

We have defined the class of **NP**-complete problems, which have the property that if there is a polynomial time algorithm for any one of these problems, there is a polynomial time algorithm for all of them. Unfortunately, nobody has found an algorithm for any **NP**-complete problem, and it is widely believed that it is impossible to do so.

This might seem like a big hint that we should just give up, and go back to solving problems like MAX-FLOW, where we can find a polynomial time solution. Unfortunately, **NP**-complete show up all the time in the real world, and people want solutions to these problems. What can we do?

What Can We Do?

Actually, there is a great deal we can do. Here are just a few possibilities:

- Restrict the inputs.

NP-completeness refers to the worst case inputs for a problem. Often, inputs are not as bad as those that arise in **NP**-completeness proofs. For example, although the general SAT problem is hard, we have seen that the cases of 2SAT and Horn formulae have simple polynomial time algorithms.

- Provide approximations.

NP-completeness results often arise because we want an exact answer. If we relax the problem so that we only have to return a good answer, then we might be able to develop a polynomial time algorithm. For example, we have seen that a greedy algorithm provides an approximate answer for the SET COVER problem.

- Develop heuristics.

Sometimes we might not be able to make absolute guarantees, but we can develop algorithms that seem to work well in practice, and have arguments suggesting why they should work well. For example, the simplex algorithm for linear programming is exponential in the worst case, but in practice it's generally the right tool for solving linear programming problems.

- Use randomness.

So far, all our algorithms have been *deterministic*; they always run the same way on the same input. Perhaps if we let our algorithm do some things randomly, we can avoid the **NP**-completeness problem?

Actually, the question of whether one can use randomness to solve an **NP**-complete problem is still open, though it appears unlikely. (As is, of course, the problem of whether one can solve an **NP**-complete problem in polynomial time!) However, randomness proves a useful tool when we try to come up with approximation algorithms and heuristics. Also, if one can assume the input comes from a suitable “random distribution”, then often one can develop an algorithm that works well on average. (Notice the distinction between *using randomness in an algorithm* and *assuming the input is random*.)

We now consider approximation algorithms and heuristics. The difference between approximation algorithms and heuristic algorithms is really pretty small; we generally call something an approximation algorithm when we can prove something about its performance, and we generally call something a heuristic when its performance is good but we can't prove why.

To begin, we will look at heuristic methods. The amount we can prove about these methods is (as yet) very limited. However, these techniques have had some success in practice, and there are arguments in favor of why they are reasonable thing to try for some problems.

Approximation Algorithm Basics

Approximation algorithms give guarantees. It is worth keeping in mind that sometimes approximation algorithms do not always perform as well as heuristic-based algorithms. Other times they provide insight into the problem, so they can help determine good heuristics.

Often when we talk about an approximation algorithm, we give an *approximation ratio*. The approximation ratio gives the ratio between our solution and the actual solution. The goal is to obtain an approximation ratio as close to 1 as possible. If the problem involves a minimization, the approximation ratio will be greater than 1; if it involves a maximization, the approximation ratio will be less than 1.

Set Cover Approximation

The inputs to the set cover problem are a finite set $X = \{x_1, \dots, x_n\}$, and a collection of subsets \mathcal{S} of X such that $\bigcup_{S \in \mathcal{S}} S = X$. The problem is to find the subcollection $\mathcal{T} \subseteq \mathcal{S}$ such that the sets of \mathcal{T} cover X , that is

$$\bigcup_{T \in \mathcal{T}} T = X.$$

Notice that such a cover exists, since \mathcal{S} is itself a cover.

A greedy heuristic would be to build a cover by repeatedly including the set in \mathcal{S} that will cover the maximum

number of as yet uncovered elements. In this case, the greedy heuristic does not yield an optimal solution. Interestingly, however, we can prove that the greedy solution is a good solution, in the sense that it is not too far from the optimal. That is, it is an example of an *approximation algorithm*.

