

In the last weeks, we have often seen that randomized algorithms admit better competitive ratios than any deterministic one. We could lower-bound the performance of a deterministic algorithm by using the perspective of an adversary that tried to make our algorithm look as badly as possible. Today, we will turn our attention to lower bounds and randomized (online) algorithms.

1 Yao's Minimax Lemma

Yao's Minimax Lemma is a very simple, yet powerful tool to prove impossibility results regarding worst-case performance of randomized algorithms, which are not necessarily online. We state it for algorithms that always do something correct but the profit or cost may vary. Such algorithms are called *Las Vegas* algorithms. We first state it for minimization problems because this is the usual way.

We assume that we have a class of deterministic algorithms \mathcal{A} and a class of instances \mathcal{X} . In order to avoid technicalities, assume that both classes are finite. Algorithm $a \in \mathcal{A}$ on instance $x \in \mathcal{X}$ incurs cost $c(a, x) \in \mathbb{R}$. A randomized algorithm is simply a probability distribution over the set of deterministic algorithms \mathcal{A} . So, let A be a randomized algorithm (which is now a random variable), then A 's worst-case cost is $\max_{x \in \mathcal{X}} \mathbb{E}c(A, x)$.

Theorem 1 (Yao's Minimax Lemma). *Let A be any random variable with values in \mathcal{A} and let X be any random variable with values in \mathcal{X} . Then,*

$$\max_{x \in \mathcal{X}} \mathbb{E}c(A, x) \geq \min_{a \in \mathcal{A}} \mathbb{E}c(a, X) .$$

Before proving the theorem, let us interpret what it means. The left-hand side of the inequality is what we will try to lower-bound: It is the worst-case performance of randomized algorithm A . The right-hand side will be easier to talk about, because algorithms are deterministic. This is a sort of average-case performance of the best deterministic algorithm in our class. The distribution over instances is arbitrary.

Roughly, the intuition for the proof is the following: if the input to a cost-minimization algorithm is chosen at random from some *known* distribution (instead of being worst-case), then we may assume W.L.O.G. that the optimal algorithm is deterministic. This is because we can choose the "best" (lest expected cost) deterministic algorithm in the support of the optimal randomized algorithm, whose cost can only be at most that of the randomized algorithm. Hence, to prove a lower bound of c on the cost of randomized algorithms, it suffices to design a distribution over inputs and argue that every deterministic algorithm incurs at least cost c for this distribution.

Proof. Let us first write the expectations as sums over all possible outcomes of X and A .

$$\mathbb{E}c(A, x) = \sum_{a \in \mathcal{A}} \Pr[A = a]c(a, x) \quad \text{and} \quad \mathbb{E}c(a, X) = \sum_{x \in \mathcal{X}} \Pr[X = x]c(a, x)$$

Now we use that the weighted average of a sequence is always upper-bounded by its maximum value. In our case, the weights are $\Pr[X = x]$. As $\sum_{x \in \mathcal{X}} \Pr[X = x] = 1$, we have

$$\max_{x \in \mathcal{X}} \mathbb{E}c(A, x) = \max_{x \in \mathcal{X}} \sum_{a \in \mathcal{A}} \Pr[A = a]c(a, x) \geq \sum_{x \in \mathcal{X}} \Pr[X = x] \sum_{a \in \mathcal{A}} \Pr[A = a]c(a, x) .$$

We can now reorder the sums and get

$$\sum_{x \in \mathcal{X}} \Pr[X = x] \sum_{a \in \mathcal{A}} \Pr[A = a]c(a, x) = \sum_{a \in \mathcal{A}} \Pr[A = a] \sum_{x \in \mathcal{X}} \Pr[X = x]c(a, x) .$$

Finally, we can use that $\sum_{a \in \mathcal{A}} \Pr[A = a] = 1$ the same way as above to obtain

$$\sum_{a \in \mathcal{A}} \Pr[A = a] \sum_{x \in \mathcal{X}} \Pr[X = x]c(a, x) \geq \min_{a \in \mathcal{A}} \sum_{x \in \mathcal{X}} \Pr[X = x]c(a, x) = \min_{a \in \mathcal{A}} \mathbb{E}c(a, X) . \quad \square$$

The analogous statement holds for maximization problems, where we have a profit $p(a, x)$ that algorithm a achieves on instance x . By setting $c(a, x) = -p(a, x)$, we get the following corollary.

Corollary 2. . *Let A be a random variable with values in \mathcal{A} and let X be a random variable with values in \mathcal{X} . Then,*

$$\min_{x \in \mathcal{X}} \mathbb{E}p(A, x) \leq \max_{a \in \mathcal{A}} \mathbb{E}p(a, X) .$$

2 Randomized Lower Bound for Online Set Cover

We will now use Yao's Lemma to show that for online Set Cover even randomized algorithms cannot be strictly $o(\log n)$ -competitive (or $o(\log m)$ -competitive), where n is the number of elements and m is the number of sets.

Theorem 3. *No algorithm for Online Set Cover is α -competitive for $\alpha < 1 + \frac{1}{2} \log_2 m$ (or $\alpha < \log_2(n + 1)$ -competitive).*

Proof. We consider m that are powers of two. Define the set system as follows. Take a complete binary tree of $n = 2m - 1$ vertices, height $h = \log_2 m + 1$, and m leaves.

