

The Wrong Direction of Jensen’s Inequality Is Algorithmically Right

Or Zamir ✉

Princeton University, NJ, USA

Abstract

Let \mathcal{A} be an algorithm with expected running time e^X , conditioned on the value of some random variable X . We construct an algorithm \mathcal{A}' with expected running time $O(e^{\mathbb{E}[X]})$, that fully executes \mathcal{A} . In particular, an algorithm whose running time is a random variable T can be converted to one with expected running time $O(e^{\mathbb{E}[\ln T]})$, which is never worse than $O(\mathbb{E}[T])$. No information about the distribution of X is required for the construction of \mathcal{A}' .

2012 ACM Subject Classification Theory of computation → Parameterized complexity and exact algorithms; Theory of computation → Algorithm design techniques; Theory of computation → Computational complexity and cryptography

Keywords and phrases algorithms, complexity, Jensen’s inequality

Digital Object Identifier 10.4230/LIPIcs.ICALP.2023.107

Category Track A: Algorithms, Complexity and Games

Acknowledgements The author would like to thank Avi Wigderson for pointing out important references.

1 Introduction

Let \mathcal{A} be a *Las Vegas*¹ randomized algorithm. Assume that conditioned on the value of some random variable X , the expected running time of \mathcal{A} is e^X . By Jensen’s inequality, $\mathbb{E}[e^X] \geq e^{\mathbb{E}[X]}$, and in fact \mathcal{A} ’s expected running time might be much larger than $e^{\mathbb{E}[X]}$: Consider for example X that gets the value $\frac{1}{p}\mathbb{E}[X]$ with probability p and 0 otherwise, for any choice of $p > 0$; While the expectation of X is always $\mathbb{E}[X]$, the expectation of e^X is $p \cdot e^{\frac{1}{p}\mathbb{E}[X]}$ which can be arbitrarily large. We show that, surprisingly, any such \mathcal{A} can be converted to a different Las-Vegas randomized algorithm \mathcal{A}' that gives the same answer yet runs in expected time $O(e^{\mathbb{E}[X]})$. Transforming \mathcal{A} to \mathcal{A}' does not require any assumption or knowledge about the distribution of X .

► **Theorem 1.** *There exists an algorithm T that receives as an input a randomized Las Vegas algorithm \mathcal{A} , and fully executes it. If the expected running time of \mathcal{A} is e^X when conditioned on the value of some random variable X , then the expected running time of $T(\mathcal{A})$ is $O(e^{\mathbb{E}[X]})$.*

As a corollary, any algorithm whose running time is a random variable T can be converted to one with expected running time $O(e^{\mathbb{E}[\ln T]})$, which is never worse than $O(\mathbb{E}[T])$.

Recently we used the following simple version of Theorem 1 in a late revision of [10] to substantially simplify the analysis in the paper. The paper improves the running time of exact exponential-time algorithms for general Constraint Satisfaction Problems.

¹ A randomized algorithm is called *Las Vegas* if it always returns the correct answer, but its running time is a random variable.



► **Lemma 2** (from an up-to-date version of [10]). *Let \mathcal{A} be an algorithm with expected running time 2^X conditioned on the value of a random variable X . There exists an algorithm \mathcal{A}' that fully executes \mathcal{A} and has an expected running time of $O(2^{E[X]} \cdot E[X])$. Transforming \mathcal{A} to \mathcal{A}' requires knowing $E[X]$.*

In this paper we focus on Theorem 1 itself, obtaining an optimal version of it.

We transform an algorithm \mathcal{A} by using a sequence of *truncated evaluations*. A *truncated evaluation* of an algorithm \mathcal{A} for t steps is the process of running algorithm \mathcal{A} and aborting its run if it did not fully execute in the first t computational steps of its run. Each of the algorithms we present is thus a sequence of values $t_1, t_2, \dots, t_i, \dots$ which we use as thresholds for truncated evaluations of \mathcal{A} . We stop at the first time \mathcal{A} is fully executed. These thresholds can be defined deterministically or be random variables. In the simpler algorithms we present, the thresholds depend on $E[X]$ or even on the entire distribution X . For the proof of Theorem 1 the thresholds are completely independent of X and \mathcal{A} .

