


3 Logic



In which our heroes move carefully through the marsh, making sure that each step follows safely from the one before it.

3-2 Logic

3.1 Why You Might Care

The thinker must think for truth, not for fame.

W. E. B. Du Bois (1868–1963)

The Souls of Black Folk (1903)

Logic is the study of truth and falsity, of theorem and proof, of valid reasoning in any context. It's also the foundation of all of computer science, the very reasoning that you use when you write the condition of an if statement in a Java program, or when you design an algorithm to beat a grandmaster at chess. More concretely, logic is also the foundation of all *computers*. At its heart, a computer is a collection of carefully arranged wires that transport electrons (which serve as a physical manifestation of information) and “gates” (which serve as physical manifestations of logical operations to manipulate those electrons).

Indeed, the processor of a computer is essentially a circuit built up from almost unthinkably simple logical components: nothing much more than wires and the physical implementations of operations like “and,” “or,” and “not.” (For example, the very same reasoning that allows you to say that $x+y$ is odd exactly when either x is odd and y not odd, or when x is not odd and y is odd, allows you to design a circuit to add two bits. And stringing together 64 copies of that small circuit—with some additional logic to handle “carrying” from one place to the next—allows you to build a circuit that adds two 64-bit numbers.)

Because logic is the study of valid reasoning, any endeavor in which one wishes to state and justify claims rigorously—such as that of this book—must at its core rely on logic. Every condition that you write in a loop is a logical statement. When you sit down to write binary search in Python, it is through a (perhaps tacit) use of logical reasoning that you ensure that your code works properly for any input. When you use a search engine to look for web pages on the topic BEATLES AND NOT JOHN OR PAUL OR GEORGE OR RINGO you've implicitly used logical reasoning to select this particular query. (Maybe you just heard the reason for the band's name for the first time, and you wanted to look it up.) Solving a Sudoku puzzle is nothing more and nothing less than following logical constraints to their conclusion. The central component of a natural language processing (NLP) system is to take an utterance by a human user that's made in a “natural” language like English and “understand” what it means—and understanding what a sentence means is essentially the same task as understanding the circumstances under which the sentence is true, and thus is a question of logic. And these are just a handful of examples; for a computer scientist, logic is the basis of the discipline.

Our main goal in this chapter will be to introduce the basic constructs of logic. We'll focus on *formal logic*, in which it is the “form” of the argument that matters, rather than the “content.” This chapter will introduce the two major types of formal logic. The first type is *propositional logic* (Sections 3.2 and 3.3), in which we will study the truth and falsity of statements, how to construct logical statements from basic logical operators (like “and” and “or”), and how to reason about those statements. The second type of logic

3.1 Why You Might Care 3-3

that we'll examine is *predicate logic* (Sections 3.4 and 3.5), which gives us a framework to write logical statements of the form “every x . . .” or “there’s some x such that . . .”

One of our main goals in this chapter will be to define a precise, formal, and unambiguous language to express reasoning—in which writer and reader agree on what each word means. But along the way, we will encounter applications of logic to natural language processing, circuits, programming languages, optimizing compilers, and building artificially intelligent systems to play chess and other games.

3-4 Logic

3.2 An Introduction to Propositional Logic

Logic is a large drawer, containing some useful instruments, and many more that are superfluous. A wise man will look into it for two purposes, to avail himself of those instruments that are really useful, and to admire the ingenuity with which those that are not so, are assorted and arranged.

Charles Caleb Colton (1777–1832)

Lacon: or, Many things in few words; addressed to those who think (1820)

A *proposition* is a statement that is either true or false—A *modern smartphone has over 100,000 times more memory than the computer on the Apollo moon-landing mission* (true, by a long shot) or *Java is a programming language that uses indentation to denote block structure* (false), for example. *Propositional logic* is the study of propositions, including how to formulate statements as propositions, how to evaluate whether a proposition is true or false, and how to manipulate propositions. The goal of this section is to introduce propositions and propositional logic.

3.2.1 Propositions and Truth Values

We'll begin, briefly, with propositions themselves:

Definition 3.1: Propositions and truth values.

A *proposition* is a statement that is either true or false. For a particular proposition p , the *truth value* of p is its truth or falsity.

A proposition is also sometimes called a *Boolean expression* or a *Boolean formula*. (See Section 2.2.1.) A proposition is written in English as a declarative sentence, the kind of sentence that usually ends with a period. (Questions and demands—like *Did you try binary search?* or *Always comment your code!*—aren't the kinds of things that are true or false, and so they're not propositions.) Here are a few examples:

Example 3.1: Some sample propositions.

The following statements are all propositions:

- 1 $2 + 2 = 4$.
- 2 33 is a prime number.
- 3 Barack Obama was the 44th person to serve as the president of the United States.
- 4 Every even integer greater than 2 can be written as the sum of two prime numbers.

(The last of these propositions is called *Goldbach's conjecture*; it's more complicated than the other propositions in this example, and we'll return to it in Section 3.4.)

Example 3.2: Determining truth values.

Let's determine the truth values of the propositions from Example 3.1. They are:

- 1 **True.** It really is the case that $2 + 2$ equals 4.
- 2 **False.** The integer 33 is not a prime number because $33 = 3 \cdot 11$. (Prime numbers are evenly divisible only by 1 and themselves; 33 is also evenly divisible by 3 and 11.)
- 3 **False.** Although Barack Obama is called president #44, Grover Cleveland was president #22 and #24. So Barack Obama was actually the 43rd person to be president of the United States, not the 44th.
- 4 **Unknown (!).** Goldbach's conjecture was first made in 1742, but has thus far resisted proof—or dis-proof. It's easy to check that a particular small even integer can be written as the sum of two prime numbers; for example, $4 = 2 + 2$ and $6 = 3 + 3$ and $8 = 3 + 5$ and $10 = 3 + 7$. But is it true for *all* even integers greater than 2? We simply don't know! Many even integers have been tested, and no violation has been found in any of these tests. But maybe it will turn out that the very next even integer that someone tests *can't* be written as the sum of two primes. (See Example 3.49.)

Before we move on from Example 3.2, there's an important point to make about statements that have an unknown truth value. Even though we don't *know* the truth value of Goldbach's conjecture, it is still a proposition and thus it *has* a truth value. That is, Goldbach's conjecture is indeed either true or false; it's just that we don't know which. (Like *The person currently sitting next to you is wearing clean underwear*: it has a truth value, you just don't know what truth value it has.)

Taking it further: Goldbach's conjecture, with its unknown truth value, stands in contrast to declarative sentences whose truth is ill-defined—for example, *This book is boring* and *Logic is fun*. Whether these claims are true or false depends on the (imprecise) definitions of words like *boring* and *fun*. “Shades of truth” questions are not the prime focus when you're studying precise, formal reasoning, as in this book, but they're important when considering the role of ambiguity in software systems that interact with humans via English language input and output. See p. 3-17.

There is also a potentially interesting philosophical puzzle that's hiding in questions about the truth values of natural-language utterances. Here's a silly (but obviously true) statement: *The sentence “snow is white” is true if and only if snow is white.* (Of course!) This claim becomes a bit less trivial if the embedded proposition is stated in a different language—Spanish or Dutch, say: *The sentence “La nieve es blanca” is true if and only if snow is white*; or *The sentence “Sneeuw is wit” is true if and only if snow is white*. But there's a troubling paradox lurking here. Surely we would like to believe that the English sentence x and the French translation of the English sentence x have the same truth value. For example, *Snow is white* and *La neige est blanche* surely are both true, or they're both false. (And, in fact, it's the former.) But this belief leads to a problem with certain self-referential sentences: for example, *This sentence starts with a ‘T’* is true, but *Cette phrase commence par un ‘T’* is, surely, false. For more on paradoxes and puzzles of translation, see, for example, [57, 111].

3.2.2 Atomic and Compound Propositions

We will distinguish between two types of propositions, those that cannot be broken down into conceptually simpler pieces and those that can be:

3-6 Logic

Definition 3.2: Atomic and compound propositions.

An *atomic proposition* is a proposition that is conceptually indivisible. A *compound proposition* is a proposition that is built up out of conceptually simpler propositions.

Here's an example of the difference:

Example 3.3: Sample atomic and compound propositions.

Consider these two propositions:

(1) *The University of Minnesota's mascot is the Badger.*

(2) *The University of Washington's mascot is the Duck or the University of Oregon's mascot is the Duck.*

The first is an atomic proposition, because it is not conceptually divisible into any simpler claim. The second is a compound proposition, because it is conceptually divisible into two simpler claims—namely *The University of Washington's mascot is the Duck* and *The University of Oregon's mascot is the Duck*.

Atomic propositions are also sometimes called *Boolean variables*; see Section 2.2.1. A compound proposition that contains Boolean variables p_1, p_2, \dots, p_k is sometimes called a *Boolean expression* or *Boolean formula over p_1, p_2, \dots, p_k* .

Example 3.4: Password validity as a compound proposition.

A certain college sends the following instructions to its users when it makes them change their password:

Your password is valid only if it is at least 8 characters long, you have not previously used it as your password, and it contains at least three different types of characters (lowercase letters, uppercase letters, digits, non-alphanumeric characters).

This compound proposition involves seven different atomic propositions:

p = the password is valid

q = the password is at least 8 characters long

r = the password has been used previously by you

s = the password contains lowercase letters

t = the password contains uppercase letters

u = the password contains digits

v = the password contains non-alphanumeric characters

The form of the compound proposition is “ p , only if q and not r and at least three of $\{s, t, u, v\}$ are true.” (Later we'll see how to write this compound proposition in standard logical notation; see Example 3.16.)

3.2.3 Logical Connectives

Logical connectives are the glue that creates the more complicated compound propositions from simpler propositions. Here are definitions of our first three of these logical connectives—*not*, *and*, and *or*:

3.2 An Introduction to Propositional Logic 3-7

Definition 3.3: Negation [not, \neg].

The proposition $\neg p$ (“not p ,” called the *negation* of the proposition p) is true when the proposition p is false, and is false when p is true.

Definition 3.4: Conjunction [and, \wedge].

The proposition $p \wedge q$ (“ p and q ,” the *conjunction* of the propositions p and q) is true when both of the propositions p and q are true, and is false when one or both of p or q is false.

Definition 3.5: Disjunction [or, \vee].

The proposition $p \vee q$ (“ p or q ,” the *disjunction* of the propositions p and q) is true when one or both of the propositions p or q is true, and is false when both p and q are false.

The prefix *con-* means “together” and *dis-* means “apart.” (*Junct* means “join.”) The *conjunction* $p \wedge q$ is true when p and q are true together; the *disjunction* $p \vee q$ is true when p is true “apart from” q , or the other way around. To help keep the symbols straight, it may be helpful to notice that the symbol \wedge is the angular version of the symbol \cap (intersection), while the symbol \vee is the angular version of the symbol \cup (union). The set $S \cap T$ is the set of all elements contained in S and T ; the set $S \cup T$ is the set of all elements contained in S or T .

In the conjunction $p \wedge q$, the propositions p and q are called *conjuncts*; in the disjunction $p \vee q$, they are called *disjuncts*.

Example 3.5: Some compound propositions.

Let p denote the proposition *Ohio State’s mascot is the Buckeye* and let q denote the proposition *Michigan’s mascot is the Wolverine*. Then we have the following compound propositions:

$\neg q$ *Michigan’s mascot is not the Wolverine.*

$p \wedge q$ *Ohio State’s mascot is the Buckeye, and Michigan’s mascot is the Wolverine.*

$p \vee q$ *Ohio State’s mascot is the Buckeye, or Michigan’s mascot is the Wolverine.*

Let’s also translate some English statements expressing compound propositions into logical notation:

Example 3.6: From English statements to compound propositions.

Translate each of the following statements into standard logical notation. (Name the atomic propositions using appropriate Boolean variables.)

- 1 Carissa is majoring in computer science and studio art.
- 2 Either Delroy took a formal logic class, or he is a quick learner.
- 3 Eli broke his hand and didn’t take the test as scheduled.
- 4 Florence knows Python or she has programmed in both C and Java.

Solution. We’ll first name the atomic propositions, and then translate them:

3-8 Logic

- 1 Write p for *Carissa is majoring in computer science* and q for *Carissa is majoring in studio art*. The given proposition is $p \wedge q$.
- 2 Write r to denote *Delroy took a formal logic class* and write s to denote *Delroy is a quick learner*. The given statement is $r \vee s$.
- 3 Denote *Eli broke his hand* by t and denote *Eli took the test as scheduled* by u . The sentence can be written as $t \wedge \neg u$.
- 4 Let v denote *Florence knows Python*, let w denote *Florence has programmed in C*, and let x denote *Florence has programmed in Java*. Then we can write $v \vee (w \wedge x)$ to represent the given sentence.

Implication (if/then)

Another important logical connective is \Rightarrow , which denotes *implication*. It expresses a familiar idea from everyday life, though one that's not quite captured by a single English word.

Consider the sentence *If you tell me how much money you make, then I'll tell you how much I make*. It's easiest to think of this sentence as a promise: I've promised that I'll tell you my financial situation *as long as you tell me yours*. I haven't promised anything about what I'll do if you don't tell me your salary—I can abstain from revealing anything about myself, or I might spontaneously tell you my salary anyway, but the point is that I haven't *guaranteed* anything. (But you'd justifiably call me a liar if you told me what you make, and I failed to disclose my salary in return.)

This kind of promise is expressed as an *implication* in propositional logic:

Definition 3.6: Implication [\Rightarrow].

The proposition $p \Rightarrow q$ is true when the truth of p implies the truth of q . In other words, $p \Rightarrow q$ is true unless p is true and q is false.

In the implication $p \Rightarrow q$, the proposition p is called the *antecedent* or the *hypothesis*, and the proposition q is called the *consequent* or the *conclusion*.

One initially confusing aspect of logical implication is that the word “implies” seems to hint at something about causation—but $p \Rightarrow q$ doesn't actually say anything about p *causing* q , only that p being true *implies that* q is true (or, in other words, p being true *lets us conclude that* q is true).

Example 3.7: Some implications.

The following propositions are all true:

- | | |
|--|----------------------------------|
| $1 + 1 = 2$ implies that $2 + 3 = 5$. | (“True implies True” is true.) |
| $2 + 3 = 4$ implies that $2 + 2 = 4$. | (“False implies True” is true.) |
| $2 + 3 = 4$ implies that $2 + 3 = 6$. | (“False implies False” is true.) |

3.2 An Introduction to Propositional Logic 3-9

But the following proposition is false:

$$2 + 2 = 4 \text{ implies that } 2 + 1 = 5. \quad (\text{“True implies False” is false.})$$

This last proposition is false because $2 + 2 = 4$ is true, but $2 + 1 = 5$ is false.

There are many different ways to express the proposition $p \Rightarrow q$ in English; see Figure 3.1a for some of them. Here’s a concrete example of multiple ways to phrase a real-world implication:

Example 3.8: Expressing implications in English.

According to United States law, people who can legally vote must be American citizens, and they must also satisfy some other various conditions that vary from state to state (for example, registering in advance or not being a felon). Thus the following compound proposition is true:

$$\text{you are a legal U.S. voter} \Rightarrow \text{you are an American citizen.}$$

All of the sentences in Figure 3.1b express this proposition in English. Most of these sentences are reasonably natural ways to express the stated implication, though the last phrasing seems awkward. But it’s easier to understand if we slightly rephrase it as “You being a legal U.S. voter *is sufficient for me to conclude that you are an American citizen.*”

Here’s another example of restating implications:

Example 3.9: More implications in English.

Consider this proposition:

The nondisclosure agreement is valid ^{p} *only if* ^{q} *you signed it.*

(This statement is different from *if you signed, then the agreement is valid*: for example, the agreement might not be valid because you’re legally a minor and thus not legally allowed to sign away rights.) We can restate $p \Rightarrow q$ as “if p then q ”:

If the nondisclosure agreement is valid, then you signed it.

- | | | | |
|-----|---|-----|---|
| (a) | “if p , then q ”
“ p implies q ”
“ p only if q ”
“ q whenever p ”
“ q , if p ”
“ q is necessary for p ”
“ p is sufficient for q ” | (b) | <i>If you are a legal U.S. voter, then you are an American citizen.</i>
<i>You being a legal U.S. voter implies that you are an American citizen.</i>
<i>You are a legal U.S. voter only if you are an American citizen.</i>
<i>You are an American citizen if you are a legal U.S. voter.</i>
<i>You are an American citizen whenever you are a legal U.S. voter.</i>
<i>You being an American citizen is necessary for you to be a legal U.S. voter.</i>
<i>You being a legal U.S. voter is sufficient for you to be an American citizen.</i> |
|-----|---|-----|---|

Figure 3.1 Expressing implications in English: (a) $p \Rightarrow q$, and (b) a particular implication [see Example 3.8].

3-10 Logic

We can also restate this implication equivalently—and perhaps more intuitively—using the so-called contrapositive $\neg q \Rightarrow \neg p$ (see Example 3.22):

The nondisclosure agreement is invalid if you didn't sign it.

Exclusive or

The four logical connectives that we have defined so far (\neg , \vee , \wedge , and \Rightarrow) are the ones that are most frequently used, but there are two other common connectives too. The first is *exclusive or*:

Definition 3.7: Exclusive or [\oplus].

The proposition $p \oplus q$ (“ p exclusive or q ” or, more briefly, “ p xor q ”) is true when one of p or q is true, but not both. Thus $p \oplus q$ is false when both p and q are true, and when both p and q are false.

(“Xor” and \oplus are usually pronounced like “ex ore,” as in *someone you used to date plus some rock with high precious-metal content.*)

When we want to emphasize the distinction between \vee and \oplus , we refer to \vee as *inclusive or*. This terminology highlights the fact that $p \vee q$ *includes* the possibility that both p and q are true, while $p \oplus q$ *excludes* that possibility. Unfortunately, the word “or” in English can mean either inclusive or exclusive or, depending on the context in which it’s being used. When you see the word “or,” you’ll have to think carefully about which meaning is intended.

Example 3.10: Inclusive versus exclusive or in English.

Translate these statements from a cover letter for a job into logical notation:

You may contact me by email or by phone. I am available for an on-site day-long interview on October 8th in Minneapolis or Hong Kong.

Use the following Boolean variables:

p = you may contact me by phone

q = you may contact me by email

r = I am physically available for an interview in Minneapolis

s = I am physically available for an interview in Hong Kong

Solution. The “or” in “email or phone” is *inclusive* (you could receive both an email and a call); the “or” in “Minneapolis or Hong Kong” is *exclusive* (it’s not physically possible to be simultaneously present in Minneapolis and Hong Kong). Thus we can translate these statements as $(p \vee q) \wedge (r \oplus s)$.

If and only if

We are now ready to define our last common logical connective:

3.2 An Introduction to Propositional Logic 3-11

Definition 3.8: If and only if [\Leftrightarrow].

The proposition $p \Leftrightarrow q$ (“ p if and only if q ”) is true when the propositions p or q have the same truth value (both p and q are true, or both p and q are false), and false otherwise.

The reason that \Leftrightarrow is read as “if and only if” is that $p \Leftrightarrow q$ means the same thing as the compound proposition $(p \Rightarrow q) \wedge (q \Rightarrow p)$. (We’ll prove this equivalence in Example 3.24.) Furthermore, the propositions $p \Rightarrow q$ and $q \Rightarrow p$ can be rendered, respectively, as “ p only if q ” and “ p , if q .” Thus $p \Leftrightarrow q$ expresses “ p if q , and p only if q ”—or, more compactly, “ p if and only if q .”

(The connective \Leftrightarrow is also sometimes called the *biconditional*, because an implication can also be called a *conditional*. You’ll also sometimes see \Leftrightarrow abbreviated in sentences as “iff” as shorthand for “if and only if”; we’ll avoid the “iff” abbreviation in this book, but it’s good to be prepared if you see it elsewhere.)

Unfortunately, just like with “or,” the word “if” is ambiguous in English. Sometimes “if” is used to express an implication, and sometimes it’s used to express an if-and-only-if definition. When you see the word “if” in a sentence, you’ll need to think carefully about whether it means \Rightarrow or \Leftrightarrow .

Example 3.11: “If” versus “if and only if” in English.

Think of a number between 10 and 1,000,000. Consider the following propositions:

p = your number is prime.

q = your number is even.

r = your number is evenly divisible by a positive integer other than 1 and itself.

Now translate the following two sentences into logical notation:

(1) *If the number you’re thinking of is even, then it isn’t prime.*

(2) *The number you’re thinking of isn’t prime if it’s evenly divisible by an integer other than 1 and itself.*

Solution. The “if” in (1) is an implication, and the “if” in (2) is “if and only if.” A correct translation of these sentences is (1) $q \Rightarrow \neg p$; and (2) $\neg p \Leftrightarrow r$.

3.2.4 Combining Logical Connectives

The six logical connectives that we’ve defined are summarized in Figure 3.2. The \neg connective is a *unary operator*, because it builds a compound proposition from a single simpler proposition. The other five connectives are *binary operators*, which build a compound proposition from two simpler propositions. (We’ll encounter the full list of binary logical connectives later; see Exercises 4.66–4.71.)

Taking it further: The unary-vs.-binary categorization of logical connectives based on how many “arguments” they accept also occurs in other contexts—for example, arithmetic and programming. In arithmetic, for example, you might distinguish between “unary minus” and “binary minus”: the former denotes negation, as in -3 ; the latter subtraction, as in $2 - 3$. The square root

3-12 Logic

negation	$\neg p$	“not p ”	highest precedence (binds tightest) ↓ lowest precedence (binds loosest)
conjunction	$p \wedge q$	“ p and q ”	
disjunction	$p \vee q$	“ p or q ”	
exclusive or	$p \oplus q$	“ p xor q ”	
implication	$p \Rightarrow q$	“if p , then q ” or “ p implies q ”	
if and only if	$p \Leftrightarrow q$	“ p if and only if q ”	

Figure 3.2 Summary of notation for propositional logic.

operator $\sqrt{}$ is also unary. In programming languages, the number of arguments that a function takes is called its *arity*. (The arity of `length` is one; the arity of `equals` is two.) You will sometimes encounter *variable arity* functions that can take a different number of arguments each time they’re invoked. Common examples include the print functions in many languages—C’s `printf` and Python’s `print`, for example, can take any number of arguments—or arithmetic in prefix languages like Scheme, where you can write an expression like `(+ 1 2 3 4)` to denote $1 + 2 + 3 + 4 (= 10)$.

Order of operations

A full description of the syntax of a programming language always includes a table of the *precedence* of operators, arranged from “binds the tightest” (highest precedence) to “binds the loosest” (lowest precedence). These precedence rules tell us when we have to include parentheses in an expression to make it mean what we want it to mean, and when the parentheses are optional.

In the same way, we’ll adopt some standard conventions regarding the precedence of our logical connectives: negation (\neg) has the highest precedence; then there is a three-way tie among \wedge , \vee , and \oplus ; then there’s \Rightarrow ; then finally \Leftrightarrow has the lowest precedence. (The horizontal lines in Figure 3.2 separate the logical connectives by their precedence.)

The word “precedence” (*pre* before, *cede* go) means “what comes first,” so the precedence rules in Figure 3.2 tell us the order in which the operators “get to go.” For example, consider the proposition $p \wedge q \Rightarrow r$. Because \wedge has higher precedence than \Rightarrow , the \wedge “goes first,” so the proposition is $(p \wedge q) \Rightarrow r$. (If \Rightarrow had “gone first,” it would have meant $p \wedge (q \Rightarrow r)$.) Although it’s a standard way of describing precedence, I find the “binds more tightly” phrasing a little weird, but here’s a way to think about it which may help make it make a bit more sense: deciding how tightly \neg binds is like deciding how small of a space to put between the operator \neg and the expression it operates on. For example, when we’re interpreting $\neg p \wedge q$, there are two possibilities:

$$\begin{array}{ccc} \neg p & \wedge & q \\ \neg \text{ binds more tightly than } \wedge & \text{and} & \neg \quad p \wedge q. \\ & & \wedge \text{ binds more tightly than } \neg \end{array}$$

Figure 3.2 says that \neg binds more tightly than \wedge , so the correct interpretation is $(\neg p) \wedge q$, the one on the left. These precedence rules match those in most programming languages; in Java, for example, the condition `! p && q` (“not p and q ” in Java syntax) will be interpreted as `(!p) && q`, because `not/~/!` binds tighter than `and/~/&&`.

Here are a couple of small examples:

Example 3.12: Precedence of logical connectives.

The propositions $p \vee \neg q$ and $p \vee q \Rightarrow \neg r \Leftrightarrow p$ mean, respectively,

$$p \vee (\neg q) \quad \text{and} \quad ((p \vee q) \Rightarrow (\neg r)) \Leftrightarrow p,$$

3.2 An Introduction to Propositional Logic 3-13

which we can see by applying the relevant precedence rules (“ \neg goes first, then \vee , then \Rightarrow , then \Leftrightarrow ”).

Figure 3.2 tells us the order in which two *different* operators are applied in an expression; for a sequence of applications of the *same* binary operator, we’ll use the convention that the operator *associates to the left*. For example, $p \wedge q \wedge r$ will mean $(p \wedge q) \wedge r$ rather than $p \wedge (q \wedge r)$. (This choice doesn’t matter for most of the logical connectives anyway: for example, the propositions $(p \wedge q) \wedge r$ and $p \wedge (q \wedge r)$ are true under exactly the same circumstances. In fact, implication is the only logical connective we’ve seen for which the order of application matters. See Exercises 3.48–3.50.)

Example 3.13: Precedence of logical connectives.

Fully parenthesize each of the following propositions. (In other words, add parentheses around each operator that respect the precedence rules and do not change the meaning.)

$$p \vee q \Leftrightarrow p$$

$$p \oplus p \oplus q \oplus q$$

$$\neg p \Leftrightarrow p \Leftrightarrow \neg(p \Leftrightarrow p)$$

Do the same for these:

$$p \wedge \neg q \Rightarrow r \Leftrightarrow s$$

$$p \Rightarrow q \Rightarrow r \wedge s.$$

Solution. Using the precedence rules from Figure 3.2 and left associativity, we get:

$$(p \vee q) \Leftrightarrow p$$

$$((p \oplus p) \oplus q) \oplus q$$

$$((\neg p) \Leftrightarrow p) \Leftrightarrow (\neg(p \Leftrightarrow p))$$

for the first set. For the second batch, we have

$$((p \wedge (\neg q)) \Rightarrow r) \Leftrightarrow s$$

$$(p \Rightarrow q) \Rightarrow (r \wedge s).$$

Writing tip: Because the order of application does matter for implication, it’s good style to include the optional parentheses so that it’s clear what you mean. Similarly, we’ll always use parentheses in propositions containing more than one of the tied-in-precedence operators \wedge , \vee , and \oplus from Figure 3.2, just as we should in programs. It’s hard to quickly recognize the result of Python code like `False and False or True`, and it’s nicer not to have to try.

