
Prediction and Learning: It's Only a Game - Homework #1

Jacob Abernethy

Due: 9/25/2013

Homework Policy: Working in groups is fine. Please write the members of your group on your solutions. There is no strict limit to the size of the group but we may find it a bit suspicious if there are more than 4 to a team. Questions labelled with **(Challenge)** are not strictly required, but you'll get some participation credit if you have something interesting to add, even if it's only a partial answer.

Preliminaries: I'm going to use a lot of big-Oh notation, including things like $O(\cdot)$, $o(\cdot)$, $\omega(\cdot)$, $\Omega(\cdot)$, $\Theta(\cdot)$. This is probably familiar to the CS folks but may be less so for students in other departments. If you're not familiar with this yet I recommend wikipedia or other online sources! These concepts will be utilized more and more in lecture.

1) **Tuning Parameters.** We are going to imagine we have some algorithm \mathcal{A} with a performance bound that depends on some input values (which can not be adjusted) and some tuning parameters (which can be optimized). We will use greek letters (α, η, ζ , etc.) for the tuning parameters and capital letters (T, D, N , etc.) for inputs. We would like the bound to be the tightest possible, up to multiplicative constants. For each of the following, tune the parameters to obtain the optimal bound. Using *big-Oh* notation is fine to hide constants, but you *must not* ignore the dependence on the input parameters. For example, assuming $M, T > 0$, imagine we have a performance guarantee of the form:

$$\text{Performance}(\mathcal{A}; M, T, \epsilon) \leq M\epsilon + \frac{T}{\epsilon} \tag{1}$$

and we know $\epsilon > 0$. Then by optimizing the above expression with respect to the free parameter we can tune $\epsilon \leftarrow \sqrt{\frac{T}{M}}$. With this value we obtain $\text{Performance}(\mathcal{A}; M, T, \epsilon) = O(\sqrt{MT})$

NOTE: I didn't have to make up this problem, I actually pulled all the bounds below from different papers!

(a) $\text{Performance}(\mathcal{A}; T, N, \eta) \leq T\eta + \frac{\log N}{\eta} + \eta^2$

(b) $\text{Performance}(\mathcal{A}; T, N, \eta) \leq \frac{\log N + \eta T}{1 - \exp(-\eta)}$

(c) $\text{Performance}(\mathcal{A}; T, \eta) \leq \max(T\eta, \eta^{-2})$

(d) $\text{Performance}(\mathcal{A}; T, \eta) \leq \frac{T}{\eta} + \exp(\eta)$. (Note: you needn't obtain the optimal choice of η here or the tightest possible bound, but try to tune in order to get a bound that is $o(T)$ – i.e. the bound should grow strictly slower than linear in T .)

(e) $\text{Performance}(\mathcal{A}; N, T, \eta, \epsilon) \leq \frac{T\epsilon}{\eta} + \frac{N}{\epsilon} + T\eta$

(f) $\text{Performance}(\mathcal{A}; T, \delta, \alpha) \leq \frac{\sqrt{T}}{\delta} + \frac{T\delta}{\alpha} + \alpha T$

2) **The Doubling Trick.** Let's imagine that an online learning algorithm that runs in T rounds has the bound in Eq. (1). By optimally tuning ϵ we obtain a bound of the form $O(\sqrt{TM})$. The problem with this approach is that it requires us to know T in advance. Is there a way around this issue?

Imagine constructing a modified algorithm \mathcal{A}' that does the following iterative procedure. \mathcal{A}' starts with an initial parameter ϵ_1 , implements algorithm \mathcal{A} using this parameter for T_1 rounds, then adjusts the parameter to ϵ_2 , and runs \mathcal{A} for T_2 rounds, and so forth. Let's say ϵ gets updated k times, where $T_1 + T_2 + \dots + T_k = T$. You can also assume that $\text{Performance}(\mathcal{A}'; T, M) = \sum_{i=1}^k \text{Performance}(\mathcal{A}; T_i, M, \epsilon_i)$.

