MAS114: Numbers and Groups
Semester 1 notes (up to week 4)

Dr James Cranch*

December 11, 2018

Introduction

Practical arrangements

- There are two lectures a week, at

  Monday 1pm (Dainton LT1)

  and

  Tuesday 1pm (Arts Tower LT4).

- Notes will be placed online on the course webpage several days before each lecture:

  `http://cranch.staff.shef.ac.uk/mas114/`

- The course webpage also has some practical advice, including on what to do if you miss a lecture.

- Each week you will have a *problem class* (on Thursday or Friday).

- You should find the *homework* online soon after your problem class. You should do these exercises (on the webpage) in your own time and hand them in at the next problem class. We will mark them, and return them to you at the beginning of the next problem class. If you'd like more feedback (on any of the solutions), please ask at the problem class.

---

*`J.D.Cranch@sheffield.ac.uk`

1

- The solutions to the *challenge problem* will make their way to me. I'll write up what you've managed.

- I offer *surgery hours* each week. During that time you can come to my office if you need extra help with the lecture material or exercises. These are at:

  Thursday 1pm–2pm

- At the end of this year there will be an exam, covering both semesters' material. This will count for

  80% of the module.

- There will be an *online test* each week, released immediately after Tuesday's lecture and due in at 2am on Monday (or, if you prefer, very late Sunday night). These will count for

  20% of the module.

- For this course you have *three hours* of contact time per week (two hours of lectures, one hour of problem classes). You're supposed to spend approximately as much time again (three more hours each week) in private study for this course, reading the notes and working on problems.

> If you do not do this, you *will not* be able to catch up in the run-up to the exams.

Things to do if you get stuck:
- Read the notes online.

- Ask your friends.

- Look in books.

- Search the web.

- Ask me.

Things not to do:
- Hope it will sort itself out.

- Leave it until the time of the exam.

What is mathematics?

It's hard to say what maths is. It is rather easier to say a few things about *what mathematics is not*.

<div style="border:1px solid black; padding:10px;">

Contrary to popular belief, mathematics is not the study of numbers.

</div>

Of course, the study of numbers is:

- *part* of mathematics,

- *very useful* in other parts of mathematics, and

- *particularly useful* in applications outside mathematics.

So we'll see lots of stuff about numbers in this course, and in other courses. But what else is there, if it's not all about numbers?

Here are a few pointers. These are just supposed to be a handful of examples rather than a big list of everything!

Ideas of space

Here are two knotted loops of string:
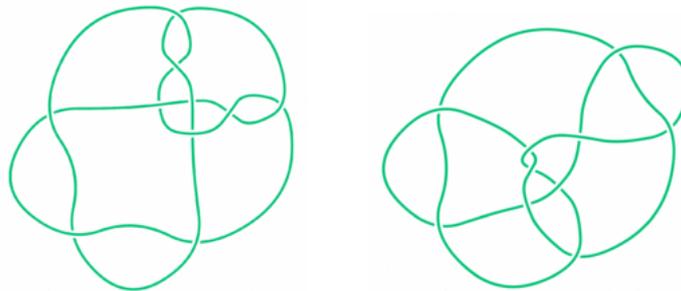
Are they the same? That is, if I had one, could I manipulate it so as to look like the other?

Knot theory — the study of problems just like this — is nowadays a thriving corner of mathematics. But knots are not numbers, and this question is not a question about numbers. Numbers might be useful in solving it, though!

This is just one of many examples of ideas of *space* in modern mathematics.

Ideas to do with space are nowadays of core importance in physics, just as numbers have. The world is made of space with interesting things in, after all.

Ideas of configuration

Is it possible to have a party of ten people, where everyone is a friend or a friend-of-a-friend of everyone else, and where everyone has exactly three friends present?
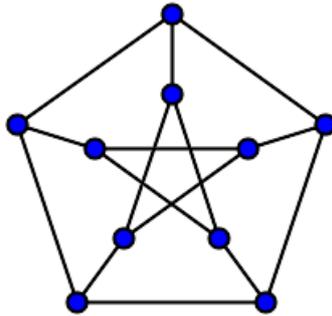
It's true that the numbers three and ten appear in this problem. But it's not really a problem about numbers: it's a problem about social networks and how they can be configured.

The study of networks (social and otherwise) has become known as *graph theory*. The subject of *combinatorics* encompasses this and many other kinds of configuration problem.

This has great application in computer science: after all, computer networks are examples of networks.

So what is mathematics?

It's hard to say! Perhaps you'll form an opinion yourselves over the next three or four years.

My working definition will be:

> Mathematics is the rigorous study of abstract systems.

Let's look at what that means.

Mathematics deals with simplifications, which are sometimes absurdly unrealistic. Mathematicians talk about a line of length 1, even if there's no ruler able to tell the difference between 0.999 and 1.001. They talk about perfect circles and lines with zero width.

Perhaps paradoxically, it's *because* of the unrealistic simplification that mathematics is able to describe the real world so well.

When mathematics models the behaviour of a spacerocket, treating the rocket as perfectly round and ignoring the dust and the small lumps of bird mess is the the right way to get an answer that's *good enough*.

One has to be very careful, but the abstraction of mathematics has been an amazing tool. For example, it *may be* true that nothing is perfectly round, but many things are so nearly round as to make their real shape irrelevant.

The purpose of choosing an abstraction is to give us something we can be completely certain about.

Every move must be fully justifiable (and fully justified, if you're trying to persuade people). There must be no risk of confusion or mistakes.

If we want to take liberties in our arguments then there's not much point in making an abstraction in the first place.

Rigour

Everyone knows that there are 64 squares on a chessboard. After all, chessboards are 8×8 grids, and $8 \times 8 = 64$.

That said, here's a square on a chessboard:
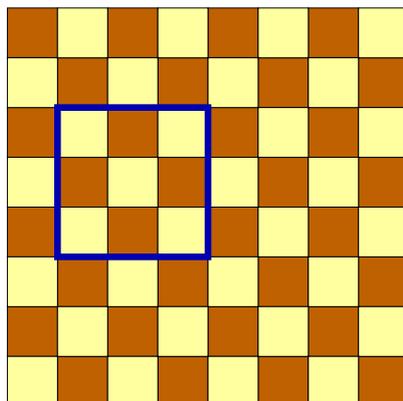
Obviously I've tricked you: I'm using the phrase "square on a chessboard" to mean two different things. But it's a good example of where mathematicians might be careful to be precise, to ensure that no mistakes are made. Things will get more complicated, and the need for care will be greater in future!

Here is a cautionary tale:

- In **1852**, Francis Guthrie asked:

  > Can every map be coloured with *only four colours* so that any two neighbouring regions have different colours?