Claim 21.1 *Let k be the size of the smallest set cover for the instance (X, \mathcal{S}) . Then the greedy heuristic finds a set cover of size at most $k \ln n$.*

Proof: Let $Y_i \subseteq X$ be the set of elements that are still not covered after i sets have been chosen with the greedy heuristic. Clearly $Y_0 = X$. We claim that there must be a set $A \in \mathcal{S}$ such that $|A \cap Y_i| \geq |Y_i|/k$. To see this, consider the sets in the optimal set cover of X . These sets cover Y_i , and there are only k of them, so one of these sets must cover at least a $1/k$ fraction of Y_i . Hence

$$|Y_{i+1}| \leq |Y_i| - |Y_i|/k = (1 - 1/k)|Y_i|,$$

and by induction,

$$|Y_i| \leq (1 - 1/k)^i |Y_0| = n(1 - 1/k)^i < ne^{-i/k},$$

where the last inequality uses the fact that $1 + x \leq e^x$ with equality iff $x = 0$. Hence when $i \geq k \ln n$ we have $|Y_i| < 1$, meaning there are no uncovered elements, and hence the greedy algorithm finds a set cover of size at most $k \ln n$. ■

Exercise: Show that this bound is tight, up to constant factors. That is, give a family of examples where the set cover has size k and the greedy algorithm finds a cover of size $\Omega(k \ln n)$.

Linear programming approach: randomized rounding. Another approach to achieve an $O(\log n)$ approximation to set cover in polynomial time is the following. We set up a linear program

$$\begin{aligned} & \min \sum_{i=1}^m z_i \\ & \text{s.t.} \\ & \forall j \in [n] \quad \sum_{S_i \ni j} z_i \geq 1 \\ & \forall i \in [m] \quad 0 \leq z_i \leq 1 \end{aligned}$$

Here the sets are S_1, \dots, S_m , and they each subsets of $[n] = \{1, \dots, n\}$. The z_i should be 0 or 1 to represent whether set i is taken or not, but forcing integrality of the z_i would make this an *integer* linear program, and integer linear programming is NP-hard. Thus we relax to the constraint $0 \leq z_i \leq 1$.

Now, suppose we obtain an optimal solution z^* to the above linear program, which we can find in polynomial time. If $\text{OPT}(LP)$ denotes $\sum_{i=1}^m z_i^*$ and $\text{OPT}(SC)$ denotes the minimum size of a set cover, then we have $\text{OPT}(LP) \leq \text{OPT}(SC)$ since any set cover solution gives a feasible solution to the linear program (by “feasible solution”, we mean a z vector satisfying all the constraints). Thus our goal is to give a set cover solution which uses at most $O(\log n) \cdot \text{OPT}(LP)$ sets.

The randomized rounding approach is as follows. Define probabilities $p_i = \min\{1, (1 + \ln n)z_i^*\}$ for $i \in [m]$. We independently include each set S_i in our output with probability p_i . We claim that: (1) the expected number of sets we take is at most $(1 + \ln n)\text{OPT}(LP) \leq (1 + \ln n)\text{OPT}(SC)$, and (2) with probability at least $1 - 1/e$, our output collection covers all of $[n]$. Applying Markov’s inequality with (1) in mind, we have that the expected number of sets is $O(\log n)\text{OPT}(SC)$ with probability, say, 99/100. Thus by a union bound, the probability that we both cover all of $[n]$ using $O(\log n)\text{OPT}(SC)$ sets is at least $1 - 1/e - 1/100 > 1/2$. Thus if we repeat this rounding approach $\lceil \log_2(1/\delta) \rceil$ times independently and return the output collection with the fewest sets that covers all of $[n]$, this will be an $O(\log n)$ approximately optimal solution to set cover with probability $1 - \delta$.

