

Learning Objectives:

- Explore the consequences of $P = NP$
- *Search-to-decision* reduction: transform algorithms that solve decision version to search version for NP -complete problems.
- Optimization and learning problems
- Quantifier elimination and solving polynomial hierarchy.
- What is the evidence for $P = NP$ vs $P \neq NP$?

17

What if P equals NP ?

"There should be no fear ... we will be protected by God.", President Donald J. Trump, inauguration speech, 2017

"No more half measures, Walter", Mike Ehrmantraut in "Breaking Bad", 2010.

"The evidence in favor of $[P \neq NP]$ and [its algebraic counterpart] is so overwhelming, and the consequences of their failure are so grotesque, that their status may perhaps be compared to that of physical laws rather than that of ordinary mathematical conjectures.", Volker Strassen, laudation for Leslie Valiant, 1986.

We have mentioned that the question of whether $P = NP$, which is equivalent to whether there is a polynomial-time algorithm for 3SAT, is the great open question of Computer Science. But why is it so important? In this lecture, we will try to figure out the implications of such an algorithm.

First, let us get one qualm out of the way. Sometimes people say, "What if $P = NP$ but the best algorithm for 3SAT takes n^{100} time?" Well, n^{100} is much larger than, say, $2^{\sqrt{n}}$ for any input shorter than 10^{60} bits, which is way, way larger than the world's total storage capacity (estimated at a "mere" 10^{21} bits or about 200 exabytes at the time of this writing). So another way to phrase this question is to say, "what if the complexity of 3SAT is exponential for all inputs that we will ever encounter, but then grows much smaller than that?" To me this sounds like the computer science equivalent of asking, "what if the

laws of physics change completely once they are out of the range of our telescopes?”. Sure, this is a valid possibility, but wondering about it does not sound like the most productive use of our time.

So, as the saying goes, we’ll keep an open mind, but not so open that our brains fall out, and assume from now on that:

- There is a mathematical god,
and
- She does not “pussyfoot around” or take “half measures”. If God decided to make 3SAT easy, then 3SAT will have a $10^6 \cdot n$ (or at worst $10^6 n^2$) -time algorithm (i.e., 3SAT will be in $\text{TIME}(cn)$ or $\text{TIME}(cn^2)$ for a not-too-large constant c). If she decided to make 3SAT hard, then for every $n \in \mathbb{N}$, 3SAT on n variables cannot be solved by a NAND program of fewer than $2^{10^{-6}n}$ lines (which, through the relations between $\text{SIZE}(T(n))$, $\text{TIME}_{++}(T(n))$ and $\text{TIME}(T(n))$ that we’ve seen in [Theorem 14.4](#) and ??, implies that $3\text{SAT} \notin \text{TIME}(2^{o(n)})$).

So far, most of our evidence points to the latter possibility of 3SAT being exponentially hard, but we have not ruled out the former possibility either. In this lecture we will explore some of its consequences.

17.1 Search-to-decision reduction

A priori, having a fast algorithm for 3SAT might not seem so impressive. Sure, it will allow us to decide the satisfiability of not just 3CNF formulas but also of quadratic equations, as well as find out whether there is a long path in a graph, and solve many other decision problems. But this is not typically what we want to do. It’s not enough to know *if* a formula is satisfiable— we want to discover the actual satisfying assignment. Similarly, it’s not enough to find out if a graph has a long path— we want to actually *find* the path.

It turns out that if we can solve these decision problems, we can solve the corresponding search problems as well:

Theorem 17.1 — Search vs Decision. Suppose that $\mathbf{P} = \mathbf{NP}$. Then for every polynomial-time algorithm V and $a, b \in \mathbb{N}$, there is a polynomial-time algorithm FIND_V such that for every $x \in \{0, 1\}^n$, if there exists $y \in \{0, 1\}^{an^b}$ satisfying $V(xy) = 1$, then $\text{FIND}_V(x)$ finds some string y' satisfying this condition.

P To understand what the statement of [Theorem 17.1](#) means, let us look at the special case of the *MAXCUT* problem. It is not hard to see that there is a polynomial-time algorithm *VERIFYCUT* such that $VERIFYCUT(G, k, S) = 1$ if and only if S is a subset of G 's vertices that cuts at least k edges. [Theorem 17.1](#) implies that if $P = NP$ then there is a polynomial-time algorithm *FINDCUT* that on input G, k outputs a set S such that $VERIFYCUT(G, k, S) = 1$ if such a set exists. This means that if $P = NP$, by trying all values of k we can find in polynomial time a maximum cut in any given graph. We can use a similar argument to show that if $P = NP$ then we can find a satisfying assignment for every satisfiable 3CNF formula, find the longest path in a graph, solve integer programming, and so and so forth.

