

Appendix A

Sets, Relations and Functions

Summary

This chapter explains the basics of formal set notation, and gives an introduction to relations and functions. The chapter ends with a short account of the principle of proof by mathematical induction.

A.1 Sets and Set Notation

Many mathematical notions — some say all mathematical notions — can be defined in terms of the fundamental concept of a set. This is good reason for starting with some basic set theory.

A set is a collection of definite, distinct objects. Examples are the set of colours of the Dutch flag, or the set of letters of the Greek alphabet. Yet another example is the set of even natural numbers greater than seven. And so on.

The elements of a set are also called its *members*. To indicate that a is an element of a set A we write $a \in A$. To deny that a is an element of a set A we write $a \notin A$. The symbol \in is the symbol for membership.

The elements of a set can be anything: words, colours, people, numbers. The elements of a set can also themselves be sets. The set consisting of the set of even natural numbers and the set of odd natural numbers is an example. This set has two elements; each of these elements has itself an infinite number of elements.

To check whether two sets are the same one has to check that they have the same elements. The fact that membership is all there is to set identity, or that sets are fully determined by their members, is called the *principle of extensionality*. It follows that to check that two sets A and B are identical, one has to check two things:

- does it hold that every element a of A is also an element of B , and
- does it hold that every element b of B is also an element of A ?

To specify a set, there are several methods: give a list of its members, as in ‘the set having the numbers 1, 2 and 3 as its only members’, give some kind of semantic description, as in ‘the set of colours of the Dutch flag’, or separate out a set from a larger set by means of a suitable restriction. This last method is called the method of *set comprehension*. Here is an example: the odd natural numbers are the natural numbers with the property that division by 2 leaves a remainder of 1. We can express this by means of the pattern $2n + 1$, as follows:

$$O = \{2n + 1 \mid n \in \mathbb{N}\}.$$

The braces are also used to list the members of a finite set:

$$D = \{\text{red, white, blue}\}.$$

Mentioning a set element more than once does not make a difference. The set

$$\{\text{white, blue, white, red}\}$$

is identical to the set D , for it has the same members.

Another way of specifying sets is by means of operations on other sets. An example is the following definition of the odd natural numbers:

$$E = \mathbb{N} - O.$$

Here $\mathbb{N} - O$ is the set of all elements of \mathbb{N} that are not members of O . Equivalent definitions are the following:

$$E = \{n \in \mathbb{N} \mid n \notin O\}$$

or

$$E = \{2n \mid n \in \mathbb{N}\}.$$

Some important sets have special names. \mathbb{N} is an example. Another example is \mathbb{Z} , for the set of integer numbers. Yet another example is the set without any members. Because of the principle of extensionality there can be only one such set. It is called \emptyset or the empty set.

If every member of a set A is also a member of set B we say that A is a subset of B , written as $A \subseteq B$. If $A \subseteq B$ and $B \subseteq A$ then it follows by the principle of extensionality that A and B are the same set. Conversely, if $A = B$ then it follows by definition that $A \subseteq B$ and $B \subseteq A$.

Exercise A.1 Explain why $\emptyset \subseteq A$ holds for every set A .

Exercise A.2 Explain the difference between \emptyset and $\{\emptyset\}$.

The *complement* of a set A , with respect to some fixed universe, or: domain, U with $A \subseteq U$, is the set consisting of all objects in U that are not elements of A . The complement set is written as \bar{A} . It is defined as the set $\{x \mid x \in U, x \notin A\}$. For example, if we take U to be the set \mathbb{N} of natural numbers, then the set of even numbers is the complement of the set of odd numbers and vice versa.

A.2 Relations

By a relation we mean a meaningful link between people, things, objects, whatever. Usually, it is quite important what kind of relationship we have in mind.