3.2.5 Truth Tables

In Section 3.2.3, we described the logical connectives \neg , \wedge , \vee , \Rightarrow , \oplus , and \Leftrightarrow , but we can more systematically define these connectives by using a *truth table* that collects the value yielded by the logical connective under every *truth assignment*.

Definition 3.9: Truth assignment.

A *truth assignment* for a proposition over variables p_1, p_2, \dots, p_k is a function that assigns a truth value to each p_i .

3-14 Logic

For example, the function f where $f(p) = \text{T}$ and $f(q) = \text{F}$ is a truth assignment for the proposition $p \vee \neg q$. (Each “T” abbreviates a truth value of true; each “F” abbreviates a truth value of false.)

For any particular proposition and for any particular truth assignment f for that proposition, we can *evaluate* the proposition under f to figure out the truth value of the entire proposition. In the previous example, the proposition $p \vee \neg q$ is true under the truth assignment with $p = \text{T}$ and $q = \text{F}$ (because $\text{T} \vee \neg \text{F}$ is $\text{T} \vee \text{T}$, which is true). A *truth table* displays a proposition’s truth value (evaluated in the way we just described) under all truth assignments:

Definition 3.10: Truth table.

A *truth table* for a proposition lists, for each possible truth assignment for that proposition (with one truth assignment per row in the table), the truth value of the entire proposition.

Example 3.14: Defining \wedge .

For example, here is the truth table that defines the logical connective “and”:

p	q	$p \wedge q$
T	T	T
T	F	F
F	T	F
F	F	F

Columns #1 and #2 correspond to the atomic propositions p and q . There is a row in the table corresponding to each possible truth assignment for $p \wedge q$ —that is, for every pair of truth values for p and q . (So there are four rows: TT, TF, FT, and FF.)

The third column corresponds to the compound proposition $p \wedge q$, and it has a T only in the first row. That is, the truth value of $p \wedge q$ is false unless both p and q are true—just as Definition 3.4 said.

Figure 3.3 shows the truth tables for the six basic logical connectives. It’s worth paying special attention to the column for $p \Rightarrow q$: the *only* truth assignment under which $p \Rightarrow q$ is false is when p is true and q is false. *False implies anything! Anything implies true!* For example, both of the following are true propositions:

- (1) *If $2 + 3 = 4$, then you will eat tofu for dinner.* (If false, then anything!)
- (2) *If the moon is a clandestine surveillance device launched by the Sumerians in 3000 BCE, then $2 + 3 = 5$.* (If anything, then true!)

To emphasize the point, observe that (1) is true *even if* you would never eat tofu if it were the last so-called food on earth; the hypothesis “ $2 + 3 = 4$ ” of the proposition wasn’t true, so the truth of the proposition doesn’t depend on what your dinner plans are.

3.2 An Introduction to Propositional Logic 3-15

p	$\neg p$	p	q	$p \wedge q$	$p \vee q$	$p \Rightarrow q$	$p \oplus q$	$p \Leftrightarrow q$
T	F	T	T	T	T	T	F	T
T	F	T	F	F	T	F	T	F
F	T	F	T	F	T	T	T	F
F	T	F	F	F	F	T	F	T

Figure 3.3 Truth tables for the basic logical connectives.

Truth tables for more complex propositions

For more complicated propositions, we can fill in a truth table by repeatedly applying the rules in Figure 3.3. For example, to find the truth table for $(p \Rightarrow q) \wedge (q \vee p)$, we compute the truth tables for $p \Rightarrow q$ and $q \vee p$, and put a “T” in the $(p \Rightarrow q) \wedge (q \vee p)$ column for precisely those rows in which the truth tables for $p \Rightarrow q$ and $q \vee p$ both had “T”s. Here’s one small example, and a somewhat more complicated one:

Example 3.15: A small truth table.

Here is a truth table for the proposition $p \wedge q \Rightarrow \neg q$:

p	q	$p \wedge q$	$\neg q$	$p \wedge q \Rightarrow \neg q$
T	T	T	F	F
T	F	F	T	T
F	T	F	F	T
F	F	F	T	T

This truth table shows that the given proposition is true precisely when at least one of p and q is false.

s	t	u	v	$s \wedge t \wedge u$	$s \wedge t \wedge v$	$s \wedge u \wedge v$	$t \wedge u \wedge v$	$(s \wedge t \wedge u) \vee (s \wedge t \wedge v) \vee (s \wedge u \wedge v) \vee (t \wedge u \wedge v)$
T	T	T	T	T	T	T	T	T
T	T	T	F	T	F	F	F	T
T	T	F	T	F	T	F	F	T
T	T	F	F	F	F	F	F	F
T	F	T	T	F	F	T	F	T
T	F	T	F	F	F	F	F	F
T	F	F	T	F	F	F	F	F
T	F	F	F	F	F	F	F	F
F	T	T	T	F	F	F	T	T
F	T	T	F	F	F	F	F	F
F	T	F	T	F	F	F	F	F
F	T	F	F	F	F	F	F	F
F	F	T	T	F	F	F	F	F
F	F	T	F	F	F	F	F	F
F	F	F	T	F	F	F	F	F
F	F	F	F	F	F	F	F	F

Why are there five rows in which the last column is true? There are four different truth assignments corresponding to exactly three of $\{s, t, u, v\}$ being true (stu , $stuv$, $stvt$, tuv), and there is one truth assignment corresponding to all four being true ($stuv$). (In Chapter 9, on counting, we’ll re-encounter this style of question. And, actually, *precisely* this reasoning will allow us to prove something interesting about error-correcting codes—see Section 4.2.5.)

Figure 3.4 A truth table for a proposition expressing “at least three of $\{s, t, u, v\}$ are true,” for Example 3.16.

3-16 Logic

Example 3.16: Three (or more) of four, formalized.

In Example 3.4 (on the validity of passwords), we had a sentence of the form “ p , only if q and not r and at least three of $\{s, t, u, v\}$ are true.” Let’s translate this sentence into propositional logic. The tricky part will be translating “at least three of $\{s, t, u, v\}$ are true.”

There are many solutions, but one relatively simple way to do it is to explicitly write out four cases, one corresponding to allowing a different one of the four variables $\{s, t, u, v\}$ to be false:

$$(s \wedge t \wedge u) \vee (s \wedge t \wedge v) \vee (s \wedge u \wedge v) \vee (t \wedge u \wedge v)$$

We can verify that we’ve gotten this proposition right with a (big!) truth table, shown in Figure 3.4. Indeed, the five rows in which the last column has a “T” are exactly the five rows in which there are three or four “T”s in the columns for s, t, u , and v .

To finish the translation, recall that “ x only if y ” means $x \Rightarrow y$, so the given sentence can be translated as $p \Rightarrow q \wedge \neg r \wedge$ (the proposition above)—that is,

$$p \Rightarrow q \wedge \neg r \wedge \left((s \wedge t \wedge u) \vee (s \wedge t \wedge v) \vee (s \wedge u \wedge v) \vee (t \wedge u \wedge v) \right).$$

3.2 An Introduction to Propositional Logic 3-17

COMPUTER SCIENCE CONNECTIONS

NATURAL LANGUAGE PROCESSING, AMBIGUITY, AND TRUTH


Our main interest in this book is in developing (and understanding) precise and unambiguous language to express mathematical notions; in this chapter specifically, we're thinking about the truth values of completely precise statements. But thinking about the truth of ambiguous or ill-defined terms is absolutely crucial to any computational system that's designed to interact with users via natural language. (A *natural language* is one like English or French or Xhosa; these languages contrast with *artificial languages* like Java or Python or, arguably, Esperanto or Klingon.)

Natural language processing (NLP) (or the roughly similar *computational linguistics*) is the subfield of computer science that lies at the discipline's interface with linguistics. In NLP, we work to develop software systems that can interact with users in a natural language. A necessary step in an NLP system is to take an utterance made by the human user and "understand it." ("Understanding what a sentence means" is more or less the same as "understanding the circumstances under which it is true"—which is fundamentally a question of logic. You can find much more in any good textbook on natural language processing, such as [45, 66].)

One major reason that NLP is hard is that there is a tremendous amount of ambiguity in natural-language utterances. We can have *lexical ambiguity*, in which two different words are spelled identically but have two different meanings; we have to determine which word is meant in a sentence. Or there's *syntactic ambiguity*, in which a sentence's structure can be interpreted very differently. (See Figure 3.5.)

But there are also subtleties about when a statement is true, even if the meaning of each word and the sentence's structure are clear. Consider, for example, designing and implementing a chatbot that's supposed to assist people with planning travel. (Many companies, including transportation providers, use some kind of automated conversational system as part of their interactions with customers.) Imagine a user who asks the chatbot to find the first flight from MSP to BOS on a particular date. That's straightforward enough, and the system would be able to generate a database query to find the result: say, that the first nonstop flight from MSP to BOS leaves at 8:45am, with a nonrefundable fare of \$472.

But now imagine that the user's response to that information is to ask: *Is there a slightly later flight that isn't too much more expensive?* This conversational system now has to be able to decide the truth of statements like *10:33am is slightly later than 8:45am* and *\$529 isn't too much more expensive than \$472*, even though the "truth" of these statements depends on heavy use of conversational context and pragmatic reasoning. Of course, even



- A: Do you want coffee or tea?
 B: Do you want cream or sugar?
 C: We ate cake with walnuts.
 D: We ate cake with forks.

Figure 3.5 Examples of *lexical ambiguity* (A and B) and *syntactic ambiguity* (C and D). The *or* of A and B can be either inclusive or exclusive; simply answering "yes" is a reasonable response to question B, but a bizarre one to question A. The *with* of C and D can either attach to the *cake* or the *eating*; the sentences' structures are consistent with using walnuts as an eating utensil in C, or the cake containing forks as an ingredient in D.

3-18 Logic

though one cannot unambiguously determine whether these sentences are true or false, they're the kind of statement made continually in natural language. So systems that process natural language must deal with this issue.

One approach for handling these statements whose truth value is ambiguous is called *fuzzy logic*, in which each proposition has a truth value that is a real number between 0 and 1. (So *10:33a is slightly later than 8:45a* is “more true” than *12:19p is slightly later than 8:45a*—so the former might have a truth value of 0.74, while the latter might have a truth value of 0.34. But *7:30a is slightly later than 8:45a* would have a truth value of 0.00, as 7:30a is unambiguously *not* slightly later than 8:45a.)

EXERCISES

What are the truth values of the following propositions?

- 3.1 $2^2 + 3^2 = 4^2$
- 3.2 The number 202 is written 11010010 in binary.
- 3.3 After executing the following C code, the variable `x` has the value 1.

```
1 int x = 202;
2 while (x > 2) {
3     x = x / 2;    /* Note: in C, a single slash denotes integer division. */
4 }               /* For example, the value of 7/2 is 3, not 3.5. */
```

Using the atomic propositions in Figure 3.6, translate the following statements about Python expressions into logical notation:

- 3.4 `x ** y` is valid Python if and only if `x` and `y` are both numeric values.
- 3.5 `x + y` is valid Python if and only if `x` and `y` are both numeric values, or they're both lists.
- 3.6 `x * y` is valid Python if and only if `x` and `y` are both numeric values, or if one of `x` and `y` is a list and the other is numeric.
- 3.7 `x * y` is a list if `x * y` is valid Python and `x` and `y` are not both numeric values.
- 3.8 if `x + y` is a list, then `x * y` is not a list.
- 3.9 `x + y` and `x ** y` are both valid Python only if `x` is not a list.
- 3.10 True story: a 29-year-old friend of mine who does not have an advance care directive was asked *If you're over 55 years old, do you have an advance care directive?* on a form at a doctor's office. What should she answer, yes or no?
- 3.11 In Example 3.16, we constructed a proposition corresponding to "at least three of $\{s, t, u, v\}$ are true." Generalize this construction by building a proposition that expresses "at least 3 of $\{p_1, \dots, p_n\}$ are true."
- 3.12 Similar to Exercise 3.11: now build a proposition that expresses "at least $n - 1$ of $\{p_1, \dots, p_n\}$ are true."
- 3.13 The *identity* of a binary operator \diamond is a value i such that, for any x , the expressions $\{x, x \diamond i, i \diamond x\}$ are all equivalent. An example from arithmetic: the identity of $+$ is 0, because $x + 0 = 0 + x = x$ for any number x . Identify the identity of \vee , and justify your answer. (Some operators do not have an identity; if there is no identity, explain why it doesn't exist.)
- 3.14 What is the identity of \wedge ? (Or, if it doesn't exist, explain why not.)
- 3.15 What is the identity of \Leftrightarrow ? (Or, if it doesn't exist, explain why not.)
- 3.16 What is the identity of \oplus ? (Or, if it doesn't exist, explain why not.)
- 3.17 The *zero* of a binary operator \diamond is a value z such that, for any x , the expressions $\{z, x \diamond z, z \diamond x\}$ are all equivalent. For example, the zero of multiplication is 0, because $x \cdot 0 = 0 \cdot x = 0$ for any number x . What is the zero of \vee ? (Or explain why \vee has no zero.)
- 3.18 What is the zero of \wedge ? (Or, if it doesn't exist, explain why not.)
- 3.19 What is the zero of \Leftrightarrow ? (Or, if it doesn't exist, explain why not.)
- 3.20 What is the zero of \oplus ? (Or, if it doesn't exist, explain why not.)
- 3.21 Because \Rightarrow is not commutative (that is, because $p \Rightarrow q$ and $q \Rightarrow p$ mean different things), it is not too surprising that \Rightarrow has neither an identity nor a zero. But there are some related concepts for this type of operator. The *left identity* of a binary operator \diamond is a value i_ℓ such that, for any x , the expressions x and $i_\ell \diamond x$ are equivalent. The *right identity* of \diamond is a value i_r such that, for any x , the expressions x and $x \diamond i_r$ are equivalent. (Again, some operators may not have left or right identities.) What are the left and right identities of \Rightarrow (if they exist)?
- 3.22 The *left zero* of \diamond is a value z_ℓ such that, for any x , the expressions z_ℓ and $z_\ell \diamond x$ are equivalent; similarly, the *right zero* is a value z_r such that, for any x , the expressions z_r and $x \diamond z_r$ are equivalent. What are the left and right zeros for \Rightarrow (if they exist)?

In many programming languages, the Boolean values `True` and `False` are actually stored as the numerical values 1 and 0, respectively. In Python, for example, both `0 == False` and `1 == True` are `True`. Thus, despite appearances, we can do arithmetic on

3-20 Logic

p : $x + y$ is valid Python u : x is a numeric value
 q : $x * y$ is valid Python v : y is a numeric value
 r : $x ** y$ is valid Python w : x is a list
 s : $x * y$ is a list z : y is a list
 t : $x + y$ is a list

(Exercises 3.4–3.9 make accurate statements about expressions involving numbers and lists in Python, but even so they do not come close to fully characterizing the set of valid Python statements, for multiple reasons: first, they're about particular variables— x and y —rather than about generic variables. And, second, they omit some common-sense facts, like the fact that it's not simultaneously possible to be both a list and a numeric value.)

Figure 3.6 Some atomic propositions in Figure 3.6 for Exercises 3.4–3.9.

x_3	x_2	x_1	x_0		x_3	x_2	x_1	x_0		x_3	x_2	x_1	x_0	
0	0	0	1		0	0	1	1		1	1	0	0	
$0 + 0 + 0 + 1 = 1$					$0 + 0 + 2 + 1 = 3$					$8 + 4 + 0 + 0 = 12$				

Figure 3.7 Representing small numbers with four Boolean variables (where 0 = False and 1 = True).

Boolean values! Furthermore, in many languages (including Python), anything that is not False (in other words, anything other than 0) acts like True for the purposes of conditionals; for example, `if 2 then X else Y` will run successfully and execute X.

Suppose that x and y are two Boolean variables in a programming language where True and False are 1 and 0. (That is, the values of x and y are both 0 or 1.) Each of the following code snippets includes a conditional statement based on an arithmetic expression using x and y . Rewrite the given condition using the standard notation of propositional logic.

- 3.23** `if x * y ...`
3.24 `if x + y ...`
3.25 `if 2 - x - y ...`
3.26 `if (x * (1 - y)) + ((1 - x) * y) ...`
3.27 We can use the Booleans-as-integers idea from Exercises 3.23–3.26 to give a simpler solution to Exercises 3.11–3.12. Assume that $\{p_1, \dots, p_n\}$ are all Boolean variables in Python—that is, their values are all 0 or 1. Write a Python conditional expressing the condition that at least 3 of $\{p_1, \dots, p_n\}$ are true.
3.28 Write a Python conditional expressing the condition that at least $n - 1$ of $\{p_1, \dots, p_n\}$ are true.

In addition to purely logical operations, computer circuitry has to be built to do simple arithmetic very quickly. Consider a number $x \in \{0, \dots, 15\}$ represented as a 4-bit binary number, and denote by x_0 the least-significant bit of x , by x_1 the next bit, and so forth. (Think of 0 as false and 1 as true.) Here you'll explore some pieces of using propositional logic and binary representation of integers to express arithmetic operations. (It's straightforward to convert your answers into circuits.) See Figure 3.7.

- 3.29** Give a proposition over $\{x_0, x_1, x_2, x_3\}$ that expresses that x is greater than or equal to 8.
3.30 Give a proposition over $\{x_0, x_1, x_2, x_3\}$ that expresses that x is evenly divisible by 4.
3.31 Give a proposition over $\{x_0, x_1, x_2, x_3\}$ that expresses that x is evenly divisible by 5. (Hint: build a truth table.)
3.32 Give a proposition over $\{x_0, x_1, x_2, x_3\}$ that expresses that x is evenly divisible by 9.
3.33 Give a proposition over $\{x_0, x_1, x_2, x_3\}$ that expresses that x is an exact power of two.
3.34 Suppose that we have two 4-bit input integers x and y , represented as in Exercises 3.29–3.33. Give a proposition over the Boolean variables $\{x_0, x_1, x_2, x_3, y_0, y_1, y_2, y_3\}$ that expresses the condition that $x = y$.
3.35 Give a proposition over $\{x_0, x_1, x_2, x_3, y_0, y_1, y_2, y_3\}$ that expresses the condition $x \leq y$.
3.36 Give a proposition over $\{x_0, x_1, x_2, x_3, y_0, y_1, y_2, y_3\}$ that expresses the condition $x = 2 \cdot y$.
3.37 Give a proposition over $\{x_0, x_1, x_2, x_3, y_0, y_1, y_2, y_3\}$ that expresses the condition $x^y = y^x$.
3.38 Consider a 4-bit input integer x represented by four Boolean variables $\{x_0, x_1, x_2, x_3\}$ as in Exercises 3.29–3.33. Let y be $x + 1$, represented again as a 4-bit value $\{y_0, y_1, y_2, y_3\}$. (Treat $15 + 1 = 0$.) For example, if $x = 11$ (which is 1011 in binary), then $y = 12$ (which is 1100 in binary). Give four propositions over $\{x_0, x_1, x_2, x_3\}$ that express the values of y_0, y_1, y_2 , and y_3 .

Exercises 3-21

The following exercises ask you to build a program to evaluate a given proposition. To make your life as easy as possible, you should consider a simple representation of propositions, based on representing any compound proposition as a list. In such a list, the first element will be the logical connective, and the remaining elements will be the subpropositions. For example, the proposition $p \Rightarrow (\neg q)$ will be represented as ["implies", ["or", "p", "r"], ["not", "q"]] using Python list notation. Using this representation of propositions, write a program, in a language of your choice, to accomplish the following operations:

- 3.39** (programming required.) Given a proposition φ , compute the set of all atomic propositions contained within φ . (We'll occasionally use lowercase Greek letters, particularly φ ["phi"] or ψ ["psi"], to denote not-necessarily-atomic propositions.) The following recursive formulation may be helpful:

$$\text{vars}(p) = \{p\}$$

$$\text{vars}(\neg\varphi) = \text{vars}(\varphi)$$

$$\text{vars}(\varphi \wedge \psi) = \text{vars}(\varphi) \cup \text{vars}(\psi) \quad \text{and similarly for } \varphi \vee \psi, \varphi \Rightarrow \psi, \varphi \oplus \psi, \text{ and } \varphi \Leftrightarrow \psi.$$

- 3.40** (programming required.) Given a proposition φ and a truth assignment for each variable in φ , evaluate whether φ is true or false under this truth assignment.
- 3.41** (programming required.) Given a proposition φ , compute the set of all truth assignments for the variables in φ that make φ true. (One good approach: use your solution to Exercise 3.39 to compute all the variables in φ , then build the full list of truth assignments for those variables, and then evaluate φ under each of these truth assignments using your solution to Exercise 3.40.)

3-22 Logic

3.3 Propositional Logic: Some Extensions

“Beauty is truth, truth beauty,”—that is all
Ye know on earth, and all ye need to know.

John Keats (1795–1821)
Ode on a Grecian Urn (1819)

With the definitions from Section 3.2 in hand, we turn to a few extensions, including some special types of propositions and some special ways of representing propositions. Along the way, we’ll see a little bit about circuits, some ways that modern programming languages use inferences about certain logical expressions to allow you to write some conditional statements a little more slickly, and we’ll get a little more practice thinking about logical implication.

3.3.1 Tautology and Satisfiability

Several important types of propositions are defined in terms of their truth tables: those that are always true (*tautologies*), sometimes true (*satisfiable* propositions), or never true (*unsatisfiable* propositions, also called *contradictions*). We will explore each of these types in turn.

Tautologies

We’ll start by considering propositions that are always true:

Definition 3.11: Tautology.

A proposition is a *tautology* if it is true under every truth assignment.

Etymologically, the word *tautology* comes from Greek *taut* “same” (*to* + *auto*) + *logy* “word.” Another meaning for the word “tautology” (in real life, not just in logic) is the unnecessary repetition of an idea, as in a phrase like “a canine dog.” (The etymology and the non-mathematical meaning are not totally removed from the usage in logic.)

One reason that tautologies are important is that we can use them to reason about logical statements, which can be particularly valuable when we’re trying to prove a claim. Here are two important tautologies:

Example 3.17: Law of the Excluded Middle.

The truth table for the proposition $p \vee \neg p$ is shown in Figure 3.8a. The last column (corresponding to the proposition $p \vee \neg p$ itself) is filled with “T”s, so $p \vee \neg p$ is a tautology.

(a)	p	$\neg p$	$p \vee \neg p$
	T	F	T
	F	T	T

(b)	p	q	$p \Rightarrow q$	$p \wedge (p \Rightarrow q)$	$p \wedge (p \Rightarrow q) \Rightarrow q$
	T	T	T	T	T
	T	F	F	F	T
	F	T	T	F	T
	F	F	T	F	T

Figure 3.8 Truth tables for two tautologies: (a) Law of the Excluded Middle, and (b) Modus Ponens.

3.3 Propositional Logic: Some Extensions 3-23

$(p \Rightarrow q) \wedge p \Rightarrow q$	Modus Ponens
$(p \Rightarrow q) \wedge \neg q \Rightarrow \neg p$	Modus Tollens
$p \vee \neg p$	Law of the Excluded Middle
$p \Leftrightarrow \neg \neg p$	Double Negation
$p \Leftrightarrow p$	
$p \Rightarrow p \vee q$	
$p \wedge q \Rightarrow p$	

$(p \vee q) \wedge \neg p \Rightarrow q$
$(p \Rightarrow q) \wedge (\neg p \Rightarrow q) \Rightarrow q$
$(p \Rightarrow q) \wedge (q \Rightarrow r) \Rightarrow (p \Rightarrow r)$
$(p \Rightarrow q) \wedge (p \Rightarrow r) \Leftrightarrow p \Rightarrow q \wedge r$
$(p \Rightarrow q) \vee (p \Rightarrow r) \Leftrightarrow p \Rightarrow q \vee r$
$p \wedge (q \vee r) \Leftrightarrow (p \wedge q) \vee (p \wedge r)$
$p \Rightarrow (q \Rightarrow r) \Leftrightarrow p \wedge q \Rightarrow r$

Figure 3.9 Some tautologies.

The proposition $p \vee \neg p$ is called the *law of the excluded middle*: for any proposition p , either p is true or p is false; there is nothing “between” true and false.

Our second example of a tautology is the proposition $[p \wedge (p \Rightarrow q)] \Rightarrow q$, called *Modus Ponens*: if we know both that (a) p is true and that (b) the truth of p implies the truth of q , then we can conclude that q is true. (“Modus Ponens” rhymes with “goad us phone-ins”; literally, it means “the mood that affirms” in Latin.)

Example 3.18: Modus Ponens.

The truth table for $p \wedge (p \Rightarrow q) \Rightarrow q$ (with a few extra columns of “scratch work,” for each of the constituent pieces of the desired final proposition) is shown in Figure 3.8b. There are only “T”s in the last column, which establishes that Modus Ponens is a tautology.

Figure 3.9 contains a number of tautologies that you may find interesting and occasionally helpful. (Exercises 3.63–3.75 ask you to verify that these propositions really are tautologies.) One terminological note from Figure 3.9: *Modus Tollens* is the proposition $[(p \Rightarrow q) \wedge \neg q] \Rightarrow \neg p$, and it’s the counterpoint to Modus Ponens: if we know both that (a) the truth of p implies the truth of q and that (b) q is not true, then we can conclude that p cannot be true either. (Modus Tollens means “the mood that denies” in Latin.)

Satisfiable and unsatisfiable propositions

The next type of propositions we’ll consider are those that are *sometimes* true, and those propositions that are *never* true:

Definition 3.12: Satisfiable propositions.

A proposition is *satisfiable* if it is true under at least one truth assignment.

Definition 3.13: Unsatisfiable propositions/contradictions.

A proposition is *unsatisfiable* if it is not satisfiable. Such a proposition is also called a *contradiction*.

If f is a truth assignment under which a proposition is true, then we say that the proposition is *satisfied* by f . Thus a proposition is satisfiable if it is satisfied by at least one truth assignment, and it’s unsatisfiable if it is not satisfied by any truth assignment. (And it’s a tautology if it is satisfied by every truth assignment.)

