summing both sides we have, $f_t(\omega_t) - f_t(\omega_{t+1}) \geq \frac{1}{\eta}\|\omega_t - \omega_{t+1}\|^2$.

Proof: FTRL with a strongly convex regularizer (cont'd)

(i) $f_t(\omega_t) - f_t(\omega_{t+1}) \leq \|\omega_{t+1} - \omega_t\| \|g_t\|_\star$.
(ii) $f_t(\omega_t) - f_t(\omega_{t+1}) \geq \frac{1}{\eta} \|\omega_t - \omega_{t+1}\|^2$.

Therefore,

$$\begin{aligned} \textit{(i), (ii)} \implies & \|\omega_t - \omega_{t+1}\|^2 \leq \eta \left(f_t(\omega_t) - f_t(\omega_{t+1})\right) \leq \eta \|\omega_{t+1} - \omega_t\| \|g_t\|_\star \\ \implies & \|\omega_t - \omega_{t+1}\| \leq \eta \|g_t\|_\star \qquad \textit{(iii)} \end{aligned}$$

$$\textit{(i), (iii)} \implies f_t(\omega_t) - f_t(\omega_{t+1}) \leq \eta \|g_t\|_\star^2.$$

$\square$

A useful corollary

FTRL bound with convex losses and strongly-convex regularizers,

$$R_T(\pi, f) = \sum_{t=1}^{T} f_t(\omega_t) - \sum_{t=1}^{T} f_t(\omega_\star) \leq \frac{1}{\eta}\Big(\lambda(\omega_\star) - \min_{\omega} \lambda(\omega)\Big) + \eta \sum_{t=1}^{T} \|g_t\|_\star^2.$$

Corollary. If $(\max_\omega \lambda(\omega) - \min_\omega \lambda(\omega)) \leq B$ and $\|g_t\|_\star \leq G$ for all t, then choosing $\eta = \sqrt{B/(TG^2)}$, we have

$$R_T \in \mathcal{O}(G\sqrt{BT}).$$

Here, $\|g_t\|_\star \leq G$ simply means that f_t is G–Lipschitz in norm $\|\cdot\|_\star$.

Example 1: Linear losses

Let $\Omega = \{\omega; \|\omega\|_2 \leq 1\}$ and $f_t(\omega_t) = \omega^\top \ell_t$ where $\|\ell_t\|_2 \leq 1$.

We will apply FTRL with $\lambda(\omega) = \frac{1}{2}\|\omega\|_2^2$ which is 1–strongly convex in $\|\cdot\|_2$.

FTRL update on round t,

$$\begin{aligned}
\omega_t &= \operatorname*{argmin}_{\omega\in\Omega} \left(\sum_{s=1}^{t-1} \omega^\top \ell_s + \frac{1}{2\eta}\|\omega\|_2^2 \right) = \operatorname*{argmin}_{\omega\in\Omega} \left(\omega^\top \left[\sum_{s=1}^{t-1} \ell_s \right] + \frac{1}{2\eta}\|\omega\|_2^2 \right) \\
&= \operatorname*{argmin}_{\omega\in\Omega} \left(\|\omega\|_2^2 + 2\eta\omega^\top \left[\sum_{s=1}^{t-1} \ell_s \right] + \eta^2 \left[\sum_{s=1}^{t-1} \ell_s \right]^2 \right) \qquad \texttt{completing the square.} \\
&= \operatorname*{argmin}_{\omega\in\Omega} \left\| \omega + \eta \sum_{s=1}^{t-1} \ell_t \right\|_2
\end{aligned}$$

That is, choose ω_t to be the projection of $-\eta \sum_{t=1}^{T} \ell_t$ onto Ω in the L_2-norm.

Example 1: Linear losses (cont'd)

- FTRL regret bound: $R_T(\pi, f) \leq \frac{1}{\eta}\left(\lambda(\omega_\star) - \lambda(\omega_t)\right) + \sum_{t=1}^{T} \|g_t\|_\star$.
- FTRL with linear losses: $\omega_t = \mathrm{argmin}_{\omega\in\Omega} \left\|\omega + \eta \sum_{s=1}^{t-1} \ell_t\right\|_2$.

This can be implemented via,

$$u_t \leftarrow u_{t-1} + \eta \ell_{t-1}, \qquad \omega_t \leftarrow \underset{\omega\in\Omega}{\mathrm{argmin}} \|\omega - u_t\|_2^2. \qquad (\star)$$

Therefore, the regret satisfies,

$$R_T \leq \frac{1}{\eta}\left(\frac{1}{2}\|\omega_\star\|_2^2 - \min_{\omega\in\Omega}\|\omega\|_2^2\right) + \eta \sum_{t=1}^{T} \|\ell_t\|^2 \leq \frac{1}{\eta}\left(\frac{1}{2}\cdot 1 - 0\right) + \eta T$$
$$\in \mathcal{O}\left(\sqrt{T}\right) \qquad \texttt{if we choose } \eta = 1/\sqrt{T}.$$

Remark. The update $(\star)$ only takes $O(1)$ computation per round.

Example 2: Online gradient descent

Let f_t be differentiable on all rounds, and let Ω be a compact convex set.

We can apply FTRL with some $\lambda(\omega)$ which is 1–strongly convex in some norm $\|\cdot\|$.

FTRL update rule

$$\omega_t = \operatorname*{argmin}_{\omega \in \Omega} \left(\sum_{s=1}^{t-1} f_s(\omega) + \frac{1}{\eta}\lambda(\omega) \right)$$

However, the complexity of this update is $O(t)$.
Ideally, we would like it to be $O(1)$.

Example 2: Online gradient descent (cont'd)

Let us take a slightly different perspective,

$$
\begin{aligned}
R_T\left(\pi, \{f_t\}_{t=1}^{t}\right) &= \sum_{t=1}^{T} f_t(\omega_t) - \min_{\omega\in\Omega} \sum_{t=1}^{T} f_t(\omega) = \max_{\omega\in\Omega}\left(\sum_{t=1}^{T} f_t(\omega_t) - \sum_{t=1}^{T} f_t(\omega)\right) \\
&\le \max_{\omega\in\Omega} \sum_{t=1}^{T} \nabla f_t(\omega_t)^\top (\omega_t - \omega) \qquad f_t(\omega) \ge f_t(\omega_t) + \nabla f_t(\omega_t)^\top (\omega - \omega_t) \texttt{ by convexity} \\
&= \sum_{t=1}^{T} \nabla f_t(\omega_t)^\top \omega_t - \min_{\omega\in\Omega} \sum_{t=1}^{T} \nabla f_t(\omega_t)^\top \omega. \\
&= R_T\left(\pi, \{\nabla f_t(\omega_t)^\top (\cdot)\}_{t=1}^{T}\right) \qquad \texttt{Linear losses with } \ell_t = \nabla f_t(\omega_t) \texttt{ !}
\end{aligned}
$$

