Purpose of the Course

CS 70 is a new course designed to complement Math 55. We will focus on fewer topics, driven by computational tasks. We hope to make the course more relevant to CS students and hence to instill a deeper and longer-lasting understanding of the underlying mathematics.

What we want to teach:

- *Precise, reliable, powerful thinking*:
  will allow you to use and develop more complex and subtle ideas in CS, well beyond the obvious "brute force" approach to every problem, and will help you to avoid silly errors on all your CS final exams.
- *The ability to state and prove nontrivial facts, in particular about programs*:
  will enable you to write rigorously correct programs, which in turn provide solid building blocks for ever-more-complex yet still reliable systems;
- *Mathematical foundations and ideas useful throughout CS*:
  will provide familiarity with logic, inductively defined structures, integer and polynomial arithmetic, and probabilities—concepts that underly all of the more advanced courses in CS.

Course outline (abbreviated).

- Propositions and Proofs
- Mathematical Induction: recursion, the stable marriage problem
- Propositional Logic: automated proof and problem-solving
- Arithmetic Algorithms: gcd, primality testing, the RSA cryptosystem
- Polynomials and their Applications: error-correcting codes, secret sharing
- Probability and Probabilistic Algorithms: load balancing, hashing, probabilistic constructions, conditional probability Bayesian inference
- Diagonalization, Self-Reference, and Uncomputability

Propositions and Proofs

Many of the unenlightened believe proofs to be pointless formal exercises in guessing a way through a maze to reach a 2,400-year-old fortune cookie.

Far from it. We all like to say things:

"This encryption system cannot be broken"
"My program works efficiently in all cases"
"There are no circumstances under which I would lie to Congress"
"It is inconceivable that our legal system would execute an innocent person"

and so on. Few of us like to say things that turn out to be false. Proof means never having to say you're sorry—it provides a means for *guaranteeing* your claims once and for all.[1] What we would like to do now is to make these concepts more precise. (Most discrete mathematics courses just get on with doing proofs; this isn't a bad idea, but does skip over some important concepts.)

Proofs in mathematics and computer science (as opposed to law and politics) require a precisely stated proposition to be proved.

PROPOSITION

A **proposition** is a sentence that is either true or false.[2]

THEOREM

A **theorem**, informally speaking, is a proposition that is guaranteed by a proof.

Examples of propositions:

(1a) Some mammals lay eggs. (1b) Some mammals have feathers.
(2) The acceleration of a rigid body is proportional to the force applied.
(3) The angles of a triangle add up to 180 degrees.
(4) For all nonnegative integers n, n^2+n+41 is prime.
(5) For all integers n, if $n>2$ then there are no positive integers a, b, c such that $a^n+b^n=c^n$.
(6) For every even integer $n>2$, there are two primes a and b such that $a+b=n$.

WORLD
MODEL

It is important to note that every proposition is true or false *with respect to a possible world*. A **world** (or a **model** in logical terminology) is, conversely, something with respect to which every proposition of interest is either true or false—that is, it is completely specified.

For physics, chemistry, biology, etc., which are *empirical* sciences, we are usually interested in truth with respect to the *real* world—the one we actually live in. But of course, we don't know which one that is. For example, (1a) happens to be true in the real world, but could easily have been otherwise; whereas (1b) may or may not be true. [Exercise: what could we mean by this? Could one prove that (1b) is false?]

(2) is one of Newton's laws. It was assumed to be incontrovertibly true for many years, and many explanations were given for why it could not possibly be otherwise; but it is in fact false in the real world.[3]

AXIOMS

Now mathematicians, physicists, and engineers have been proving theorems in Newtonian mechanics for centuries and continue to do so. As a matter of physical fact, most of these theorems are false in the real world. They are, however, *still theorems* in Newtonian mechanics. An incorrect proof is not a proof, but a false theorem may still be a theorem *provided it follows logically from a specified set of* **axioms**. An axiom is a proposition that is assumed to be true without proof. In Newtonian mechanics, Newton's laws are taken as axiomatic.

PROOF

Proof, therefore, is a means of showing that a theorem follows logically from a set of axioms.

[1]"What about incorrect proofs?" you may ask. An incorrect proof is not a proof, any more than artificial grass is grass.

[2]Sentences that are not propositions include questions and commands—these cannot be true or false, although they can be perceptive or absurd.

[3]Many people state that Newton's laws were disproved by Einstein. This is not the case; Einstein merely proposed laws that are inconsistent with Newton's. Empirical laws such as Newton's can only be disproved by observations or by discovering an internal inconsistency.

Now we have to explain what is meant by "follows logically."

