Logic 3

Today:

- ☐ Doing stuff with circuits
- ☐ Tautologies
- ☐ Formalizing proofs — an important theorem
- ☐ Adding quantifiers — first-order logic
- ☐ The math/English gap
- ☐ Negating quantified formulas

1 Doing Stuff with Circuits

We already know that an arbitrary boolean function $f\colon \mathbb{B}^n \to \mathbb{B}$ can be realized with a limitless supply of two-input gates—indeed we know that just having AND, OR, and NOT gates would suffice. So would having just NAND gates. It follows that an arbitrary boolean function $F\colon \mathbb{B}^n \to \mathbb{B}^m$ can likewise be realized with a limitless supply two-input gates, as long as you start with a set that's functionally complete. Yet the generic way of realizing an arbitrary boolean function $F\colon \mathbb{B}^n \to \mathbb{B}^m$ might be totally impractical, because our generic method—to transform the truth table into DNF—could use more than $mn2^n$ gates. If we have a practical design problem, we may need to do way better than that. Here is an example showing how we often can.

Suppose you are given two 64-bit binary numbers, $a_{63}a_{62}\dots a_1a_0$ and $b_{63}b_{62}\dots b_1b_0$, which we would like to add to get a 65-bit sum, $s_{64}s_{63}s_{62}\dots s_1s_0$. The DNF approach would ask us to make 65 circuits (one for each output bit), each of size about 2^{128}. You can forget doing that!

We can do far better by following grade-school arithmetic, adapted to base-2. This is a standard example, which I will work out in class. See zyBook 1.10, the Wikipedia page "adders (electronics)", or the PowerPoint pictures I just drew.

2 Tautologies

Now for a different problem. Suppose I give you a formula $\phi(x_1,\dots,x_n)$ and I ask you: *is it true*? The customary answer would be: it depends! Namely, it depends on the truth assignment. For most formulas, some truth assignments will make them true and some will make them false.

But not every formula is really a mixture like that. Some formulas ϕ are *always* true. Some are always false. And some are sometimes true and sometimes false.

A boolean formula that is always true—it is true for *every* truth assignment—is called a *tautology*. We write $\models \phi$ to indicate that ϕ is a tautology. I've heard the $\models$ symbol called a "double turnstile."

Figure 1: A possible list of axioms that, together with *modus ponens*, suffice to establish completeness: that which is true will admit a proof.

A boolean formula that is sometimes true—it is true for *some* truth assignment—is called a *satisfiable* formula. There's no special symbol for it that I know.

Can you look at a formula ϕ and figure out if it's a tautology? The natural approach is to make a truth table for it. If you get a 1 (true) for each and every row, the formula is a tautology. If you see even one 0 (false), then it's not.

The same approach works to figure out if a formula ϕ is satisfiable: make a truth table and check if there's at least one "true".

The methods above are very inefficient. If the formula has 100 variables than your truth table will have 2^{100} rows to check out, which is not remotely practical. But maybe there is a way to *prove* that a formula is true *without* having to construct a huge truth table.

Before taking up proofs, let me relate tautologies to equivalence. When we write that $\phi \equiv \psi$ we are asserting that $\phi \leftrightarrow \psi$ is a tautology, which we could write as $\models \phi \leftrightarrow \psi$. Be careful not to say that $\phi \equiv \psi$ mean $\phi \leftrightarrow \psi$. The latter formula can be true or false, but what we are trying to say is that $\phi \equiv \psi$ when $\phi \leftrightarrow \psi$ is *always* true—that is, when it is a tautology.

3 Formalizing Proofs — Two Famous Theorems

A *proof system* is a formalized way to prove statements in logic. There are numerous proof systems, but let's just sketch how one of them looks. The proof itself is a formal object—just a sequence of lines, each of which is a string. You might want to number them. Each line may be one of three things:

1. You can write an assumption, which is an arbitrary boolean formula.

2. You can write an axiom from a given list of axioms. The list of Figure 1 is an example. The list will be described using symbols such as ϕ or ψ. These represent arbitrary boolean formulas. When you copy an axiom from the list you are welcome to substitute in any boolean formula you like in place of each of those variables.

3. Finally, if a line above you in the proof has a formula ϕ, and another line above you has a formula $\phi \to \psi$, then you are welcome to enter into the current line the formula ψ. If you like, you may whisper the magic incantation *modus ponens* when you do this.

At some point you declare your proof to be done. Suppose it's final line is the boolean formula ψ. Then we would say that you have just proven ψ from assumptions $\phi_1, \ldots, \phi_n$, where $\phi_1, \ldots, \phi_n$ are all the formulas you entered into your proof as assumptions. We would write this as $\{\phi_1, \ldots, \phi_n\} \vdash \psi$. You may alternatively write $\vdash (\phi_1 \wedge \cdots \wedge \phi_n) \to \psi$.

