CS260: Machine Learning Theory
Lecture 11: Follow the Regularized Leader
October 31, 2011

Lecturer: Jennifer Wortman Vaughan

1 Last Time...

In the last class, we introduced the expert advice framework.

Learning from Expert Advice

At each round $t \in \{1, 2, \cdots, T\}$,

- The learner chooses a distribution $\vec{p}_t$.
- Each expert $i \in \{1, \cdots, n\}$ suffers loss $\ell_{i,t} \in [0, 1]$.
- The learner suffers expected loss $\vec{p}_t \cdot \vec{\ell}_t$.

The regret of the learning algorithm is then defined to be

$$\sum_{t=1}^{T} \vec{p}_t \cdot \vec{\ell}_t - \min_{i \in \{1,\ldots,n\}} \sum_{t=1}^{T} \ell_{i,t}.$$

We first discussed the Randomized Weighted Majority algorithm with parameter η. At each round t, RWM chooses $\vec{p}_t$ by setting

$$p_{i,t} = \frac{e^{-\eta L_{i,t-1}}}{\sum_{j=1}^{n} e^{-\eta L_{j,t-1}}}$$

for all experts $i \in \{1, \cdots, n\}$. We also introduced the class of Follow the Regularized Leader (FTRL) algorithms, that use weights of the form

$$\vec{p}_t = \arg\min_{\vec{p} \in \Delta_n} \left(\eta \sum_{s=1}^{t-1} \vec{\ell}_s \cdot \vec{p} + R(\vec{p}) \right)$$

where $R(\cdot)$ is a convex function called the regularizer, and $\eta > 0$ is a parameter that allows us to adjust the relative impact of the two terms.

Finally, we mentioned a useful fact called Jensen's inequality, which comes up surprisingly often in machine learning.

Theorem 1 (Jensen's Inequality)**.** *For any convex function f and any random variable X, $f(E[X]) \leq E[f(X)]$. Conversely, for any concave function f and any random variable X, $E[f(X)] \leq f(E[X])$.*

One trick for remembering which way the inequalities go is to keep in mind the following pictures.

All CS260 lecture notes build on the scribes' notes written by UCLA students in the Fall 2010 offering of this course. Although they have been carefully reviewed, it is entirely possible that some of them contain errors. If you spot an error, please email Jenn.

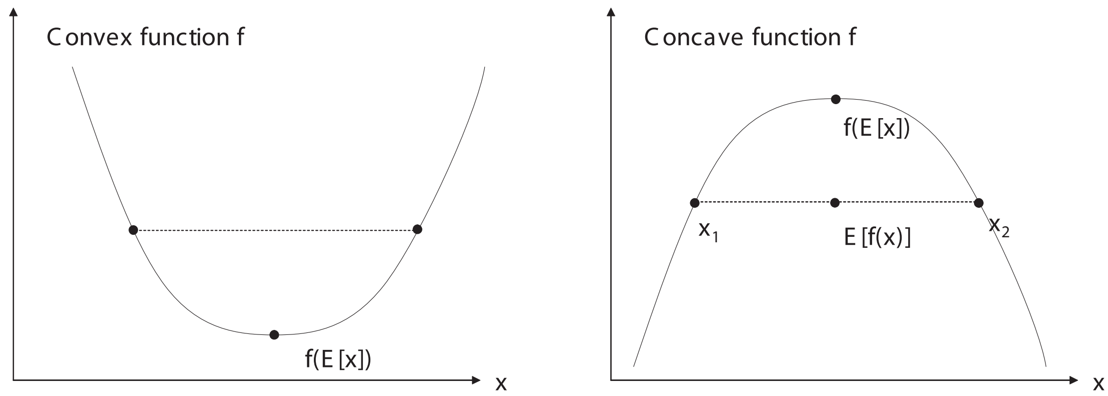

(a) convex graph (b) concave graph

2 Weighted Majority and Entropy

We can prove that RWM is a Follow the Regularized Leader algorithm with

$$R(\vec{p}) = -H(\vec{p}) = -\sum_{i=1}^{n} p_i \log \frac{1}{p_i}.$$

To show this, it is sufficient to show that the distribution $\vec{p}_t$ chosen by RWM at time t is the distribution $\vec{p}$ that minimizes

$$\eta \sum_{s=1}^{t-1} \vec{l}_s \cdot \vec{p} - H(\vec{p}). \tag{1}$$

First note that for any $\vec{p} \in \Delta_n$,

$$\begin{aligned} \eta \sum_{s=1}^{t-1} \vec{l}_s \cdot \vec{p} - H(\vec{p}) &= \eta \vec{L}_{t-1} \cdot \vec{p} - H(\vec{p}) \\ &= \eta \sum_{i=1}^{n} L_{i,t-1} p_i - \sum_{i=1} p_i \cdot \log \frac{1}{p_i} \\ &= \sum_{i=1}^{n} p_i \left(\eta L_{i,t-1} - \log \frac{1}{p_i} \right) \\ &= -\sum_{i=1}^{n} p_i \log \left(\frac{e^{-\eta L_{i,t-1}}}{p_i} \right). \end{aligned} \tag{2}$$