Proof Idea: The idea behind the proof of [Theorem 17.1](#) is simple; let us demonstrate it for the special case of 3SAT. (In fact, this case is not so “special”—since 3SAT is NP-complete, we can reduce the task of solving the search problem for *MAXCUT* or any other problem in NP to the task of solving it for 3SAT.) Suppose that $P = NP$ and we are given a satisfiable 3CNF formula φ , and we now want to find a satisfying assignment y for φ . Define $3SAT_0(\varphi)$ to output 1 if there is a satisfying assignment y for φ such that its first bit is 0, and similarly define $3SAT_1(\varphi) = 1$ if there is a satisfying assignment y with $y_0 = 1$. The key observation is that both $3SAT_0$ and $3SAT_1$ are in NP, and so if $P = NP$ then we can compute them in polynomial time as well. Thus we can use this to find the first bit of the satisfying assignment. We can continue in this way to recover all the bits.

Proof of Theorem 17.1. If $P = NP$ then for every polynomial-time algorithm V and $a, b \in \mathbb{N}$, there is a polynomial-time algorithm $STARTSWITH_V$ that on input $x \in \{0, 1\}^*$ and $z \in \{0, 1\}^\ell$, outputs 1 if and only if there exists some $y \in \{0, 1\}^{an^b}$ such that the first ℓ bits of y are equal to z and $V(xy) = 1$. Indeed, we leave it as an exercise to verify that the $STARTSWITH_V$ function is in NP and hence can be solved in polynomial time if $P = NP$.

Now for any such polynomial-time V and $a, b \in \mathbb{N}$, we can implement $FIND_V(x)$ as follows:

1. For $\ell = 0, \dots, an^b - 1$ do the following:
 - a. Let $b_0 = STARTSWITH_V(xz_0 \cdots z_{\ell-1}0)$ and $b_1 = STARTSWITH_V(xz_0 \cdots z_{\ell-1}1)$

b. If $b_0 = 1$ then $z_\ell = 0$, otherwise $z_\ell = 1$. 2. Output z_0, \dots, z_{an^b-1} .

To analyze the *FIND* algorithm, note that it makes $2an^{b-1}$ invocations to *STARTSWITH_V* and hence if the latter is polynomial-time, then so is *FIND_V*. Now suppose that x is such that there exists *some* y satisfying $V(xy) = 1$. We claim that at every step $\ell = 0, \dots, an^b - 1$, we maintain the invariant that there exists $y \in \{0, 1\}^{an^b}$ whose first ℓ bits are z s.t. $V(xy) = 1$. Note that this claim implies the theorem, since in particular it means that for $\ell = an^b - 1$, z satisfies $V(xz) = 1$.

We prove the claim by induction. For $\ell = 0$, this holds vacuously. Now for every $\ell > 0$, if the call *STARTSWITH_V*($xz_0 \dots z_{\ell-1}0$) returns 1, then we are guaranteed the invariant by definition of *STARTSWITH_V*. Now under our inductive hypothesis, there is $y_\ell, \dots, y_{an^b-1}$ such that $P(xz_0, \dots, z_{\ell-1}y_\ell, \dots, y_{an^b-1}) = 1$. If the call to *STARTSWITH_V*($xz_0 \dots z_{\ell-1}0$) returns 0 then it must be the case that $y_\ell = 1$, and hence when we set $z_\ell = 1$ we maintain the invariant. ■

17.2 Optimization

Theorem 17.1 allows us to find solutions for **NP** problems if $\mathbf{P} = \mathbf{NP}$, but it is not immediately clear that we can find the *optimal* solution. For example, suppose that $\mathbf{P} = \mathbf{NP}$, and you are given a graph G . Can you find the *longest* simple path in G in polynomial time?

P This is actually an excellent question for you to attempt on your own. That is, assuming $\mathbf{P} = \mathbf{NP}$, give a polynomial-time algorithm that on input a graph G , outputs a maximally long simple path in the graph G .

It turns out the answer is *Yes*. The idea is simple: if $\mathbf{P} = \mathbf{NP}$ then we can find out in polynomial time if an n -vertex graph G contains a simple path of length n , and moreover, by **Theorem 17.1**, if G does contain such a path, then we can find it. (Can you see why?) If G does not contain a simple path of length n , then we will check if it contains a simple path of length $n - 1$, and continue in this way to find the largest k such that G contains a simple path of length k .

