Set Theory And Basic Logic

A Typical reader of this text is likely to be motivated by a need to deal with formal mathematics in his or her professional career. But I hope that in addition there will be some readers who will simply take pleasure in a mathematical journey toward a higher level of sophistication. There are many who would enjoy this trip, just as there are many who might enjoy listening to a symphony with a clear melodic line.

Mathematicians must frequently deal with assertions known as *propositions*; these are statements that are either true or false. Sometimes it is very easy to determine truth or falsity (for example, 3<5 is true) sometimes it may be very complicated. But if we have some propositions whose truth value is known, there is a mechanical procedure to determine the truth value of new propositions that are formed from the old using the connectives “or, “and” and “not”. The method uses *truth tables*, which we now describe.

OR Proposition:

A	B	A or B
T	T	T
T	F	T
F	T	T
F	F	F

The only time (A or B) is false is when both A and B are false.

A	B	C	A or B or C
T	T	T	T
T	T	F	T
T	F	T	T
T	F	F	T
T	T	T	T
T	T	F	T
T	F	T	T
F	F	F	F

The only time (A or B or C) is false, is when all three variables A, B, and C are false.

AND Connective:

A	B	A and B
T	T	T
T	F	F
F	T	F
F	F	F

The only time (A and B) is true is when both A and B are true.

A	B	C	A or B or C
T	T	T	T
T	T	F	F
T	F	T	F
T	F	F	F
T	T	T	F
T	T	F	F
T	F	T	F
F	F	F	F

The only time (A and B and C) is true, is when all three variables A, B, and C are true.

NOT Connective:

A	Not A
T	F
F	T

Always the opposite!

A implies B, is the same as (if A, then B)

A	B	A implies B
T	T	T
T	F	F
F	T	T
F	F	T

Given the statement: if A, then B, then the converse of the statement would be: if B, then A!

Note: if an assertion is true, it does not follow that the converse is also true.

If A implies B, and B implies A are both true, then we say that A and B are equivalent propositions.

So, A is equivalent to B mean, A implies B and B implies A, and they should have the same truth-value.

A	B	A is equivalent to B
T	T	T
T	F	F
F	T	F
F	F	T

A is equivalent to B is true, only when A and B both have the same value.

Some standard abbreviations:

- $(A \vee B)$ for (A or B)
- $(A \wedge B)$ for (A and B)
- $(A \Rightarrow B)$ for (A implies B)
- $(A \Leftrightarrow B)$ for (A is equivalent to B)
- $(\neg A)$ for (not A)

The notation for equivalence can be extended to *compound propositions*, which are constructed from individual propositions A, B, C, … by using the connectives "or", "and", and "not".

DeMorgan Laws:

D1	$[\neg(A \wedge B)] \Leftrightarrow [(\neg A) \vee (\neg B)]$
D2	$[\neg(A \vee B)] \Leftrightarrow [(\neg A) \wedge (\neg B)]$

What this means is that regardless of the truth or falsity of A and B, the left and right side have exactly the same truth-value.

A	**B**	**ØA**	**ØB**	**[(ØA) Ú (ØB)]**	**A Ù B**	**Ø(A Ù B)**
T	T	F	F	F	T	F
T	F	F	T	T	F	T
F	T	T	F	T	F	T
F	F	T	T	T	F	T

An implication that is guaranteed to be true because the hypothesis is always false is sometimes said to be *vacuously true*.

Quantifiers:

Not all mathematical statements are propositions, which have a definite truth-value. For example, consider the statement: x is less than 5, where x is a real number. The assertion is true for some values of x and false for other. It is possible to convert such a statement into a proposition by using quantifiers, which are of two types, *existential* and *universal*.

The *existential quantifier* ∃ means "there exists"; other equivalent phrases are "there is at least one" and "for some".

The *universal quantifier* ∀ means "for all", equivalently, "for every".

Thus again, thinking of x as a real number, ∃x (x<5), means that there exists an x such that x<5; in other words, there is at least one real number that is less than 5. Which is certainly true.

On the other hand, ∀x (x<5), means that for all x, x<5; in other words, every real number is less than 5. This is definitely false.