By Jensen's inequality, we then have that for any $\vec{p} \in \Delta_n$,

$$\eta \sum_{s=1}^{t-1} \vec{l}_s \cdot \vec{p} - H(\vec{p}) = -\sum_{i=1}^{n} p_i \log\left(\frac{e^{-\eta L_{i,t-1}}}{p_i}\right)$$
$$\geq -\log\left(\sum_{i=1}^{n} p_i \frac{e^{-\eta L_{i,t-1}}}{p_i}\right)$$
$$= -\log\left(\sum_{i=1}^{n} e^{-\eta L_{i,t-1}}\right) \tag{3}$$

and so this is a lower bound on the quantity that the FTRL algorithm will minimize.

Plugging the RWM distribution $\vec{p}_t$ into Equation 2, we get

$$\eta \sum_{s=1}^{t-1} \vec{l}_s \cdot \vec{p}_t - H(\vec{p}_t) = -\sum_{i=1}^{n} p_{i,t} \log\left(\frac{e^{-\eta L_{i,t-1}}}{p_{i,t}}\right)$$
$$= -\sum_{i=1}^{n} p_{i,t} \log\left(\sum_{j=1}^{n} e^{-\eta L_{j,t-1}}\right)$$
$$= -\log\left(\sum_{i=1}^{n} e^{-\eta L_{i,t-1}}\right) \tag{4}$$

Since the final expression in Equation 4 is equal to the final expression in Equation 3, it must be the case that the Randomized Weighted Majority algorithm chooses the distribution that minimizes Equation 1.

3 Regret Bounds for Follow the Regularized Leader

We will now prove a regret bound that holds for the class of Follow the Regularized Leader algorithms. We start with a useful lemma.

3.1 The Advantage of Knowing the Future

We first prove a lemma that can be viewed as a regret bound for a hypothetical algorithm that chooses the distribution $\vec{p}$ that minimizes

$$\eta \sum_{s=1}^{t} \vec{\ell}_s \cdot \vec{p} + R(\vec{p})$$

at each time t; that is, a hypothetical algorithm that uses the distribution $\vec{p}_{t+1}$ instead of $\vec{p}_t$ at time t. Note that it is not actually possible to run such an algorithm since $\vec{\ell}_t$ is not known to the algorithm at the time when $\vec{p}_t$ is chosen. However, we will be able to use this bound to derive the regret bound for FTRL.

Lemma 1 (Be-the-Regularized-Leader Lemma)**.** *Let $\vec{p}_t$ be the distribution chosen by Follow the Regularized Leader at time t. For any $\vec{p} \in \Delta_n$, for any $\eta > 0$,*

$$\sum_{t=1}^{T} \vec{\ell}_t \cdot \vec{p}_{t+1} - \sum_{t=1}^{T} \vec{\ell}_t \cdot \vec{p} \leq \frac{1}{\eta}\left(R(\vec{p}) - R(\vec{p}_1)\right) . \tag{5}$$

Proof: This proof is by induction on T.

Consider the base case of $T = 0$. The left hand side of the equation is 0. By definition of the algorithm, we know that

$$\vec{p}_1 = \arg\min_{\vec{p}\in\Delta_n} (\eta \cdot 0 + R(\vec{p})) = \arg\min_{\vec{p}\in\Delta_n} R(\vec{p}) \ .$$

Therefore, for any $\vec{p}$,

$$RHS = \frac{1}{\eta}(R(\vec{p}) - R(\vec{p_1})) \geq 0 = LHS,$$

and so our base case holds.