3-24 Logic

(a)	p	q	$p \Leftrightarrow q$	$p \oplus q$	$(p \Leftrightarrow q) \wedge (p \oplus q)$	(b)	p	q	$p \vee q$	$\neg p$	$\neg q$	$\neg p \wedge \neg q$	$p \vee q \Rightarrow \neg p \wedge \neg q$
	T	T	T	F	F		T	T	T	F	F	F	F
	T	F	F	T	F		T	F	T	F	T	F	F
	F	T	F	T	F		F	T	T	T	F	F	F
	F	F	T	F	F		F	F	F	T	T	T	T

Figure 3.10 Truth tables for two propositions, for Examples 3.19 and 3.20.

Example 3.19: Contradiction of $p \Leftrightarrow q$ and $p \oplus q$.

The truth table for $(p \Leftrightarrow q) \wedge (p \oplus q)$ is shown in Figure 3.10a. The last column of the truth table is all False, which means that $(p \Leftrightarrow q) \wedge (p \oplus q)$ is unsatisfiable. Though it might not have been immediately apparent when they were defined, the logical connectives \oplus and \Leftrightarrow demand precisely opposite things of their arguments: the proposition $p \oplus q$ is true when p and q have *different* truth values, while $p \Leftrightarrow q$ is true when p and q have the *same* truth values. Because p and q cannot simultaneously have the same and different truth values, the conjunction $(p \Leftrightarrow q) \wedge (p \oplus q)$ is a contradiction.

Example 3.20: Demanding satisfaction: $p \vee q \Rightarrow \neg p \wedge \neg q$.

Is the proposition $p \vee q \Rightarrow \neg p \wedge \neg q$ satisfiable?

Solution. We'll answer the question by building a truth table for the given proposition. Because there is at least one "T" in the last column in the truth table shown in Figure 3.10b, the proposition is satisfiable. Specifically, this proposition is satisfied by the truth assignment $p = \text{False}$, $q = \text{False}$. (Under this truth assignment, the hypothesis $p \vee q$ is false; because false implies anything, the entire implication is true.)

It may be helpful to think about the general relationship between tautology and satisfiability. Let ϕ be *any* proposition. (We occasionally denote generic propositions by lowercase Greek letters, particularly ϕ ["phi"] or ψ ["psi"].) Then ϕ is a tautology exactly when $\neg\phi$ is unsatisfiable: ϕ is a tautology when the truth table for ϕ is all "T"s, which happens exactly when the truth table for $\neg\phi$ is all "F"s. And that's precisely the definition of $\neg\phi$ being unsatisfiable!

Taking it further: While satisfiability seems like a pretty precise technical definition that wouldn't matter all that much, the *satisfiability problem*—given a proposition ϕ , determine whether ϕ is satisfiable—turns out to be at the heart of the biggest open question in computer science today. If you figure out how to solve the satisfiability problem efficiently (or prove that it's impossible to solve efficiently), then you'll be the most famous computer scientist of the century. See p. 3-32.

3.3.2 Logical Equivalence

We'll now consider a relationship between *pairs* of propositions. When two propositions "mean the same thing" (that is, they are true under precisely the same circumstances), they are called *logically equivalent*:

3.3 Propositional Logic: Some Extensions 3-25

Definition 3.14: Logical equivalence.

Two propositions φ and ψ are *logically equivalent*, written $\varphi \equiv \psi$, if they have exactly identical truth tables (in other words, their truth values are the same under every truth assignment).

To state it differently: two propositions φ and ψ are logically equivalent whenever $\varphi \Leftrightarrow \psi$ is a tautology. Here's a small example of logical equivalence:

Example 3.21: $\neg(p \wedge q)$ and $(p \wedge q) \Rightarrow \neg q$ are logically equivalent.

In Example 3.15, we found that $(p \wedge q) \Rightarrow \neg q$ is true except when p and q are both true. Thus $\neg(p \wedge q)$ is logically equivalent to $(p \wedge q) \Rightarrow \neg q$, as this truth table shows:

p	q	$(p \wedge q) \Rightarrow \neg q$	$\neg(p \wedge q)$
T	T	F	F
T	F	T	T
F	T	T	T
F	F	T	T

Writing tip: Now that we have a reasonable amount of experience in writing truth tables, we will permit ourselves to skip columns when they're both obvious and not central to the point of a particular example. When you're writing *anything*—whether as a food critic or a Shakespeare scholar or a computer scientist—you should always think about the intended audience, and how much detail is appropriate for them.

Implication, converse, contrapositive, inverse, and mutual implication

We'll now turn to an important question of logical equivalence that involves the proposition $p \Rightarrow q$ and three other implications derived from it:

Definition 3.15: Converse, contrapositive, and inverse.

The *converse* of an implication $p \Rightarrow q$ is the proposition $q \Rightarrow p$.

The *contrapositive* of an implication $p \Rightarrow q$ is the proposition $\neg q \Rightarrow \neg p$.

The *inverse* of an implication $p \Rightarrow q$ is the proposition $\neg p \Rightarrow \neg q$.

These three new implications derived from the original implication $p \Rightarrow q$ —particularly the converse and the contrapositive—will arise frequently. Let's compare the three new implications to the original in light of logical equivalence:

Example 3.22: Implications, contrapositives, converses, inverses.

Consider the implication $p \Rightarrow q$. Which of the converse, contrapositive, and inverse of $p \Rightarrow q$ are logically equivalent to the original proposition $p \Rightarrow q$?

Solution. To answer this question, let's build the truth table, which is shown in Figure 3.11. Thus $p \Rightarrow q$ is logically equivalent to its contrapositive $\neg q \Rightarrow \neg p$, but *not* to its inverse or its converse.

Here's a real-world example to make these results more intuitive:

3-26 Logic

p	q	proposition	converse	contrapositive	inverse
		$p \Rightarrow q$	$q \Rightarrow p$	$\neg q \Rightarrow \neg p$	$\neg p \Rightarrow \neg q$
T	T	T	T	T	T
T	F	F	T	F	T
F	T	T	F	T	F
F	F	T	T	T	T

Figure 3.11 The truth table for an implication and its converse, its contrapositive, and its inverse.

Example 3.23: Contrapositives, converses, inverses, and Turing Awards.

Consider the following (true!) proposition, of the form $p \Rightarrow q$:

If ^{p} you won the Turing Award in 2008, then ^{q} your name is Barbara.

(The 2008 Turing Award winner was Barbara Liskov, who was honored for her groundbreaking work on principles of programming and programming languages, among other things.) The contrapositive of this proposition is $\neg q \Rightarrow \neg p$, which is also true:

If ^{$\neg q$} your name isn't Barbara, then ^{$\neg p$} you didn't win the Turing Award in 2008.

But the converse $q \Rightarrow p$ and the inverse $\neg p \Rightarrow \neg q$ are both blatantly false:

If your name is Barbara, then you won the Turing Award in 2008.

If you didn't win the Turing Award in 2008, then your name isn't Barbara.

Consider Barbara Kingsolver (best-selling author), Barbara Lee (long-time U.S. House member from California), Barbara McClintock (geneticist and 1983 Nobel Prize winner), and Barbara Walters (acclaimed television interviewer): all named Barbara, and not one of them a Turing Award winner.

It's worth emphasizing the results from Example 3.22: any implication $p \Rightarrow q$ is logically equivalent to its contrapositive $\neg q \Rightarrow \neg p$, but it is *not* logically equivalent to its converse $q \Rightarrow p$ or its inverse $\neg p \Rightarrow \neg q$. (You might notice, though, that the inverse and the converse *are* logically equivalent to each other. In fact, the converse's contrapositive—that is, the contrapositive of $q \Rightarrow p$ —is $\neg p \Rightarrow \neg q$, which just *is* the inverse.)

Here's another example of the concepts of tautology and satisfiability, as they relate to implications and converses:

Example 3.24: Mutual implication.

Consider the conjunction of the implication $p \Rightarrow q$ and its converse—that is, $(p \Rightarrow q) \wedge (q \Rightarrow p)$. Is this proposition a tautology? Satisfiable? Unsatisfiable? Is there a simpler proposition to which it's logically equivalent?

Solution. We can answer this question with a truth table:

3.3 Propositional Logic: Some Extensions 3-27

p	q	$p \Rightarrow q$	$q \Rightarrow p$	$(p \Rightarrow q) \wedge (q \Rightarrow p)$
T	T	T	T	T
T	F	F	T	F
F	T	T	F	F
F	F	T	T	T

Because there is a “T” in its column, $(p \Rightarrow q) \wedge (q \Rightarrow p)$ is satisfiable (and thus isn’t a contradiction). But that column does contain an “F” as well, and therefore $(p \Rightarrow q) \wedge (q \Rightarrow p)$ is *not* a tautology.

Notice that the truth table for $(p \Rightarrow q) \wedge (q \Rightarrow p)$ is identical to the truth table for $p \Leftrightarrow q$. (See Figure 3.3.) Thus $p \Leftrightarrow q$ and $(p \Rightarrow q) \wedge (q \Rightarrow p)$ are logically equivalent. (And \Leftrightarrow is called *mutual implication* for this reason: p and q imply each other.)

Some other logically equivalent statements

Figure 3.12 contains a large collection of logical equivalences. (It’s worth taking a few minutes to think about why each logical equivalence holds. See also Exercises 3.76–3.85.)

Here’s a brief description of some of the terminology. (Figure 3.12 contains a lot of new words to handle all at once—sorry!) Informally, an operator is *commutative* if the order of its arguments doesn’t matter; an operator is *associative* if the way we parenthesize successive applications doesn’t matter; and an operator is *idempotent* (Latin: *idem* “same” + *potent* “strength”) if applying it to the same argument twice gives that argument back. (There are two other frequently discussed concepts: the *identity* and the *zero* of the operator; logical equivalences involving identities and zeros were left to you, in Exercises 3.13–3.22.)

Taking it further: There are at least two ways in which logical equivalences (like those in Figure 3.12) play an important role in programming. First, most modern languages have a feature called *short-circuit evaluation* of logical expressions—they evaluate conjunctions and disjunctions from left to right, and stop as soon as the truth value of the logical expression is known—and programmers can exploit this feature to make their code cleaner or more efficient. Second, in compiled languages, an optimizing compiler can make use of logical equivalences to simplify the machine code that ends up being executed. See p. 3-34.

3.3.3 Representing Propositions: Circuits and Normal Forms

Now that we’ve established the core concepts of propositional logic, we’ll turn to some bigger and more applied questions. We’ll spend the rest of this section exploring two specific ways of representing propositions: *circuits*, the wires and connections from which physical computers are built; and two *normal forms*, in which the structure of propositions is restricted in a particular way.

The approach we’re taking with normal forms is a commonly used idea to make reasoning about some language L easier: we define a *subset* S of L , with two goals: (1) any statement in L is equivalent to some statement in S ; and (2) S is “simple” in some way. Then we can consider any statement from the “full” language L , which we can then “translate” into a simple-but-equivalent statement of S . Defining this subset

3-28 Logic

Commutativity	$p \vee q \equiv q \vee p$	Distribution of \wedge over \vee	$p \wedge (q \vee r) \equiv (p \wedge q) \vee (p \wedge r)$
	$p \wedge q \equiv q \wedge p$		$p \vee (q \wedge r) \equiv (p \vee q) \wedge (p \vee r)$
	$p \oplus q \equiv q \oplus p$	Contrapositive	$p \Rightarrow q \equiv \neg q \Rightarrow \neg p$
	$p \Leftrightarrow q \equiv q \Leftrightarrow p$		$p \Rightarrow q \equiv \neg p \vee q$
Associativity	$p \vee (q \vee r) \equiv (p \vee q) \vee r$	Mutual Implication	$p \Rightarrow (q \Rightarrow r) \equiv p \wedge q \Rightarrow r$
	$p \wedge (q \wedge r) \equiv (p \wedge q) \wedge r$		$p \Leftrightarrow q \equiv \neg p \Leftrightarrow \neg q$
	$p \oplus (q \oplus r) \equiv (p \oplus q) \oplus r$		$(p \Rightarrow q) \wedge (q \Rightarrow p) \equiv p \Leftrightarrow q$
	$p \Leftrightarrow (q \Leftrightarrow r) \equiv (p \Leftrightarrow q) \Leftrightarrow r$		
Idempotence	$p \vee p \equiv p$	De Morgan's Laws	$\neg(p \wedge q) \equiv \neg p \vee \neg q$
	$p \wedge p \equiv p$		$\neg(p \vee q) \equiv \neg p \wedge \neg q$

Figure 3.12 Some logically equivalent propositions. De Morgan's Laws are named after Augustus De Morgan (1806–1871), a British mathematician who not only made major contributions to the study of logic but was also the math tutor of Ada Lovelace (1815–1852), “the first computer programmer.” (See Figure 2.15.)

and its accompanying translation will make it easier to accomplish some task for *all* expressions in L , while still making it easy to write statements clearly.

Taking it further: The idea of translating all propositions into a particular form has a natural analogue in designing and implementing programming languages. For example, every `for` loop can be expressed as a `while` loop instead, but it would be very annoying to program in a language that doesn't have `for` loops. A nice compromise is to allow `for` loops, but behind the scenes to translate each `for` loop into a `while` loop. This compromise makes the language easier for the “user” programmer to use (for loops exist!) and also makes the job of the programmer of the compiler/interpreter easier (she can worry exclusively about implementing and optimizing `while` loops!).

In programming languages, this translation is captured by the notion of *syntactic sugar*. (The phrase is meant to suggest that the addition of `for` to the language is a bonus for the programmer—“sugar on top,” maybe—that adds to the syntax of the language.) The programming language Scheme is perhaps the pinnacle of syntactic sugar; the core language is almost unbelievably simple. Here's one illustration: `(and x y)` (Scheme for “ $x \wedge y$ ”) is syntactic sugar for `(if x y #f)` (that's “if x then y else false”). So a Scheme programmer can use `and`, but there's no “real” `and` and that has to be handled by the interpreter.

Circuits

We'll introduce the idea of circuits by using the proposition $(p \wedge \neg q) \vee (\neg p \wedge q)$ as a particular example. (Note, by the way, that this proposition is logically equivalent to $p \oplus q$.)

This proposition is a disjunction of two smaller propositions, $p \wedge \neg q$ and $\neg p \wedge q$. Similarly, $p \wedge \neg q$ is a conjunction of two even simpler propositions, namely p and $\neg q$. A representation of a proposition called




Figure 3.13 The proposition $(p \wedge \neg q) \vee (\neg p \wedge q)$ both (a) as a tree, and (b) as a circuit.

3.3 Propositional Logic: Some Extensions 3-29


a *tree* continues to break down every compound proposition embedded within it. (We'll talk about trees in detail in Chapter 11.) See Figure 3.13a. The tree-based view isn't much of a change from our usual notation $(p \wedge \neg q) \vee (\neg p \wedge q)$; all we've done is use the parentheses and order-of-operation rules to organize the logical connectives, and then we've arranged that organization in a hierarchical layout.

This tree representation is closely related to another important way of viewing logical expressions: *circuits*, which represent a proposition as a collection of *wires* and *gates*. Wires carry a truth value from one physical location to another; gates are physical implementations of logical connectives. See Figure 3.13b. In the circuit-based view, we can think of truth values “flowing in” as inputs to the left side of each gate, and a truth value “flowing out” as output from the right side of the gate. (The only substantive difference between the tree and the circuit—aside from which way is up—is whether the two p inputs come from the same wire, and likewise whether the two q inputs do.)

Example 3.25: Using and and not for or.

Build a circuit for $p \vee q$ using only \wedge and \neg gates.

Solution. We'll use one of De Morgan's Laws, which says that $p \vee q \equiv \neg(\neg p \wedge \neg q)$:



This basic idea—replacing one logical connective by another one (or by multiple other ones)—is crucial to the construction of computers themselves; we'll return to this idea in Section 4.4.1.

Conjunctive and Disjunctive Normal Forms

In the rest of this section, we'll consider a way to simplify the structure of propositions: *conjunctive* and *disjunctive normal forms*, which constrain propositions to have a particular format.

To define these restricted types of propositions, we first need a basic definition. A *literal* is a Boolean variable (a.k.a. an atomic proposition) or the negation of a Boolean variable. (So p and $\neg p$ are both literals.)

Definition 3.16: Conjunctive normal form.

A proposition is in *conjunctive normal form (CNF)* if it is the conjunction of one or more *clauses*, where each *clause* is the disjunction of one or more literals.

Definition 3.17: Disjunctive normal form.

A proposition is in *disjunctive normal form (DNF)* if it is the disjunction of one or more *clauses*, where each *clause* is the conjunction of one or more literals.

3-30 Logic

Less formally, a proposition in conjunctive normal form is “the *and* of a bunch of *ors*,” and a proposition in disjunctive normal form is “the *or* of a bunch of *ands*.”

Taking it further: In computer architecture and digital electronics, people usually refer to a proposition in CNF as being a *product of sums*, and a proposition in DNF as being a *sum of products*. (There is a deep way of thinking about formal logic based on \wedge as multiplication, \vee as addition, 0 as False, and 1 as True; see Exercises 3.23–3.26.)

Here are small examples of both CNF and DNF:

Example 3.26: Small propositions in CNF and DNF.

Here are two propositions, one in conjunctive normal form and one in disjunctive normal form:

conjunctive normal form: $\left(\neg p \vee q \vee r \right) \wedge \left(\neg q \vee \neg r \right) \wedge \left(r \right)$
clause #1 clause #2 clause #3

disjunctive normal form: $\left(\neg p \wedge q \wedge r \right) \vee \left(\neg q \wedge \neg r \wedge s \right) \vee \left(r \right) \vee \left(p \wedge \neg r \right)$
clause #1 clause #2 clause #3 clause #4

While conjunctive and disjunctive normal forms seem like heavy restrictions on the format of propositions, *every* proposition is logically equivalent to a CNF proposition and to a DNF proposition:

Theorem 3.18: All propositions are expressible in CNF.

For any proposition φ , there is a proposition φ_{cnf} over the same Boolean variables and in conjunctive normal form such that $\varphi \equiv \varphi_{\text{cnf}}$.

Theorem 3.19: All propositions are expressible in DNF.

For any proposition φ , there is a proposition ψ_{dnf} over the same Boolean variables and in disjunctive normal form such that $\varphi \equiv \psi_{\text{dnf}}$.

These two theorems are perhaps the first results that we’ve encountered that are unexpected, or at least unintuitive. There’s no particular reason for it to be clear that they’re true—let alone how we might prove them. But we can, and we will: we’ll prove both theorems in Section 4.4.1 and again in Section 5.4.3, after we’ve introduced some relevant proof techniques. But, for now, here are a few examples of translating propositions into DNF/CNF.

Problem-solving tip: A good strategy when you’re trying to prove a not-at-all-obvious claim is to test out some small examples, and then try to start to figure a general pattern. In Examples 3.27 and 3.28, we’ll try out a few examples of converting (relatively simple) propositions into logically equivalent propositions that are in CNF/DNF. We’ll figure out how to generalize this technique to *any* proposition in Section 4.4.1.

Example 3.27: Translating basic connectives into DNF.

Give propositions in disjunctive normal form that are logically equivalent to each of the following:

1 $p \vee q$

3.3 Propositional Logic: Some Extensions 3-31

2 $p \wedge q$

3 $p \Rightarrow q$

4 $p \Leftrightarrow q$

Solution.1 This question is boring: $p \vee q$ is *already* in DNF, with two clauses (clause #1 is p ; clause #2 is q).2 This question is boring, too: $p \wedge q$ is also already in DNF, with a single clause ($p \wedge q$).3 Figure 3.12 tells us that $p \Rightarrow q \equiv \neg p \vee q$, and $\neg p \vee q$ is in DNF.4 The proposition $p \Leftrightarrow q$ is true when p and q are either both true or both false. So we can rewrite $p \Leftrightarrow q$ as $(p \wedge q) \vee (\neg p \wedge \neg q)$. We can check that we've gotten it right with a truth table:

p	q	$p \wedge q$	$\neg p \wedge \neg q$	$(p \wedge q) \vee (\neg p \wedge \neg q)$	$p \Leftrightarrow q$
T	T	T	F	T	T
T	F	F	F	F	F
F	T	F	F	F	F
F	F	F	T	T	T

Thus $p \Leftrightarrow q \equiv (p \wedge q) \vee (\neg p \wedge \neg q)$.

And here's the analogous task of translating a few of the logical connectives into CNF. (As with DNF, both $p \vee q$ and $p \wedge q$ are already in CNF, so we'll skip them in this example.)

Example 3.28: Translating basic connectives into CNF.

Give propositions in conjunctive normal form that are logically equivalent to each of the following:

1 $p \Rightarrow q$

2 $p \Leftrightarrow q$

3 $p \oplus q$

Solution.1 As in Example 3.27, we know that $p \Rightarrow q \equiv \neg p \vee q$, and $\neg p \vee q$ is in CNF (with one clause), too.2 We can rewrite $p \Leftrightarrow q$ as follows:

$$\begin{aligned}
 p \Leftrightarrow q &\equiv (p \Rightarrow q) \wedge (q \Rightarrow p) && \text{mutual implication (Example 3.24)} \\
 &\equiv (\neg p \vee q) \wedge (\neg q \vee p). && x \Rightarrow y \equiv \neg x \vee y \text{ (Figure 3.12), used twice}
 \end{aligned}$$

The proposition $(\neg p \vee q) \wedge (\neg q \vee p)$ is in CNF.3 Because $p \oplus q$ is true as long as one of $\{p, q\}$ is true and one of $\{p, q\}$ is false, we can verify via truth table that $p \oplus q \equiv (p \vee q) \wedge (\neg p \vee \neg q)$, which is in CNF.

COMPUTER SCIENCE CONNECTIONS

COMPUTATIONAL COMPLEXITY, SATISFIABILITY, AND A MILLION DOLLARS

Complexity theory is the subfield of computer science devoted to understanding the computational resources—time and memory, usually—necessary to solve particular problems. It’s the subject of a great deal of fascinating current research in theoretical computer science. Here’s a little bit of a flavor of that work.

One of the central problems of complexity theory is the *satisfiability problem*, in which you’re given a proposition φ over n variables, and asked to determine whether φ is satisfiable. This problem is pretty simple to solve. In fact, we’ve implicitly described an algorithm for it already: you just construct the truth table for the n -variable proposition φ , and then check to see whether there are any “T”s in φ ’s column of the table. But this algorithm is not very fast, because the truth table for φ has lots and lots of rows— 2^n rows, to be precise. Even a moderate value of n means that this algorithm will not terminate in your lifetime; 2^{300} exceeds the number of particles in the known universe.

So, it’s clear that there is an algorithm that solves the SAT problem. What’s not clear is whether there is a substantially more efficient algorithm to solve the SAT problem. It’s so unclear, in fact, that literally nobody knows the answer, and this question is one of the biggest open problems in computer science and mathematics today. (Arguably, it’s *the* biggest.) The Clay Mathematics Institute will even give a \$1,000,000 prize to anyone who solves it.

Why is this problem so important? The reason is that, in a precise technical sense, SAT is *just as hard* as a slew of other problems that have a plethora of unspeakably useful applications: the traveling salesperson problem, protein folding, optimally packing the trunk of a car with suitcases. (Or see Figure 3.14.) This slew is a class of computational problems known as NP (“nondeterministic polynomial time”), for which it is easy to “verify” correct answers. In the context of SAT, that means that whenever you’ve got a satisfiable proposition φ , it’s very easy for you to (efficiently) convince me that φ is satisfiable. Here’s how: you’ll simply tell me a truth assignment under which φ evaluates to true. And I can make sure that you didn’t try to fool me by plugging and chugging: I substitute your truth assignment in for every variable, and then I make sure that the final truth value of φ is indeed True.

One of the most important results in theoretical computer science in the 20th century—that’s saying something for a field that was founded in the 20th century!—is the *Cook–Levin Theorem*: *if you can solve SAT efficiently, then you can solve any problem in NP efficiently*. The major open question is what’s known as the *P-versus-NP question*. A problem that’s in P is easy to solve from scratch. A problem that’s in NP is easy to verify (in the way described above). So the question is:

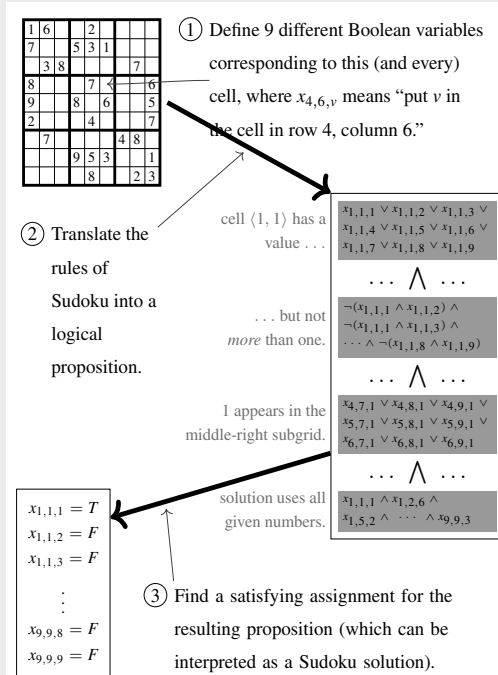


Figure 3.14 Solving Sudoku using Satisfiability. This type of translation is the main reason that SAT is so important: a huge range of (important and practical) problems can be “translated” into satisfiability, so that an efficient solution to SAT implies an efficient solution to all of them, too.

3.3 Propositional Logic: Some Extensions 3-33

does $P = NP$? Is verifying an answer to a problem no easier than solving the problem from scratch? (It seems intuitively “clear” that the answer is no—but nobody has been able to prove it!)

(You can read more about complexity theory in general, and the P-versus-NP question addressed here in particular, in most books on algorithms or the theory of computing. The original results on NP-completeness are from Stephen Cook [31] and Leonid Levin [79]; the first translations of other problems are due to Richard Karp [68].)