Potential elements e to be covered correspond to nodes in the tree. Only some of these potential elements will have to be covered by our online algorithm, in particular elements along an (unknown) root-leaf path U^* will have to be covered.

The sets $S \in \mathcal{S}$ correspond to the m leaves of the tree, where S_k contains all elements corresponding to the root-leaf path for the k -th leaf. We set $c_S = 1$ for all $S \in \mathcal{S}$.

to the left or to the right subtree. If the element in step t comes from this subtree, the algorithm does not choose another set and does not incur any cost. If the element comes from a different subtree, then the algorithm chooses a new set and incurs cost of 1. Both events happen with probability $\frac{1}{2}$, so $\mathbb{E}C_t = \frac{1}{2}$.

Overall, we get for any $a \in \mathcal{A}$

$$\mathbb{E}c(a, X) = \sum_{t=1}^h \mathbb{E}C_t = 1 + (h-1)\frac{1}{2} = 1 + \frac{1}{2} \log_2 m . \quad \square$$

Remark: The above proof also implies the same hardness result for Online Fractional Set Cover since making a fractional decision at step t doesn't add power to the algorithm.

3 Randomized Lower Bound for Ski Rental

We now consider the Ski Rental Problem. We will show that no randomized algorithm can be better than $\left(1 - \left(1 - \frac{1}{B}\right)^B\right)^{-1}$ -competitive. Note that this matches the guarantee of the randomized algorithm that we derived from the fractional primal-dual one.

Theorem 4. *For a fixed B , no Ski Rental algorithm is α -competitive for $\alpha < \left(1 - \left(1 - \frac{1}{B}\right)^B\right)^{-1}$.*

Also note that $\lim_{B \rightarrow \infty} \left(1 - \left(1 - \frac{1}{B}\right)^B\right)^{-1} = \frac{e}{e-1} \approx 1.58$.

Proof. The instances have a very simple structure, they are simple non-negative integers representing the number of skiing days. For the offline optimum OPT, we now have

$$c(\text{OPT}, x) = \begin{cases} x & \text{if } x < B \\ B & \text{otherwise} \end{cases}$$

For any random variable X , this implies

$$\mathbb{E}c(\text{OPT}, X) = \sum_{t=1}^{B-1} t \Pr[X = t] + B \Pr[X \geq B] = \sum_{t=1}^B \Pr[X \geq t] .$$

Deterministic algorithms are also easy to describe. They only differ in how long they wait until buying the skis. So a is again a non-negative integer, which means that the algorithm rents skis for a days before buying them on day $a + 1$ (if this day exists). The cost of a is

$$c(a, x) = \begin{cases} x & \text{if } x \leq a \\ a + B & \text{otherwise} \end{cases}$$

So, the expected cost for a random X is

$$\mathbb{E}c(a, X) = \sum_{t=1}^a t \Pr[X = t] + (a + B) \Pr[X > a] = \sum_{t=1}^a \Pr[X \geq t] + B \Pr[X > a] .$$

Now, we have to find a distribution such that all deterministic algorithms a perform badly. The idea is to make all algorithms incur the same cost in expectation. This means that algorithms a and $a - 1$ have the same cost

$$\mathbb{E}c(a - 1, X) = \mathbb{E}c(a, X) ,$$

which means that

$$\sum_{t=1}^{a-1} \Pr[X \geq t] + B \Pr[X > a - 1] = \sum_{t=1}^{a-1} \Pr[X \geq t] + \Pr[X \geq a] + B \Pr[X > a] .$$

So we need

$$B \Pr[X \geq a] = \Pr[X \geq a] + B \Pr[X \geq a + 1] \quad \Leftrightarrow \quad \Pr[X \geq a + 1] = \left(1 - \frac{1}{B}\right) \Pr[X \geq a] .$$

This is fulfilled if we set

$$\Pr[X \geq t] = \left(1 - \frac{1}{B}\right)^{t-1} \quad \text{for all } t \geq 1 .$$

For this choice of X , we have

$$\mathbb{E}c(\text{OPT}, X) = \sum_{t=1}^B \left(1 - \frac{1}{B}\right)^{t-1} = B \left(1 - \left(1 - \frac{1}{B}\right)^B\right) ,$$

whereas for all a

$$\mathbb{E}c(a, X) = \sum_{t=1}^a \left(1 - \frac{1}{B}\right)^{t-1} + B \left(1 - \frac{1}{B}\right)^a = B .$$

Now suppose a randomized algorithm A is α -competitive for $\alpha < \left(1 - \left(1 - \frac{1}{B}\right)^B\right)^{-1}$. Then this means that

$$\mathbb{E}c(A, X) \leq \alpha \mathbb{E}c(\text{OPT}, X) .$$

As we know $\mathbb{E}c(A, X) = \sum_{a \in \mathcal{A}} \Pr[A = a] \mathbb{E}c(a, X) = B$ and $\mathbb{E}c(\text{OPT}, X) = B \cdot \left(1 - \left(1 - \frac{1}{B}\right)^B\right)$. This is a contradiction. \square

Notes

The lecture is partly based on Thomas Kesselheim's notes¹.

References

- [1] Andrew Chi-Chih Yao: Probabilistic Computations: Toward a Unified Measure of Complexity (Extended Abstract). FOCS 1977 (original paper introducing Yao's Minimax Lemma)

¹<http://tcs.cs.uni-bonn.de/lib/exe/fetch.php?media=teaching:ss20:vl-aau:lecturenotes05.pdf>