Math 300: Foundations of Higher Mathematics

Northwestern University, Lecture Notes

Written by Santiago Cañez

These are notes which provide a basic summary of each lecture for Math 300, "Foundations of Higher Mathematics", taught by the author at Northwestern University. The book used as a reference is the 3rd edition of *Book of Proof* by Hammack. Watch out for typos! Comments and suggestions are welcome.

Contents

Lecture 1: Direct Proofs

Welcome! This is a course dedicated to understanding how to read, write, and *think* about higher-level mathematics. The aim is two-fold: to become comfortable with writing rigorous mathematical arguments, and, more importantly, to get used to the thought process which goes into coming up with said arguments in the first place. This will be quite different from the computational courses you're likely more used to, but better reflects the approach with modern mathematics follows.

Cardinality. Just to get a sense for the types of things which a more rigorous approach to mathematics allows us to do, we'll give a brief introduction to the topic of *cardinality*, which will be one of the final things we'll look at in this course. The notion of cardinality gives us a precise way of talking about the "size" of a set, in the sense of the number of elements it has. The point is that if we want to answer questions like: "How large is $\mathbb{R}$, the set of real numbers?" or "Is $\mathbb{R}$ larger than $\mathbb{Z}$, the set of integers?", we had better have a precise definition of what "large" means in this context. This is meant to illustrate the fundamental idea that, in the end, definitions are absolutely crucial, and that *everything* we do in mathematics arises from precise definitions.

Whatever "size" means, it makes sense to say that $\mathbb{R}$ and $\mathbb{Z}$ are both *infinite* sets, since they each contain infinitely many numbers. However, leaving the answer at that, that $\mathbb{R}$ and $\mathbb{Z}$ are both infinite, is not the end of the story, because it turns out that nonetheless we can give meaning to the idea that $\mathbb{R}$ is *larger* than $\mathbb{Z}$; that is, even both $\mathbb{R}$ and $\mathbb{Z}$ are infinite, it will turn out that the cardinality of $\mathbb{R}$ is larger than that of $\mathbb{Z}$, which intuitively means that $\mathbb{R}$ has more elements than does $\mathbb{Z}$. Thus, the point is that once we give precise meaning to the notion of the "size" of set, it will make sense to talk about different "sizes" of infinity.

As another example, just going by intuition, it makes sense at first glance to say that $\mathbb{Z}$ is larger than $\mathbb{N}$, which is the set of positive integers. Indeed, since $\mathbb{Z}$ contains all elements of $\mathbb{Z}$ *along* with their negative, it might be tempting to say that $\mathbb{Z}$ is "twice" as large as $\mathbb{N}$. However, consider the following lists:

0	1	-1	2	-2	3	-3	$\dots$
1	2	3	4	5	6	7	$\dots$

In the first we list all integers, alternating (after the initial 0) between a positive and its negative, and in the second we list all positive integers. The point here is that these lists show there is a way to pair off elements of $\mathbb{Z}$ and $\mathbb{N}$ in a one-to-one manner so that nothing is left over in either set. Intuitively, this suggests there should be as many things in the first list as in the second, so that $\mathbb{Z}$ and $\mathbb{N}$ should actually have the same size. Indeed, this will be true, but again of course depends on giving precise meaning to the word "size".

Importance of definitions. The brief discussion of cardinality above is meant to emphasize that we ask some fairly strange and interesting questions in math, but which all depend on having precise definitions and techniques available. To look at something more concrete, suppose we were consider the following claim:

> If n is an even integer, then n^2 is even.

Is this true? Certainly if we consider different examples of even integers—2,4,6,8 for instance—the claim appears to be true since squaring each of these still results in an even integer. However, note that the claim is not saying the square of some specific even integer will still be even, but rather that the square of *any* even integer should still be even. In other words, the given claim should really be read as saying:

For all integers n, if n is even, then n^2 is even,

making it clear that the claim should hold for any n which happens to be an even integer.

Proving this thus requires we consider an arbitrary even integer n, with the goal being to show that n^2 is then even as well. To do so requires that we understand what "even" actually means, since it is only through working with a precise definition of even that we have any hope of proving our claim. Intuitively, an even integer is one which is "evenly divisible" by 2, but this doesn't work as a definition since we haven't yet given meaning to the phrase "evenly divisible". Here, then, are two possible ways of defining what it means for an integer n to be even:

First definition: n is even if $\frac{n}{2}$ is an integer.
Second definition: n is even if it can be written as $n = 2m$ for some integer m.

We can use either one, but the second definition should actually be preferred since the first still has some ambiguity built into it: if we want to say that $\frac{n}{2}$ is an integer, we would have to know what "integer" actually means in a more precise way. The second definition avoids this ambiguity.

Direct proofs. Nonetheless, here are proofs of our given claim using either proposed definition. In either case we must verify that n^2 also satisfies whichever definition of "even" we're working with.

Claim. If n is an even integer, then n^2 is even.

Proof 1. Suppose n is an even integer. Then $\frac{n}{2}$ is an integer, so

$$\frac{n^2}{2} = n\left(\frac{n}{2}\right)$$

is an integer since it is the product of two integers. Thus, since $\frac{n^2}{2}$ is an integer, n^2 even. □

Proof 2. Suppose n is an even integer. Then we can write n as $n = 2m$ for some integer m. Hence

$$n^2 = (2m)^2 = 4m^2 = 2(2m^2).$$

Thus since we can write n^2 as 2 times an integer, we conclude that n^2 is even. □

These are both examples of "direct proofs", in that they proceed directly from the given assumption to the desired conclusion, using only definitions and other manipulations. These are not the only possible proofs—for instance, we might say in the first attempt that since $\frac{n}{2}$ is an integer, $(\frac{n}{2})(\frac{n}{2}) = \frac{n^2}{4}$ is also an integer, which can only happen when n^2 is a multiple of 4, which also implies that n^2 is even. This is all true, but depends on some additional notions and facts we haven't made explicit yet, such as what it means for an integer to be a "multiple of 4", and why being a multiple of 4 implies being even. There is nothing wrong with this approach, but we should be mindful of the additional complexities it brings into play.

Now consider the claim that if n is an odd integer, then n^2 is odd. Again, we should first recognize the implicit "for all n" hiding in the setup: the claim is really "for all integers n, if n is odd, then n^2 is odd." Thus, proving this requires that we work with some arbitrary odd integer n without making any additional assumptions as to what it is. In addition, now we need a precise definition of "odd". One possible definition of odd is "not even", or that $\frac{n}{2}$ is not an integer. However, this doesn't give us much to work with, since there's not much we can say simply from

knowing that $\frac{n}{2}$ is not an integer; for instance, this doesn't really tell us anything about what $\frac{n}{2}$ can actually look like.

So, we look instead for a better definition odd. Here's one: an integer n is odd if it can be written as $n = 2m + 1$ for some integer m. This is simply saying that odd integers which are one more than an even integer. This is a good definition, since it gives us something concrete to work with, namely an explicit form for what n must look like. Here then is our proof.

Claim. If n is an odd integer, then n^2 is odd.

Proof. Suppose n is an odd integer. Then we can write it as $n = 2m+1$ for some integer m. Hence

$$n^2 = (2m + 1)^2 = 4m^2 + 4m + 1 = 2(2m^2 + 2m) + 1.$$

Thus since we can write n^2 as one plus twice an integer, we conclude that n^2 is odd. □

So far these are all pretty simple proofs, but they give a good introduction to the thought process behind working with proofs in general, where using precise definitions is key. Note also the structure: each proof begins with a marker showing the start of the proof, as evidenced by the *Proof* at the beginning, and each ends with a marker showing the end of the proof, as indicated by the □ symbol. The □ symbol is a very common way of indicating the end of a proof, and you should get in the habit of using it yourself. Note also that all proofs here are written using *complete sentences* with all thoughts spelled out in full. Again, this is something you should get in the habit of doing yourself. The goal is produce a clearly written proof which anyone reading can follow; the onus is on you, and not the reader, to make your ideas as clear as possible.

Another example. Here is one more example illustrating the ideas above. The claim is that the product of two rational numbers is itself rational. Of course, we first need a definition: a real number r is *rational* if it can be written as as the quotient $\frac{a}{b}$ of two integers a and b with $b \neq 0$. Thus, things like $\frac{1}{2}, -\frac{8}{3}, \frac{3}{17}$ are rational. (We'll look at examples of non-rational things later on.) Now that we have the required definition, proving our claim should be fairly straightforward; again, the point is simply to go wherever the definition takes us.

Claim. The product of two rational numbers is rational. To spell this out more explicitly, the claim is that for all x and y, if x and y are each rational, then xy is rational.

Proof. Suppose x and y are rational numbers. Then

$$x = \frac{a}{b} \text{ and } y = \frac{c}{d}$$

for some integers a, b, c, d with b and d nonzero. This gives

$$xy = \left(\frac{a}{b}\right)\left(\frac{c}{d}\right) = \frac{ac}{bd}.$$

Since the result is a fraction of two integers with nonzero denominator (we are taking for granted the fact that multiplying integers always results in an integer), we conclude that xy is rational as claimed. □

Following along in the book. As I said in class, we'll be jumping around in the book a bit in order to present things in a (hopefully) more natural manner. For instance, the book introduces the notion of a direct proof in Chapter 4, which you should definitely go through in order to see

more examples of proofs worked out. Just keep in mind that by this point the book has already introduced more material, so some things you'll see in Chapter 4 are things we have yet to discuss. Moving forward, it should not be too difficult to find the portions of the book which correspond to a specific topic given here, but feel free to ask if you're having trouble doing so.

My opinion is that it is better to jump into proofs right away and introduce required logical concepts (such as "negation", "contrapositive", etc) more organically as they are actually needed, as opposed to presenting it all at the start and saving proofs until afterwards. In class we'll be focusing on the key points to takeaway and on the overall thought process, which I think is simpler to get a handle on using our approach.

Lecture 2: More on Direct Proofs

Warm-Up 1. We show that the sum of an even integer and an odd integer is always odd. This is meant to be another simple example of a direct proof, which just requires working with the definitions of even and odd. Note how we use these definitions to show us both what it is we have to work with—when writing out what information our assumption gives us—and also to guide us towards what it is we want to establish. The given statement that "the sum of an even integer and an odd integer is always odd" can be rephrased as "if x is an even integer and y an odd integer, then $x + y$ is odd", which makes it a bit clearer to see what it is we have to do.

