

“Why Are We Learning This?” : The Importance of Theoretical Computer Science

IEEE Canadian Atlantic Computer Society

Taylor J. Smith

Department of Computer Science
St. Francis Xavier University
Antigonish, Nova Scotia, Canada

December 9, 2024

- ▶ Assistant professor, StFX
Dept. of Computer Science
- ▶ Director, FLAReLab
- ▶ Research specialties:
 - ▶ Automata theory
 - ▶ Formal languages
- ▶ Teaching specialties:
 - ▶ Theory of computing
 - ▶ Algorithm design & analysis
 - ▶ Computational logic
- ▶ PhD '21, Queen's University
- ▶ IEEE membership:
 - ◆ GS'16–21
 - ◆ M'22–present



Why are we learning TCS?

- Traditional TCS courses

- What do students think?

- What do academics think?

We're *really* learning TCS...

- ... To know how to do something

- ... To understand what we can't do

- ... Because it's fun!

How do we make learning TCS better?

Why are we learning TCS?

Traditional TCS courses

What do students think?

What do academics think?

We're *really* learning TCS...

... To know how to do something

... To understand what we can't do

... Because it's fun!

How do we make learning TCS better?

- ▶ There's a reason this question is the title of the talk—it's one of the most common refrains from students!

- ▶ There's a reason this question is the title of the talk—it's one of the most common refrains from students!

Officially?

- ▶ Because we have to.
 - ▶ Every Canadian U15 university has some kind of TCS course on the books, usually required to earn a CS degree
 - ▶ Ditto for the Maple League universities

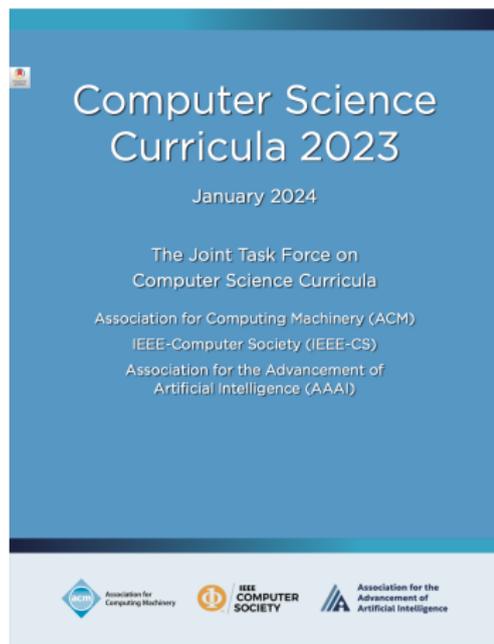
- ▶ There's a reason this question is the title of the talk—it's one of the most common refrains from students!

Officially?

- ▶ Because we have to.
 - ▶ Every Canadian U15 university has some kind of TCS course on the books, usually required to earn a CS degree
 - ▶ Ditto for the Maple League universities

Unofficially?

- ▶ Because it's the foundation of all of computer science!
 - ▶ There isn't a single result in any other CS course that doesn't rely on theory in some way



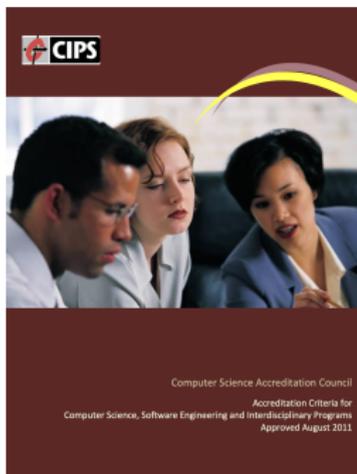
AL-Models: Computational Models and Formal Languages

CS Core:

1. Formal automata
 - a. Finite State
 - b. Pushdown
 - c. Linear Bounded
 - d. Turing Machine
2. Formal languages, grammars and Chomsky Hierarchy
(See also: [FPL-Translation](#), [FPL-Syntax](#))
 - a. Regular (Type-3)
 - i. Regular Expressions
 - b. Context-Free (Type-2)
 - c. Context-Sensitive (Type-1)
 - d. Recursively Enumerable (Type-0)
3. Relations among formal automata, languages, and grammars
4. Decidability, (un)computability, and halting
5. The Church-Turing thesis
6. Algorithmic correctness
 - a. Invariants (e.g., in iteration, recursion, tree search)

KA Core:

7. Deterministic and nondeterministic automata
8. Pumping Lemma proofs
 - a. Proof of Finite State/Regular-Language limitation
 - b. Pushdown Automata/Context-Free-Language limitation
9. Decidability
 - a. Arithmetization and diagonalization
10. Reducibility and reductions
11. Time complexity based on Turing Machine
12. Space complexity (e.g., Pspace, Savitch's Theorem)
13. Equivalent models of algorithmic computation
 - a. Turing Machines and Variations (e.g., multi-tape, non-deterministic)
 - b. Lambda Calculus (See also: [FPL-Functional](#))
 - c. Mu-Recursive Functions



7.3.1 Key Areas of Computer Science

All students in all accredited programs, including computer science, software engineering and interdisciplinary programs should be required to take courses which satisfy graduate attributes 1 (a), (b), (c) (d) and (e) described in Section 4.1. The following paragraphs expand on the range of knowledge and skills that students might acquire in the respective areas. Introductory courses in computer programming and computer usage would not be expected to meet these requirements. See sections 7.3.2 to 7.3.4 regarding specifics of how to apply this section to computer science, software engineering and interdisciplinary programs respectively.

This classification of five subtopics is not intended to limit the scope of computing nor the areas to be included or excluded at particular institutions, but rather it is illustrative of a range of topics that an accredited program would be expected to provide.

Graduate Attribute 1a: Software engineering, including requirements specification, software design and software architecture, software development, software testing, software maintenance, and other topics related to software process (For Software Engineering programs this list is expanded in Section 7.3.3).

Graduate Attribute 1b: Algorithms and data structures, including data structures such as stacks, trees, lists, queues, etc.; abstract data types, established solutions to classical problems (e.g., sorting and searching), and analysis of algorithms.

Graduate Attribute 1c: Systems software, including operating systems concepts, virtual memory management, management of distributed, parallel, and concurrent processes, transaction processing, logging, security, and computer networking.

Graduate Attribute 1d. Computer elements and architectures, including computer organization, digital device and communications technology, logical and physical hardware design.

Graduate Attribute 1e. Theoretical foundations of computing, including models of computation, analysis of algorithms, fundamentals of program specification and verification, computational complexity, grammars and automata.