Truncated evaluations are frequently used in complexity theory (for example, see the proof of the time and space hierarchies in [2]). The first algorithmic use of such a sequence of truncated evaluations that we are aware of, is by Alt, Guibas, Mehlhorn, Karp and Wigderson [1]. They used it to convert *Las Vegas* randomized algorithms to *Monte Carlo* randomized algorithms, with success probability larger than what Markov's inequality gives. Luby, Sinclair and Zuckerman [5] then introduced a universal strategy for truncated evaluations. That is a sequence that is guaranteed to run in time $O(s \log s)$ if there exists any sequence that runs in time $O(s)$ for the same algorithm. Our contribution thus is two-fold: first, we prove the existence of good strategies in terms of $E[X]$, and second, we show that these strategies can be explicitly constructed (i.e., without paying additional logarithmic factors). Not paying additional factors guarantees, due to Jensen's inequality, that our transformed algorithm is never worse than the original algorithm.

A natural use for such theorems is the regime of exponential-time algorithms. Consider the following toy example. As part of the classic algorithm of Schönig [9] it was shown that given an assignment α and a 3-SAT formula φ it is possible to test in $O(2^r)$ expected time whether φ has a satisfying assignment α_0 with $HAM(\alpha_0, \alpha) \leq r$, where $HAM(\cdot, \cdot)$ is the standard Hamming distance. This claim was also derandomized later [6]. We can use this primitive naively to obtain a non-trivial 3-SAT algorithm: Pick a random assignment α , and then run the above procedure of Schönig. Let X be the Hamming distance between α and a satisfying assignment α_0 of the input formula φ , this is a random variable. Conditioned on X , the expected running time of our algorithm is 2^X . The expected running time of our algorithm is thus $E[2^X] = \sum_{r=0}^n \binom{n}{r} 2^{-n} 2^r = 2^{-n} (1+2)^n = (\frac{3}{2})^n$. Using this paper's main Theorem, on the other hand, we can notice that $E[X] = \frac{n}{2}$ and thus we can convert the above algorithm in a black-box manner into one with expected running time $2^{E[X]} = (\sqrt{2})^n < (\frac{3}{2})^n$. We note that Schönig already presented an algorithm using this procedure that is faster than both of the algorithms above.

Another similar example is the famous PPSZ algorithm for solving k -SAT, including its recent improvements, and generalizations for CSPs [7, 4, 3, 10]. In these algorithms, a randomly chosen permutation determines the number of input variables that we need to *guess* the values of. The expectation of this number of variables is then analyzed. The success probability or running time is exponential in this number. In the original PPSZ algorithm the analyzed quantity is the success probability and thus Jensen's inequality is applicable to bound this probability from below. In other variations (including [10]), the analyzed quantity is the running time and then Jensen's inequality is no longer applicable and either a more complicated analysis or the statement of this paper is necessary. Further discussion on possible applications and in particular possible implications for SAT algorithms appears in Section 3.

1.1 Preliminaries

We use standard notation throughout the paper. The notation $\ln x$ is used for the natural logarithm, and $\log x$ is used for the base two logarithm.

► **Definition 3** (Iterated functions). *Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a function. We define the iterated functions $f^{(k)} : \mathbb{R} \rightarrow \mathbb{R}$ recursively as follows. $f^{(0)}(x) := x$, and for any $k > 0$ we let $f^{(k)}(x) := f(f^{(k-1)}(x))$.*

► **Definition 4** (Star functions). *Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a function. Assume f is strictly increasing and strictly shrinking² for all $x \geq x_0$. The star function of f , defined with respect to x_0 for every $x \geq x_0$, is*

$$f^*(x) = \min\{k \mid f^{(k)}(x) \leq x_0\}.$$

The (general) Tower function $Tower_b(n, x) : \mathbb{N} \times \mathbb{R} \rightarrow \mathbb{R}$ is defined as $f^{(n)}(x)$ where $f(x) = b^x$. The standard Tower function $Tower : \mathbb{N} \rightarrow \mathbb{N}$ is defined as $Tower(n) = Tower_2(n, 1)$. The discrete inverse of the Tower function is \log^* , defined with respect to $x_0 = 2$. That is, $\log^* n$ is the smallest integer such that $Tower(\log^* n) \geq n$.

2 Proof of Theorem 1

We begin by presenting a simple proof of Lemma 2.