Formally, we can describe a *relation* between two sets A and B as a collection of ordered pairs (a, b) such that $a \in A$ and $b \in B$. An ordered pair is, as the name already gives away, a collection of two distinguishable objects, in which the order plays a role. E.g., we use $(Bill, Hillary)$ to indicate the ordered pair that has *Bill* as its first element and *Hillary* as its second element. This is different from the pair $(Hillary, Bill)$ where *Bill* plays second fiddle.

The notation for the set of all ordered pairs with their first element taken from A and their second element taken from B is $A \times B$. This is called the *Cartesian product* of A and B . A relation between A and B is a subset of $A \times B$.

The Cartesian product of the sets $A = \{a, b, \dots, h\}$ and $B = \{1, 2, \dots, 8\}$, for example, is the set

$$A \times B = \{(a, 1), (a, 2), \dots, (b, 1), (b, 2), \dots, (h, 1), (h, 2), \dots, (h, 8)\}.$$

This is the set of positions on a chess board. And if we multiply the set of chess colours $C = \{\text{White, Black}\}$ with the set of chess figures,

$$F = \{\text{King, Queen, Knight, Rook, Bishop, Pawn}\},$$

we get the set of chess pieces $C \times F$. If we multiply this set with the set of chess positions, we get the set of piece positions on the board, with $(\text{White, King}, (e, 1))$ indicating that the white king occupies square $e1$. To get the set of moves on a chess board, take $((C \times F) \times ((A \times B) \times (A \times B)))$, and read $((\text{White, King}, ((e, 1), (f, 2)))$ as ‘white king moves from $e1$ to $f2$ ’, but bear in mind that not all moves in $((C \times F) \times ((A \times B) \times (A \times B)))$ are legal in the game.

$A \times A$ is sometimes also denoted by A^2 . Similarly for $A \times A \times A$ and A^3 , and so on.

As an example of a relation as a set of ordered pairs consider the relation of authorship between a set A of authors and a set B of books. This relation associates with every author the book(s) he or she wrote.

Sets of ordered pairs are called binary relations. We can easily generalize this to sets of triples, to get so-called ternary relations, to sets of quadruples, and so on. An example

of a ternary relation is that of borrowing something from someone. This relation consists of triples, or: 3-tuples, (a, b, c) , where a is the borrower, b is the owner, and c is the thing borrowed. In general, an n -ary relation is a set of n -tuples (ordered sequences of n objects). We use A^n for the set of all n -tuples with all elements taken from A .

Unary relations are called *properties*. A property can be represented as a set, namely the set that contains all entities having the property. For example, the property of being divisible by 3, considered as a property of integer numbers, corresponds to the set $\{\dots, -9, -6, -3, 0, 3, 6, 9, \dots\}$.

An important operation on binary relations is composition. If R and S are binary relations on a set U , i.e. $R \subseteq U^2$ and $S \subseteq U^2$, then the composition of R and S , notation $R \circ S$, is the set of pairs (x, y) such that there is some z with $(x, z) \in R$ and $(z, y) \in S$. E.g., the composition of $\{(1, 2), (2, 3)\}$ and $\{(2, 4), (2, 5)\}$ is $\{(1, 4), (1, 5)\}$.

Exercise A.3 What is the composition of $\{(n, n + 2) \mid n \in \mathbb{N}\}$ with itself?

Another operation on binary relations is converse. If R is a binary relation, then its converse (or: inverse) is the relation given by $R^\sim = \{(y, x) \mid (x, y) \in R\}$. The converse of the relation ‘greater than’ on the natural numbers is the relation ‘smaller than’ on the natural numbers. If a binary relation has the property that $R^\sim \subseteq R$ then R is called *symmetric*. It is also denoted as

$$\forall x \forall y (Rxy \rightarrow Ryx). \quad (\text{A.1})$$

Exercise A.4 Show that it follows from $R^\sim \subseteq R$ that $R = R^\sim$.