So, suppose x is an even integer and y an odd integer. Our goal is to show that $x + y$ is odd, which requires showing that we can write $x + y$ in the form required of an odd integer, namely as 2 times some integer plus 1. To get to this point, we use the definition of even to say that there exists an integer k such that $x = 2k$, and the definition of odd to say that there exists an integer ℓ such that $y = 2\ell + 1$. (Note that we should not use $y = 2k + 1$ here since we've already introduced k to mean something different previously; in other words, there is no reason why the k which satisfies $y = 2k + 1$ has to be the asme as the one which satisfies $x = 2k$, so we should use a different letter ℓ for the integer showing up in the statement of what it means for y to be odd.) Then

$$x + y = 2k + (2\ell + 1) = 2(k + \ell) + 1,$$

which is the required form of an odd integer.

We're done, but let us now write out the proof more cleanly without the additional parenthetical thoughts I put in above, to give a sense for how you would normally see it written:

Claim. If x is an even integer and y an odd integer, then $x + y$ is odd.

Proof. Suppose x is an even integer and y an odd integer. Since x is even and y is odd, there exist integers k and ℓ such that $x = 2k$ and $y = 2\ell + 1$. Then

$$x + y = 2k + (2\ell + 1) = 2(k + \ell) + 1,$$

which is the form required of an odd integer. Thus $x + y$ is odd as claimed. □

Warm-Up 2. We say that an integer a *divides* an integer b (or equivalently that b is divisible by a, or that b is a multiple of a) if there exists an integer k such that $b = ak$. We show that if a divides b and b divides c, then a divides c. Again this is an example of using basic definitions to carry us through; in particular, our end goal is to write c as a times some integer.

Proof. Suppose that a divides b and b divides c. Then there exists an integer k such that $b = ak$ and there exists an integer ℓ such that $c = b\ell$. Hence

$$c = b\ell = (ak)\ell = a(k\ell),$$

which shows that a divides c as desired. □

Example. Here is yet another example, only in this case we reach a point where working with definitions alone is not enough and we have to make use of another realization. The claim is that any even integer n can be written as $n = 4k$ or $n = 4k + 2$ or some integer k; in other words, any even integer is either a multiple of 4 or two more than a multiple of 4.

We start off, simple enough, by writing n as $n = 2m$ for some integer m from the fact that n is even. Our goal is to write n either in the form $4k$ or the form $4k + 2$, but now it is not a matter of simply manipulating the $n = 2m$ expression itself without bringing in some additional property. Indeed, to move from $2m$ to $4k$ or $4k + 2$ really requires knowing something about m itself, and the point is that m, being itself an integer, is either even or odd. Thus the "additional property" we need to consider here is that any integer is either even or odd, and this is what will allow our proof to move forward. Now, this is not a deep observation, but illustrates the idea that "direct proofs" still often require a good conceptual understanding of what we're dealing with, even if they are "direct". This is an example of a proof by cases, where we consider two cases—m being even vs m being odd—separately, and show that our conclusion holds in either case. To be clear, verifying our conclusion that "$n = 4k$ or $n = 4k + 2$ for some integer k" only requires that we show *one* of the statements $n = 4k$ or $n = 4k + 2$ holds and not both simultaneously; in general, an "or" statement is true when at least one of the claimed conclusions holds.

Here then is our proof:

Proof. Suppose n is an even integer. Then $n = 2m$ for some integer m. If m is even, there exists an integer k such that $m = 2k$, in which case $n = 2m = 2(2k) = 4k$. Otherwise m is odd, in which case there exists an integer k such that $m = 2k + 1$, so that $n = 2m = 2(2k + 1) = 4k + 2$. Thus, $n = 4k$ or $n = 4k + 2$ for some integer k as claimed. □

Non-divisibility example. Now we look at an example which does not deal with divisibility, evenness, or oddness at all, but is simply a statement about positive real numbers. The claim is that if x is a positive real number, then there exists a real number y such that

$$y < x < 2y.$$

In this case, there are no definitions we need to make use of, and the point is simply to make sense of what it is we're actually trying to show. In words, the claim is that no matter what positive number we take, we can always find another which is smaller but such that doubling it gives something larger the original. The fact that we are trying to prove "there exists a real number y" means that all we need to do is *produce* at least one y which satisfies the required property. Now, it may be that there are multiple y's which work, but again this type of *existence* proof only requires the existence of one such y.

Certainly if we take specific values of x we can find specific y's which work: for $x = 3$ for instance, $y = 2$ satisfies the requirement that $y < x < 2y$, and for $x = 5$ we can take $y = 4$ as one possible y. But this is not enough since we want to produce such a y for *any* positive x. Our choice of y should depend on the arbitrary x we're looking at, and our description of what y is should not depend on any information not given in the setup. The proof will take the structure of "Suppose $x > 0$. Set $y =$ (whatever value we claim is going to work), and then we'll verify that it does work."

To think about which y will work, we'll do some scratch work to think about what the claim means visually. We draw x to the right of 0 on a number line, and we're looking for y and $2y$ which look like:

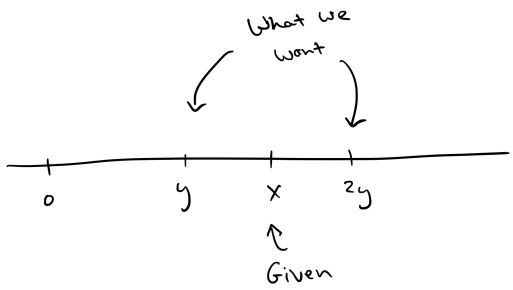

Think about what kinds of values, visually, to the left of x have the property that doubling them gives something to the right. We should convince ourselves that such values are those which occur strictly between $\frac{1}{2}x$ and x, since doubling anything smaller than $\frac{1}{2}x$ will still give something smaller than x:

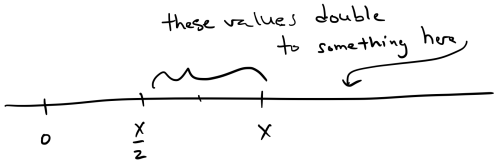

So, all we need is to pick a value for y which falls in this range. For instance, $y = \frac{3}{4}x$ will fall in this range, and so will tons of other things ($y = \frac{7}{8}x$, $y = \frac{\pi}{4}x$, etc), but all we need is one. So, in our proof we will set y to be $\frac{3}{4}x$, and then verify that this does indeed satisfy the property we want.

Note that the scratch work we went through here to determine which y will work is *not* something which will show up in our final proof, and was only our way of working through the thought process required to finish our argument. This is a crucial part of proof writing which cannot be emphasized enough: whenever you see a proof written out in a book or elsewhere, what you are seeing is the final presentable argument verifying the claim at hand, but which does not indicate the work which went into coming up with that argument in the first place. It is important to understand that this "scratch work" is really where the bulk of the difficulty lies; once we know what to do, writing out the actual formal proof is usually straightforward, but getting to that point is the hard part.

Claim. If $x > 0$, then there exists y such that $y < x < 2y$.

Proof. Suppose $x > 0$ and set $y = \frac{3}{4}x$. Then

$$y = \frac{3}{4}x < x < \frac{3}{2}x = 2y,$$

where we use the fact that x is positive to guarantee that the inequalities in $\frac{3}{4} < 1 < \frac{3}{2}$ are maintained after we multiply through by x. Thus $y = \frac{3}{4}x$ satisfies the required property. □

To emphasize once more, the proof does not illustrate where we came up with $y = \frac{3}{4}x$ in the first place, it only says "here is the y which will work and let's make sure it does". Finding a suitable choice of y came from thinking about what $y < x < 2y$ would mean visually.

Sets and subsets. The types of basic examples we've seen dealing with divisibility, evenness, oddness, and inequalities are a good thing to start off, but are not indicative of the more elaborate types of concepts you see in higher-level course. So, we will spend the next few days introducing new mathematical concepts on which we can try out the basic proof techniques we're building up.

One of the most fundamental notions in all of mathematics is that of a *set*, which for our purposes just means a collection of objects. It could be a collection of numbers, as in the case of the set of integers or the set of real numbers, a collection of people, or a collection of who knows what else. When A is a set, the notation "$x \in A$" means that x is an element of A, or that x is in A. For instance, $n \in \mathbb{Z}$ means that n is an integer. If A and B are sets, we say A is a *subset* of B if every element of A is an element of B, or in other words,

> A being a subset of B means that if $x \in A$, then $x \in B$.

Thus, A consists of elements of B, just maybe not all elements of B. We use the notation $A \subseteq B$ to mean that A is a subset of B. Verifying that one set is a subset of another requires verifying the definition directly: if $x \in A$, then $x \in B$.

Example. Let A be the set of all integers which are divisible by 4. In set notation we can express this as

$$A = \{n \in \mathbb{Z} \mid 4 \text{ divides } n\}.$$

Here, the braces { and } indicate that we are looking at a set, and the portions before and after the dividing | define the set in question: the "$n \in \mathbb{Z}$" to the left tells us what types of objects we are looking at, integers in this case, and the "4 divides n" to the right tells us what property they are required to satisfy in order to belong to the given set. Thus, here we are looking at the set of all $n \in \mathbb{Z}$ with the property that 4 divides n.

Let $B = \{m \in \mathbb{Z} \mid 2 \text{ divides } m\}$, which is the set of all integers which are divisible by 2, or in other words the set of all even integers. We claim that $A \subseteq B$, meaning that A is a subset of B, or that any integer which is divisible by 4 is also divisible by 2. To show this we start with an arbitrary $x \in A$, and work towards verifying that $x \in B$. Along the way we use the defining properties of what it means for x to be an element of A or B. Here is our proof:

Proof. Let $x \in A$. Then 4 divides x by definition of A, so there exists $k \in \mathbb{Z}$ such that $x = 4k$. This gives

$$x = 4k = 2(2k),$$

which shows that x is divisible by 2, and hence that $x \in B$. Thus $x \in A$ implies $x \in B$, so $A \subseteq B$ as claimed. □

Note again that we start with an arbitrary $x \in A$, use what it actually means for x to be in A in order to say that we can write x as $x = 4k$, and then manipulate to show that x is divisible by 2, which is what it means for x to be an element of B. Never lose sight of that fact that sets are defined in a way which tells us what it means for something to be or not be an element of that set; in other words, $x \in A$ in the above example gives us important information because we know what it means for something to be in A. Similarly, if we want to show that $x \in B$, all we have to do is verify that x satisfies the defining property required of elements of B. Definitions are key!

Lecture 3: Upper Bounds

Warm-Up. Define A and B to be the sets

$$A = \{n \in \mathbb{Z} \mid \text{ there exists } k \in \mathbb{Z} \text{ such that } n = 4k + 1\}$$

and

$$B = \{n \in \mathbb{Z} \mid \text{ there exists } k \in \mathbb{Z} \text{ such that } n = 4k + 9\}.$$

That is, A is the set of all integers which can be written in the form $4k + 1$ and B the set of all integers which can be written in the form $4k + 9$. We show that $A \subseteq B$ and $B \subseteq A$. Note that this is what it means to say that A and B are actually the same set: by definition, $A = B$ if it is true that $A \subseteq B$ and $B \subseteq A$.