The above reasoning was not specifically tailored to finding paths in graphs. In fact, it can be vastly generalized to proving the following result:

Theorem 17.2 — Optimization from $P = NP$. Suppose that $P = NP$. Then for every polynomial-time computable function $f : \{0,1\}^* \rightarrow \{0,1\}^*$, we can compute in $\text{poly}(n)$ time $\max_{x \in \{0,1\}^n} f(x)$ (where we identify the output of $f(x)$ with a natural number via the binary representation), and moreover find some $x^* \in \{0,1\}^n$ that achieves this maximum.

Proof Idea: Since f is polynomial-time computable, if $x \in \{0,1\}^n$ then $f(x)$ has at most $\text{poly}(n)$ bits and so we can think of $f(x)$ as a number between 0 and N where $N \leq 2^{\text{poly}(n)}$. If $P = NP$ then we can easily obtain an algorithm for computing $\max_{x \in \{0,1\}^n} f(x)$ that runs in time $N \cdot \text{poly}(n)$ -time just as above, by using the fact that $P = NP$ to obtain a polynomial-time algorithm checking whether there exists an x with $f(x) \geq k$ for $k = N, N-1, N-2, \dots, 0$. But if N is exponentially large in n , a running of $N \cdot \text{poly}(n)$ might not be good enough. The crucial observation is that we can use *binary search*: rather than checking whether there is x with $f(x) \geq k$ for $k = N, N-1, N-2, \dots$, we first check for $k = \lfloor N/2 \rfloor$, then (based on the answer) for either $k = \lfloor 3N/4 \rfloor$ or $k = \lfloor N/4 \rfloor$ and so on and so forth.

Proof of Theorem 17.2. For every such f , we can define the following Boolean function: $F : \{0,1\}^* \rightarrow \{0,1\}$: $F(1^n, k) = 1$ iff there exists $x \in \{0,1\}^n$ s.t. $f(x) \geq k$. Since f is computable in polynomial time, F is in NP , and so, under our assumption that $P = NP$, F itself can be computed in polynomial time. Now, for every n , we can compute the largest k such that $F(1^n, k) = 1$ by a binary search. We maintain two numbers a, b such that we are guaranteed that $a \leq \max_{x \in \{0,1\}^n} f(x) < b$. Initially we set $a = 0$ and $b = 2^{T(n)}$ where $T(n)$ is the running time of f . At each point in time, we compute the midpoint $c = \lfloor (a+b)/2 \rfloor$ and let $y = F(1^n, c)$. If $y = 1$ then we set $a = c$ and leave b as it is. If $y = 0$ then we set $b = c$ and leave a as it is. Since $|b-a|$ shrinks by a factor of 2, within $\log_2 2^{T(n)} = T(n)$ steps, we will get to the point at which $b \leq a+1$, and then we can simply output a . We can also use Theorem 17.1 to obtain the actual x that achieves the maximum. ■

For example, if G is a *weighted* graph, and every edge of G is given a weight which is a number between 0 and 2^k , then Theorem 17.2 shows that we can find the maximum-weight simple path in G (i.e., simple path maximizing the sum of the weights of its edges) in time polynomial in the number of vertices and in k .

17.2.1 Example: Supervised learning

One classical optimization task is *supervised learning*. In supervised learning we are given a list of *examples* x_0, x_1, \dots, x_{m-1} (where we can think of each x_i as a string in $\{0, 1\}^n$ for some n) and the *labels* for them y_0, \dots, y_{m-1} (which we will think of simply bits, i.e., $y_i \in \{0, 1\}$). For example, we can think of the x_i 's as images of either dogs or cats, for which $y_i = 1$ in the former case and $y_i = 0$ in the latter case. Our goal is to come up with a *hypothesis* or *predictor* $h : \{0, 1\}^n \rightarrow \{0, 1\}$ such that if we are given a new example x that has an (unknown to us) label y , then with high probability h will *predict* the label. That is, with high probability it will hold that $h(x) = y$. The idea in supervised learning is to use the *Occam's Razor principle*: the simplest hypothesis that explains the data is likely to be correct. There are several ways to model this, but one popular approach is to pick some fairly simple function $H : \{0, 1\}^{k+n} \rightarrow \{0, 1\}$. We think of the first k inputs as the *parameters* and the last n inputs as the example data. (For example, we can think of the first k inputs of H as specifying the weights and connections for some neural network that will then be applied on the latter n inputs.) We can then phrase the supervised learning problem as finding, given a set of labeled examples $S = \{(x_0, y_0), \dots, (x_{m-1}, y_{m-1})\}$, the set of parameters $\theta_0, \dots, \theta_{k-1} \in \{0, 1\}$ that minimizes the number of errors made by the predictor $x \mapsto H(\theta, x)$.