Here's a vintage map of England and Wales (and, bizarrely, the Isle of Man) coloured in this way:
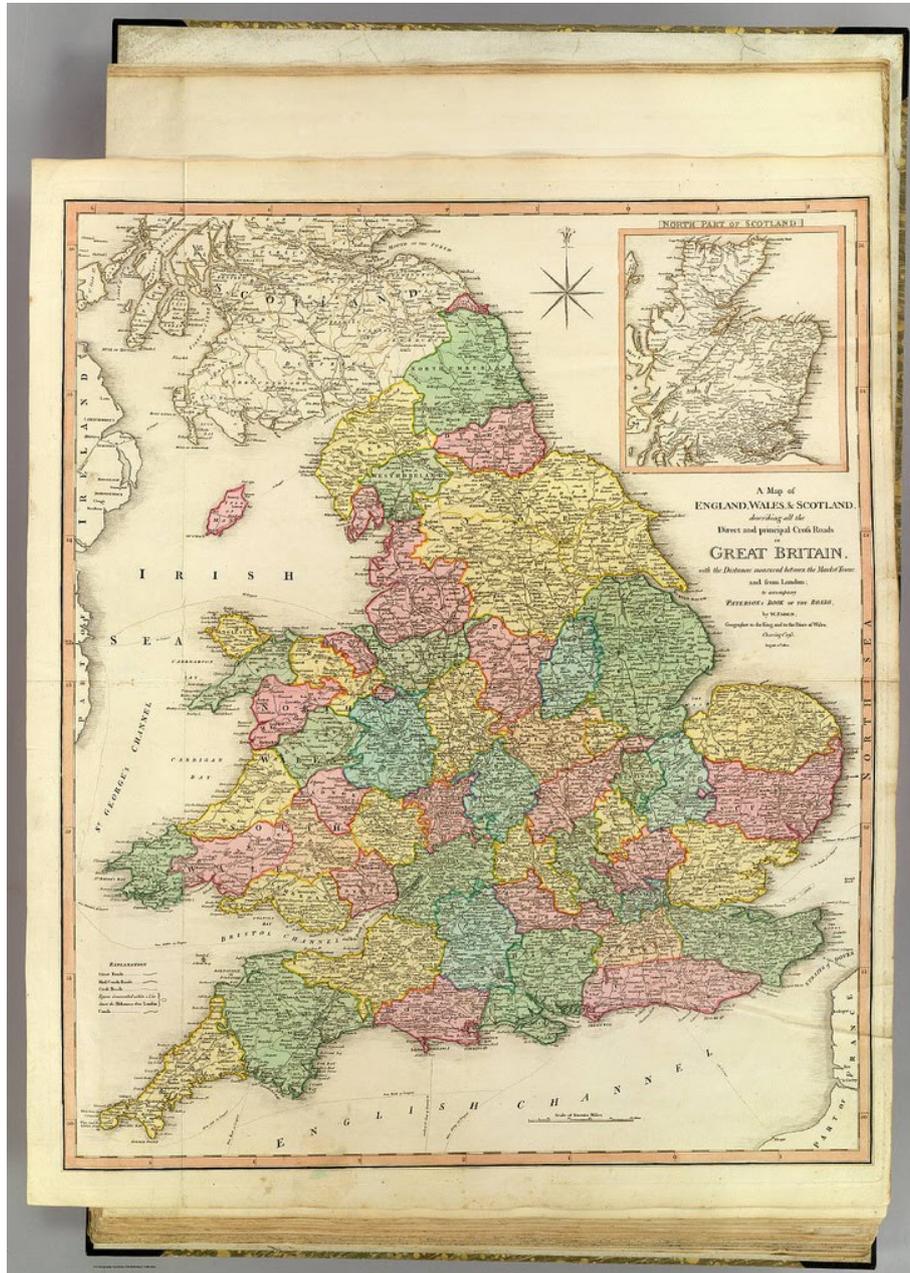
That's *not a proof*, but it is some evidence that it might be possible.

- In **1879**, Alfred Kempe offered a proof that the answer is *yes*.

- In **1880**, Peter Tait offered another, different, proof that the answer was *yes*. Mathematicians were satisfied, and stopped trying to prove it!

- In **1890**, Percy Heawood pointed out that Kempe's proof contained a big mistake.

- In **1891**, Julius Petersen pointed out that Tait's proof also contained a big mistake. Now, after twelve years spent believing the problem had been solved, and the answer was *yes*, mathematicians realised that in fact, they still had no idea.

- In **1976**, Kenneth Appel and Wolfgang Haken offered a new proof that the answer was indeed *yes*!

- As of **2018**, the argument of Appel and Haken has been checked many times, and is accepted as a complete solution.

The mistakes of Kempe and Tait are not particularly complicated. What caused this 12-year period of confusion was a *lack of sufficient rigour*.

In this course we will learn the basics of correct, logical argument. If you get the right final answer but your justification is incorrect or incomprehensible, you will deserve (and you will probably get) *very few marks*. Kempe and Tait had the right final answer too, but they had no way of knowing that.

At times this may seem like an unnecessary burden: especially when you feel that the right answer is "obvious". However, if you don't spend time in shallow water learning how to swim, you'll never be comfortably able to swim in deep water.

Sets and functions

Having spent the best part of an hour trying to persuade you that mathematics is not about numbers, most of this semester will now be about numbers.

But in order to study them properly, we'll need to start at the beginning, by talking about *sets*.

Sets of numbers

One problem with numbers is that there are several different sorts of them.

We're probably used to several sorts already:

natural numbers

integers

rational numbers

real numbers

complex numbers

We'll go into more detail later in the course.

We often need to say which we mean, in order to avoid confusion and error. For example, it's certainly possible that I might invite 3 friends over for dinner, but it's hard to invite -5 friends or $3/4$ friends or $\sqrt{2}$ friends over.

The natural numbers

The natural numbers are all the numbers you might find by counting.

The set of natural numbers is written $\mathbb{N}$ (that's just a letter N, written in a style called "blackboard bold"); in set notation, we might write

$$\mathbb{N} = \{0, 1, 2, 3, \ldots\}.$$

That "..." is often pronounced "dot-dot-dot", but it has the meaning of "and so on". When we use this symbol, we must be sure that the reader will be sure *how to go on* from here. Here I hope it really is clear: we go on with 4, 5, 6, adding one each time as we go, and we are to go on without end.

We'll see many more of those curly brackets later!

Actually, some mathematicians use the phrase "natural numbers" slightly differently, to denote the set

$$\{1, 2, 3, \ldots\}.$$

In other words, they leave out 0.

Our convention in this module will be that $0 \in \mathbb{N}$: that zero is a natural number.

If we're trying to work inside the natural numbers, we can add and multiply all we want, but subtraction and division are a pain: for example we can't do

$$3 - 5, \quad \text{or} \quad 2/7.$$

Working with a bigger system of numbers can cure this.

The integers

The *integers* are all the whole numbers, positive, negative and zero. The set of integers is denoted by $\mathbb{Z}$ (why Z? The German word for "number" is "Zahl"). So we might write

$$\mathbb{Z} = \{\ldots, -2, -1, 0, 1, 2, \ldots\}.$$

Here we have to go on in *both* directions without end.

Every natural number is an integer. This means that

$$\mathbb{N} \subset \mathbb{Z},$$

or, in words, "the set of naturals is contained in the set of integers".

We often use the handy words *non-negative*, meaning "not negative" (in other words, positive or zero) and *non-positive*, meaning "not positive" (in other words, negative or zero).

So the natural numbers are the same thing as the nonnegative integers.

If we're working in the integers, we can add, subtract and multiply all we want. Division is still a problem: for example, we can't do

$$-4/9.$$

The rational numbers

The *rational numbers* (sometimes just called the *rationals*), are the numbers that can be written as fractions $\frac{a}{b}$, where a and b are integers and $b \neq 0$.

Fractions can be written in many different ways: for example, we have

$$\frac{1}{2} = \frac{2}{4} = \frac{5}{10} = \frac{-3}{-6}.$$

In general, fractions $\frac{a}{b} = \frac{c}{d}$ are equal if

$$ad = bc.$$

We write $\mathbb{Q}$ for the set of rational numbers (Q stands for "quotient", which is a name for what you get when you do division).