More than one quantifier can appear in a sentence. For example, with x and y real, ∀x∃y (x + y = 13), says that for every x there exists a y such that x + y = 13. This is true; if you give an x, we can solve for y: y = 13 – x.

However, consider the assertion ∃x∀y (x + y = 13), this says that for some x, every y satisfies x + y = 13, definitely false.

There is a very convenient mechanical procedure for finding the negation of a statement involving quantifiers. Here is how it works in a typical case; let's analyze the assertion Ø**[(**$**x)(x<5)]**. We are saying that it is not the case that some x is less than 5; in other words, every x is greater than or equal to 5. Thus we have "**x(x**³**5)**.

So to go from ∃x (x<5) to its negation, we have reversed the quantifiers (from ∃ to ∀), and changed the main statement x<5 to its negative.

Similarly to negate "**x (x**³**5)**, i.e., to say that "every x is at least 5 or greater" is false; we assert that some x is less than 5. So we arrive at $**x (x<5)**.

Taking the negation twice brings us right back to the original statement!

This process works also when there is more than one quantifier. For example, to negate "**x**$**y (x + y = 13)**, notice that it is of the form ∀x P, so that its negative is ∃x (not P). But P is ∃y (x + y = 13), so (not P) is ∀y (x + y ≠ 13). Thus the negation of "**x**$**y (x + y = 13)** is $**x**"**y (x + y** ¹ **13)**, which illustrates how the method works in general. Proceeding from left to right, reversing quantifiers as you go, until you reach the main statement, which you change to its negative.

The truth or falsity of statement involving quantifiers depends crucially on the allowable values of the variables. For example, **"x ($y>0)(x + y = 13)**, which says that for all x there exists a positive real y such that x + y = 13. This is false, e.g., take x to be 13 or larger. Taking the negation of the statement gives **$x ("y>0)(x + y ¹ 13)**, this is true, again, take x to be 13 or larger. If you add a positive number to x, you certainly cannot get 13.

Caution: The phrase "> 0" in the statements refer to the allowable values of the variable y. It is not changed when we go from a statement to its negation.

Interesting:

Perhaps you recall the definition of a limit from calculus, and perhaps you remember being baffled by it. The limit statement is actually a very complicated mathematical sentence using three quantifiers, and it is not surprising that it causes confusion. What does it mean to say the sequence $x_1, x_2, \ldots, x_n$ of real numbers converges to the real number x? Intuitively, as n gets large, x_n gets very close to x. How close you ask? As close as you wish! For example, suppose we want the sequence to be within 10^{-9} of x, i.e., $x - 10^{-9} < x_n < x + 10^{-9}$. How large must n be? It might turn out in a particular case that all x_n's from $n = 10^{15}$ onward satisfy this inequality. Thus for every degree of closeness you might give, there is some point in the sequence so that every x_n from that point on is that close to x. The degree of closeness is measured by a small positive number $\in$, 10^{-9} in this case. The point in the sequence that achieves the desired closeness is measured by a positive integer N; in this case, $N = 10^{15}$. We can then write the formal definition of convergence of x_n to x:

x_n ® x means ("Î > 0)($N)("n ³ N)(|$x_n$ – x| < Î)

For every positive real number $\in$ there is a positive integer N such that for every positive integer n = N, N+1, N+2, … , x_n differs from x by less than $\in$. If you did not understand the definition of a limit in calculus and it makes a bit more sense now, fine. If it is still obscure, there is no problem; we will not be using it. I just wanted to make a point that quantifiers will be encountered in all areas of mathematics!

Functions:

We have worked with functions many times. For example, if $f(x) = x^2$, where x is a real number, then if I give you the value of x, you can square it to produce the value of f(x), a non-negative real number. In general, a function or a mapping from a set A to a set B is a rule that assign to each element x in A an element f(x) in B. We write $f:A \rightarrow B$; A is called the *domain* of the function f, and B the *co-domain*. In the above example A is the set of reals, and B the set of non-negative reals.

Note: Since a non-negative real number is in particular a real number, we can if we like change B to the set of all reals. This does not change the fact that $f(x) = x^2$, but technically, B is part of the description of f, and a fussy mathematician would insist that we now have a slightly different function.