- ▶ Usually, a traditional TCS course builds upward from a very simple model of computation.
 - ▶ finite automata (= regular expressions)
 - ▶ pushdown automata (= context-free grammars)
 - ▶ linear-bounded automata
 - ▶ Turing machines

- ▶ Usually, a traditional TCS course builds upward from a very simple model of computation.
 - ▶ finite automata (= regular expressions)
 - ▶ pushdown automata (= context-free grammars)
 - ▶ linear-bounded automata
 - ▶ Turing machines
- ▶ Once we reach Turing machines, we can start to talk more about what computers *do* rather than what computers *are*.
 - ▶ decidable problems
 - ▶ undecidable problems

- ▶ Usually, a traditional TCS course builds upward from a very simple model of computation.
 - ▶ finite automata (= regular expressions)
 - ▶ pushdown automata (= context-free grammars)
 - ▶ linear-bounded automata
 - ▶ Turing machines
- ▶ Once we reach Turing machines, we can start to talk more about what computers *do* rather than what computers *are*.
 - ▶ decidable problems
 - ▶ undecidable problems
- ▶ Finally, once we have a collection of problems, we can start to draw connections between problems.
 - ▶ reductions
 - ▶ complexity theory
 - ▶ hardness and completeness

COMPUTER SCIENCE 3331A/B FOUNDATIONS OF COMPUTER SCIENCE I



Languages as sets of strings over an alphabet; operations on languages; finite automata, regular expressions; language hierarchy; Turing machines; models of computation.

Prerequisite(s): Computer Science 2214A/B or Mathematics 2155F/G.

Extra Information: 3 lecture hours.

COURSE WEIGHT: 0.50

[More details](#)

CS360 – Introduction to the Theory of Computing

Description

Models of computers including finite automata and Turing machines. Basics of formal languages with applications to the syntax of programming languages. Alternate characterizations of language classes. Proving unrecognizability. Unsolvable problems and their relevance to the semantics of programming.

CISC 223 Software Specifications Units: 3.00

Introduction to techniques for specifying the behaviour of software, with applications of these techniques to design, verification and construction of software. Logic-based techniques such as loop invariants and class invariants. Automata and grammar-based techniques, with applications to scanners, parsers, user-interface dialogs and embedded systems. Computability issues in software specifications.

356 Theory of Computing

An introduction to the theoretical foundations of computer science, examining finite automata, context-free grammars, Turing machines, decidability and undecidability, and complexity theory. Strategies will be developed to help categorize problems as tractable or intractable. Prerequisites: CSCI 255, 277. Three credits.

Actual evaluation comments:

- ▶ “at the end of the semester I still can’t understand the real life application of this subject, I never feel I can apply my any coding skills or even mathematical skills here.”
- ▶ “I think it would help if the supplemental material was presented in a way that’s a bit easier to visualize.”
- ▶ “Maybe [add] some application based questions! I know it is theoretical computer science, but something to contextualize what we’re learning.”
- ▶ “This subject is mostly on theory, it would be better if add any practical examples or projects.”

Actual evaluation comments:

- ▶ “I started the semester enrolled in 6 courses with the intention of doing a vibe check and cutting it down to 4. You passed the vibe check by asking if anyone had played Dwarf Fortress. I am not sure how repeatable this strategy is, but it worked.”

Actual evaluation comments:

- ▶ “I started the semester enrolled in 6 courses with the intention of doing a vibe check and cutting it down to 4. You passed the vibe check by asking if anyone had played Dwarf Fortress. I am not sure how repeatable this strategy is, but it worked.”

Key takeaway:

- ▶ Talk about **Dwarf Fortress** 🐉



- ▶ At STOC 1996, a group of researchers presented a report assessing the state of the field and making recommendations for the future.

In order for [the Theory of Computing] to prosper in the coming years, it is essential to strengthen our communication with the rest of computer science and with other disciplines, and to increase our impact on key application areas.

- ▶ At STOC 1996, a group of researchers presented a report assessing the state of the field and making recommendations for the future.

In order for [the Theory of Computing] to prosper in the coming years, it is essential to strengthen our communication with the rest of computer science and with other disciplines, and to increase our impact on key application areas.

- ▶ This report proved to be contentious, and some other researchers quickly drafted a rebuttal.

In order for the Theory of Computing to prosper in the future it is essential that [it] attracts the same calibre of researchers, that Theoretical Computer Scientists concentrate their research efforts in Theory of Computing and that they enjoy the freedom to do so.

- ▶ At STOC 1996, a group of researchers presented a report assessing the state of the field and making recommendations for the future.

In order for [the Theory of Computing] to prosper in the coming years, it is essential to strengthen our communication with the rest of computer science and with other disciplines, and to increase our impact on key applications.

Why not both?

- ▶ This report proved to be controversial, and some other researchers quickly drafted a rebuttal.

In order for the Theory of Computing to prosper in the future it is essential that [it] attracts the same calibre of researchers, that Theoretical Computer Scientists concentrate their research efforts in Theory of Computing and that they enjoy the freedom to do so.

The *real* key takeaways:

- ▶ TCS is (admittedly) a dry subject to students—but it doesn't have to be this way.
- ▶ We should appreciate TCS as a subject in itself.
 - ▶ There's nothing wrong with appreciating the fascinating results as they are
 - ▶ That's what attracted me to the field!
- ▶ At the same time, we should be mindful that TCS is more than just a bunch of theorems and proofs.
 - ▶ TCS is a broadly applicable method of reasoning about how to solve/compute problems
 - ▶ And it has a rich history!

Why are we learning TCS?

Traditional TCS courses

What do students think?

What do academics think?

We're *really* learning TCS...

... To know how to do something

... To understand what we can't do

... Because it's fun!

How do we make learning TCS better?

- ▶ Most traditional TCS courses follow a rigid “definition–theorem–proof” lecture style.
 - ▶ Probably the main source of student consternation!
- ▶ There’s a myriad of interesting examples, applications, and anecdotes that can help reveal to students why they’re *really* learning TCS.
 - ▶ How can we use TCS concepts to do something?
 - ▶ How can we use TCS concepts to know our limits?
 - ▶ Can we have fun with this?

Why are we learning TCS?

Traditional TCS courses

What do students think?

What do academics think?

We're *really* learning TCS...

... To know how to do something

... To understand what we can't do

... Because it's fun!

How do we make learning TCS better?