Let \mathcal{A} be an algorithm whose expected running time is e^X when we condition on the value of some non-negative random variable X . We observe, by Markov's inequality, that

$$\Pr(X > \mathbb{E}[X] + 1) \leq \frac{\mathbb{E}[X]}{\mathbb{E}[X] + 1} = 1 - \frac{1}{\mathbb{E}[X] + 1}.$$

Hence, consider the following algorithm.

■ **Algorithm 1** Simple repetition.

Input: \mathcal{A} , $\mathbb{E}[X]$.

- 1: **repeat**
 - 2: Run \mathcal{A} for $2e^{\mathbb{E}[X]+1}$ computational steps.
 - 3: **until** \mathcal{A} completed a run.
-

► **Lemma 5.** *Algorithm 1 is expected to terminate in $O(e^{\mathbb{E}[X]} \cdot \mathbb{E}[X])$ time.*

Proof. If $X \leq \mathbb{E}[X] + 1$ then the expected running time of \mathcal{A} is at most $e^{\mathbb{E}[X]+1}$, and hence by Markov's inequality a truncated evaluation of \mathcal{A} for $2e^{\mathbb{E}[X]+1}$ steps concludes with probability at least $\frac{1}{2}$. By another application of Markov's inequality we got $\Pr(X \leq \mathbb{E}[X] + 1) > \frac{1}{\mathbb{E}[X]+1}$. Hence, the expected number of iterations until the truncated evaluations concludes is at most $2(\mathbb{E}[X] + 1)$. Each iteration takes $O(e^{\mathbb{E}[X]})$ time. ◀

² That is, $f(x) < x$.

2.1 Optimal bound when the distribution of X is known

The bound given by Markov's inequality in

$$\Pr(X \geq \mathbb{E}[X] + 1) \leq \frac{\mathbb{E}[X]}{\mathbb{E}[X] + 1} = 1 - \frac{1}{\mathbb{E}[X] + 1}$$

is attained *only* by the following distribution of X :

$$\Pr(X = k) := \begin{cases} \frac{1}{\mathbb{E}[X] + 1} & k = 0 \\ 1 - \frac{1}{\mathbb{E}[X] + 1} & k = \mathbb{E}[X] + 1 \end{cases}$$

In this distribution, on the other hand, the value of X is very small with a relatively high probability. In particular, in the case where $X < \mathbb{E}[X] + 1$ we need to run \mathcal{A} for much less than $e^{\mathbb{E}[X] + 1}$ computational steps. Hence, it is sensible to hope that every distribution X has some threshold other than $\mathbb{E}[X] + 1$ for which an algorithm similar to Algorithm 1 results in a better bound. Consider the following algorithm, which is a generalization of Algorithm 1 in which the threshold can be arbitrary.

■ **Algorithm 2** Simple repetition with variable threshold.

Input: \mathcal{A} , t .

- 1: **repeat**
 - 2: Run \mathcal{A} for $2e^t$ computational steps.
 - 3: **until** \mathcal{A} completed a run.
-

► **Lemma 6.** *Let X be a non-negative random variable. There exists $t \in [0, \mathbb{E}[X] + 1]$ such that $\frac{e^t}{\Pr(X < t)} \leq e^{\mathbb{E}[X] + 1}$.*

Proof. Assume by contradiction that $\frac{e^t}{\Pr(X < t)} > e^{\mathbb{E}[X] + 1}$ for every $t \in [0, \mathbb{E}[X] + 1]$. Equivalently,

$$\Pr(X \geq t) = 1 - \Pr(X < t) > 1 - e^{t - (\mathbb{E}[X] + 1)}.$$

Therefore,

$$\begin{aligned} \mathbb{E}[X] &= \int_0^\infty \Pr(X \geq t) dt \geq \int_0^{\mathbb{E}[X] + 1} \Pr(X \geq t) dt \\ &> \int_0^{\mathbb{E}[X] + 1} \left(1 - e^{t - (\mathbb{E}[X] + 1)}\right) dt \\ &= (\mathbb{E}[X] + 1) - \left(1 - e^{-(\mathbb{E}[X] + 1)}\right) = \mathbb{E}[X] + e^{-(\mathbb{E}[X] + 1)} \\ &> \mathbb{E}[X], \end{aligned}$$

which is a contradiction. ◀

Lemma 6 implies the following.