In other words, we can define for every set S as above the function $F_S : \{0, 1\}^k \rightarrow [m]$ such that $F_S(\theta) = \sum_{(x,y) \in S} |H(\theta, x) - y|$. Now, finding the value θ that minimizes $F_S(\theta)$ is equivalent to solving the supervised learning problem with respect to H . For every polynomial-time computable $H : \{0, 1\}^{k+n} \rightarrow \{0, 1\}$, the task of minimizing $F_S(\theta)$ can be “massaged” to fit the form of [Theorem 17.2](#) and hence if $\mathbf{P} = \mathbf{NP}$, then we can solve the supervised learning problem in great generality. In fact, this observation extends to essentially any learning model, and allows for finding the optimal predictors given the minimum number of examples. (This is in contrast to many current learning algorithms, which often rely on having access to an extremely large number of examples— far beyond the minimum needed, and in particular far beyond the number of examples humans use for the same tasks.)

17.2.2 Example: Breaking cryptosystems

We will discuss *cryptography* later in this course, but it turns out that if $\mathbf{P} = \mathbf{NP}$ then almost every cryptosystem can be efficiently broken.

One approach is to treat finding an encryption key as an instance of a supervised learning problem. If there is an encryption scheme that maps a “plaintext” message p and a key θ to a “ciphertext” c , then given examples of ciphertext/plaintext pairs of the form $(c_0, p_0), \dots, (c_{m-1}, p_{m-1})$, our goal is to find the key θ such that $E(\theta, p_i) = c_i$ where E is the encryption algorithm. While you might think getting such “labeled examples” is unrealistic, it turns out (as many amateur homebrew crypto designers learn the hard way) that this is actually quite common in real-life scenarios, and that it is also possible to relax the assumption to having more minimal prior information about the plaintext (e.g., that it is English text). We defer a more formal treatment to our lecture on cryptography.

17.3 Finding mathematical proofs

In the context of Gödel’s Theorem, we discussed the notion of a *proof system* (see [Definition 12.7](#)). Generally speaking, a *proof system* can be thought of as an algorithm $V : \{0,1\}^* \rightarrow \{0,1\}$ (known as the *verifier*) such that given a *statement* $x \in \{0,1\}^*$ and a *candidate proof* $w \in \{0,1\}^*$, $V(x, w) = 1$ if and only if w encodes a valid proof for the statement x . Any type of proof system that is used in mathematics for geometry, number theory, analysis, etc., is an instance of this form. In fact, standard mathematical proof systems have an even simpler form where the proof w encodes a *sequence* of lines w^0, \dots, w^m (each of which is itself a binary string) such that each line w^i is either an *axiom* or follows from some prior lines through an application of some *inference rule*. For example, **Peano’s axioms** encode a set of axioms and rules for the natural numbers, and one can use them to formalize proofs in number theory. Also, there are some even stronger axiomatic systems, the most popular one being **Zermelo–Fraenkel with the Axiom of Choice** or ZFC for short. Thus, although mathematicians typically write their papers in natural language, proofs of number theorists can typically be translated to ZFC or similar systems, and so in particular the existence of an n -page proof for a statement x implies that there exists a string w of length *poly*(n) (in fact often $O(n)$ or $O(n^2)$) that encodes the proof in such a system. Moreover, because verifying a proof simply involves going over each line and checking that it does indeed follow from the prior lines, it is fairly easy to do that in $O(|w|)$ or $O(|w|^2)$ (where as usual $|w|$ denotes the length of the proof w). This means that for every reasonable proof system V , the following function $SHORTPROOF_V : \{0,1\}^* \rightarrow \{0,1\}$ is in **NP**, where for every input of the form $x1^m$, $SHORTPROOF_V(x, 1^m) = 1$ if and only if

there exists $w \in \{0,1\}^*$ with $|w| \leq m$ s.t. $V(xw) = 1$. That is, $SHORTPROOF_V(x, 1^m) = 1$ if there is a proof (in the system V) of length at most m bits that x is true. Thus, if $\mathbf{P} = \mathbf{NP}$, then despite Gödel's Incompleteness Theorems, we can still automate mathematics in the sense of finding proofs that are not too long for every statement that has one. (Frankly speaking, if the shortest proof for some statement requires a terabyte, then human mathematicians won't ever find this proof either.) For this reason, Gödel himself felt that the question of whether $SHORTPROOF_V$ has a polynomial time algorithm is of great interest. As he wrote in a letter to John von Neumann in 1956 (before the concept of \mathbf{NP} or even "polynomial time" was formally defined):