Example 2: Online gradient descent (cont'd)

We can now apply FTRL on the linear losses $\nabla f_t(\omega_t)^\top(\cdot)$ with $\lambda(\omega) = \frac{1}{2}\|\omega\|_2^2$ which is 1–strongly convex in $\|\cdot\|_2$.

$$\begin{aligned} \omega_t &= \operatorname*{argmin}_{\omega\in\Omega} \left(\omega^\top \left[\sum_{s=1}^{t-1} \nabla f_s(\omega_s) \right] + \frac{1}{2\eta}\|\omega\|_2^2 \right) \\ &= \operatorname*{argmin}_{\omega\in\Omega} \left\| \omega + \eta \sum_{s=1}^{t-1} \nabla f_s(\omega_s) \right\|_2 \qquad \texttt{completing the square.} \end{aligned}$$

This is the projection of $-\eta \sum_{t=1}^{T} \ell_t$ onto Ω in the L_2-norm. Can be implemented as follows in $\mathcal{O}(1)$ time,

$$u_t \leftarrow u_{t-1} - \eta \nabla f_{t-1}(\omega_{t-1}), \qquad \omega_t \leftarrow \operatorname*{argmin}_{\omega\in\Omega} \|\omega - u_t\|_2^2.$$

Regret: If $\|\nabla f_t(\cdot)\|_2 \le G$ and $\max_\omega \lambda(\omega) - \min_\omega \lambda(\omega) \le B$, then

$$R_T\left(\pi, \{f_t\}_{t=1}^{t}\right) \le R_T\left(\pi, \{\nabla f_t(\omega_t)^\top(\cdot)\}_{t=1}^{T}\right) \le \frac{B}{\eta} + \eta T G^2 \in \mathcal{O}(G\sqrt{BT}).$$

Example 2: Connections to projected/stochastic gradient descent:

$$u_t \leftarrow u_{t-1} - \eta \nabla f_{t-1}(\omega_{t-1}), \qquad \omega_t \leftarrow \operatorname*{argmin}_{\omega \in \Omega} \|\omega - u_t\|_2^2. \qquad (*)$$

1. Suppose we are interested in finding the minimum $\omega_\star$ of a fixed function f, i.e $\omega_\star = \operatorname{argmin}_{\omega \in \Omega} f(\omega)$. If we take $f_t = f$, we obtain the standard *projected gradient descent* update,

$$u_t \leftarrow u_{t-1} - \eta \nabla f(\omega_{t-1}), \qquad \omega_t \leftarrow \operatorname*{argmin}_{\omega \in \Omega} \|\omega - u_t\|_2^2.$$

   We also obtain the following guarantee for projected gradient descent,

$$\min_{\omega_t} f(\omega_t) - f(\omega_\star) \leq \frac{1}{T} \sum_{t=1}^{T} (f(\omega_t) - f(\omega_\star)) = \frac{R_T}{T} \in \mathcal{O}\left(\frac{G\sqrt{B}}{\sqrt{T}}\right).$$

2. In ML, $(*)$ is similar to the (projected) SGD update, where f_t is the loss on the current data point.

Example 3: The experts problem revisited

Here, $\Omega = \Delta([K])$, losses $f_t(p) = \ell_t^\top p$, where $\ell_t \in [0,1]^K$.

Let us try $\lambda(\omega) = \frac{1}{2}\|\omega\|_2^2$. We have the following bound on the regret, (You can try the update rule at home)

$$R_T \le \frac{1}{\eta} \underbrace{\left(\max_{\omega} \lambda(\omega) - \min_{\omega} \lambda(\omega)\right)}_{=\frac{1}{2}\cdot 1 - \frac{1}{2}\cdot\frac{1}{K} \le \frac{1}{2}} + \sum_{t=1}^{T} \|g_t\|_\star \le \frac{1}{2\eta} + \eta K T \qquad \texttt{As } \|\ell_t\|_2 \le \sqrt{K}$$

$$\in \mathcal{O}\left(\sqrt{KT}\right) \qquad \texttt{Choosing } \eta = 1/\sqrt{KT}$$

But, we saw that Hedge achieved $\mathcal{O}(\sqrt{T\log(K)})$. This is because $(\|\cdot\|_2, \|\cdot\|_2)$ does not capture the geometry of the problem.

Instead, let us try $(\|\cdot\|_1, \|\cdot\|_\infty)$.
- We know that $\|\ell_t\|_\infty \le 1$ (does not scale with K).
- And $\Delta([K])$ is a subset of the L_1-ball of radius 1.

Example 3: The experts problem revisited (cont'd)

Experts problem: $\Omega = \Delta([K])$, $f_t(p) = \ell_t^\top p$, where $\ell_t \in [0,1]^K$.

Let us try the negative entropy for regularization,

$$\lambda(p) = -\mathrm{H}(p) = \sum_{i=1}^{K} p(i) \log(p(i))$$

Recall, from Chapter 0, that $\lambda(p)$ is 1–strongly convex in $\|\cdot\|_1$.

We have the following bound on the regret,

$$\begin{aligned} R_T &\leq \frac{1}{\eta}\left(\max_\omega \lambda(\omega) - \min_\omega \lambda(\omega)\right) + \sum_{t=1}^{T} \|g_t\|_\star \\ &= \frac{1}{\eta}\left(\underbrace{\max_\omega \mathrm{H}(\omega)}_{\leq \log(K)} - \min_\omega \mathrm{H}(\omega)\right) + \sum_{t=1}^{T} \underbrace{\|\ell_t\|_\infty}_{\leq 1} \leq \frac{\log(K)}{\eta} + \eta T \\ &\in \mathcal{O}\left(\sqrt{T \log(K)}\right) \qquad \texttt{Choosing } \eta = \sqrt{\log(K)/T} \end{aligned}$$

Example 3: The experts problem revisited (cont'd)

Let us now derive the update rule,

$$p_t = \underset{p \in \Delta([K])}{\operatorname{argmin}} \left(\sum_{s=1}^{t-1} \ell_s^\top p + \frac{1}{\eta} \sum_{i=1}^{K} p(i) \log(p(i)). \right)$$

We can write this as the following optimization problem,

$$\text{minimize}_p \; \sum_{s=1}^{t-1} \ell_s^\top p + \frac{1}{\eta} \sum_{i=1}^{K} p(i) \log(p(i)). \qquad \text{sub. to } \; \mathbf{1}^\top p = 1, \; p \geq 0$$