► **Corollary 7.** *For every distribution X there exists a value of $t = t(X)$ for which Algorithm 2 runs in $O(e^{\mathbb{E}[X]})$ time.*

We note that the additive constant $+1$ in the exponent in Lemma 6 is necessary. For a parameter E , consider the random variable X supported on $[0, E + 1 + \ln(1 + e^{-(E+1)})]$ and distributed with density $f(x) := e^{x-(E+1)}$; Its expectation is

$$\begin{aligned} \mathbb{E}[X] &= \int_0^{E+1+\ln(1+e^{-(E+1)})} x f(x) dx \\ &= \left((x-1) e^{x-(E+1)} \right) \Big|_0^{E+1+\ln(1+e^{-(E+1)})} \\ &= \left(E + \ln(1 + e^{-(E+1)}) \right) \cdot (1 + e^{-(E+1)}) + e^{-(E+1)} \\ &= E + O(e^{-(E+1)}) = E + o(1), \end{aligned}$$

where the $o(1)$ term is vanishing when $E \rightarrow \infty$. On the other hand, for any $t \geq 0$ we have

$$\frac{e^t}{\Pr(X < t)} = \frac{e^t}{\min(1, e^{t-(E+1)} - e^{-(E+1)})} > \frac{e^t}{e^{t-(E+1)}} = e^{E+1}.$$

2.2 Optimal algorithm when the distribution of X is unknown

If the only thing known about the distribution of X is its expectation $\mathbb{E}[X]$, then there is no fixed value of t for which Algorithm 2 is better than Algorithm 1. Fix a value of $\mathbb{E}[X]$ and a choice of t . If $t < \mathbb{E}[X]$ then with the constant distribution $X \equiv \mathbb{E}[X]$ Algorithm 2 never terminates. Otherwise, $t \geq \mathbb{E}[X]$ and we consider the following distribution X :

$$\Pr(X = k) := \begin{cases} 1 - \frac{\mathbb{E}[X]}{t+1} & k = 0 \\ \frac{\mathbb{E}[X]}{t+1} & k = t + 1 \end{cases}.$$

For this distribution, the expected running time of Algorithm 2 is

$$\frac{e^t}{1 - \frac{\mathbb{E}[X]}{t+1}} = e^{\mathbb{E}[X]} \cdot \frac{e^s}{\left(\frac{s+1}{\mathbb{E}[X]+s+1}\right)} = e^{\mathbb{E}[X]} \left(1 + \frac{\mathbb{E}[X]}{s+1}\right) e^s \geq e^{\mathbb{E}[X]} \mathbb{E}[X] \cdot \frac{e^s}{s+1} \geq e^{\mathbb{E}[X]} \mathbb{E}[X],$$

where $s := t - \mathbb{E}[X] \geq 0$ and the last inequality follows as $e^s \geq s + 1$ for any s .

To improve Algorithm 1 then, we need to consider *several thresholds*. We demonstrate this idea with the following Lemma.

► **Lemma 8.** *Let X be a non-negative random variable. It holds that either $\Pr(X \leq \mathbb{E}[X] - \ln \mathbb{E}[X]) > \frac{1}{\mathbb{E}[X]+1}$ or $\Pr(X \leq \mathbb{E}[X] + 2) > \frac{1}{\ln \mathbb{E}[X]+2}$.*

Proof. Assume that $p := \Pr(X \leq \mathbb{E}[X] - \ln \mathbb{E}[X]) \leq \frac{1}{\mathbb{E}[X]+1}$. We observe that

$$\begin{aligned} \mathbb{E}[X] &= p \mathbb{E}[X \mid X \leq \mathbb{E}[X] - \ln \mathbb{E}[X]] + (1-p) \mathbb{E}[X \mid X > \mathbb{E}[X] - \ln \mathbb{E}[X]] \\ &\geq (1-p) \mathbb{E}[X \mid X > \mathbb{E}[X] - \ln \mathbb{E}[X]], \end{aligned}$$

and hence

$$\begin{aligned} \mathbb{E}[X \mid X > \mathbb{E}[X] - \ln \mathbb{E}[X]] &\leq \frac{1}{1-p} \mathbb{E}[X] \\ &\leq \frac{1}{1 - \frac{1}{\mathbb{E}[X]+1}} \mathbb{E}[X] \\ &= \mathbb{E}[X] + 1. \end{aligned}$$