Of course, any integer n can be regarded as a rational (we can take $\frac{n}{1}$), so

$$\mathbb{Z} \subset \mathbb{Q}.$$

If we're working in the rationals, we can add, subtract, multiply *and* divide all we want. (Well, we can't divide by zero, but who wants to divide by zero?).

There are still many things we might want to do but can't do in the rationals though: square roots, logarithms, trigonometry, and suchlike.

The real numbers

The *real numbers* $\mathbb{R}$ are perhaps the most general sort of numbers you'll have used by now (or perhaps not). They contain lots of the numbers you care about, for example:

$$\pi \in \mathbb{R}, \qquad \log 1729 \in \mathbb{R}, \qquad \sqrt{5} \in \mathbb{R}, \qquad \sin(37^\circ) \in \mathbb{R}.$$

One could define $\mathbb{R}$ as the set of all possible decimal expansions, but there are problems with this:

- It requires some adjustment, because

$$0.999999\cdots = 1.000000\cdots.$$

- Proving things about decimal expansions — even simple things like arithmetic — is a big pain.

- The idea of digits is, mathematically, an unnatural one. It is okay for the way we *write* mathematics to depend on the fact that we have ten fingers, but our *understanding* of fundamental mathematical constructions shouldn't depend on how many fingers we have.

Producing a good and useful definition of $\mathbb{R}$ is quite tricky, and there wasn't one until about 1870. We'll see one later in the course.

Sets in general

Now we have all these collections of numbers, it's good to have a language to discuss them with.

A *set* is a collection of objects. The objects in a set are often called its *elements*.

Given a set S, we write:

- $a \in S$ to mean "a is in S".

- $a \notin S$ to mean "a is not in S".

- $|S|$ to denote the *size* of S: the number of elements in it. (Of course, some sets are infinite, but this works well for finite ones, at least.)

Listing elements

If we have a small set, it might be practical to define it by listing its elements; we do so in curly brackets. Here's an example set:

$$T = \{\text{Tinky Winky}, \text{Dipsy}, \text{Laa-Laa}, \text{Po}\}.$$

Let's write some examples of facts about T using our notation:

$$\text{Po} \in T, \qquad \text{Noo-noo} \notin T, \qquad |T| = 4$$

Note that sets don't have any ordering on them. If we find it more convenient to list Teletubbies according to alphabetic order, we can write

$$T = \{\text{Dipsy}, \text{Laa-Laa}, \text{Po}, \text{Tinky Winky}\},$$

and in doing so we are defining exactly the same set T.

Also note that an element is either in a set, or not in it. So we could, if we wanted, define exactly the same set again by writing

$$T = \{\text{Dipsy}, \text{Laa-Laa}, \text{Po}, \text{Po}, \text{Po}, \text{Po}, \text{Po}, \text{Tinky Winky}, \text{Dipsy}\}.$$

However, there are few good reasons to write something like that.

Empty sets

The empty set, which could be written $\{\}$, is more commonly written $\varnothing$. It has size given by $|\varnothing| = 0$.

Note that $\varnothing$ is very different to $\{\varnothing\}$. The former, as I mentioned, has no elements; the latter has exactly one element.

That shouldn't confuse you. They're different for pretty much the same reason that "an empty bag" is not the same thing as "a bag which contains an empty bag and nothing else".

Containment

If A and B are sets, we write $A \subset B$ to mean "if x is a member of A then x is also a member of B". We say that A is a *subset* of B, or that A is *contained* in B.

The symbols "$\in$" and "$\subset$" are different, and using the wrong one tends to result in nonsense.

For example, we might write

$$\text{Mathematicians} \subset \text{People},$$

which says "all mathematicians are people". If we used the symbol "∈" instead, that would mean that "mathematicians is a person". It's not a mistake you'd make speaking English, and if you're using symbols you should aim to be no less precise.

Notice that, for every set A we have

$$A \subset A \qquad \text{and} \qquad \varnothing \subset A.$$

Set operations

Let A and B be sets. We define their *union* $A \cup B$ to contain exactly the things that are in one set or the other (or both):

$$A \cup B = \{x \mid x \in A \text{ or } x \in B\}.$$

That notation is called a *set comprehension*: the thing on the left of the vertical bar are the things we want to put in the set, and the things on the right of the vertical bar are the conditions under which we put them in. We'll use them a lot.

Similarly, we define the *intersection $A \cap B$* to contain exactly the things that are in both sets:

$$A \cap B = \{x \mid x \in A \text{ and } x \in B\}.$$

Lastly, we define the *difference $A \backslash B$* to be the things which are in A but not in B:

$$A \backslash B = \{x \mid x \in A \text{ and } x \notin B\}.$$

Equality

Two sets A and B are equal if they have the same members.

A straightforward way of proving this is often to show that $A \subset B$ and $B \subset A$. That is, in words, two sets A and B are equal if every element of A is an element of B and every element of B is an element of A.

Here's an example of this proof strategy:

Proposition 3.1. [1]
Let A, B and C be three sets. We have

$$A \cap (B \cup C) = (A \cap B) \cup (A \cap C).$$

---

[1]In this course, we'll be numbering results by lecture, so that Theorem 15.3 will be the third result in the 15th lecture.

Proof. Let's show firstly that $A \cap (B \cup C) \subset (A \cap B) \cup (A \cap C)$.

Suppose that $x \in A \cap (B \cup C)$; we must prove that $x \in (A \cap B) \cup (A \cup C)$.

Since $x \in A \cap (B \cup C)$, we have both $x \in A$ and $x \in B \cup C$ by the definition of intersection. But that means that $x \in B$ or $x \in C$, by the definition of union. In either case, the desired result holds:

- If $x \in B$, then since $x \in A$ also, then $x \in A \cap B$, and so $x \in (A \cap B) \cup (A \cap C)$ by the definition of intersections and unions respectively.

- If $x \in C$, then, similarly, since $x \in A$, we have $x \in A \cap C$, and hence $x \in (A \cap B) \cup (A \cap C)$.

So we've proved that containment.

Now let's prove the other containment: namely, that $(A \cap B) \cup (A \cap C) \subset A \cap (B \cup C)$. Let's suppose accordingly that $x \in (A \cap B) \cup (A \cap C)$; we must prove that $x \in A \cap (B \cup C)$.

Since $x \in (A \cap B) \cup (A \cap C)$, we have either $x \in A \cap B$ or $x \in A \cap C$ by the definition of union. In either case, we get what we want:

- If $x \in A \cap B$, then $x \in A$ and $x \in B$ by the definition of intersection. From the latter, we get that $x \in B \cup C$, by the definition of the union, and hence $x \in A \cap (B \cup C)$ by the definition of intersection.

- If $x \in A \cap C$, then, as before, $x \in A$ but now $x \in C$. However, it's still true that $x \in B \cup C$, and so $x \in A \cap (B \cup C)$. $\square$

That was the first example of a formal proof in this course. You'll have to write many proofs like this yourself, in assessed homework and in the exam. Though we'll discuss it in depth later, it may be worth observing the style from the beginning. One big mistake that many beginner mathematicians make is *not using words to explain the flow of the argument*.

A warning

What we are practising here is called *naïve set theory*. What's so naïve about it?