For (1), let X_i for $i \in [m]$ be 1 if set i is included in our output collection, and 0 otherwise. Then the number of sets we output is $\sum_i X_i$. Note

$$\mathbb{E} \sum_{i=1}^m X_i = \sum_{i=1}^m \mathbb{E} X_i = \sum_{i=1}^m p_i \leq (1 + \ln n)\text{OPT}(LP).$$

For (2), we want to upper bound the quantity

$$\mathbb{P}(\exists j \in [n] : j \text{ not covered by output set}),$$

which by a union bound is at most

$$\sum_{j=1}^n \mathbb{P}(j \text{ not covered by output set}).$$

We focus on one of these summands. Fix some $j \in [n]$. If any of the p_i for $i \in S_j$ has $p_i = 1$, then j will definitely be covered. Otherwise, the probability that j is *not* covered is at most

$$\begin{aligned} \prod_{S_i \ni j} (1 - p_i) &\leq \prod_{S_i \ni j} e^{-p_i} \quad (\forall x \in \mathbb{R}, 1 + x \leq e^x) \\ &= e^{-\sum_{S_i \ni j} p_i} \\ &= e^{-(1 + \ln n) \sum_{S_i \ni j} z_i} \\ &\leq e^{-(1 + \ln n)} \quad (\text{LP constraint}) \\ &\leq 1/(en), \end{aligned}$$

and thus (2) is established.

Vertex Cover Approximations

In the Vertex Cover problem, we wish to find a set of vertices of minimal size such that every edge is adjacent to some vertex in the cover. That is, given an undirected graph $G = (V, E)$, we wish to find $U \subseteq V$ such that every edge $e \in E$ has an endpoint in U . We have seen that Vertex Cover is **NP**-complete.

A natural greedy algorithm for Vertex Cover is to repeatedly choose a vertex with the highest degree, and put it into the cover. When we put the vertex in the cover, we remove the vertex and all its adjacent edges from the graph, and continue. Unfortunately, in this case the greedy algorithm gives us a rather poor approximation, as can be seen with the following example:

We create a bipartite graph with vertex set $V = L \cup R$. We have $|L| = r$. Furthermore $R = R_2 \cup R_3 \cup \dots \cup R_r$. We have $|R_i| = \lfloor r/i \rfloor$, each vertex in R_i has degree i pointing to vertices in L , and the neighbors of any pair of vertices in R_i for fixed i are distinct. Thus, for example, R_r contains a single vertex with an edge to every vertex in L . R_1 contains r vertices each with an edge to a distinct vertex in L , etc.

Now, one vertex cover solution is to simply output L , which has size r . Meanwhile, greedy would first take the vertex in R_r (since it has degree r , whereas all vertices in L have degree $r-1$: each vertex in L has at most one edge to a vertex in each R_i , and there are only $r-1$ R_i 's since $2 \leq i \leq r$). Once the vertex in R_r is removed, the greedy algorithm will take all vertices in R_{r-1} , then all those in R_{r-2} , etc. Thus greedy takes a number of vertices which is

$$\sum_{i=2}^r |R_i| = \sum_{i=2}^r \lfloor r/i \rfloor \geq \sum_{i=2}^r [(r/i) - 1] = \Omega(r \log r) - r = \Omega(r \log r).$$

Since $|L| + |R| = n$, we have that $r \log r = \Theta(n)$, and $r = \Theta(n/\log n)$. Thus greedy outputs a vertex cover whose size is a $\Theta(\log n)$ factor larger than the size of the optimal vertex cover.

A better algorithm for vertex cover is the following: repeatedly choose an edge, and throw *both* of its endpoints into the cover. Throw the vertices and its adjacent edges out of the graph, and continue.

It is easy to show that this second algorithm uses at most twice as many vertices as the optimal vertex cover. This is because each edge that gets chosen during the course of the algorithm must have one of its endpoints in the cover; hence we have merely always thrown two vertices in where we might have gotten away with throwing in 1.