107:6 The Wrong Direction of Jensen's Inequality Is Algorithmically Right

Denote by $Y := X - (\mathbb{E}[X] - \ln \mathbb{E}[X])$. The above can now be rephrased as $\mathbb{E}[Y \mid Y > 0] \leq \ln \mathbb{E}[X] + 1$. Applying Markov's inequality to Y conditioned on $Y > 0$, we get

$$\Pr(Y > \ln \mathbb{E}[X] + 2 \mid Y > 0) \leq \frac{\ln \mathbb{E}[X] + 1}{\ln \mathbb{E}[X] + 2} = 1 - \frac{1}{\ln \mathbb{E}[X] + 2}.$$

We conclude by noting that

$$\Pr(X > \mathbb{E}[X] + 2) = \Pr(Y > \ln \mathbb{E}[X] + 2) \leq \Pr(Y > \ln \mathbb{E}[X] + 2 \mid Y > 0). \quad \blacktriangleleft$$

Consider the following Algorithm.

■ **Algorithm 3** Two thresholds algorithm.

Input: \mathcal{A} , $\mathbb{E}[X]$.

- 1: **repeat**
 - 2: **for** $\lceil \mathbb{E}[X] + 1 \rceil$ **times do**
 - 3: Run \mathcal{A} for $2e^{\mathbb{E}[X] - \ln \mathbb{E}[X]}$ computational steps.
 - 4: **for** $\lceil \ln \mathbb{E}[X] + 2 \rceil$ **times do**
 - 5: Run \mathcal{A} for $2e^{\mathbb{E}[X] + 2}$ computational steps.
 - 6: **until** \mathcal{A} completed a run.
-

Due to Lemma 8, each iteration of the outermost loop of Algorithm 3 succeeds to fully execute \mathcal{A} with probability larger than $1 - e^{-1}$. Thus, in expectation we run this loop for a constant number of iterations. The first **for** loop takes $O(\mathbb{E}[X] \cdot e^{\mathbb{E}[X] - \ln \mathbb{E}[X]}) = O(e^{\mathbb{E}[X]})$ expected time, and the second takes $O(e^{\mathbb{E}[X]} \ln \mathbb{E}[X])$. We conclude the following.

► **Corollary 9.** *Algorithm 3 runs in expected time $O(e^{\mathbb{E}[X]} \ln \mathbb{E}[X])$.*

Intuitively, the *proof* of Lemma 8 can be viewed as a reduction from the variable X to the variable $Y \mid (Y > 0)$, that has a much lower expectation: $\mathbb{E}[Y \mid Y > 0] \leq \ln \mathbb{E}[X] + 1$. We can thus hope that iterating the proof for $\ln^* \mathbb{E}[X]$ times would result in reducing X to a variable with constant expectation. A natural implementation of this idea would result in an algorithm that runs in expected time $O(e^{\mathbb{E}[X]} \ln^* \mathbb{E}[X])$. We next formalize this intuition, and do so in a more careful manner to avoid the $\ln^* \mathbb{E}[X]$ factor.

► **Definition 10.** *Let $\lambda(x) := 3 \ln(x)$ and note it is strictly increasing and shrinking for all $x \geq 5$. We define $\lambda^*(x)$, for $x \geq 5$, to be the smallest $k \in \mathbb{N}$ such that $\lambda^{(k)}(x) \leq 5$.*

▷ **Claim 11.** The following hold for all $x \geq 5$:

1. $\lambda^*(x) = \Theta(\log^* x)$.
2. $\lambda^{(\lambda^*(x))}(x) > 4$.
3. $\sum_{i=0}^{\lambda^*(x)} \frac{1}{\lambda^{(i)}(x)} < 2$.

Proof. (1) We have that $\lambda^{(2)}(x) \leq \log x \leq \lambda(x)$ for all $x \geq 410$. In particular, $\log^* x \leq \lambda^*(x) \leq 2 \log^* x + \lambda^*(410)$.

(2) $\lambda^{(\lambda^*(x)-1)}(x) > 5$ and hence $\lambda^{(\lambda^*(x))}(x) > \lambda(5) > 4.82$.