One can obviously easily construct a Turing machine, which for every formula F in first order predicate logic and every natural number n , allows one to decide if there is a proof of F of length n (length = number of symbols). Let $\psi(F, n)$ be the number of steps the machine requires for this and let $\varphi(n) = \max_F \psi(F, n)$. The question is how fast $\varphi(n)$ grows for an optimal machine. One can show that $\varphi \geq k \cdot n$ [for some constant $k > 0$]. If there really were a machine with $\varphi(n) \sim k \cdot n$ (or even $\sim k \cdot n^2$), this would have consequences of the greatest importance. Namely, it would obviously mean that in spite of the undecidability of the Entscheidungsproblem,¹ the mental work of a mathematician concerning Yes-or-No questions could be completely replaced by a machine. After all, one would simply have to choose the natural number n so large that when the machine does not deliver a result, it makes no sense to think more about the problem.

For many reasonable proof systems (including the one that Gödel referred to), $SHORTPROOF_V$ is in fact \mathbf{NP} -complete, and so Gödel can be thought of as the first person to formulate the \mathbf{P} vs \mathbf{NP} question. Unfortunately, the letter was only discovered in 1988.

2

¹ The undecidability of Entscheidungsproblem refers to the uncomputability of the function that maps a statement in first order logic to 1 if and only if that statement has a proof.

² TODO: Maybe add example on finding Nash equilibrium

17.4 Quantifier elimination

So, if $\mathbf{P} = \mathbf{NP}$ then we can solve all \mathbf{NP} search problems in polynomial time. But can we do more? Yes we can!

An NP decision problem can be thought of as the task of deciding the truth of a statement of the form

$$\exists_x P(x) \tag{17.1}$$

for some NAND program P . But we can think of more general statements such as

$$\exists_x \forall_y P(x, y) \tag{17.2}$$

or

$$\exists_x \forall_y \exists_z P(x, y, z) . \tag{17.3}$$

For example, given an n -input NAND program P , we might want to find the *smallest* NAND program P' that is computes the same function as P . The question of whether there is such a P' of size at most k can be phrased as

$$\exists P' \forall x |P'| \leq k \wedge P(x) = P'(x) . \tag{17.4}$$

It turns out that if $P = NP$ then we can solve these kinds of problems as well.³

Theorem 17.3 — Polynomial hierarchy collapse. If $P = NP$ then for every $a \in \mathbb{N}$ there is a polynomial-time algorithm that on input a NAND program P on an inputs, returns 1 if and only if

$$\exists_{x_1 \in \{0,1\}^n} \forall_{x_2 \in \{0,1\}^n} \cdots Q_{x_a \in \{0,1\}^n} P(x_1, \dots, x_a) \tag{17.5}$$

where Q is either \exists or \forall depending on whether a is odd or even, respectively.

³ Since NAND programs are equivalent to Boolean circuits, this is known as the **circuit minimization problem** and is widely studied in Engineering.

Proof. We prove the theorem by induction. We assume that there is a polynomial-time algorithm $SOLVE_{a-1}$ that can solve the problem Eq. (17.5) for $a - 1$ and use that to solve the problem for a . On input a NAND program P , we will create the NAND program S_P that on input $x_1 \in \{0,1\}^n$, outputs $1 - SOLVE_{a-1}(1 - P_{x_1})$ where P_{x_1} is a NAND program that on input $x_2, \dots, x_a \in \{0,1\}^n$ outputs $P(x_1, \dots, x_n)$. Now note that by the definition of $SOLVE$

$$\begin{aligned} \exists_{x_1 \in \{0,1\}^n} S_P(x_1) &= \\ \exists_{x_1} \overline{SOLVE_{a-1}(\overline{P_{x_1}})} &= \\ \exists_{x_1} \overline{\exists_{x_2} \cdots Q'_{x_a} \overline{P(x_1, \dots, x_a)}} &= \\ \exists_{x_1} \forall_{x_2} \cdots Q_{x_a} P(x_1, \dots, x_a) . \end{aligned} \tag{17.6}$$

Hence we see that if we can solve the satisfiability problem for S_P , then we can solve Eq. (17.5). ■

This algorithm can also solve the search problem as well: find the value x_1 that certifies the truth of Eq. (17.5). We note that while this algorithm is in polynomial time, the exponent of this polynomial blows up quite fast. If the original NANDSAT algorithm required $\Omega(n^2)$ time, solving a levels of quantifiers would require time $\Omega(n^{2^a})$.⁴

17.4.1 Approximating counting problems

Given a NAND program P , if $\mathbf{P} = \mathbf{NP}$ then we can find an input x (if one exists) such that $P(x) = 1$. But what if there is more than one x like that? Clearly we can't efficiently output all such x 's; there might be exponentially many. But we can get an arbitrarily good multiplicative approximation (i.e., a $1 \pm \epsilon$ factor for arbitrarily small $\epsilon > 0$) for the number of such x 's, as well as output a (nearly) uniform member of this set. We will defer the details to later in this course, when we learn about *randomized computation*.