3-34 Logic

COMPUTER SCIENCE CONNECTIONS

SHORT-CIRCUIT EVALUATION, OPTIMIZATION, AND MODERN COMPILERS

The logical equivalences in Figure 3.12 may seem far removed from “real” programming, but logical equivalences are actually central in modern programming. Here are two ways in which they play an important role:

Short-circuit evaluation: In most modern programming languages, a logical expression involving **ands** and **ors** will only be evaluated until the truth value of the expression can be determined. For an example in Java, see Figure 3.15a. Like most modern languages, Java evaluates an \wedge expression from left to right and stops as soon as it finds a false conjunct. Similarly, Java evaluates an \vee expression from left to right and stops as soon as it finds a true disjunct. These simplifications rely on facts from propositional logic: $\text{False} \wedge \text{anything} \equiv \text{False}$ and $\text{True} \vee \text{anything} \equiv \text{True}$. This style of evaluation is called *short-circuit evaluation*.

Two slick ways in which programmers can take advantage of short-circuit evaluation are shown in Figure 3.15b. First, lines 1–3 use short-circuit evaluation to avoid deeply nested **if** statements to handle exceptional cases. When $x = 0$, evaluating the second disjunct would cause a divide-by-zero error—but the second disjunct isn’t evaluated when $x = 0$ because the first disjunct was true! Second, lines 4–7 use short-circuit evaluation to make code faster. If the second conjunct typically takes much longer to evaluate (or if it is much more frequently true) than the first conjunct, then careful ordering of conjuncts avoids a long and usually fruitless computation.

Compile-time optimization: For a program written in a compiled language like C, the source code is translated into machine-readable form by the *compiler*. But this translation is not verbatim; instead, the compiler streamlines your code (when it can!) to make it run faster.

One of the simplest types of compiler optimizations is *constant folding*: if some of the values in an arithmetic or logical expression are constants—known to the compiler at “compile time,” and thus unchanged at “run time”—then the compiler can “fold” those constants together. Using the rules of logical or arithmetic equivalence broadens the types of code that can be folded in this way. For example, in C, when you write an assignment statement like $y = x + 2 + 3$, most compilers will translate it into $y = x + 5$. Similarly, for $z = 7 * x * 8$, a modern compiler will be able

```
1 if (2 > 3 && x + y < 9)
2   ...
3 } else {
4   ...
5 }
```

&& is Java notation for \wedge ;
|| is Java notation for \vee .

(a) The second conjunct of the **if** will never be evaluated: $2 > 3$ is false, and $\text{False} \wedge \text{anything} \equiv \text{False}$.

```
1 if (x == 0 || (x-1)/x > 0.5) {
2   ...
3 }

4 if (simpleOrOftenFalse(x)
5     && complexOrOftenTrue(x)) {
6   ...
7 }
```

(b) Two handy ways to rely on short-circuit evaluation.

Figure 3.15 Short-circuit evaluation, illustrated in Java.

```
1 if (p || !p)
2   x = 51;
3 } else {
4   x = 63;
5 }
```

($p \parallel !p$) denotes $p \vee \neg p$ in C.

```
1 x = 51;
```

Figure 3.16 Two snippets of C code. When this code is compiled on a modern optimizing compiler (a recent version of gcc, with optimization turned on), the machine code that is produced is *exactly* identical for both snippets.

3.3 Propositional Logic: Some Extensions 3-35

to optimize it into $z = x * 56$, using the commutativity of multiplication. Because the compiler can reorder the multiplicands without affecting the value, and this reordering allows the 7 and 8 to be folded into 56, the compiler does the reordering and the folding.

An example using logical equivalences is shown in Figure 3.16. Because $p \vee \neg p$ is a tautology—the law of the excluded middle—*no matter what the value of p* , the “then” clause is executed, not the “else” clause. Thus the compiler doesn’t even have to waste time checking whether p is true or false, and this optimization can be applied.

3-36 Logic

EXERCISES

The operators \wedge and \vee are idempotent—in other words, $p \wedge p \equiv p \vee p \equiv p$. But \Rightarrow , \oplus , and \Leftrightarrow are not idempotent:

- 3.42** Simplify the expression $p \Rightarrow p$ (that is, give an as-simple-as-possible logically equivalent proposition).
3.43 Repeat for $p \oplus p$.
3.44 Repeat for $p \Leftrightarrow p$.
3.45 Add parentheses to the proposition $p \Rightarrow \neg p \Rightarrow p \Rightarrow q$, so that the resulting proposition is a tautology.
3.46 Add parentheses to the proposition $p \Rightarrow \neg p \Rightarrow p \Rightarrow q$, so that the resulting proposition is logically equivalent to q .
3.47 Give as simple as possible a proposition that is logically equivalent to the (unparenthesized) original proposition $p \Rightarrow \neg p \Rightarrow p \Rightarrow q$.
3.48 Unlike the logical connectives \wedge , \vee , \oplus , and \Leftrightarrow , implication is not associative. In other words, $p \Rightarrow (q \Rightarrow r)$ and $(p \Rightarrow q) \Rightarrow r$ are not logically equivalent. Prove it: give a truth assignment in which $p \Rightarrow (q \Rightarrow r)$ and $(p \Rightarrow q) \Rightarrow r$ have different truth values.
3.49 Consider the propositions $p \Rightarrow (q \Rightarrow q)$ and $(p \Rightarrow q) \Rightarrow q$. One of these is a tautology; the other isn't. Which is which? Explain.
3.50 Consider the propositions $p \Rightarrow (p \Rightarrow q)$ and $(p \Rightarrow p) \Rightarrow q$. Is either one a tautology? Satisfiable? Unsatisfiable? What is the simplest proposition to which each is logically equivalent?

On an exam, I once asked students to write a proposition logically equivalent to $p \oplus q$ using only the logical connectives \Rightarrow , \neg , and \wedge . Here are some of the students' answers. Which ones are right?

- 3.51** $\neg(p \wedge q) \Rightarrow (\neg p \wedge \neg q)$
3.52 $(p \Rightarrow \neg q) \wedge (q \Rightarrow \neg p)$
3.53 $(\neg p \Rightarrow q) \wedge \neg(p \wedge q)$
3.54 $\neg[(p \wedge \neg q \Rightarrow \neg p \wedge q) \wedge (\neg p \wedge q \Rightarrow p \wedge \neg q)]$
3.55 Write a proposition logically equivalent to $p \oplus q$ using only the logical connectives \Rightarrow , \neg , and \vee .
3.56 Simplify the code in Figure 3.17a as much as possible. (For example, if $p \Rightarrow q$, it's a waste of time to test whether q holds in a block where p is known to be true.)
3.57 Repeat for the code in Figure 3.17b.
3.58 Repeat for the code in Figure 3.17c.
3.59 Simplify $(\neg p \Rightarrow q) \wedge (q \wedge p \Rightarrow \neg p)$ as much as possible.
3.60 Repeat for $(p \Rightarrow \neg p) \Rightarrow ((q \Rightarrow (p \Rightarrow p)) \Rightarrow p)$.
3.61 Repeat for $(p \Rightarrow p) \Rightarrow (\neg p \Rightarrow \neg p) \wedge q$.
3.62 Claim: Every proposition over the single variable p is either logically equivalent to p or it's logically equivalent to $\neg p$. Is this claim true or false? Prove your answer.

Show using truth tables that these propositions (mostly from Figure 3.9) are tautologies:

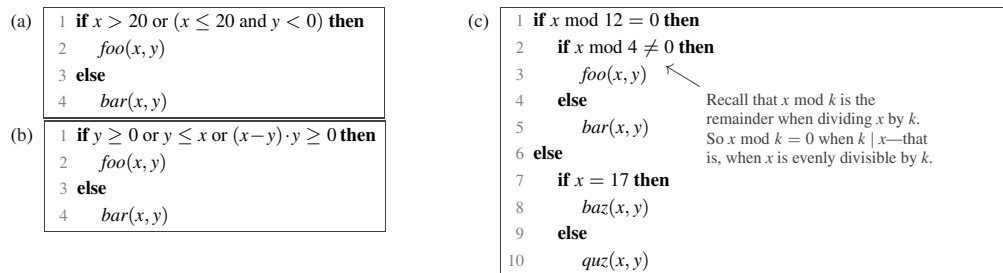


Figure 3.17 Some code that uses nested conditionals, or compound propositions as conditions.

Exercises 3-37

- 3.63 $(p \Rightarrow q) \wedge \neg q \Rightarrow \neg p$ (Modus Tollens) 3.70 $(p \Rightarrow q) \vee (p \Rightarrow r) \Leftrightarrow p \Rightarrow q \vee r$
 3.64 $p \Rightarrow p \vee q$ 3.71 $p \wedge (q \vee r) \Leftrightarrow (p \wedge q) \vee (p \wedge r)$
 3.65 $p \wedge q \Rightarrow p$ 3.72 $p \Rightarrow (q \Rightarrow r) \Leftrightarrow p \wedge q \Rightarrow r$
 3.66 $(p \vee q) \wedge \neg p \Rightarrow q$ 3.73 $p \vee (p \wedge q) \Leftrightarrow p$
 3.67 $(p \Rightarrow q) \wedge (\neg p \Rightarrow q) \Rightarrow q$ 3.74 $p \wedge (p \vee q) \Leftrightarrow p$
 3.68 $(p \Rightarrow q) \wedge (q \Rightarrow r) \Rightarrow (p \Rightarrow r)$ 3.75 $p \oplus q \Rightarrow p \vee q$
 3.69 $(p \Rightarrow q) \wedge (p \Rightarrow r) \Leftrightarrow p \Rightarrow q \wedge r$

Show using truth tables that the following logical equivalences from Figure 3.12 hold:

- 3.76 $\neg(p \wedge q) \equiv \neg p \vee \neg q$ (De Morgan's Law) 3.81 $p \Leftrightarrow (q \Leftrightarrow r) \equiv (p \Leftrightarrow q) \Leftrightarrow r$ (associativity of \Leftrightarrow)
 3.77 $\neg(p \vee q) \equiv \neg p \wedge \neg q$ (De Morgan's Law) 3.82 $p \Rightarrow q \equiv \neg p \vee q$
 3.78 $p \vee (q \vee r) \equiv (p \vee q) \vee r$ (associativity of \vee) 3.83 $p \Rightarrow (q \Rightarrow r) \equiv p \wedge q \Rightarrow r$
 3.79 $p \wedge (q \wedge r) \equiv (p \wedge q) \wedge r$ (associativity of \wedge) 3.84 $p \Leftrightarrow q \equiv \neg p \Leftrightarrow \neg q$
 3.80 $p \oplus (q \oplus r) \equiv (p \oplus q) \oplus r$ (associativity of \oplus) 3.85 $\neg(p \Rightarrow q) \equiv p \wedge \neg q$

- 3.86 Optimizing compilers (see p. 3-34) will perform the following optimization, transforming the first block of C code into the second:

```
1 if (p || !p) { /* "p or not p" */
2     x = 51;
3 } else {
4     x = 63;
5 }
```

```
1 x = 51;
```

The compiler performs this transformation because $p \vee \neg p$ is a tautology: no matter the truth value of p , the proposition $p \vee \neg p$ is true. But there are situations in which this code translation actually changes the behavior of the program, if p can be an arbitrary expression (rather than just a Boolean variable)! Describe such a situation. (Hint: why do (some) people watch auto racing?)

- 3.87 A very mild pet peeve of mine: code that uses Option A for some Boolean-valued expression p instead of Option B.

Option A
1 if $p = \text{True}$ then A else B

Option B
1 if p then A else B

State and prove the tautology that establishes the equivalence of Option A and Option B.

See Figure 3.18 for the description of a circuit. For each of the following exercises, draw a circuit with at most 3 gates that is consistent with the listed behavior. The light's status is unknown for unlisted inputs. (If multiple circuits are consistent with the given behavior, choose one with as few gates as possible.)

- 3.88 The light is on when the true inputs are $\{q\}$ or $\{r\}$. It is off when the true inputs are $\{p\}$ or $\{p, q\}$ or $\{p, q, r\}$.
 3.89 The light is on when the true inputs are $\{p, q\}$ or $\{p, r\}$. It is off when the true inputs are $\{p\}$ or $\{q\}$ or $\{r\}$.
 3.90 The light is on for at least one input setting. It is off when the true inputs are $\{p\}$ or $\{q\}$ or $\{r\}$ or $\{p, q, r\}$.
 3.91 The light is on for at least two input settings. It is off when the true inputs are $\{p, q\}$ or $\{p, r\}$ or $\{q, r\}$ or $\{p, q, r\}$.
 3.92 The light is always off, no matter what the input.
 3.93 Consider a simplified class of circuits like those from Exercises 3.88–3.92: there are two inputs $\{p, q\}$ and at most two gates, each of which is \wedge , \vee , or \neg . There are a total of $2^4 = 16$ distinct propositions over inputs $\{p, q\}$: four different input configurations, each of which can turn the light on or leave it off. Which, if any, of these 16 propositions cannot be expressed using up to two $\{\wedge, \vee, \neg\}$ gates?
 3.94 (programming required.) Consider the class of circuits from Exercises 3.88–3.92: inputs $\{p, q, r\}$, and at most three gates chosen from $\{\wedge, \vee, \neg\}$. There are a total of $2^8 = 256$ distinct propositions over inputs $\{p, q, r\}$: eight different input configurations, each of which can turn the light on or leave it off. Write a program to determine how many of these 256 propositions can be represented by a circuit of this type. (If you design it well, your program will let you check your answers to Exercises 3.88–3.93.)

3-38 Logic




Figure 3.18 An unknown circuit that takes three inputs $\{p, q, r\}$, and either turns on a light bulb (output of the circuit = true) or leaves it off (output = false). The circuit contains at most three gates, each of which is an \wedge , \vee , or \neg gate.

3.95 Consider a set $S = \{p, q, r, s, t\}$ of Boolean variables. Let $\varphi = p \oplus q \oplus r \oplus s \oplus t$. Describe *briefly* the conditions under which φ is true. Use English and, if appropriate, standard (nonlogical) mathematical notation. (Hint: look at the symbol \oplus itself. What's $p + q + r + s + t$, if you treat true as 1 and false as 0 as in Exercises 3.23–3.26?)

3.96 *Dithering* is a technique for converting grayscale images to black-and-white images (for printed media like newspapers). The classic dithering algorithm proceeds as follows. For every pixel in the image, going from top to bottom (“north to south”), and from left to right (“west to east”):

- “Round” the current pixel to black or white. (If it’s closer to black, make it black; if it’s closer to white, make it white.)
- This alteration to the current pixel has some created “rounding error”: we have added some “whiteness units” by making it white, or removed some “whiteness units” by making it black. We compensate for this error by adding/removing the same total number of “whiteness units,” distributed (in a particular way) among the neighboring pixels.

I assigned a dithering exercise in an introductory CS class, and I got, more or less, the code in Figure 3.19 from one student. This code is correct, but it is very repetitious. Reorganize this code so that it’s not so repetitive—in particular, ensure that each “distribute the error” operation (E, S, SE, and SW) appears *only once* in your solution.

Recall Definition 3.16: a proposition φ is in conjunctive normal form (CNF) if φ is the conjunction of one or more clauses, where each clause is the disjunction of one or more literals, and where a literal is an atomic proposition or its negation. Definition 3.17 says that φ is in disjunctive normal form (DNF) if φ is the disjunction of one or more clauses, where each clause is the conjunction of one or more literals. Give a proposition that’s logically equivalent to each of the following, in the specified normal form.

3.97 DNF: $\neg(p \wedge q) \Rightarrow r$

3.101 CNF: $\neg(p \wedge q) \Rightarrow r$

3.98 DNF: $p \wedge (q \vee r) \Rightarrow (q \wedge r)$

3.102 CNF: $p \wedge (q \Rightarrow (r \Rightarrow q \oplus r))$

3.99 DNF: $p \vee \neg(q \Leftrightarrow p \wedge r)$

3.103 CNF: $(p \Rightarrow q) \Rightarrow (q \Rightarrow r \wedge p)$

3.100 DNF: $p \oplus (\neg p \Rightarrow (q \Rightarrow r) \wedge \neg r)$

3.104 CNF: $p \Leftrightarrow (q \vee r \vee \neg p)$

```

1 for y = 1, 2, ..., height:
2   for x = 1, 2, ..., width:
3     if P[x, y] is more white than black then
4       error := "white" - P[x, y]
5       P[x, y] := "white"
6       if x > 1 then
7         if x < width and y < height then distribute E.
8         if x < width and y < height then distribute S, SE, SW, and E.
9         if x < width and y < height then distribute S and SW.
10      else
11        if x < width and y < height then distribute E.
12        if x < width and y < height then distribute S, SE, and E.
13        if x < width and y < height then distribute S.
14      else
15        error := "black" - P[x, y]
16        P[x, y] := "black"
17        [exact duplication of lines 6–13]
```

“Distributing the error” from pixel (x, y) means taking the amount by which $P[x, y]$ changed, and counteracting that change so that the total brightness of the image is unchanged. Specifically, the amount of change (stored in *error*, which may be negative) is distributed $\frac{7}{16}$ to the eastern neighboring pixel, $\frac{5}{16}$ to the south, $\frac{3}{16}$ to the south east, and $\frac{1}{16}$ to the south west. If any of these neighboring pixels don’t exist (because the current pixel is on the border of the image), simply ignore the corresponding fraction of *error* (and don’t add it anywhere).

Distribute E:
increase $P[x + 1, y]$ by $\frac{7}{16} \cdot \text{error}$

Distribute S:
increase $P[x, y + 1]$ by $\frac{5}{16} \cdot \text{error}$

Distribute SE:
increase $P[x + 1, y + 1]$ by $\frac{3}{16} \cdot \text{error}$

Distribute SW:
increase $P[x, y + 1]$ by $\frac{1}{16} \cdot \text{error}$

Figure 3.19 Some (inefficient) dithering code.

Exercises 3-39

A CNF proposition φ is in 3CNF if each clause contains exactly three distinct literals. (Note that p and $\neg p$ are distinct literals.)

Similarly, a proposition φ is in 3DNF if it is the disjunction of one or more clauses, each of which is the conjunction of exactly three distinct literals.

- 3.105** In terms of the number of clauses, what's the smallest 3CNF formula that's a tautology?
- 3.106** In terms of the number of clauses, what's the smallest 3CNF formula that's not satisfiable?
- 3.107** In terms of the number of clauses, what's the smallest 3DNF formula that's a tautology?
- 3.108** In terms of the number of clauses, what's the smallest 3DNF formula that's not satisfiable?

Consider 3CNF propositions over the variables $\{p, q, r\}$ for which no clause appears more than once. (The following exercises are boring if we omit the restriction about repeating clauses; we could repeat the same clause as many times as we please: $(p \vee q \vee r) \wedge (p \vee q \vee r) \wedge (p \vee q \vee r) \cdots$) Two clauses that contain precisely the same literals (in any order) do not count as distinct. (But recall that a single clause can contain both a variable and its negation.)

- 3.109** In terms of the number of clauses, what's the largest 3-variable distinct-clause 3CNF proposition (at all)?
- 3.110** In terms of the number of clauses, what's the largest 3-variable distinct-clause 3CNF proposition that's a tautology?
- 3.111** In terms of the number of clauses, what's the largest 3-variable distinct-clause 3CNF proposition that's satisfiable?

3.4 An Introduction to Predicate Logic

Ont pû tromper quelques hommes, ou les tromper tous dans certains lieux & en certains tems, mais non pas tous les hommes, dans tous les lieux & dans tous les siècles.

One can fool some men, or fool all men in some places and times, but one cannot fool all men in all places and ages.

Jacques Abbadie (c. 1654–1727)

Traité de la Vérité de la Religion Chrétienne (1684)

Propositional logic gives us formal notation to encode Boolean expressions. But these expressions are relatively simple, a sort of “unstructured programming” style of logic. *Predicate logic* is a more general type of logic that allows us to write function-like logical expressions called *predicates*, and to express a broader range of notions than in propositional logic.

3.4.1 Predicates

Informally, a predicate is a property that a particular entity might or might not have; for example, *being a vowel* is a property that some letters have (A, E, ...) and some letters do not have (B, C, ...). A predicate isn't the kind of thing that's true or false, so predicates are different from propositions; rather, a predicate is like a “proposition with blanks” waiting to be filled in. Here are a few examples of predicates:

Example 3.29: Some predicates.

- “The integer ___ is prime.”
- “The string ___ is a palindrome.”
- “The person ___ costarred in a movie with Kevin Bacon.”
- “The string ___ is alphabetically after the string ___.”
- “The integer ___ evenly divides the integer ___.”

Once the blanks of a predicate are filled in, the resulting expression is a proposition. Here are some propositions—some true, some false—derived from the predicates in Example 3.29:

Example 3.30: Some propositions derived from Example 3.29.

- “The integer 57 is prime.”
- “The string TENET is a palindrome.”
- “The person Helen Mirren costarred in a movie with Kevin Bacon.”
- “The string PYTHON is alphabetically after the string PYTHAGOREAN.”
- “The integer 17 evenly divides the integer 42.”

Here is the formal definition of predicates:

Definition 3.20: Predicate.

A *predicate* P is function that assigns the value True or False to each element of a set U . (That is, a predicate P is a function $P : U \rightarrow \{\text{True}, \text{False}\}$.) The set U is called the *universe* or *domain of discourse*, and we say that P is a predicate *over* U .

Although we didn't use the name at the time, we've already encountered predicates, when talked about sets in Chapter 2: Definition 2.20 introduced the notation $\{x \in U : P(x)\}$ to denote the set of those objects $x \in U$ for which P is true. This set abstraction notation “selects” those elements of the set U for which the predicate P is true.

Here are a few more sample predicates based on arithmetic:

Example 3.31: Some more example predicates.

- 1 $\text{isPrime}(n)$: the positive integer n is a prime number.
- 2 $\text{isPowerOf}(n, k)$: the integer n is an exact power of k : that is, $n = k^i$ for some exponent $i \in \mathbb{Z}^{\geq 0}$.
- 3 $\text{onlyPowersOfTwo}(S)$: every element of the set S is an exact power of two.
- 4 $Q(n, a, b)$: the positive integer n satisfies $n = a + b$, and integers a and b are both prime.
- 5 $\text{sumOfTwoPrimes}(n)$: the positive integer n is equal to the sum of two prime numbers.

To highlight Definition 3.20: the isPrime predicate is a function $\text{isPrime} : \mathbb{Z}^{>0} \rightarrow \{\text{True}, \text{False}\}$, which means that the universe for isPrime is $\mathbb{Z}^{>0}$. (Sometimes the universe requires a little more thought: for example, onlyPowersOfTwo answers questions about sets of integers, which means that its universe is the set of sets of integers—in other words, $\mathcal{P}(\mathbb{Z})$.) When the universe is clear from context, we will sometimes be a little sloppy in our notation and not bother writing it down.

Deriving propositions from predicates

Again, by plugging particular values into the predicates from Example 3.31, we get propositions, each of which has a truth value:

Example 3.32: Propositions derived from predicates.

Using the predicates in Example 3.31, let's figure out the truth values of the propositions $\text{isPrime}(261)$, $\text{isPrime}(262)$, $Q(8, 3, 5)$, and $Q(9, 3, 6)$. For each, we'll simply plug the given arguments into the definition of the predicate and figure out the truth value of the resulting proposition.

- 1 A little arithmetic shows that $261 = 3 \cdot 87$; thus $\text{isPrime}(261)$ is false.
- 2 Similarly, we have $262 = 2 \cdot 131$, so $\text{isPrime}(262)$ is false.
- 3 To compute the truth value of $Q(8, 3, 5)$, we plug $n = 8$, $a = 3$, and $b = 5$ into the definition of $Q(n, a, b)$. The proposition $Q(8, 3, 5)$ requires that *the positive integer 8 satisfies $8 = 3 + 5$, and the integers 3 and 5 are both prime*. All of the requirements are met, so $Q(8, 3, 5)$ is true.

3-42 Logic

4 On the other hand, $Q(9, 3, 6)$ is false, because $Q(9, 3, 6)$ requires that $9 = 3 + 6$, and that the integers 3 and 6 are both prime. But 6 isn't prime.

Just as with propositional logical connectives (and as with most functions in your favorite programming language), each predicate takes a fixed number of arguments. So a predicate might be *unary* (taking one argument, like the predicate *isPrime*); or *binary* (taking two arguments, like *isPowerOf*); or *ternary* (taking three arguments, like Q from Example 3.31); and so forth.

Here are a few more examples:

Example 3.33: More propositions derived from predicates.

Using the predicates in Example 3.31, find the truth values of these propositions:

- 1 $sumOfTwoPrimes(17)$
- 2 $sumOfTwoPrimes(34)$
- 3 $isPowerOf(16, 2)$
- 4 $isPowerOf(2, 16)$
- 5 $onlyPowersOfTwo(\{1, 2, 8, 128\})$

Solution. As before, we plug the given values into the definition, and see if the resulting statement is true:

- 1 False. The only way to get an odd number n by adding two prime numbers is for one of those prime numbers to be 2 (the only even prime)—but $17 - 2 = 15$, and 15 isn't prime.
- 2 True, because $34 = 17 + 17$, and 17 is prime. (And the other 17 is prime, too.)
- 3 True, because $2^4 = 16$ (and the exponent 4 is an integer).
- 4 False, because $16^{1/4} = 2$ (and $\frac{1}{4}$ is not an integer).
- 5 True. Every element of $\{1, 2, 8, 128\}$ is a power of two: $\{1, 2, 8, 128\} = \{2^0, 2^1, 2^3, 2^7\}$.

These brief examples may already be enough to begin to give you a sense of the power of logical abstraction that predicates give us: we can now consider the same logical “condition” applied to two different “arguments.” In a sense, propositional logic is like programming without functions; letting ourselves use predicates allows us to write two related propositions using related notation, and to reason simultaneously about multiple propositions—just like writing a function in Java allows you to think simultaneously about the same function applied to different arguments.

Taking it further: Predicates also give a convenient way of representing the state of play of multiplayer games like Tic-Tac-Toe, checkers, and chess. The basic idea is to define a predicate $P(B)$ that expresses “Player 1 will win from board position B if both players play optimally.” For more on this idea, and on the application of logic to playing these kinds of games, see p. 3-54.

3.4.2 Quantifiers

We've seen that we can form a proposition from a predicate by applying that predicate to a particular argument. But we can also form a proposition from a predicate using *quantifiers*, which allow us to formalize

3.4 An Introduction to Predicate Logic 3-43

statements like *every Java program contains at least four for loops* (false!) or *there is a proposition that cannot be expressed using only the connectives \wedge and \vee* (true! See Exercise 4.71). These types of statements are expressed by the two standard quantifiers, the *universal* (“every”) and *existential* (“some”) quantifiers;

Definition 3.21: Universal quantifier [for all, \forall].