In class I will work out an example. But I don't actually expect you to do any example of formal proofs. Because, honestly, they are extremely tedious.

It was a dream of 19th century mathematics that complex mathematical claims be produced, or at least verified, with some sort of formal proof. A famous book, *Principia Mathematica*, was published in 1910–1913 by Alfred Whitehead and Bertrand Russell. It spanned three volumes and only managed to prove some basic facts in arithmetic. And even there there were errors, I understand. Formal proofs—at least before the age of machine-assisted ones—are not what mathematicians routinely do.

All the same, the *idea* of a formal proof is fundamental in mathematics, as in the following theorem, first established by Kurt Gödel for a specific proof system.

Theorem 1 *Let ϕ be a boolean formula. Then*

- **Soundness:** *If $\vdash \phi$ then $\models \phi$.*
- **Completeness:** *If $\models \phi$ then $\vdash \phi$.*

It is often stated more generally as follows. Let ϕ be a boolean formula and let Γ be a set of boolean formulas. Then

- **Soundness:** *If $\Gamma \vdash \phi$ then $\Gamma \models \phi$.*
- **Completeness:** *If $\Gamma \models \phi$ then $\Gamma \vdash \phi$.*

When we write $\Gamma \vdash \phi$ it means that ϕ can be proven from assumptions in Γ. When we write $\Gamma \models \phi$ it means that every truth assignment that makes every formula in Γ true also makes ϕ true. The first form is a special case of the second, where the $\Gamma = \emptyset$.

This is a deep theorem. It is saying that we can come up with simple proof systems that are so powerful that everything that is true about sentential logical can be proven, while nothing that is untrue can be proven. And, in fact, its not only for sentential logic. The statement holds for first-order logic, which we are about to introduce.

I remember learning this theorem as an undergraduate at UCB in a class by a famous logician, Leon Henkin (1921–2006). I don't remember the proof, but I remember that it felt extremely tedious, and took several lectures. In fact, proving the theorem was pretty much the main goal of the class. I think I didn't appreciate the project back then—and maybe how ambitious Henkin was being to do this in a lower-division philosophy class. I suspect that I didn't give Henkin's program the respect it deserved.

4 Adding Quantifiers — First Order Logic

Our predicate calculus is rather limited, incapable of capturing things like "the sum of odd numbers is even" or "there are an infinite number of primes." To encode such statements we need to substantially embellish our propositional logic.

We start by imagining some universe $\mathcal{U}$ of points—the *things* we are discussing. The points in the universe $\mathcal{U}$ might be sets, integers, real numbers, people, animals, colors, lions named "Tony"—we don't much care.

We will have *variables*: things like x or y_3 that represent arbitrary points from our universe.

We will have *predicates*: things that take in some number of points from our universe and return true or false. Like: PRIME(n), perhaps to mean that n is a prime. Or $n < m$, perhaps meaning that the number n is less than the number m.

We will have *functions*: things that take in some number of points from our universe and return a new point in our universe. Things like + or square-root.

We will have *constant*: symbols that represent specific points in our universe. Like 0 or $\emptyset$ or BILLY.

We will have two spacial symbol, $\forall$ and $\exists$, which are the *universal quantifier* and the *existential quantifier*. When we write $(\forall x)(\cdot)$ we will be asserting that the claim in the parentheses holds for any x in our universe $\mathcal{U}$. When we write $(\exists x)(\cdot)$ we will be asserting that the claim in the parentheses holds for some x in our universe $\mathcal{U}$.

Let us give some fanciful example.

1. "All apples are bad." We might render this as

   $$(\forall x)(A(x) \to B(x)).$$

   Here $A(x)$ is supposed to mean that x is an apple, while $B(x)$ is supposed to mean that x is something bad. What is our universe $\mathcal{U}$? Who knows. Maybe all fruit. Maybe all objects. It's not something that the symbols tell us.

2. "Billy has beat up every boy at Caesar-Chavez elementary school." (Billy should not have done that.) Perhaps we would render this as something like

   $$(\forall x)((\text{CCstudent}(x) \wedge \text{Boy}(x) \wedge (x \neq \text{BILLY})) \to \text{HasBeatenUp}(\text{BILLY}, x)).$$

Here we have a constant symbol BILLY and predicates CCstudent, Boy, and HasBeatenUp, the first two being unary predicates and the last being a binary predicate. Perhaps the universe of discourse is Davis school children. People it is people living on Earth. Who knows.

3. "All lions are fierce." This might be translated to $(\forall x)(L(x) \rightarrow F(x))$. "Some lions do not drink coffee." This might be translated to $(\exists x)(L(x) \wedge \neg C(x))$. "Some fierce creatures do not drink coffee." We might write this as $(\exists x)(F(x) \wedge \neg C(x))$. Does the conjunct of the first two statements imply the third? In answering something like this, you do not want to rely on your knowledge of lions and coffee. It should hold regardless of our universe and the meaning of predicates L, F, and C.

