Lecture 19: Multiplicative Weights Update

Objectives of this lecture

- Formalize the multiplicative weights update algorithm (MWU).
- Use MWU to prove LP duality

1 Philosophy

If we allow fractional/randomized strategies, the minimum regret in learning from experts can approach the best expert, plus $\log n$. This is done by multiplying weights by $\exp(-\beta)$ per mistake, instead of halving.

The multiplicative weights update algorithm can also be used to solve zero-sum games.

2 Assumed Knowledge

1. Prediction with Experts framework.
2. Weighted majority algorithm and its analysis:
   (a) Initialize weights to 1.
   (b) Everyone with mistakes have weights halved.
   (c) Make prediction on weighted majority.
   (d) If mistake is made, total weight decrease by factor of at least 1/4.
   (e) If expert i made m_i mistakes, the weight is 2^{-m_i}, this needs to be at most the overall potential, so we get bound on total number of mistakes from this.
3. LP duality statement.
4. 0-sum games: mixed / pure strategies, best response.

3 Randomized Experts Algorithm

Recall the prediction with expert advice framework discussed last time: we are trying to make a sequence of online binary (0/1, true/false) predictions with help from n experts. That is, at each round, the algorithm has access to the past truth records, plus the predictions of all experts up til this round. The goal is to make mistakes comparable to the best expert. Let the number of mistakes made by expert i be m_i, our goal is to guarantee at most $C \cdot \min_{i=1}^{n} m_i$ mistakes for a C that's as small as possible.

This lecture we will utilize fractional strategies. This is because any deterministic / pure strategies played at each round (on 0 or 1)

Lemma 1: Lower Bound for Prediction with Experts

For any algorithm, with two experts who predict 0 all the time and 1 all the time respectively, there exists an input sequence where the algorithm is > 2-competitive.

Proof. Generate a truth sequence that's opposite of what the algorithm predicts at each step.

After T steps, the algorithm makes T mistakes, but one of all 0 or all 1 makes at most $T/2$ mistakes. So the competitive ratio of the algorithm is at least $T/(T/2) = 2$. □

Note however that a coin toss here makes $T/2$ mistakes on average. So randomized strategies can achieve significantly fewer errors.

We can take this idea to the weighted majority algorithm: if the total weight on 0 is $w0$ and the total weight on 1s is $w1$, then we play a 'mixed strategy' where we predict 0 with probability $\frac{w0}{w0+w1}$ and 1 with probability $\frac{w1}{w0+w1}$.

Now consider how much the total potential drops in either case.

WOLOG by symmetry first consider the case where the truth is 0. Then the weights in the $w1$ side gets halved, so the total weight decreases additively by $\frac{w1}{2}$, or a multiplicative decrease by a factor of

$$\frac{w1}{2(w_0 + w_1)}.$$

On the other hand, the probability of making an error in this case (where truth is 0) is $\frac{w1}{w0+w1}$. Note that this is almost the exact form as above. So we can relate the overall potential function to the number of errors more closely algebraically. Specifically, if the expected number of mistakes made at step t is $M^{(t)}$, we get

$$\Phi^{(t+1)} = \left(1 - \frac{M^{(t)}}{2}\right)\Phi^{(t)}$$

where per entry we still have

$$w_i^{(t)} = \left(1 - \frac{1}{2}\right)^{\sum_{k=1}^{t} m_i^{(k)}}$$

The 2s are both in the denominators, so all we really need to do is to move everyting into the exponent. We do so via the following fact.

Fact 1: Taylor expansion of exponentials

As $\delta \to 0$, $\exp(\delta) \approx 1+\delta$. (we drop the second order term throughout these presentations to simplify analyses)

With this fact, we can modify the algorithm by picking some parameter $\beta \to 0$, and:

1. set the decrease factor per mistake to $\exp(-\beta)$
2. still play fractional strategies based on the probability distributions.

With these modifications, the potential per entry becomes

$$w_i^{(t)} = \exp\left(-\beta \sum_{1 \le k \le t} m_i^{(t)}\right) = \exp\left(-\beta\, m_i\right),$$

while the overall potential simplifies to

$$\Phi^{(t)} \approx \left(1 - \beta M^{(t)}\right)\Phi^{(t-1)} \approx \exp\left(-\beta M^{(t)}\right)\Phi^{(t-1)} = n \cdot \exp\left(-\beta M\right).$$

So we get (up to hand waving due to dropping second order terms) the inequality

$$\exp\left(-\beta\, m_i\right) \le_{\approx} n \exp\left(-\beta M\right)$$

or that

$$M \le_{\approx} m_i + \frac{\ln n}{\beta}.$$

The actual rigorous analyses have an extra factor of $1+\beta$ in the competitive ratio from the second order terms. We might or might not cover it in a future optional class.

4 Proof of LP Duality

This is meant to be an extreme example of what one can do with MWU. It is for enrichment purposes only.