The Welsh mathematician Bertrand Russell realised in 1901 that there are serious problems with being allowed to form sets carelessly:

Paradox 3.2. *Suppose there is a set S of all sets which are not elements of themselves:*

$$S = \{A \mid A \notin A\}.$$

This creates a contradiction.

Proof.   Is S a member of itself? If $S \in S$, then by the definition of S, we have $S \notin S$. On the other hand, if $S \notin S$, then again by the definition of S we have $S \in S$.

$\square$

As a result of this paradox, modern set theorists impose strict rules on what sets can be formed, with the aim of banning this particular beast and everything like it.

However, you probably won't need to worry about this, unless you take a course in set theory later in your mathematical careers.

Functions

A function is to be thought of as a machine that takes an element of one set and gives you an element of another. Here's a formal definition:

Definition 3.3. Given sets A and B, a *function* (sometimes called a *map*) $f : A \to B$ gives for each element $a \in A$ a unique element $f(a) \in B$.

Examples include:

- the function $f : \mathbb{Z} \to \mathbb{Z}$ defined by $f(n) = n^2 - 7$.

- the function $g : \mathbb{Q} \to \{4, 6\}$ defined by

$$g(x) = \begin{cases} 4, & \text{if } x = 3/7 \text{ or } x = -14/17; \\ 6, & \text{otherwise.} \end{cases}$$

The set A is called the *domain* of f, and B is called the *codomain* of f. We call $f(a)$ the *value* of f at a, or the *image* of a under f.

Consider the "age in years" function from the set of people in this room to the natural numbers.

The *domain* of this function is the set of values you're permitted to apply it to. This is the set of people in this room, because I said so.

The *codomain* of this function is the set of values it is *permitted* to take. This is the set $\mathbb{N}$ of natural numbers, because I said so.

Some people like to talk about the *image* of this function, being the set of values it actually takes in practice. This might (perhaps) be the set

$$\{18, 19, 20, 35\}\,.$$

The *range* is not a phrase that's used consistently:

- some people use it to mean the codomain;

- some people use it to mean the image;

- some (confused) people, who don't know the difference, use it inconsistently to mean both.

When you're trying to work out whether something's a function, there are three bits of the definition where things can go wrong:

"each $a \in A$" A function must be defined for every single element of the domain. Why does $\alpha(x) = 1/x$ not define a function $\alpha : \mathbb{Q} \to \mathbb{Q}$?

α is not defined at zero

"unique element" A function must have only one value at any given element of the domain. If we set $\beta(n)$ to be the real number x whose square is n, why does that not define a function $\beta : \mathbb{N} \to \mathbb{R}$?

$\beta(3)$ could be $+\sqrt{3}$ or $-\sqrt{3}$.

"$f(a) \in B$" A function must return values within its codomain. Why does $\gamma(n) = n - 7$ not define a function $\gamma : \mathbb{N} \to \mathbb{N}$?

$\gamma(4) = -3$ does not lie inside $\mathbb{N}$.

Two functions are equal if:

- they have the same domain and codomain, as $f, g : A \to B$; and

- their values are equal, for every point in the domain: in other words, for all $a \in A$, we have $f(a) = g(a)$.

Given two functions $f : A \to B$ and $g : B \to C$, we can define their *composite* $g \circ f : A \to C$ by the rule:

$$g \circ f(x) = g(f(x)).$$

Functions don't have to be described by formulae (as they are in the examples, and non-examples, above).

For example, if the domain is finite we can define them pictorially. Accordingly, here is a function from the set T of Teletubbies, as considered earlier, to the set of colours:
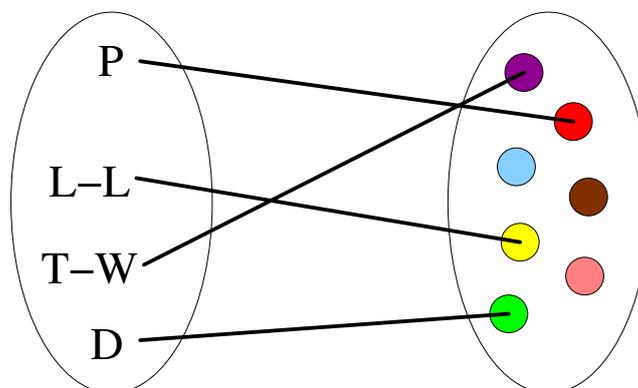
Now we're well-equipped to describe functions, we can start describing their properties.

Here are some useful words.

Definition 4.4. A function $f : A \to B$ is said to be *injective* if, for any two elements $a_1, a_2 \in A$ with $a_1 \neq a_2$, then $f(a_1) \neq f(a_2)$. I think of this as saying that "nothing is hit twice", or equivalently that "no two elements of the domain have the same image".

Definition 4.5. A function $f : A \to B$ is said to be *surjective* if, for every element $b \in B$, there is some element $a \in A$ with $f(a) = b$. I think of this as saying that "every element of the codomain is hit at least once".

Definition 4.6. A function $f : A \to B$ is said to be *bijective* if it is both injective and surjective. I think of this as saying that "every element of the codomain is hit exactly once".

For example, let's consider our function assigning colours to Teletubbies.
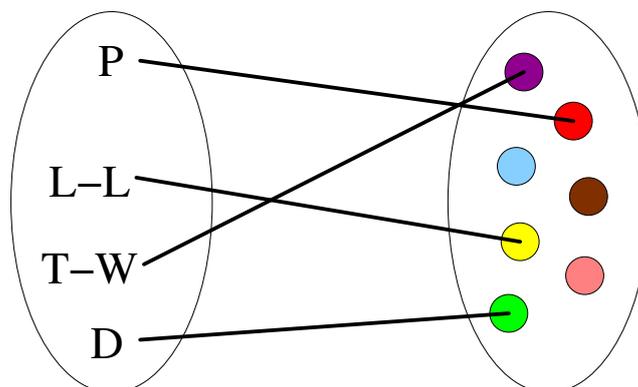
It is injective, because each one of the Teletubbies has a different colour.

However, it is not surjective, because there are no pink Teletubbies in all of Teletubbyland. Hence it is also not bijective.

Note that these properties (injective, surjective, bijective) don't just depend on the rule that defines it: they depend on the domain and codomain.

For example, consider the rule $f(n) = n^2$. Is this injective, considered as a function $f : \mathbb{N} \to \mathbb{N}$?

> Yes! Every natural number has a different square

Is it injective as a function $f : \mathbb{Z} \to \mathbb{Z}$?

> No; $f(3) = f(-3)$.

Similarly, consider the rule $g(n) = n + 100$. Is this surjective, considered as a function $g : \mathbb{N} \to \mathbb{N}$?

> No; there is no $a \in \mathbb{N}$ with $g(a) = 50$.

Is it surjective as a function $g : \mathbb{Z} \to \mathbb{Z}$?

> Yes it is! For any n we have $g(n - 100) = n$.

Remark 4.7. Note that that function also has an *inverse*, a function which undoes g. Namely, we can take the inverse $g^{-1} : \mathbb{Z} \to \mathbb{Z}$ to be $g^{-1}(n) = n - 100$.

You'll see a lot more about inverses next semester.

By the way, I'd like you to try to remove phrases like "one-to-one" from your vocabulary: it's never clear whether "f is one-to-one" means "f is injective", "f is bijective", or simply "f is a function", and many people who use it treat this as an excuse for not explaining which they mean.

Logic

Now we'll briefly move on to another building block of mathematics: logic. Logic studies the properties of statements, which can be either *true* or *false*.