Somewhat surprisingly, this simple algorithm is still the best known approximation algorithm for the vertex cover problem. That is, no algorithm has been proven to do better than within a factor of 2.

Linear programming approach: deterministic rounding. Here we set up the linear program

$$\begin{aligned} & \min \sum_{i=1}^n z_i \\ & \text{s.t.} \\ & \forall e = (u, v) \in E \quad z_u + z_v \geq 1 \\ & \forall i \in [n] \quad 0 \leq z_i \leq 1 \end{aligned}$$

where the vertices are $[n]$. We solve the linear program to obtain a vector z^* , which ideally should tell us whether i is in the vertex cover or not (z_i^* would be 1 if i is in the vertex cover and 0 otherwise). However, since this is a linear program and not an integer linear program, each z_i^* may be some arbitrary fraction in $[0, 1]$.

The key insight is that since each edge $e = (u, v)$ enforces $z_u^* + z_v^* \geq 1$, it must be that at least one of z_u^*, z_v^* is at least $1/2$. Thus we take as our vertex cover: $\{v : z_v^* \geq 1/2\}$. This at most doubles some z_v^* (from $1/2$ to 1), and thus we at most double the value of $\sum_i z_i^*$. Thus, this solution provides a 2-approximation to minimum vertex cover.

Maximum Cut Approximation

We will provide both a randomized and a deterministic approximation algorithm for the MAX CUT problem. The MAX CUT problem is to divide the vertices in a graph into two disjoint sets so that the numbers of edges between vertices in different sets is maximized. This problem is **NP-hard**. Notice that the MIN CUT problem can be solved in polynomial time by repeated using the min cut-max flow algorithm. (Exercise: Prove this!)

The randomized version of the algorithm is as follows: we divide the vertices into two sets, HEADS and TAILS. We decide where each vertex goes by flipping a (fair) coin.

What is the probability an edge crosses between the sets of the cut? This will happen only if its two endpoints lie on different sides, which happens $1/2$ of the time. (There are 4 possibilities for the two endpoints – HH, HT, TT, TH – and two of these put the vertices on different sides.) So, on average, we expect $1/2$ the edges in the graph to cross the cut. Since the most we could have is for all the edges to cross the cut, this random assignment will, on average, be within a factor of 2 of optimal.

A local search approach. We now examine a deterministic algorithm with the same “approximation ratio”. (In fact, the two algorithms are intrinsically related– but this is not so easy to see!) The algorithm implements the hill climbing approximation heuristic, which is an example of what is called *local search*: we greedily do improving steps until no more improvement can be made.. We will split the vertices into sets S_1 and S_2 . Start with all vertices

on one side of the cut. Now, if you can switch a vertex to a different side so that it increases the number of edges across the cut, do so. Repeat this action until the cut can no longer be improved by this simple switch.

We switch vertices at most $|E|$ times (since each time, the number of edges across the cut increases). Moreover, when the process finishes we are within a factor of 2 of the optimal, as we shall now show. In fact, when the process finishes, at least $|E|/2$ edges lie in the cut.

We can count the edges in the cut in the following way: consider any vertex $v \in S_1$. For every vertex w in S_2 that it is connected to by an edge, we add $1/2$ to a running sum. We do the same for each vertex in S_2 . Note that each edge crossing the cut contributes $1/2$ for each vertex of the edge.