FOLLOWS LOGICALLY

> A proposition B **follows logically** from another proposition A if B is true in all possible worlds in which A is true.

This relationship is often written as $A \models B$, which can be pronounced "A logically entails B." If one draws a picture of the set of worlds where A holds and the set of worlds where B holds, the set where A holds will be contained within the set where B holds whenever $A \models B$. More mathematically, let $M(A)$ be the set of worlds where A holds and $M(B)$ the set of worlds where B holds (we call these worlds the **models** of A and B). Then

MODELS

$$A \models B \text{ if and only if } M(A) \subseteq M(B)$$

This relationship is shown in Figure 1(a). The direction of the $\subseteq$ may be somewhat counterintuitive: normally we think of A being "stronger" or "bigger" than B if A entails B. One route to reain intuition is to remember that every axiom added to A *reduces* the set of possible worlds where A holds, so makes it *more* feasible for this set to be contained within the set of B-worlds (Figure 1(b)).

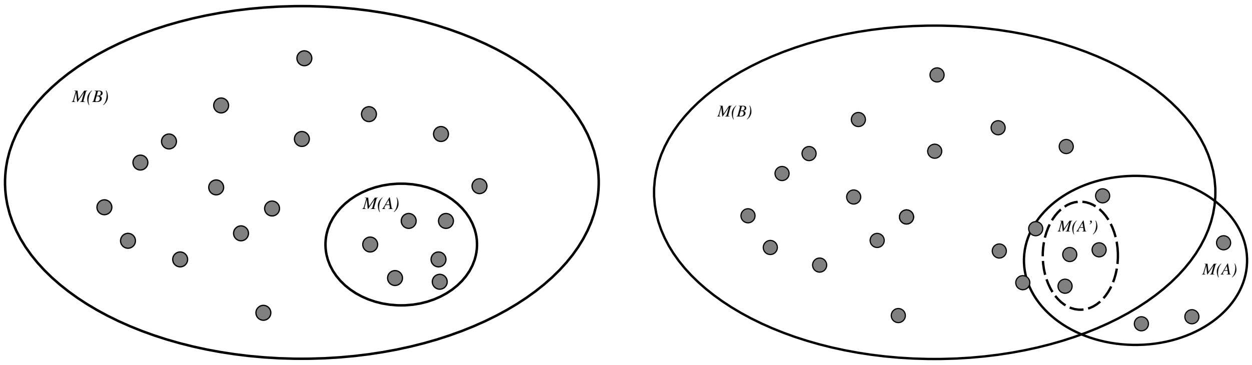

(a) (b)

Figure 1: (a) A entails B iff B is true in every world where A is true. (b) Adding axioms to A to make A' *reduces* the set of possible worlds; the figure shows a case where this allows A' to entail B.

We said earlier that a proof guarantees a proposition. More precisely, using the definition of logical entailment, we see that a proof guarantees the truth of a theorem to the extent that the axioms themselves are true. We will see in the next section some of the methods by which this guarantee is provided.

Let us return to the consideration of the propositions in our list above. Proposition (3), "The angles of a triangle add up to 180 degrees," is one of Euclid's axioms. It is simply postulated to be true. In fact, it is true only in planar geometry. A triangle inscribed on the surface of a sphere can violate the axiom; and general relativity says that space itself is curved, so geometries violating the axiom are actually more realistic. As with Newtonian mechanics, the theorems that Euclid (and many generations of schoolchildren after him) derived are true only in an idealized "flat" universe.

When we come to purely mathematical propositions, the situation is a little different: mathematical axioms are taken as *defining* the world under discussion rather than attempting to *describe* it. For example, Peano's axioms define what it means to be a natural number (nonnegative integer):

> 0 is a natural number
> If n is a natural number, $s(n)$ is a natural number

where $s(\cdot)$ is the *successor function*. We usually write $s(0)$ as 1, $s(s(0))$ as 2, and so on. From this simple beginning, we can build up more definitions—addition, multiplication, subtraction, division, prime numbers, and so on. Theorems in mathematics are (usually) true because the axioms of mathematics are (usually) true by definition.[4] Mathematical texts usually talk about proving that a proposition is true or false, which is really shorthand for "entailed by the standard axioms" or "inconsistent with the standard axioms." We will do the same thing, but we will try to be careful about citing the axioms that are required. Why do this? First, it's a good practice because it eliminates some errors that occur when one accidentally invokes a false axiom; second, it often reveals opportunities to prove a more general theorem because some axioms may not be required for the proof; and third, the axioms you're relying on may later turn out to be inconsistent, so it's good to keep a record. (This third possibility is *extremely* unlikely for most of the proofs we will do.)