We will write out the Lagrangian for the $\mathbf{1}^\top p = 1$ constraint and then verify that the solution is non-negative.

$$\mathcal{L} = \sum_{s=1}^{t-1} \ell_s^\top p + \frac{1}{\eta} \sum_{i=1}^{K} p(i) \log(p(i)) + \mu(p^\top \mathbf{1} - 1).$$

Example 3: The experts problem revisited (cont'd)

$$\mathcal{L} = \sum_{s=1}^{t-1} \ell_s^\top p + \frac{1}{\eta} \sum_{i=1}^{K} p(i) \log(p(i)) + \mu(p^\top \mathbf{1} - 1).$$

Taking derivatives and setting to 0,

$$\frac{\partial \mathcal{L}}{\partial p(i)} = \sum_{s=1}^{t-1} \ell_s(i) + \frac{1}{\eta}(1 + \log(p(i))) + \mu \stackrel{\Delta}{=} 0$$

$$\implies p_t(i) = e^{-\eta\mu} \exp\left(-\eta \sum_{s=1}^{t-1} \ell_s(i)\right)$$

As $\sum_i p(i) = 1$, we obtain,

$$p_t(i) = \frac{\exp\left(-\eta \sum_{s=1}^{t-1} \ell_s(i)\right)}{\sum_{j=1}^{K} \exp\left(-\eta \sum_{s=1}^{t-1} \ell_s(j)\right)}$$

This is precisely the Hedge algorithm!

Ch 6.3: Follow the Perturbed Leader

Follow the perturbed leader (FTPL)

- Given: time horizon T, a distribution D.
- Sample $f_0 \sim D$
- for $t = 1, \ldots, T$
  - Choose $\omega_t = \operatorname*{argmin}_{\omega \in \Omega} \sum_{s=0}^{t-1} f_s(\omega)$

We will assume an oblivious adversary for simplicity. With an adaptive adversary, you just need to sample $f_0 \sim D$ on every round.

FTPL vs FTRL

FTRL

- Given: horizon T, regularizer Λ.
- for $t = 1, \ldots, T$

$$\omega_t \leftarrow \operatorname*{argmin}_{\omega \in \Omega} \sum_{s=1}^{t-1} f_s(\omega) + \Lambda(\omega)$$

FTPL

- Given: horizon T, distribution D.
- Sample $f_0 \sim D$
- for $t = 1, \ldots, T$

$$\omega_t \leftarrow \operatorname*{argmin}_{\omega \in \Omega} \sum_{s=0}^{t-1} f_s(\omega)$$

- FTRL is deterministic, while FTPL is stochastic.
- In FTPL, we replace the regularizer Λ with an f_0 sampled from a distribution D.

Why FTPL over FTRL?

- In some usecases, if f_0 "looks like" f_t's, the optimization may be simpler.

On the flipside,

- Choice of D is not always straightforward.
- Analysis techniques differ from use case to use case.

Example: online linear optimization in a convex polytope

Let $f_t(\omega) = \ell_t^\top \omega$ be linear and $\Omega = \{\omega : A\omega \leq b\}$ where $A \in \mathbb{R}^{N\times d}, b \in \mathbb{R}^n$.

Say we run FTRL with a quadratic regularizer $\Lambda(\omega) = \frac{1}{2\eta}\|\omega\|_2^2$. Then, from a previous example,

$$\omega_t = \operatorname*{argmin}_{\omega\in\Omega} \left(\sum_{s=1}^{t-1} \omega^\top \ell_s + \frac{1}{2\eta}\|\omega\|_2^2\right) = \operatorname*{argmin}_{\omega\in\Omega} \left\|\omega + \eta \sum_{s=1}^{t-1} \ell_t\right\|_2$$

This can be implemented via,

$$u_t \leftarrow u_{t-1} + \eta\ell_{t-1}, \qquad \omega_t \leftarrow \operatorname*{argmin}_{\omega\in\Omega} \|\omega - u_t\|_2^2.$$

This projection operation, while still a convex optimization problem, can still be computationally expensive in a polytope.

Example: online linear optimization in a convex polytope (cont'd)

Let $f_t(\omega) = \ell_t^\top \omega$ be linear and $\Omega = \{\omega : A\omega \leq b\}$ where $A \in \mathbb{R}^{N \times d}$, $b \in \mathbb{R}^n$.

Let us try FTPL. We will sample $\ell_0 \sim D$ (for some appropriately chosen D), and then on each round

$$\omega_t \leftarrow \operatorname*{argmin}_{\omega \in \Omega} \sum_{s=0}^{t-1} \ell_s^\top \omega$$

This can be implemented via the following linear program, which is computationally cheaper than a projection

$$\min_{\omega} \omega^\top \left[\sum_{s=0}^{t-1} \ell_s \right] \qquad \text{subject to } A\omega \leq b.$$

A preliminary bound for FTPL

Lemma (FTPL). Let $f = (f_1, \ldots, f_T)$ be a sequence of losses. Then FTPL satisfies,

$$R_T(\pi^{\mathrm{FTPL}}, f) \triangleq \mathbb{E}\left[\sum_{t=1}^{T} f_t(\omega_t)\right] - \min_{\omega\in\Omega}\sum_{t=1}^{T} f_t(\omega)$$
$$\leq \sum_{t=1}^{T}\mathbb{E}\left[f_t(\omega_t) - f_t(\omega_{t+1})\right] + \mathbb{E}\left[\max_{\omega\in\Omega} f_0(\omega) - \min_{\omega\in\Omega} f_0(\omega)\right]$$

where the expectation is with respect to $f_0 \sim D$.

N.B. Does not assume convexity of Ω, f_t or Λ.
Recall we are assuming an oblivious adversary.

Proof of FTPL lemma

Recall the following bound for FTRL. For all $u \in \Omega$,

$$\sum_{t=1}^{T} f_t(\omega_t) - \sum_{t=1}^{T} f_t(u) \leq \Lambda(u) - \min_{\omega \in \Omega} \Lambda(\omega) + \sum_{t=1}^{T} \left(f_t(\omega_t) - f_t(\omega_{t+1})\right).$$

For a given f_0, let us apply the above lemma with $\Lambda = f_0$ and $u = \omega_\star = \mathrm{argmin}_{\omega \in \Omega} \sum_{t=1}^{T} f_t(\omega)$.

$$\sum_{t=1}^{T} f_t(\omega_t) - \sum_{t=1}^{T} f_t(\omega_\star) \leq f_0(\omega_\star) - \min_{\omega \in \Omega} f_0(\omega) + \sum_{t=1}^{T} \left(f_t(\omega_t) - f_t(\omega_{t+1})\right).$$

The claim follows by noting that $f_0(\omega_\star) \leq \max_\omega f_0(\omega)$ and then taking expectation.