If U is a set, then the relation $I = \{(x, x) \mid x \in U\}$ is called the identity relation on U . If a relation R on U has the property that $I \subseteq R$, i.e. if every element of U stands in relation R to itself, then R is called *reflexive*. The relation \leq (‘less than or equal’) on the natural numbers is reflexive, the relation $<$ (‘less than’) is not. The relation $A^2 - I$ is the set of all pairs $(x, y) \in A^2$ with $x \neq y$. If A is the set $\{a, bc\}$, then $A^2 - I$ gives the following relation:

$$\{(a, b), (a, c), (b, a), (b, c), (c, a), (c, b)\}.$$

A relation R is called *transitive* if it holds for all x, y, z that if $(x, y) \in R$ and $(y, z) \in R$, then also $(x, z) \in R$. To say that the relation of friendship is transitive boils down to saying that it holds for anyone that the friends of their friends are their friends.

Exercise A.5 Which of the following relations are transitive?

- (1) $\{(1, 2), (2, 3), (3, 4)\}$
- (2) $\{(1, 2), (2, 3), (3, 4), (1, 3), (2, 4)\}$
- (3) $\{(1, 2), (2, 3), (3, 4), (1, 3), (2, 4), (1, 4)\}$
- (4) $\{(1, 2), (2, 1)\}$

$$(5) \{(1, 1), (2, 2)\}$$

The next exercise shows that transitivity can be expressed in terms of relational composition.

Exercise A.6 Check that a relation R is transitive if and only if it holds that $R \circ R \subseteq R$.

Exercise A.7 Can you give an example of a transitive relation R for which $R \circ R = R$ does not hold?


A.3 Back and Forth Between Sets and Pictures

A domain of discourse with a number of 1-place and 2-place predicates on is in fact a set of entities with certain designated subsets (the 1-place predicates) and designated sets of pairs of entities (the 2-place predicates).


Relations are sets of pairs, and it is useful to acquire the skill to mentally go back and forth between sets-of-pairs representation and picture representation. Take the following simple example of a relation on the set $\{1, 2, 3\}$.

$$\{(1, 2), (1, 3), (2, 3)\}. \quad (\text{A.2})$$


Here is the corresponding picture:



Exercise A.8 Give the set of pairs that constitutes the relation of the following picture:




For another example, consider the picture:



No arrowheads are drawn, which indicates that the pictured relation is symmetric. Here is the representation of the same relation as a set of pairs:


$$\{(1, 2), (2, 1), (1, 3), (3, 1), (2, 3), (3, 2)\}.$$

Exercise A.9 Give the representation of the pictured relation as a set of pairs:



A.4 Relational Properties

Talking about pictures with predicate logic is very useful to develop a clear view of what relational properties the predicate logical formulas express. Predicate logic is very precise, and it takes practice to get used to this precision. Consider the following picture.




A relation is *transitive* if from the facts that there are links from x to y and a link from y to z it follows that there is a link from x to z . Here is a formula for this:


$$\forall x \forall y \forall z ((Rxy \wedge Ryz) \rightarrow Rxz). \quad (\text{A.3})$$

Let us check whether the link relation in the last picture is transitive. It may seem at first sight that it is, for what transitivity expresses is that if you can go from x to z by first taking an R step from x to y and next another R step from y to z , between, then there is also a direct R step from x to z . This seems indeed to be the case in the picture. But there is a snag. In reasoning like this, we assume that the three points x , y and z are all *different*. But this is not what the formula says. Take any two different points in the picture. Surely there is a link from the first point to the second. But the linking relation is symmetric: it goes in both directions. Therefore there also is a link from the second point back to the first. But this means that the first point has to be R related to itself, and it isn't. So the relation in the picture is not transitive after all.

Can we also come up with a picture of three points with a symmetric linking relation, where the relation is transitive? Yes, there are several possibilities. Here is the first one:




But there is another possibility. Take the following picture:



This is a picture where the link relation is empty. There are no links, so it trivially holds that if one can get from a point to a point via two links, then one can also get there with a single link. So the empty relation is transitive.

