CS861: Theoretical Foundations of Machine Learning Lecture 24 - 11/4/2024
University of Wisconsin–Madison, Fall 2024

Lecture 26: FTPL Laplace perturbation, Intro game theory

Lecturer: Kirthevasan Kandasamy *Scribed by: Yupeng Zhang*

In this lecture we will be looking at FTPL with Laplace Perturbation.

1 FTPL for experts: Laplace perturbation

1.1 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 $t = 1, \ldots, T$,
- $A_t \leftarrow \arg\min_{a\in[K]} \sum_{s=0}^{t-1} \ell_s(a)$. $\omega_t = \arg\min_{\omega\in\Omega} \sum_{s=0}^{t-1} f_s(\omega)$

We will now try $D(\eta) = \text{Lap}(1/\eta)$.
The $\text{Lap}(1/\eta)$ distribution has pdf ψ:

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

1.2 Maximum of K i.i.d Laplace RVs

Let $Z = (Z(1), \ldots, Z(K))$ where $Z(i) \sim \text{Lap}(1/\eta)$,

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

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

We have that,

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

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

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

A common trick: For $Z \geq 0$, $E[Z] = \int_0^\infty zp(z)dz = \int_0^\infty p(z) \int_0^z dvdz = \int_0^\infty \int_v^\infty p(z)dzdv = \int_0^\infty P(Z \geq v)dv$

1.3 FTPL Lemma and Analysis

Lemma 1 (FTPL). *Let $f = (f_1, \dots, f_T)$ be a sequence of losses. Then FTPL satisfies,*

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

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

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

We showed this exact same step for Geometric perturbation. Let us first bound the second term. By symmetry of the Laplace distribution,

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

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

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

Note: $\eta > 0$, so RHS is bigger. But, η will also be small, so it says that your distribution round $t + 1$ does not change from t.

We therefore have,

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

Now we bound both terms of

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

Therefore the regret can be bounded,

$$\begin{aligned}
R_T &\le \eta T + \frac{2}{\eta}(1 + \log(K)) \\
&= 3\sqrt{T(1 + \log(K))} \\
&\in \mathcal{O}(\sqrt{T \log(K)})
\end{aligned}$$

by choosing $\eta = \sqrt{(1 + \log(K))/T}$

1.4 Proof of claim 1

It remains prove the claim 1 above. 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,

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

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

$$P(A_t = a) = \int_{\mathbb{R}^K} \mathbf{1}\left(a = \arg\min_{j\in[K]} \tilde{\ell}_0(j) + \ell_t^a(j) + \sum_{s=1}^{t-1} \ell_s(j) \right) \psi(\tilde{\ell}_0 + \ell_t^a) d\tilde{\ell}_0.$$

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

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

Therefore,

$$P(A_t = a) \leq e^{\eta} \int_{\mathbb{R}^K} \mathbf{1}\left(a = \arg\min_{j\in[K]} \tilde{\ell}_0(j) + \ell_t^a(j) + \sum_{s=1}^{t-1} \ell_s(j) \right) \psi(\tilde{\ell}_0) d\tilde{\ell}_0.$$

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

$$\mathbf{1}\left(a = \arg\min_{j\in[K]} \tilde{\ell}_0(j) + \ell_t^a(j) + \sum_{s=1}^{t-1} \ell_s(j) \right) \leq \mathbf{1}\left(a = \arg\min_{j\in[K]} \tilde{\ell}_0(j) + \ell_t(j) + \sum_{s=1}^{t-1} \ell_s(j) \right)$$

(Note, if a is the minimizer when you add ℓ_t^a, it has to be the case that it minimizes with ℓ_t as the other indices are increasing.)
Therefore,

$$\begin{aligned}
P(A_t = a) &\leq e^{\eta} \int_{\mathbb{R}^K} \mathbf{1}\left(a = \arg\min_{j\in[K]} \tilde{\ell}_0(j) + \ell_t(j) + \sum_{s=1}^{t-1} \ell_s(j) \right) \psi(\tilde{\ell}_0) d\tilde{\ell}_0 \\
&\leq e^{\eta} \int_{\mathbb{R}^K} \mathbf{1}\left(a = \arg\min_{j\in[K]} \sum_{s=0}^{t-1} \ell_s(j) \right) \psi(\ell_0) d\ell_0 \\
&= e^{\eta} P(A_{t+1} = a)
\end{aligned}$$

2 FTPL Summary

2.1 Proof Strategy

FTPL lemma:

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

Key steps:

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

Important notes:

- 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

In fact, Hedge is FTPL with Gumbel perturbation.