17.5 What does all of this imply?

So, what will happen if we have a 10^6n algorithm for 3SAT? We have mentioned that NP-hard problems arise in many contexts, and indeed scientists, engineers, programmers and others routinely encounter such problems in their daily work. A better 3SAT algorithm will probably make their lives easier, but that is the wrong place to look for the most foundational consequences. Indeed, while the invention of electronic computers did of course make it easier to do calculations that people were already doing with mechanical devices and pen and paper, the main applications computers are used for today were not even imagined before their invention.

An exponentially faster algorithm for all NP problems would be no less radical an improvement (and indeed, in some sense would be more) than the computer itself, and it is as hard for us to imagine what it would imply as it was for Babbage to envision today's world. For starters, such an algorithm would completely change the way we program computers. Since we could automatically find the "best" (in any measure we chose) program that achieves a certain task, we would not need to define *how* to achieve a task, but only specify tests as to what would be a good solution, and could also ensure that a program satisfies an exponential number of tests without actually running them.

⁴ We do not know whether such loss is inherent. As far as we can tell, it's possible that the *quantified boolean formula* problem has a linear-time algorithm. We will, however, see later in this course that it satisfies a notion known as PSPACE-hardness that is even stronger than NP-hardness.

The possibility that $P = NP$ is often described as “automating creativity”. There is something to that analogy, as we often think of a creative solution as one that is hard to discover but that, once the “spark” hits, is easy to verify. But there is also an element of hubris to that statement, implying that the most impressive consequence of such an algorithmic breakthrough will be that computers would succeed in doing something that humans already do today. In fact, creativity already is to a large extent automated or minimized (e.g., just see how much popular media content is mass-produced), and as in most professions we should expect to see the need for humans diminish with time even if $P \neq NP$.

Nevertheless, artificial intelligence, like many other fields, will clearly be greatly impacted by an efficient 3SAT algorithm. For example, it is clearly much easier to find a better Chess-playing algorithm when, given any algorithm P , you can find the smallest algorithm P' that plays Chess better than P . Moreover, as we mentioned above, much of machine learning (and statistical reasoning in general) is about finding “simple” concepts that explain the observed data, and if $NP = P$, we could search for such concepts automatically for any notion of “simplicity” we see fit. In fact, we could even “skip the middle man” and do an automatic search for the learning algorithm with smallest generalization error. Ultimately the field of Artificial Intelligence is about trying to “shortcut” billions of years of evolution to obtain artificial programs that match (or beat) the performance of natural ones, and a fast algorithm for NP would provide the ultimate shortcut.⁵

More generally, a faster algorithm for NP problems would be immensely useful in any field where one is faced with computational or quantitative problems— which is basically all fields of science, math, and engineering. This will not only help with concrete problems such as designing a better bridge, or finding a better drug, but also with addressing basic mysteries such as trying to find scientific theories or “laws of nature”. In a [fascinating talk](#), physicist Nima Arkani-Hamed discusses the effort of finding scientific theories in much the same language as one would describe solving an NP problem, for which the solution is easy to verify or seems “inevitable”, once found, but that requires searching through a huge landscape of possibilities to reach, and that often can get “stuck” at local optima:

“the laws of nature have this amazing feeling of inevitability... which is associated with local perfection.”

⁵ One interesting theory is that $P = NP$ and evolution has already discovered this algorithm, which we are already using without realizing it. At the moment, there seems to be very little evidence for such a scenario. In fact, we have some partial results in the other direction showing that, regardless of whether $P = NP$, many types of “local search” or “evolutionary” algorithms require exponential time to solve 3SAT and other NP -hard problems.

“The classical picture of the world is the top of a local mountain in the space of ideas. And you go up to the top and it looks amazing up there and absolutely incredible. And you learn that there is a taller mountain out there. Find it, Mount Quantum. . . they’re not smoothly connected . . . you’ve got to make a jump to go from classical to quantum . . . This also tells you why we have such major challenges in trying to extend our understanding of physics. We don’t have these knobs, and little wheels, and twiddles that we can turn. We have to learn how to make these jumps. And it is a tall order. And that’s why things are difficult.”

Finding an efficient algorithm for **NP** amounts to always being able to search through an exponential space and find not just the “local” mountain, but the tallest peak.