At this point, we'll introduce a little bit more notation. Let us write proposition (4) as

(4) $\forall n\ n^2+n+41$ is prime

UNIVERSAL QUANTIFIER

The $\forall$ symbol is the **universal quantifier**; here, it binds the variable n, and means "for all n ..." Strictly speaking, the proposition should be written

(4) $\forall n \in \mathbf{N}\ n^2+n+41$ is prime

where **N** is the set of natural numbers. This qualification is too often omitted when the context makes it only somewhat obvious.

COUNTEREXAMPLE

How does one prove a universally quantified statement? We can check that it is true for lots of examples: $n=0$, $n=1$, and so on all the way up to $n=39$. Does this constitute a proof? Of course not! The proposition is false because $40^2+40+41$ is not prime. The case $n=40$ is called a **counterexample** for the proposition.

Proposition (5) is Fermat's last theorem. It has been known for over 300 years, and called a theorem for much of that time because Fermat claimed to have a proof and no counterexample was ever found. It recently *became* a proper theorem when Andrew Wiles developed a proof (several hundred pages long).

CONJECTURE

Proposition (6) is Goldbach's **conjecture**. A conjecture is a proposition that has not been proved or disproved. Using quantifier notation, it is written

$\forall n$ if n is even then $\exists a,b$ such that a and b are prime and $a+b=n$

EXISTENTIAL QUANTIFIER

Here, $\exists$ is the **existential quantifier**, meaning "For some ..." or "There exists ...." For any particular n, the existentially quantified statement can be proved simply by finding any particular a and b, whereas of course the universal quantification over n cannot be proved by exhibiting examples. It can be *disproved* by exhibiting a counterexample, but none has been found despite testing up to enormous values of n; we suspect the conjecture is true.

One may say, "Surely something so simple ought to be provable easily!" But Fermat's last theorem has turned out (so far) to have a very complex proof; and Kurt Gödel's famous Incompleteness Theorem showed that there are true propositions that have *no* proof in arithmetic. Unfortunately there is also no way to tell if Goldbach's conjecture is one of those.[5]

We have gone on long enough about propositions and proofs. Let's start proving some propositions.

[4] We say "usually" because occasionally a proposed set of axioms for some part of mathematics is shown to be inconsistent—that is, self-contradictory—and therefore *cannot* be true.

[5] Gödel's Incompleteness Theorem is not generally a good excuse for being unable to find a proof in a homework question.

P	Q	$\neg P$	$P \wedge Q$	$P \vee Q$	$P \implies Q$	$P \Leftrightarrow Q$
$False$	$False$	$True$	$False$	$False$	$True$	$True$
$False$	$True$	$True$	$False$	$True$	$True$	$False$
$True$	$False$	$False$	$False$	$True$	$False$	$False$
$True$	$True$	$False$	$True$	$True$	$True$	$True$

Table 1: Truth-table definitions of the standard logical operators.

Proof by enumeration

We begin with a very simple proof method based on the definition of logical entailment. Consider the following trivial example:

Given: Roses are red and violets are blue
Prove: Roses are red.

This is "obviously correct" because of the meaning of "and." A proposition "P and Q," which in mathematical notation is written $P \wedge Q$, is true just when P is true and Q is true. Thus, we can view "$\wedge$" as an operator that constructs a complex proposition called a **conjunction** out of two simple ones. This operator is defined by the truth value of the complex proposition for all possible truth values of its constituents, as shown in Table 1.

CONJUNCTION

To prove formally that roses are red (P), given that roses are red and violets are blue ($P \wedge Q$), we first identify all the possible cases where $P \wedge Q$ is true. There is just one such case, namely the fourth line of the table. And indeed, in that case, "roses are red" (P) is also true. This completes the proof.

Duh. This seems rather daft, but actually it illustrates very well the fundamental concept of formal proof by enumeration of possible cases. We will see later in the course that the same basic idea can be instantiated in a program that performs feats of deduction well beyond the powers of human beings.

Proof by enumeration is tedious beyond belief, so we'll look at only one more case:

Given: If Dewey isn't elected, then I'll eat my hat
Dewey isn't elected
Prove: I'll eat my hat

IMPLICATION

ANTECEDENT
CONSEQUENT

The first sentence here has the form of an **implication**. In mathematical notation, this is written as $P \implies Q$. The truth table for "$\implies$" is given in Table 1. Notice that the implication $P \implies Q$ is only false in the case where its **antecedent** P is true and its **consequent** Q is false. The implication is *true* in those cases where the antecedent is *false*. This is fine, because in such cases the implication simply doesn't apply.