Background:

A *one-dimensional word* is a string of symbols from some alphabet Σ . For example, if $\Sigma = \{a, b\}$, one possible 1D word is abbaba.

A *two-dimensional word* is an array of symbols from some alphabet Σ . For example, if $\Sigma = \{0, 1\}$, one possible 2D word is:

$$\begin{bmatrix} 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 \end{bmatrix}$$

- ▶ Words (and collections of words, called *languages*) form the foundations of much of computer science.
- ▶ All data we can store in a computer's memory/disk is essentially a binary word!

- ▶ Some data we care about exhibits certain special properties.
- ▶ For example, *periodicity*: the property of some segment of data repeating as a pattern.

Example:

A periodic 1D word over $\Sigma = \{a, b\}$ is abbaabbaabba. The pattern is abba, repeated three times.

A periodic 2D word over $\Sigma = \{0, 1\}$ is

$$\begin{bmatrix} 0 & 1 & 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 & 1 & 0 \end{bmatrix}$$

The pattern is $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$, repeated 2×3 times.

- ▶ Some data we care about exhibits certain special properties.
- ▶ For example, *periodicity*: the property of some segment of data repeating as a pattern.
- ▶ Periodicity is a fundamental notion in formal languages and combinatorics on words.
- ▶ Patterns show up everywhere in the real world!
 - ▶ E.g., biology, genomics, astronomy, meteorology, earth sciences, oceanography, etc.

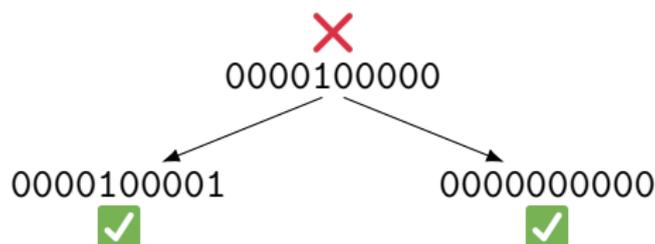
- ▶ Some data we care about exhibits certain special properties.
- ▶ For example, *periodicity*: the property of some segment of data repeating as a pattern.
- ▶ Periodicity is a fundamental notion in formal languages and combinatorics on words.
- ▶ Patterns show up everywhere in the real world!
 - ▶ E.g., biology, genomics, astronomy, meteorology, earth sciences, oceanography, etc.
- ▶ For this reason, many researchers have come up with efficient algorithms to detect periodicity.
 - ▶ 1D: $O(n)$ time, word of length n^*
 - ▶ 2D: $O(mn)$ time, word of size $m \times n$ (Smith et al., 2017)

*R. Kolpakov and G. Kucherov. Finding maximal repetitions in a word in linear time. In Proc. of FOCS 1999, pages 596–604. 1999. 

- ▶ Of course, data storage isn't perfect.
- ▶ Data can be corrupted with *errors*.
- ▶ How can we *detect* and *correct* such errors if they occur?
 - ▶ In the context of periodicity, how do we *recover the pattern* from an erroneous word?

- ▶ Of course, data storage isn't perfect.
- ▶ Data can be corrupted with *errors*.
- ▶ How can we *detect* and *correct* such errors if they occur?
 - ▶ In the context of periodicity, how do we *recover the pattern* from an erroneous word?

In one dimension:

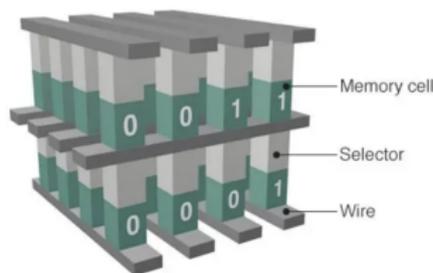


- ▶ Even with only one error, there could be many candidates.
- ▶ However, we just need to check $O(\log(n))$ possibilities.

- ▶ Of course, data storage isn't perfect.
- ▶ Data can be corrupted with *errors*.
- ▶ How can we *detect* and *correct* such errors if they occur?
 - ▶ In the context of periodicity, how do we *recover the pattern* from an erroneous word?

In more than one dimension?

- ▶ Yes, multidimensional storage media exists!
 - ▶ E.g., Intel's *3D XPoint*



- ▶ Suppose our data has size $m \times n$; the pattern has size $p \times q$.
- ▶ As long as the number of errors is bounded above by $\left\lfloor \frac{1}{2+\epsilon} \lfloor \frac{n}{p} \rfloor \lfloor \frac{m}{q} \rfloor \right\rfloor$, we need only check $O(\log(mn))$ possibilities!
- ▶ How do we correct errors?
 - ▶ Perform Karp–Miller–Rosenberg (KMR) renaming to assign “labels” to unique subwords
 - ▶ Find a candidate/approximate pattern having size $p \times q$
 - ▶ Verify that this candidate/approximate pattern is primitive (not periodic)
 - ▶ Compute the Hamming distance between pattern and data; if distance is close enough, correct the error

- ▶ The same error correction procedure can be extended to $d \geq 3$ dimensions.
- ▶ As we might expect, we need to check $O(\log^d(N))$ possibilities.
 - ▶ Here, $N = \prod_{i=1}^d n_i$
- ▶ With a small tweak, however, it is possible to make the number of possibilities independent of the dimensions; i.e., just $O(\log(N))$ possibilities!

- ▶ The same error correction procedure can be extended to $d \geq 3$ dimensions.
- ▶ As we might expect, we need to check $O(\log^d(N))$ possibilities.
 - ▶ Here, $N = \prod_{i=1}^d n_i$
- ▶ With a small tweak, however, it is possible to make the number of possibilities independent of the dimensions; i.e., just $O(\log(N))$ possibilities!
- ▶ In practice, we must also account for $d \cdot 2^d$ overhead.
- ▶ But still, simple techniques (periodicity checking) allow us to do crucial real-world tasks!

Why are we learning TCS?

Traditional TCS courses

What do students think?

What do academics think?

We're *really* learning TCS...

... To know how to do something

... To understand what we can't do

... Because it's fun!

How do we make learning TCS better?

Background:

The *Boolean satisfiability problem* (or SAT) is defined as follows:

SATISFIABILITY

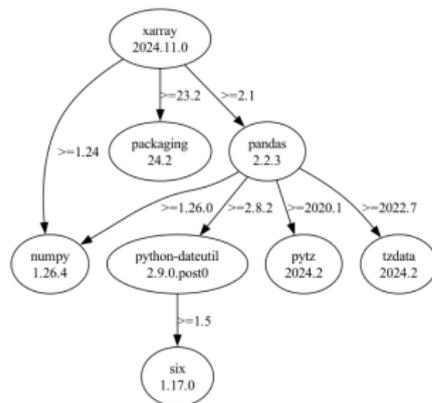
Given a Boolean logic formula in conjunctive normal form, does there exist an assignment of true/false values to each variable that makes the formula true?