Exercise A.10 Give all the transitive link relations on a domain consisting of three individuals, on the assumption that the link relation is symmetric. We have already seen two examples: the empty relation (no points are linked) and the total relation (all points are linked). What are the other possibilities? Draw pictures!

The relations in the pictures above were all symmetric: links were the same in both directions. The following picture with arrows gives a relation that is not symmetric. We need the arrows, for now the directions of the links matter:



Again we use R to refer to the binary relation in the picture. Again we can ask if the relation of the picture is transitive. This time the answer is ‘yes’. If we can get from x to y with two \rightarrow steps, then we can also get from x to y with a single step.

Not only is the relation in the picture not symmetric, but something stronger holds:

$$\forall x \forall y (Rxy \rightarrow \neg Ryx). \quad (\text{A.4})$$

Formula (A.4) expresses that the relation R is *asymmetric*.

Exercise A.11 Give an example of a binary relation on a domain of three objects that it is neither symmetric nor asymmetric.

The relation in the current picture also has another property, called *irreflexivity*:


$$\forall x \neg Rxx. \quad (\text{A.5})$$

This expresses that the R relation does not have self loops. We say: the relation is *irreflexive*.

The dual to *irreflexivity* is the property of having *all* self loops. This is called *reflexivity*:

$$\forall x Rxx. \quad (\text{A.6})$$

Here is an example of a reflexive relation:



Exercise A.12 Show that any asymmetric relation has to be irreflexive. (Hint: assume that a relation is asymmetric, and suppose it contains a loop (x, x) . Why is this impossible?)

A binary relation R is called an *equivalence relation* if R has the following three properties: (i) R is reflexive, (ii) R is symmetric, (iii) R is transitive.

Exercise A.13 Give all equivalence relations on a domain consisting of three objects. Draw pictures!

Exercise A.14 Consider the three predicate logical sentences (A.1), (A.3) and (A.6). These sentences together express that a certain binary relation R is an equivalence relation: symmetric, transitive and reflexive. Show that none of these sentences is semantically entailed by the other ones by choosing for each pair of sentences a model (situation) that makes these two sentences true but makes the third sentence false. In other words: find three examples of binary relations, each satisfying just two of the properties in the list (A.1), (A.3) and (A.6). This shows, essentially, that the definition of being an equivalence cannot be simplified (why?).

Exercise A.15 Consider the following predicate logical formulas:

- $\forall x \forall y (Rxy \rightarrow \neg Ryx)$ (R is asymmetric)
- $\forall x \exists y Rxy$ (R is serial)
- $\forall x \forall y \forall z ((Rxy \wedge Ryz) \rightarrow Rxz)$ (R is transitive).

Take any situation with a non-empty domain of discourse, with a binary relation on it. Show: if the three formulas are true of this situation, then the domain of discourse must be *infinite*. (Hint: start with a domain consisting of a single individual d_1 . Then by seriality there has to be an R -successor to d_1 . Suppose we take d_1 as its own R -successor. Then this would get us in conflict with we are in conflict with asymmetry, by Exercise ???. So there has to be a d_2 with (d_1, d_2) in R . And so on ...)

Exercise A.16 Consider again the three properties of asymmetry, seriality and transitivity of the previous exercise.

- (1) Give a picture of a finite situation with a relation R that is asymmetric and serial but not transitive.
- (2) Give a picture of a finite situation with a relation R that is serial and transitive but not asymmetric.
- (3) Give a picture of a finite situation with a relation R that is transitive and asymmetric but not serial.

A.5 Functions

Functions are relations with the following special property: for any (a, b) and (a, c) in the relation it has to hold that b and c are equal. Thus a *function* from a set A (called *domain*) to a set B (called *range*) is a relation between A and B such that for each $a \in A$ there is

one and only one associated $b \in B$. In other words, a function is a mechanism that maps an input value to a uniquely determined output value. Looking at the relation *author of* from above, it is immediately clear that it is not a function, because the input *Michael Ende* is not mapped to a unique output but is related to more than one element from the set B of books.