$\square$

The experts problem revisited

We will apply FTPL with $f_0(\cdot) = \ell_0^\top(\cdot)$, and where $\ell_0(a) \sim D(\eta)$ for all $a \in [K]$.

On round t we will choose $p_t \in \operatorname{argmin}_{p\in\Delta([K])} \sum_{s=0}^{t-1} p^\top \ell_t$, which is equivalent to choosing $A_t \in \operatorname{argmin}_{a\in[K]} \sum_{s=0}^{t-1} \ell_t(a)$.

This gives rise to the following algorithm

FTPL for the experts problem

- Given: time horizon T, parameter η
- Sample $\ell_0(a) \sim D(\eta)$ for $a \in [K]$. $\quad \ell_0 \sim D$
- for $t = 1, \ldots, T$,
  - $A_t \leftarrow \operatorname{argmin}_{a\in[K]} \sum_{s=0}^{t-1} \ell_s(a)$. $\quad \omega_t = \operatorname{argmin}_{\omega\in\Omega} \sum_{s=0}^{t-1} f_s(\omega)$

FTPL for experts: negative geometric perturbation

For all a, we will sample $\ell_0(a) = -Z(a)$ where $Z(a) \sim \mathrm{Geom}(\eta)$.

The $\mathrm{Geom}(\eta)$ distribution: the distribution of the number Z of $\mathrm{Bern}(\eta)$ coin flips to get the first 1.

$$\textbf{pmf:} \text{ for } k \in \{1, 2, \dots\} \qquad p(k) = \mathbb{P}(Z = k) = (1-\eta)^{k-1}\eta.$$

Some useful properties:

1. $\mathbb{P}(Z \geq k+1 | Z \geq k) = \frac{\mathbb{P}(Z \geq k+1,\, Z \geq k)}{\mathbb{P}(Z \geq k)} = \frac{\mathbb{P}(Z \geq k+1)}{\mathbb{P}(Z \geq k)} = \frac{(1-\eta)^k}{(1-\eta)^{k-1}} = 1 - \eta.$
2. Let $Z(a) \sim \mathrm{Geom}(\eta)$ for $a \in [K]$. Then (try at home)

$$\mathbb{E}[\|Z\|_\infty] = \mathbb{E}\left[\max_{a \in [K]} Z(a)\right] \leq 1 + \frac{H_K}{\eta}, \qquad \text{where, } H_K = 1 + \frac{1}{2} + \cdots + \frac{1}{K}.$$

FTPL for experts: negative geometric perturbation (cont'd)

Lemma (FTPL). Let $f = (f_1, \ldots, f_T)$ be a sequence of losses. Then FTPL satisfies,

$$R_T(\pi^{\mathrm{FTPL}}, f) \leq \sum_{t=1}^{T} \mathbb{E}\left[f_t(\omega_t) - f_t(\omega_{t+1})\right] + \mathbb{E}\left[\max_{\Omega} f_0(\omega) - \min_{\Omega} f_0(\omega)\right]$$

Note that for the experts problem $\omega \in \Delta([K])$ is such that $\omega_t(a) = \mathbb{1}(A_t = a)$. Moreover,

$$\max_{\omega \in \Delta([K])} f_0(\omega) = \max_{\omega \in \Delta([K])} \ell_0^\top \omega = \max_{a \in [K]} \ell_0(a), \qquad \min_{\omega \in \Delta([K])} f_0(\omega) = \min_{a \in [K]} \ell_0(a),$$

we have,

$$R_T(\pi^{\mathrm{FTPL}}, f) \leq \sum_{t=1}^{T} \mathbb{E}[\ell_0(A_t) - \ell_0(A_{t+1})] + \mathbb{E}\left[\max_{a \in [K]} \ell_0(a) - \min_{a \in [K]} \ell_0(a)\right]$$

FTPL for experts: negative geometric perturbation (cont'd)

$$R_T(\pi^{\mathrm{FTPL}}, f) \leq \sum_{t=1}^{T} \mathbb{E}[\ell_0(A_t) - \ell_0(A_{t+1})] + \mathbb{E}\left[\max_{a\in[K]} \ell_0(a) - \min_{a\in[K]} \ell_0(a)\right]$$

Let us first bound the second term. Recall that $\ell_0(a) = -Z(a)$ where $Z(a) \sim \mathrm{Geom}(\eta)$. As $Z(a) \geq 1$, we have

$$\max_a \ell_0(a) = \max_a -Z(a) \leq -1.$$

$$-\min_a \ell_0(a) = \max_a Z(a) = \|Z\|_\infty$$

Therefore,

$$\mathbb{E}\left[\max_{a\in[K]} \ell_0(a) - \min_{a\in[K]} \ell_0(a)\right] \leq -1 + \mathbb{E}[\|Z\|_\infty] \leq -1 + 1 + \frac{H_K}{\eta} = \frac{H_K}{\eta}.$$

FTPL for experts: negative geometric perturbation (cont'd)

$$R_T(\pi^{\mathrm{FTPL}}, f) \le \sum_{t=1}^{T} \mathbb{E}[\ell_0(A_t) - \ell_0(A_{t+1})] + \underbrace{\mathbb{E}\left[\max_{a\in[K]} \ell_0(a) - \min_{a\in[K]} \ell_0(a)\right]}_{\le H_K/\eta}$$

Let us now bound the first term. Shortly, we will prove the following claim.

Claim. $\mathbb{P}(A_{t+1} = a | A_t = a) \ge 1 - \eta$ for all $a \in [K]$, where $\mathbb{P}$ is with respect to ℓ_0.