Conjunctive normal form means the formula consists of conjuncted clauses of disjuncted variables. For example:

$$\phi = (x_1 \vee \overline{x_2} \vee x_3) \wedge (\overline{x_1} \vee x_4) \wedge (x_2 \vee \overline{x_3}) \wedge (\overline{x_5}).$$

- ▶ There are many heuristic SAT-solving algorithms for practical problems (and competitions).
- ▶ However, no efficient algorithm is known to solve instances of SAT in general!

- ▶ Many programmers rely on packages made by other programmers to perform common tasks.
 - ▶ E.g., if you're doing some kind of numerical computation in Python, you won't start from scratch... you'll `import numpy`
- ▶ These packages themselves don't always start from scratch; some packages depend on other packages.



An example of a dependency tree produced by `pipdeptree`

- ▶ How might we automatically determine/resolve dependencies?

- ▶ How might we automatically determine/resolve dependencies?
- ▶ Let's use Boolean logic: formulate the problem in terms of the SATISFIABILITY problem.

- ▶ How might we automatically determine/resolve dependencies?
- ▶ Let's use Boolean logic: formulate the problem in terms of the SATISFIABILITY problem.

Example:

Suppose we want to install v1.0 of the foo package.

(foo 1.0)

- ▶ How might we automatically determine/resolve dependencies?
- ▶ Let's use Boolean logic: formulate the problem in terms of the SATISFIABILITY problem.

Example:

foo v1.0 depends on the bar package: either v1.5 or v2.0.

$$(\text{foo } 1.0) \wedge (\neg \text{foo } 1.0 \vee \text{bar } 1.5 \vee \text{bar } 2.0)$$

- ▶ How might we automatically determine/resolve dependencies?
- ▶ Let's use Boolean logic: formulate the problem in terms of the SATISFIABILITY problem.

Example:

bar v1.5 depends on v3.4 of the baz package.

$$(\text{foo } 1.0) \wedge (\neg \text{foo } 1.0 \vee \text{bar } 1.5 \vee \text{bar } 2.0) \wedge (\neg \text{bar } 1.5 \vee \text{baz } 3.4)$$

- ▶ How might we automatically determine/resolve dependencies?
- ▶ Let's use Boolean logic: formulate the problem in terms of the SATISFIABILITY problem.

Example:

bar v2.0 depends on v3.6 of the baz package.

$$\begin{aligned} &(\text{foo } 1.0) \wedge (\neg \text{foo } 1.0 \vee \text{bar } 1.5 \vee \text{bar } 2.0) \wedge (\neg \text{bar } 1.5 \vee \text{baz } 3.4) \\ &\quad \wedge (\neg \text{bar } 2.0 \vee \text{baz } 3.6) \end{aligned}$$

- ▶ How might we automatically determine/resolve dependencies?
- ▶ Let's use Boolean logic: formulate the problem in terms of the SATISFIABILITY problem.

Example:

We can't install both bar v1.5 and v2.0 at the same time.

$$\begin{aligned} &(\text{foo } 1.0) \wedge (\neg \text{foo } 1.0 \vee \text{bar } 1.5 \vee \text{bar } 2.0) \wedge (\neg \text{bar } 1.5 \vee \text{baz } 3.4) \\ &\quad \wedge (\neg \text{bar } 2.0 \vee \text{baz } 3.6) \wedge (\neg \text{bar } 1.5 \vee \neg \text{bar } 2.0) \end{aligned}$$

- ▶ How might we automatically determine/resolve dependencies?
- ▶ Let's use Boolean logic: formulate the problem in terms of the SATISFIABILITY problem.

Example:

We can't install both baz v3.4 and v3.6 at the same time.

$$\begin{aligned} & (\text{foo } 1.0) \wedge (\neg \text{foo } 1.0 \vee \text{bar } 1.5 \vee \text{bar } 2.0) \wedge (\neg \text{bar } 1.5 \vee \text{baz } 3.4) \\ & \wedge (\neg \text{bar } 2.0 \vee \text{baz } 3.6) \wedge (\neg \text{bar } 1.5 \vee \neg \text{bar } 2.0) \wedge (\neg \text{baz } 3.4 \vee \neg \text{baz } 3.6) \end{aligned}$$

- ▶ How might we automatically determine/resolve dependencies?
- ▶ Let's use Boolean logic: formulate the problem in terms of the SATISFIABILITY problem.

Example:

The baz developers just took down v3.6 due to a bug.

$$\begin{aligned} & (\text{foo } 1.0) \wedge (\neg \text{foo } 1.0 \vee \text{bar } 1.5 \vee \text{bar } 2.0) \wedge (\neg \text{bar } 1.5 \vee \text{baz } 3.4) \\ & \wedge (\neg \text{bar } 2.0 \vee \text{baz } 3.6) \wedge (\neg \text{bar } 1.5 \vee \neg \text{bar } 2.0) \wedge (\neg \text{baz } 3.4 \vee \neg \text{baz } 3.6) \\ & \wedge (\neg \text{baz } 3.6) \end{aligned}$$

- ▶ How might we automatically determine/resolve dependencies?
- ▶ Let's use Boolean logic: formulate the problem in terms of the SATISFIABILITY problem.