Proof: First, we identify all those cases where $P \implies Q$ and P are both true. There is just one such case, namely the fourth line of the table. And indeed, in that case, "I'll eat my hat" (Q) is also true. $\square$

Notice the $\square$ marking the end of a proof.[6] Some people like to read texts and papers just looking at the parts between the "**Proof**:" and the $\square$, regarding the ordinary text as useless filler. Other people do the exact opposite, regarding the proof text as tedious detail. You can usually guess which category a person belongs to.

[6] In better days, people wrote QED instead, standing for for *quod erat demonstrandum*—Latin for *which was the thing to be demonstrated*.

Proof by application of inference rules

Notice that, once we have done a proof by enumeration, we can extract a general pattern called an **inference**
INFERENCE RULE **rule**. The two patterns in the preceding section are

For any propositions P and Q, the proposition P can be inferred from the proposition $P \wedge Q$.

For any propositions P and Q, the proposition Q can be inferred from the propositions P and $P \implies Q$.

These can also be written using the following notation:

$$\frac{P \wedge Q}{P} \qquad \frac{P,\ P \implies Q}{Q}$$

AND-ELIMINATION MODUS PONENS The former rule is called **and-elimination**. The latter rule is called **modus ponens** (Latin for "placing method") and is one of the most common steps used in proofs. There are many more such rules, some of which we will see in subsequent sections.

Rules can be chained together. For example, if we know $A \wedge B$ and $B \implies C$, we can apply and-elimination to derive B and then modus ponens to derive C. A full truth-table proof would require eight rows to handle the three propositions.

Proof by contrapositive

Consider the propositions

If John is at work, he's logged in.
If John isn't logged in, he's not at work.

LOGICALLY EQUIVALENT Clearly, each of these propositions can be proved from the other. They are **logically equivalent** propositions—that is, their truth values are the same in all possible worlds. Writing the first as $P \implies Q$ and the second as
NEGATION $\neg Q \implies \neg P$ (where we have used the **negation** symbol $\neg$), it is easy to verify this equivalence using truth tables. (We will see many such equivalences in later lectures.)

CONTRAPOSITIVE The proposition $\neg Q \implies \neg P$ is the **contrapositive** of $P \implies Q$. (This is not to be confused with the
CONVERSE **converse**, which is $Q \implies P$ and is *not* equivalent.) Proof by contrapositive (also called indirect proof) means proving that $P \implies Q$ by proving $\neg Q \implies \neg P$ instead. Since the two are equivalent, proving one of them automatically establishes the other. Here is a very simple example:

Theorem 1.1: *For any integer n, if n^2 is even then n is even.*

Proof: We will prove the contrapositive: if n is odd then n^2 is odd.

1. If n is odd, then (by definition) $n = 2a + 1$ for some integer a.
2. For any numbers x and y, we know that $x = y \implies x^2 = y^2$. Hence
$$n^2 = (2a+1)^2 = 4a^2 + 4a + 1 = 2(2a^2 + 2a) + 1$$
3. Since a is an integer, $(2a^2 + 2a)$ is an integer.
4. Hence (by the definition of oddness), n^2 is odd.

□

Non-proof

Failure to note the justification for each step can lead easily to non-proofs. Consider the following example.

Theorem 1.2: *(not!)* $1 = -1$

Proof: $1 = \sqrt{1} = \sqrt{(-1)(-1)} = \sqrt{-1}\sqrt{-1} = \sqrt{-1}^2 = -1$ $\square$

Presumably, at least one of these steps is false. But each one (presumably) looked "reasonable" to the author of the proof. Writing out the full justifying axioms for each step quickly reveals an axiom that is false: it is simply untrue that, for any numbers x and y, $\sqrt{xy} = \sqrt{x}\sqrt{y}$.

Other classic errors:

- Dividing both sides of an equation by a variable. For example:

  $$ax = bx \text{ hence } a = b$$

  The "axiom" to which this step implicitly appeals is false, because if $x = 0$ the claim $a = b$ does not follow. Some extra work may be needed to prove $x \neq 0$.

- Dividing both sides of an inequality by a variable. This is even worse! For example:

  $$ax < bx \text{ hence } a < b$$

  Here the claim $a < b$ is false if $x < 0$, and unproven if $x = 0$.

