
Lecture 5

1 Overview

This lecture discusses applications of the experts framework to finding approximate solutions of two person zero-sum games, boosting and solving linear programs. The lecture is based on the survey paper by Arora, Hazan, Kale linked on the webpage.

1.1 Recap of experts framework

The experts framework consists of n experts who each make predictions every day, and an online algorithm that makes a decision after seeing all the expert predictions. At the end of the day the loss/gain for each expert is revealed. The goal of the online algorithm is to perform not much worse than the single best expert over a long period of time.

We analyzed the algorithm where the initial weights of all experts is 1, the online algorithm follows expert i with probability proportional to the weight w_i , and the weights are updated as $w_i(t+1) = w_i(t)(1 - \epsilon)^{\ell_i^{(t)}}$. The expected loss L for this online algorithm is bounded in terms of the loss L^* for the best expert,

$$L \leq (1 + \epsilon)L^* + \frac{\ln N}{\epsilon} \quad (1)$$

Recall that the algorithm was analyzed by bounding the potential function $W(t) = \sum_i w_i(t)$, the bound on expected loss was obtained using,

$$(1 - \epsilon)^{L^*} \leq W(T) \leq ne^{-\epsilon L} \quad (2)$$

The lower bound was established by noting that $W(T)$ is at least the weight of the best expert while the upper bound was derived by showing that $W(t)(1 - \epsilon L_t) \geq W(t+1)$ where L_t is the expected loss for the online algorithm in round t .

2 Two person zero sum games

A two person zero sum game is specified by a $m \times n$ payoff matrix A , where A_{ij} is the amount paid by the row player to the column player if strategies (i, j) are played. The expected payoff for a pair of mixed strategies (x, y) is $x^t A y$.

The value of the game is $R = \max_y \min_x x^t A y$ if the row player plays second and $C = \min_x \max_y x^t A y$ if the column player plays second. In particular, there is a column strategy y^* that achieves payoff at least R against any row strategy, and there is a row strategy x^* that pays at most C against any column strategy.

Playing (x^*, y^*) against each other we have $R \leq x^* A y^* \leq C$, the minimax theorem from lecture 2 asserts that $C = R$. We will use the experts framework to compute ϵ optimal strategies for the game.

Wlog we consider payoff matrices A with positive weights $a_{ij} \in [0, 1]$. Recall that an equilibrium consists of a pair of mixed strategies (x, y) , where neither player can gain by changing strategies. An ϵ approximate equilibrium is a strategy pair such that no player can gain more than ϵ by deviating i.e. each of the inequalities below is tight within ϵ ,

$$\min_{i \in [m]} (Ay)_i < x^t A y < \max_{j \in [n]} (x^t A)_j \quad (3)$$

The experts algorithm for finding ϵ approximate equilibria is the following:

1. For $T = O(\frac{\log n}{\epsilon^2})$ rounds repeat steps 2 and 3:
2. The m strategies of the row player are the experts, the initial weights are $w_i = 1$. The mixed strategy $x_i(t) = \frac{w_i(t)}{W(t)}$ is played in round t and the weights are updated to $w_i \rightarrow w_i(1 - \epsilon)^{l_i^t}$, where l_i^t is the loss for the i -th strategy in round (t) .
3. The column player plays the best response $y(t)$ to the row player's strategy $x(t)$.
The 'oracle' : The best response against strategy x is found by computing the largest coordinate $\max_{j \in [n]} (x^t A)_j$.
4. Output: The average strategies $x = \frac{1}{T} \sum_t x(t)$ and $y = \frac{1}{T} \sum_t y(t)$ are a pair of 2ϵ -optimal strategies.

Analysis: Let M be the total loss suffered by the algorithm over T rounds. In each round the column player plays the best response against the row strategy, in particular the column player can guarantee a payoff at least R for every round,

$$RT \leq M \quad (4)$$

On the other hand, the analysis of the expert's algorithm bounds M in terms of the loss of the best expert i.e. in terms of $\min Ay$ the value of the row player's best response to y ,

$$M \leq (1 + \epsilon) \min_{i \in [m]} (Ay)_i \cdot T + \frac{\ln n}{\epsilon} \quad (5)$$

The best response of the row player guarantees that the payoff is at most C , combining equations (4) and (5) we have,

$$R \leq M \leq (1 + \epsilon)C + \frac{\ln n}{\epsilon T} \leq C + 2\epsilon \quad (6)$$

We used $C \leq 1$ which follows as the entries of the payoff matrix belong to $[0, 1]$. (If the entries of the payoff matrix were in $[0, \rho]$ then the same analysis goes through if we scale loss l_i^t by ρ in step 2, the time T gets multiplied by ρ). It remains to show that the strategies x and y are 2ϵ -optimal, the strategy will be to bound the difference between $\min Ay$ and $\max x^t A$ as in equation (3).