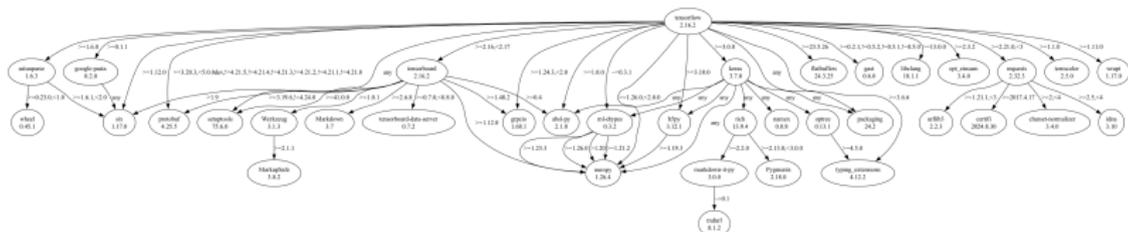
Example:

So, how do we solve this dependency problem?

$$\begin{aligned} & (\text{foo } 1.0) \wedge (\neg \text{foo } 1.0 \vee \text{bar } 1.5 \vee \text{bar } 2.0) \wedge (\neg \text{bar } 1.5 \vee \text{baz } 3.4) \\ & \wedge (\neg \text{bar } 2.0 \vee \text{baz } 3.6) \wedge (\neg \text{bar } 1.5 \vee \neg \text{bar } 2.0) \wedge (\neg \text{baz } 3.4 \vee \neg \text{baz } 3.6) \\ & \qquad \qquad \qquad \wedge (\neg \text{baz } 3.6) \end{aligned}$$

- ▶ $\text{baz } 3.6 = \text{False}$ (clause 7; we can't get this package)
- ▶ $\text{bar } 2.0 = \text{False}$ (clause 4, since we can't get $\text{baz } 3.6$)
- ▶ $\text{bar } 1.5 = \text{True}$ (clause 2, since we need some bar package)
- ▶ $\text{baz } 3.4 = \text{True}$ (clause 3, since $\text{bar } 1.5$ depends on this)
- ▶ $\text{foo } 1.0 = \text{True}$ (clause 1; we're installing this package!)

- ▶ Unfortunately, the example we just saw was an easy one.
- ▶ Some package dependency trees are... less easy to follow.



- ▶ Since we have no efficient algorithm for SATISFIABILITY, there is no efficient way in general to resolve dependencies!
- ▶ In practice, package managers like pip use *backtracking algorithms* and *heuristics*.
 - ▶ Still, there are some cases where untangling dependencies will take a *long* time!

Background:

The *halting problem* is defined as follows:

HALT

Given a description of a program and an input to that program, does the program halt when run on that input?

The halting problem is *undecidable*; that is, there does not exist an algorithm that always returns a yes/no answer for every pair of program and input.

- ▶ The halting problem is one of the most famous problems in all of computer science.
- ▶ Note that while the problem is *undecidable*, this doesn't mean it's *impossible*.
 - ▶ If a program halts on an input, we can definitely detect that; if it gets caught in an infinite loop, we'll never get an answer

- ▶ If we want to show that other decision problems (i.e., “yes/no” problems) are undecidable, then we can use something known as a *reduction*.
- ▶ How does a reduction work?
 - ▶ Let's suppose we know problem X is undecidable
 - ▶ We want to learn more about a new problem Y

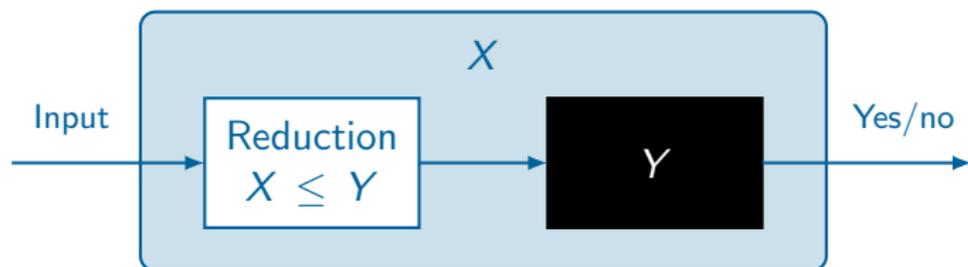
- ▶ If we want to show that other decision problems (i.e., “yes/no” problems) are undecidable, then we can use something known as a *reduction*.
- ▶ How does a reduction work?
 - ▶ Let's suppose we know problem X is undecidable
 - ▶ We want to learn more about a new problem Y
 - ▶ If we assume that we can write an algorithm to decide Y ...



- ▶ If we want to show that other decision problems (i.e., “yes/no” problems) are undecidable, then we can use something known as a *reduction*.
- ▶ How does a reduction work?
 - ▶ Let's suppose we know problem X is undecidable
 - ▶ We want to learn more about a new problem Y
 - ▶ If we assume that we can write an algorithm to decide Y ...
 - ▶ ...and show that we can turn X into Y ...



- ▶ If we want to show that other decision problems (i.e., “yes/no” problems) are undecidable, then we can use something known as a *reduction*.
- ▶ How does a reduction work?
 - ▶ Let's suppose we know problem X is undecidable
 - ▶ We want to learn more about a new problem Y
 - ▶ If we assume that we can write an algorithm to decide Y ...
 - ▶ ...and show that we can turn X into Y ...
 - ▶ ...then we just created an algorithm to “decide” X !
 - ▶ But this can't be, so Y must be undecidable like X



- ▶ There are all kinds of undecidable problems beyond HALT.

- ▶ There are all kinds of undecidable problems beyond HALT.
- ▶ Some of these problems are computational:
 - ▶ Given a Turing machine and an input word, does it accept that word?
 - ▶ Given a second-order λ -calculus expression, can we detect its type?

- ▶ There are all kinds of undecidable problems beyond HALT.
- ▶ Some of these problems are computational:
 - ▶ Given a Turing machine and an input word, does it accept that word?
 - ▶ Given a second-order λ -calculus expression, can we detect its type?
- ▶ Some of these problems are mathematical:
 - ▶ Given a set of $n \times n$ integral matrices, can we obtain the zero matrix by taking some product of matrices from this set?
 - ▶ Given two finite group presentations, are these groups isomorphic?

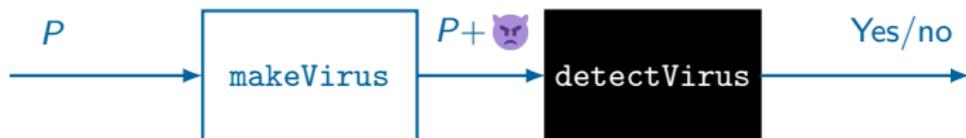
- ▶ There are all kinds of undecidable problems beyond HALT.
- ▶ Some of these problems are computational:
 - ▶ Given a Turing machine and an input word, does it accept that word?
 - ▶ Given a second-order λ -calculus expression, can we detect its type?
- ▶ Some of these problems are mathematical:
 - ▶ Given a set of $n \times n$ integral matrices, can we obtain the zero matrix by taking some product of matrices from this set?
 - ▶ Given two finite group presentations, are these groups isomorphic?
- ▶ Some of these problems are just plain interesting:
 - ▶ Does one player have a winning strategy in *Magic: The Gathering*?
 - ▶ **Can we detect a virus in a computer file?**

- ▶ Virus detecting software exists, but it suffers occasionally from false positives/false negatives.
- ▶ Why don't we have 100% flawless virus detection software?