- (a) Can you construct a schedule for the sequence of (ϵ_i, T_i) that achieves the same bound (up to a multiplicative constant factor) as the optimally tuned bound (namely, $O(\sqrt{MT})$) that requires knowing T in advance? (Note: this is referred to as a "doubling trick".) In other words, you want to choose the sequence of sequence of T_1, T_2, \dots , with the associated $\epsilon_1, \epsilon_2, \dots$ so that whenever the game truly ends, at a previously unknown T , the bound $\text{Performance}(\mathcal{A}'; T, M) = O(\sqrt{MT})$ will always hold.
- (b) **(Challenge)** How about a similar doubling trick to obtain the optimal bound from problem 1 part (e) above? This requires modifying two parameters!

3) **Exponential Weights Algorithm with a Prior.** The Exponential Weights Algorithm had the initial weights w_1^1, \dots, w_n^1 all set to 1. What if instead we imagine an algorithm \mathcal{A}' where we set these weights according to $w_i^1 = p_i$ where \vec{p} is some distribution (i.e. $p_i \geq 0$ for all i and $\sum_i p_i = 1$). We will do the same multiplicative update rule as before.

- (a) Prove that with this modified algorithm we achieve the following bound: for any expert i we have that

$$\text{Loss}_T(\mathcal{A}) \leq \frac{\log p_i^{-1} + \eta \text{Loss}_T(\text{expert } i)}{1 - e^{-\eta}}$$

- (b) **(Challenge)** We would love to obtain a bound that has optimally tuned η for the best choice of i . That is, we would hope that we can guarantee the following:

$$\text{Loss}_T(\mathcal{A}) \leq \min_i \left(\text{Loss}_T(\text{expert } i) + \log p_i^{-1} + \sqrt{\log p_i^{-1} T} \right).$$

The problem is that we can't determine the best i in advance. Is it possible to get something like this, or close, using a doubling trick? (Note: I don't have a full answer to this question myself! But I do know there are some difficulties here :-))

4) **Subsets as Experts.** We saw that when we wanted to do "predictive sorting" it's not easy to apply the halving algorithm to the class of all permutations (rankings) as experts. But this isn't the case for all classes of "complex experts".

Imagine a setting where we have N experts but our goal is not to choose one but $k < N$ of them on each round! We can imagine having a "hyperexpert" for each subset $S \subset [N]$, with $|S| = k$, of which there are clearly $\binom{N}{k}$. Let \mathcal{S}_k^N be the set of all k -sized subsets of $[N]$. On round t , each "base" expert i suffers loss ℓ_i^t which implies that the hyperexpert S suffers loss

$$\text{loss}_t(S) := \sum_{i \in S} \ell_i^t,$$

that is, the hyperexpert loss is additive across the base experts in the subset. Our aim is to run the Exponential Weights Algorithm on these subset hyperexperts. With this well-defined loss on each hyperexpert and a given parameter η , we can define the weight update $w_S^{t+1} = w_S^t \exp(-\eta \text{loss}_t(S))$.

In many scenarios in which we are dealing with hyperexperts, it's suitable to compute the "projected" weights for each base expert i . That is, assume we can run our algorithm by simply knowing $u_i^t := \sum_{S \in \mathcal{S}_k^N: i \in S} w_S^t$. In other words, if our weights define a distribution over \mathcal{S}_k^N , then the value u_i^t corresponds to the (unnormalized) marginal probability of i being in a randomly drawn subset. Can we obtain these values efficiently? If so, how?

Hint: Given a vector of positive values $v_{1:n} := \langle v_1, \dots, v_n \rangle$, we can define

$$\text{SumProd}(v_{1:n}, k) := \sum_{S \in \mathcal{S}_k^N} \prod_{i \in S} v_i$$

Naively this requires $O(n^k)$ computation, but perhaps there is a faster way to compute SumProd?