Let P be a predicate over S . The proposition $\forall x \in S : P(x)$ is true if, for *every* possible $x \in S$, we have that $P(x)$ is true.

Definition 3.22: Existential quantifier [there exists, \exists].

Let P be a predicate over S . The proposition $\exists x \in S : P(x)$ is true if, for *at least one* possible $x \in S$, we have that $P(x)$ is true.

(The proposition $\forall x \in S : P(x)$ is read “for all x in S , $P(x)$ ” and the proposition $\exists x \in S : P(x)$ is read “there exists an x in S such that $P(x)$.”)

A hint to remember the symbols \forall and \exists : the *for all* notation is \forall , an upside-down ‘A’ as in “all”; the *exists* notation is \exists , a backward ‘E’ as in “exists.” (Annoyingly, they had to be flipped in different directions: a backward ‘A’ is still an ‘A,’ and an upside-down ‘E’ is still an ‘E.’)

Here’s an example of two propositions about prime numbers that use quantifiers:

Example 3.34: Some propositions about primes using quantifiers.

What are the truth values of the following two propositions?

- 1 $\forall n \in \mathbb{Z}^{\geq 2} : \text{isPrime}(n)$
- 2 $\exists n \in \mathbb{Z}^{\geq 2} : \text{isPrime}(n)$

Solution. The proposition $\forall n \in \mathbb{Z}^{\geq 2} : \text{isPrime}(n)$ says “every integer $n \geq 2$ is prime.” This statement is false because, for example, the integer 32 is greater than or equal to 2 and is not prime.

The proposition $\exists n \in \mathbb{Z}^{\geq 2} : \text{isPrime}(n)$, on the other hand, says “there exists an integer $n \geq 2$ that is prime.” This statement is true because, for example, the integer 31 is prime (and $31 \geq 2$).

We can use quantifiers to make precise many intuitive statements. For example, let’s formalize the predicates from Example 3.31:

Example 3.35: The example predicates from Example 3.31, formalized.

$\text{isPrime}(n)$: the positive integer $n \in \mathbb{Z}^{>0}$ is a prime number.

An integer $n \in \mathbb{Z}^{>0}$ is prime if and only if $n \geq 2$ and the only positive integers that evenly divide n are 1 and n itself. (See Definition 2.14.) In other words, we can describe the primality of n in terms of a condition on every candidate divisor d : either $d \in \{1, n\}$, or d doesn’t evenly divide n . Using the

3-44 Logic

“divides” notation (see Definition 2.12), we can formalize $isPrime(n)$ as

$$isPrime(n) = n \geq 2 \wedge [\forall d \in \mathbb{Z}^{\geq 1} : (d \mid n \Rightarrow d = 1 \vee d = n)].$$

$isPowerOf(n, k)$: the integer n is an exact power of k .

An integer n is an exact power of k if and only if n can be written as k^i , for some integer i :

$$isPowerOf(n, k) = \exists i \in \mathbb{Z}^{\geq 0} : n = k^i.$$

$onlyPowersOfTwo(S)$: every element of the set S is an exact power of two.

Because $isPowerOf(n, 2)$ expresses “ n is a power of two,” we can be lazy and just use $isPowerOf$:

$$onlyPowersOfTwo(S) = \forall x \in S : isPowerOf(x, 2).$$

$Q(n, a, b)$: the positive integer n satisfies $n = a + b$, and integers a and b are both prime.

If we use $isPrime$, then formalizing Q actually doesn’t require a quantifier at all:

$$Q(n, a, b) = (n = a + b) \wedge isPrime(a) \wedge isPrime(b).$$

$sumOfTwoPrimes(n)$: the positive integer n is equal to the sum of two prime numbers.

This predicate requires that *there exist* prime numbers a and b that sum to n , and we can use the predicate Q to express this idea:

$$sumOfTwoPrimes(n) = \exists \langle a, b \rangle \in \mathbb{Z} \times \mathbb{Z} : Q(n, a, b).$$

(“There exists a pair of integers $\langle a, b \rangle$ such that $Q(n, a, b)$.”) Alternatively, we could instead write $sumOfTwoPrimes(n)$ by *nesting* one quantifier within the other (see Section 3.5):

$$sumOfTwoPrimes(n) = \exists a \in \mathbb{Z} : [\exists b \in \mathbb{Z} : Q(n, a, b)].$$

Here’s one further example, regarding the *prefix* relationship between two strings:

Example 3.36: Prefixes, formalized.

A binary string $x \in \{0, 1\}^k$ is a *prefix* of the binary string $y \in \{0, 1\}^n$, for $n \geq k$, if y is x with some extra bits added on at the end. For example, 01 and 0110 are both prefixes of 01101010, but 1 is not a prefix of 01101010. If we write $|x|$ and $|y|$ to denote the length of x and y , respectively, then

$$isPrefixOf(x, y) = |x| \leq |y| \wedge [\forall i \in \{i \in \mathbb{Z} : 1 \leq i \leq |x|\} : x_i = y_i].$$

That is, y is no shorter than x , and the first $|x|$ characters of y equal their corresponding characters in x .

Quantifiers as loops

One useful way of thinking about these quantifiers is by analogy to loops in programming. If we ever encounter a $y \in S$ for which $\neg P(y)$ holds, then we immediately know that $\forall x \in S : P(x)$ is false. Similarly, any $y \in S$ for which $Q(y)$ holds is enough to demonstrate that $\exists x \in S : Q(x)$ is true. But if we “loop through” all candidate values of x and fail to encounter any x with $\neg P(x)$ or $Q(x)$, we know that $\forall x \in S : P(x)$ is true or $\exists x \in S : Q(x)$ is false. By this analogy, we might think of the two standard quantifiers as executing the programs in Figure 3.20.

Another intuitive and useful way to think about \forall and \exists is as a supersized version of \wedge and \vee :

$$\begin{aligned} \forall x \in \{x_1, x_2, \dots, x_n\} : P(x) & \quad \exists x \in \{x_1, x_2, \dots, x_n\} : P(x) \\ \equiv P(x_1) \wedge P(x_2) \wedge \dots \wedge P(x_n) & \quad \equiv P(x_1) \vee P(x_2) \vee \dots \vee P(x_n) \end{aligned}$$

The first of these propositions is true only if *every one* of the $P(x_i)$ terms is true; the second is true if *at least one* of the $P(x_i)$ terms is true.

There is one way in which these analogies are loose, though: just as for \sum (summation) and \prod (product) notation (from Section 2.2.7), the loop analogy only makes sense when the domain of discourse is finite! The \forall “program” for a true proposition $\forall x \in \mathbb{Z} : P(x)$ would have to complete an infinite number of iterations before returning True. But the intuition may still be helpful.

Precedence and parenthesization

As in propositional logic, we’ll adopt standard conventions regarding order of operations so that we don’t overdose on parentheses. We treat the quantifiers \forall and \exists as binding tighter than the propositional logical connectives. (For the sake of clarity, though, we’ll err on the side of using too many parentheses in quantified statements, rather than too few.)

Thus $\forall x \in S : P(x) \Rightarrow \exists y \in S : P(y)$ will be understood to mean $[\forall x \in S : P(x)] \Rightarrow [\exists y \in S : P(y)]$. To express the other reading (which involves nested quantifiers; see Section 3.5), we can use parentheses explicitly, by writing $\forall x \in S : [P(x) \Rightarrow \exists y \in S : P(y)]$.




Figure 3.20 For loops analogous to (a) $\forall x \in S : P(x)$, and (b) $\exists x \in S : Q(x)$.

3-46 Logic

Free and bound variables

Consider the variables x and y in the expressions

$$3 \mid x \quad \text{and} \quad \forall y \in \mathbb{Z} : 1 \mid y.$$

Understanding the first of these expressions requires knowledge of what x means, whereas the second is a self-contained statement that can be understood without any outside knowledge. The variable x is called a *free* or *unbound variable*: its value is not fixed by the expression. In contrast, the variable y is a *bound variable*: its value is defined within the expression itself. We say that the quantifier *binds* the variable y , and the *scope* or *body* of the quantifier is the part of the expression in which it has bound y . (We've encountered bound variables before; they arise whenever a variable name is assigned a value within an expression. For example, the variable i is bound in the arithmetic expression $\sum_{i=1}^{10} i^2$, as is the variable n in the set-defining expression $\{n \in \mathbb{Z} : |n| \leq |n^2|\}$.)

A single expression can contain both free variables and bound variables: for example, the expression $\exists y \in \mathbb{Z}^{\geq 0} : x \geq y$ contains a bound variable y and a free variable x . Here's another example:

Example 3.37: Free and bound variables.

Which variables are free in the following expression?

$$[\forall x \in \mathbb{Z} : x^2 \geq y] \wedge [\forall z \in \mathbb{Z} : y = z \vee z^y = 1]$$

Solution. The variable y doesn't appear as the variable bound by either of the quantifiers in this expression, so y is a free variable. Both x and z are bound by the universal quantifiers. (Incidentally, this expression is true if and only if $y = 0$.)

To test whether a particular variable x is free or bound in an expression, we can (consistently) replace x by a different name in that expression. If the meaning stays the same, then x is bound; if the meaning changes, then x is free. For example:

Example 3.38: Testing for free and bound variables.

Consider the following pairs of propositions:

$$\exists x \in S : x > 251 \quad \text{and} \quad \exists y \in S : y > 251 \quad (\text{A})$$

$$x \geq 42x \quad \text{and} \quad y \geq 42y \quad (\text{B})$$

The expressions in (A) express precisely the same condition, namely: *some element of S is greater than 251*. Thus, the variables x and y in these two expressions are bound.

But the expressions in (B) mean different things, in that there are contexts in which these two statements have different truth values (for example, $x = 3$ and $y = -2$). The first expression states a condition on the value of x ; the second states a condition on the value of y . So x is a free variable in " $x \geq 42x$."

3.4 An Introduction to Predicate Logic 3-47

Taking it further: The free-versus-bound-variable distinction is also something that may be familiar from programming, at least in some programming languages. There are some interesting issues in the design and implementation of programming languages that center on how free variables in a function definition, for example, get their values. See p. 3-56.

An expression of predicate logic that contains no free variables is called *fully quantified*. For expressions that are not fully quantified, we adopt a standard convention that any unbound variables in a stated claim are *implicitly* universally quantified. For example, consider these claims:

Claim A: If $x \geq 1$, then $x^2 \leq x^3$.

Claim B: For all $x \in \mathbb{R}$, if $x \geq 1$, then $x^2 \leq x^3$.

When we write a (true) claim like Claim A, we will implicitly interpret it to mean Claim B. (Claim B also explicitly notes \mathbb{R} as the domain of discourse, which was left implicit in Claim A.)

3.4.3 Theorem and Proof in Predicate Logic

Recall that a *tautology* is a proposition that is always true—in other words, it is true no matter what each Boolean variable p in the proposition “means” (that is, whether p is true or false). In this section, we will be interested in the corresponding notion of always-true statements of predicate logic, which are called *theorems*. A statement of predicate logic is “always true” when it’s true no matter what its predicates mean. (Formally, the “meaning” of a predicate P is the set of elements of the universe U for which the predicate is true—that is, $\{x \in U : P(x)\}$.)

Definition 3.23: Theorems in predicate logic.

A fully quantified expression of predicate logic is a *theorem* if and only if it is true for every possible meaning of each of its predicates.

Analogously, two fully quantified expressions are *logically equivalent* if, for every possible meaning of their predicates, the two expressions have the same truth values.

We’ll begin with a pair of related examples, one theorem and one nontheorem:

Example 3.39: A theorem of predicate logic.

Let S be any set. The following claim is true *regardless of what the predicate P denotes*:

$$\forall x \in S : [P(x) \vee \neg P(x)].$$

Indeed, this claim simply says that every $x \in S$ either makes $P(x)$ true or $P(x)$ false. And that assertion is true if the predicate $P(x)$ is “ $x \geq 42$ ” or “ x has red hair” or “ x prefers programming in Python to playing Parcheesi”—indeed, it’s true for any predicate P .

3-48 Logic

Example 3.40: A nontheorem.

Let's show that the following proposition is not a theorem:

$$[\forall x \in S : P(x)] \vee [\forall x \in S : \neg P(x)].$$

A theorem must be true regardless of P 's meaning, so we can establish that this proposition isn't a theorem by giving an example predicate that makes it false. Here's one example: consider the predicate $P = \text{isPrime}$ (with the universe $S = \mathbb{Z}$). Observe that $\forall x \in \mathbb{Z} : \text{isPrime}(x)$ is false because, for example, $\text{isPrime}(4) = \text{False}$; and $\forall x \in \mathbb{Z} : \neg \text{isPrime}(x)$ is false because, for example, $\neg \text{isPrime}(5) = \text{False}$. Thus the given proposition is false when P is isPrime , and so it is not a theorem.

Note the crucial difference between Example 3.39 ("every element of S either makes P true or makes P false") and Example 3.40 ("either every element of S makes P true, or every element of S makes P false"). (Intuitively, it's the difference between "Every letter is either a vowel or a consonant" and "Every letter is a vowel or every letter is a consonant." The former is true; the latter is false.)

Example 3.40 establishes that $[\forall x \in S : P(x)] \vee [\forall x \in S : \neg P(x)]$ fails to be globally true for every meaning of the predicate P , but it *is* true for some meanings of the predicate. For example, if the predicate $P(x)$ is $x^2 \geq 0$ (with $S = \mathbb{R}$), then this disjunction is true (because $\forall x \in \mathbb{R} : x^2 \geq 0$ is true).

The challenge of proofs in predicate logic

The remainder of this section states some theorems of predicate logic, along with an initial discussion of how we might prove that they're theorems. (A *proof* of a statement is simply a convincing argument that the statement is a theorem.) Much of the rest of the book will be devoted to developing and writing proofs of theorems like these, and Chapter 4 will be devoted exclusively to some techniques and strategies for proofs. (This section will preview some of the ideas we'll see there.) Some theorems of predicate logic are summarized in Figure 3.21; we'll prove a few here, and you'll return to some of the others in the exercises.

While predicate logic allows us to express claims that we couldn't state without quantifiers, that extra expressiveness comes with a cost. For a quantifier-free proposition (like all propositions in Sections 3.2 and 3.3), there is a straightforward—if tedious—algorithm to decide whether a given proposition is a tautology: first, build a truth table for the proposition; and, second, check to make sure that the proposition is true in every row. It turns out that the analogous question for predicate logic is much more difficult—in fact, *impossible* to solve in general: there's no algorithm that's guaranteed to figure out whether a given fully quantified expression is a theorem! Demonstrating that a statement in predicate logic is a theorem will require you to *think* in a way that demonstrating that a statement in propositional logic is a tautology did not.

Taking it further: The fact that there's no algorithm guaranteed to determine whether a given proposition is a theorem follows from a mind-numbing 1931 result by Kurt Gödel (1906–1978). See p. 3-58. The absence of such an algorithm sounds like bad news; it means that proving predicate-logic statements is harder, because you can't just use a simple "plug and chug" technique

3.4 An Introduction to Predicate Logic 3-49

$\forall x \in S : [P(x) \vee \neg P(x)]$	
$\neg[\forall x \in S : P(x)] \Leftrightarrow [\exists x \in S : \neg P(x)]$	De Morgan's Laws (quantified form)
$\neg[\exists x \in S : P(x)] \Leftrightarrow [\forall x \in S : \neg P(x)]$	
$[\forall x \in S : P(x)] \Rightarrow [\exists x \in S : P(x)]$	<i>if the set S is nonempty</i>
$\forall x \in \emptyset : P(x)$	Vacuous quantification
$\neg \exists x \in \emptyset : P(x)$	
$[\exists x \in S : P(x) \vee Q(x)] \Leftrightarrow [\exists x \in S : P(x)] \vee [\exists x \in S : Q(x)]$	
$[\forall x \in S : P(x) \wedge Q(x)] \Leftrightarrow [\forall x \in S : P(x)] \wedge [\forall x \in S : Q(x)]$	
$[\exists x \in S : P(x) \wedge Q(x)] \Rightarrow [\exists x \in S : P(x)] \wedge [\exists x \in S : Q(x)]$	
$[\forall x \in S : P(x) \vee Q(x)] \Leftarrow [\forall x \in S : P(x)] \vee [\forall x \in S : Q(x)]$	
$[\forall x \in S : P(x) \Rightarrow Q(x)] \wedge [\forall x \in S : P(x)] \Rightarrow [\forall x \in S : Q(x)]$	
$[\forall x \in \{y \in S : P(y)\} : Q(x)] \Leftrightarrow [\forall x \in S : P(x) \Rightarrow Q(x)]$	
$[\exists x \in \{y \in S : P(y)\} : Q(x)] \Leftrightarrow [\exists x \in S : P(x) \wedge Q(x)]$	
$\varphi \wedge [\exists x \in S : P(x)] \Leftrightarrow [\exists x \in S : \varphi \wedge P(x)]$	<i>if x does not appear as a free variable in φ</i>
$\varphi \vee [\forall x \in S : P(x)] \Leftrightarrow [\forall x \in S : \varphi \vee P(x)]$	<i>if x does not appear as a free variable in φ</i>

Figure 3.21 A few theorems involving quantification.

to figure out whether a given statement is actually always true. But this fact is also precisely the reason that creativity plays a crucial role in proofs and in theoretical computer science more generally—and why, arguably, proving things can be fun! (For me, this difference is exactly why I find Sudoku less interesting than crossword puzzles: when there's no algorithm to solve a problem, we have to embrace the creative challenge in attacking it.)

3.4.4 A Few Examples of Theorems and Proofs

In the rest of this section, we will see a few further theorems of predicate logic, with proofs. As we've said, there's no formulaic approach to prove these theorems; we'll need to employ a variety of strategies in this endeavor.

Negating quantifiers: a first example

Suppose that your egomaniacal, overconfident partner from Intro CS wanders into the lab and says *For any array A that you give me, partner, my implementation of insertion sort correctly sorts A*. You know, though, that your partner is wrong. (You spot a bug in his egomaniacal code.) What would that mean? Well, you might reply, gently but firmly: *There's an array A for which your implementation of insertion sort does not correctly sort A*. The equivalence that you're using is a theorem of predicate logic:

Example 3.41: Negating universal quantifiers.

Let's prove the equivalence you're using to debunk your partner's claim:

$$\neg[\forall x \in S : P(x)] \Leftrightarrow [\exists x \in S : \neg P(x)]. \quad (*)$$

3-50 Logic

It's probably easiest to view $(*)$ as a quantified version of the tautology $\neg(p \wedge q) \Leftrightarrow \neg p \vee \neg q$, which was one of De Morgan's Laws from propositional logic. Here's the (oversimplified) intuition:

$$\begin{aligned} \neg[\forall x \in S : P(x)] &\approx \neg[P(x_1) \wedge P(x_2) \wedge \cdots \wedge P(x_n)] && \text{if we imagine the universe } S = \{x_1, x_2, \dots, x_n\}, \text{ then} \\ &\quad \forall x \in S : P(x) \text{ means } P(x_1) \wedge P(x_2) \wedge \cdots \wedge P(x_n) \\ &\equiv [\neg P(x_1) \vee \neg P(x_2) \vee \cdots \vee \neg P(x_n)] && \text{by De Morgan's Laws (propositional version)} \\ &\approx \exists x \in S : \neg P(x), && \exists x \in S : Q(x) \text{ means } Q(x_1) \vee Q(x_2) \vee \cdots \vee Q(x_n) \end{aligned}$$

There is something slightly more subtle about $(*)$ because S might be infinite, but the idea is the same:

- If there's an $a \in S$ such that $P(a) = \text{False}$, then $\exists x \in S : \neg P(x)$ is true (because a is an example) and $\forall x \in S : P(x)$ is false (because a is a counterexample).
- On the other hand, if $P(a) = \text{True}$ for every $a \in S$, then we know that $\exists x \in S : \neg P(x)$ is false and $\forall x \in S : P(x)$ is true.

The analogous claim for the negation of $\exists x \in S : P(x)$ is also a theorem:

Example 3.42: Negating existential quantifiers.

Let's prove that this claim is a theorem, too:

$$\neg[\exists x \in S : P(x)] \Leftrightarrow [\forall x \in S : \neg P(x)]. \quad (\dagger)$$

To see that (\dagger) is true for an arbitrary predicate P , we start with claim $(*)$ from Example 3.41, but using the predicate $Q(x) = \neg P(x)$. (Note that Q is also a predicate—so Example 3.41 holds for Q too!) Thus we know that

$$\neg[\forall x \in S : Q(x)] \Leftrightarrow [\exists x \in S : \neg Q(x)],$$

and, because $p \Leftrightarrow q \equiv \neg p \Leftrightarrow \neg q$, we therefore also know that

$$[\forall x \in S : Q(x)] \Leftrightarrow \neg[\exists x \in S : \neg Q(x)].$$

But $Q(x)$ is just $\neg P(x)$ and $\neg Q(x)$ is just $P(x)$, by definition of Q , and so we have

$$[\forall x \in S : \neg P(x)] \Leftrightarrow \neg[\exists x \in S : P(x)].$$

Thus we've now shown that (\dagger) is true for any predicate P , so it is a theorem.

All implies some: a proof of an implication

The entirety of Chapter 4 is devoted to proofs and proof techniques; there's lots more there about how to approach proving or disproving new claims. But here we'll preview a particularly useful proof strategy for proving an implication, and use it to establish another theorem of predicate logic:

3.4 An Introduction to Predicate Logic 3-51

Definition 3.24: Proof by assuming the antecedent.

Suppose that we must prove an implication $\varphi \Rightarrow \psi$. Because the only way for $\varphi \Rightarrow \psi$ to *fail* to be true is for φ to be true and ψ to be false, we will prove that the implication $\varphi \Rightarrow \psi$ is always true by ruling out the one scenario in which it wouldn't be. Specifically, we *assume* that φ is true, and then *prove* that ψ must be true too, under this assumption.

(Recall from the truth table of \Rightarrow that the only way for the implication $\varphi \Rightarrow \psi$ to be false is when φ is true but ψ is false. Also recall that φ is called the *antecedent* of the implication $\varphi \Rightarrow \psi$; hence this proof technique is called *assuming the antecedent*.) Here are two examples of proofs that use this technique:

(An example from propositional logic.) Let's prove that $p \Rightarrow p \vee q$ is a tautology: we assume that the antecedent p is true, and we must prove that the consequent $p \vee q$ is true too. But that's obvious, because p is true (by our assumption), and $\text{True} \vee q \equiv \text{True}$.

(An example from arithmetic.) Let's prove that *if x is a perfect square, then $4x$ is a perfect square*: assume that x is a perfect square, that is, assume that $x = k^2$ for an integer k . Then $4x = 4k^2 = (2k)^2$ is a perfect square too, because $2k$ is also an integer.

Finally, here's a theorem of predicate logic that we can prove using this technique:

Example 3.43: If everybody's doing it, then somebody's doing it.

Consider the following proposition, for an arbitrary nonempty set S :

$$[\forall x \in S : P(x)] \Rightarrow [\exists x \in S : P(x)].$$

We'll prove this claim by assuming the antecedent. Specifically, we assume $\forall x \in S : P(x)$, and we need to prove that $\exists x \in S : P(x)$.

Because the set S is nonempty, we know that there's at least one element $a \in S$. By our assumption, we know that $P(a)$ is true. But because $P(a)$ is true, then it's immediately apparent that $\exists x \in S : P(x)$ is true too—because we can just pick $x = a$.

Problem-solving tip: When you're facing a statement that contains a lot of mathematical notation, try to understand it by rephrasing it as an English sentence. Restating the assertion from Example 3.43 in English makes it pretty obvious that it's true: *if everyone in S satisfies P —and there's actually someone in S —then of course someone in S satisfies P !*

Vacuous quantification

Consider the proposition *All even prime numbers greater than 12 have a 3 as their last digit*. Write N to denote the set of all even prime numbers greater than 12; formalized, then, this claim can be written as $\forall n \in N : n \bmod 10 = 3$. Is this claim true or false? It has to be true! The point is that N actually contains no elements (there *are* no even prime numbers other than 2, because an even number is by definition divisible

3-52 Logic

by 2). Thus this claim says: “for every $n \in \emptyset$, some silly thing is true of n .” But there *is* no n in \emptyset , so the claim has to be true! The general statement of the theorem is

$$\forall x \in \emptyset : P(x). \quad (\ddagger)$$

Quantification over the empty set is called *vacuous quantification*, and (\ddagger) is said to be *vacuously true*.

Here’s another way to see that (\ddagger) is a theorem, using the De Morgan–like view of quantification. What would it mean for (\ddagger) to be false? There would have to be some $x \in \emptyset$ such that $\neg P(x)$ —but there never exists *any* element $x \in \emptyset$, let alone an $x \in \emptyset$ such that $\neg P(x)$. Thus $\exists x \in \emptyset : \neg P(x)$ is false, and therefore its negation $\neg \exists x \in \emptyset : \neg P(x)$, which is equivalent to (\ddagger) , is true.

Disjunctions and quantifiers

Here’s one last example for this section, in which we’ll figure out when the “or” of two quantified statements can be expressed as one single quantified statement:

Example 3.44: Disjunctions and quantifiers.

Consider the following two propositions, for an arbitrary set S :

$$[\forall x \in S : P(x) \vee Q(x)] \Leftrightarrow [\forall x \in S : P(x)] \vee [\forall x \in S : Q(x)] \quad (\text{A})$$

$$[\exists x \in S : P(x) \vee Q(x)] \Leftrightarrow [\exists x \in S : P(x)] \vee [\exists x \in S : Q(x)] \quad (\text{B})$$

Is either (A) or (B) a theorem? Both? Prove your answers.

Solution. Claim (B) is a theorem. To prove it, we’ll show that the left-hand side implies the right-hand side, and vice versa. (That is, we’re proving $p \Leftrightarrow q$ by proving both $p \Rightarrow q$ and $q \Rightarrow p$, which is a legitimate proof because $p \Leftrightarrow q \equiv (p \Rightarrow q) \wedge (q \Rightarrow p)$.) Both proofs will use the technique of assuming the antecedent.