To show that $A \subseteq B$, we must show that if $n \in A$, then $n \in B$. Thus we start with $n \in A$, which by definition of A means that we can write n as $n = 4k + 1$ for some $k \in \mathbb{Z}$. The goal is to show that we can write n in the form $4(\text{integer}) + 9$. Note that we use the same k in the defining expressions of A and B, but this is not meant to suggest that the same value of k is needed; i.e. we are not claiming that $n = 4k+1$ also equal to $n = 4k+9$ for the same k, but rather that $n = 4k+1$ can be written as $n = 4\ell + 9$ for some *other* integer ℓ.

To see what ℓ we thus need, we use what we want to show as a guide, namely that $4k + 1$ can be written as $4\ell + 9$. We are given k, so we need ℓ satisfying

$$4k + 1 = 4\ell + 9.$$

But now we can figure out precisely what ℓ must be solving for ℓ, and we see that ℓ must be $k - 2$. Thus, our scratch work shows that if we want to write $4k + 1$ in the form $4(\text{integer}) + 9$ instead, the "integer" term we need is $k - 2$. Thus in our proof we will simply verify that $\ell = k - 2$ is the integer which expresses $4k + 1$ as $4\ell + 9$. A similar scratch works for the other containment $B \subseteq A$ we need to show, which requires showing that any $4k + 9$ in B can be written as $4\ell + 1$ for some choice of ℓ, which we determine ahead of time by solving $4k + 9 = 4\ell + 1$ for ℓ. Here, then, is our final proof:

Claim. For the sets A and B defined above, we have $A = B$, or in other words, $A \subseteq B$ and $B \subseteq A$.

Proof. First we show that $A \subseteq B$. Let $n \in A$. Then there exists $k \in \mathbb{Z}$ such that $n = 4k + 1$. Hence:

$$4(k - 2) + 9 = 4k - 8 + 9 = 4k + 1 = n,$$

so $n = 4(k - 2) + 9$ is in the form required of an element of B. Thus $n \in B$, so $A \subseteq B$.

Second we show that $B \subseteq A$. Let $n \in B$. Then there exists $k \in \mathbb{Z}$ such that $n = 4k + 9$, so

$$4(k + 2) + 1 = 4k + 8 + 1 = 4k + 9 = n,$$

and thus $n = 4(k+2)+1$ is in the form required of an element of A. Hence $B \subseteq A$, so since $A \subseteq B$ and $B \subseteq A$, we conclude that $A = B$. □

Upper bounds. We now move to introducing a new mathematical concept—the notion of an upper bound of a set of real numbers. On the one hand, for those of you planning on taking a course in real analysis, this is a crucial concept related to properties of real numbers. On the other hand, and the main reason why we're introducing it in this course, it's a notion which is interesting and fairly simple to understand, but provides good practice in working with definitions

and mathematical reasoning. Indeed, understanding various properties of the set of real numbers $\mathbb{R}$ is a theme we'll return to again and again, as applications of the techniques we'll develop. This material is NOT in the book we are using.

Here is the definition. Suppose $S \subseteq \mathbb{R}$, so that S is a set consisting of real numbers. We say that a real number u is an *upper bound* of S if for all $s \in S$, $s \leq u$. Thus an upper bound of S is a real number which is bigger than or equal to everything in S, as the name "upper bound" is meant to suggest. Note that upper bounds are not unique in that if a set has an upper bound, it will have many of them. Indeed, 2 is an upper bound of the *closed interval* $[0, 2]$ defined by

$$[0, 2] = \{x \in \mathbb{R} \mid 0 \leq x \leq 2\},$$

but so are 3, 4, and tons of other things. Also note that not all subsets of $\mathbb{R}$ have upper bounds; for instance, $\mathbb{Z} \subseteq \mathbb{R}$ does not have an upper bound since there is no restriction as to how large integers can be.

Now, among all upper bounds of a set, there is one in particular which is worth singling out: the *smallest* upper bound. This notion is important enough in mathematics that it is given a special name: *supremum*. Here is the precise definition: the *supremum* (or *least upper bound*) of S is an upper bound b of S such that for any other upper bound u of S, we have $b \leq u$. Thus, this definition precisely says that the supremum is an upper bound which is smaller than or equal to any other upper bound, as the alternate term "least upper bound" is meant to suggest. We use the notion $\sup S$ to denote the supremum of S, if it exists.

Example. We claim that $\sup [0, 2] = 2$. According to the definition of supremum, showing that $\sup [0, 2] = 2$ requires two things: showing that 2 is an upper bound of $[0, 2]$, and showing that 2 is smaller than or equal to any other upper bound. First, by definition of the interval $[0, 2]$, if $x \in [0, 2]$ we have $0 \leq x \leq 2$, so 2 is larger than or equal to anything in $[0, 2]$, and hence 2 is an upper bound of $[0, 2]$.

Now, let $u \in \mathbb{R}$ be another upper bound of $[0, 2]$. We must show that $2 \leq u$. Since u is an upper bound of $[0, 2]$, we know that u is larger than or equal to anything in $[0, 2]$. But $2 \in [0, 2]$, so in particular u must be larger than or equal to 2, which is precisely what we want to show. Thus if u is any other upper bound of $[0, 2]$, we have $2 \leq u$, so 2 is the supremum of $[0, 2]$ as claimed.

What about $(0, 2)$? The key realization above in showing that $2 \leq u$ came from recognizing that 2 is an element of the set of which u is an upper bound, so $2 \leq u$ simply the fact that u is an upper bound. However, note that this reasoning does not work if we consider our set to be the *open* interval $(0, 2)$ defined by

$$(0, 2) = \{x \in \mathbb{R} \mid 0 < x < 2\}$$

instead. It is still true, at least intuitively, that $\sup (0, 2) = 2$ since visualizing this on a number line does suggest that 2 is the smallest upper bound of $(0, 2)$. And showing that 2 is an upper bound of $(0, 2)$ is just as simple as in the example above since, by definition, anything in $(0, 2)$ is smaller than 2.

But in order to show that 2 is the *least* upper bound requires a new approach. In this case, 2 is NOT in $(0, 2)$, so if u is another upper bound of $(0, 2)$ we cannot say immediately that $2 \leq u$ simply by the fact that u is an upper bound: we only know that u is larger than or equal to everything in $(0, 2)$, but now 2 is not such an element in $(0, 2)$. What we need to do in this case is show that nothing *smaller* than 2 can be an upper bound of $(0, 2)$; if nothing smaller than 2 is an upper bound, and 2 itself is *an* upper bound, then it makes sense to conclude that 2 is indeed the *smallest* upper bound of $(0, 2)$.

But how exactly do we show that nothing smaller than 2 is an upper bound of $(0, 2)$? Showing this requires knowing precisely what it means for something to *not* be an upper bound of set, which requires *negating* the definition "for all $s \in S$, $s \leq u$" of an upper bound. We will come back to this later after we discuss *negations*, which will require us to understand a bit more basic logic.

Uniqueness of supremums. In the definition of supremum we referred to a real number being *the* supremum of set, which suggests that if a set has a supremum then it can only have one. This is true but requires justification, since uniqueness of supremums is not built into the definition of supremum itself, but will instead be a consequence. This is important in order to make the notation $\sup S$ unambiguous: if it was possible to have more than one upper bound, the notation $\sup S$ would not be enough to specify to which one we were referring.

So, how we show that something is unique? The standard way of doing so is to suppose you have two such things and then show that they actually have to be the same. In our case, we claim that if b and b' are both supremums of a set S of real numbers, then $b = b'$. We must show that $b = b'$ using only the fact that b and b' are supremums of S, and if we look at the definition of supremum we see that this definition involves statements saying that certain inequalities will hold. Since inequalities are then all we have to work with, we must think about how to show that two numbers are the same using only inequalities. For instance, one way to do this is to show that each number is smaller than or equal to the other: if $b \leq b'$ and $b' \leq b$, then we we will be able to conclude that $b = b'$.

Thus we now have a strategy: show that $b \leq b'$ and $b' \leq b$ using the fact that b and b' are both supremums of S. The fact that b is a supremum means, by definition, that it will be smaller than or equal to any other upper bound. Thus if we want to show that some number x is larger than or equal to b, all we need to know is that x is an upper bound of S since this alone will guarantee that $x \geq b$. But in our situation, we know that b' is an upper bound of S since being an upper bound is part of the definition of being a supremum, so this will give us one of the inequalities $b \leq b'$ we need. The other inequality will follow from the same reasoning after switching the roles of b and b'. Here, then, is our final proof:

Claim. If a set $S \subseteq \mathbb{R}$ has a supremum, then it has only one.

Proof. Suppose b and b' are both supremums of S. We will show that $b = b'$. Since b is a supremum of S and b' is an upper bound of S, we have that $b \leq b'$ since b by definition is smaller than or equal to any other upper bound of S. Similarly, since b' is a supremum of S and b is an upper bound of S, we have that $b' \leq b$ since b', being a supremum, is smaller than or equal to any other upper bound of S. Thus since $b \leq b'$ and $b' \leq b$, we conclude that $b = b'$, showing that supremums are unique. □

Final example. As a final example, suppose that $S \subseteq T \subseteq \mathbb{R}$, so that S and T are both sets of real numbers with S contained in T. Suppose also that both S and T have supremums. We claim that $\sup S \leq \sup T$. Visually this makes sense if you imagine S and T on a number line. To show this we again use the definition of supremum as a guide. The number $\sup S$ is smaller than or equal to any upper bound of S, so if we want to show that $\sup S \leq \sup T$ all we need to show is that $\sup T$ is *an* upper bound of S. Why is this true? Well, we know that $\sup T$ is an upper bound of T, meaning that for all $x \in T$ we have $x \leq \sup T$. But in particular, since $S \subseteq T$, anything in S is also in T, and so anything in S will thus be less than or equal to $\sup T$ as well. Here is a cleanly written proof:

Claim. If $S \subseteq T \subseteq \mathbb{R}$ and both S and T have supremums, then $\sup S \leq \sup T$.