- ▶ Virus detecting software exists, but it suffers occasionally from false positives/false negatives.
- ▶ Why don't we have 100% flawless virus detection software?
- ▶ Let's suppose we do (and see what goes wrong)!
- ▶ Assume we have a function `detectVirus(P)` that takes as input a program's source code P and returns "yes" if it's a virus, and "no" otherwise.

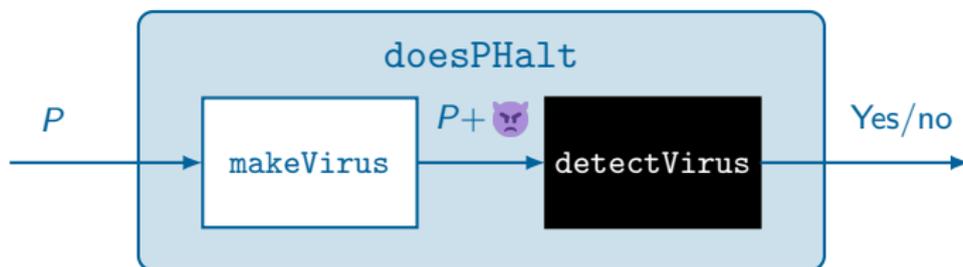


- ▶ Let's write a new function `makeVirus(P)` that takes a program's source code P and appends `"sudo rm -rf /*"` to the end of the program, then pass this to `detectVirus`.
- ▶ If `detectVirus` returns "yes", then P (whatever it was) must have finished running before executing `rm -rf`.
- ▶ If `detectVirus` returns "no", then we never reached `rm -rf`, and so P must not have finished running.



- ▶ But wait... what did we just do?
- ▶ We implicitly decided whether or not P halts!
 - ▶ We only reach the bad code inserted by `makeVirus` if P halts
 - ▶ Thus, the answer from `detectVirus` will be the same as the answer for `doesPHalt`
- ▶ Since `HALT` is undecidable, so is virus detection!

(Note. It's possible that P may already contain a virus and also never halt. This is easy for us to fix: just modify P so it's incapable of infecting anything on the computer itself; e.g., redirect P 's output to a VM, etc.)



Why are we learning TCS?

Traditional TCS courses

What do students think?

What do academics think?

We're *really* learning TCS...

... To know how to do something

... To understand what we can't do

... Because it's fun!

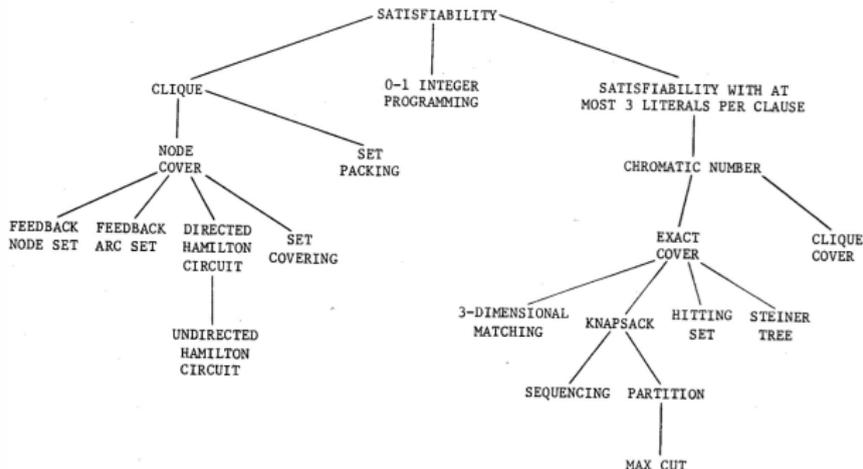
How do we make learning TCS better?

Background:

A problem is *NP-complete* if:

- ▶ using a deterministic computer, we can efficiently *verify* whether a solution is correct (i.e., *polynomial time* in the input size); and
 - ▶ every other problem that can be verified in polynomial time can be transformed into this problem efficiently (i.e., via a *polynomial-time reduction*).
-
- ▶ The first criterion indicates that the problem is in NP.
 - ▶ It's “easy” to *check* a solution, but “difficult” to *find* a solution
 - ▶ The second criterion indicates that the problem is *NP-hard*.
 - ▶ Every other problem in NP can be represented as an instance of this problem
 - ▶ Alternatively, this problem is at least as difficult to *solve* as any other problem in NP

- ▶ There are many, many examples of NP-complete problems!



- ▶ However, some of these problems aren't exactly run-of-the-mill...

February 4, 2024:



My car, pre-shovelling

February 4, 2024:



My car, mid-shovelling

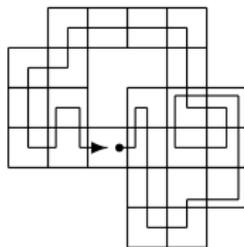
- ▶ Could this job have been easier if I had a snowblower?
- ▶ Preferably, an automatic snowblower that routes itself around my driveway?

- ▶ Could this job have been easier if I had a snowblower?
- ▶ Preferably, an automatic snowblower that routes itself around my driveway?
- ▶ Let's formulate this as a problem:

SNOWBLOWING

Does there exist a path to allow a snowblower to clear a given polygonal region without piling snow deeper than some depth d within this region?

Example:



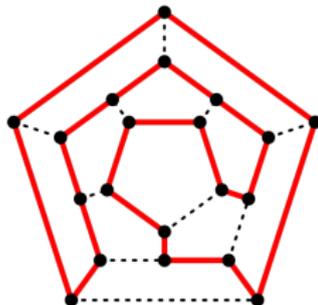
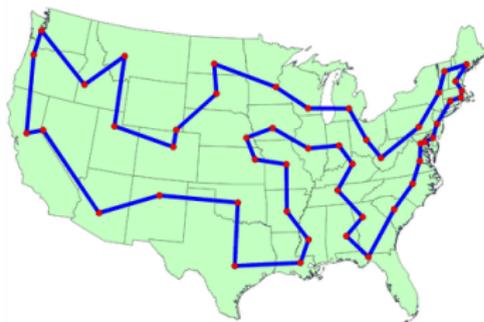
- ▶ This looks quite similar to two very well-studied (and NP-complete) problems!

TRAVELLING-SALESPERSON

Does there exist a route of length at most ℓ that visits a list of cities and also returns you to your home city?

HAMILTONIAN-CYCLE

Does there exist a path in a graph that visits each vertex exactly once and returns to the starting vertex?