The row player would suffer a loss of $T \max_{j \in [n]} (x^t A)_j$ if the column player always played the strategy that is an optimal response to x , however the column player does better than playing a fixed strategy by choosing the optimal response for each round,

$$\max_{j \in [n]} (x^t A)_j \leq \frac{M}{T} \quad (7)$$

Using the experts analysis from equation (5) we have,

$$\max_{j \in [n]} (x^t A)_j - \min_{i \in [m]} (A y)_i \leq \frac{M}{T} - \min_{i \in [m]} (A y)_i \leq 2\epsilon \quad (8)$$

It follows that the output of the experts algorithm is a pair of 2ϵ optimal strategies.

3 Boosting

The AdaBoost algorithm is an enormously influential application of the multiplicative weights framework to machine learning. Applications of the multiplicative weights method can be seen as using a weak primitive encapsulated as a computational oracle to solve a harder problem, thereby ‘boosting’ the power of the weak primitive. This will be a recurrent theme as we see further applications of the paradigm.

Learning: We are given a set of n labeled examples $(x_1, l_1), \dots, (x_n, l_n)$ where the labels $l_i \in \{\pm 1\}$ are binary valued. A weak learner takes as input a distribution D on the examples and produces a ‘hypothesis’ $h(\cdot)$, such that $\Pr_{x \sim D}[h(x) = l(x)] \geq \frac{1}{2} + \gamma$. We will boost the weak learner using the multiplicative weights framework to produce a hypothesis that correctly outputs the labels for $1 - \nu$ fraction of the points.

The simplest hypothesis is good: What should be the size of the training set S such that if a hypothesis $h \in \mathcal{H}$ agrees with $1 - \nu$ fraction of the examples in S , then with probability $(1 - \delta)$ h agrees with $(1 - \epsilon)$ fraction of all data points? As samples in S are drawn uniformly at random, the probability that there exists a hypothesis $h \in \mathcal{H}$ with error probability more than ϵ that correctly classifies $(1 - \nu)$ fraction of the points in S is at most $|\mathcal{H}| \cdot (1 - \epsilon)^{(1 - \nu)|S|}$ by the union bound. The size of the training set can therefore be bounded in terms of the representation size of a hypothesis from \mathcal{H} ,

$$(1 - \nu)|S| \geq \frac{-1}{\ln(1 - \epsilon)} (\ln |\mathcal{H}| + \ln(1/\delta))$$

Better bounds on $|S|$ can be obtained where $\ln |\mathcal{H}|$ gets replaced by the VC -dimension, for our purposes it suffices to note that finding a hypothesis that agrees with a $(1 - \nu)$ fraction of a sufficiently large training set solves the learning problem.

Boosting: The experts algorithm for boosting is the following:

1. For $T = \frac{2}{\gamma^2} \log \frac{1}{\nu}$ rounds repeat the following:
2. The experts are the n examples, we start with a uniform distribution on the examples. The column player plays a hypothesis h and the payoff is 1 if $h(x) = l(x)$ and 0 otherwise. The row player follows the experts algorithm on the examples.

3. The column player uses the weak learner as an oracle to produce hypothesis h with expected payoff more than $1/2 + \gamma$.
4. The output hypothesis $h(x)$ is the majority of $h_1(x), h_2(x), \dots, h_T(x)$.

Analysis: Using the basic form (2) of the analysis of the experts algorithm,

$$W(T) \leq e^{-\epsilon L} n \leq e^{-\epsilon(\frac{1}{2} + \gamma)T} n \quad (9)$$

The inequality follows as $h_t(x)$ predicts correctly with probability at least $\frac{1}{2} + \gamma$ for all t , hence the total loss $L \geq (\frac{1}{2} + \gamma)T$.