Proof. For any $x \in S$, $x \in T$ so $x \leq \sup T$ since $\sup T$ is an upper bound of T. Thus since $x \leq \sup T$ for all $x \in S$, we conclude that $\sup T$ is an upper bound of S and hence that $\sup S \leq \sup T$ since $\sup S$ is smaller than or equal to any other upper bound of S. □

Summary. We'll continue using the notion of a supremum in the coming weeks to illustrate more properties of $\mathbb{R}$ and to give more examples of things on which we can apply the techniques we'll soon develop. The key point to takeaway in the examples we say today was that in the end everything came down to working with definitions: some arguments were a little more straightforward than others, while others required thinking about strategies for showing what it is we wanted to show, but always we used definitions as a guide for what to do.

Lecture 4: Unions & Intersections

Warm-Up. Suppose $A, B \subseteq \mathbb{R}$ have supremums. Define $A + B$ to be the set of all numbers obtained by adding something in A to something in B:

$$A + B = \{a + b \in \mathbb{R} \mid a \in A \text{ and } b \in B\}.$$

We show that $\sup(A + B) = \sup A + \sup B$. This makes sense if you consider some examples: in the case where, say, $A = [0, 2]$ and $B = [-1, 3]$, we have $A + B = [-1, 5]$ since adding numbers in the interval $[0, 2]$ to those in the interval $[-1, 3]$ results in numbers in the interval $[-1, 5]$, and the supremum of $A + B = [-1, 5]$ is indeed $5 = \sup A + \sup B$.

In order to show that $\sup(A + B) = \sup A + \sup B$, we can directly show that $\sup A + \sup B$ satisfies the defining properties of the supremum of $A + B$—namely, that $\sup A + \sup B$ is an upper bound of $A + B$ and that it is smaller than or equal to any other upper bound of $A + B$. Since supremums are unique, this alone will guarantee that $\sup A + \sup B = \sup(A + B)$. Now, the first requirement comes from the fact that $\sup A$, being an upper bound of A, is bigger than or equal to anything in A, and similarly $\sup B$ is bigger than or equal to anything in B: for any $a + b$ with $a \in A$ and $b \in B$, we have $a \leq \sup A$ and $b \leq \sup B$, so $a + b \leq \sup A + \sup B$.

However, the second requirement, that $\sup A + \sup B$ is smaller than or equal to any other upper bound of $A + B$, requires more thought. If u is an upper bound of $A + B$, we need to show that

$$\sup A + \sup B \leq u.$$

How do we get to this point? We need a way of manipulating and reformulating this inequality in a way which will allow us to use our assumption that u is an upper bound of $A + B$. Note that we can rewrite the given inequality as

$$\sup A \leq u - \sup B.$$

But now we have something we can work with: our goal is to show that $u - \sup B$ is larger than or equal to $\sup A$, and by the fact $\sup A$ is the *smallest* upper bound of A, it is enough to show that $u - \sup B$ is an upper bound of A as well. In other words, if we know that $u - \sup B$ is *an* upper bound of A, the definition of supremum alone will guarantee that $u - \sup B$ is larger than or equal to the smallest upper bound $\sup A$ of A.

The point is that we've now rephrased the inequality $\sup A \leq u - \sup B$ we want to establish as the claim that $u - \sup B$ is an upper bound of A. Justifying this latter claim requires showing that

$$a \leq u - \sup B \text{ for all } a \in A$$

since this, by definition, is what it means for $u - \sup B$ to be an upper bound of A. This is good: our assumption that u is an upper bound of $A + B$ only tells us something about inequalities of the form

$$a + b \in u \text{ for } a \in A \text{ and } b \in B$$

which do not explicitly mention supremums, so we have to find a way to rephrase the inequality we want $\sup A \leq u - \sup B$ also in a way which does not mention supremums. So far we're at

$$a \leq u - \sup B \text{ for all } a \in A,$$

by the same type of reasoning will give us a way to rephrase this without using $\sup B$: we can rearrange this inequality as

$$\sup B \leq u - a \text{ for all } a \in A,$$

and to justify this all we need to show is that $u - a$ is an upper bound of B for any $a \in A$, since if so it must be larger than or equal to the smallest upper bound $\sup B$ of B as desired.

The thought process above is all a part of our scratch work, where we take what it is we want to show and "work backwards" to see how we can get to that point. It then becomes a question of rephrasing statements and unpacking definitions to get things to work out. Here, then, is our final proof:

Proof. Since $\sup A, \sup B$ are upper bonds of A, B respectively, we have

$$a \leq \sup A \text{ for all } a \in A \text{ and } b \leq \sup B \text{ for all } b \in B.$$

Thus, for any $a + b \in A + B$, where $a \in A$ and $b \in b$, we have:

$$a + b \leq \sup A + b \leq \sup A + \sup B,$$

showing that $\sup A + \sup B$ is an upper bound of $A + B$.

Now, suppose u is any other upper bound of $A + B$. Then

$$a + b \leq u \text{ for any } a \in A \text{ and any } b \in B.$$

Rearranging this gives that

$$\text{for any } a \in A, b \leq u - a \text{ for any } b \in B.$$

But this means that for any $a \in A$, $u - a$ is an upper bound of B, so

$$\sup B \leq u - a \text{ for any } a \in A$$

by the fact that $\sup B$ is the smallest upper bound of B. Rearranging once more gives

$$a \leq u - \sup B \text{ for all } a \in A,$$

which shows that $u - \sup B$ is an upper bound of A. Thus $\sup A \leq u - \sup B$ since $\sup A$ is the smallest upper bound of A, and thus we conclude that $\sup A + \sup B \leq u$. Hence $\sup A + \sup B$ is an upper bound of $A + B$ which is smaller than or equal to any other upper bound of $A + B$, so $\sup A + \sup B$ is the supremum of $A + B$ as claimed. □

Unions and intersections. We now move to studying properties of abstract sets, and in particular constructions which allow us to construct new sets out of old sets. The two constructions we'll consider first are that of *union* and *intersection*. This material is first introduced in Chapter 6 of the book, but Chapter 7 is where it really begins to be developed.

Suppose A, B are both subsets of some larger set U. The *union* of A and B, denoted by $A \cup B$, is the set of all things we get by throwing in all elements of A together with all elements of B. More precisely, the union can be defined as

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

Thus, to say that x is in $A \cup B$ means that $x \in A$ *or* $x \in B$. The *intersection* of A and B, denoted by $A \cap B$, is the set of all things A and B have in common, or more precisely:

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

Thus, to say that x is in $A \cap B$ means that $x \in A$ *and* $x \in B$.

For a simple example, let $A = [-1, 2]$ and $B = [0, 4]$. Then $A \cup B = [-1, 4]$ since throwing in all numbers in the interval $[-1, 2]$ together with all numbers in the interval $[0, 4]$ gives all numbers in the interval $[-1, 4]$. The only numbers which the intervals A and B have in common are those between 0 and 2 inclusive, so $A \cap B = [0, 2]$. For another general example, take A to be any set and $\emptyset$ to be the *empty set*, which is the set which contains no elements at all. Then $A \cup \emptyset = A$ since $\emptyset$ contributes no additional elements, and $A \cap \emptyset = \emptyset$ since $x \in A \cap \emptyset$ would means that $x \in A$ and $x \in \emptyset$, but there is no such x for which $x \in \emptyset$ can be true since $\emptyset$ contains no elements.

Example. We will prove that for any sets A, B, C:

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

This is a basic property of sets describing how the operations of taking unions and intersections relate to one another, and is in some a "distributive" property for sets. The claim, in words, is that if we take the elements of A that also belong to $B \cup C$, we get the same thing as elements of A that belong to B, or elements of A that belong to C.

Being a statement that two sets are equal, we prove this by showing that each side is a subset of the other. We do this with a so-called "element chase", which is a type of proof where we start with an element on one side, and "chase it through" an unwinding of various definitions until we see that the same element belongs to the other side as well. This can get a bit tedious, but provides great practice in working with definitions and structuring proofs appropriately.

Proof. Let $x \in A \cap (B \cup C)$. Then $x \in A$ and $x \in B \cup C$ by definition of intersection. Since $x \in B \cup C$, we have $x \in B$ or $x \in C$ by definition of union. Hence we have two possibilities to consider: $x \in B$ or $x \in C$. If $x \in B$, then since $x \in A$ and $x \in B$, we get $x \in A \cap B$. If $x \in C$, then since $x \in A$ and $x \in C$, we get $x \in A \cap C$. Thus in either case we have $x \in A \cap B$ or $x \in A \cap C$, so $x \in (A \cap B) \cup (A \cap C)$. Hence $A \cap (B \cup C) \subseteq (A \cap B) \cup (A \cap C)$.

Now let $x \in (A \cap B) \cup (A \cap C)$. Then $x \in A \cap B$ or $x \in A \cap C$ by definition of union. If $x \in A \cap B$, then $x \in A$ and $x \in B$; since $x \in B$ we get $x \in B \cup C$. If $x \in A \cap C$, then $x \in A$ and $x \in C$; since $x \in C$, we get $x \in B \cup C$. Hence in either case we have $x \in A$ and $x \in B \cup C$, so $x \in A \cap (B \cup C)$. Thus $(A \cap B) \cup (A \cap C) \subseteq A \cap (B \cup C)$, so since $A \cap (B \cup C) \subseteq (A \cap B) \cup (A \cap C)$ and $(A \cap B) \cup (A \cap C) \subseteq A \cap (B \cup C)$, we conclude that $A \cap (B \cup C) = (A \cap B) \cup (A \cap C)$ as claimed. □

Lecture 5: More on Unions & Intersections

Warm-Up 1. Suppose A and B are sets. We show that $A \subseteq B$ if and only if $A \cap B = A$. First, a bit of logic: "if and only if" means that both sides imply each other. That is, a statement of the form "P if and only if Q" means "if P, then Q" *and* "if Q, then P". Thus, proving an if and only if statement requires proving two implications. As a result of this if and only if statement we say that $A \subseteq B$ and $A \cap B = A$ are *logically equivalent*, meaning that both statements mean the same thing and are just ways of rephrasing one another. Intuitively this should make sense: $A \cap B$ takes everything that A and B have in common, and if this results in A itself, then it should have been the case that everything in A was already in B to start with.

Proof. We first prove the forward direction. Suppose $A \subseteq B$. We want to show that $A \cap B = A$, which requires us to show that $A \cap B \subseteq A$ and $A \subseteq A \cap B$. If $x \in A \cap B$, then $x \in A$ and $x \in B$. In particular, $x \in A$ so we conclude that $A \cap B \subseteq A$. Now let $x \in A$. Since $A \subseteq B$, we then know that $x \in B$ as well. Hence $x \in A$ and $x \in B$, so $x \in A \cap B$. Thus $A \subseteq A \cap B$, which together with $A \cap B \subseteq A$ means that $A \cap B = A$.