- ▶ It's possible to reduce the HAMILTONIAN-CYCLE problem to the SNOWBLOWING problem.
- ▶ This means that the SNOWBLOWING problem is NP-complete!
 - ▶ I.e., we can easily *check* whether a snowblower path clears the region, but it's much harder to *find* a path from scratch
- ▶ Does this mean all is lost?

- ▶ It's possible to reduce the HAMILTONIAN-CYCLE problem to the SNOWBLOWING problem.
- ▶ This means that the SNOWBLOWING problem is NP-complete!
 - ▶ I.e., we can easily *check* whether a snowblower path clears the region, but it's much harder to *find* a path from scratch
- ▶ Does this mean all is lost?
- ▶ Actually, no: for the simplest variant of the SNOWBLOWING problem, there exists an *8-approximation algorithm*.
 - ▶ This means that we can find a path that is at most *eight* times as long as the optimal path

Background:

A problem is *PSPACE-complete* if:

- ▶ we can *solve* the problem using a “reasonable” amount of memory (i.e., *polynomial space* in the input size); and
 - ▶ every other problem that can be solved using polynomial space can be transformed into this problem efficiently (i.e., PSPACE-hard via a *polynomial-time reduction*).
-
- ▶ Note that we're changing our perspective from *time* to *space*.
 - ▶ We know that $\text{NP} \subseteq \text{PSPACE}$, since there exist problems whose solutions can be found using more than polynomial time, but only polynomial space.

- ▶ There are many problems in CS/mathematics that are PSPACE-complete.
- ▶ Polynomial space also frequently arises when we play games!



GEOGRAPHY



HEX



OTHELLO



CHESS



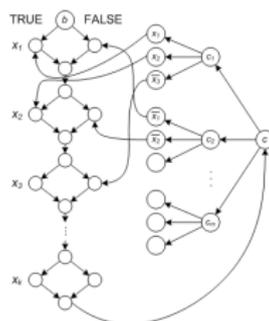
Go

(Actually PSPACE-hard)

- ▶ To obtain the reduction that is vital to proving a problem is PSPACE-complete, one often uses *gadgets*.
 - ▶ Gadgets can also be used in NP-completeness proofs
- ▶ A gadget is like a “piece” of one problem; multiple “pieces” can be assembled into an instance of another problem we’re reducing to.

Example:

We can reduce to GEOGRAPHY from another PSPACE-complete problem called FORMULA-GAME using the collection of gadgets shown at right.



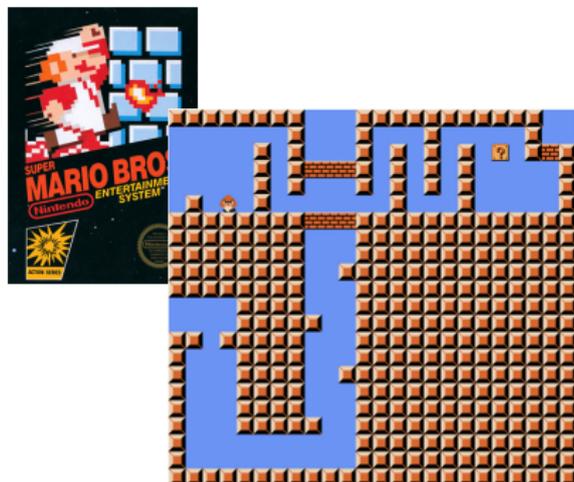
- ▶ If gadgets are like “pieces” of problems, where might these gadgets appear in something we’re familiar with?

- ▶ If gadgets are like “pieces” of problems, where might these gadgets appear in something we’re familiar with?
- ▶ How about video games? (Specifically, game levels/objects)

- ▶ If gadgets are like “pieces” of problems, where might these gadgets appear in something we’re familiar with?
- ▶ How about video games? (Specifically, game levels/objects)

Super Mario Bros.:

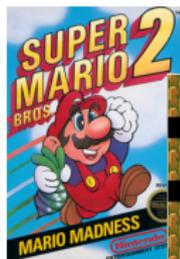
Deciding whether we can get from the start of a level to the goal is NP-complete
(Reduce from 3SAT)



- ▶ If gadgets are like “pieces” of problems, where might these gadgets appear in something we’re familiar with?
- ▶ How about video games? (Specifically, game levels/objects)

Super Mario Bros. 2:

Deciding whether we can get from the start of a level to the goal is NP-complete
(Reduce from 3SAT)



- ▶ If gadgets are like “pieces” of problems, where might these gadgets appear in something we’re familiar with?
- ▶ How about video games? (Specifically, game levels/objects)

Super Mario Bros. 3:

Deciding whether we can get from the start of a level to the goal is NP-complete
(Reduce from 3SAT)



- ▶ If gadgets are like “pieces” of problems, where might these gadgets appear in something we’re familiar with?
- ▶ How about video games? (Specifically, game levels/objects)

Donkey Kong Country:

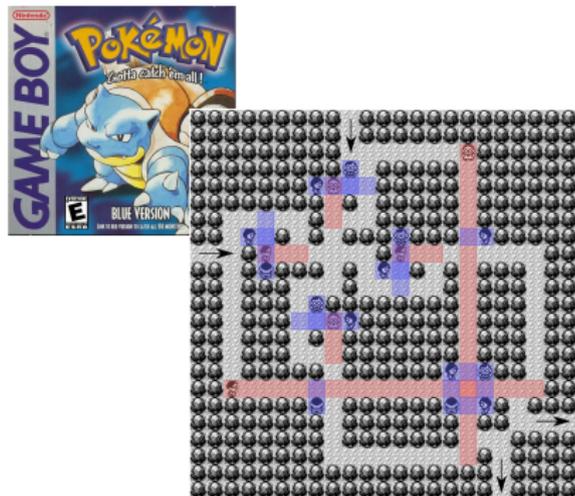
Deciding whether we can get from the start of a level to the goal is PSPACE-complete (Reduce from TQBF)



- ▶ If gadgets are like “pieces” of problems, where might these gadgets appear in something we’re familiar with?
- ▶ How about video games? (Specifically, game levels/objects)

Pokémon:

Deciding whether we can reach a given target location from a given start location is PSPACE-complete (Reduce from TQBF)



- ▶ If gadgets are like “pieces” of problems, where might these gadgets appear in something we’re familiar with?
- ▶ How about video games? (Specifically, game levels/objects)