But perhaps more than any computational speedups, a fast algorithm for **NP** problems would bring about a *new type of understanding*. In many of the areas where **NP**-completeness arises, it is not as much a barrier for solving computational problems as it is a barrier for obtaining “closed-form formulas” or other types of more constructive descriptions of the behavior of natural, biological, social and other systems. A better algorithm for **NP**, even if it is “merely” $2^{\sqrt{n}}$ -time, seems to require obtaining a new way to understand these types of systems, whether it is characterizing Nash equilibria, spin-glass configurations, entangled quantum states, or any of the other questions where **NP** is currently a barrier for analytical understanding. Such new insights would be very fruitful regardless of their computational utility.

17.6 Can $\mathbf{P} \neq \mathbf{NP}$ be neither true nor false?

The **Continuum Hypothesis** is a conjecture made by Georg Cantor in 1878, positing the non-existence of a certain type of infinite cardinality.⁶ This was considered one of the most important open problems in set theory, and settling its truth or falseness was the first problem put forward by Hilbert in the 1900 address we mentioned before. However, using the theories developed by Gödel and Turing, in 1963 Paul Cohen proved that both the Continuum Hypothesis and its negation are consistent with the standard axioms of set theory (i.e., the Zermelo-Fraenkel axioms + the Axiom of choice, or “ZFC” for short).⁷

⁶ One way to phrase it is that for every infinite subset S of the real numbers \mathbb{R} , either there is a one-to-one and onto function $f : S \rightarrow \mathbb{R}$ or there is a one-to-one and onto function $f : S \rightarrow \mathbb{N}$.

⁷ Formally, what he proved is that if ZFC is consistent, then so is ZFC when we assume either the continuum hypothesis or its negation.

Today, many (though not all) mathematicians interpret this result as saying that the Continuum Hypothesis is neither true nor false, but rather is an axiomatic choice that we are free to make one way or the other. Could the same hold for $P \neq NP$?

In short, the answer is *No*. For example, suppose that we are trying to decide between the “3SAT is easy” conjecture (there is an 10^6n time algorithm for 3SAT) and the “3SAT is hard” conjecture (for every n , any NAND program that solves n variable 3SAT takes $2^{10^{-6}n}$ lines). Then, since for $n = 10^8$, $2^{10^{-6}n} > 10^6n$, this boils down to the finite question of deciding whether or not there is a 10^{13} -line NAND program deciding 3SAT on formulas with 10^8 variables.

If there is such a program then there is a finite proof of its existence, namely the approximately 1TB file describing the program, and for which the verification is the (finite in principle though infeasible in practice) process of checking that it succeeds on all inputs.⁸ If there isn’t such a program, then there is also a finite proof of that, though any such proof would take longer since we would need to enumerate over all *programs* as well. Ultimately, since it boils down to a finite statement about bits and numbers; either the statement or its negation must follow from the standard axioms of arithmetic in a finite number of arithmetic steps. Thus, we cannot justify our ignorance in distinguishing between the “3SAT easy” and “3SAT hard” cases by claiming that this might be an inherently ill-defined question. Similar reasoning (with different numbers) applies to other variants of the P vs NP question. We note that in the case that 3SAT is hard, it may well be that there is no *short* proof of this fact using the standard axioms, and this is a question that people have been studying in various restricted forms of proof systems.

⁸ This inefficiency is not necessarily inherent. Later in this course we may discuss results in program-checking, interactive proofs, and average-case complexity, that can be used for efficient verification of proofs of related statements. In contrast, the inefficiency of verifying *failure* of all programs could well be inherent.

17.7 Is $P = NP$ “in practice”?

The fact that a problem is NP -hard means that we believe there is no efficient algorithm that solve it in the *worst case*. It does not, however, mean that every single instance of the problem is hard. For example, if all the clauses in a 3SAT instance φ contain the same variable x_i (possibly in negated form), then by guessing a value to x_i we can reduce φ to a 2SAT instance which can then be efficiently solved. Generalizations of this simple idea are used in “SAT solvers”, which are algorithms that have solved certain specific interesting SAT formulas with thousands of variables, despite the fact that we believe SAT to be exponentially hard in the worst case. Similarly, a lot of problems arising in economics and machine learning are NP -hard.⁹

⁹ Actually, the computational difficulty of problems in economics such as finding optimal (or any) equilibria is quite subtle. Some variants of such problems are NP -hard, while others have a certain “intermediate” complexity.

