On the impossibility to beat random assignment in k-SAT

Bogdan Grechuk • 13 Feb 2026

This post continues my series on favorite theorems of the 21st century. For an overview of the categories and my previous selections, see this earlier post. In the category *Between the Centuries*—that is, theorems proved in the late 20th century but published in the early 21st century—my favorite result in *Algorithms and Complexity* is Johan Håstad's celebrated hardness-of-approximation theorem (Håstad 2001), which shows that for k-SAT one cannot beat the trivial random assignment algorithm (assuming $P \neq NP$).

A central question in the theory of algorithms is: which computational problems can be solved efficiently? Here "efficiently'' typically means in time polynomial in the size of the input. To prove that a problem *is* efficiently solvable, it suffices to exhibit a concrete algorithm and analyze its running time. In contrast, proving that a problem *cannot* be solved efficiently is notoriously difficult. Indeed, we currently lack unconditional techniques capable of ruling out all possible polynomial-time algorithms for a natural problem. There are simply too many conceivable algorithms to eliminate one by one.

A powerful workaround is the theory of reductions. If we can efficiently transform instances of one problem into instances of another, then an efficient algorithm for the latter would yield one for the former. Using such reductions, researchers discovered thousands of problems that are computationally equivalent in a very strong sense: if any one of them admits a polynomial-time algorithm, then they all do. These problems are called *NP-complete*.

The famous conjecture $P \neq NP$—for which a million-dollar Clay Millennium Prize is offered—predicts that no NP-complete problem can be solved in polynomial time. Thus, proving that a problem is NP-complete gives compelling evidence of its intractability. Even stronger is to show that a problem cannot be *approximately* solved beyond a certain threshold, unless $P = NP$. Results of this type belong to the rich and beautiful theory of hardness of approximation.

A classical NP-complete problem is k-*SAT*. Let $x_1, x_2, \dots, x_n$ be Boolean variables. A k-*clause* is an expression of the form

$$y_1 \vee y_2 \vee \cdots \vee y_k,$$

where each y_i is either x_j or $\neg x_j$ for some j. An instance of k-SAT consists of m such clauses, and the task is to decide whether there exists a truth assignment to the variables that satisfies all clauses. For every fixed $k \geq 3$, this problem is NP-complete.

If deciding full satisfiability is too hard, a natural relaxation is to maximize the number of satisfied clauses. Given m clauses, what fraction can we guarantee to satisfy efficiently?

A simple randomized algorithm assigns each variable independently to `true` or `false` with probability $1/2$. Each k-clause is violated only if all its k literals evaluate to false, an event of probability 2^{-k}. Thus, in expectation, a random assignment satisfies a fraction $1 - 2^{-k}$ of the clauses. By standard derandomization techniques, one can even achieve this guarantee deterministically in polynomial time.

The natural question is: can we do better?

Håstad's theorem gives a strikingly definitive answer.

Theorem 1 (Håstad, 2001) *Assume that $P \neq NP$. Then for every $\epsilon > 0$ and every fixed $k \geq 3$, there is no polynomial-time algorithm that, given an instance of k-SAT with m clauses, can distinguish between:*

- *the case where all m clauses are simultaneously satisfiable, and*
- *the case where at most $(1 - 2^{-k} + \epsilon)m$ clauses can be satisfied.*

In other words, unless $P = NP$, no polynomial-time algorithm can achieve an approximation ratio strictly better than $1 - 2^{-k}$. The trivial random assignment algorithm is, in a precise and optimal sense, best possible.

What makes this result particularly remarkable is its sharpness. Hardness-of-approximation results often leave a gap between what is algorithmically achievable and what is provably impossible. Håstad's theorem closes this gap exactly for k-SAT: the approximation threshold matches the performance of the simplest conceivable algorithm.

The proof introduced powerful new techniques in probabilistically checkable proofs (PCP) and Fourier analysis of Boolean functions. Beyond settling the approximability of k-SAT, Håstad's ideas reshaped the landscape of approximation complexity and influenced a long line of subsequent work establishing optimal hardness results for many other constraint satisfaction problems.

References

Håstad, Johan. 2001. “Some Optimal Inapproximability Results.” *J. ACM* 48 (4): 798–859. https://doi.org/10.1145/502090.502098.