(3) For all $x \geq 17$ it holds that $\lambda(x) \leq \frac{x}{2}$. Let k' be the smallest integer such that $\lambda^{(k')}(x) < 17$. We thus have

$$\sum_{i=0}^{k'-1} \frac{1}{\lambda^{(i)}(x)} < \frac{1}{17} \sum_{i=0}^{\infty} 2^{-i} = \frac{2}{17}.$$

On the other hand, there are at most $\lambda^*(17)$ summands that are strictly larger than $\frac{1}{17}$, thus by (2) we have

$$\sum_{i=k'}^{\lambda^*(x)} \frac{1}{\lambda^{(i)}(x)} < \frac{\lambda^*(17)}{4} = \frac{5}{4}. \quad \triangleleft$$

We are now ready to prove a generalized version of Lemma 8, that is going to be the core of our final algorithm.

► **Lemma 12.** *Let X be a non-negative distribution and $E \geq \max(\mathbb{E}[X], 5)$ be an upper bound on its expectation. There either exists $1 \leq k \leq \lambda^*(E)$ such that $\Pr(X < E - \lambda^{(k)}(E)) \geq \left((\lambda^{(k-1)}(E) + 2)^2 + 1\right)^{-1}$, or it holds that $\Pr(X < E + 10) \geq \frac{1}{2}$.*

Proof. We recursively denote by $E_0 := E$ and by $E_k := \lambda^{(k)}(E) + \sum_{i=0}^{k-1} \frac{1}{E_i}$ for $1 \leq k \leq \lambda^*(E)$. Note that $E_k \geq \lambda^{(k)}(E)$ and hence also

$$E_k = \lambda^{(k)}(E) + \sum_{i=0}^{k-1} \frac{1}{E_i} \leq \lambda^{(k)}(E) + \sum_{i=0}^{k-1} \frac{1}{\lambda^{(i)}(E)} < \lambda^{(k)}(E) + 2,$$

where the last inequality follows from Claim 11. In particular, $\lambda^{(k)}(E) \leq E_k < \lambda^{(k)}(E) + 2$.

Assume that $\Pr(X < E - \lambda^{(k)}(E)) < \left((\lambda^{(k-1)}(E) + 2)^2 + 1\right)^{-1} < \frac{1}{(E_{k-1})^2 + 1}$ for every $1 \leq k \leq \lambda^*(E)$.

Denote by $Y_k := X - (E - \lambda^{(k)}(E))$ for $k \geq 0$. We prove by induction on k that $\mathbb{E}[Y_k | Y_k \geq 0] \leq E_k$. For $k = 0$ the claim is straightforward as $Y_0 = X$ and $E_0 = E$. For the inductive step, we assume the hypothesis holds for $k - 1$ and show it holds for k . We note that $Y_{k-1} > Y_k$ and hence if $Y_k \geq 0$ then $Y_{k-1} \geq 0$ as well. Hence,

$$\begin{aligned} \mathbb{E}[Y_{k-1} | Y_{k-1} \geq 0] &\geq \Pr(Y_k \geq 0 | Y_{k-1} \geq 0) \mathbb{E}[Y_{k-1} | Y_k \geq 0] \\ &\geq \Pr(Y_k \geq 0) \mathbb{E}[Y_{k-1} | Y_k \geq 0]. \end{aligned}$$

Thus, by the induction hypothesis we have

$$\begin{aligned} \mathbb{E}[Y_{k-1} | Y_k \geq 0] &\leq \frac{\mathbb{E}[Y_{k-1} | Y_{k-1} \geq 0]}{\Pr(Y_k \geq 0)} \\ &\leq \frac{E_{k-1}}{1 - \frac{1}{(E_{k-1})^2 + 1}} \\ &= E_{k-1} + \frac{1}{E_{k-1}}. \end{aligned}$$

Therefore,

$$\begin{aligned} \mathbb{E}[Y_k | Y_k \geq 0] &= \mathbb{E}[Y_{k-1} | Y_k \geq 0] + \lambda^{(k)}(E) - \lambda^{(k-1)}(E) \\ &\leq E_{k-1} + \frac{1}{E_{k-1}} + \lambda^{(k)}(E) - \lambda^{(k-1)}(E) \\ &= E_k. \end{aligned}$$

In particular, we have that $\mathbb{E}[Y_{\lambda^*(E)} | Y_{\lambda^*(E)} \geq 0] \leq E_{\lambda^*(E)} < \lambda^{(\lambda^*(E))}(E) + 2 \leq 7$.