One of the most important operations of function is that of *composition*, and again this is a familiar idea. For example, if $h(x) = \sin^2(x)$, then given the value of x, we first compute sin(x) and then square the result. Thus if f(x) = sin(x) and $g(y) = y^2$, we have h(x) = g(f(x)).

In general, if $f:A \rightarrow B$ and $g:B \rightarrow C$, the composition of f and g is the function $h:A \rightarrow C$ defined by $h(x) = g(f(x)) \ \forall \ x \in A$. We write $h = g \bullet f$, indicating that to compute h, we first calculate f and then g.

In mathematics we very often ask whether a particular function has a certain property. For example, if $f(x) = x^2$, x real, and I give you the value of f(x), can you determine x? Consider the following; if f(x) = 16, then x can be 4, but it can also be –4. In general, if y is any positive real number, then there are two x's such that $x^2 = y$, namely $x = \sqrt{y}$ and $x = -\sqrt{y}$.

Drawing the graph of f, notice that the horizontal line at height y will intersect the graph in two points. However, when y = 0, there is only one value of x, x = 0; the graph touches the x-axis at only one point.

Here is another question. If $f(x) = x^2$ and we take the co-domain B to be the non-negative reals, is B completely covered? In other words given a y in B, is there at least one x such that f(x) = y? The answer is YES, because of the analysis we just did.

However, consider the following, suppose we had taken the co-domain to be the entire set of reals. Now the answer is NO, because if y is a negative number, there is no way in the world for you to find an x such that $y = x^2$.

Definition:

(1) The function $f:A \rightarrow B$ is said to be *injective* or *one-to-one* if no two distinct x's in A yield the same value of f(x). In other words, if $x_1 \neq x_2$, then $f(x_1) \neq f(x_2)$. Equivalently, take the contra positive, if $f(x_1) = f(x_2)$, then $x_1 = x_2$.

(2) The function $f:A \rightarrow B$ is said to be *surjective* or *onto*, if given any y in B, there is at least one x in A such that f(x)=y.

(3) The function $f:A \rightarrow B$ is said to be *bijective* or a *bijection* or *one-to-one onto* or a *one-to-one correspondence* if f is both injective and surjective.

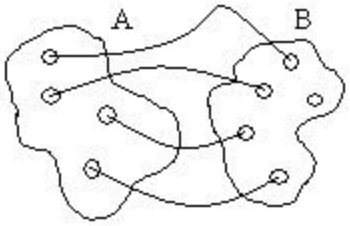

Injective or one-to-one

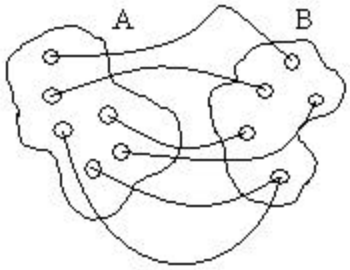

Surjective or onto

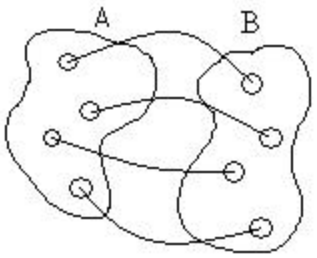

Bijective or one-to-one onto

If the statement that f(x) = y is represented by an arrow from x to y, then if f is not injective, there will be more than one arrow landing at the same point y. If f is not surjective, there will be points of B that are never hit by an arrow.

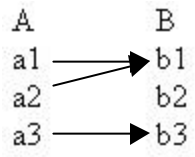

In this example we have $f(a_1)=f(a_2)=b_1$ and $f(a_3)=b_3$. The function f is not injective because $f(a_1)=f(a_2)=b_1$, and f is not surjective because the is no $a \in A$ such that $f(a)=b_2$.

Going back to $f(x) = x^2$ from that reals to the non-negative reals, then f is surjective but not injective. If f(x) = x from the non-negative reals to the reals, then f is injective but not surjective. But *for finite set of the same size*, this situation cannot occur.

Theorem:
If $f:A \rightarrow B$ where A and B each have n elements for some positive integer n, then f is injective iff f is surjective.