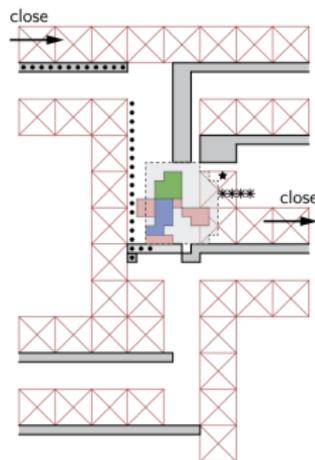
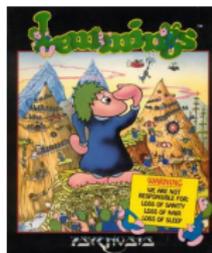
And many others:

- ▶ Donkey Kong Country 2 and 3
- ▶ Legend of Zelda
 - ▶ A Link to the Past
 - ▶ Link’s Awakening
 - ▶ Ocarina of Time
 - ▶ Majora’s Mask
- ▶ Metroid

- ▶ If gadgets are like “pieces” of problems, where might these gadgets appear in something we’re familiar with?
- ▶ How about video games? (Specifically, game levels/objects)

Lemmings:

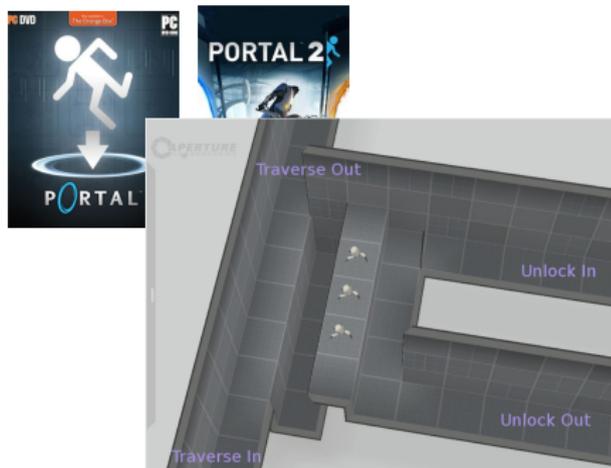
Deciding whether we can save one Lemming using Builders and Bashers is PSPACE-complete
(Even harder if we try to maximize the number of Lemmings saved!)



- ▶ If gadgets are like “pieces” of problems, where might these gadgets appear in something we’re familiar with?
- ▶ How about video games? (Specifically, game levels/objects)

Portal 1 and 2:

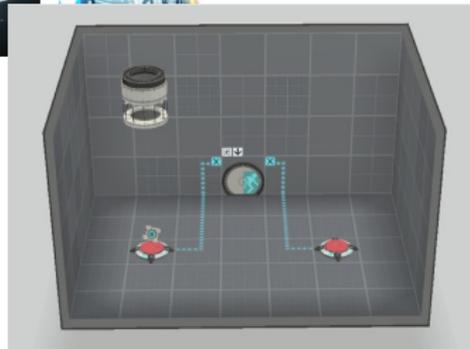
Deciding whether we can get from the start of a level to the goal is NP-hard (if we just have turrets)



- ▶ If gadgets are like “pieces” of problems, where might these gadgets appear in something we’re familiar with?
- ▶ How about video games? (Specifically, game levels/objects)

Portal 1 and 2:

Deciding whether we can get from the start of a level to the goal is PSPACE-complete (if we have cubes, weighted buttons, and doors)



Why are we learning TCS?

Traditional TCS courses

What do students think?

What do academics think?

We're *really* learning TCS...

... To know how to do something

... To understand what we can't do

... Because it's fun!

How do we make learning TCS better?

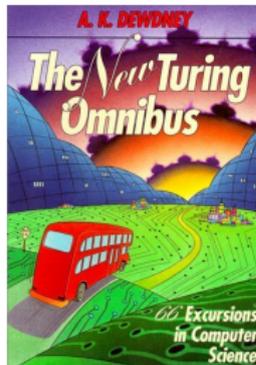
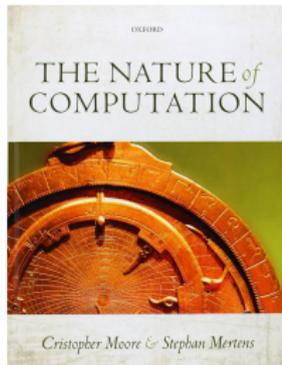
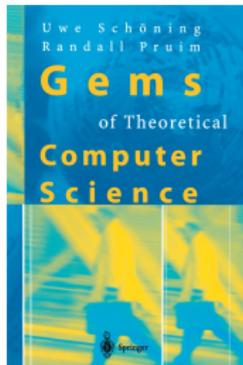
- ▶ Let's return to our two key takeaways:
 - ▶ We should appreciate TCS as a subject in itself
 - ▶ We should be mindful that TCS is more than just a bunch of theorems and proofs

- ▶ Let's return to our two key takeaways:
 - ▶ We should appreciate TCS as a subject in itself
 - ▶ We should be mindful that TCS is more than just a bunch of theorems and proofs
- ▶ The majority of TCS courses seem to lean toward the former direction.
 - ▶ I.e., the rigid “definition–theorem–proof” lecture style
- ▶ The majority of students seem to want to move in the latter direction.
 - ▶ “I still can't understand the real life application of this subject”
 - ▶ “something to contextualize what we're learning”
 - ▶ “add any practical examples or projects”

- ▶ Let's return to our two key takeaways:
 - ▶ We should appreciate TCS as a subject in itself
 - ▶ We should be mindful that TCS is more than just a bunch of theorems and proofs
- ▶ The majority of TCS courses seem to lean toward the former direction.
 - ▶ I.e., the **Why not both?** pure style
- ▶ The majority **Why not both?** ve in the latter direction.
 - ▶ “I still can't understand the real life application of this subject”
 - ▶ “something to contextualize what we're learning”
 - ▶ “add any practical examples or projects”

- ▶ Presenting the theorems, the proofs, etc., is inescapable in any TCS course. (Sorry, students.)
 - ▶ Abstraction and mathematical reasoning are vital skills for any computer scientist
 - ▶ It's important to get students comfortable with this mode of thinking early
- ▶ But we should couch these results in real-world applications and notions with which students are familiar.
 - ▶ Don't “meet the students where they are”, but meet them halfway
 - ▶ If you can present reductions and completeness in terms of Pokémon, *do it!*

Some inspirational literature



Thank you!



taylorjsmith.xyz taylorjsmith.xyz/flarelab/