Functions are an important kind of relations, because they allow us to express the concept of *dependence*. For example, we know that the gravitational potential energy of a wrecking ball depends on its mass and the height we elevated it to, and this dependence is most easily expressed in a functional relation.

Functions can be viewed from different angles. On the one hand, they can be seen as sets of data, represented as a collection of pairs of input and output values. This tells us something about the behaviour of a function, i.e. what input is mapped to which output.

The function converting temperatures from Kelvin to Celsius can be seen as a set of pairs $\{(0, -273.15), \dots\}$, and the function converting temperatures from Celsius to Fahrenheit as a set $\{(-273.15, -459.67), \dots\}$. Determining the output of the function, given some input, simply corresponds to a table lookup. Any function can be viewed as a – possibly infinite – database table. This is called the extensional view of functions. Another way to look at functions is as *instructions for computation*. This is called the intensional view of functions. In the case of temperature conversion the intensional view is more convenient than the extensional view, for the function mapping Kelvin to Celsius can easily be specified as a simple subtraction

$$x \mapsto x - 273.15$$

This is read as ‘an input x is mapped to x minus 273.15’. Similarly, the function from Celsius to Fahrenheit can be given by

$$x \mapsto x \times \frac{9}{5} + 32$$

For example, if we have a temperature of 37 degrees Celsius and want to convert it to Fahrenheit, we replace x by 37 and compute the outcome by multiplying it with $\frac{9}{5}$ and then adding 32.

$$37 \times \frac{9}{5} + 32 \rightarrow 66.6 + 32 \rightarrow 98.6$$

The example shows that the intensional view of functions can be made precise by representing the function as an expression, and specifying the principles for simplifying (or: rewriting) such functional expressions. Rewriting functional expressions is a form of simplification where part of an expression is replaced by something simpler, until we arrive at an expression that cannot be simplified (or: reduced) any further. This rewriting corresponds to the computation of a function. For example, the function converting Celsius to Fahrenheit applied to the input 37 is the expression $37 \times \frac{9}{5} + 32$. This expression denotes the output, and at the same time it shows how to arrive at this output: First, $37 \times \frac{9}{5}$ is

rewritten to 66.6 , according to the rewriting rules for multiplication. The result of this simplification is $66.6 + 32$, which is then rewritten to 98.6 , in accordance with the rewriting rules for addition.

Functions can be composed, as follows. Let g be the function that converts from Kelvin to Celsius, and let f be the function that converts from Celsius to Fahrenheit. Then $f \cdot g$ is the function that converts from Kelvin to Fahrenheit, and that works as follows. First convert from Kelvin to Celsius, then take the result and convert this to Fahrenheit. It should be clear from this explanation that $f \cdot g$ is defined by

$$x \mapsto f(g(x)),$$

which corresponds to

$$x \mapsto (x - 273.15) \times \frac{9}{5} + 32$$

Exercise A.17 The successor function $s : \mathbb{N} \rightarrow \mathbb{N}$ on the natural numbers is given by $n \mapsto n + 1$. What is the composition of s with itself?

A special function which is simple yet very useful is the *characteristic function* of a set. The characteristic function of subset A of some universe (or: domain) U is a function that maps all members of A to the truth-value **True** and all elements of U that are not members of A to **False**. E.g. the function representing the property of being divisible by 3, on the domain of integers, would map the numbers

$$\dots, -9, -6, -3, 0, 3, 6, 9, \dots$$

to **True**, and all other integers to **False**. Characteristic functions characterize membership of a set. Since we specified relations as sets, this means we can represent every relation as a characteristic function.

Exercise A.18 \leq is a binary relation on the natural numbers. What is the corresponding characteristic function?

Exercise A.19 Let $f : A \rightarrow B$ be a function. Show that the relation $R \subseteq A^2$ given by $(x, y) \in R$ if and only if $f(x) = f(y)$ is an equivalence relation (reflexive, transitive and symmetric) on A .