Then, we can write

$$\mathbb{E}[\ell_t(A_t) - \ell_t(A_{t+1})] = \mathbb{E}[\underbrace{\ell_0(A_t) - \ell_0(A_{t+1})}_{=0} | A_t = A_{t+1}]\mathbb{P}(A_t = A_{t+1}) +$$

$$\mathbb{E}[\underbrace{\ell_0(A_t) - \ell_0(A_{t+1})}_{\le 1 \text{ as } \ell_t \in [0,1]^K} | A_t \ne A_{t+1}]\mathbb{P}(A_t \ne A_{t+1})$$

$$\le \mathbb{P}(A_t \ne A_{t+1}) = \sum_{a=1}^{K} \underbrace{\mathbb{P}(A_{t+1} \ne a | A_t = a)}_{\le \eta} \mathbb{P}(A_t = a) \le \eta.$$

FTPL for experts: negative geometric perturbation (cont'd)

$$R_T(\pi^{\mathrm{FTPL}}, f) \leq \sum_{t=1}^{T} \underbrace{\mathbb{E}[\ell_0(A_t) - \ell_0(A_{t+1})]}_{\leq \eta} + \underbrace{\mathbb{E}\left[\max_{a \in [K]} \ell_0(a) - \min_{a \in [K]} \ell_0(a)\right]}_{\leq H_K/\eta}$$

Therefore,

$$\begin{aligned} R_T &\leq \eta T + \frac{H_K}{\eta} \\ &= 2\sqrt{TH_K} \quad \texttt{By choosing } \eta = \sqrt{H_K/T} \\ &\in \mathcal{O}(\sqrt{T\log(K)}). \end{aligned}$$

FTPL for experts: negative geometric perturbation (cont'd)

We will now prove the claim.

Claim. $\mathbb{P}(A_{t+1} = a | A_t = a) \geq 1 - \eta$ for all $a \in [K]$, where $\mathbb{P}$ is with respect to ℓ_0.

Proof. Recall that $\ell_0(a) = -Z(a)$ where $Z(a) \sim \mathrm{Geom}(\eta)$.

Let a be given. We will show that for every realization of $\{Z(j)\}_{j \neq a}$, we have

$$\mathbb{P}(A_{t+1} = a | A_t = a, \{Z(j)\}_{j \neq a}) \geq 1 - \eta \implies \mathbb{P}(A_{t+1} = a | A_t = a) \geq 1 - \eta.$$

Fix the values of $\{Z(j)\}_{j \neq a}$. First observe,

$$A_t = a \iff \sum_{s=0}^{t-1} \ell_s(a) = \sum_{s=1}^{t-1} \ell_s(a) - Z(a) \leq \sum_{s=1}^{t-1} \ell_s(j) - Z(j) \quad \forall j \neq a$$

Now define $L_{t-1}(j) \triangleq \sum_{s=1}^{t-1} \ell_s(j)$ and $J_{t-1} \triangleq \min_{j \neq a} \left(L_{t-1}(j) - Z(j)\right)$. Therefore,

$$A_t = a \iff Z(a) \geq L_{t-1}(a) - J_{t-1}.$$

FTPL for experts: negative geometric perturbation (cont'd)

$$L_{t-1}(j) \triangleq \sum_{s=1}^{t-1} \ell_s(j), \qquad J_{t-1} \triangleq \min_{j \neq a} L_{t-1}(j) - Z(j).$$

Now let us consider $A_{t+1} = j$. We can write,

$$\begin{aligned}
A_{t+1} = a \iff & \sum_{s=0}^{t} \ell_s(a) = \sum_{s=1}^{t} \ell_s(a) - Z(a) \leq \sum_{s=1}^{t} \ell_s(j) - Z(j) \quad \forall j \neq a \\
\iff & Z(a) \geq L_{t-1}(a) + \ell_t(a) - \left(\sum_{s=1}^{t-1} \ell_s(j) + \ell_t(j) - Z(j) \right) \quad \forall j \neq a \\
\iff & Z(a) \geq L_{t-1}(a) - \underbrace{\left(\sum_{s=1}^{t-1} \ell_s(j) - Z(j) \right)}_{\geq J_{t-1}} + \underbrace{\ell_t(a) - \ell_t(j)}_{\leq 1 \text{ as } \ell_t \in [0,1]} \quad \forall j \neq a \\
\Longleftarrow \; & Z(a) \geq L_{t-1}(a) - J_{t-1} + 1
\end{aligned}$$

FTPL for experts: negative geometric perturbation (cont'd)

$A_t = a \iff Z(a) \geq L_{t-1}(a) - J_{t-1}$. $\quad A_{t+1} = a \impliedby Z(a) \geq L_{t-1}(a) - J_{t-1} + 1$

For $Z \sim \mathrm{Geom}(\eta)$, $\mathbb{P}(Z \geq k+1 | Z \geq k) = 1 - \eta$.

We therefore have,

$$
\begin{aligned}
&\mathbb{P}(A_{t+1} = a | A_t = a, \{Z(j)\}_{j \neq a}) \\
&\quad \geq \mathbb{P}\left(Z(a) \geq L_{t-1}(a) - J_{t-1} + 1 \,\middle|\, Z(a) \geq L_{t-1}(a) - J_{t-1}, \{Z(j)\}_{j \neq a}\right) \\
&\quad = 1 - \eta
\end{aligned}
$$

Hence, $\mathbb{P}(A_{t+1} = a | A_t = a) \geq 1 - \eta$.

$\square$

FTPL for experts: Laplace perturbation

FTPL for the experts problem

- Given: time horizon T, parameter η
- Sample $\ell_0(a) \sim D(\eta)$ for $a \in [K]$. $\ell_0 \sim D$
- for $t = 1, \dots, T$,
  - $A_t \leftarrow \mathrm{argmin}_{a \in [K]} \sum_{s=0}^{t-1} \ell_s(a)$. $\omega_t = \mathrm{argmin}_{\omega \in \Omega} \sum_{s=0}^{t-1} f_s(\omega)$

We will now try $D(\eta) = \mathrm{Lap}\,(1/\eta)$.

The $\mathrm{Lap}\,(1/\eta)$ distribution has pdf ψ:

$$\psi(z) = \frac{\eta}{2} e^{-\eta |z|},$$

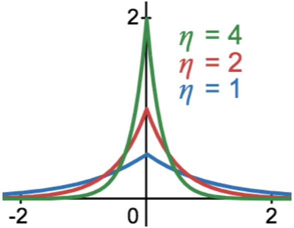

FTPL for experts: Laplace perturbation (cont'd)

Maximum of K i.i.d Laplace RVs. Let $Z = (Z(1), \ldots, Z(K))$ where $Z(i) \sim \mathrm{Lap}(1/\eta)$,

$$\mathbb{E}[\|Z\|_\infty] = \int_0^\infty \mathbb{P}(\|Z\|_\infty \geq t)\mathrm{d}t \quad \texttt{by identity below.}$$

$$= \int_0^a \underbrace{\mathbb{P}(\|Z\|_\infty \geq t)}_{\leq 1}\mathrm{d}t + \int_a^\infty \underbrace{\mathbb{P}(\|Z\|_\infty \geq t)}_{=\mathbb{P}(\exists\, i, |Z(i)| \geq t)}\mathrm{d}t \;\leq a + \sum_{i=1}^K \int_a^\infty \mathbb{P}(|Z(i)| \geq t)\mathrm{d}t$$