Implication

Much of logic involves deductive reasoning. Here's the definition that encapsulates that:

Definition 4.8. Let A and B be statements. We say that A *implies B*, written $A \Rightarrow B$, to mean "if A is true, then B also has to be true".

There are many common ways of saying the same thing, used by mathematicians. These include:

- *A* implies *B*;

- $A \Rightarrow B$;

- If *A*, then *B*;

- *A* only if *B*;

- *A* is sufficient for *B*;

- *B* is necessary for *A*.

Remark 4.9. Notice that $A \Rightarrow B$ cannot be used to mean "*A* is true, and so *B* is also true".

For example, let *A* be the statement "I visited Cardiff last week", and *B* be the statement "I've been to Wales this month". We can all agree that $A \Rightarrow B$ is true, which says

> "If I visited Cardiff last week, then I must have been to Wales in the last month."

However, as it happens, neither of these is true: in fact, I haven't been in Wales for a while longer than that.

As such, it is quite different to saying "*A* is true, and therefore *B* is also true". Beginning students often get these confused.

Remark 5.10. The implication $A \Rightarrow B$ only says something interesting about *B* if *A* happens to be true! If *A* is false, then we have $A \Rightarrow B$ no matter whether *B* is true or false.

For example, it's correct to say

> "If $2 + 2 = 337$, then this course is lectured by Dr Cranch".

or indeed

> "If $2 + 2 = 337$, then this course is lectured by Robert de Niro".

This may be a surprise if you're basing your intuition on ordinary English, where people use the words "if... then" in several different ways, sometimes slightly ambiguously.

Remark 5.11. Another important point is that if $A \Rightarrow B$ does not say that *B* is true *only* if *A* is true (that would be $B \Rightarrow A$, in fact).

So it's correct to say

18

If it rains next Wednesday, then $2 + 2 = 4$.

In fact, it's true that $2 + 2 = 4$ no matter whether it rains next Wednesday, but that's not a problem. Whether or not it's *helpful* to say that is another question!

We can sum up the comments above by giving a *truth table* for implication. We write 0 for false and 1 for true.

First off, note that if Q is true, then $P \Rightarrow Q$ is definitely true no matter whether P is true.

Also, if P is false, then $P \Rightarrow Q$ is true no matter whether Q is true.

In fact, it's only if P is true and Q is false that $P \Rightarrow Q$ is false.

P	Q	$P \Rightarrow Q$
0	0	1
0	1	1
1	0	0
1	1	1

We've seen in the above, that $A \Rightarrow B$ and $B \Rightarrow A$ are different statements. It's helpful to have a word relating them:

Definition 5.12. Consider a statement of the form $A \Rightarrow B$. Then the *converse* of that statement is the statement $B \Rightarrow A$.

The truth of an implication has very little to do with the truth of its converse, as we'll see.

Example 1. Let P be the statement "X lives in England" and Q be the statement "X lives in Sheffield".

Is the statement $P \Rightarrow Q$ always true?

No. (X could live in Newcastle.)

Is its converse always true?

Yes, it is.

Example 2. Let C be the statement "Y is English", and let D be the statement "Y lives in Sheffield".

Is the statement $C \Rightarrow D$ always true?

No. (Y could live in New York.)

Is the converse always true?

No. (Y could be German.)

Equivalence

Equivalence is the relationship between two statements of being both true, or both false.

Definition 5.13. Let P and Q be statements. We write $P \Leftrightarrow Q$, pronounced "P is equivalent to Q" for the statement that P is true if and only if Q is true.

Sometimes people shorten "if and only if" to "iff".

Negation

Negation is a type of opposite:

Definition 5.14. Let P be a statement. The *negation* of P, written $\neg P$ and often pronounced "not P", is the statement "P is false".

Remark 5.15. We must be careful in thinking of negation as an opposite. For example, the negation of "Richard is happy" is "Richard is not happy".

Most people would say that the "opposite" is "Richard is sad". That's not quite the same thing!

Similarly, the negation of "Alice is in front of Bob" is not "Alice is behind Bob", but:

"Alice is not in front of Bob".

Note that double negation doesn't do anything: the statement $\neg(\neg P)$ is equivalent to P. Since statements are either true or false, if it's not "not true", it's true.

The negation allows us to make sense of something very important about implication:

Definition 5.16. Consider a statement of the form $P \Rightarrow Q$. Its *contrapositive* is the statement $(\neg Q) \Rightarrow (\neg P)$.

The main useful fact about the contrapositive is that it's equivalent to the original implication. This is something very familiar from everyday life. If my friend Mel says

"If I can come visit you this evening, then I'll text you at lunchtime,"

then I might rephrase it to myself as:

"I don't get a text at lunchtime, then Mel won't visit this evening."

But here's a formal statement and proof, anyway:

Proposition 5.17. *Let P and Q be statements. The statement* $P \Rightarrow Q$ *is equivalent to its contrapositive* $(\neg Q) \Rightarrow (\neg P)$.

Proof. I'll prove this using a truth table, showing what happens in all possibilities.

P	Q	$P \Rightarrow Q$
0	0	1
0	1	1
1	0	0
1	1	1

We now provide a table for $(\neg Q) \Rightarrow (\neg P)$ (using the table above to fill in the implications):

P	Q	$\neg Q$	$\neg P$	$(\neg Q) \Rightarrow (\neg P)$
0	0	1	1	1
0	1	0	1	1
1	0	1	0	0
1	1	0	0	1

The reader can see from this that $(\neg Q) \Rightarrow (\neg P)$ is true exactly when $P \Rightarrow Q$ is, and this proves that they're equivalent.

$\square$

"And" and "Or"

There are other ways to combine statements, other than implication:

Definition 5.18. Let P and Q be statements. The statement $P \wedge Q$, pronounced "P and Q", is the statement that both P and Q are true.

The statement $P \vee Q$, pronounced "P or Q", is the statement that at least one of P or Q (and possibly both) is true.

We can make truth tables for both of them:

P	Q	$P \wedge Q$
0	0	0
0	1	0
1	0	0
1	1	1

P	Q	$P \vee Q$
0	0	0
0	1	1
1	0	1
1	1	1

Remark 5.19. While you're all used to these words from everyday life, there can be vagueness about how "or" is used in English.

For example,

you can have your pie with chips or with mashed potatoes

is probably intended to mean "but not both". In mathematical argument when we use "or" and mean "but not both", we have to say so explicitly.

It is sometimes worth knowing that implication can be defined in terms of "or":

Proposition 5.20. *The statement $P \Rightarrow Q$ is equivalent to $(\neg P) \vee Q$.*

Proof. The only way that the first statement can be false is if P is true and Q is false. But that's also the only way that the second statement can be false, so they're equivalent. □

Note that that shows another style of proof of logical statements: by analysis rather than the "case bash" used in truth tables.

Quantifiers

Lastly, we need to discuss how to make statements about *general* situations and *particular* examples. The phrases we use again and again are "for all" and "there exists": these are called *quantifiers*.

The following symbols are in common use:

$$\forall \qquad \text{for "for all"};$$
$$\exists \qquad \text{for "there exists"};$$
$$\text{s.t.} \qquad \text{for "such that"}.$$

So, for example,

$$\forall n \in \mathbb{Z}, \quad n^2 - 1 = (n+1)(n-1)$$

is to be read as

"For all integers n, we have $n^2 - 1 = (n+1)(n-1)$."