And yet vendors and customers manage to figure out market-clearing prices (as economists like to point out, there is milk on the shelves) and mice succeed in distinguishing cats from dogs. Hence people (and machines) seem to regularly succeed in solving interesting instances of **NP**-hard problems, typically by using some combination of guessing while making local improvements.

It is also true that there are many interesting instances of **NP**-hard problems that we do *not* currently know how to solve. Across all application areas, whether it is scientific computing, optimization, control or more, people often encounter hard instances of **NP** problems on which our current algorithms fail. In fact, as we will see, all of our digital security infrastructure relies on the fact that some concrete and easy-to-generate instances of, say, 3SAT (or, equivalently, any other **NP**-hard problem) are exponentially hard to solve.

Thus it would be wrong to say that **NP** is easy “in practice”, nor would it be correct to take **NP**-hardness as the “final word” on the complexity of a problem, particularly when we have more information about how any given instance is generated. Understanding both the “typical complexity” of **NP** problems, as well as the power and limitations of certain heuristics (such as various local-search based algorithms) is a very active area of research. We will see more on these topics later in this course.

10

17.8 What if $\mathbf{P} \neq \mathbf{NP}$?

So, $\mathbf{P} = \mathbf{NP}$ would give us all kinds of fantastical outcomes. But we strongly suspect that $\mathbf{P} \neq \mathbf{NP}$, and moreover that there is no much-better-than-brute-force algorithm for 3SAT. If indeed that is the case, is it all bad news?

One might think that impossibility results, telling you that you *cannot* do something, is the kind of cloud that does not have a silver lining. But in fact, as we already alluded to before, it does. A hard (in a sufficiently strong sense) problem in **NP** can be used to create a code that *cannot be broken*, a task that for thousands of years has been the dream of not just spies but of many scientists and mathematicians over the generations. But the complexity viewpoint turned out to yield much more than simple codes, achieving tasks that people had previously not even dared to dream of. These include the notion of *public key cryptography*, allowing two people to communicate securely without ever having exchanged a secret key; *electronic cash*, allowing

¹⁰ Talk more about coping with **NP** hardness. Main two approaches are *heuristics* such as SAT solvers that succeed on *some* instances, and *proxy measures* such as mathematical relaxations that instead of solving problem X (e.g., an integer program) solve program X' (e.g., a linear program) that is related to that. Maybe give compressed sensing as an example, and least square minimization as a proxy for maximum a posteriori probability.

private and secure transaction without a central authority; and *secure multiparty computation*, enabling parties to compute a joint function on private inputs without revealing any extra information about it. Also, as we will see, computational hardness can be used to replace the role of *randomness* in many settings.

Furthermore, while it is often convenient to pretend that computational problems are simply handed to us, and that our job as computer scientists is to find the most efficient algorithm for them, this is not how things work in most computing applications. Typically even formulating the problem to solve is a highly non-trivial task. When we discover that the problem we want to solve is **NP**-hard, this might be a useful sign that we used the wrong formulation for it.

Beyond all these, the quest to understand computational hardness – including the discoveries of lower bounds for restricted computational models, as well as new types of reductions (such as those arising from “probabilistically checkable proofs”) – has already had surprising *positive* applications to problems in algorithm design, as well as in coding for both communication and storage. This is not surprising since, as we mentioned before, from group theory to the theory of relativity, the pursuit of impossibility results has often been one of the most fruitful enterprises of mankind.

17.9 Lecture summary

- The question of whether $P = NP$ is one of the most important and fascinating questions of computer science and science at large, touching on all fields of the natural and social sciences, as well as mathematics and engineering.
- Our current evidence and understanding supports the “SAT hard” scenario that there is no much-better-than-brute-force algorithm for 3SAT or many other **NP**-hard problems.
- We are very far from *proving* this, however. Researchers have studied proving lower bounds on the number of gates to compute explicit functions in *restricted forms* of circuits, and have made some advances in this effort, along the way generating mathematical tools that have found other uses. However, we have made essentially no headway in proving lower bounds for *general* models of computation such as NAND and NAND++ programs. Indeed, we currently do not even know how to rule out the possibility that for every $n \in \mathbb{N}$, SAT restricted to n -length inputs has a NAND

program of $10n$ lines (even though there *exist* n -input functions that require $2^n / (10n)$ lines to compute).

- Understanding how to cope with this computational intractability, and even benefit from it, comprises much of the research in theoretical computer science.

17.10 Exercises

17.11 Bibliographical notes

11

¹¹ TODO: Scott's two surveys

17.12 Further explorations

Some topics related to this lecture that might be accessible to advanced students include: (to be completed)

- Polynomial hierarchy hardness for circuit minimization and related problems, see for example [this paper](#).

17.13 Acknowledgements