Hence the cut C satisfies

$$C = \frac{1}{2} \left(\sum_{v \in S_1} |\{w : (v, w) \in E, w \in S_2\}| + \sum_{v \in S_2} |\{w : (v, w) \in E, w \in S_1\}| \right).$$

Since we are using the local search algorithm, at least half the edges from any vertex v must lie in the set opposite from v ; otherwise, we could switch what side vertex v is on, and improve the cut! Hence, if vertex v has degree $\delta(v)$, then

$$\begin{aligned} C &= \frac{1}{2} \left(\sum_{v \in S_1} |\{w : (v, w) \in E, w \in S_2\}| + \sum_{v \in S_2} |\{w : (v, w) \in E, w \in S_1\}| \right) \\ &\geq \frac{1}{2} \left(\sum_{v \in S_1} \frac{\delta(v)}{2} + \sum_{v \in S_2} \frac{\delta(v)}{2} \right) \\ &= \frac{1}{4} \sum_{v \in V} \delta(v) \\ &= \frac{1}{2} |E|, \end{aligned}$$

where the last equality follows from the fact that if we sum the degree of all vertices, we obtain twice the number of edges, since we have counted each edge twice.

In practice, we might expect that the local search approach may often do better than just getting a cut within a factor of 2, so it is also good to do as a heuristic, even aside from provable guarantees.

Euclidean Travelling Salesperson Problem Approximation

In the Euclidean Travelling Salesman Problem, we are given n points (cities) in the $x - y$ plane, and we seek the tour (cycle) of minimum length that travels through all the cities. This problem is NP-complete (showing this is somewhat difficult).

Our approximation algorithm involves the following steps:

1. Find a minimum spanning tree T for the points.
2. Create a *pseudo tour* by walking around the tree. The pseudo tour may visit some vertices twice.
3. Remove repeats from the tour by *short-cutting* through the repeated vertices. (See Figure 21.1.)




Figure 21.1: Building an approximate tour. Start at X , move in the direction shown, short-cutting repeated vertices.

We now show the following inequalities:

$$\begin{aligned}
 \text{length of tour} &\leq \text{length of pseudo tour} \\
 &\leq 2(\text{size of } T) \\
 &\leq 2(\text{length of optimal tour})
 \end{aligned}$$

Short-cutting edges can only decrease the length of the tour, so the tour given by the algorithm is at most the length of the pseudo tour. The length of our pseudo tour is at most twice the size of the spanning tree, since this pseudo tour consists of walking through each edge of the tree at most twice. Finally, the length of the optimal tour is at least the size of the minimum spanning tree, since any tour contains a spanning tree (plus an edge!).

Using a similar idea, one can come up with an approximation algorithm that returns a tour that is within a factor of $3/2$ of the optimal. Also, note that this algorithm will work in any setting where short-cutting is effective. More specifically, it will work for any instance of the travelling salesperson problem that satisfies the *triangle inequality* for distances: that is, if $d(x,y)$ represents the distance between vertices x and y , and $d(x,z) \leq d(x,y) + d(y,z)$ for all x,y and z .

Linear Programming Relaxation

As we saw with set cover and vertex cover, many **NP**-complete problems can be easily described by a natural Integer Programming problem. (Of course, all **NP**-complete problems can be transformed into some Integer Programming problem, since Integer Programming is **NP**-complete; but what we mean here is in many cases the transformation is quite natural.) Even though we cannot solve the related Integer Program, *if we pretend it is a linear program*, then we can solve it, using (for example) the simplex method. This idea is known as *relaxation*, since we are relaxing the constraints on the solution; we are no longer requiring that we get a solution where the variables take on integer values.

If we are extremely lucky, we might find a solution of the linear program where all the variables are integers, in which case we will have solved our original problem. Usually, we will not. In this case we will have to try to somehow take the linear programming solution, and modify it into a solution where all the variables take on integer values. As we saw with set cover, *randomized rounding* is one technique for doing this.

MAX-SAT

We may formulate MAX-SAT as an integer programming problem in a straightforward way (in fact, we have seen a similar reduction before, back when we examined reductions; it is repeated here). Suppose the formula contains variables x_1, x_2, \dots, x_n which must be set to TRUE or FALSE, and clauses C_1, C_2, \dots, C_m . For each variable x_i we associate a variable y_i which should be 1 if the variable is TRUE, and 0 if it is FALSE. For each clause C_j we have a variable z_j which should be 1 if the clause is satisfied and 0 otherwise.