First, let’s prove that $[\exists x \in S : P(x) \vee Q(x)]$ implies $[\exists x \in S : P(x)] \vee [\exists x \in S : Q(x)]$:

Suppose that $[\exists x \in S : P(x) \vee Q(x)]$ is true. Then there is some particular $x^* \in S$ for which either $P(x^*)$ or $Q(x^*)$. But in either case, we’re done: if $P(x^*)$ then $\exists x \in S : P(x)$ because x^* satisfies the condition; if $Q(x^*)$ then $\exists x \in S : Q(x)$, again because x^* satisfies the condition.

Second, let’s prove that $[\exists x \in S : P(x)] \vee [\exists x \in S : Q(x)]$ implies $[\exists x \in S : P(x) \vee Q(x)]$:

Suppose that $[\exists x \in S : P(x)] \vee [\exists x \in S : Q(x)]$ is true. Thus either there’s an $x^* \in S$ such that $P(x^*)$ or an $x^* \in S$ such that $Q(x^*)$. That x^* suffices to make the left-hand side of (B) true.

On the other hand, (A) is not a theorem, for much the same reason as in Example 3.40. (In fact, if $Q(x)$ is $\neg P(x)$, then Examples 3.39 and 3.40 directly establish that (A) is not a theorem.) The set \mathbb{Z} and the

3.4 An Introduction to Predicate Logic 3-53

predicates *isOdd* and *isEven* make (A) false: the left-hand side is true (“all integers are either even or odd”) but the right-hand side is false (“either (i) all integers are even, or (ii) all integers are odd”).

Problem-solving tip: In thinking about whether a particular quantified expression is a theorem—like the question of whether (A) from Example 3.44 is a theorem—it’s often useful to get intuition by plugging in a few sample values for the universe S and the predicates P and Q . It’s especially helpful to consider sample values that are relatively easy to describe and think about, like “odd” and “even.”

Although the mutual implication (A) from Example 3.44 is not a theorem, one direction of it is. Let’s prove this implication as another example:

Example 3.45: Disjunction, quantifiers, and one-way implications.

The \Leftarrow direction of (A) from Example 3.44 is a theorem:

$$[\forall x \in S : P(x) \vee Q(x)] \Leftarrow [\forall x \in S : P(x)] \vee [\forall x \in S : Q(x)].$$

Let’s prove it. We assume the antecedent $[\forall x \in S : P(x)] \vee [\forall x \in S : Q(x)]$. This assumption means that either $\forall x \in S : P(x)$ or $\forall x \in S : Q(x)$.

If we’re in the first case, in which $\forall x \in S : P(x)$, then $P(x)$ is true for every $x \in S$. But then it’s certainly the case that $P(x) \vee Q(x)$ is true for any $x \in S$ too: if $P(x)$ is true, then $P(x) \vee \text{whatever}$ is true too.

If we’re in the case in which $\forall x \in S : Q(x)$ holds, then $Q(x)$ is true for any $x \in S$. Similarly, then, $P(x) \vee Q(x)$ is true for any $x \in S$: if $Q(x)$ is true, then $\text{whatever} \vee Q(x)$ is true too.

In both cases, we’ve argued that $P(x) \vee Q(x)$ is true for all $x \in S$ —that is, that $[\forall x \in S : P(x) \vee Q(x)]$.

You’ll have a chance to consider a number of other theorems of predicate logic in the exercises, including the \wedge -analogy to Examples 3.44 and 3.45 (in Exercises 3.137 and 3.138).

COMPUTER SCIENCE CONNECTIONS

GAME TREES, LOGIC, AND WINNING TIC-TAC(-TOE)

In 1997, Deep Blue, a chess-playing program developed by IBM, beat the chess Grandmaster Garry Kasparov in a six-match series. This event was a turning point in the public perception of computation and artificial intelligence (AI); it was the first time that a computer had outperformed the best humans at something that most people tended to identify as a “human endeavor.” Ten years later, a research group developed a program called Chinook, a perfect checkers-playing system: from any game position arising in its games, Chinook chooses *the* best possible legal move. (And a few years after that came another major turning point in the public perception of computation and AI: IBM Watson, designed to play the quiz show *Jeopardy!*, defeated some of the best-ever human players in a match.)

While chess and checkers are very complicated games, the basic ideas of playing them—ideas based on logic—are shared with simpler games. Consider *Tic-Tac*, a 2-by-2 version of Tic-Tac-Toe. (Thanks to Jon Kleinberg for suggesting the game.) See Figure 3.22.

Note that—unless O is tremendously dull—O will win the game, but we will use a *game tree* (as in Figure 3.22a), which represents all possible moves, to systematize this reasoning. Here’s the basic idea. Let’s define a predicate $P(B)$ to mean “Player O wins under optimal play starting from board B .” (“Optimal play” means that both Player O and Player X always make the very best possible move in every turn.) For example, $P(\begin{smallmatrix} \text{X} & \text{O} \\ \text{O} & \text{O} \end{smallmatrix}) = \text{True}$ because O has already won; and $P(\begin{smallmatrix} \text{O} & \text{X} \\ \text{X} & \text{O} \end{smallmatrix}) = \text{False}$ because the game has ended in a draw. The answer to the question “does O win Tic-Tac if both players play optimally?” is the truth value of $P(\begin{smallmatrix} _ & _ \\ _ & _ \end{smallmatrix})$. If it’s O’s turn in board B , then $P(B)$ is true if and only if *there exists* a possible move for O leading to a board B' in which $P(B')$; if it’s X’s turn, then $P(B)$ is true if and only if *every* possible move made by X leads to a board B' in which $P(B')$. So

$$\begin{aligned} P(\begin{smallmatrix} _ & \text{O} \\ _ & _ \end{smallmatrix}) & \quad \text{and} \quad P(\begin{smallmatrix} _ & _ \\ _ & _ \end{smallmatrix}) \\ = P(\begin{smallmatrix} \text{X} & \text{O} \\ _ & _ \end{smallmatrix}) \wedge P(\begin{smallmatrix} _ & \text{O} \\ \text{X} & _ \end{smallmatrix}) \wedge P(\begin{smallmatrix} _ & _ \\ \text{O} & \text{X} \end{smallmatrix}) & \quad = P(\begin{smallmatrix} \text{O} & _ \\ _ & _ \end{smallmatrix}) \vee P(\begin{smallmatrix} \text{O} & _ \\ _ & _ \end{smallmatrix}) \vee P(\begin{smallmatrix} _ & _ \\ \text{O} & _ \end{smallmatrix}) \vee P(\begin{smallmatrix} _ & _ \\ _ & \text{O} \end{smallmatrix}). \end{aligned}$$

The game tree, labeled appropriately, is shown in Figure 3.22b. If we view the truth values from the leaves as “bubbling up” from the bottom of the tree, then a board B gets assigned the truth value True if and only if Player O can guarantee a win from the board B .

Some serious complications arise in writing a program to play more complicated games like checkers or chess. Here are just a few of the issues that one must confront in building a system like Deep Blue or Chinook. First, there are $\approx 500,000,000,000,000,000,000$ different checkers positions—and $\approx 10^{40}$ chess positions!—so we can’t afford to

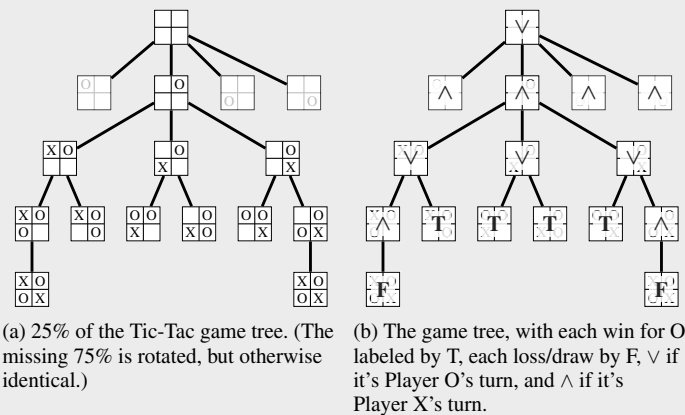


Figure 3.22 Tic-Tac. Two players, O and X, make alternate moves, starting with O; a player wins by occupying a complete row or column. Diagonals don't count, and if the board is filled without O or X winning, then the game is a draw.

3.4 An Introduction to Predicate Logic 3-55

represent them all. (Luckily, our system only needs to find *one* good move at each stage, and we can choose moves so most positions are never reached.) Second, approximately one bit per trillion is written incorrectly *merely in copying data on current hard disk technologies*. So a program constructing a massive structure like the checkers game tree must “check its work.” Third, for a game as big as chess, we can’t afford to compute all the way to the bottom of the tree; instead, we *estimate* the quality of each position after computing a handful of layers deep in the game tree. (There are many more complications, too. For more on game trees and algorithms for exploring large search spaces, see any good AI text, like [110]. You can also read more specifically about Deep Blue [26], or about Chinook [112].)

COMPUTER SCIENCE CONNECTIONS

NONLOCAL VARIABLES AND LEXICAL VS. DYNAMIC SCOPING

In a function f written in a programming language—say, C or Python, though nearly any language has the relevant features—we can use several different types of variables that store values. There are *local variables*, whose values are defined completely within the body of f . There are *parameters* to the function, inputs to f whose value is specified when f is invoked. And there are *nonlocal variables*, which get their value from other contexts. The most common type of these “other” variables is a *global variable*, which persists throughout the execution of the entire program.

For an example function (written in both C and Python as illustrative examples) that uses all three types of variables, see Figure 3.23. (The variable a is a parameter, the variable `result` is a local variable, and the variable b is a nonlocal variable.)

In the body of the `addB` function, the variable a is a *bound* variable; specifically, it is bound when the function is invoked with an actual parameter. But the variable b is *unbound*. (Just as with a quantified expression, an unbound variable is one for which the meaning of the function could change if we replaced that variable with a different name. If we changed the a to an x in both Line 1 and Line 2, then the function would behave identically, but if we changed the b to a y , then the function would behave differently.) In this function, the variable b has to somehow get a value from somewhere if we are going to be able to invoke the function `addB` without causing an error. Often b will be a global variable, but it is also possible in Python or C (with appropriate compiler settings) to *nest* function definitions—just as quantifiers can be nested. (See Section 3.5.)

One fundamental issue in the design and implementation in programming languages is illustrated in Figure 3.24. Suppose x is an unbound variable in the definition of a function f . Generally, programming languages either use *lexical scope*, where x ’s value is found by looking “outward” where f is *defined*; or *dynamic scope*, where x ’s value is found by looking where f is *called*.

Almost all modern programming languages use lexical scope, though *macros* in C and other languages use dynamic scope. (And some languages only use dynamic scope.) While we’re generally used to lexical scope and therefore it feels more intuitive, there are some circumstances in which macros can be tremendously useful and convenient.

(The differences between lexical versus dynamic scope, and other related issues, are classical topics in

```
1 int addB(int a) {
2     int result = a + b;
3     return result;
4 }
```

```
1 def addB(a):
2     result = a + b
3     return result
```

Figure 3.23 A function `addB` written in C and analogous function `addB` written in Python. Here `addB` takes one (integer) parameter a , accesses a nonlocal variable b , defines a local variable `result`, and returns $a + b$.

```
1 int b = 17;
2
3 int addB(int a) {
4     return a + b;
5 }
6
7 int test() {
8     int b = 128;
9     return addB(3);
10 }
11
12 /* Here addB() is a function, and */
13 test(3); /* the value of test(3) is 20. */
```

A function in C finds values for unbound variables in the defining environment.

```
13 int b = 17;
14
15 #define addB(a) a + b
16
17 int test() {
18     int b = 128;
19     return addB(3);
20 }
21
22 /* Here addB() is a macro, and */
23 test(3); /* the value of test(3) is 131. */
```

A macro in C finds values for unbound variables in the calling environment.

Figure 3.24 Two C snippets defining `addB`, where the nonlocal variable b gets its value from different places.

3.4 An Introduction to Predicate Logic 3-57

the design and implementation of programming languages. One of the other interesting issues is that there are actually multiple paradigms for passing parameters to a function, too; we're discussing *call-by-value* parameter passing, which probably is the most common, but there are other choices. See any good textbook on programming languages for much more.)

COMPUTER SCIENCE CONNECTIONS

GÖDEL'S INCOMPLETENESS THEOREM

Given a fully quantified proposition φ , is φ a theorem? This apparently simple question drove the development of some of the most profound and mind-numbing results of the last hundred years. In the early 20th century, there was great interest in the “formalist program,” advanced especially by the German mathematician David Hilbert (1862–1943). The formalist approach aimed to turn all of mathematical reasoning into a machine: one could feed in a mathematical statement φ as input, turn a hypothetical crank, and the machine would spit out a proof or disproof of φ as output. But this program was shattered by two closely related results—two of the greatest intellectual achievements of the 20th century.

The first blow to the formalist program was the proof by Kurt Gödel (1906–1978), in 1931, of what became known as *Gödel's Incompleteness Theorem*. Gödel's incompleteness theorem is based on the following two important and desirable properties of logical systems. First, a logical system is called *consistent* if only true statements can be proven. (In other words, if there is a proof of φ in the system, then φ is true.) Second, a logical system is called *complete* if every true statement can be proven. (In other words, if φ is true, then there is a proof of φ in the system.)

See Figure 3.25.

Both consistency and completeness are Good Things for a logical system. If the system is inconsistent, then there is a false statement φ that can be proven (which means that anything can be proven, as false implies anything!). And if the system is incomplete, then there is a true statement φ that cannot be proven. But what Gödel was able to prove is this troubling result:

Gödel's (First) Incompleteness Theorem:

Any sufficiently powerful logical system is either inconsistent or incomplete.

(Here “sufficiently powerful” just means “capable of expressing multiplication”; predicate logic as described here is certainly “sufficiently powerful.”) Gödel's proof proceeds by constructing a self-referential logical expression φ that means “ φ is not provable.” (So if φ is true, then the system is incomplete; and if φ is false, then the system is inconsistent.)

The second strike against the formalist program was the proof of the *undecidability of the halting problem*, shown independently by Alan Turing and Alonzo Church in the 1930s. We can think of the halting problem as asking the following question: given a function f written in Python and an input x , does running $f(x)$ get stuck in an infinite loop?

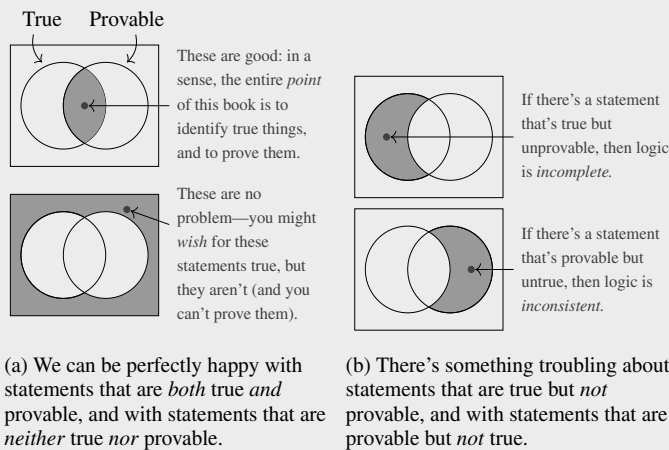


Figure 3.25 A Venn diagram dividing the universe of logical statements into those that are *true* (or not), and those that are *provable* (or not).

3.4 An Introduction to Predicate Logic 3-59

(Or does it eventually terminate?) The *undecidability* of this problem means that *there is no algorithm that solves the halting problem*. A corollary of this result is that our problem—given a fully quantified proposition φ , is φ a theorem?—is also undecidable. We'll discuss uncomputability in detail in Section 4.4.4.

If you're interested in reading more: undecidability, incompleteness, and their profound consequences are the focus of a number of excellent textbooks on the theory of computation (for example, [73, 120])—and also Douglas Hofstadter's fascinating masterpiece *Gödel, Escher, Bach* [56]. There's also a compelling biography of Kurt Gödel, who led a deeply troubled life in many ways despite his professional success, that describes much more about the person behind the result [23].

3-60 Logic

EXERCISES

Using the characteristics in Figure 3.26, define a predicate that's true for each of the following lists of languages, and false for every other language in the table. For example, the predicate $P(x) = "x \text{ has strong typing and } x \text{ is not functional}"$ makes $P(\text{Pascal})$ and $P(\text{Java})$ true, and makes $P(x)$ false for every $x \in \{C, C++, \text{\LaTeX}, \text{ML}, \text{Perl}, \text{Scheme}\}$.

3.112 Java

3.113 C, Pascal

3.114 ML, Pascal

3.115 Pascal, Scheme, Perl

3.116 \LaTeX , Java, C++, Perl

3.117 C, Pascal, ML, C++, \LaTeX , Scheme, Perl

Examples 3.4 and 3.16 construct a proposition corresponding to “the password contains at least three of four character types (digits, lowercase letters, uppercase letters, other).” In that example, we took “the password contains at least one digit” (and its analogues) as an atomic proposition. But we could give a lower-level characterization of valid passwords. Let isDigit , isLower , and isUpper be predicates that are true of single characters of the appropriate type. Use standard arithmetic notation and these predicates to formalize the following conditions on a password $x = \langle x_1, \dots, x_n \rangle$, where x_i is the i th character in the password:

3.118 x is at least 8 characters long.

3.119 x contains at least one lowercase letter.

3.120 x contains at least one non-alphanumeric character. (Remember: isDigit , isLower , and isUpper are the only predicates available!)

3.121 (Inspired by a letter to the editor in *The New Yorker* by Alexander George [51].) Steve Martin, the great comedian, reports in *Born Standing Up: A Comic's Life* [85] that, inspired by Lewis Carroll, he started closing his shows with the following line. (It got big laughs.) “I’m not going home tonight; I’m going to Bananaland, a place where only two things are true, only two things: One, all chairs are green; and two, no chairs are green.” Steve Martin describes the joke as a contradiction—but, in fact, these two true things are not contradictory! Explain. How can “all chairs in Bananaland are green” and “no chairs in Bananaland are green” be simultaneously true?

As a rough approximation, we can think of a database as a two-dimensional table, where rows correspond to individual entities, and columns correspond to fields (data about those entities). A database query defines a predicate $Q(x)$ that consists of tests of the values from various columns, joined by the basic logical connectives. The database system then returns a list of rows/entities for which the predicate is true. We can think of this type of database access as involving predicates: in response to query Q , the system returns the list of all rows x for which $Q(x)$ is true. An example is shown in Figure 3.27. Each of the following predicates $Q(x)$ uses tests on particular columns in x 's row. For each, give a logically equivalent predicate in which each column's name appears at most once. You may also use the symbols True and False and the four most common logical connectives (\neg , \wedge , \vee , \Rightarrow) as many times as you please. Use a truth table to prove that your answer is logically equivalent to the given predicate.

3.122 $[\text{age}(x) < 18] \vee ([\text{age}(x) \geq 18] \wedge [\text{gpa}(x) \geq 3.0])$

3.123 $\text{takenCS}(x) \Rightarrow \neg ([\text{home}(x) = \text{Hawaii}] \Rightarrow ([\text{home}(x) = \text{Hawaii}] \wedge \text{takenCS}(x)))$

	paradigm	typing	scope	(continued)	paradigm	typing	scope
C	imperative	weak	lexical	ML	functional	strong	lexical
C++	object-oriented	weak	lexical	Pascal	imperative	strong	lexical
Java	object-oriented	strong	lexical	Perl	scripting	weak	either
\LaTeX	scripting	weak	dynamic	Scheme	functional	weak	either

Figure 3.26 Some well-known programming languages, with some characteristics for each.

name	GPA	CS taken?	home	age	school year	on campus?	has a major?
Alice	4.0	yes	Alberta	20	3	yes	yes
Bob	3.14	yes	Bermuda	19	3	yes	no
Charlie	3.54	no	Cornwall	18	1	no	yes
Desdemona	3.8	yes	Delaware	17	2	no	no

Figure 3.27 A sample database. Here, to find a list of all students with grade point averages over 3.4 who have taken at least one CS course if they're from Hawaii, we could query $[GPA(x) \geq 3.4] \wedge (takenCS(x) \Leftarrow [home(x) = Hawaii])$. For this database, this query would return Charlie (and not Alice, Bob, or Desdemona).

$$3.124 \quad (hasMajor(x) \wedge [schoolYear(x) \neq 3] \wedge onCampus(x)) \vee (hasMajor(x) \wedge [schoolYear(x) \neq 3] \wedge \neg onCampus(x)) \\ \vee (hasMajor(x) \wedge [schoolYear(x) = 3] \wedge \neg onCampus(x))$$

3.125 Following the last few exercises, you might begin to think that any query can be rewritten without duplication. Can it? Consider a unary predicate that is built up from the predicates $P(x)$ and $Q(x)$ and the propositional symbols $\{\text{True, False, } \wedge, \vee, \neg, \Rightarrow\}$. Decide whether the following claim is true or false, and prove your answer:

Claim: Every such predicate is logically equivalent to a predicate that uses only the following symbols:

- (i) $\{\text{True, False, } \wedge, \vee, \neg, \Rightarrow\}$, all of which can be used as many times as you please; and
- (ii) the predicates $\{P(x), Q(x)\}$, which can appear *only one time each*.

Search engines allow users to specify Boolean conditions in their queries. For example, “social OR networks” will return only web pages containing either the word “social” or the word “networks.” You can view a query as a predicate Q : the search engine returns (in some order) the list of pages p for which $Q(p)$ is true. Consider the following queries:

Query A: “java AND program AND NOT computer”

Query B: “(computer OR algorithm) AND java”

Query C: “java AND NOT (computer OR algorithm OR program)”

Give an example of a web page—or a sentence—that would be returned as specified in these problems:

3.126 Returned by query A but not by B or C.

3.127 Returned by query B but not by A or C.

3.128 Returned by query C but not by A or B.

3.129 Returned by query A and B but not by C.

3.130 Prove or disprove: $[\forall n \in \mathbb{Z} : isPrime(n) \Rightarrow \frac{n}{2} \notin \mathbb{Z}]$.

3.131 Translate this quote (attributed to Groucho Marx) into logical notation: *It isn't necessary to have relatives in Kansas City in order to be unhappy.* Let $P(x)$ be “x has relatives in Kansas City” and $Q(x)$ be “x is unhappy,” and view the statement as implicitly making a claim that a particular kind of person exists.

Write an English sentence that expresses the logical negation of each given sentence. (Don't just say “It is not the case that ...”; give a genuine negation.) Some of the given sentences are ambiguous in their meaning; if so, describe all of the interpretations of the sentence that you can find, then choose one and give its negation.

3.132 Every entry in the array A is positive.

3.133 Every decent programming language denotes block structure with parentheses or braces.

3.134 There exists an odd number that is evenly divisible by a different odd number.

3.135 There is a point in Minnesota that is farther than ten miles from a lake.

3.136 Every sorting algorithm takes at least $n \log n$ steps on some n -element input array.

3-62 Logic

In Examples 3.44 and 3.45, we proved that

$$\begin{aligned} [\exists x \in S : P(x) \vee Q(x)] &\Leftrightarrow [\exists x \in S : P(x)] \vee [\exists x \in S : Q(x)] \\ [\forall x \in S : P(x) \vee Q(x)] &\Leftarrow [\forall x \in S : P(x)] \vee [\forall x \in S : Q(x)] \end{aligned}$$

are theorems. Argue that the following \wedge -analogies to these statements are also theorems:

$$\begin{aligned} \mathbf{3.137} \quad & [\exists x \in S : P(x) \wedge Q(x)] \Rightarrow [\exists x \in S : P(x)] \wedge [\exists x \in S : Q(x)] \\ \mathbf{3.138} \quad & [\forall x \in S : P(x) \wedge Q(x)] \Leftrightarrow [\forall x \in S : P(x)] \wedge [\forall x \in S : Q(x)] \end{aligned}$$

Explain why the following are theorems of predicate logic:

$$\begin{aligned} \mathbf{3.139} \quad & [\forall x \in S : P(x) \Rightarrow Q(x)] \wedge [\forall x \in S : P(x)] \Rightarrow [\forall x \in S : Q(x)] \\ \mathbf{3.140} \quad & [\forall x \in \{y \in S : P(y)\} : Q(x)] \Leftrightarrow [\forall x \in S : P(x) \Rightarrow Q(x)] \\ \mathbf{3.141} \quad & [\exists x \in \{y \in S : P(y)\} : Q(x)] \Leftrightarrow [\exists x \in S : P(x) \wedge Q(x)] \end{aligned}$$

Explain why the following propositions are theorems of predicate logic, assuming that x does not appear as a free variable in the expression φ (and assuming that S is nonempty):

$$\begin{aligned} \mathbf{3.142} \quad & \varphi \Leftrightarrow [\forall x \in S : \varphi] \\ \mathbf{3.143} \quad & \varphi \vee [\forall x \in S : P(x)] \Leftrightarrow [\forall x \in S : \varphi \vee P(x)] \\ \mathbf{3.144} \quad & \varphi \wedge [\exists x \in S : P(x)] \Leftrightarrow [\exists x \in S : \varphi \wedge P(x)] \\ \mathbf{3.145} \quad & (\varphi \Rightarrow [\exists x \in S : P(x)]) \Leftrightarrow [\exists x \in S : \varphi \Rightarrow P(x)] \\ \mathbf{3.146} \quad & ([\exists x \in S : P(x)] \Rightarrow \varphi) \Leftrightarrow [\forall x \in S : P(x) \Rightarrow \varphi] \end{aligned}$$

3.147 Give an example of a predicate P , a nonempty set S , and an expression φ containing x as a free variable such that the proposition from Exercise 3.143 is false. Because x has to get its meaning from somewhere, we will imagine a universal quantifier for x wrapped around the entire expression. Specifically, give an example of P , φ , and S for which

$$\forall x \in S : [\varphi \vee [\forall x \in S : P(x)]] \quad \text{is not logically equivalent to} \quad \forall x \in S : [[\forall x \in S : \varphi \vee P(x)]] .$$

3.5 Predicate Logic: Nested Quantifiers

Who knows what beautiful and winged life, whose egg has been buried for ages under many concentric layers of woodenness in the dead dry life of society ... may unexpectedly come forth from amidst society's most trivial and handselled furniture, to enjoy its perfect summer life at last!

Henry David Thoreau (1817–1862)

Walden; or, Life in the Woods (1854)

Just as we can place one loop inside another in a program, we can place one quantified statement inside another in predicate logic. In fact, the most interesting quantified statements almost always involve more than one quantifier. (For example: *during every semester, there's a computer science class that every student on campus can take*.) In formal notation, such a statement typically involves *nested quantifiers*—that is, multiple quantifiers in which one quantifier appears inside the scope of another. We've encountered statements involving nested quantification before, although so far we've discussed them using English rather than mathematical notation. Let's start by formalizing one of these informal uses of nested quantification:

Example 3.46: Onto functions, formalized.