We have that,

$$\mathbb{P}(|Z(i)| \geq t) = \int_t^\infty \frac{\eta}{2} e^{-\eta z} + \int_{-\infty}^{-t} \frac{\eta}{2} e^{\eta z} = e^{-\eta t}.$$

Therefore, choosing $a = \frac{1}{\log(K)}$, we have

$$\mathbb{E}[\|Z\|_\infty] \leq a + K \int_a^\infty e^{-\eta t}\mathrm{d}t \leq a + \frac{K}{\eta} e^{-\eta a} \leq \frac{1}{\eta}\left(1 + \log(K)\right).$$

---

[0]For $Z \geq 0$, $\mathbb{E}[Z] = \int_0^\infty z p(z)\mathrm{d}z = \int_0^\infty p(z) \int_0^z \mathrm{d}v\mathrm{d}z = \int_0^\infty \int_v^\infty p(z)\mathrm{d}z\mathrm{d}v = \int_0^\infty \mathbb{P}(Z \geq v)\mathrm{d}v$.

FTPL for experts: Laplace perturbation (cont'd)

Lemma (FTPL). Let $f = (f_1, \dots, f_T)$ be a sequence of losses. Then FTPL satisfies,

$$R_T(\pi^{\mathrm{FTPL}}, f) \le \sum_{t=1}^{T} \mathbb{E}\left[f_t(\omega_t) - f_t(\omega_{t+1})\right] + \mathbb{E}\left[\max_{\Omega} f_0(\omega) - \min_{\Omega} f_0(\omega)\right]$$

Using a similar argument as before (i.e for geometric perturbation), we have

$$R_T(\pi^{\mathrm{FTPL}}, f) \le \sum_{t=1}^{T} \mathbb{E}[\ell_0(A_t) - \ell_0(A_{t+1})] + \mathbb{E}\left[\max_{a\in[K]} \ell_0(a) - \min_{a\in[K]} \ell_0(a)\right]$$

Let us first bound the second term. By symmetry of the Laplace distribution,

$$\mathbb{E}\left[\max_{a\in[K]} \ell_0(a) - \min_{a\in[K]} \ell_0(a)\right] = 2\mathbb{E}\left[\max_{a\in[K]} \ell_0(a)\right] \le \frac{2}{\eta}\left(1 + \log(K)\right).$$

FTPL for experts: Laplace perturbation (cont'd)

$$R_T(\pi^{\mathrm{FTPL}}, f) \leq \sum_{t=1}^{T} \mathbb{E}[\ell_0(A_t) - \ell_0(A_{t+1})] + \underbrace{\mathbb{E}\left[\max_{a\in[K]} \ell_0(a) - \min_{a\in[K]} \ell_0(a)\right]}_{\leq \frac{2}{\eta}(1+\log(K)).}$$

To bound the first term, we will use the following claim.

Claim. $\mathbb{P}(A_t = a) \leq e^{\eta}\mathbb{P}(A_{t+1} = a)$ for all $a \in [K]$, where $\mathbb{P}$ is w.r.t ℓ_0.

We therefore have,

$$\begin{aligned}
\mathbb{E}[\ell_t(A_t) - \ell_t(A_{t+1})] &= \sum_{a=1}^{K} \ell_t(a)\mathbb{P}(A_t = a) - \sum_{a=1}^{K} \ell_t(a)\mathbb{P}(A_{t+1} = a) \\
&= \sum_{a=1}^{K} \ell_t(a)\left(\mathbb{P}(A_t = a) - \mathbb{P}(A_{t+1} = a)\right) \\
&\leq \sum_{a=1}^{K} \underbrace{\ell_t(a)}_{\leq 1} \underbrace{\left(1 - e^{-\eta}\right)}_{\leq \eta} \mathbb{P}(A_t = a) \leq \eta.
\end{aligned}$$

FTPL for experts: Laplace perturbation (cont'd)

$$R_T(\pi^{\mathrm{FTPL}}, f) \leq \sum_{t=1}^{T} \underbrace{\mathbb{E}[\ell_0(A_t) - \ell_0(A_{t+1})]}_{\leq \eta} + \underbrace{\mathbb{E}\left[\max_{a \in [K]} \ell_0(a) - \min_{a \in [K]} \ell_0(a)\right]}_{\leq \frac{2}{\eta}(1+\log(K)).}$$

Therefore,

$$\begin{aligned} R_T &\leq \eta T + \frac{2}{\eta}(1 + \log(K)) \\ &= 3\sqrt{T(1+\log(K))} \quad \texttt{By choosing } \eta = \sqrt{(1+\log(K))/T} \\ &\in \mathcal{O}(\sqrt{T \log(K)}). \end{aligned}$$

FTPL for experts: Laplace perturbation (cont'd)

We will now prove the following claim.

Claim. $\mathbb{P}(A_t = a) \leq e^{\eta}\mathbb{P}(A_{t+1} = a)$ for all $a \in [K]$, where $\mathbb{P}$ is w.r.t ℓ_0.

Let a be given. Let ψ be the pdf of ℓ_0. Therefore,

$$\psi(\ell_0) = \prod_{j=1}^{K} \frac{\eta}{2} e^{-\eta|\ell_0(j)|} = \frac{\eta^K}{2^K} e^{-\eta\|\ell_0\|_1}.$$

We can write,

$$\mathbb{P}(A_t = a) = \int_{\mathbb{R}^K} \mathbb{1}\left(a = \operatorname*{argmin}_{j\in[K]} \sum_{s=0}^{t-1} \ell_s(j)\right) \psi(\ell_0)\mathrm{d}\ell_0.$$

Let $\ell_t^a \in [0,1]^K$ such that $\ell_t^a(j) = \mathbb{1}(j = a)\ell_t(a)$. That is $\ell_t^a = [0, \dots, \ell_t(j), \dots 0]$.
Now, let us use the substitution $\widetilde{\ell}_0 = \ell_0 - \ell_t^a$. We have,

$$\mathbb{P}(A_t = a) = \int_{\mathbb{R}^K} \mathbb{1}\left(a = \operatorname*{argmin}_{j\in[K]} \widetilde{\ell}_0(j) + \ell_t^a(j) + \sum_{s=1}^{t-1} \ell_s(j)\right) \psi(\widetilde{\ell}_0 + \ell_t^a)\mathrm{d}\widetilde{\ell}_0.$$