4.1 Equivalence of Zero-Sum Games and LP Duality

After:

1. turning everything into inequalities,
2. putting an arbitrarily large linear term on the sum,
3. rescaling all inequalities so the RHS are 1 (slightly pad the $b_i = 0$s away from 0),

LP duality is equivalent to

Theorem 1: LP Duality as Zero Sum Game

For any payoff matrix $A \in \mathbb{R}^{m \times n}$, there is a probability vector $x \in \mathbb{R}^n$ such that $Ax \leq \mathbb{1}$ (entrywise) if and only if there does not exist a probability vector $y \in \mathbb{R}^m$ such that $A^\top y > \mathbb{1}$ (entrywise).

Furthermore, we can assume all entries in A are positive by adding x to all its entries.

4.2 MWU Algorithm

Consider playing this game for T rounds against a dual player. That is, we provide a pure strategy response $x^{(t)}$ for $t = 1 \ldots T$, against a dual player playing a mixed strategy $y^{(1)}, y^{(2)}, \ldots, y^{(t)}$.

Each step, the strategy picked by the dual player is to look at the sum of the xs so far, and pick vector proportional to the exponential of the vector Ax. That is, the weights in the MWU at iteration t are set to

$$w_i^{(t)} \leftarrow \exp\left(\beta \sum_{k=1}^{t-1} \left(Ax^{(k)}\right)_i\right)$$

then the $y^{(t)}$ vector is created by normalizing the $w^{(t)}$ so that it's a probability vector:

$$y_i^{(t)} \leftarrow \frac{w_i^{(t)}}{\sum_j w_j^{(t)}}$$

At each step, two possibilities occur:

1. $A^\top y^{(t)} \geq \mathbb{1}$, at which point we found the dual probability vector.
2. or there is some entry $j(t)$ such that $(A^\top y^{(t)})_{j(t)} < 1$. By picking

   $$x^{(t)} \leftarrow \mathbb{1}\left(j(t)\right),$$

   the vector that's 1 in location $j(t)$ and 0 everywhere else, we can guarantee that $(y^{(t)})^\top A x^{(t)} < 1$.

This is the key condition. Specifically, the bound on dot product enables us to control the growth of the potential function,

$$\Phi^{(t)} = \sum_{i=1}^{m} w_i^{(t)}.$$

For this, we use the first order Taylor expansion of $\exp(w + \Delta w) \approx \exp(w)(1 + \Delta w)$ to get

$$w_i^{(t+1)} = w_i^{(t)} \exp\left(\beta \left(Ax^{(t)}\right)_i\right) \approx w_i^{(t)} \left(1 + \beta \left(Ax^{(t)}\right)_i\right)$$

Summing across all i, plugging in the short hand of the sum of the w_is equaling to Φ ($\Phi^{(t)} = \sum_j w_j^{(t)}$) gives

$$\Phi^{(t+1)} \approx \Phi^{(t)} + \beta \Phi^{(t)} \sum_i \left(y_i^{(t)} \left(Ax^{(t)}\right)_i\right) \leq (1 + \beta)\Phi^{(t)}$$

where for the last line we used the bound on the payoff of $x^{(t)}$ against the mixed strategy $y^{(t)}$.

But $1+\beta \approx \exp(\beta)$, and the initial potential is $\exp(0)=1$ on all entries. So we get the bound of

$$\Phi^{(t)} \leq m \exp\left(\beta t\right)$$

at all times.

Now think of the entry-wise value of $A(x^{(1)}+x^{(2)}+\ldots x^{(t)})$, specifically $(A(x^{(1)}+x^{(2)}+\ldots x^{(t)}))_i$. This value contributes to $w_i^{(t+1)}$, so is at most $\Phi^{(t+1)}$.

So we get the max of this modified mixed strategy is at most

$$\frac{\ln m + t\beta}{t\beta} = 1 + \frac{\ln m}{t\beta}.$$

β is a fixed parameter, so as long as we let t to be much bigger than $\beta^{-1}\ln m$, we can get things to arbitrarily approach 1.

4.3 Why we don't actually solve LPs like this

This is slow. The Taylor approximation of $\exp(\delta) \approx 1+\delta$ (Fact 1) only kick in when $\delta \to 0$. To invoke it on $\exp(\beta(Ax)_i)$, β needs to be less than the inverse of the max entry in Ax. Because it's possible for x to be concentrated on one entry, $(x = \mathbb{1}(j))$, the only 'safe' choice of β is

$$\beta \leq \frac{1}{\max_{ij} A_{ij}}.$$

When plugged back into the reductions from general LPs to 0-sum games, the resulting iteration count from $\log n/\beta$ is often enormous (and are indeed seen in practice).

However, such MWU based ideas do lead to reasonable algorithms on graphs. You will see how playing the path interception game as oralhw3p2 leads to a flow algorithm in this Friday's recitation.