A.6 Recursion and Induction

A recursive definition is a recipe for constructing objects from a finite number of ingredients in a finite number of ways. An example is the following recursive definition of natural numbers:

- 0 is a natural number.

- adding 1 to a natural number n gives a new natural number $n + 1$.
- nothing else is a natural number.

This recipe gives rise to an important method for proving things: proof by *mathematical induction*.

As an example, we prove the fact about natural numbers that the sum of the first n odd natural numbers equals n^2 . For example $1 + 3 = 2^2$, $1 + 3 + 5 + 7 = 4^2$, and so on. More formally and generally, we have for all natural numbers n :

$$\sum_{k=0}^{n-1} (2k + 1) = n^2.$$

Here is a proof of this fact by mathematical induction.

Basis. For $n = 0$, we have $\sum_{k=0}^0 (2k + 1) = 1 = 1^2$, so for this case the statement holds.

Induction step. We assume the statement holds for some particular natural number n and we show that it also holds for $n + 1$. So assume $\sum_{k=0}^{n-1} (2k + 1) = n^2$. This is the **induction hypothesis**. We have to show: $\sum_{k=0}^n (2k + 1) = (n + 1)^2$. Indeed,


$$\sum_{k=0}^n (2k + 1) = \sum_{k=0}^{n-1} (2k + 1) + 2n + 1.$$

Now use the induction hypothesis to see that this is equal to $n^2 + 2n + 1$, which in turn equals $(n + 1)^2$. Therefore we have:

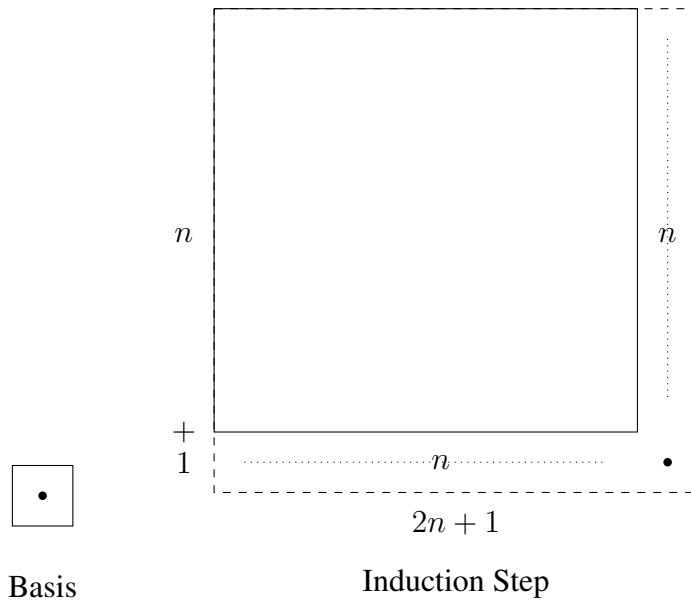
$$\sum_{k=0}^n (2k + 1) = \sum_{k=0}^{n-1} (2k + 1) + 2n + 1 \stackrel{ih}{=} n^2 + 2n + 1 = (n + 1)^2.$$

The equality $\stackrel{ih}{=}$ is the step where the induction hypothesis was used. We have checked two cases: the case 0 and the case $n + 1$. By the recursive definition of natural numbers, we have covered *all* cases, for these are the two possible shapes of natural numbers. So we have proved the statement for all natural numbers n .


The procedure of proof by mathematical induction does not help to *find* interesting patterns, but once such a pattern is found it is very helpful to *check* whether the pattern really holds. So how can one find a pattern like $\sum_{k=0}^{n-1} (2k + 1) = n^2$ in the first place? By imagining the following way to build up a square with side n :



Such a picture is what the ancient Greeks called a *gnomon* (“thing by which one knows”). The structure of the inductive proof can now be pictured as follows:



Exercise A.20 Consider the following *gnomon*:



What does this suggest for the sum of the first n even numbers? Give a form for $\sum_{k=0}^n 2k$, and prove with induction that your form is correct.