Now suppose that the equation holds for every value up to $T - 1$. Then for all $\vec{p} \in \Delta_n$,

$$\sum_{t=1}^{T-1} \vec{\ell}_t \cdot \vec{p}_{t+1} - \sum_{t=1}^{T-1} \vec{\ell}_t \cdot \vec{p} \leq \frac{1}{\eta}(R(\vec{p}) - R(\vec{p_1})),$$

and

$$\sum_{t=1}^{T-1} \vec{\ell}_t \cdot \vec{p}_{t+1} + \frac{1}{\eta}R(\vec{p_1}) \leq \sum_{t=1}^{T-1} \vec{\ell}_t \cdot \vec{p} + \frac{1}{\eta}R(\vec{p}) \ .$$

Since the inequality holds for *any* distribution $\vec{p}$, we can plug in $\vec{p} = \vec{p}_{T+1}$ and get

$$\sum_{t=1}^{T-1} \vec{\ell}_t \cdot \vec{p}_{t+1} + \frac{1}{\eta}R(\vec{p_1}) \leq \sum_{t=1}^{T-1} \vec{\ell}_t \cdot \vec{p}_{T+1} + \frac{1}{\eta}R(\vec{p}_{T+1}) \ .$$

Adding $\vec{\ell}_T \cdot \vec{p}_{T+1}$ to both sides we get

$$\sum_{t=1}^{T} \vec{\ell}_t \cdot \vec{p}_{t+1} + \frac{1}{\eta}R(\vec{p_1}) \leq \sum_{t=1}^{T} \vec{\ell}_t \cdot \vec{p}_{T+1} + \frac{1}{\eta}R(\vec{p}_{T+1}) \ .$$

By definition, $\vec{p}_{T+1}$ is the distribution $\vec{p}$ that minimizes the right hand side of this equation. Therefore, any other $\vec{p}$ can only increase the value of RHS. Therefore, for *any* $\vec{p} \in \Delta_n$,

$$\sum_{t=1}^{T} \vec{\ell}_t \cdot \vec{p}_{t+1} + \frac{1}{\eta}R(\vec{p_1}) \leq \sum_{t=1}^{T} \vec{\ell}_t \cdot \vec{p} + \frac{1}{\eta}R(\vec{p}).$$

Rearranging terms, we see that Equation 5 holds for T, proving the lemma. $\square$

3.2 The Regret Bound for Follow the Regularized Leader

Using this lemma, we can bound the regret of FTRL. Rearranging Equation 5, we get that for any $\vec{p} \in \Delta_n$,

$$-\sum_{t=1}^{T} \vec{\ell}_t \cdot \vec{p} \leq -\sum_{t=1}^{T} \vec{\ell}_t \cdot \vec{p}_{t+1} + \frac{1}{\eta}\left(R(\vec{p}) - R(\vec{p_1})\right) \ .$$

Adding $\sum_{t=1}^{T} \vec{\ell}_t \cdot \vec{p}_t$ to both sides yields

$$\begin{aligned} \sum_{t=1}^{T} \vec{\ell}_t \cdot \vec{p}_t - \sum_{t=1}^{T} \vec{\ell}_t \cdot \vec{p} &\le \sum_{t=1}^{T} \vec{\ell}_t \cdot \vec{p}_t - \sum_{t=1}^{T} \vec{\ell}_t \cdot \vec{p}_{t+1} + \frac{1}{\eta}(R(\vec{p}) - R(\vec{p}_1)) \\ &= \sum_{t=1}^{T} (\vec{\ell}_t \cdot \vec{p}_t - \vec{\ell}_t \cdot \vec{p}_{t+1}) + \frac{1}{\eta}(R(\vec{p}) - R(\vec{p}_1)) \; . \end{aligned} \tag{6}$$

We can see that the left hand side of this inequality is the regret of FTRL with respect to the function $\vec{p}$. Since this holds for any $\vec{p}$, it holds for the $\vec{p}$ with minimal cumulative loss, which we know will put all of its weight on a single expert. Therefore, this gives us a regret bound if we can bound the terms on the right. We have proved the following theorem.

Theorem 2. *For any sequence of losses $\vec{\ell}_1, \cdots, \vec{\ell}_T$, let $\vec{p}_1, \cdots, \vec{p}_T$ be the distributions chosen by Follow the Regularized Leader with parameter η and regularizer R. Then*

$$\sum_{t=1}^{T} \vec{\ell}_t \cdot \vec{p}_t - \min_{i \in \{1, \cdots, n\}} \sum_{t=1}^{T} \ell_{i,t} \le \sum_{t=1}^{T} (\vec{\ell}_t \cdot \vec{p}_t - \vec{\ell}_t \cdot \vec{p}_{t+1}) + \frac{1}{\eta}\left(\max_{\vec{p} \in \Delta_n} R(\vec{p}) - \min_{\vec{p} \in \Delta_n} R(\vec{p}) \right) \; .$$

Both terms in this bound depend on the choice of regularizer. The first term in the bound measures how quickly the algorithm's distribution changes from one time step to the next, that is, how *stable* the algorithm is. (Remember, as we saw with the Follow the Leader algorithm, instability is bad in this setting.) The first term is a measure of how wide a range of values the regularizer can take on.

Next time we will see how to apply this theorem to obtain a regret bound of $O(\sqrt{T \log n})$ for Randomized Weighted Majority.