And

$$\exists x \in \mathbb{R} \quad \text{s.t.} \quad x^2 - 3x - 12 = 0$$

is to be read as

22

"There exists a real number x such that $x^2 - 3x - 12 = 0$."

It's important that you get used to this notation. This is not because there is anything amazing about it, but because mathematics involves lots of general rules and particular examples: much of the mathematics you do for the next few years will require you to be able to deal with these things.

One thing you'll have to get used to is situations with two or three quantifiers. These happen very frequently: "in general, there is always a particular example of such-and-such", or "there is a particular amazing example which has the general property of such-and-such".

For example, the statement

$$\forall n \in \mathbb{N}, \quad \exists x \in \mathbb{R} \quad \text{s.t.} \quad x^2 = n$$

says that every natural number n has a square root x.

Remark 6.21. The order of quantifiers is very important. If we swap over the two quantifiers in the last example, we get

$$\exists x \in \mathbb{R} \quad \text{s.t.} \quad \forall n \in \mathbb{N}, \quad x^2 = n.$$

This says that there's a particular number x which has the property that x is the square root of *every natural number*. And that's nonsense.

Another thing that mathematicians have to do every day is understanding how negation interacts with quantifiers.

The negation of "all Teletubbies are red" is "not all Teletubbies are red", which is equivalent to "there exists a Teletubby which is not red".

Similarly, the negation of "there exists a dolphin who likes Beethoven" is "there does not exist a dolphin who likes Beethoven", and that's equivalent to "all dolphins do not like Beethoven".

In symbols,

$$\neg \left(\forall x \in X, \quad P(x) \right) \qquad \text{is equivalent to} \qquad \exists x \in X \quad \text{s.t.} \quad \neg P(x).$$
$$\neg \left(\exists x \in X \quad \text{s.t.} \quad P(x) \right) \qquad \text{is equivalent to} \qquad \forall x \in X, \quad \neg P(x).$$

Perhaps you may want to remember that "negation swaps $\forall$ and $\exists$." But being able to *do it correctly by remembering what's going on* is much more important than remembering a slogan. After a while it should come to seem natural.

Suppose I am wondering whether all fish are slippery. If it's true, I need to find some general reason why *every single* fish is slippery. If it's false, I only need to find *one single* fish which isn't slippery, and then I've proved it.

In general, if you have a general statement and you don't know if whether it's true or false, then it could either be:

- *true*, in which case you need to prove it in general (that's a statement with a "∀" in);

- *false*, in which case you need to find a counterexample (that's a statement with a "∃" in).

Induction

The basics of induction

The most obvious interesting thing about the natural numbers is that it's natural to start listing them, one after the other:

$$\mathbb{N} = \{0, 1, 2, 3, 4, \ldots\}.$$

This, of course, is how *counting* works.

It turns out that this way of thinking about the integers gives us a very powerful tool for proving things one integer at a time: the *principle of mathematical induction*, usually known to mathematicians simply as *induction*.

Informally, I like to think of the following example:

> *If* I can reach the bottom (rung number zero?) of a ladder, *and* if I'm on any rung I can reach the next rung up, *then* I can reach any rung on the ladder.

Why, for example, can I reach the fourth rung? One can imagine a detailed proof of this, as follows:

- I can reach rung zero;

- Because I can reach rung zero, I can reach rung one;

- Because I can reach rung one, I can reach rung two;

- Because I can reach rung two, I can reach rung three;

- Because I can reach rung three, I can reach rung four.

The connection with counting is obvious: our proof visibly counts up to four.

Writing that out was okay, but you are probably glad I didn't write out a proof that we could reach the hundred and seventy-eighth rung. I suppose that we could do so, writing "and so on" at some point: but that's a little vague (what about situations where it isn't obvious what "and so on" means)?

It's helpful to have a way which isn't vague.

So here's a formal version:

Definition 6.22 (Induction). Let $P(n)$ be a statement that depends on a natural number n. Then, if

  (i) the statement $P(0)$ is true, and

  (ii) for all $k \in \mathbb{N}$, if $P(k)$ is true, then $P(k+1)$ is true,

then the statement $P(n)$ is true for all $n \in \mathbb{N}$.

  Here are some useful words:

Remark 6.23. We call part (i) the *base case*, and part (ii) the *induction step*. These words agree quite well with our mental picture of a ladder!
  When we are trying to prove the induction step $P(k) \Rightarrow P(k+1)$ we refer to $P(k)$ as the *induction hypothesis*.

Example of induction

We'll prove many things by induction in this course, but this is one:

Proposition 6.24. *For any natural number n, we have the following formula for the sum of the first n positive integers:*

$$1 + 2 + \cdots + (n-1) + n = \frac{n(n+1)}{2}.$$

  Let $P(n)$ be the statement above for some particular n.
  So $P(3)$ is the statement that says $1 + 2 + 3 = 3 \times 4/2$, and $P(10)$ is the statement that

$$1 + 2 + 3 + 4 + 5 + 6 + 7 + 8 + 9 + 10 = (10 \times 11)/2.$$

Notice that $P(n)$ is *not a number*, it's a *statement*.

Proof. We will prove $P(n)$, which says that

$$1 + 2 + \cdots + (n-1) + n = \frac{n(n+1)}{2}.$$

for all n by induction.

> For our base case, $P(0)$ says that the sum of *no integers at all* is $0 \times 1/2$, which is true, as the sum of no integers is zero.
> Now we will do our induction step, proving $P(k) \Rightarrow P(k+1)$ for all k. Suppose $P(k)$ is true: we need to show that $P(k+1)$ is true.

The statement $P(k)$ tells us that

$$1 + 2 + \cdots + (k - 1) + k = \frac{k(k + 1)}{2}.$$

We need to prove $P(k + 1)$, which would say that

$$1 + 2 + \cdots + (k - 1) + k + (k + 1) = \frac{(k + 1)(k + 2)}{2}.$$

Now note that

$$
\begin{aligned}
&1 + 2 + \cdots + (k - 1) + k + (k + 1) \\
=\ & (1 + 2 + \cdots + (k - 1) + k) + (k + 1) \\
=\ & \tfrac{k(k+1)}{2} + (k + 1) \quad \text{(by the induction hypothesis)} \\
=\ & \tfrac{k(k+1)+2(k+1)}{2} \\
=\ & \tfrac{(k+1)(k+2)}{2}.
\end{aligned}
$$

This is exactly the statement $P(k + 1)$, which is what we needed for the induction step, and that completes the proof.

$\square$

Remark 6.25. You may know other ways of proving that. (I can think of a few.) But I hope you're impressed with this as a strong potential method for proving identities.

Nonexamples of induction

Let's now try proving some *completely false* statements using induction. The plan is (of course) not to succeed, but to understand where we need to be careful.

Example 3. We'll try proving using induction that for all n, we have

$$n = n + 83.$$

Clearly this statement is complete and utter rubbish.

> If you believe that induction is a reliable method of proof (and I do, and I hope you do too), then it had better be the case that we're not using induction correctly.

Anyway, here's an induction "proof". Suppose that $k = k + 83$ for some k. We'll prove that $(k + 1) = (k + 1) + 83$. But we have

$$
\begin{aligned}
k + 1 &= (k + 83) + 1 &&\text{(by assumption)} \\
&= (k + 1) + 83 &&\text{(by rearrangement)}.
\end{aligned}
$$

This completes the proof.

What's the problem with the argument above?