For the backwards direction, suppose $A \cap B = A$. We want to show that $A \subseteq B$. Hence, let $x \in A$. Since $A = A \cap B$, we know that $x \in A \cap B$ as well. Hence $x \in A$ and $x \in B$ by definition of intersection. In particular $x \in B$, so anything in A is in B and hence $A \subseteq B$ as desired. □

Warm-Up 2. The *Cartesian product* $S \times T$ (pronounced "S cross T") of sets S and T is defined to be the set of all ordered pairs (x, y) of an element x of S and y in T. Concretely:

$$S \times T = \{(x, y) \mid x \in S \text{ and } y \in T\}.$$

For instance, $\mathbb{R} \times \mathbb{R} = \{(x, y) \mid x, y \in \mathbb{R}\}$ is the ordinary xy-plane and is often denoted by $\mathbb{R}^2$. Similarly, ordinary 3-dimensional space $\mathbb{R}^3$ denotes the Cartesian product $\mathbb{R} \times \mathbb{R} \times \mathbb{R}$ of three sets. (The book doesn't talk about Cartesian products until Chapter 9.)

Suppose A, B, C are sets. We show that

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

Again this is meant to be an example of an "element chase" argument, where we just unpack definitions. Think of this stated equality as also a type of "distributive" property for sets, just as in the equality $A \cap (B \cup C) = (A \cap B) \cup (A \cap C)$ we proved last time.

Proof. Let $(x, y) \in A \times (B \cup C)$. Then $x \in A$ and $y \in B \cup C$ by definition of Cartesian product. Since $y \in B \cup C$, $y \in B$ or $y \in C$. In the case where $y \in B$, since $x \in A$ and $y \in B$ we have $(x, y) \in A \times B$. In the case where $y \in C$, we have $(x, y) \in A \cap C$ since $x \in A$ and $y \in C$. Thus in either case we have $(x, y) \in A \times B$ or $(x, y) \in A \times C$, so $(x, y) \in (A \times B) \cup (A \times C)$. Hence $A \times (B \cup C) \subseteq (A \times B) \cup (A \times C)$.

Conversely let $(x, y) \in (A \times B) \cup (A \times C)$. Then $(x, y) \in A \times B$ or $(x, y) \in A \times C$. If $(x, y) \in A \times B$, then $x \in A$ and $y \in B$; since $yinB$, $y \in B \cup C$. If $(x, y) \in A \times C$, then $x \in C$ and $y \in C$; since $y \in C$, $y \in B \cup C$ as well. Hence in either case we have $x \in A$ and $y \in B \cup C$, so $(x, y) \in A \times (B \cup C)$. Thus $(A \times B) \cup (A \times C) \subseteq A \times (B \cup C)$, so we conclude that $A \times (B \cup C) = (A \times B) \cup (A \times C)$ as claimed. □

Warning. Here is a fact: if $S \subseteq A$ or $S \subseteq B$, then $S \subseteq A \cup B$. That is, being a subset of A or B alone guarantees being a subset of the union $A \cup B$. However, the *converse* of this claim, namely the statement that if $S \subseteq A \cup B$ then $S \subseteq A$ or $S \subseteq B$, is not true. For instance, taking $S = \mathbb{Z}$,

$A = \{\text{set of even integers}\}$ and $B = \{\text{set of odd integers}\}$ provides one possible *counterexample*; in this case, $A \cup B = S$ so S is a subset of $A \cup B$ but it is not true that S is a subset of A or of B. For another counterexample, we can take $S = \{1, 3\}$, $A = \{1, 2\}$, and $B = \{2, 3\}$; here $A \cup B = \{1, 2, 3\}$, so S is a subset of $A \cup B$, but it is not true that $S \subseteq A$ or $S \subseteq B$.

Nonetheless, here is a purported proof that $S \subseteq A \cup B$ does imply $S \subseteq A$ or $S \subseteq B$, and so the point is to understand why this proof fails. Being able to recognize such mistakes is important in building up intuition and getting used to working with logic and rigor, so we'll see more examples of "false proofs" later on.

False claim. If $S \subseteq A \cup B$, then $S \subseteq A$ or $S \subseteq B$.

"Proof". Let $x \in S$. Since $S \subseteq A \cup B$, we have $x \in A \cup B$. Thus $x \in A$ or $x \in B$. If $x \in A$, then we have that $x \in S$ implies $x \in A$, so in this case $S \subseteq A$. If $x \in B$, then $x \in S$ implies $x \in B$, so in this case $S \subseteq B$. Hence $S \subseteq A$ or $S \subseteq B$ as claimed. □

We know this proof cannot be correct since we previously gave counterexamples to the claim being made, so what exactly is wrong? The issue comes in stating that "$x \in S$ implies $x \in A$" is true in the case where $x \in A$, or in stating that $x \in S$ implies $x \in B$" is true in the case where $x \in B$. To be able to say that "$x \in S$ implies $x \in A$" for instance, we would have to know that *for all* x, $x \in S$ implies that $x \in A$. However, we do not know that this is in fact true for all $x \in S$—we definitely know that any $x \in S \subseteq A \cup B$ is either in A or B, but which of $x \in A$ or $x \in B$ occurs can change depending on which x we're looking at. In other words, we know that *some* $x \in S$ also satisfy $x \in A$, and that *some* $x \in S$ also satisfy $x \in B$, but we don't know that *all* $x \in S$ also satisfy $x \in A$, nor that *all* $x \in S$ also satisfy $x \in B$ as would be required in order to conclude that $S \subseteq A$ or $\subseteq B$ respectively.

This is subtle point which is the type of thing which one can quickly gloss over when constructing a proof involving sets. We should be mindful that whatever we're claiming to be true is in fact true and that we've provided adequate justifications.

General unions and intersections. So far we've defined the union and intersection of two sets at a time, but there's no reason why we couldn't look at the union or intersection of three, four, or more sets. For that matter, there's no reason why we couldn't look at the union or intersection of *infinitely many* sets.

For instance, suppose that we have an infinitely collection of sets:

$$A_1,\ A_2,\ A_3,\ \dots$$

indexed by positive integers. The union and intersection of these sets are often denoted by

$$\bigcup_{n=1}^{\infty} A_n \text{ and } \bigcap_{n=1}^{\infty} A_n$$

respectively, which should be viewed as analogous to the notation $\sum_{n=1}^{\infty}$ for infinite sums. More generally, we consider can consider the union and intersection of sets *indexed* by other infinite collections apart from positive integers. In the book, indexed sets are introduced in Chapter 8.

Example. For any $r > 0$ define D_r to be the *open disk* of radius r centered at the origin, which is the set of all points in $\mathbb{R}^2$ whose distance to the origin is less than r:

$$D_r = \left\{(x, y) \in \mathbb{R}^2 \;\middle|\; \sqrt{x^2 + y^2} < r\right\}.$$

Visually, this indeed looks like a disk (i.e. the region enclosed by a circle), and it is "open" since it does not contain the boundary circle $\sqrt{x^2+y^2}=r$ itself. We can view the D_r as a collection of sets indexed by positive numbers, which describe the possible radii.

We compute, with justification, the union and intersection of these sets:

$$\bigcup_{r>0} D_r \text{ and } \bigcap_{r>0} D_r.$$

To be clear, the union consists of all points which belong to *some* disk in our collection, and the intersection consists of all points which belong to *all* disks in our collection; we can write this out precisely as:

$$\bigcup_{r>0} D_r = \left\{(x,y) \in \mathbb{R}^2 \;\middle|\; \text{there exists } r>0 \text{ such that } (x,y) \in D_r\right\}$$

and

$$\bigcap_{r>0} D_r = \left\{(x,y) \in \mathbb{R}^2 \;\middle|\; \text{for all } r>0, (x,y) \in D_r\right\}$$

To see what these should intuitively be we simply think about what they look like when drawn in the xy-plane. Never be afraid to use pictures as a way to develop intuition!

For the union, we draw one disk, then another of a larger radius, then another, and so on—everything covered by all disks you could possibly draw should be included in the union, so since visually we can keep drawing larger and larger disks to cover the entire xy-plane, we see that the union should be all of $\mathbb{R}^2$:

$$\bigcup_{r>0} D_r = \mathbb{R}^2.$$

Now, to actually prove this we must show, since we are claiming that certain sets are equal, that each side is a subset of the other. The forward containment does not require much since each D_r is already a subset of $\mathbb{R}^2$, but the backwards containment requires some care: the claim is that if (x,y) is any point of $\mathbb{R}^2$, then (x,y) is in the union of all the D_r, which requires us to show that it belongs to *some* D_r of an appropriate radius. We should be clear about which radius we are taking if we want to be precise, and visually the point is that we need a radius which will extend further from the origin than $\sqrt{x^2+y^2}$, which is the distance from (x,y) to $(0,0)$.

Proof that the claimed union equality is correct. Let $(x,y) \in \bigcup_{r>0} D_r$. Then there exists $r>0$ such that $(x,y) \in D_r$ by the definition of union. But $D_r \subseteq \mathbb{R}^2$ by construction of D_r as a set of points in $\mathbb{R}^2$ satisfying some condition, so $(x,y) \in \mathbb{R}^2$ as well. Hence $\bigcup_{r>0} D_r \subseteq \mathbb{R}^2$.

Conversely suppose $(x,y) \in \mathbb{R}^2$ and set $s = \sqrt{x^2+y^2}+1$. Since

$$\sqrt{x^2+y^2} < \sqrt{x^2+y^2}+1 = s,$$

the point (x,y) satisfies the requirement needed to belong to the disk D_s, so $(x,y) \in D_s$. Hence since (x,y) is in some disk D_s, we conclude that $(x,y) \in \bigcup_{r>0} D_r$. Thus $\mathbb{R}^2 \subseteq \bigcup_{r>0} D_r$, so we have that $\bigcup_{r>0} D_r = \mathbb{R}^2$ as claimed. □

Now, for the intersection, we are looking for points which belong to all disks D_r simultaneously, no matter the radius. If you draw disks of smaller and smaller radius, it should seem intuitively clear that the only point which all disks have in common is $(0,0)$, so we guess that

$$\bigcap_{r>0} D_r = \{(0,0)\}.$$

Again, to prove this requires showing that each side is a subset of the other. The forward direction here is the one which requires some thought, since we have to know that if (x, y) belongs to all disks, then it must be $(0, 0)$, and it is not so clear at the start how to actually show this. The key is using the definition of D_r to say that (x, y) satisfies

$$\sqrt{x^2 + y^2} < r \text{ for all } r > 0,$$

and thinking about what number $\sqrt{x^2 + y^2}$, which is nonnegative but smaller than everything positive, could then possibly be.