4. Last non-mathematical example: "Nobody like a sore loser." Perhaps our universe of discourse is all human beings. Perhaps we have a predicate $L(x, y)$ if person x likes person y. And another predicate $S(x)$ if x is a sore loser. In such a case, we might try translating our English sentence into $(\forall x)(S(x) \rightarrow (\forall y)(\neg L(y, x)))$. In this translation the sore loser doesn't even like himself.

We are not going to formally define the syntax for first-order logic, but doing it wouldn't be so different from how we defined the syntax for boolean formulas. We would have logical connectives consisting of (,), $\neg$, $\wedge$, $\vee$, and maybe other logical connectives like $\rightarrow$ and $\leftrightarrow$. It is customary to include the equality symbol, $=$, treating it differently from other binary predicates. We would have the symbols $\forall$ and $\exists$. There would be a set of variables, such as $x_1, x_2, \ldots$; a finite set of constant symbols; a finite set of function symbols; and a finite set of predicate symbols. You get to fix all of these choices when you are working in first-order logic.

We can do some really interesting things with first order logic, like defining group theory or set theory. Before we get there, though, a couple things I'd like to do. One is to call out the fiction that logic is good for dealing with "logical" English-language utterance. I'd also like to talk about how to negate quantified assertions.

5 The Math/English Gap

It is common to introduce propositional logic by telling students that we are formalizing a way to capture and then reason about the veracity of English language proposition—statements like "it is raining today," or "all politicians are crooks." From "All men are mortal" and "Socrates is a man" we will be concluding that "Socrates is mortal." That sort of thing. But I think that this desire to meld logic and natural language is misguided. And not just for the reasons usually cited, but because natural-language utterances are almost *never* used to communicate logical claims.

The usual complaints begin by noting that the English language *or* and *if ... then* often don't translate to the mathematical OR and IMPLIES. For example, if you go to a wedding and the host asks if you would like the vegetarian or the meat entrée, you ought not to answer *yes* unless you are trying to be cute. Because the *or* was asking for you to select one of those alternatives. Sometimes it is said that this particular "or" was really VEG xor MEAT. But that's not true, either. If the host had offered meal choices that were vegetarian, beef, of fish, you would still have been cheeky to select all three (after all, the xor of three ones is also one). The intent for this "or" would seem to be to select exactly one from the list, regardless of the list's length. That's unlike a logical-or *and* unlike a logical-xor.

And the problems really run deeper. Propositions, many books teach, are factual assertions about the world, like "it is raining," or "lions are fierce". But do such statements actually have clear, incontestable truth value? They do not. It is raining *where*? It is raining *when*? And how specific is this place and time to be? If we shrink Alice to a micrometer in height then it may pour within a foot of her, yet not a drop will fall upon her pretty little head. If we shrink Alice's lifespan to a microsecond, then it could be pouring from the point of view of *my* timescale, but highly unlikely to bother Alice within hers. And even if we do not play with the scale of time or space, surely rainfall happens within a continuum, from a scorchingly hot and sunny day to a robust Indian monsoon. I have no idea where a wet-fog ends and a light rain begins. Not just because I am limited, but because such distinctions are invariably artificial.

You could argue that these are specious complaints because people generally agree as to when and where it is and is not raining. But math isn't about what people do or don't agree to. And if the only English-language statements that we are comfortable to assign truth values to are the mathematical ones, then we have relinquished any claim that logic is meaningfully connected to human discourse.

In natural languages, logical connectives like *or* and *if* ... *then* are usually structural claims about our social reality. If I tell you

> Either Facebook lives or liberal democracy dies.

then I am expressing a belief (a valid one, I think) that a certain pair of things are incompatible not for *logical* reasons but because of the social, political, structural reality that surrounds them. Were I to instead say

> Either Facebook lives or the square root of ten exceeds three.

then you would not be wrong to brand me as a lunatic, although the logical truth of this statement should be less contestable than before. Similarly, the sentence

> If the forecast calls for rain then I will grab my umbrella.

would seem to most people to make sense, while

> If the forecast calls for rain then I am older than my son.

seems ludicrous.

There was a time when philosophers imagined logic to be entwined with language and human reasoning. I think we have outgrew such views. But I'm not sure that textbook authors, or professors, got the memo.

If a book you are reading makes claims or connections that seem suspect, don't allow the authority of authorship—or professorship!—to blunt your skepticism or your critical thinking.

6 Negating Quantified Formulas

A simple point: if you want to negate a quantified logical expression, how do you do this? It is quite simple:

$$\neg(\forall x)\phi \equiv (\exists x)(\neg\phi)$$

and

$$\neg(\exists x)\phi \equiv (\forall x)(\neg\phi).$$

Combine this with De Morgan's laws, and rules like $P \to Q \equiv \neg P \vee Q$, and you can negate anything you like. I'll do some examples in class, or a TA will do some examples in discussion section.