107:8 The Wrong Direction of Jensen's Inequality Is Algorithmically Right

Therefore,

$$\begin{aligned} \Pr(X \geq E + 10) &\leq \Pr\left(X \geq E + 10 \mid X \geq E - \lambda^{(\lambda^*(E))}(E)\right) \\ &= \Pr\left(Y_{\lambda^*(E)} \geq \lambda^{(\lambda^*(E))}(E) + 10 \mid Y_{\lambda^*(E)} \geq 0\right) \\ &\leq \Pr\left(Y_{\lambda^*(E)} \geq 14 \mid Y_{\lambda^*(E)} \geq 0\right) < \frac{7}{14} = \frac{1}{2}. \end{aligned}$$

Consider the following algorithm.

■ **Algorithm 4** Multiple thresholds algorithm.

Input: \mathcal{A} , E .

```

1: repeat
2:   for  $k = 1$  to  $\lambda^*(E)$  do
3:     for  $2\lceil(\lambda^{(k-1)}(E) + 2)^2 + 1\rceil$  times do
4:       Run  $\mathcal{A}$  for  $2e^{E-\lambda^{(k)}(E)}$  computational steps.
5:   for 2 times do
6:     Run  $\mathcal{A}$  for  $2e^{E+10}$  computational steps.
7: until  $\mathcal{A}$  completed a run.

```

► **Corollary 13** (of Lemma 12). *Each **repeat** loop of Algorithm 4 fully executes \mathcal{A} with probability at least $\frac{3}{4}$.*

► **Lemma 14.** *Let $E \geq \max(\mathbb{E}[X], 5)$, Algorithm 4 runs in $O(e^E)$ expected time.*

Proof. By Corollary 13 we enter the **repeat** loop a constant number of times in expectation. We thus analyze the computational cost of a single **repeat** loop. The evaluations in Lines 5–6 take $O(e^E)$ time. The evaluations in Lines 2–4 take

$$\sum_{k=1}^{\lambda^*(E)} 2\lceil(\lambda^{(k-1)}(E) + 2)^2 + 1\rceil \cdot 2e^{E-\lambda^{(k)}(E)} = O\left(e^E \cdot \sum_{k=1}^{\lambda^*(E)} (\lambda^{(k-1)}(E))^2 e^{-\lambda^{(k)}(E)}\right)$$

time. By the definition of $\lambda(x)$, we have $e^{-\lambda^{(k)}(x)} = e^{-3\ln(\lambda^{(k-1)}(x))} = (\lambda^{(k-1)}(x))^{-3}$. In particular,

$$\sum_{k=1}^{\lambda^*(E)} (\lambda^{(k-1)}(E))^2 e^{-\lambda^{(k)}(E)} = \sum_{k=1}^{\lambda^*(E)} (\lambda^{(k-1)}(E))^{-1} < 2,$$

where the last inequality follows from Claim 11. ◀

Finally, we also get rid of the necessity to provide the algorithm with E or $\mathbb{E}[X]$.

► **Theorem 15.** *Algorithm 5 runs in expected time $O(e^{\mathbb{E}[X]})$.*

Proof. By Lemma 14 the iteration of the outermost **for** loop corresponding to E takes at most $C \cdot e^E$ time, for some global constant C . All iterations in which $E < \mathbb{E}[X]$ thus take $O(e^{\mathbb{E}[X]})$ time. By Corollary 13, each subsequent iteration succeeds with probability at least $\frac{3}{4}$. Thus the expected running time is bounded by

$$C e^{\mathbb{E}[X]} \cdot \sum_{t=0}^{\infty} e^t \left(\frac{1}{4}\right)^t = O\left(e^{\mathbb{E}[X]}\right). \quad \blacktriangleleft$$

■ **Algorithm 5** Final algorithm.

Input: \mathcal{A} .

- 1: **for** $E = 5$ to ∞ **do**
- 2: **for** $k = 1$ to $\lambda^*(E)$ **do**
- 3: **for** $2\lceil(\lambda^{(k-1)}(E) + 2)^2 + 1\rceil$ **times do**
- 4: Run \mathcal{A} for $2e^{E-\lambda^{(k)}(E)}$ computational steps.
- 5: **for 2 times do**
- 6: Run \mathcal{A} for $2e^{E+10}$ computational steps.
- 7: **return if** \mathcal{A} completed a run.