FTPL for experts: Laplace perturbation (cont'd)

$$\psi(\ell_0) = \frac{\eta^K}{2^K} e^{-\eta\|\ell_0\|_1}. \qquad \mathbb{P}(A_t = a) = \int_{\mathbb{R}^K} \mathbb{1}\left(a = \operatorname*{argmin}_{j\in[K]} \widetilde{\ell}_0(j) + \ell_t^a(j) + \sum_{s=1}^{t-1} \ell_s(j)\right) \psi(\widetilde{\ell}_0 + \ell_t^a)\mathrm{d}\widetilde{\ell}_0.$$

Now we will upper bound $\psi(\widetilde{\ell}_0 + \ell_t^a)$ as follows,

$$\begin{aligned}
\psi(\widetilde{\ell}_0 + \ell_t^a) &= \frac{\eta^K}{2^K} e^{-\eta\|\widetilde{\ell}_0 + \ell_t^a\|_1} \\
&\leq \frac{\eta^K}{2^K} e^{-\eta\|\widetilde{\ell}_0\|_1 + \eta\|\ell_t^a\|_1} \qquad \texttt{As } \|\widetilde{\ell}_0\|_1 \leq \|\widetilde{\ell}_0 + \ell_t^a\|_1 + \|-\ell_t^a\|_1 \\
&\leq e^{\eta} \frac{\eta^K}{2^K} e^{-\eta\|\widetilde{\ell}_0\|_1} = e^{\eta}\psi(\widetilde{\ell}_0). \qquad \texttt{As } \|\ell_t^a\|_1 = \ell_t(a) \leq 1.
\end{aligned}$$

Therefore,

$$\mathbb{P}(A_t = a) \leq e^{\eta} \int_{\mathbb{R}^K} \mathbb{1}\left(a = \operatorname*{argmin}_{j\in[K]} \widetilde{\ell}_0(j) + \ell_t^a(j) + \sum_{s=1}^{t-1} \ell_s(j)\right) \psi(\widetilde{\ell}_0)\mathrm{d}\widetilde{\ell}_0.$$

FTPL for experts: Laplace perturbation (cont'd)

$$\mathbb{P}(A_t = a) \le e^{\eta} \int_{\mathbb{R}^K} \mathbb{1}\left(a = \operatorname*{argmin}_{j\in[K]} \widetilde{\ell}_0(j) + \ell_t^a(j) + \sum_{s=1}^{t-1} \ell_s(j) \right) \psi(\widetilde{\ell}_0) \mathrm{d}\widetilde{\ell}_0.$$

Recall $\ell_t^a(j) = \mathbb{1}(j = a)\ell_t(a)$. Therefore, $\ell_t^a(a) = \ell_t(a)$ and $\ell_t^a(j) \le \ell_t(j)$ for all $j \neq a$. Hence,

$$\mathbb{1}\left(a = \operatorname*{argmin}_{j\in[K]} \widetilde{\ell}_0(j) + \ell_t^a(j) + \sum_{s=1}^{t-1} \ell_s(j) \right) \le \mathbb{1}\left(a = \operatorname*{argmin}_{j\in[K]} \widetilde{\ell}_0(j) + \ell_t(j) + \sum_{s=1}^{t-1} \ell_s(j) \right)$$

Therefore,

$$\begin{aligned}
\mathbb{P}(A_t = a) &\le e^{\eta} \int_{\mathbb{R}^K} \mathbb{1}\left(a = \operatorname*{argmin}_{j\in[K]} \widetilde{\ell}_0(j) + \ell_t(j) + \sum_{s=1}^{t-1} \ell_s(j) \right) \psi(\widetilde{\ell}_0) \mathrm{d}\widetilde{\ell}_0 \\
&\le e^{\eta} \int_{\mathbb{R}^K} \mathbb{1}\left(a = \operatorname*{argmin}_{j\in[K]} \sum_{s=0}^{t-1} \ell_s(j) \right) \psi(\ell_0) \mathrm{d}\ell_0 = e^{\eta}\mathbb{P}(A_{t+1} = a)
\end{aligned}$$

$\square$

FTPL Summary

- Proof strategy:

  FTPL lemma: $$R_T(\pi^{\mathrm{FTPL}}, f) \leq \sum_{t=1}^{T} \mathbb{E}\left[f_t(\omega_t) - f_t(\omega_{t+1})\right] + \mathbb{E}\left[\max_{\Omega} f_0(\omega) - \min_{\Omega} f_0(\omega)\right]$$

  1. Choose $D(\eta)$ so that $\mathbb{E}\left[\max_{\Omega} f_0(\omega) - \min_{\Omega} f_0(\omega)\right] \leq \mathcal{O}\left(\frac{1}{\eta}\right)$.
  2. Show that ω_t and ω_{t+1} have similar distributions.
     Hence argue that $\mathbb{E}\left[f_t(\omega_t) - f_t(\omega_{t+1})\right] \leq \mathcal{O}(\eta^m)$.

  Proof technique for step 2 can depend on D and the problem instance.

- Although high-level intuitions are similar across all FTPL instances, we do not usually have a unified analysis (like FTRL). However, the computational advantages can sometimes make FTPL worthwhile.

- FTPL also does not assume convexity of Ω, f_t.

Ch 6.4: Case study: Online shortest paths

We have a graph with M edges. There are fixed source and and destination vertices.

- There are K possible paths $\mathcal{A} = \{a_1, \ldots, a_K\}$ from the source to the destination. We will represent each path as $a_j \in \{0,1\}^M$ where $a_j(i) = 1$ means that edge i is on a_j. Say that the maximum path length is m, i.e $a_j^\top \mathbf{1}_M \leq m$.
- On each round a learner chooses a path $A_t \in \mathcal{A}$ from source to destination.
- Simultaneously, the adversary chooses losses $\ell_t \in [0,1]^M$ for each edge.
- The learner incurs loss $A_t^\top \ell_t$, but observes ℓ_t (losses on all edges).

Application: packet routing in a network.

We can write the regret as,

$$R_T(\pi, \ell) = \sum_{t=1}^{T} A_t^\top \ell_t - \min_{a_j \in \mathcal{A}} \sum_{t=1}^{T} a_j^\top \ell_t.$$

Attempt 1: Applying Hedge

- Treat each path in $\mathcal{A} = \{a_1, \ldots, a_K\}$ as an expert, and scale the losses by $\frac{1}{m}$.
- The regret for the *scaled losses* will be $\mathcal{O}(\sqrt{T \log(K)})$. Hence,

$$R_T \in \mathcal{O}\left(m\sqrt{T \log(K)}\right) \in \mathcal{O}\left(m\sqrt{mT \log(M/m)}\right) \qquad \texttt{As } K \leq \binom{M}{m} \asymp \left(\frac{M}{m}\right)^m.$$

Attempt 1: Applying Hedge

- Treat each path in $\mathcal{A} = \{a_1, \dots, a_K\}$ as an expert, and scale the losses by $\frac{1}{m}$.
- The regret for the *scaled losses* will be $\mathcal{O}(\sqrt{T \log(K)})$. Hence,

$$R_T \in \mathcal{O}\left(m\sqrt{T\log(K)}\right) \in \mathcal{O}\left(m\sqrt{mT\log(M/m)}\right) \qquad \texttt{As } K \le \binom{M}{m} \asymp \left(\frac{M}{m}\right)^m.$$

- Per-iteration run time is $\mathcal{O}(K)$, which can be large.