We wish to maximize the number of satisfied clauses s , or

$$\sum_{j=1}^m z_j.$$

The constraints include that that $0 \leq y_i, z_j \leq 1$; since this is an integer program, this forces all these variables to be either 0 or 1. Finally, we need a constraint for each clause saying that its associated variable z_j can be 1 if and

only if the clause is actually satisfied. If the clause C_j is $(x_2 \vee \bar{x}_4 \vee x_6 \vee \bar{x}_8)$, for example, then we need the restriction:

$$y_2 + y_6 + (1 - y_4) + (1 - y_8) \geq z_j.$$

This forces z_j to be 0 unless the clause can be satisfied. In general, we replace x_i by y_i , \bar{x}_i by $1 - y_i$, \vee by $+$, and set the whole thing $\geq z_j$ to get the appropriate constraint.

When we solve the linear program, we will get a solution that might have $y_1 = 0.7$ and $z_1 = 0.6$, for instance. This initially appears to make no sense, since a variable cannot be 0.7 TRUE. But we can still use these values in a reasonable way. If $y_1 = 0.7$, it suggests that we would prefer to set the variable x_1 to TRUE (1). In fact, we could try just rounding each variable up or down to 0 or 1, and use that as a solution! This would be one way to turn the non-integer solution into an integer solution. Unfortunately, there are problems with this method. For example, suppose we have the clause $C1 = (x_1 \vee x_2 \vee x_3)$, and $y_1 = y_2 = y_3 = 0.4$. Then by simple rounding, this clause will not be TRUE, even though it “seems satisfied” to our linear program (that is, $z_1 = 1$). If we have a lot of these clauses, regular rounding might perform very poorly.

It turns out that there is an interpretation for 0.7 that suggests a better way than simple rounding. We think of the 0.7 as a *probability*. That is, we interpret $y_1 = 0.7$ as meaning that x_1 would like to be true with probability 0.7. So we take each variable x_i , and independently we set it to 1 with the probability given by y_i (and with probability $1 - y_i$ we set x_i to 0). This process is known as *randomized rounding*. One reason randomized rounding is useful is it allows us *to prove* that the expected number of clauses we satisfy using this rounding is within a constant factor of the true optimum.

First, note that whatever the maximum number of clauses s we can satisfy is, the value found by the linear program, or $\sum_{j=1}^m z_j$, is at least as big as s . This is because the linear program could achieve a value of at least s simply by using as the values for y_i the truth assignment that makes satisfying s clauses possible.

Now consider a clause with k variables; for convenience, suppose the clause is just $C_1 = (x_1 \vee x_2 \dots \vee x_k)$. Suppose that when we solve the linear program, we find $z_1 = \beta$. Then we claim that the probability that this clause is satisfied after the rounding is at least $(1 - 1/e)\beta$. This can be checked (using a bit of sophisticated math), but it follows by noting (with experiments) that the worst possibility is that $y_1 = y_2 \dots = y_k = \beta/k$. In this case, each x_i is FALSE with probability $(1 - \beta/k)$, and so C_1 ends up being unsatisfied with probability $(1 - \beta/k)^k$. Hence the probability it is satisfied is at least (again using some math) $1 - (1 - \beta/k)^k \geq (1 - 1/e)\beta$.

Hence the i th clause is satisfied with probability at least $(1 - 1/e)z_i$, so the expected number of satisfied clauses after randomized rounding is at least $(1 - 1/e)\sum_{j=1}^m z_j$. This is within a factor of $(1 - 1/e)$ of our upper bound on the maximum number of satisfiable clauses, $\sum_{j=1}^m z_j$. Hence we expected to get within a constant factor of the maximum.