Proof:
If f is not injective, there will be at least two arrows from distinct points of A to a single point of B. But then for f to be surjective, the remaining n-1 points of B must be covered by at most n-2 remaining arrows, which is not possible. We conclude that f can not be surjective. Conversely, assume that f is not surjective, so that some point of B is not covered. Then we have n arrows from A to B landing on at most n-1 points, so there must be at least two arrows that land on the same point, so that f is not injective.

If $f:A \rightarrow B$ and C is a subset of B, very often we will be interested in the set of points in A that map into C.

Definition:
Let $f:A \rightarrow B$, with $C \subseteq B$. The pre-image of C under f is defined by $f^{-1}(C)=\{x \in A: f(x) \in C\}$

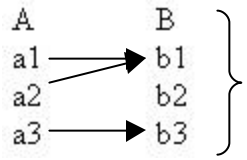

$f^{-1}\{b_1\}=f^{-1}\{b_1, b_2\}$, $f^{-1}\{a_1,a_2\}$, $f^{-1}\{b_2\}=\phi$,
$f^{-1}\{b_1,b_2,b_3\}=\{a_1,a_2,a_3\}$

In general the pre-image behaves extremely well with respect to the operation of union, intersection and complement.

Theorem:
Let $f:A\rightarrow B$, and suppose we have an arbitrary collection of subset B_i of B. Then

$f^{-1}(\cup_i B_i) = \cup_i f^{-1}(B_i)$
$f^{-1}(\cap_i B_i) = \cap_i f^{-1}(B_i)$
$f^{-1}(B_i^c) = [f^{-1}(B_i)]^c$

Thus the pre-image of a union is the union of the pre-images. The pre-image of an intersection, is the intersection of the pre-images. And the pre-image of the complements, is the complement of the pre-image.

Proof:
We have the following for the argument of the union

$x \in f^{-1}(\cup_i B_i)$ iff $f(x) \in \cup_i B_i$, iff $f(x) \in B_i$ for at least one i, iff $x \in f^{-1}(B_i)$ for at least one i, iff $x \in \cup_i f^{-1}(B_i)$

for the argument of the intersection is as follows

$x \in f^{-1}(\cap_i B_i)$ iff $f(x) \in \cap_i B_i$, iff $f(x) \in B_i$ " i, iff $x \in f^{-1}(B_i)$ " i, iff $x \in \cap_i f^{-1}(B_i)$

and for the complement

$x \in f^{-1}(B_i^c)$ iff $f(x) \in B_i^c$, iff $f(x) \notin B_i$, iff $x \notin f^{-1}(B_i)$, iff $x \in [f^{-1}(B_i)]^c$

If $f:A\rightarrow B$ and C a subset of A, the image or direct image of C under f is defined as the set of all possible values f(x) generated by allowing x to range over C. Thus **$f(C)=\{f(x): x \in C\}$**, and $y \in f(x)$ for some x belonging to C. The image of f is defined as f(A). Unfortunately, we cannot prove a theorem for direct images; however, there are a few reasonably nice properties.

Theorem: Let $f:A\rightarrow B$, then
(a) if $C \subseteq A$ then $C \subseteq f^{-1}(f(C))$
(b) if $D \subseteq B$ then $f[f^{-1}(D)] \subseteq D$
(c) $f(\cup_i A_i) = \cup_i f(A_i)$
(d) $f(\cap_i A_i) = \cap_i f(A_i)$

Proof:
(a) if $x \in C$ then $f(x) \in f(C)$, in other words, $x \in f^{-1}[f(C)]$
(b) if $y \in f[f^{-1}(D)]$, then $y=f(x)$ for some $x \in f^{-1}(D)$. But then $f(x) \in D$, and since $y=f(x)$ we have $y \in D$.
(c) we have $y \in f(\cup_i A_i)$ iff $y=f(x)$ for some $x \in \cup_i A_i$, iff for at least one i, $y=f(x)$ for some $x \in A_i$, iff for at least one i, $y \in f(A_i)$, iff $y \in \cup_i f(A_i)$
(d) if $y \in f(\cap_i A_i)$ then $y=f(x)$ for some $x \in \cap_i A_i$, and therefore y belongs to $f(A_i)$ for all i, that is, $y \in \cap_i f(A_i)$