Proof that the claimed intersection equality is correct. Let $(x, y) \in \bigcap_{r>0} D_r$. Then $(x, y) \in D_r$ for all $r > 0$ by definition of intersection, so

$$\sqrt{x^2 + y^2} < r \text{ for all } r > 0$$

by definition of D_r. But now, this says that $\sqrt{x^2 + y^2}$ is a nonnegative number which is smaller than every positive number, and since 0 is the only such nonnegative number we conclude that $\sqrt{x^2 + y^2} = 0$. But this then requires that x^2 and y^2 both be 0 as well, so $x = 0$ and $y = 0$. Thus $(x, y) = (0, 0)$, so $(x, y) \in \{(0, 0)\}$. Hence $\bigcap_{r>0} D_r \subseteq \{(0, 0)\}$.

Conversely, take $(0, 0) \in \{(0, 0)\}$. Since $\sqrt{0^2 + 0^2} = 0 < r$ for any $r > 0$, the definition of D_r says that $(0, 0) \in D_r$ for any $r > 0$. Thus $(0, 0) \in \bigcap_{r>0} D_r$, so $\{(0, 0)\} \subseteq \bigcap_{r>0} D_r$. Hence the claimed equality holds. □

Lecture 6: Negations

Warm-Up. We determine, with proof, the following infinite union and intersection:

$$\bigcup_{n\in\mathbb{N}} \left(-\tfrac{1}{n}, \tfrac{1}{n}\right) \text{ and } \bigcap_{n\in\mathbb{N}} \left(-\tfrac{1}{n}, \tfrac{1}{n}\right).$$

To be clear, these are respectively the union and intersection of all intervals of the form $\left(-\frac{1}{n}, \frac{1}{n}\right)$ as n ranges among all positive integers; that is, we are considering the intervals

$$(-1, 1), \left(-\tfrac{1}{2}, \tfrac{1}{2}\right), \left(-\tfrac{1}{3}, \tfrac{1}{3}\right), \ldots.$$

Thinking about these intervals drawn on a number line, it appears that the union should be the interval $(-1, 1)$ since varying the intervals in question will cover all numbers strictly between -1 and 1, and it appears that the intersection should consist only of the number 0 since all other numbers are excluded as the intervals in question become smaller and smaller. Thus, we conjecture that

$$\bigcup_{n\in\mathbb{N}} \left(-\tfrac{1}{n}, \tfrac{1}{n}\right) = (-1, 1) \text{ and } \bigcap_{n\in\mathbb{N}} \left(-\tfrac{1}{n}, \tfrac{1}{n}\right) = \{0\}.$$

We now prove these two equalities. First, let $x \in \bigcup_{n\in\mathbb{N}} \left(-\frac{1}{n}, \frac{1}{n}\right)$. By definition of union, this means there exists $n \in \mathbb{N}$ such that $x \in \left(-\frac{1}{n}, \frac{1}{n}\right)$, so that $-\frac{1}{n} < x < \frac{1}{n}$. But $n \geq 1$, so

$$-1 \leq -\frac{1}{n} < x < \frac{1}{n} \leq 1,$$

which shows that $x \in (-1, 1)$ as well. Hence $\bigcup_{n\in\mathbb{N}} \left(-\frac{1}{n}, \frac{1}{n}\right) \subseteq (-1, 1)$. Conversely, let $x \in (-1, 1)$. Since $(-1, 1)$ is one of the intervals of which we are taking the union (namely the interval $\left(-\frac{1}{n}, \frac{1}{n}\right)$

when $n = 1$), we thus get that $x \in \bigcup_{n\in\mathbb{N}} \left(-\frac{1}{n}, \frac{1}{n}\right)$. Hence $(-1, 1) \subseteq \bigcup_{n\in\mathbb{N}} \left(-\frac{1}{n}, \frac{1}{n}\right)$, so we conclude that $\bigcup_{n\in\mathbb{N}} \left(-\frac{1}{n}, \frac{1}{n}\right) = (-1, 1)$.

Now, pick $0 \in \{0\}$. Since

$$-\frac{1}{n} < 0 < \frac{1}{n} \text{ for all } n \in \mathbb{N},$$

we have that $0 \in \left(-\frac{1}{n}, \frac{1}{n}\right)$ for all $n \in \mathbb{N}$. Thus $0 \in \bigcap_{n\in\mathbb{N}} \left(-\frac{1}{n}, \frac{1}{n}\right)$ by definition of intersection, so $\{0\} \subseteq \bigcap_{n\in\mathbb{N}} \left(-\frac{1}{n}, \frac{1}{n}\right)$. Conversely, let $x \in \bigcap_{n\in\mathbb{N}} \left(-\frac{1}{n}, \frac{1}{n}\right)$. We claim that $x = 0$. Since x is in $\bigcap_{n\in\mathbb{N}} \left(-\frac{1}{n}, \frac{1}{n}\right)$, we have that

$$x \in \left(-\tfrac{1}{n}, \tfrac{1}{n}\right) \text{ for all } n \in \mathbb{N}$$

by definition of intersection, so $-\frac{1}{n} < x < \frac{1}{n}$ for all $x \in \mathbb{N}$. To justify that this implies $x = 0$, we make use of some properties of limits from calculus. Namely, the limit of the left side $-\frac{1}{n}$ of given inequality as n goes to ∞ is 0, as is the limit of the right side $\frac{1}{n}$, so taking limits gives

$$0 \leq x \leq 0,$$

and hence $x = 0$ as desired. Thus $x \in \{0\}$, so $\bigcap_{n\in\mathbb{N}} \left(-\frac{1}{n}, \frac{1}{n}\right) \subseteq \{0\}$. We conclude that $\bigcap_{n\in\mathbb{N}} \left(-\frac{1}{n}, \frac{1}{n}\right)$ is $\{0\}$ as claimed.

Working towards contrapositives. The key observation in the last bit of the proof above was that the following statement is true:

> If $-\frac{1}{n} < x < \frac{1}{n}$ for all $n \in \mathbb{N}$, then $x = 0$,

which we justified in a bit of a hand-wavy way using limits. Of course, this can be made precise by more formally justifying the properties of limits we used, but this would take us beyond the scope of this course. So, we ask, is there another way to justify the implication above?

The issue is that there is no way to directly move from the given inequality to knowing precisely what x must be. To approach this, we need a way to rephrase the given implication. The key point is that we can instead ask ourselves: what if x wasn't zero? If $x \neq 0$, then if the given claim is true it should also be true that x cannot satisfy the given inequality for all $n \in \mathbb{N}$; that is, it should be true that

> If $x \neq 0$, then x does not satisfy $-\frac{1}{n} < x < \frac{1}{n}$ for all $n \in \mathbb{N}$.

Indeed, if x did satisfy this given inequality for all $n \in \mathbb{N}$, our original implication would imply that x must have been zero. This new implication is called the *contrapositive* of the original implication, and the basic fact of logic is that an implication is always logically equivalent to its contrapositive, which means that proving one is equivalent to proving the other.

As another example, consider the claim for an integer n:

> If n^2 is even, then n is even.

This is true, but trying to prove this directly leads nowhere: if n^2 is even we can write it as $n^2 = 2k$ for some $k \in \mathbb{Z}$, but now there is no way to go from this to an expression where we have n written as twice an integer; in particular, we can take square roots to get $n = \sqrt{2k}$, but we have no way of knowing whether $\sqrt{2k}$ can be written as 2ℓ for some $\ell \in \mathbb{Z}$. So, proving this claim always requires something new. The contrapositive in this case is

> If n is odd, then n^2 is odd,

which we already know to be true. Since this contrapositive is true, "If n^2 is even, then n is even" is also true.

We'll talk more about contrapositives later, but the reason for introducing them now is to point out that we're at a point where we need to better understand what it means for a statement to be *false*. In particular, going back to the first example above what exactly does it mean for x to not satisfy "$-\frac{1}{n} < x < \frac{1}{n}$ for all $n \in \mathbb{N}$"? The answer to this question requires that we be able to *negate* the statement "$-\frac{1}{n} < x < \frac{1}{n}$ for all $n \in \mathbb{N}$", and indeed working with contrapositives in general requires working with negations.

Negations. Given a statement P, the *negation* of P, which we'll denote by $\sim P$, is the statement saying what it means for P to be false; that is, negating a statement changes its truthiness/falsensess. To be clear, we can negate statements which are originally true or false, in which case the negation is false or true respectively. In the book, material on negations can be found in Chapter 2. WARNING: I think the book is way too formal here, and focuses too much on "truth tables" and other things which I don't think shed much light on how to actually think through negations. Keep this in mind as you go through the chapter.

As a start, consider the statement "For all $n \in \mathbb{N}$, $-\frac{1}{n} < x < \frac{1}{n}$". Since this statement is claiming something should be true for *all* $n \in \mathbb{N}$, showing that this is false only requires that the given condition fail for *at least one* $n \in \mathbb{N}$. In other words, in order for a "for all" statement to be false only requires that there *exist* an instance in which it is false, not that it be false in all possible instances. In our case, this means that the negation of

> For all $n \in \mathbb{N}$, $-\frac{1}{n} < x < \frac{1}{n}$

is

> There exists $n \in \mathbb{N}$ such that "$-\frac{1}{n} < x < \frac{1}{n}$" is not true.

The key point is that negating a "for all" always gives a "there exists": $\sim\forall = \exists$

Next we have to understand it means for "$-\frac{1}{n} < x < \frac{1}{n}$" to not be true. This inequality is really two inequalities in one: $-\frac{1}{n} < x$ *and* $x < \frac{1}{n}$. Since this claims that both $-\frac{1}{n} < x$ and $x < \frac{1}{n}$ are true, in order for this to be false only requires that at least *one* of the given inequalities fail; that is, "$-\frac{1}{n} < x$ *and* $x < \frac{1}{n}$" is false when "$-\frac{1}{n} \geq x$ *or* $x \geq \frac{1}{n}$". The key point here is that negating an "and" statement always gives an "or": $\sim(P \text{ and } Q) = (\sim P \text{ or } \sim Q)$.

Thus, we can finally write out the full negation of our original implication:

> The negation of "For all $n \in \mathbb{N}$, $-\frac{1}{n} < x < \frac{1}{n}$" is "There exists $n \in \mathbb{N}$ such that $-\frac{1}{n} \geq x$ or $x \geq \frac{1}{n}$".

As this is meant to suggest, negating should be a mechanical step-by-step process: we simply start on the left and negate everything we see down the way.

