15-451/651 Algorithm Design & Analysis, Spring 2026

Quiz 6 Solutions

Weather in Pittsburgh is a ternary variable from the set $\{\text{RAIN}, \text{SUMMER}, \text{SNOW}\}$.

There are n **really bad** experts (aka. groundhogs) making weather predictions.

We would like to be as bad as the worst of them. That is, we want to be right about the weather at most 10 times as often as the worst expert, plus some additive $\log n$s.

Formally, let the ground truth over t days be $x(1)\dots x(t)$, and the predictions of the n experts on day t be $y(t)_1, \dots, y(t)_n$ respectively. So the number of times expert i is correct after the t days is

$$r(t,i) = \left|\left\{1 \le j \le t \mid |y(j)_i == x(j)\right\}\right|.$$

Give and prove an algorithm that generates predictions $a(t)$ based only on:

1. the outcomes of previous days,
2. the predictions/ up til day t previous predictions,
3. its previous predictions,

such after t days, the algorithm is right at most $10(\min_{1\le i\le n} r(t,i)) + 10\log_2 n$ times.

Note: unlike the in-class quizzes/midterms, your solution must be entirely self-contained, and cannot cite lecture notes / homework solutions. However, this quiz is open book/notes.

Solution. We initialize a system with n experts, each with weight 1. On each day t we:

1. sum the weights of the experts predicting each of the 3 possible weather outcomes, predicting the outcome with the highest total weight.
2. halve the weights of the experts who predicted correctly.

Define potential function $\Phi(t)$ be the total weight of all experts after t days.

$$\Phi(t) = \sum_{i=1}^{n} w(t)_i = \sum_{i=1}^{n} \left(\frac{1}{2}\right)^{r(t,i)}$$

Initially, no predictions have been made, so $\Phi(0) = n$.

Since there are $k = 3$ weather outcomes, the algorithm's prediction (the outcome with the maximum total weight) must represent at least a 1/3 fraction of the total weight Φ_t.

When the algorithm made a correct prediction, all the experts who voted for that outcome was correct, and have their weights halved. So the total weight goes down by a factor of at least $\frac{1}{2}\cdot\frac{1}{3} = \frac{1}{6}$, or in other words

$$\Phi(t+1) \le \frac{5}{6}\Phi(t)$$

Also, note that on other days, the potential does not increase, since we only decrease the weights of experts (when they make a correct prediction).

We let C be the total number of times the algorithm is correct over t days. Then we have

$$\Phi(t) \leq n\left(\frac{5}{6}\right)^{C},$$

while on the other hand $\Phi(t) \geq w(t)_i$ gives

$$\Phi(t) \geq \left(\frac{1}{2}\right)^{r(t,i)}.$$

Combining these gives

$$\left(\frac{1}{2}\right)^{r(t,i)} \leq n\left(\frac{5}{6}\right)^{C},$$

or

$$C \leq \log_{\frac{6}{5}} n + r(t,i)\log_{\frac{6}{5}} 2 \leq 10r(t,i) + 10\lg n$$

Aside: in an earlier version that the TAs thankfully caught, we tried to be as wrong as the worst ground hog. This turns out to be impossible: imagine a hog that always predicts 'summer', and another one that always predicts 'snow', and then nature conspires against us by ensuring all of our predictions are correct. Some hog is at least wrong half the time, while we make 0 errors.