The inclusions in (a), (b), and (d) may be proper, again taking a look at the following:

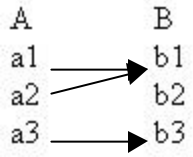

$f^{-1}[f(\{a_2\})] = f^{-1}(\{b_1\}) = \{a_1,a_2\}$
$f[f^{-1}\{b_1,b_2\}] = f(\{a_1,a_2\}) = \{b_1\}$
$f(\{a_1,a_3\} \cap \{a_2,a_3\}) = f(\{a_3\}) = \{b_3\}$
but $f(\{a_1,a_3\}) \cap f(\{a_2,a_3\}) = \{b_1,b_3\} \cap \{b_1,b_3\} = \{b_1,b_3\}$

Relations:

If f is a function from A to B, then given any x in the domain A, the rule defining f produces an element f(x) in the co-domain B. There is no ambiguity; once x is specified f(x) is determined. What would happen if we were to allow more than one from some of the points in A? We would not have a function anymore, but we would arrive at a new mathematical object called a *relation*. To define this formally, we need some terminology.

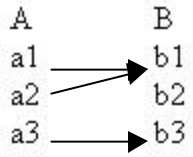

Definition: An *ordered pair* (a,b) is roughly a set of two elements in which order counts, so that for example, (3,5) is not the same as (5,3). The key idea is that two ordered pairs (a,b) and (c,d) are regarded as identical if and only if a=c and b=d.

One way to capture this notion is to define an ordered pair as function
$f : \{1,2\} \rightarrow \{a,b\}$ with f(1)=a and f(2)=b.

Similarly, an ordered n-tuple $(a_1, \ldots, a_n)$ may be defined as a function
$f : \{1, 2, \ldots, n\} \rightarrow \{a_1,\ldots,a_n\}$ with $f(i)=a_i$, $1 \le i \le n$.

If A and B are sets, the *Cartesian product* of A and B, written A x B, is the set of all ordered pairs (a,b), $a \in A$, $b \in B$. Similarly, the Cartesian product of n sets $A_1,\ldots,A_n$, written $A_1 \times \ldots \times A_n$, is the set of all ordered n-tuples $(a_1, \ldots, a_n)$, where $a_i \in A_i$, $1 \le i \le n$. If $A_i=A$ for all i, the Cartesian product is often written as A^n.

A *relation* between two sets A and B is the subset of A x B; a *relation on A* is a subset of A x A.

Similarly, an n-ary relation among sets $A_1, \ldots, A_n$ is a subset of $A_1 \times \ldots \times A_n$; an n-ary relation on A is a subset of A^n.

Example:
We define a relation $\Re$ on the set A = {1,2,3,4,5,6,7,8} as follows: the ordered pair (a,b) will belong to $\Re$ (sometimes denoted $a\Re b$) iff a<b and a divides b. Here is a list of all the ordered pairs in $\Re$.

(1,2), (1,3), (1,4), (1,5), (1,6), (1,7), (1,8), (2,4), (2,6), (2,8), (3,6), (4,8)

Note: that $(3,3)\notin\Re$; although 3 divides 3, 3 is not less than 3, this is also true for (1,1), (2,2), (4,4), (5,5), (6,6), (7,7), and (8,8).

You can see the reason for the term relation. For (a,b) to belong to $\Re$, a must be related to b in a certain way, namely a must be less than b, and it must also be a divisor of b as well.

A relation between A and B can be represented by a mapping diagram, but in this case, there can be more than one arrow leaving a point of A. For our relation example given, there will be seven arrows leaving the point 1, three arrows from point 2, but single arrows leaving from points 3 and 4.

We are going to concentrate on the two types of relations that are most important in all areas of mathematics: equivalence relations and partial orderings.

Equivalence Relations:
Let A = $\mathbb{Z}$, the set of all integers, and define a relation $\Re$ on a as follows: $a\Re b$ iff a-b is divisible by 4. In this case we say that a is congruent to b modulo 4, written as $a\equiv b \bmod 4$. In a similar fashion we may define the relation of congruence modulo m for any positive integer $m\geq 2$.