Proof by cases

Sometimes we don't know which of a set of possible cases is true, but we know that at least one of the cases is true. If we can prove our result in each of the cases, then we have a proof. The English phrase "damned if you do and damned if you don't" sums up this proof method. Here's a nice example:

Theorem 1.3: *For some irrational numbers x and y, x^y is rational.*

Proof:

1. Since the theorem is existentially quantified, we need only prove the existence of at least one example. Consider the case $x = \sqrt{2}$ and $y = \sqrt{2}$. It must be true that

   (a) $\sqrt{2}^{\sqrt{2}}$ is rational

   or (b) $\sqrt{2}^{\sqrt{2}}$ is irrational.

2. In case (a), we have shown irrational numbers x and y such that x^y is rational.

3. In case (b), consider the values $x = \sqrt{2}^{\sqrt{2}}$ and $y = \sqrt{2}$. We have

   $$\begin{aligned} x^y &= (\sqrt{2}^{\sqrt{2}})^{\sqrt{2}} \\ &= \sqrt{2}^{\sqrt{2}\sqrt{2}} \text{ by the axiom } (x^y)^z = x^{yz} \\ &= \sqrt{2}^2 = 2 \end{aligned}$$

   Hence we have shown irrational numbers x and y such that x^y is rational.

4. Since one of cases (a) and (b) must be true, it follows that for some irrational numbers x and y, x^y is rational.

□

NONCONSTRUCTIVE

Notice that even after the proof, we still don't know which of the two cases is true, so we can't actually exhibit any irrational numbers satisfying the theorem. This is an example of a **nonconstructive** proof, one in which an existential theorem is proved without constructing an example.

Proof by contradiction

Also called *reductio ad absurdum* (reduction to an absurd thing), proof by contradiction is closely related to proof by contrapositive. The idea is to assume the opposite of what one is trying to prove and then show that, when combined with the axioms, this leads to a contradiction.

RATIONAL NUMBER

REDUCED FORM

Consider the following example. A **rational number** is a number that can be expressed as the ratio of two integers. For example, $2/3$, $3/5$, and $9/16$ are all rationals. The **reduced form** of a rational is a fraction in which the numerator and denominator share no factors other than 1. For example, the reduced form of $3/6$ is $1/2$. Note that any number with a finite or recurring decimal representation is a rational.

Theorem 1.4: $\sqrt{2}$ *is irrational.*

Proof:

1. Assume $\sqrt{2}$ is rational.
2. By the definition of rational numbers, there are integers a and b with no common factor other than 1, such that $\sqrt{2} = a/b$.
3. For any numbers x and y, we know that $x = y \implies x^2 = y^2$. Hence $2 = a^2/b^2$.
4. Multiplying both sides by b^2, we have $a^2 = 2b^2$.
5. b is an integer, hence b^2 is an integer, hence a^2 is even (by the definition of evenness).
6. Hence, by the theorem proved earlier, a is even.
7. Hence, by the definition of evenness, there is an integer c such that $a = 2c$.
8. Hence $2b^2 = 4c^2$, hence $b^2 = 2c^2$.
9. Since c is an integer, c^2 is an integer, hence b^2 is even.
10. Hence, by the theorem proved earlier, b is even.
11. Hence a and b have a common factor 2, contradicting step 1.
12. Hence, step 1 is false, i.e., $\sqrt{2}$ is irrational.

□

Style and substance in proofs

Our proofs justify each step by appealing to a definition or general axiom. The depth to which one must do this in practice is a matter of taste. For example, we could break down the step, "Since a is an integer, $(2a^2 + 2a)$ is an integer," into several more steps. [Exercise: what are they?] A justification can be stated without proof only if you are absolutely confident that (1) it is correct and (2) the reader will automatically agree that it is correct.

Notice that in the proof that $\sqrt{2}$ is irrational, we used the earlier result, "For any integer n, if n^2 is even then n is even," twice. This suggests that it may be a useful fact in many proofs. A subsidiary result that is useful in a more complex proof is called a **lemma**. It is often a good idea to break down a long proof into several lemmata. This is similar to the way in which large programming tasks should be divided up into smaller subroutines. Furthermore, make each lemma (like each subroutine) as general as possible so it can be reused elsewhere.

LEMMA

The dividing line between lemmata and theorems is not clear-cut. Usually, when writing a paper, the theorems are those propositions that you want to "export" from the paper to the rest of the world, whereas the lemmata are propositions used in the proofs of your theorems. There are, however, some lemmata (for example, the Pumping Lemma and the Lifting Lemma) that are perhaps more famous and important than the theorems they were used to prove.

For further reading on proofs and proof style, see the paper by De Millo, Lipton, and Perlis, "Social processes and proofs of theorems and programs," In *Communications of the ACM*, May 1979, vol.22, , pp. 271–80. Also Don Knuth's *Mathematical Writing*, Mathematical Association of America, 1989.