Example. We negate the statement: there exists $x \in \mathbb{R}$ such that $x^2 = -1$. This is of course false, meaning that its negation should be true. But regardless of whether the original statement is true or false, if it *were* to be true all we would need is an instance of a single $x \in \mathbb{R}$ satisfying $x^2 = -1$. Thus, in order for the given claim to be false would require that there is no such x, or in other words that any $x \in \mathbb{R}$ we take will not satisfy the given requirement. This means that negating this existence gives a "for all":

> The negation of "There exists $x \in \mathbb{R}$ such that $x^2 = -1$" is "For all $x \in \mathbb{R}$, x does not satisfy $x^2 = -1$", or in other words "For all $x \in \mathbb{R}$, $x^2 \neq -1$".

The key point is that negating a "there exists" gives a "for all": $\sim\exists = \forall$

Final example. We a final example, we negate the statement that:

> For all $n \in \mathbb{Z}$, there exists $k \in \mathbb{Z}$ such that $n = 2k$ or $n = 2k + 1$.

Note that this statement is true, since it just amounts to saying that any integer is either even or odd. To negate it, we start at the beginning: the statement claims that "for all $n \in \mathbb{Z}$", some property holds, so negating gives that there exists $n \in \mathbb{Z}$ such that the given property does not hold:

> There exists $n \in \mathbb{Z}$ such that $\sim$(there exists $k \in \mathbb{Z}$ such that $n = 2k$ or $n = 2k + 1$).

Next we must negate "there exists $k \in \mathbb{Z}$ such that $n = 2k$ or $n = 2k + 1$". But this is a statement saying that there is some $k \in \mathbb{Z}$ satisfying some property, so negating requires that no matter which $k \in \mathbb{Z}$ we take, the given property will not hold:

> $\sim$(there exists $k \in \mathbb{Z}$ such that $n = 2k$ or $n = 2k + 1$) = for all $k \in \mathbb{Z}$, $\sim$($n = 2k$ or $n = 2k + 1$).

Putting this into the negation we're building up gives so far:

> There exists $n \in \mathbb{Z}$ such that for all $k \in \mathbb{Z}$, $\sim$($n = 2k$ or $n = 2k + 1$).

Finally, we must negate "$n = 2k$ or $n = 2k + 1$". This statement only requires that at least one of $n = 2k$ or $n = 2k + 1$ be true, so negating requires that *both* $n = 2k$ and $n = 2k + 1$ be false; in other words, negating an "or" statement gives an "and": $\sim(P$ or $Q) = (\sim P$ and $\sim Q)$. Thus, the negation of "$n = 2k$ or $n = 2k + 1$" is "$n \neq 2k$ and $n \neq 2k + 1$".

All together then, the complete negation of

> For all $n \in \mathbb{Z}$, there exists $k \in \mathbb{Z}$ such that $n = 2k$ or $n = 2k + 1$.

is

> There exists $n \in \mathbb{Z}$ such that for all $k \in \mathbb{Z}$, $n \neq 2k$ and $n \neq 2k + 1$.

Note again that coming up with this negation was a purely a mechanical step-by-step process: we negate everything from the start on down, turning "for all"s into existences, existences into "for all"s, and's into or's and or's into and's.

Lecture 7: Contrapositives

Warm-Up 1. Recall that to say $u \in \mathbb{R}$ is an *upper bound* of $S \subseteq \mathbb{R}$ means that for all $s \in S$, $s \leq u$. Negating this gives what it means for u to *not* be an upper bound of S. When negating, "for all" becomes "there exists" and $s \leq u$ becomes $s > u$, so the negation is

> "there exists $s \in S$ such that $\sim(s \leq u)$", or "there exists $s \in S$ such that $s > u$".

Thus, to be concrete, saying that $u \in \mathbb{R}$ is not an upper bound of $S \subseteq \mathbb{R}$ means that there exists $s \in S$ such that $u < s$. We'll come back to this when we talk more about upper bounds later on.

Warm-Up 2. We negate the statement that

> There exists $n \in \mathbb{Z}$ such that for all $k \in \mathbb{Z}$, $n \leq k$.

Note that this is saying that there exists an integer which is smaller than every other integer, which we know is false since there is no such smallest integer. Thus, the negation should be true.

Performing our step-by-step negation process gives:

$$\begin{aligned}
&\sim(\text{There exists } n \in \mathbb{Z} \text{ such that for all } k \in \mathbb{Z},\ n \leq k) \\
&= \text{For all } n \in \mathbb{Z},\ \sim(\text{for all } k \in \mathbb{Z},\ n \leq k) \\
&= \text{For all } n \in \mathbb{Z}, \text{ there exists } k \in \mathbb{Z} \text{ such that } \sim(n \leq k) \\
&= \text{For all } n \in \mathbb{Z}, \text{ there exists } k \in \mathbb{Z} \text{ such that } n > k.
\end{aligned}$$

Thus, the negation of "There exists $n \in \mathbb{Z}$ such that for all $k \in \mathbb{Z}$, $n \leq k$" is "For all $n \in \mathbb{Z}$, there exists $k \in \mathbb{Z}$ such that $n > k$". In other words, there is no smallest integer since given any integer n, which we can find another k which is smaller than it.

Negating implications. As a final example, we negate the following:

For all $\epsilon > 0$, there exists $\delta > 0$ such that if $|x - 2| < \delta$, then $|f(x) - f(2)| < \epsilon$.

Here, f is some single-variable function (for instance, $f(x) = x^2$, of $f(x) = \sin x$), and in fact the statement given here is the precise definition of what it means for f to be *continuous* at 2. What is this definition actually saying and why does it capture the intuitive notion of what "continuous" should mean from a calculus course? Answering these questions is beyond the score of this course and belong instead to a course in *real analysis* such as Math 320. For us, the point is that regardless of whether we understand what this definition is saying or not we should still be able to negate it and write out what it means for a function to *not* be continuous at 2. This again emphasizes the point that negation should be a mechanical process which doesn't actually require we understand the intricacies of the statements being made.

We can form much of the negation as we've done previously, and we get:

There exists $\epsilon > 0$ such that for all $\delta > 0$, $\sim$(if $|x - 2| < \delta$, then $|f(x) - f(2)| < \epsilon$).

But now the new thing is that we have to negate "if $|x - 2| < \delta$, then $|f(x) - f(2)| < \epsilon$", which requires that we understand what it means for an implication to be false.

Let's start with a simpler example. Consider the statement: if $x > 1$, then $x > 5$. This is false, but why exactly is it false? How would you convince me or someone else that it is false? You might say "it is not true that every x satisfying $x > 1$ also satisfies $x > 5$"; this is correct, but now how would you convince me of that? To convince me you would have to produce an example of x satisfying $x > 1$ but not $x > 5$, which is easy to do; for instance $x = 3$ works. But taking a step back, what you've done in order to convince me that "if $x > 1$, then $x > 5$" is false is show that it is *possible* for $x > 1$ to be true with $x > 5$ being false, by showing the *existence* of such an x. That is, the negation of "if $x > 1$, then $x > 5$" is

there exists x such that $x > 1$ but $x \geq 5$.

In general, to show that an implication "if P, then Q" (also written symbolically as $P \Rightarrow Q$, pronounced "P implies Q") is false requires showing that P can be true with Q being false at the same time. That is, $\sim(P \Rightarrow Q) = (P \text{ and } \sim Q)$. This is a crucial point: negating an implication does NOT produce another implication, but rather produces the statement that the assumption P is true with the conclusion Q being false.

Now, where does the "there exists" at the start of "there exists x such that $x > 1$ but $x \geq 5$" come from? Recall that our original implication "if $x > 1$, then $x > 5$" really has an *implicit* "for all" at the start:

for all x, if $x > 1$, then $x > 5$.

Thus, negating should indeed give an existence, namely the existence of an x satisfying $x > 1$ but not $x > 5$. This again emphasizes the idea that $P \Rightarrow Q$ is false when it is *possible* for P to be true with Q being false; the "it is possible" says that we only require at least one instance (i.e. "there exists") where P is true but Q is false.

Going back to our continuity negation, we thus see that the negation of "if $|x - 2| < \delta$, then $|f(x) - f(2)| < \epsilon$" is "there exists x such that $|x - 2| < \delta$ but $|f(x) - f(2)| \geq \epsilon$". Thus, to say that a function f is NOT continuous at 2 concretely means that

> There exists $\epsilon > 0$ such that for all $\delta > 0$, there exists x such that $|x - 2| < \delta$ and $|f(x) - f(2)| \geq \epsilon$.

Again, understanding precisely what this means and how it captures an intuitive notion of "not continuous" is left to an analysis course.

Truth tables. We can summarize the claim that $P \Rightarrow Q$ is false precisely when P is true but Q is false by writing out the *truth table* for $P \Rightarrow Q$:

P	Q	$P \Rightarrow Q$
T	T	T
T	F	F
F	T	T
F	F	T

This tables lists the different possibilities as to whether P/Q are true or false, and gives the truthness/falseness of $P \Rightarrow Q$ in each case. The second line is the one we figured out in the discussion above: $P \Rightarrow Q$ should be false when P is true and Q is false.

Now, the third and fourth lines might seem strange at first, since they say that we consider $P \Rightarrow Q$ to be true whenever P is false *regardless* of whether Q is true or not. If we agree that, as in our discussion above, the given implication should be false *only* when you can show me that P is true and Q is false, then we are left concluding that $P \Rightarrow Q$ should be true in *all* other scenarios, as the table suggests. But, we can also find a better reason as to why we should consider $P \Rightarrow Q$ to be true whenever P is false. Consider the implication:

> If a unicorn is larger than 5, then I am 10 feet tall.

Is this true? Is this false? The point is that to convince me that this is false you would have to show that it is possible for "unicorn larger than 5" to be true while "I am 10 feet tall" is false. You cannot possibly do this, since there are no such unicorns, or indeed unicorns at all! The intuition is that this implication only claims that something should happen (me being 10 feet tall) as a *consequence* of something else (there being a unicorn larger than 5), so if that "something else" cannot possibly happen, the implication holds by default because, if there are no unicorns at all, then I am not lying if I say that "If a unicorn is larger than 5, then I am 10 feet tall". I am NOT 10 feet tall, but I am only claiming to be 10 feet tall under the assumption that there is a unicorn larger than 5, so I am not lying.

So, it makes to consider $P \Rightarrow Q$ to be true whenever P is false. This seemingly strange conclusion actually has some important consequences. For instance, we can now prove that any set whatsoever contains the empty set as a subset. The claim is that for any set $S, \emptyset \subseteq S$. To verify this requires, according to the definition of subset, that we show "if $x \in \emptyset$, then $x \in S$". But in this case the assumption $x \in \emptyset$ is never true, so we are in the scenario of an implication where the

hypothesis is false, in which case we consider the implication to be true! Thus, it is true that "if $x \in \emptyset$, then $x \in S$", so $\emptyset$ is indeed a subset of S. The point is that it is true that every element of $\emptyset$ is also an element of S, simply because there are not elements in $\emptyset$ on which to test condition! In this setting, we say that "if $x \in \emptyset$, then $x \in S$" is *vacuously true*, meaning that it is true simply because there is nothing on which to actually test it on.