Definition 2.51 said: *a function $f : A \rightarrow B$ is called onto if, for every $b \in B$, there exists at least one $a \in A$ for which $f(a) = b$* . We can rewrite this definition as:

A function $f : A \rightarrow B$ is onto if $\forall b \in B : [\exists a \in A : f(a) = b]$.

Another example that we've seen: the concept of a partition of a set (Definition 2.32) also used nested quantification, just without the \forall and \exists notation. (See Exercise 3.168.) Here are two other examples:

Example 3.47: No unmatched elements in an array.

Let's express the condition that every element of an array $A[1 \dots n]$ is a “double”—that is, appears at least twice in A . (For example, the array $[3, 2, 1, 1, 4, 4, 2, 3, 1]$ satisfies this condition.) This condition requires that, for every index i , there exists another index j such that $A[i] = A[j]$ —in other words,

$\forall i \in \{1, 2, \dots, n\} : [\exists j \in \{1, 2, \dots, n\} : i \neq j \wedge A[i] = A[j]]$.

Example 3.48: Alphabetically later.

Let's formalize the predicate “The string ____ is alphabetically after the string ____” from Example 3.29. For two letters $a, b \in \{A, B, \dots, Z\}$, write $a < b$ if a is earlier in the alphabet than b ; we'll use this ordering on *letters* to define an ordering on *strings*. Let x and y be strings over $\{A, B, \dots, Z\}$. There are two ways for x to be alphabetically later than y :

- y is a (proper) prefix of x . (See Example 3.36.) For example, FORTRAN is after FORT.

3-64 Logic

- x and y share an initial prefix of zero or more identical letters, and the first i for which $x_i \neq y_i$ has x_i later in the alphabet than y_i . For example, PASTOR comes after PASCAL because T comes after C.

Formally, then, $x \in \{A, B, \dots, Z\}^n$ is alphabetically after $y \in \{A, B, \dots, Z\}^m$ if

$$\begin{aligned} & [m < n \wedge [\forall j \in \{1, 2, \dots, m\} : x_j = y_j]] \vee && y \text{ is a proper prefix of } x \\ & [\exists i \in \{1, \dots, \min(n, m)\} : x_i > y_i \wedge [\forall j \in \{1, 2, \dots, i-1\} : x_j = y_j]]. \end{aligned}$$

the first $i-1$ characters match, and $x_i > y_i$

“Sorting alphabetically” is usually called *lexicographic ordering* in computer science—that is, the way in which words are organized in the dictionary (also known as the *lexicon*). There’s a surprisingly fascinating history of alphabetical order, with a much more interesting set of stories than you’d have ever guessed; if you’re intrigued, see the recent book by Judith Flanders [47].

Here is one more example of a statement that we’ve already seen—Goldbach’s conjecture—that implicitly involves nested quantifiers; we’ll formalize it in predicate logic. (Part of the point is to illustrate how complex even some apparently simple concepts are; there’s a good deal of complexity hidden in words like “even” and “prime,” which at this point seem pretty intuitive!)

Example 3.49: Goldbach’s Conjecture.

Recall Goldbach’s conjecture, from Example 3.1: *Every even integer greater than 2 can be written as the sum of two prime numbers.* Formalize this proposition using nested quantifiers.

Solution. Using the *sumOfTwoPrimes* predicate from Example 3.35, we can write this statement as either of the following:

$$\forall n \in \{n \in \mathbb{Z} : n > 2 \wedge 2 \mid n\} : \text{sumOfTwoPrimes}(n) \quad (\text{A})$$

$$\forall n \in \mathbb{Z} : [n > 2 \wedge 2 \mid n \Rightarrow \text{sumOfTwoPrimes}(n)] \quad (\text{B})$$

In (B), we quantify over *all* integers, but the implication $n > 2 \wedge 2 \mid n \Rightarrow \text{sumOfTwoPrimes}(n)$ is trivially true for an integer n that’s not even or not greater than 2, because false implies anything! Thus the only values of n for which the implication has any “meat” are even integers greater than 2. As such, these two formulations are equivalent. (See Exercise 3.140.) Expanding the definition of *sumOfTwoPrimes*(n) from Example 3.35, we can also rewrite (B) as

$$[\forall n \in \mathbb{Z} : n > 2 \wedge 2 \mid n \Rightarrow (\exists p \in \mathbb{Z} : \exists q \in \mathbb{Z} : [\text{isPrime}(p) \wedge \text{isPrime}(q) \wedge [n = p + q]])] \quad (\text{C})$$

We’ve also already seen that the predicate *isPrime* implicitly contains quantifiers too (“for all potential divisors d , d does not evenly divide p ”)—and, for that matter, so does the “evenly divides” predicate \mid . In Exercises 3.185–3.187, you’ll show how to rewrite Goldbach’s Conjecture in a few different ways, including using yet further layers of nested quantifiers.

3.5 Predicate Logic: Nested Quantifiers 3-65

Writing tip: Just as with nested loops in programs, the deeper the nesting of quantifiers, the harder an expression is for a reader to follow. Using well-chosen predicates (like *isPrime*, for example) in a logical statement can make it much easier to read—just like using well-chosen (and well-named) functions makes your software easier to read!

3.5.1 Order of Quantification

In expressions that involve nested quantifiers, the order of the quantifiers matters! As a frivolous example, take the title of the 1947 hit song “Everybody Loves Somebody” (sung by Dean Martin). There are two plausible interpretations of the title:

$$\forall x : \exists y : x \text{ loves } y \quad \text{and} \quad \exists y : \forall x : x \text{ loves } y.$$

The former is the more natural reading; it says that every person x has someone whom they love, but each different x can love a different person. (As in: “every child loves their mother.”) The latter says that there is one single person loved by *every* x . (As in: “Everybody loves Raymond.”) These claims are different!

Taking it further: Disambiguating the order of quantification in English sentences is one of the daunting challenges in natural language processing (NLP) systems. Compare *Every student received a diploma* and *Every student heard a commencement address*: there are, surely, many diplomas and only one address, but building a software system that understands that fact is tremendously challenging. There are many other vexing types of ambiguity in NLP systems, too. (See p. 3-17.) Human listeners are able to use pragmatic knowledge about the world to disambiguate, but doing so properly in an NLP system is very difficult.

Figure 3.28 shows a visual representation of the importance of this order of quantification. In this visualization, for example, we have the following correspondences:

$$\forall r : \exists c : P(r, c) \longleftrightarrow \text{every row has at least one column with a filled cell in it.}$$

$$\exists c : \forall r : P(r, c) \longleftrightarrow \text{there is a \textit{single} column for which every row has that column's cell filled.}$$

Now look at Figure 3.28d and Figure 3.28f. The proposition $\exists c : \forall r : P(r, c)$ is true in Figure 3.28f but *not* true in Figure 3.28d, in which every row has a filled cell but which cell is filled varies from row to row. But the proposition $\forall r : \exists c : P(r, c)$ is true in *both* Figure 3.28d and Figure 3.28f.

Here’s a mathematical example that illustrates the difference even more precisely.




Figure 3.28 An illustration of order of quantification. Let r index a *row* of the grid, and let c index a *column*. If $P(r, c)$ is true in each filled cell, then the corresponding proposition is true.

3-66 Logic

Example 3.50: The largest real number.

One of the following propositions is true; the other is false. Which is which?

$$\exists y \in \mathbb{R} : \forall x \in \mathbb{R} : x < y \quad (\text{A})$$

$$\forall x \in \mathbb{R} : \exists y \in \mathbb{R} : x < y \quad (\text{B})$$

Solution. Translating these propositions into English helps resolve this question. (A) says that there is a real number y for which the following property holds: every real number is less than y . (“There is a largest real number.”) But there isn’t a largest real number! So (A) is false. (If someone tells you that y^* satisfies $\forall x \in \mathbb{R} : x < y^*$, then you can convince him he’s wrong by choosing $x = y^* + 1$.) On the other hand, (B) says that, for every real number x , there is a real number greater than x . And that’s true: for any $x \in \mathbb{R}$, the number $x + 1$ is greater than x .

In fact, (B) is nearly the negation of (A). (Before you read through the derivation, can you figure out why we had to say “nearly” in the last sentence?)

$$\begin{aligned} \neg(\text{A}) &\equiv \neg[\exists y \in \mathbb{R} : \forall x \in \mathbb{R} : x < y] \equiv \forall y \in \mathbb{R} : \neg[\forall x \in \mathbb{R} : x < y] && \text{De Morgan's Laws (quantified form)} \\ &\equiv \forall y \in \mathbb{R} : \exists x \in \mathbb{R} : \neg(x < y) && \text{De Morgan's Laws (quantified form)} \\ &\equiv \forall y \in \mathbb{R} : \exists x \in \mathbb{R} : y \leq x && \neg(x < y) \Leftrightarrow y \leq x \\ &\equiv \forall x \in \mathbb{R} : \exists y \in \mathbb{R} : x \leq y. && \text{renaming the bound variables} \end{aligned}$$

So (B) and the negation of (A) are almost—but not quite—identical: the latter has a \leq where the former has a $<$. But both (B) and $\neg(\text{A})$ are theorems.

Commutativity of \forall and commutativity of \exists

Although the order of quantifiers does matter when universal and existential quantifiers both appear in a proposition, the order of consecutive universal quantifiers doesn’t matter. Neither does the order of consecutive existential quantifiers. That is, these quantifiers are *commutative*. Thus the following statements are theorems of predicate logic:

$$\forall x \in S : \forall y \in T : P(x, y) \Leftrightarrow \forall y \in T : \forall x \in S : P(x, y) \quad (*)$$

$$\exists x \in S : \exists y \in T : P(x, y) \Leftrightarrow \exists y \in T : \exists x \in S : P(x, y) \quad (**)$$

Why is (*) true? See Figure 3.28a. Both sides of (*) require that all the cells be filled: one says that, in every row, every column is filled; the other says that, in every column, every row is filled. But each of these statements expresses the condition that $P(x, y)$ is true for every pair $\langle x, y \rangle \in S \times T$.

The argument for (**) is similar. Each side is true if and only if $P(x, y)$ holds for at least one pair $\langle x, y \rangle \in S \times T$ (that is, when at least one cell is filled anywhere in the grid; see Figure 3.28b).

3.5 Predicate Logic: Nested Quantifiers 3-67

Because of these equivalences, as notational shorthand we'll sometimes write $\forall x, y \in S : P(x, y)$ instead of $\forall x \in S : \forall y \in S : P(x, y)$. We'll use $\exists x, y \in S : P(x, y)$ analogously.

Nested quantification and nested loops

Just as it can be helpful to think of a quantifier in terms of a corresponding loop, it can be helpful to think of nested quantifiers in terms of nested loops. And a useful way to think about the importance of the order of quantification is through the way in which changing the order of nested loops changes what they compute. (See Exercises 3.198–3.203 for some examples.)

Here's one example about how thinking about the nested-loop analogy for nested quantifiers can be helpful. Imagine writing a nested loop to examine every element of a 2-dimensional array. As long as iterations don't depend on each other, it doesn't matter whether we proceed through the array in *row-major order* (from top to bottom, going from left to right within each row) or *column-major order* (from left to right, going from top to bottom within each column): The code segments in Figure 3.29 always have the same return value, which is another way of illustrating the logical equivalence expressed by (**), above. (The graphical view is that both programs check every cell of the “grid” of possible inputs to A , as in Figure 3.28b—just in different orders.)

3.5.2 Negating Nested Quantifiers

Recall the rules for negating quantifiers from Examples 3.41 and 3.42:

$$\neg [\forall x \in S : P(x)] \Leftrightarrow [\exists x \in S : \neg P(x)] \quad \text{and} \quad \neg [\exists x \in S : P(x)] \Leftrightarrow [\forall x \in S : \neg P(x)].$$

Informally, these theorems say that “*everybody is P* is false” is equivalent to “*somebody isn't P*”; and, similarly, “*somebody is P* is false” is equivalent to “*everybody isn't P*.”

Here we will consider negating a sequence of *nested* quantifiers. Negating nested quantifiers proceeds in precisely the same way as negating a single quantifier, just acting on one quantifier at a time. (We already saw this idea in Example 3.50, where we repeatedly applied these quantified versions of De Morgan's Laws to a sequence of nested quantifiers.) For example:

1 for $j = 1$ to m :	1 for $i = 1$ to n :
2 for $i = 1$ to n :	2 for $j = 1$ to m :
3 if $A[i, j]$ then	3 if $A[i, j]$ then
4 return True	4 return True
5 return False	5 return False





Figure 3.29 Going through a matrix in row-major order or column-major order. If every cell in the grid is inspected, then the final result is identical (regardless of which order is used).

3-68 Logic

Example 3.51: No cell is filled \equiv every cell is empty.

What does it mean for the statement $\exists r : \exists c : P(r, c)$ to be *false*? From the visual perspective in Figure 3.28, the statement is true if any cell in the grid is filled. There are two different ways of thinking about what it means for $\exists r : \exists c : P(r, c)$ to be false, then:

- (1) *there does not exist a filled cell* $\neg [\exists c : \exists r : P(c, r)]$
 (2) *every cell is unfilled.* $\forall c : \forall r : [\neg P(c, r)]$



(These two sentences probably don't sound very different to you, precisely because you probably already have some sense of what it means to negate quantifiers.) To write this equivalence formally, we repeatedly apply the rule for negating an existential quantifier (Example 3.42):

$$\begin{aligned} & \neg [\exists r : \exists c : P(r, c)] && \text{"There does not exist a filled cell."} \\ \equiv & \forall r : \neg [\exists c : P(r, c)] \\ \equiv & \forall r : \forall c : \neg P(r, c). && \text{"Every cell is unfilled."} \end{aligned}$$

Similarly, let's rewrite $\neg [\exists r : \forall c : P(r, c)]$:

$$\begin{aligned} & \neg [\exists r : \forall c : P(r, c)] && \text{"It's not the case that there's a row with all columns filled."} \\ \equiv & \forall r : \neg [\forall c : P(r, c)] \\ \equiv & \forall r : \exists c : \neg P(r, c). && \text{"Every row has at least one unfilled column."} \end{aligned}$$

Here's an example of negating a longer sequence of nested quantifiers:

Example 3.52: Triple negations.

For a ternary predicate P , we have:

$$\begin{aligned} \neg \exists x : \forall y : \exists z : P(x, y, z) &\equiv \forall x : \neg \forall y : \exists z : P(x, y, z) \\ &\equiv \forall x : \exists y : \neg \exists z : P(x, y, z) \\ &\equiv \forall x : \exists y : \forall z : \neg P(x, y, z). \end{aligned}$$

Here's a last example, which requires translation from English into logical notation:

Example 3.53: Negating nested quantifiers.

Negate the following sentence:

For every iPhone user, there's an iPhone app that every one of that user's iPhone-using friends has downloaded.

Solution. First, let's reason about how the statement would be false: there would be some iPhone user—we'll call him Steve—such that, for every iPhone app, Steve has a friend who didn't download that app.

Write U and A for the sets of iPhone users and apps, respectively. In (pseudo)logical notation, the original claim looks like

$$\forall u \in U : \exists a \in A : \forall v \in U : [(u \text{ and } v \text{ are friends}) \Rightarrow (v \text{ downloaded } a)].$$

To negate this statement, we repeatedly apply the quantified De Morgan's laws, once per quantifier:

$$\begin{aligned} & \neg \forall u \in U : \exists a \in A : \forall v \in U : [(u \text{ and } v \text{ are friends}) \Rightarrow (v \text{ downloaded } a)] \\ \equiv & \exists u \in U : \neg \exists a \in A : \forall v \in U : [(u \text{ and } v \text{ are friends}) \Rightarrow (v \text{ downloaded } a)] \\ \equiv & \exists u \in U : \forall a \in A : \neg \forall v \in U : [(u \text{ and } v \text{ are friends}) \Rightarrow (v \text{ downloaded } a)] \\ \equiv & \exists u \in U : \forall a \in A : \exists v \in U : \neg [(u \text{ and } v \text{ are friends}) \Rightarrow (v \text{ downloaded } a)]. \end{aligned}$$

Using $\neg(p \Rightarrow q) \equiv p \wedge \neg q$ (Exercise 3.85), we can further write this expression as:

$$\equiv \exists u \in U : \forall a \in A : \exists v \in U : [(u \text{ and } v \text{ are friends}) \wedge \neg(v \text{ downloaded } a)].$$

This last proposition, translated into English, matches the informal description above as to why the original claim would be false: *there's some person such that, for every app, that person has a friend who hasn't downloaded that app.*

3.5.3 Two New Ways of Considering Nested Quantifiers

We'll close this section with two different but useful ways to think about nested quantification. As a running example, consider the following (true) proposition

$$\forall x \in \mathbb{Z} : \exists y \in \mathbb{Z} : x = y + 1, \quad (\dagger)$$

which says that the number that's one less than every integer is an integer too. We'll discuss two ways of thinking about propositions like (\dagger) with nested quantifiers: as a “game with a demon” in which you battle against an all-knowing demon to try to make the innermost quantifier's body true; and as a single quantifier whose body is a predicate, but a predicate that just happens to be expressed using quantifiers.

Nested quantifiers as demon games

One way to think about any proposition involving nested quantifiers is as a “game” played between you and a demon. (Thanks to Dexter Kozen for teaching me this way of thinking of nested quantifiers [73].) Here are the rules of the game:

3-70 Logic

- Your goal is to make the innermost statement— $x = y + 1$ for our running example (\dagger)—turn out to be true; the demon’s goal is to make that statement false.
- Every “for all” quantifier is a choice that the demon makes; every “there exists” quantifier is a choice that you get to make:
 - in the expression $\forall a \in S : \dots$, the demon chooses a particular value of $a \in S$, and the game continues in the “ \dots ” part of the expression.
 - in the expression $\exists b \in S : \dots$, you choose a particular value of $b \in S$, and, again, the game continues in the “ \dots ” part.
- Your choices and the demon’s choices are made in the left-to-right order of the quantifiers (starting with the outside-most quantifier, and moving inward).
- You win the game—in other words, the proposition in question is true—if, no matter how cleverly the demon plays, you make the innermost statement true.

Here are two examples of viewing quantified statements as demon games, one for a true statement and one for a false statement:

Example 3.54: Showing that (\dagger) is true.

We’ll use a “game with the demon” to argue that $\forall x \in \mathbb{Z} : \exists y \in \mathbb{Z} : x = y + 1$ is true.

- 1 The outermost quantifier is \forall , so the demon picks a value for $x \in \mathbb{Z}$.
- 2 Now you get to pick a value $y \in \mathbb{Z}$. A good choice for you is $y = x - 1$.
- 3 Because you chose $y = x - 1$, indeed $x = y + 1$. You win!

(For example, if the demon picks 16, you pick 15. If the demon picks -19 , you pick -20 . And so forth.)

No matter what the demon picks, your strategy will make you win—and therefore (\dagger) is true!

By contrast, consider (\dagger) with the order of quantification reversed:

Example 3.55: A losing demon game.

Consider playing a demon game for the proposition

$$\exists y \in \mathbb{Z} : \forall x \in \mathbb{Z} : x = y + 1.$$

Unfortunately, the \exists is first, which means that you have to make the first move. But when you pick a number y , the demon *then* gets to pick an x —and there are an infinitude of x values that the demon can choose so that $x \neq y + 1$. (You pick 42? The demon picks 666. You pick 17? The demon picks 666. You pick 665? The demon picks 616.) Therefore you can’t guarantee that you win the game, so we haven’t established this claim.

3.5 Predicate Logic: Nested Quantifiers 3-71

By the way, you *could* win a demon game to prove the negation of the claim in Example 3.55:

$$\neg(\text{the claim from Example 3.55}) \equiv \forall y \in \mathbb{Z} : \exists x \in \mathbb{Z} : x \neq y + 1.$$

First, the demon picks some unknown $y \in \mathbb{Z}$. Then you have to pick an integer x such that $x \neq y + 1$ —but that’s easy: for any y the demon picks, you pick $x = y$. You win!

Nested quantifiers as single quantifiers

Here’s a reminder of our running example, with a portion of it highlighted:

$$\forall x \in \mathbb{Z} : \exists y \in \mathbb{Z} : x = y + 1. \quad (\dagger)$$

What kind of thing is the highlighted piece? It can’t be a proposition, because x is a free variable in it. But once we plug in a value for x , the expression becomes true or false. In other words, the expression $\exists y \in \mathbb{Z} : x - 1 = y$ is itself a (unary) predicate: once we are given a value of x , we can compute the truth value of the expression. Similarly, the expression $x - 1 = y$ is also a predicate—but a binary predicate, taking both x and y as arguments. Let’s name these predicates: let’s write $P(x, y)$ for $x - 1 = y$, and $\text{hasIntPredecessor}(x)$ for $\exists y \in \mathbb{Z} : x - 1 = y$.

Using this perspective, we can write (\dagger) as follows:

$$\begin{aligned} \forall x \in \mathbb{Z} : \exists y \in \mathbb{Z} : x = y + 1 &\equiv \forall x \in \mathbb{Z} : \exists y \in \mathbb{Z} : P(x, y) \\ &\equiv \forall x \in \mathbb{Z} : \text{hasIntPredecessor}(x). \end{aligned} \quad (\ddagger)$$

One implication of this view is that negating nested quantifiers is really just the same as negating non-nested quantifiers. For example:

Example 3.56: Negating nested quantifiers.

We can view the negation of (\dagger) , as written in (\ddagger) , as follows:

$$\begin{aligned} \neg(\dagger) &\equiv \neg \forall x \in \mathbb{Z} : \text{hasIntPredecessor}(x) \\ &\equiv \exists x \in \mathbb{Z} : \neg \text{hasIntPredecessor}(x). \end{aligned}$$

And, re-expanding the definition of hasIntPredecessor and again applying the quantified De Morgan’s Law, we have that

$$\begin{aligned} \neg \text{hasIntPredecessor}(x) &\equiv \neg \exists y \in \mathbb{Z} : P(x, y) \\ &\equiv \forall y \in \mathbb{Z} : \neg P(x, y) \\ &\equiv \forall y \in \mathbb{Z} : x - 1 \neq y. \end{aligned}$$

3-72 Logic

Together, these two negations show

$$\begin{aligned}\neg(\dagger) &\equiv \exists x \in \mathbb{Z} : \neg \textit{hasIntPredecessor}(x) \\ &\equiv \exists x \in \mathbb{Z} : \forall y \in \mathbb{Z} : \neg P(x, y) \\ &\equiv \exists x \in \mathbb{Z} : \forall y \in \mathbb{Z} : x - 1 \neq y.\end{aligned}$$

Taking it further: This view of nested quantifiers as a single quantifier whose body just happens to express its condition using quantifiers has a close analogy with writing a particular kind of function in a programming language. If we look at a two-argument function in the right light, we can see it as a function that takes one argument *and returns a function that takes one argument*. This approach is called *Currying*; see p. 3-73.

3.5 Predicate Logic: Nested Quantifiers 3-73

COMPUTER SCIENCE CONNECTIONS

CURRYING

For a binary predicate $P(x, y)$ in an expression like $\forall y : \forall x : P(x, y)$, we can think of this expression as first plugging in a value for y , which then yields a unary predicate $\forall x : P(x, y)$ (which takes an argument x). There's an interesting parallel between this view of nested quantifiers and a way of writing functions in some programming languages.

For concreteness, let's think about a small function that takes two arguments and just returns their sum. Figure 3.30 shows implementations of this function in three different programming languages, and then uses `sum` to actually compute $3 + 2$ and $99 + 12$.

But now suppose that we want to define a function that takes one argument and adds 3 to it. Can we make use of the `sum` function to do so? (The analogy to predicates is that taking a two-argument predicate and applying it to one argument gives one-argument predicate; here we're trying to take a two-argument function in a programming language and apply it to one argument to yield a one-argument function.) The answer is yes—and it turns out that creating the “add 3” function using `sum` is very easy in ML: we simply apply `sum` to one argument, and the result is a function that “still wants” one more argument. To do this, we execute `val add3 = sum 3`; which defines a *value* (that's the `val` part of the syntax) whose name is `add3` and whose value is `sum` applied to 3. That makes `add3` a one-argument function, and now `add3 0` has the value 3; `add3 108` has the value 111; and `add3 199` has the value 202.

A function like `sum` in ML, which takes its multiple arguments “one at a time,” is said to be *Curried*—after the logician Haskell Curry (1900–1982). (As usual, there's a richer history than a concept named after one person would suggest: the idea dates back at least to Moses Schönfinkel (1888–1942) and Gottlob Frege (1848–1925), two earlier logicians.) The programming language Haskell is also named in Curry's honor. Thinking about Curried functions is a classical topic in the study of programming languages. (Currying can be very handy in programming—particularly when you have one parameter to a function that stays the same for many different calls

```
1 (* ML *)
2 fun sum a b = a + b;
3 sum 3 2;      (* returns 5 *)
4 sum 99 12;    (* returns 111 *)
```

A quick note on ML syntax: `fun` is a keyword that says we're defining a function; `sum` is the name of it; `a b` is the list of arguments; and that function is defined to return the value of `a + b`.

```
1 # Python
2 def sum(a,b):
3     return a + b
4 sum(3,2)      # returns 5
5 sum(99,12)    # returns 111
```

For Scheme: `(lambda args body)` denotes the function that takes arguments `args` and returns the value of the function body `body`. Also, Scheme is a *prefix language*, so applying the function `f` to arguments `arg1, arg2, ..., argN` is written `(f arg1 arg2 ... argN)`; for example, `(+ 1 2)` has the value 3.

```
1 ; Scheme
2 (define sum
3   (lambda (a b) (+ a b)))
4 (sum 3 2)      ;; returns 5
5 (sum 99 12)    ;; returns 111
```

Figure 3.30 A sum function, implemented in three languages.

```
1 fun sum a b = a + b;
2 val add3 = sum 3;
3
4 sum 3 2;      (* returns 5 *)
5 add3 2;       (* returns 5 *)
```

```
1 def sum(a):
2     def sumA(b):
3         return a + b
4     return sumA
5
6 add3 = sum(3)
7
8 sum(3)(2)      # returns 5
9 add3(2)        # returns 5
```

Define `sum` to be a function that takes an argument `a` and returns ...

```
1 (define sum (lambda (a)
2   (lambda (b) (+ a b))))
3
4 (define add3 (sum 3))
5
6 ((sum 3) 2)    ;; returns 5
7 (add3 2)       ;; returns 5
```

... the function that takes an argument `b` and returns `a + b`. (This function [taking `b`, returning `a+b`] is called `sumA` in the Python code, and is unnamed in the Scheme code.)

Figure 3.31 Curried implementations of `sum`.

3-74 Logic

to a multiparameter function. For example, you might Curry a function to look up data in a large dataset: `query(dataset)(item)`. (See [1], or any good textbook on programming languages.)

While writing Curried functions is almost automatic in ML, you can also write Curried functions in other programming languages, too. Examples of a Curried version of `sum` in Python and Scheme are in Figure 3.31; it's even possible to write Curried functions in C or Java, though it's much less natural than in ML, Python, and Scheme.

EXERCISES

- 3.148** Let F denote the set of all functions $f : \mathbb{R} \rightarrow \mathbb{R}$ taking real numbers as input and producing real numbers as output. (For one example, $\text{plusOne}(x) = x + 1$ is a function $\text{plusOne} : \mathbb{R} \rightarrow \mathbb{R}$, so $\text{plusOne} \in F$.) Is the proposition $\forall c \in \mathbb{R} : [\exists f \in F : f(0) = c]$ true or false? Justify your answer.
- 3.149** What about $\exists f \in F : [\forall c \in \mathbb{R} : f(0) = c]$? (F is defined in Exercise 3.148.) Again, justify your answer.
- 3.150** What about $\forall c \in \mathbb{R} : [\exists f \in F : f(c) = 0]$? (F is defined in Exercise 3.148.) Justify.
- 3.151** What about $\exists f \in F : [\forall c \in \mathbb{R} : f(c) = 0]$? (F is defined in Exercise 3.148.) Again, justify.

Under many operating systems, users can schedule a task to be run at a specified time in the future. In Unix-like operating systems, this type of scheduled job is called a cron job. (Greek: chron- “time.”) For example, a backup might run nightly at 2:00am, and a scratch drive might be emptied out weekly on Friday night at 11:50pm. Let $T = \{1, 2, \dots, t_{\max}\}$ be a set of times (measured in minutes, let’s say), and let J be a set of jobs. Let scheduledAt be a predicate where $\text{scheduledAt}(j, t)$ means “job j is scheduled at time t .” (Assume that jobs do not last more than one minute. The same job can be scheduled at multiple times.) Formalize the following conditions using only standard quantifiers, arithmetic operators, logical connectives, and the scheduledAt predicate.

- 3.152** There is never more than one job scheduled at the same time.
- 3.153** Every job is scheduled at least once.
- 3.154** Job A is never run twice within two minutes.
- 3.155** Job B is run at least three times.
- 3.156** Job C is run at most twice.
- 3.157** Job D is run sometime after the last time that Job E is run.
- 3.158** Job F is run at least once between consecutive executions of Job G .
- 3.159** Job H is run at most once between consecutive executions of Job I .

Let $P[1 \dots n, 1 \dots m]$ be a 2-dimensional array of the pixels of a black-and-white image; for every x and y , the value of $P[x, y] = 0$ if the $\langle x, y \rangle$ th pixel is black, and $P[x, y] = 1$ if it’s white. Translate these statements into predicate logic:

- 3.160** Every pixel in the image is black.
- 3.161** There is at least one white pixel.
- 3.162** Every row has at least one white pixel.
- 3.163** There are never two consecutive white pixels in the same column.

A standard American crossword puzzle is a 15-by-15 grid, which can be represented as a two-dimensional 15-by-15 array G , where $G[i, j] = \text{True}$ if and only if the cell in the i th row and j th column is “open” (otherwise known as “unfilled” or “not a black square”). If i and j are out of range—that is, for any $i \leq 0$ or $i > 15$ or $j \leq 0$ or $j > 15$ —assume $G[i, j] = \text{False}$. A contiguous horizontal or vertical sequence of two or more open squares surrounded by black squares is called a word. (The assumption that $G[i, j]$ is False when i or j is out of range is equivalent to us pretending that our real grid is surrounded by black squares. In CS, this style of structure is called a sentinel, wherein we introduce boundary values to avoid having to write out verbose special cases.)

- 3.164** There are certain customs that G must obey to be a standard American puzzle (all of which are met by the example grid in Figure 3.32). Rewrite the following informally stated condition as a fully formal definition. *No unchecked letters*: every open cell appears in both a down word and an across word.
- 3.165** Repeat for the *no two-letter words* condition: every word has length at least 3.
- 3.166** Repeat for *rotational symmetry*: if the entire grid is rotated by 180° , then the rotated grid is identical to the original grid.
- 3.167** Repeat for *overall interlock*: for any two open squares, there is a path of open squares that connects the first to the second. (That is, we can get from *here* to *there* through words.) Your answer should include a formal definition of what it means for there to be a path from one cell to another (“there exists a sequence of squares such that ...”).

3-76 Logic




Figure 3.32 A standard American crossword puzzle.

- 3.168** According to Definition 2.32, a *partition* of a set S is a set $\{A_1, A_2, \dots, A_k\}$ of sets such that (i) A_1, A_2, \dots, A_k are all nonempty; (ii) $A_1 \cup A_2 \cup \dots \cup A_k = S$; and (iii) for any i and $j \neq i$, the sets A_i and A_j are disjoint. Formalize this definition using nested quantifiers and basic set notation.

- 3.169** Consider the “maximum” problem: given an array of numbers, return the maximum element of that array:

Input: An array $A[1 \dots n]$, where each $A[i] \in \mathbb{Z}$.

Output: An integer $x \in \mathbb{Z}$ such that . . .

Complete the formal specification for this problem by finishing the specification for the output.



Let $T = \{1, \dots, 12\} \times \{0, 1, \dots, 59\}$ be the set of numbers that a digital clock in twelve-hour mode can display. We can think of a clock as a function $c : T \rightarrow T$: if the real time is $t \in T$, the clock displays the time $c(t)$. (For example, if the clock *fastby7* runs seven minutes fast, then $\text{fastby7}(12:00) = 12:07$.) Formalize the following predicates using only standard quantifiers, equality symbols, and the function $\text{add} : T \times \mathbb{Z}^{\geq 0} \rightarrow T$, where $\text{add}(t, x)$ is the time that is x minutes later than t . (See Exercise 2.243.)

- 3.170** A clock is *right* if it always displays the correct time. Formalize *right*.
- 3.171** A clock *keeps time* if there’s some fixed offset by which it is always off from being right. (For example, *fastby7* above correctly keeps time.) Formalize *keepsTime*.
- 3.172** A clock is *close enough* if it always displays a time that’s within two minutes of the correct time. Formalize *closeEnough*.
- 3.173** A clock is *broken* if there’s some fixed time that it always displays, regardless of the real time. Formalize *broken*.
- 3.174** “Even a broken clock is right twice a day,” they say. (They mean: “even a broken clock displays the correct time at least once per T .”) Using the *broken* predicate from Exercise 3.173, formalize the adage and prove it true.

A classic topic of study for computational biologists is genomic distance measures: given two genomes, we’d like to report a single number that represents how different those two genomes are. These distance computations are useful in, for example, reconstructing the evolutionary tree of a collection of species. Consider two genomes A and B of bacterium. Label the n genes that appear in A ’s chromosome, in order, as $\pi_A = \langle 1, 2, \dots, n \rangle$. The same genes appear in B —say, in the order $\pi_B = \langle r_1, r_2, \dots, r_n \rangle$. A particular model of genomic distance will define a specific way in which this list of numbers can mutate; the question is to find a minimum-length sequence of mutations to explain the difference between the orders π_A and π_B . One type of biologically motivated mutation is the prefix reversal, shown in Figure 3.33a. This model defines what’s called the pancake-flipping problem (coincidentally, the subject of the lone academic paper with Bill Gates as an author [50]). See Figure 3.33b.

- 3.175** You are given a sequence of pancake radii $\langle r_1, r_2, \dots, r_n \rangle$, listed from top to bottom, where $\{r_1, r_2, \dots, r_n\} = \{1, 2, \dots, n\}$ (but not necessarily in order). Give a fully quantified logical expression expressing the condition that the given pancakes are sorted.
- 3.176** Again, you are given a sequence of pancake radii $\langle r_1, r_2, \dots, r_n \rangle$, listed from top to bottom. Give a fully quantified logical expression expressing the condition that the given pancakes can be sorted with exactly one flip.
- 3.177** Again, you are given a sequence of pancake radii $\langle r_1, r_2, \dots, r_n \rangle$, listed from top to bottom. Give a fully quantified logical expression expressing the condition that the given pancakes can be sorted with exactly two flips. (Hint: consider writing a program to verify that your indices aren’t off by one.)

Exercises 3-77



(b) The pancake-flipping problem: you are given a stack of pancakes, which you want to arrange in order of size by repeatedly flipping some number of pancakes at the top of the pile, using as few flips as possible.

Figure 3.33 Genome rearrangements: prefix reversals and the pancake-flipping problem.

3.178 Let P be a set of people, and let T be a set of times. Let $\text{friends}(x, y)$ be a predicate denoting that $x \in P$ and $y \in P$ are friends. Let $\text{bought}(x, t)$ be a predicate denoting that $x \in P$ bought an iPad at time $t \in T$. Formalize this statement in predicate logic: “Everyone who bought an iPad has a friend who bought one before they did.”

3.179 Is the claim from Exercise 3.178 true (in the real world)? Justify your answer.

*In programming, an assertion is a logical statement that announces (“asserts”) a condition ϕ that the programmer believes to be true. For example, a programmer who is about to access the 202nd element of an array A might assert that $\text{length}(A) \geq 202$ before accessing this element. When an executing program reaches an **assert** statement, the program aborts if the condition in the statement isn’t true. (Using assertions can be an extremely valuable way of documenting and debugging programs, particularly because liberally including assertions will allow the revelation of unexpected data values much earlier in the execution of a program. And these languages have a global toggle that allows the testing of assertions to be turned off, so once the programmer is satisfied that the program is working properly, she doesn’t have to worry about any running-time overhead for these checks.)*

3.180 Give a nonempty input array $A[1 \dots n]$ that would cause the assertion in Figure 3.34a to fail. (That is, identify an array A that would cause for the asserted condition to be false.)

3.181 Give a nonempty input array $A[1 \dots n]$ that would cause the assertion in Figure 3.34b to fail.

3.182 Give a nonempty input array $A[1 \dots n]$ that would cause the assertion in Figure 3.34c to fail.

While the quantifiers \forall and \exists are by far the most common, there are some other quantifiers that are sometimes used. For each of the following quantifiers, write an expression that is logically equivalent to the given statement that uses only the quantifiers \forall and \exists ; standard propositional logic notation ($\wedge, \neg, \vee, \Rightarrow$); standard equality/inequality notation ($=, \geq, \leq, <, >$); and the predicate P in the question.

(a)

```

1 last := 0
2 for index := 1, ..., n - 1:
3   if A[index] > A[index + 1] then
4     last := index
5 assert last ≥ 1 and last ≤ n - 1
6 swap A[last] and A[last + 1]
```

(b)

```

1 total := A[1]
2 i := 1
3 for i := 2, ..., n - 1:
4   if A[i + 1] > A[i] then
5     total := total + A[i]
6   assert total > A[1]
7 return total
```

(c)

```

1 for start := 1, ..., n - 1:
2   min := start
3   for i := start + 1, ..., n:
4     assert start = 1 or A[i] > A[start - 1]
5     if A[min] > A[i] then
6       min := i
7   swap A[start] and A[min]
```

Figure 3.34 Three pieces of pseudocode containing assertions.

3-78 Logic

3.183 Write an equivalent expression to $\exists! x \in \mathbb{Z} : P(x)$ (“there exists a unique $x \in \mathbb{Z}$ such that $P(x)$ ”), which is true when there is one and only one value of x in the set \mathbb{Z} such that $P(x)$ is true.

3.184 Write an equivalent expression to $\exists_{\infty} x \in \mathbb{Z} : P(x)$ (“there exist infinitely many $x \in \mathbb{Z}$ such that $P(x)$ ”), which is true when there are infinitely many different values of $x \in \mathbb{Z}$ such that $P(x)$ is true.

3.185 Prove that the following two formulations of Goldbach’s conjecture are logically equivalent:

$$\forall n \in \mathbb{Z} : [n > 2 \wedge 2 \mid n \Rightarrow (\exists p \in \mathbb{Z} : \exists q \in \mathbb{Z} : [isPrime(p) \wedge isPrime(q) \wedge n = p + q])] \quad (1)$$

$$\forall n \in \mathbb{Z} : \exists p \in \mathbb{Z} : \exists q \in \mathbb{Z} : [n \leq 2 \vee 2 \nmid n \vee [isPrime(p) \wedge isPrime(q) \wedge n = p + q]] \quad (2)$$

3.186 Rewrite Goldbach’s conjecture without using *isPrime*—that is, using only quantifiers, the \mid predicate, and standard arithmetic ($+$, \cdot , \geq , etc.).

3.187 Even the \mid predicate implicitly involves a quantifier: $p \mid q$ is equivalent to $\exists k \in \mathbb{Z} : p \cdot k = q$. Rewrite Goldbach’s conjecture without using \mid either—that is, use only quantifiers and standard arithmetic symbols ($+$, \cdot , \geq , etc.).

3.188 (*programming required.*) As we discussed, the truth value of Goldbach’s conjecture is currently unknown. In a project that concluded in 2012, the conjecture has been verified for all even integers from 4 up to 4×10^{18} , through a massive distributed computational effort led by Tomás Oliveira e Silva. Write a program to test Goldbach’s conjecture, in a programming language of your choice, for a much smaller range: all even integers up to 10,000.

Let S be an arbitrary nonempty set and let P be an arbitrary binary predicate. Decide whether the following statements are always true (for any P and S), or whether they can be false. Prove your answers.

3.189 $[\exists y \in S : \forall x \in S : P(x, y)] \Rightarrow [\forall x \in S : \exists y \in S : P(x, y)]$

3.190 $[\forall x \in S : \exists y \in S : P(x, y)] \Rightarrow [\exists y \in S : \forall x \in S : P(x, y)]$

Consider any unary predicate $P(x)$ over a nonempty set S . It turns out that both of the following propositions are theorems of propositional logic. Prove them both.

3.191 $\forall x \in S : [P(x) \Rightarrow (\exists y \in S : P(y))]$

3.192 $\exists x \in S : [P(x) \Rightarrow (\forall y \in S : P(y))]$

Most real-world English utterances are ambiguous—that is, there are multiple possible interpretations of the given sentence. A particularly common type of ambiguity involves order of quantification. For each of the following English sentences, find as many different logical readings based on order of quantification as you can. Write down those interpretations using pseudological notation, and also write a sentence that expresses each meaning unambiguously.

3.193 A computer crashes every day.

3.194 Every prime number except 2 is divisible by an odd integer greater than 1.

3.195 Every student takes a class every term.

3.196 Every submitted program failed on a case submitted by a student.

3.197 You should have found two different logical interpretations in Exercise 3.194. One of these interpretations is a theorem, and one of them is not. Decide which is which, and prove your answers.

3.198 The code in Figure 3.35a uses nested loops to compute some fact about a predicate P . Write a fully quantified statement of predicate logic whose truth value matches the value returned by the code. (Assume that S is a finite universe.)

3.199 Do the same for the code in Figure 3.35b.

3.200 Do the same for the code in Figure 3.35c.

3.201 Do the same for the code in Figure 3.35d.

3.202 Do the same for the code in Figure 3.35e.

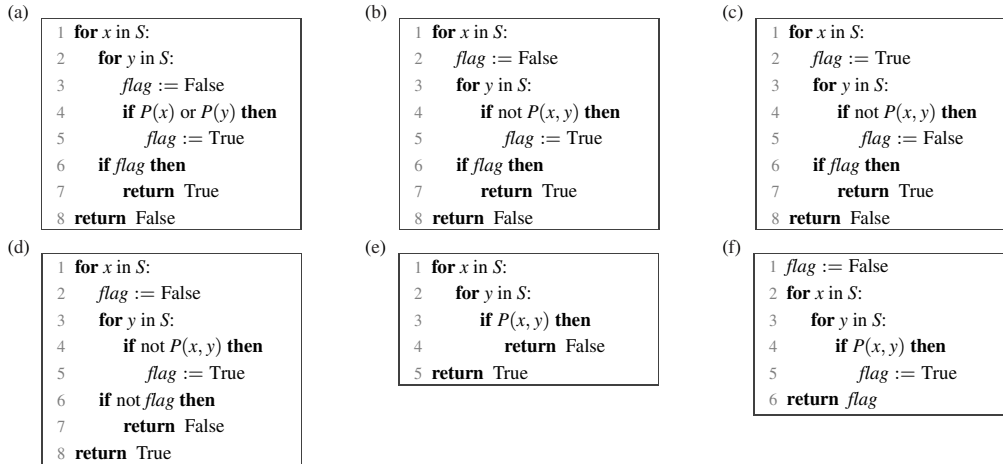


Figure 3.35 Some nested loops that compute some facts about a predicate P . (Assume that the universe S is finite.)

3.203 Do the same for the code in Figure 3.35f.

3.204 As we've discussed, there is no algorithm that can decide whether a given fully quantified proposition φ is a theorem of predicate logic. But there are several specific types of fully quantified propositions for which we *can* decide whether a given statement is a theorem. Here you'll show that, when quantification is only over a *finite* set, it is possible to give an algorithm to determine whether φ is a theorem. Suppose that you are given a fully quantified proposition φ , where the domain for every quantifier is a finite set—say $S = \{0, 1\}$. Describe an algorithm that is guaranteed to figure out whether φ is a theorem.

3.6 Chapter at a Glance

Propositional Logic

A *proposition* is the kind of thing that is either true or false. An *atomic proposition* (or *Boolean variable*) is a conceptually indivisible proposition. A *compound proposition* (or *Boolean formula*) is one built up using a *logical connective* and one or more simpler propositions. The most common logical connectives are shown in Figure 3.36. A proposition that contains the atomic propositions p_1, \dots, p_k is sometimes called a *proposition over p_1, \dots, p_k* , or a *Boolean formula* or a *Boolean expression over p_1, \dots, p_k* .

The *truth value* of a proposition is its truth or falsity. (The truth value of a Boolean formula over p_1, \dots, p_k is determined only by the truth values of each of p_1, \dots, p_k .) Each logical connective is defined by how the truth value of the compound proposition formed using that connective relates to the truth values of the constituent propositions. A *truth table* defines a connective by listing, for each possible assignment of truth values for the constituent propositions, the truth value of the entire compound proposition. See Figure 3.37.

Consider a Boolean formula over variables p_1, \dots, p_k . A *truth assignment* is a setting to true or false for each variable. (So a truth assignment corresponds to a row of the truth table for the proposition.) A truth assignment *satisfies* the proposition if, when the values from the truth assignment are plugged in, the proposition is true. A Boolean formula is a *tautology* if *every* truth assignment satisfies it; it's *satisfiable* if *some* truth assignment satisfies it; and it's *unsatisfiable* or a *contradiction* if no truth assignment does. Two Boolean propositions are *logically equivalent* if they're satisfied by exactly the same truth assignments (that is, they have identical truth tables).

Consider an implication $p \Rightarrow q$. Note from Figure 3.37 that the proposition $p \Rightarrow q$ is true if, whenever p is true, q is too. So the only situation in which $p \Rightarrow q$ is false is when p is true and q is false. False implies anything! Anything implies true! The *antecedent* or *hypothesis* of the implication $p \Rightarrow q$ is p ; the

negation	$\neg p$	"not p "
disjunction	$p \vee q$	" p or q "
(inclusive or: " p , q , or both")		
conjunction	$p \wedge q$	" p and q "
implication	$p \Rightarrow q$	"if p , then q " or " p implies q "
equivalence	$p \Leftrightarrow q$	" p if and only if q "
exclusive or	$p \oplus q$	" p xor q "
("p or q, but not both")		

Figure 3.36 The six basic logical connectives of propositional logic.

p	$\neg p$	p	q	$p \wedge q$	$p \vee q$	$p \Rightarrow q$	$p \Leftrightarrow q$	$p \oplus q$
T	F	T	T	T	T	T	T	F
T	F	T	F	F	T	F	F	T
F	T	F	T	F	T	T	F	T
F	T	F	F	F	F	T	T	F

Figure 3.37 Truth tables for the logical connectives in Figure 3.36.

consequent or *conclusion* of the implication is q . The *converse* of the implication $p \Rightarrow q$ is the implication $q \Rightarrow p$. The *contrapositive* is the implication $\neg q \Rightarrow \neg p$. Any implication is logically equivalent to its contrapositive. But an implication is *not* logically equivalent to its converse.

A *literal* is a Boolean variable or the negation of a Boolean variable. A proposition is in *conjunctive normal form (CNF)* if it is the conjunction (and) of a collection of clauses, where a clause is a disjunction (or) of a collection of literals. A proposition is in *disjunctive normal form (DNF)* if it is the disjunction of a collection of clauses, where a clause is a conjunction of a collection of literals. Every proposition is logically equivalent to a proposition that is in CNF, and to another that is in DNF.

Predicate Logic

A *predicate* is a statement containing some number of “blanks” (variables), and that has a truth value once values are plugged in for those variables. (Alternatively, a *predicate* is a Boolean-valued function.) Once particular values for these variables are plugged in, the resulting expression is a proposition. A proposition can also be formed from a predicate through *quantifiers*:

- The *universal quantifier* \forall (“for all”): the proposition $\forall x \in U : P(x)$ is true if, for every $x \in U$, we have that $P(x)$ is true.
- The *existential quantifier* \exists (“there exists”): the proposition $\exists x \in U : P(x)$ is true if, for at least one $x \in U$, we have that $P(x)$ is true.

The set U is called the *universe* or *domain of discourse*. When the universe is clear from context, it may be omitted from the notation.

In the expression $[\forall x : \underline{\quad}]$ or $[\exists x : \underline{\quad}]$, the *scope* or *body* of the quantifier is the underlined blank, and the variable x is *bound* by the quantifier. A *free* or *unbound* variable is one that is not bound by any quantifier. A *fully quantified* expression is one with no free variables.

A *theorem* of predicate logic is a fully quantified expression that is true for all possible meanings of the predicates in it. Two expressions are *logically equivalent* if they are true under precisely the same set of meanings for their predicates. (Alternatively, two expressions φ and ψ are *logically equivalent* if $\varphi \Leftrightarrow \psi$ is a theorem.) Two useful theorems of predicate logic are De Morgan’s laws:

$$\neg \forall x \in S : P(x) \Leftrightarrow \exists x \in S : \neg P(x) \quad \text{and} \quad \neg \exists x \in S : P(x) \Leftrightarrow \forall x \in S : \neg P(x).$$

There is no general algorithm that can test whether any given expression is a theorem. If we wish to prove that an implication $\varphi \Rightarrow \psi$ is an theorem, we can do so with a *proof by assuming the antecedent*: to prove that the implication $\varphi \Rightarrow \psi$ is always true, we will rule out the one scenario in which it wouldn’t be; specifically, we *assume* that φ is true, and then *prove* that ψ must be true too, under this assumption.

3-82 Logic

A *vacuously quantified* statement is one in which the domain of discourse is the empty set. The vacuous universal quantification $\forall x \in \emptyset : P(x)$ is a theorem; the vacuous existential quantification $\exists x \in \emptyset : P(x)$ is always false.

Quantifiers are *nested* if one quantifier is inside the scope of another quantifier. Nested quantifiers work in precisely the same way as single quantifiers, applied in sequence. A proposition involving nested quantifier like $\forall x \in S : \exists y \in T : R(x, y)$ is true if, for every choice of x , there is some choice of y (which can depend on the choice of x) for which $R(x, y)$ is true. Order of quantification matters in general; the expressions $\forall x : \exists y : R(x, y)$ and $\exists y : \forall x : R(x, y)$ are *not* logically equivalent.

Key Terms and Results

Key Terms

Propositional Logic

- proposition
- truth value
- atomic and compound propositions
- logical connectives:
 - negation (\neg)
 - conjunction (\wedge)
 - disjunction (\vee)
 - implication (\Rightarrow)
 - exclusive or (\oplus)
 - if and only if (\Leftrightarrow)
- truth assignments and truth tables
- tautology
- satisfiability/unsatisfiability
- logical equivalence
- antecedent and consequent
- converse, contrapositive, and inverse
- conjunctive normal form (CNF)
- disjunctive normal form (DNF)

Predicate Logic

- predicate
- quantifiers:
 - universal quantifier (\forall)
 - existential quantifier (\exists)
- free and bound variables
- fully quantified expression
- theorems of predicate logic
- logical equivalence in predicate logic
- proof by assuming the antecedent
- vacuous quantification
- nested quantifiers

Key Results

Propositional Logic

- 1 We can build a truth table for any proposition by repeatedly applying the definitions of each of the logical connectives, as shown in Figure 3.3.
- 2 Two propositions φ and ψ are logically equivalent if and only if $\varphi \Leftrightarrow \psi$ is a tautology.
- 3 An implication $p \Rightarrow q$ is logically equivalent to its contrapositive $\neg q \Rightarrow \neg p$, but not to its converse $q \Rightarrow p$.
- 4 There are many important propositional tautologies and logical equivalences, some of which are shown in Figures 3.9 and 3.12.
- 5 We can show that propositions are logically equivalent by showing that every row of their truth tables are the same.
- 6 Every proposition is logically equivalent to one that is in disjunctive normal form (DNF) and to one that is in conjunctive normal form (CNF).

Predicate Logic

- 1 We can build a proposition from a predicate $P(x)$ by plugging in a particular value for x , or by quantifying over x as in $\forall x \in S : P(x)$ or $\exists x \in S : P(x)$.
- 2 Unlike with propositional logic, there is no algorithm that is guaranteed to determine whether a given fully quantified predicate-logic expression is a theorem.
- 3 There are many important predicate-logic theorems, some of which are shown in Figure 3.21.
- 4 The statements $\neg \forall x : P(x)$ and $\exists x : \neg P(x)$ are logically equivalent. So are $\neg \exists x : P(x)$ and $\forall x : \neg P(x)$.
- 5 We can think of nested quantifiers as a sequence of single quantifiers, or as “games with a demon.”