3 Online Shortest Paths: A Case Study

3.1 Problem Setting

Given:

- Graph with M edges
- Fixed source and destination vertices
- K possible paths $A = \{a_1, \ldots, a_K\}$ from source to destination
- Each path $a_j \in \{0,1\}^M$ where $a_j(i) = 1$ means edge i is on path a_j
- Maximum path length m, i.e., $a_j^\top 1_M \leq m$

On each round:

- Learner chooses path $A_t \in A$
- Adversary chooses losses $\ell_t \in [0,1]^M$ for each edge
- Learner incurs loss $A_t^\top \ell_t$, but observes ℓ_t (losses on all edges)

Application: packet routing in a network.
Regret:

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

3.2 Attempt 1: Applying Hedge (FTRL)

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

$$R_T \in O(m\sqrt{T \log(K)}) \in O(m\sqrt{mT \log(M/m)})$$

  as $K \leq \binom{M}{m} \sim (\frac{M}{m})^m$
- Per-iteration run time is $O(K)$, which can be large

3.3 Attempt 2: Applying FTPL

Algorithm:

- Given: time horizon T, parameter η
- Sample $\ell_0(e) \sim D(\eta)$ for each edge e
- For $t = 1, \ldots, T$:
  - Choose path $A_t \leftarrow \arg\min_{a_j \in A} \sum_{s=0}^{t-1} \ell_s^\top a_j$

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 $\binom{M}{m}$

Note: the following proof (idea) is similar to Hedge with Laplace, so you can try them out at home.

3.4 Regret Analysis for FTPL with Online Shortest Paths

Using Laplace perturbations $\ell_0(e) \sim \text{Lap}(1/\eta)$ for each edge:

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

By symmetry of the Laplace distribution:

$$\mathbb{E}[\max_{a \in A} \ell_0^\top a - \min_{a \in A} \ell_0^\top a] = 2\,\mathbb{E}[\max_{a \in A} \ell_0^\top a] \leq \frac{2m}{\eta}(1 + \log(M))$$

Claim 2. *$P(A_t = a) \leq e^{m\eta} P(A_{t+1} = a)$ for all $a \in [K]$, where P is w.r.t ℓ_0.*

Final regret bound:

$$R_T \leq m^2 \eta T + \frac{2m}{\eta}(1 + \log(M)) = 3m\sqrt{mT(1 + \log(M))}$$

by choosing $\eta = \sqrt{(1 + \log(M))/(mT)}$.
Comparison:

- For Hedge: $R_T \in O(m\sqrt{mT \log(M/m)})$
- Similar regret, but FTPL has $O(M)$ computation per round, while Hedge has $O(K)$, where K could be as large as $\binom{M}{m}$

3.5 Proof of Claim 2

Claim 3. *$P(A_t = a) \leq e^{m\eta} P(A_{t+1} = a)$ for all $a \in [K]$, where P is w.r.t ℓ_0.*

Proof sketch: The proof is similar to Laplace perturbations for Hedge. Let a path $a_j \in A$ be given. Then:

$$P(A_t = a_j) = \int_{R^K} 1(a = \arg\min_{a_j \in A} \sum_{s=0}^{t-1} \ell_s^\top a_j)\psi(\ell_0) d\ell_0$$

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

4 Learning in Games

Definition: Two-player normal form game.

- In a TPNFG, player 1 has m actions and player 2 has n actions.
- Each player chooses an action (Player 1: a_1, Player 2: a_2) and receive utility/payoff $Q^{(1)}(a_1, a_2)$, $Q^{(2)}(a_1, a_2)$.

e.g. Rock-paper-scissors, the following table is $(Q^{(1)}(a_1, a_2), Q^{(2)}(a_1, a_2))$ when two player choose action R,P,S respectively

$P_1 \backslash P_2$	R	P	S
R	$(0,0)$	$(-1,1)$	$(1,-1)$
P	$(1,-1)$	$(0,0)$	$(-1,1)$
S	$(-1,1)$	$(1,-1)$	$(0,0)$

Acknowledgements

These notes are based on lecture materials prepared in Fall 2024 by Prof Kirthevasan Kandasamy.