Surprisingly, by combining the simple coin flipping algorithm with the randomized rounding algorithm, we can get an even better algorithm. The idea is that the coin flipping algorithm does best on long clauses, since each literal in the clause makes it more likely the clause gets set to TRUE. On the other hand, randomized rounding does best on short clauses; the probability the clause is satisfied $(1 - (1 - \beta/k)^k)$ decreases with k . It turns out that if we try both algorithms, and take the better result, on average we will satisfy 3/4 of the clauses.

We also point out that there are even more sophisticated approximation algorithms for MAX-SAT, with better approximation ratios.

Hardness of Approximation

As we've seen, many optimization problems that are NP-hard to solve can still be well-approximated in polynomial time. However, it is natural to ask whether the approximation ratios achieved are the best possible (such as $\ln n$ for SET COVER and 2 for VERTEX COVER) or one can have arbitrarily good approximations in polynomial time (like we have for EUCLIDEAN TSP). For the first 20 years of NP-completeness, there was almost no understanding of the hardness of approximation. The breakthrough came from a very unexpected direction — the study of *probabilistic proof systems*, where we incorporate randomness (and sometimes interaction) into the classical notion of an NP proof:

Theorem 21.2 (PCP Theorem) *Every language in NP has a randomized polynomial-time verifier V , such that:*

1. *“Extreme” Efficiency: on input $\langle x, y \rangle$, V reads x , tosses $O(\log n)$ coins, reads $O(1)$ bits of y , and then decides whether to accept or reject.*
2. *Completeness: if $x \in L$, then there exists $y \in \{0, 1\}^{\text{poly}(|x|)}$ such that $\mathbb{P}[V(\langle x, y \rangle) = 1] = 1$.*
3. *Soundness: if $x \notin L$, then for all $y \in \{0, 1\}^{\text{poly}(|x|)}$, we have $\mathbb{P}[V(\langle x, y \rangle) = 1] \leq 1/2$.*

What is remarkable is that arbitrary NP statements can be verified by probabilistically checking only a *constant* number of bits of a witness y ! Since SHORT PROOF is NP-complete, this means that any mathematical proof (in a standard logical system like ZFC Set Theory) can be transformed into one that has this probabilistic verification property. Like with RP, the error probability of 1/2 can be made an arbitrarily small constant by repeating the verifier V several times.

The study of PCPs and related concepts (“interactive proofs” and “zero-knowledge proofs”) emerged from work on cryptography, where the goal was to enable distrustful parties to convince each other of various facts (e.g. I

have an encrypted and digitally signed bank statement showing that I have at least \$10,000 in my account), without revealing their secrets (e.g. a decryption key). But they turned out to have great significance for computational complexity, including unlocking the mysteries of approximation. Indeed, the following turns out to be *equivalent* to the PCP Theorem.

Theorem 21.3 (PCP Theorem, restated as a Gap-Producing Reduction) *There is a constant $\epsilon > 0$ and a polynomial-time reduction $\phi \mapsto \phi'$ from 3-SAT to 3-SAT such that:*

1. *If ϕ is satisfiable, then so is ϕ' .*
2. *If ϕ is unsatisfiable, then no assignment satisfies even a $(1 - \epsilon)$ fraction of the clauses of ϕ' .*

Note that such a gap-producing reduction implies that there is no $(1 - \epsilon)$ -approximation algorithm for 3-SAT. The equivalence of the two versions of the PCP Theorem is not too hard to establish (can you see why Gap-Producing Reductions imply the first statement of the PCP Theorem?), but proofs of the PCP Theorem itself are quite involved and use many deep and beautiful ideas (encoding computations using algebra/polynomials and/or gap-amplification using random walks on expander graphs). If you're interested in this area, see a graduate-level textbook on computational complexity (like the one by Goldreich or the one by Arora and Barak), or take a graduate course on complexity (like CS 221). While now many tight inapproximability results are known (like $\ln n$ for SET COVER and $7/8$ for MAX 3-SAT), there are still major open problems like the Unique Games Conjecture.