There's no base case.

If you don't have a base case, such as $P(0)$, then it's of no use to prove that $P(k) \Rightarrow P(k + 1)$ for all k. It's no use to be able to climb a ladder if the bottom of the ladder is unreachable.

Here's another, more subtle example:

Example 4. We'll try proving using induction that all horses are the same colour.

Again, we find ourselves hoping very strongly that there's a mistake in the use of induction in what follows. I'll write it out and we can see if we can spot it.

In order to do this, we'll let $P(n)$ be the statement "Given any n horses, all of them have the same colour". We'll prove $P(n)$ for all n by induction: that will give us what we want, because we can take n to be the number of horses in the world.

We'll take $P(1)$ as the base case of the induction. This is the statement "Given any one horse, all of them have the same colour": this is obviously true.

Now we'll prove the induction step. We will assume that $P(k)$ is true ("given any k horses, all of them have the same colour"): our job is to prove that $P(k + 1)$ is true ("given any $(k + 1)$ horses, all of them have the same colour").

So suppose we have $(k + 1)$ horses. Name two of them Alice and Zebedee.
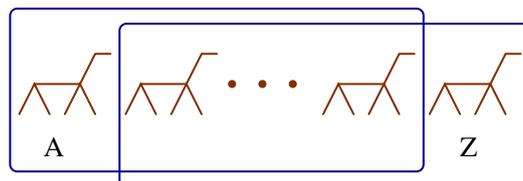
Excluding Alice, there are k horses, which all have the same colour, by the induction hypothesis. So all the horses except Alice have the same colour as Zebedee.

Also, excluding Zebedee, there are k horses, which all have the same colour, again by the induction hypothesis. So all the horses except Zebedee have the same colour as Alice.

Hence all the horses except Alice and Zebedee have the same colour as both Alice and Zebedee, which says that all the horses have the same colour. That ends the proof.

What's wrong with this?

> The particular case $P(1) \Rightarrow P(2)$ doesn't work.

I find this surprisingly subtle.

In fact, it's a parody of a *valid* style of argument. If it is the case that *any two things are the same*, then we could prove using exactly this method that they're *all the same*. In fact, this is something you already know, since "all are alike" and "no two differ" are synonymous phrases.

Variants

There are other techniques which use the same *idea* of induction, but not quite the same formal principle as I've written out above.

We can start with a base case which isn't $P(0)$. For example, if

(i) $P(15)$ is true, and

(ii) $P(k)$ implies $P(k+1)$ for all $k \geqslant 15$,

then $P(n)$ is true for all $n \geqslant 15$.

Perhaps you want to think of that as saying "if have a door which leads to the fifteenth rung of a ladder, and you know how to climb ladders, then you can get to every rung above the fifteenth".

Actually, this is really just ordinary induction in disguise.

Indeed, if we define $Q(n)$ to be $P(n+15)$, then proving $Q(n)$ for all $n \in \mathbb{N}$ by induction is the same as proving $P(n)$ for all integers $n \geqslant 15$.

Perhaps you want to think of that as "ignoring all the bits of the ladder below the fifteenth rung, imagining the ladder starts outside your door, and starting counting rungs from there". Or perhaps you're bored of the ladder analogy now.

Actually, you should have been prepared for this variant: my induction proof that "all horses have the same colour" started with 1, not 0. (Okay, that proof was

wrong. But there was nothing wrong with *that bit* of the proof: there's nothing wrong with induction starting from 1. It was something else that was wrong).

Another common variant is known as *strong induction*. This allows you to assume all previous cases in your induction step, rather than just the last one.

It looks like the following:

Definition 7.26 (Strong induction). Let $P(n)$ be a statement that depends on a natural number n. Then if

(i) $P(0)$ is true, and

(ii) for all k, if $P(0), P(1), \ldots P(k)$ are all true, then $P(k + 1)$ is also true,

then $P(n)$ is true for all $n \in \mathbb{N}$.

Again, this can be viewed as a cleverly disguised version of ordinary induction: if we define the statement $Q(n)$ as follows:

$$Q(n) = P(0) \wedge P(1) \wedge \cdots \wedge P(n - 1) \wedge P(n)$$
$$= \text{"all the statements } P(0), \ldots, P(n) \text{ are true",}$$

then proving $Q(n)$ by ordinary induction works out to be effectively the same thing as proving $P(n)$ by strong induction.

Indeed, the base case $Q(0)$ is the same thing as the base case $P(0)$. The induction step $Q(k) \Rightarrow Q(k + 1)$ looks like

$$P(0) \wedge \cdots \wedge P(k) \Rightarrow P(0) \wedge \cdots \wedge P(k) \wedge P(k + 1).$$

In order to prove this, we assume $P(0), \ldots, P(k)$ are all true and have to prove that $P(0), \ldots, P(k + 1)$ are all true. But then all of these except the last are assumptions: what is left is to prove $P(k + 1)$ assuming $P(0), \ldots, P(k)$, and that's exactly the induction step of a strong induction.

Here's an example of strong induction in practice. First we'll need a definition:

Definition 7.27. The *Fibonacci numbers* $F_0, F_1, \ldots$ are defined by taking $F_0 = 0$ and $F_1 = 1$, and then for $n \geqslant 2$ by

$$F_n = F_{n-1} + F_{n-2}.$$

Now for the result in question:

Proposition 7.28. *For all $n \in \mathbb{N}$ we have $F_n < 2^n$.*

Proof. Let $P(n)$ be the statement $F_n < 2^n$. We know that $P(0)$ is true, since

$$F_0 = 0 < 1 = 2^0.$$

We also know that $P(1)$ is true, since

$$F_1 = 1 < 2 = 2^1.$$

Now we'll show that for all $k \geqslant 1$ we have that if $P(0), \ldots, P(k)$ are all true, then $P(k + 1)$ is true. So suppose that $P(0), \ldots, P(k)$ are all true.

We now have that

$$
\begin{aligned}
F_{k+1} &= F_k + F_{k-1} \\
&< 2^k + 2^{k-1} \qquad \text{(using } P(k) \text{ and } P(k-1)\text{)} \\
&< 2^k + 2^k \\
&= 2^{k+1},
\end{aligned}
$$

which is exactly $P(k + 1)$. This completes the induction step, and so finishes the proof. $\square$

That strong induction argument really had *two base cases* before the induction step.

So we proved $P(0)$, and we proved $P(0) \Rightarrow P(1)$ by proving $P(1)$, and then we proved $P(0) \wedge \cdots \wedge P(k) \Rightarrow P(k + 1)$ by proving $P(k - 1) \wedge P(k) \Rightarrow P(k + 1)$.

I like to think that the proof was arranged according to the shape of the definition of the Fibonacci numbers: that definition has two base cases $F_0 = 0$ and $F_1 = 1$, and a step $F_{n+2} = F_{n+1} + F_n$. This is not a rare coincidence. Lecture 8

Why induction works

In this section we make a few comments on why induction works. They may be helpful in thinking about when you can and when you can't generalise induction to other settings.

Let's introduce a definition:

Definition 8.29. A set (of numbers) is *well-ordered* if every nonempty subset has a least element.

For now, our main use of that is to say this:

Definition 8.30. The *well-ordering principle for* $\mathbb{N}$ says that $\mathbb{N}$ is well-ordered.

This is a very special property of $\mathbb{N}$. The integers $\mathbb{Z}$, for example, are not well-ordered. Indeed, the subset $\mathbb{Z} \subset \mathbb{Z}$ of all integers does not have a smallest element: there is no smallest integer.