3 Conclusions and Open Problems

We showed that a Las-Vegas algorithm with expected running e^X conditioned on the value of some random variable X , can always be converted to a Las-Vegas algorithm with expected running time $O(e^{\mathbb{E}[X]})$. In particular, an algorithm whose running time is a random variable T can be converted to one with expected running time $O(e^{\mathbb{E}[\ln T]})$, which is never worse than $O(\mathbb{E}[T])$.

We demonstrated a use of this theorem to simplify a proof in the regime of exponential time algorithms [10]. It is interesting to try applying it to other exponential and non-exponential time algorithms and see if it can simplify or even improve the analysis.

3.1 Considering the variance

In terms of $\mathbb{E}[X]$ only, we can not get any better than $O(e^{\mathbb{E}[X]})$ as the distribution of X might be constant. In that case though, the variance of X is zero. Can we get a better bound just by assuming that the variance of X is large? Unfortunately, with the standard definition of variance this is not the case. For any choice of E and $V \geq 2E^2e^{-E}$ consider the following distribution:

$$\Pr(X = k) := \begin{cases} e^{-E} & k = 0 \\ 1 - \frac{Ve^{-E}}{V - E^2e^{-E}} & k = E \\ \frac{(Ee^{-E})^2}{V - E^2e^{-E}} & k = \frac{V}{Ee^{-E}} \end{cases}.$$

Its expectation is E , its variance is V , which can be arbitrarily large, and nevertheless $\Pr(X < E) = e^{-E}$ so no strategy can beat $O(e^E)$.

On the other hand, the wishful thinking above is true with some other notions of deviation. For example, if we consider mean absolute deviation instead of standard deviation (i.e., $\mathbb{E}[|X - \mathbb{E}[X]|]$), then *it is* true that if the deviation is large then we can get a better running time. It is intriguing to find useful notion of deviation for which such a statement is true, with the goal of improving the running time of algorithms by analyzing the deviation of X .

In particular, consider the PPSZ algorithm for solving k -SAT [7] [4]. The algorithm uses randomization in two ways: first, a random permutation of the variables in the input formulas is drawn; Then, the chosen permutation determines the number of variables we need to *guess* the value of. In a recent improvement of the PPSZ analysis, Scheder [8] showed that in some large subset of permutations the number of guessed variables is smaller than what we expect when taking a uniformly random permutation. In particular, this implies that there is some non-negligible variance in the original algorithm's running time. Can we get better SAT algorithms by analyzing this variance?

References

- 1 Helmut Alt, Leonidas Guibas, Kurt Mehlhorn, Richard Karp, and Avi Wigderson. A method for obtaining randomized algorithms with small tail probabilities. *Algorithmica*, 16(4):543–547, 1996.
- 2 Sanjeev Arora and Boaz Barak. *Computational complexity: a modern approach*. Cambridge University Press, 2009.
- 3 Thomas Dueholm Hansen, Haim Kaplan, Or Zamir, and Uri Zwick. Faster k-sat algorithms using biased-pps. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 578–589, 2019.
- 4 Timon Hertli. 3-sat faster and simpler—unique-sat bounds for ppsz hold in general. *SIAM Journal on Computing*, 43(2):718–729, 2014.
- 5 Michael Luby, Alistair Sinclair, and David Zuckerman. Optimal speedup of las vegas algorithms. *Information Processing Letters*, 47(4):173–180, 1993.
- 6 Robin A Moser and Dominik Scheder. A full derandomization of schöning's k-sat algorithm. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 245–252, 2011.
- 7 Ramamohan Paturi, Pavel Pudlák, Michael E Saks, and Francis Zane. An improved exponential-time algorithm for k-sat. *Journal of the ACM (JACM)*, 52(3):337–364, 2005.
- 8 Dominik Scheder. Ppsz is better than you think. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 205–216. IEEE, 2022.
- 9 T Schoning. A probabilistic algorithm for k-sat and constraint satisfaction problems. In *40th Annual Symposium on Foundations of Computer Science (Cat. No. 99CB37039)*, pages 410–414. IEEE, 1999.
- 10 Or Zamir. Faster algorithm for unique $(k, 2)$ -csp. *ESA*, 2022.