Let S be the set of points for which the majority hypothesis is incorrect, a point in S incurs a loss of at most $T/2$ over the course of the algorithm.

$$|S|(1 - \epsilon)^{T/2} \leq W(T) \quad (10)$$

Combining the upper bound and the lower bound, choosing $\epsilon = \gamma$ and taking logarithms we have

$$\ln\left(\frac{|S|}{n}\right) + \frac{T}{2} \ln(1 - \gamma) \leq -\gamma T \left(\frac{1}{2} + \gamma\right) \quad (11)$$

Using $-\gamma - \gamma^2 \leq \ln(1 - \gamma)$ (proved in lecture 4), we have,

$$\ln\left(\frac{|S|}{n}\right) \leq -\frac{\gamma^2 T}{2} = \log \nu \quad (12)$$

The fraction of mistakes is not more than ν , hence the fraction of points that are correctly classified is least $1 - \nu$.

4 Additional material

Some additional applications of the experts framework that were not covered in class.

4.1 Set cover

The set cover problem is to cover all the elements in $[n]$ using the smallest number of sets from a collection \mathcal{C} . The problem can be modeled as a two player game where the row player plays an element $x \in [n]$, the column player plays sets from \mathcal{C} . The payoff is equal to 1 if $x \in C$ and is 0 otherwise.

1. The row player follows the experts algorithm with $\epsilon = 1$, elements that get covered are never played again.
2. The column player plays the best response which is the greedy strategy of choosing the set that covers the largest number of remaining elements.

Analysis: If the size of the optimal set cover is k , for every distribution of weights on n there is a set that covers at least $1/k$ fraction of the weights. For every round the total weight decreases by at least $(1 - 1/k)$,

$$W(t+1) \leq W(t)(1 - 1/k) \Rightarrow W(T) \leq ne^{-T/k}$$

The weight decreases to 0 when the sets played by the column player cover all the n elements. Taking logarithms we have $T \leq k \log n$ if $W(T) \geq 1$. It follows that the greedy algorithm produces an $O(\log n)$ approximation for set cover.

4.2 Solving linear programs

The problem is to produce an approximately feasible solution for the linear program $Ax \geq b, x \in P$ where $P \in \mathbb{R}^n$ is a convex set. The setup is that the m constraints given by $Ax \geq b$ are hard to satisfy, while it is easy to check membership in P . Finding approximately feasible solutions suffices to optimize any linear objective function using the binary search trick.

The experts framework reduces the problem of checking feasibility of m constraints to the problem of checking the feasibility of a single constraint.

1. There are m experts corresponding to the constraints $A_i \cdot x \geq b_i$, if the column player plays $x \in P$ the loss function for the i -th expert is $A_i \cdot x - b_i$.
2. The row player follows the experts algorithm, if the weights of the row player are w_i the column player queries the oracle with input $\sum (w_i A_i) x_i \geq \sum w_i \cdot b_i$ and plays the x returned by the oracle.

The ‘Oracle’: Given a constraint $c^T x \geq d$, the oracle finds $x \in P$ satisfying the constraint and outputs ‘no answer’ if the constraint can not be satisfied for $x \in P$.

3. The output $x = \frac{1}{T} \sum_{t \in T} x_t$ satisfies all the LP constraints within a margin of ϵ for sufficiently large $T = \frac{2\rho \ln n}{\epsilon^2}$.

Analysis: The convex set P is assumed to be bounded so there is a width parameter ρ such that $A_i \cdot x - b_i \in [-\rho, \rho]$ for all i , the value of ρ can be guessed by binary search. If the program is feasible, the oracle query returns an answer x_t for every round as the input to the oracle is a linear combination of feasible inequalities. For example, any feasible solution x satisfies these inequalities.

If the oracle returns ‘no answer’ for any round the program is infeasible. For each round the expected loss of the row player is $\sum w_i (A_i \cdot x - b_i) \geq 0$ by design of the oracle. The output $x \in P$ by convexity and for $T = \frac{2\rho \ln n}{\epsilon^2}$ the analysis of the expert algorithm guarantees,

$$0 \leq LT \leq \sum_t (A_i \cdot x_t - b_i) + \epsilon \quad \forall i \in [n]$$

The average solution $x = \frac{1}{T} \sum_{t \in T} x_t$ satisfies all the LP constraints with a margin of ϵ .

4.3 The AHK paper

Variations of the experts algorithm can be used to deal with simultaneous gains and losses. This can be done in a variety of ways, and lead to slightly different error bounds. For positive losses in the range $[0, \rho]$, one can scale the terms by setting $T \rightarrow T\rho$ to average out the additional ρ factor at every step, refer to the *AHK* the paper for ways to deal with negative losses and more details.

The paper also has an application to complexity theory, they show that if functions are a little difficult to compute, that there must be a set of inputs where they are very difficult to compute.