

Yao's Principle

Thomas Kesselheim

Last Update: October 24, 2025

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 Principle

Yao's principle 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 $\max_{x \in \mathcal{X}} \mathbf{E} [c(A, x)]$.

Theorem 6.1 (Yao's Principle). *Let A be a random variable with values in \mathcal{A} and let X be a random variable with values in \mathcal{X} . Then,*

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

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.

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

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

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

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

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

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

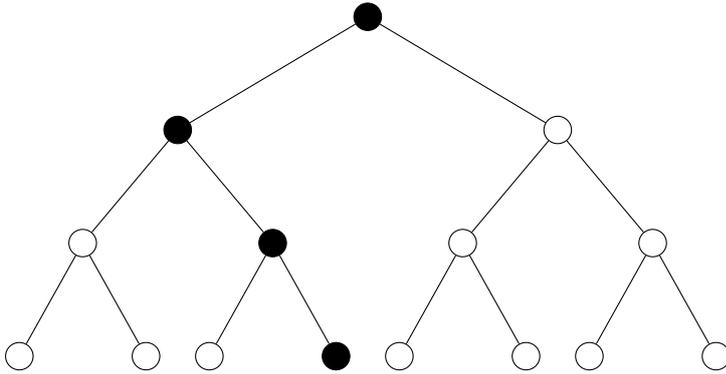


Figure 1: The binary tree used in the construction of the lower bound. The black nodes correspond to a potential sequence.

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

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

2 Randomized Lower Bound for Online Set Cover

We will now use Yao's principle to show that indeed for online Set Cover even randomized algorithms cannot be strictly $o(\log n)$ -competitive.

Theorem 6.3. No algorithm for Online Set Cover is strictly α -competitive for $\alpha < 1 + \frac{1}{2} \log_2 n$.

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

Potential elements e to be covered correspond to nodes in the tree. Only some of these potential elements will be present in our universe U of elements which have to be covered. The set U corresponds to the elements on one root-leaf path.

The sets $S \in \mathcal{S}$ correspond to leaves of the tree, where S contains all ancestors in the tree that are in U . We set $c_S = 1$ for all $S \in \mathcal{S}$.

Our sequence σ reveals the set U , starting from the root and going downward. Note that any such σ can be covered by a single set $S \in \mathcal{S}$, namely the one that corresponds to the leaf where the path ends. That is, $\text{cost}(\text{OPT}(\sigma)) = 1$.

However, the algorithm initially cannot distinguish the n different sets in \mathcal{S} because we only get to know what sets an element is contained in when it arrives. So, initially, we see only n placeholders in \mathcal{S} . Then the root arrives, which is only contained in half of the sets and so on. Only in the very last step, we will know which set would have covered all of U .

We now claim that for any randomized online algorithm $\mathbf{E} [\text{cost}(\text{ALG}(\sigma))] \geq 1 + \frac{1}{2} \log_2 n$ for some sequence σ of this form. In the notation of Yao's principle, this means that $\max_{x \in \mathcal{X}} \mathbf{E} [c(A, x)] \geq 1 + \frac{1}{2} \log_2 n$, where now A is an arbitrary randomized algorithm and \mathcal{X} is the set of sequences we are considering.

Yao's principle tells us that $\max_{x \in \mathcal{X}} \mathbf{E}[c(A, x)] \geq \min_{a \in \mathcal{A}} \mathbf{E}[c(a, X)]$ for any choice of the random variable X , so it is sufficient to show how to choose X such that $\min_{a \in \mathcal{A}} \mathbf{E}[c(a, X)] \geq 1 + \frac{1}{2} \log_2 n$.

In our case, we will draw one leaf uniformly at random and take the path ending at this leaf. Equivalently, we can determine this leaf as follows: Start at the root and flip a coin whether to go left or right. At every inner node continue this procedure. Note that this is exactly the way the online algorithm gets to the path. It always only gets to know whether the following elements are in the left or the right subtree. The elements to further distinguish the leaves have not yet arrived.

Without loss of generality, consider a *lazy* (deterministic) algorithm. This means that the algorithm only buys an $S \in \mathcal{S}$ if the current element is not yet covered. If so, it buys only a single one. In principle, an algorithm may buy multiple sets at a time but it is easy to see that this cannot be cheaper.

Let C_t be the cost of the algorithm in the t -th step. Clearly $C_1 = 1$. Now consider any later step $t > 1$. We will show that $\mathbf{E}[C_t] = \frac{1}{2}$. To this end, observe that the algorithm has chosen exactly one set that covers the element appearing in step $t - 1$. This can either belong 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 $\mathbf{E}[C_t] = \frac{1}{2}$.

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

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

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 6.4. *For a fixed B , no algorithm for Ski Rental 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

$$\begin{aligned} \mathbf{E}[c(\text{OPT}, X)] &= \sum_{t=1}^{B-1} t \Pr[X = t] + B \Pr[X \geq B] = \sum_{t=1}^{B-1} \sum_{i=1}^t \Pr[X = t] + B \Pr[X \geq B] \\ &= \sum_{i=1}^{B-1} \sum_{t=i}^{B-1} \Pr[X = t] + B \Pr[X \geq B] = \sum_{i=1}^{B-1} \Pr[i \leq X < B] + B \Pr[X \geq B] = \sum_{t=1}^B \Pr[X \geq t] . \end{aligned}$$

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

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

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

$$\mathbf{E}[c(a - 1, X)] = \mathbf{E}[c(a, X)] ,$$

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

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

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

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

whereas for all a

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

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

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

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

References

- Andrew Chi-Chih Yao: Probabilistic Computations: Toward a Unified Measure of Complexity (Extended Abstract). FOCS 1977 (original paper introducing Yao's principle)
- I currently do not have references for the two lower-bound proofs.