Attempt 2: Applying FTPL

FTPL for online shortest paths

- Given: time horizon T, parameter η
- Sample $\ell_0(e) \sim D(\eta)$ for each *edge* e.
- for $t = 1, \ldots, T$,
  - Choose path

$$A_t \leftarrow \operatorname*{argmin}_{a_j \in \mathcal{A}} \sum_{s=0}^{t-1} \ell_s^\top a_j .$$

Run time per iteration:

- Updating losses on each edge (incrementally): $O(M)$
- Computing shortest path via Dijkstra's: $O(M)$. (not convex, but still efficient)
- Much cheaper than $O(K)$ where K could be as large as $\binom{M}{m}$.

In fact, you do not even need to construct $\mathcal{A}$ explicitly.

Bounding the regret of FTPL for online shortest paths

Recall FTPL lemma,

$$R_T(\pi^{\mathrm{FTPL}}, f) \leq \sum_{t=1}^{T} \mathbb{E}\left[f_t(\omega_t) - f_t(\omega_{t+1})\right] + \mathbb{E}\left[\max_{\Omega} f_0(\omega) - \min_{\Omega} f_0(\omega)\right]$$

We will try Laplace perturbations, i.e $\ell_0(e) \sim \mathrm{Lap}\left(1/\eta\right)$ for each edge e.

Applying the FTPL lemma we obtain

$$R_T(\pi^{\mathrm{FTPL}}, f) \leq \sum_{t=1}^{T} \mathbb{E}[\ell_0^\top A_t - \ell_0^\top (A_{t+1})] + \mathbb{E}\left[\max_{a\in\mathcal{A}} \ell_0^\top a - \min_{a\in\mathcal{A}} \ell_0^\top a\right]$$

Let us first bound the second term. By symmetry of the Laplace distribution,

$$\mathbb{E}\left[\max_{a\in\mathcal{A}} \ell_0^\top a - \min_{a\in\mathcal{A}} \ell_0^\top a\right] = 2\mathbb{E}\left[\max_{a\in\mathcal{A}} \ell_0^\top a\right] \leq 2m\mathbb{E}\left[\max_{e\in\mathrm{Edges}} \ell_0(e)\right] \leq \frac{2m}{\eta}\left(1 + \log(M)\right).$$

FTPL for OSP: Laplace perturbation (cont'd)

$$R_T(\pi^{\mathrm{FTPL}}, f) \leq \sum_{t=1}^{T} \mathbb{E}[\ell_0^\top A_t - \ell_0^\top (A_{t+1})] + \underbrace{\mathbb{E}\left[\max_{a \in \mathcal{A}} \ell_0^\top a - \min_{a \in [K]} \ell_0^\top a\right]}_{\leq \frac{2m}{\eta}(1+\log(M)).}$$

To bound the first term, we will use the following claim.

Claim. $\mathbb{P}(A_t = a) \leq e^{m\eta}\mathbb{P}(A_{t+1} = a)$ for all $a \in [K]$, where $\mathbb{P}$ is w.r.t ℓ_0.

We therefore have,

$$\begin{aligned}
\mathbb{E}[\ell_0^\top A_t - \ell_0^\top A_{t+1}] &= \sum_{a_j \in \mathcal{A}} \ell_t^\top a_j \left(\mathbb{P}(A_t = a_j) - \mathbb{P}(A_{t+1} = a_j)\right) \\
&\leq \sum_{a_j \in \mathcal{A}} \underbrace{\ell_t^\top a_j}_{\leq m \times 1} \underbrace{\left(1 - e^{-m\eta}\right)}_{\leq m\eta} \mathbb{P}(A_t = a_j) \leq \eta m^2.
\end{aligned}$$

FTPL for OSP: Laplace perturbation (cont'd)

$$R_T(\pi^{\mathrm{FTPL}}, f) \leq \sum_{t=1}^{T} \underbrace{\mathbb{E}[\ell_0^\top A_t - \ell_0^\top (A_{t+1})]}_{\leq \eta m^2} + \underbrace{\mathbb{E}\left[\max_{a \in \mathcal{A}} \ell_0^\top a - \min_{a \in [K]} \ell_0^\top a\right]}_{\leq \frac{2m}{\eta}(1+\log(M)).}$$

Therefore,

$$\begin{aligned} R_T &\leq m^2 \eta T + \frac{2m}{\eta}(1 + \log(M)) \\ &= 3m\sqrt{mT(1 + \log(K))} \qquad \texttt{By choosing } \eta = \sqrt{(1 + \log(M))/(mT)} \\ &\in \mathcal{O}(m\sqrt{mT\log(M)}). \end{aligned}$$

C.f. For Hedge, $R_T \in \mathcal{O}\left(m\sqrt{mT\log(M/m)}\right)$. While the regret is similar, FTPL has $\mathcal{O}(M)$ computation per round, while Hedge has $\mathcal{O}(K)$, where K could be as large as $\binom{M}{m}$.

FTPL for OSP: Laplace perturbation (cont'd)

Claim. $\mathbb{P}(A_t = a) \le e^{m\eta}\mathbb{P}(A_{t+1} = a)$ for all $a \in [K]$, where $\mathbb{P}$ is w.r.t ℓ_0.

Proof sketch. The proof is similar to Laplace perturbations for Hedge.

Let a path $a_j \in \mathcal{A}$ be given. Then, we can write We can write,

$$\mathbb{P}(A_t = a_j) = \int_{\mathbb{R}^K} \mathbb{1}\left(a = \operatorname*{argmin}_{a_j \in \mathcal{A}} \sum_{s=0}^{t-1} \ell_s^\top a_j\right) \psi(\ell_0)\mathrm{d}\ell_0.$$

Next define $\ell_t^{a_j} \in [0,1]^M$ so that $\ell_t^{a_j}(i) = \ell_t(i) \times a_j(i)$. Use the substitution $\widetilde{\ell}_0 = \ell_0 - \ell_t^a$ and proceed in a similar fashion.