The main interest is this:

Theorem 8.31. *The well-ordering principle for $\mathbb{N}$ and the principle of strong induction(Definition 7.26) are equivalent.*

Proof.   We'll show first that we can derive the principle of strong induction from the well-ordering principle.

So suppose we had a statement $P(n)$ for each $n \in \mathbb{N}$, and we had a base case (that $P(0)$ was true) and an induction step (that, for all $k \in \mathbb{N}$ if we have $P(i)$ for all $i < k$, we also have $P(k)$). We need to show that $P(n)$ is true for all n.

We might argue as follows. Let A be the set of natural numbers n for which $P(n)$ does not hold:

$$A = \{n \in \mathbb{N} \mid \neg P(n)\}.$$

So A is the set of "counterexamples".

If A has any elements at all, it has a smallest element a. But a can't be 0, because we have $P(0)$. But a can't be bigger than 0 either: because a is minimal, we have $P(i)$ for all $i < a$. Hence we have $P(a)$ also, by the induction step. But that's a contradiction: we assumed that $\neg P(a)$.

Hence A doesn't have any elements, which is the same as saying that $P(n)$ holds for all n.

Now we'll show the other half of the equivalence: that we can derive the well-ordering principle from the principle of strong induction.

Let $Q(n)$ be the statement, "any subset of $\mathbb{N}$ which contains n has a smallest element". We'll prove $Q(n)$ for all n by strong induction.

Firstly, for a base case, we must prove $Q(0)$ ("any subset of $\mathbb{N}$ which contains 0 has a smallest element"). This is clearly true, as 0 is the smallest natural number of all, so any such subset has 0 as its smallest element.

Now we must prove the induction step: we assume for some k that $Q(i)$ is true for all $i < k$, and we prove that $Q(k)$ is true. Consider a subset $S \subset \mathbb{N}$ which contains k. If it contains some element $i < k$, then by $Q(i)$ it has a least element. If, however, it contains no element $i < k$, then k is its least element: so in particular, it has a least element. This proves $Q(k)$.

Hence we have $Q(n)$ for all n by strong induction. So if S is a nonempty subset of $\mathbb{N}$, it has at least one element: call it n. But then by $Q(n)$, the set S has a least element: this proves the well-ordering principle.

$\square$

Part of the reason this is such good news is that there are other well-ordered sets, besides the natural numbers. Whenever you find a well-ordering, you get a notion of induction for free.

For example, consider the set of pairs (m, n) of naturals, where we say that $(m, n) < (m', n')$ if $m < m'$, or if $m = m'$ and $n < n'$. (This is called the *lexicographic* ordering, because it's inspired by the way that words in dictionaries are ordered).

It is not too hard to prove that that is well-ordered: any set of pairs of natural numbers has a "least" element with respect to this ordering. Hence we can do (strong) induction on pairs of naturals!

Advice on writing proofs

In this section I simply give a few thoughts on how to write a good proof.

Think of it as a literary form

Writing proofs which are easy to understand, enlightening, or fun to read is very difficult.

Like with other forms of literature, the best way to learn to write well is to read, and to think about what you read.

If you see a proof which you consider to be particularly well written, try to *learn from it* as an example of what's good to do. Similarly, if you see a proof which you consider to be particularly badly written, maybe you can learn something from the experience about what to avoid!

Also, *read it back* to yourself (or better yet, get a colleague to read it). You're writing it so others can read: it's good to test it to make sure this is possible. This goes particularly for proofs containing large amounts of symbols: these can be very hard to read, and reading it back to yourself is probably the best way of detecting this.

Use words and symbols

A common mistake that beginners make is to use very few words at all. Words are fantastic for *explaining* what you're doing.

In particular, too many people overuse the symbol "$\therefore$" (meaning "therefore"), and the symbol "$\because$" (meaning "because"). I won't use them at all in my proofs in these notes. There are so many good phrases which do its job, and choosing one helps you write what you're actually thinking. Also, it distracts the reader's eye from the symbols which matter — the actual maths you're doing — to things which don't.

Here are some phrases which do the job of "$\therefore$":

so, hence, therefore, thus, consequently, as a result, accordingly, for this reason, and so, and in particular, as a consequence, because of that, and then, and from that, ...

Here are some which do the job of "$\because$":

because, since, as a result of, as we have, as we know, as we have seen that, because we saw, ...

On the other hand, if you do write words, don't write them all in one paragraph. Try to write in a structured fashion, and if you are producing paragraph after paragraph after paragraph, ask yourself if you could use symbols to make your life easier in places.

The clearest pieces of mathematical writing I've read often make good use of words and symbols, working together.

When it's really important, I like to use both to make sure the point is clear. I might write,

Let $A = \sum_{i=1}^{n} a(i)$ be the sum of the first n values of a, and let $p > A$ be a prime number greater than A.

There are also ways of abusing symbols. Consider the following sentence:

The square of $5 = 25$.

Many novices write this, wanting it to mean:

(The square of 5) $= 25$.

However, experienced readers will read it as

The square of $(5 = 25)$.

This is of course nonsense: equations don't really have squares, and $5 = 25$ is an invalid equation anyway.

Some of he worst notational abuses are possible with induction. When writing an induction proof, don't explain what you need to prove as

Let $P(n) = a_n = 3^n + 1$.

This has $P(n)$ as a *number*, not a *statement*. Instead write:

Let $P(n)$ be the statement that $a_n = 3^n + 1$.

33

Make the structure obvious

Every sentence (and equation) of a proof needs to be justified:

- Usually, most sentences simply follow on from the one before. It is helpful to say so (using the connection words above, for example).

- Whenever a sentence doesn't just follow on from the one before (but is a radical new idea, or draws on things said a while earlier), it's even more important to say so. For example, you could:

  - use words which introduce a change of pace ("now we do this...");
  - name or number something earlier, and refer back to it by name or number;
  - leave a paragraph break;
  - leave a space;
  - have a descriptive section heading.

It is particularly important to avoid unstated assumptions.

For example, if a proof contains an assertion that some construction is a function, then the definition of a function gives you some things to check: that *every* element of the domain gives a *unique* element lying *inside* the codomain. Unless they're obviously true, it could be that these checks are the hardest and most interesting part of the proof. They could even be lies.

Explain beforehand and afterwards

Here is a classic account of a public speaking strategy:

> First I tell them what I'm going to tell them. Then I tell them. And then I tell them what I've told them.

This approach is even better in proofs than it is in public speaking.

Some explanation at the beginning is important. It may well be that your reader doesn't know even what you're aiming to do, and it's even more likely that your reader doesn't know how you're planning to do it.

Also, some explanation at the end is important. When you reach a conclusion, why do you think that what you have written actually means you have finished the proof?

If the proof is long, then regard it as being made of several parts. Give each part an explanation when you start and when you finish it.

I quite like proofs with the following kind of framework:

We'll prove this by induction on n.

The base case, where $n = 0$, is the statement "blah blah blah". But that's true, because blah blah blah blah blah.

Now we need the induction step: we assume the statement true for $n = k$, and prove it for $n = k + 1$. So we're assuming that blah blah blah, and have to prove that blah blah blah.

But blah blah blah blah blah. Also, blah blah blah blah blah. So, in conclusion, blah blah blah, which is what we had to prove. That finishes off the induction step, and so completes the proof.

Index