Note: Negative integers can be used, but they do not contribute anything significant since a-b is divisible by –m iff it is divisible by m. The case m=1 is legal but uninteresting since any two integers are congruent modulo 1.

Now a-b is divisible by 4 iff a and b leave the same remainder when divided by 4. For example, if a=-5 and b=7, then $a\Re b$ since a-b = -12 = 4(-3). Also a = 4(-2) + 3 and b = 4(1) + 3, so both a and b leave remainder 3.

In general, if a = 4s + i and b = 4t + i, then a-b = 4(s-t), a multiple of 4. Conversely, if a-b = 4s and b = 4t + i, then a = 4(s+t) + i, so both a and b leave remainder i. The relation $\Re$ has the following properties:

1. $\Re$ is *reflexive*, $a\Re a\ \forall\ a\in A$ (since a-a=0, which is divisible by 4)
2. $\Re$ is *symmetric*, if $a\Re b$ then $b\Re a$ (if a-b=4s, then b-a=4(-s))
3. $\Re$ is *transitive*, if $a\Re b$ and $b\Re c$ then $a\Re c$ (if a-b=4s and b-c=4t then a-c=(a-b) + (b-c) = 4(s+t))

A relation that is reflexive, symmetric and transitive is called an *equivalence relation*.

If $a \in A$, the set S(a) of elements equivalent to a (that is, the set of all $b \in A$ such that $b \Re a$) is called the equivalence class of a. Lets find the equivalence class for our explicit example, congruence modulo 4. S(0) is the set of integers that leave remainder 0 when divided by 4, so that S(0) consists of all multiples of 4. Thus

$$S(0) = \{ \ldots, -8, -4, 0, 4, 8, \ldots \}$$

Similarly

$$S(1) = \{ \ldots, -7, -3, 1, 5, 9, \ldots \}$$
$$S(2) = \{ \ldots, -6, -2, 2, 6, 10, \ldots \}$$
$$S(3) = \{ \ldots, -5, -1, 3, 7, 11, \ldots \}$$

Notice that the sets S(i), i=0, 1, 2, 3, from a partition of A, that is, they are disjoint and their union is A. But what about the sets S(i) where $i \neq 0, 1, 2, 3$? We won't get any new results! For example, S(5) is the set of all elements equivalent to 5, i.e., the set of all elements that leave remainder 1 when divided by 4. Thus S(5) coincides with S(1). In general, if $b \in S(a)$, then $S(b) \in S(a)$, as you will see in a moment.

Let's now prove that if $\Re$ is any equivalence relation on A, then the equivalence classes form a partition of A. For any $a \in A$ we have $a \in S(a)$ by reflexive, so that the union of the equivalence classes is A. Now considering two equivalence classes S(a) and S(b). There are two cases:

Case 1: b is equivalent to a, so that $b \in S(a)$. Then $S(b)=S(a)$. For if c belongs to S(b), then $c \Re a$ by transitivity, and $c \in S(a)$. Conversely, if $c \in S(a)$ then $c \Re a$; but $b \Re a$, so that $a \Re b$ by symmetry, and therefore $c \Re b$ by transitivity, proving that $c \in S(b)$.

Case 2: b is not equivalent to a, so that $b \notin S(a)$. Then $S(b) \cap S(a) = \phi$. For if c belongs to both S(b) and S(a), then $c \Re b$ and $c \Re a$; but then $b \Re c$ by symmetry, hence $b \Re a$ by transitivity. Thus $b \in S(a)$, which is contradiction.

We have shown that the equivalence classes are disjoint sets whose union is A, as required.

You may not have seen congruence modulo m previously, but there is one equivalence relation that is extremely familiar. Let $\Re$ be the equality relation on A, that is, $a \Re b$ iff a=b. Equality is reflexive (a=a), symmetric (if a=b, then b=a), and transitive (if a=b and b=c, then a=c). The equivalence class of the element $a \in A$ is the set {a} consisting of a alone.

Partial Orderings: