Chapter 8

Mathematical Logic

Perhaps the most distinguishing characteristic of mathematics is its reliance on logic. Explicit training in mathematical logic is essential to a mature understanding of mathematics. Familiarity with the concepts of logic is also a prerequisite to studying a number of central areas of computer science, including databases, compilers, and complexity theory.

8.1 Propositions and predicates

A *proposition* is a statement that is either true or false. For example, "It will rain tomorrow" and "It will not rain tomorrow" are propositions, but "It will probably rain tomorrow" is not, pending a more precise definition of "probably".

A *predicate* is a statement that contains a variable, such that for any specific value of the variable the statement is a proposition. Usually the allowed values for the variable will come from a specific set, sometimes called the *universe* of the variable, which will be either explicitly mentioned or clear from context. A simple example of a predicate is $x \geq 2$ for $x \in \mathbb{R}$. Clearly, for any real value of x, this statement is either true or false. We denote predicates in a similar way to functions, as in $P(x)$. In fact, the connection to functions runs deep: A predicate $P(x)$ can be considered a function, $P : \mathcal{U} \to \{0, 1\}$, where $\mathcal{U}$ is the universe of the variable x, 1 represents truth, and 0 represents falsehood.

A predicate may have more than one variable, in which case we speak of predicates in two variables, three variables, and so on, denoted as $Q(x, y)$, $S(x, y, z)$, etc.

8.2 Quantifiers

Given a predicate $P(x)$ that is defined for all elements in a set A, we can reason about whether $P(x)$ is true for all $x \in A$, or if it's at least true for some $x \in A$. We can state propositions to this effect using the *universal quantifier* $\forall$ and the *existential quantifier* $\exists$.

- $\forall x \in A : P(x)$ is true if and only if $P(x)$ is true for all $x \in A$. This proposition can be read "For all $x \in A$, $P(x)$."

- $\exists x \in A : P(x)$ is true if and only if $P(x)$ is true for at least one $x \in A$. This proposition can be read "There exists $x \in A$ such that $P(x)$."

Given a predicate in more than one variable we can quantify each (or some) of the variables. For example, the statement "For every real x and y, it holds that $x^2 - y^2 = (x - y)(x + y)$" can be formalized as

$$\forall x, y \in \mathbb{R} : x^2 - y^2 = (x - y)(x + y).$$

Somewhat more interestingly, the statement "There is no greatest integer" might be formulated as

$$\forall n \in \mathbb{Z}\ \exists m \in \mathbb{Z} : m > n.$$

It is crucial to remember that the meaning of a statement may change if the existential and universal quantifiers are exchanged. For example, $\exists m \in \mathbb{Z}\ \forall n \in \mathbb{Z} : m > n$ means "There is an integer strictly greater than all integers." This is not only contrary to the spirit of the original statement, but is patently wrong as it asserts in particular that there is an integer that is strictly greater than itself.

Exchanging the order of two quantifiers of the same type (either universal or existential) does not change the truth value of a statement. We do not prove this here.

8.3 Negations

Given a proposition P, the *negation* of P is the proposition "P is false". It is true if P is false, and false if P is true. The negation of P is denoted by $\neg P$, read as "not P." If we know the meaning of P, such as when P stands for "It will rain tomorrow," the proposition $\neg P$ can be stated more naturally than "not P," as in "It will not rain tomorrow." The truth-value of $\neg P$ can be represented by the following *truth table*:

P	$\neg P$
true	false
false	true

A truth table simply lists the truth values of particular statements in all possible cases. Something interesting can be observed in we consider the truth values of $\neg\neg Q$, which can be obtained by using the above table once with $P = Q$ and once with $P = \neg Q$:

Q	$\neg Q$	$\neg\neg Q$
true	false	true
false	true	false

We see that the statements Q and $\neg\neg Q$ have the same truth values. In this case we say that the two statements are *equivalent*, and write $Q \Leftrightarrow \neg\neg Q$. If $A \Leftrightarrow B$ we can freely use B in the place of A, or A instead of B in our logical derivations.

Negation gets really interesting when the negated proposition is quantified. Then we can assert that

$$\begin{aligned} \neg\forall x \in A : P(x) &\Leftrightarrow \exists x \in A : \neg P(x) \\ \neg\exists x \in A : P(x) &\Leftrightarrow \forall x \in A : \neg P(x) \end{aligned}$$

These can be interpreted as the claim that if $P(x)$ is not true for all $x \in A$ then it is false for some $x \in A$ and vice versa, and the claim that if $P(x)$ is not false for any $x \in A$ then it is true for all $x \in A$ and vice versa. What this means, in particular, is that if we want to disprove a statement that asserts something for all $x \in A$, it is sufficient to demonstrate *one* such x for which the statement does not hold. On the other hand, if we need to disprove a statement that asserts the existence of an $x \in A$ with a certain property, we actually need to show that for *all* such x this property does not hold.

Looked at another way, the above equivalences imply that if we negate a quantified statement, the negation can be "pushed" all the way inside, so that no negated quantifiers are left. Indeed, leaving any negated quantifiers is often considered a mark of poor style. Here is how how this elimination is done in a particular example:

$$\begin{aligned} \neg\forall n \in \mathbb{Z}\ \exists m \in \mathbb{Z} : m > n &\quad \Leftrightarrow \\ \exists n \in \mathbb{Z}\ \neg\exists m \in \mathbb{Z} : m > n &\quad \Leftrightarrow \\ \exists n \in \mathbb{Z}\ \forall m \in \mathbb{Z} : m \leq n & \end{aligned}$$

This can be read as "There exists an integer that is greater or equal to any other integer," which is the proper negation of the original statement.

8.4 Logical connectives

The symbol $\neg$ is an example of a *connective*. Other connectives combine two propositions (or predicates) into one. The most common are $\wedge$, $\vee$, $\oplus$, $\rightarrow$ and $\leftrightarrow$. $P \wedge Q$ is read as "P and Q"; $P \vee Q$ as "P or Q"; $P \oplus Q$ as "P xor Q"; $P \rightarrow Q$ as "P implies Q" or "if P then Q"; and $P \leftrightarrow Q$ as "P if and only if Q". The truth-value of these *compound propositions* (sometimes called *sentences*) depends on the truth values of P and Q (which are said to be the *terms* of these sentences), in a way that is made precise in the truth-table below.

We will not concern ourselves much with the $\oplus$ and $\leftrightarrow$ connectives, as they are encountered somewhat less frequently.

One interesting thing about the above table is that the proposition $P \rightarrow Q$ is false only when P is true and Q is false. This is what we would expect: If P is true but Q is false then, clearly, P does not imply Q. The important thing to remember is that if

P	Q	$P \wedge Q$	$P \vee Q$	$P \oplus Q$	$P \rightarrow Q$	$P \leftrightarrow Q$
T	T	T	T	F	T	T
T	F	F	T	T	F	F
F	T	F	T	T	T	F
F	F	F	F	F	T	T

P is false, then $P \rightarrow Q$ is true. One way this can be justified is by remembering that we expect a proposition to be either false or true. Now, $P \rightarrow Q$ being false says that P does not imply Q, which means precisely that P is true but Q is still false. In all other cases we expect $P \rightarrow Q$ to be true. (Did I succeed in turning something obvious into a confusing mess? Well, we all know what is paved with good intentions...)

Now, there is another statement involving P and Q that is false precisely when P is true and Q is false. It is, of course, $\neg P \vee Q$. As the following truth table demonstrates, the proposition $\neg P \vee Q$ is equivalent to $P \rightarrow Q$:

P	Q	$P \rightarrow Q$	$\neg P \vee Q$
T	T	T	T
T	F	F	F
F	T	T	T
F	F	T	T

This means something rather interesting: We can replace a proposition that involves implication by an equivalent one that instead has negation ($\neg$) and disjunction ($\vee$). Also, since $P \rightarrow Q$ is false only when P is true and Q is false, the proposition $\neg(P \rightarrow Q)$ is equivalent to $P \wedge \neg Q$:

$$\neg(P \rightarrow Q) \Leftrightarrow P \wedge \neg Q.$$

This means that in a negated implication, the negation can be "pushed inside", somewhat like with quantifiers. In fact, similar equivalences exist for other negated compound statements, as can be verified using truth tables (do it!):

$$\begin{aligned} \neg(P \vee Q) &\Leftrightarrow \neg P \wedge \neg Q \\ \neg(P \wedge Q) &\Leftrightarrow \neg P \vee \neg Q \end{aligned}$$

These are the famous *DeMorgan's laws.* What they mean is that we can eliminate negated compounds (sounds like a military operation, doesn't it?) just as we can eliminate negated quantifiers.

Here is another important logical equivalence: The implication $P \rightarrow Q$ is equivalent to the *contrapositive* implication $\neg Q \rightarrow \neg P$:

$$(P \rightarrow Q) \Leftrightarrow (\neg Q \rightarrow \neg P).$$

This is demonstrated by the following truth table:

P	Q	$P \to Q$	$\neg Q$	$\neg P$	$\neg Q \to \neg P$
T	T	T	F	F	T
T	F	F	T	F	F
F	T	T	F	T	T
F	F	T	T	T	T

Indeed, an implication of the form "If P then Q" is sometimes proved by assuming that Q does not hold and showing that under this assumption P does not hold. This is called a proof by contrapositive. (Despite the similarity, it is different from a proof by contradiction.)

8.5 Tautologies and logical inference

A sentence that is true regardless of the values of its terms is called a *tautology*, while a statement that is always false is a *contradiction*. Another terminology says that tautologies are *valid* statements and contradictions are *unsatisfiable* statements. All other statements are said to be *satisfiable*, meaning they can be either true or false.

Easy examples of a tautology and a contradiction are provided by $P \vee \neg P$ and $P \wedge \neg P$, as demonstrated by the following truth table:

P	$\neg P$	$P \vee \neg P$	$P \wedge \neg P$
T	F	T	F
F	T	T	F

Note that by our definition of logical equivalence, all tautologies are equivalent. It is sometimes useful to keep a "special" proposition $\mathcal{T}$ that is always true, and a proposition $\mathcal{F}$ that is always false. Thus any tautology is equivalent to $\mathcal{T}$ and any contradiction is equivalent to $\mathcal{F}$.

Here is another tautology: $(P \wedge Q) \to P$:

P	Q	$P \wedge Q$	$(P \wedge Q) \to P$
T	T	T	T
T	F	F	T
F	T	F	T
F	F	F	T

The statement $(P \wedge Q) \to P$ is read "P and Q implies P". The fact that this is a tautology means that the implication is always true. Namely, if we know the truth of $P \wedge Q$, we can legitimately conclude the truth of P. In such cases the symbol $\Rightarrow$ is used, and we can write $(P \wedge Q) \Rightarrow P$. There is a crucial difference between $(P \wedge Q) \to P$ and $(P \wedge Q) \Rightarrow P$. The former is a single statement, while the latter indicates a relationship between two statements. Such a relationship is called an *inference rule*. A similar inference rule, $P \Rightarrow P \vee Q$ can be established analogously.

In general, any tautology of the form $A \rightarrow B$ can be used to "manufacture" the inference rule $A \Rightarrow B$ that says that if we know A we can conclude B. Similarly, a tautology of the form $A \leftrightarrow B$ can be converted into the equivalence $A \Leftrightarrow B$, which can be regarded as two inference rules, $A \Rightarrow B$ and $B \Rightarrow A$. A particularly important inference rule is called *modus ponens*, and says that if we know that P and $P \rightarrow Q$ are both true, we can conclude that Q is true. It follows from the tautology $(P \wedge (P \rightarrow Q)) \rightarrow Q$:

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

We've already seen a number of inference rules above, like $(P \rightarrow Q) \Leftrightarrow (\neg Q \rightarrow \neg P)$, without calling them that. Here are three others, all corresponding to tautologies that you are invited to verify using truth tables:

$$\begin{aligned}
(\neg P \rightarrow \mathcal{F}) &\Leftrightarrow P \\
(P \leftrightarrow Q) &\Leftrightarrow (P \rightarrow Q) \wedge (Q \rightarrow P) \\
(P \leftrightarrow Q) &\Leftrightarrow (P \rightarrow Q) \wedge (\neg P \rightarrow \neg Q)
\end{aligned}$$

These three rules are of particular importance. The first formally establishes the validity of proofs by contradiction, and the second and third provide two means for proving "if and only if" statements. We've been using these all along, but now we know why they are justified.