Contrapositives. Given an implication $P \Rightarrow Q$, its *contrapositive* is the implication $\sim Q \Rightarrow \sim P$. As we alluded to last time, the point of contrapositives is that they give a way to rephrase various statements, coming from the fact that an implication and its contrapositive are logically equivalent. To be logically equivalent means that they imply each other, or that they are both true and false in precisely the same scenarios. This can be checked with a truth table: the truth table for $\sim Q \Rightarrow \sim P$ looks like

P	Q	$\sim Q$	$\sim P$	$\sim Q \Rightarrow \sim P$
T	T	F	F	T
T	F	T	F	F
F	T	F	T	T
F	F	T	T	T,

and the point is that, according to the final column, the contrapositive is true precisely in the same scenarios as when $P \Rightarrow Q$ is true, and is false in the same scenarios as when $P \Rightarrow Q$ is false. This guarantees that if the contrapositive is true, the original implication is true as well.

We can also reason that this is the case without making use of a truth table. Suppose that $\sim Q \Rightarrow \sim P$ is true, and we want to show that $\sim P \Rightarrow \sim Q$ is then true as well. Well, suppose P is true. If Q were false, $\sim Q$ would be true and the contrapositive we are assuming to be true would thus imply that $\sim P$ is true as well, and hence that P is false. But P is true, so this is not possible and hence Q must have been true as well, so $P \Rightarrow Q$ is true also. It might take a few times reading through this to understand what it is saying, but the upshot, as we've said, is that proving the contrapositive gives a valid way of proving an implication. In the book, contrapositives are covered in Chapter 3, although not to the full extent I think they should be covered.

Example 1. We explained last time that proving "if n^2 is even, then n is even" directly is not feasible, whereas proving the contrapositive "if n is not even, then n^2 is not even" is much simpler. Similarly, to prove that "if n^2 is odd, then n is odd" you can prove instead that "if n is even, then n^2 is even".

Example 2. Suppose $a, b \geq 0$. We show that if $a^2 < b^2$, then $a < b$. The point is that we cannot simply do this by taking square roots of both sides of $a^2 < b^2$, since doing so requires knowing that the process of taking square roots preserves inequalities, which is precisely what this problem is meant to justify! So, without using square roots, we see that it is not possible to prove from $a^2 < b^2$ directly that $a < b$.

Contrapositives to the rescue! Here is our proof:

Proof. We prove the contrapositive, which says that if $a \geq b$, then $a^2 \geq b^2$. Since $a \geq 0$, multiplying both sides of $a \geq b$ by a preserves the inequality to give $a^2 \geq ab$. But since $b \geq 0$, multiplying both sides of $a \geq b$ by b also preserves the inequality to give $ab \geq b^2$. Thus

$$a^2 \geq ab \geq b^2,$$

so $a^2 \geq b^2$ as claimed. □

Example 3. Suppose $a, b \in \mathbb{R}$. We show that if $a \leq b + \epsilon$ for all $\epsilon > 0$, then $a \leq b$. This makes sense intuitively: as ϵ varies through all possible positive numbers, $b + \epsilon$ varies through all possible numbers larger than b, so $a \leq b + \epsilon$ for all $\epsilon > 0$ is saying that a is less than or equal to all numbers larger than b; clearly, a itself should be then smaller than or equal to b, which is what we claim. However, there is no direct say to move form $a \leq b + \epsilon$ for all $\epsilon > 0$ to $a \leq b$ since we cannot easily get rid of ϵ, so we instead look at the contrapositive.

The contrapositive is: if $a > b$, then it is not true that (for all $\epsilon > 0$, $a \leq b + \epsilon$). But fortunately we know how to form this final negation (which is the whole reason why we spoke about negations in the first place), so the contrapositive becomes:

> If $a > b$, then there exists $\epsilon > 0$ such that $a > b + \epsilon$.

To prove this requires that we produce some positive ϵ satisfying $a > b + \epsilon$. Drawing a and b on a number line, with b to the left of a, suggests that any number in between should be a valid $b + \epsilon$, and we can describe such a number concretely by taking the midpoint between a and b. The distance between a and b is $a - b$, so this midpoint is $b + \frac{a-b}{2}$, and thus $\epsilon = \frac{a-b}{2}$ should satisfy our needs. In our proof we simply set ϵ to be this value and then verify that it has the property we want:

Claim. If $a \leq b + \epsilon$ for all $\epsilon > 0$, then $a \leq b$.

Proof. By way of contrapositive, suppose $a > b$. Then $a - b > 0$, so $\epsilon = \frac{1}{2}(a - b)$ is positive. For this positive number we have

$$b + \epsilon = b + \frac{1}{2}(a - b) < b + (a - b) = a$$

since $\frac{1}{2}(a - b) < a - b$. Thus $b + \epsilon < a$, which proves the contrapositive of the given claim. □

Lecture 8: More on Sets

Warm-Up 1. For any $r > 0$, let D_r again denote the open disk of radius r centered at $(0, 0)$:

$$D_r = \left\{ (x, y) \in \mathbb{R}^2 \;\middle|\; \sqrt{x^2 + y^2} < r \right\}.$$

We previously showed that the intersection of all such disks is $\{(0, 0)\}$, where the bulk of the work came down to showing that

> If $(x, y) \in D_r$ for all $r > 0$, then $(x, y) = (0, 0)$.

This is the implication needed to be able to say that the given intersection is a subset of $\{(0, 0)\}$, since "$(x, y) \in D_r$ for all $r > 0$" is precisely what it means to say that $(x, y) \in \bigcap_{r>0} D_r$. Previously we proved this implication directly, and now we give a proof by contrapositive.

The contrapositive is the statement:

> If $(x, y) \neq (0, 0)$, then there exists $r > 0$ such that $(x, y) \notin D_r$

since the negation of "$(x, y) \in D_r$ for all $r > 0$" is "there exists $r > 0$ such that $(x, y) \notin D_r$. Thus, to prove this requires that we produce a radius r so that the corresponding disk does not contain (x, y). To see which r should work, draw a non-origin point (x, y) anywhere in the xy-plane and imagine a disk which should not contain it; clearly such a disk should have radius smaller than the distance from (x, y) to the origin, meaning r should satisfy $r < \sqrt{x^2 + y^2}$. Taking r to be half this distance should thus work.

Proof. Suppose $(x, y) \neq (0, 0)$ and set $r = \frac{1}{2}\sqrt{x^2 + y^2}$. Since $(x, y) \neq (0, 0)$, at least one of x or y is nonzero, so $\sqrt{x^2 + y^2}$ is strictly positive and hence $r > 0$. Since

$$\sqrt{x^2 + y^2} \not\leq r$$

for this value of r, we have that $(x, y) \notin D_r$ as desired. □

Warm-Up 2. We show that

> If $-\frac{1}{n} < x < \frac{1}{n}$ for all $n \in \mathbb{N}$, then $x = 0$

by proving the contrapositive. This is also something we saw previously when showing that the intersection of all intervals of the form $\left(-\frac{1}{n}, \frac{1}{n}\right)$ consists only of 0, and indeed it was this claim which motivated our introduction of contrapositives a few days ago. Previously we proved the given implication using some unjustified properties of limits, but now we can give a proof which does not require anything fancy.

First, note that the given inequality $-\frac{1}{n} < x < \frac{1}{n}$ can be phrased as $|x| < \frac{1}{n}$ where $|x|$ is the absolute value of x. Thus the contrapositive is:

> If $x \neq 0$, there exists $n \in \mathbb{N}$ such that $|x| \geq \frac{1}{n}$,

which says that for any nonzero number, we can find a reciprocal of the form $\frac{1}{n}$ which is smaller than its absolute value. Here $|x|$ will be positive, so the real claim is that given any positive number there is a reciprocal $\frac{1}{n}$ smaller than it. Intuitively, this is what tells us that the reciprocals $\frac{1}{n}$ get closer and closer to 0 as n increases, and indeed this type of statement is what is needed when actually trying to prove that the sequence $\frac{1}{n}$ *converges* to 0.

So, assuming $x \neq 0$, we have to produce some $n \in \mathbb{N}$ satisfying $|x| \geq \frac{1}{n}$. To see what n might work, note that we can rearrange this inequality to get

$$n \geq \frac{1}{|x|}$$

by multiplying through by n and dividing through by $|x|$, both of which are operations which do not affect the inequality since n and $|x|$ are positive. Thus, all we need is n satisfying $n \geq \frac{1}{|x|}$, and we thus only need to make use of the fact that no matter what $\frac{1}{|x|}$, there will certainly be some positive integer larger than it since positive integers grow without restriction. Thus, we pick $n \in \mathbb{N}$ such that $n \geq \frac{1}{|x|}$ and verify that this n has the property we want. Note that our proof will not show how we came up with n in the first place, which was the result of some scratch work.

Proof. If $x \neq 0$, then $|x| > 0$. Thus $\frac{1}{|x|}$ is a positive real number, so we can pick some $n \in \mathbb{N}$ such that $n \geq \frac{1}{|x|}$. Rearranging this gives $|x| \geq \frac{1}{n}$, which is the desired result. □

Contrapositives and equivalences. Suppose we want to show that an integer x is divisible by 3 if and only if $x^2 - 1$ is not divisible by 3. This is saying that "x is divisible by 3" and "$x^2 - 1$ is not divisible by 3" are equivalent statements, and we know that prove something like this we need to prove that both sides imply each other.

The forward direction is straightforward and similar to things we've done before. If x is divisible by 3, we can write it as $x = 3k$ for some $k \in \mathbb{Z}$, in which case $x^2 - 1 = 9k^2 - 1 = 3(3k^2) - 1$. This expression is not as written in the form of an integer divisible by 3, so we would want to conclude that $x^2 - 1$ is not divisible by 3 as required. But a little care is required here: just because $x^2 - 1 = 3(3k^2) - 1$ is not as written in the form 3(integer) does not immediately rule out that it

can't be rewritten in that form nonetheless. In fact, this is not possible, but justifying why is also something we should think about. We'll leave this for now and take it for granted that something in the form 3(integer) $-$ 1 cannot be rewritten in the form 3(integer), but this is something we'd be able to prove after we talk about proofs by *contradiction*: namely, suppose $3(3k^2) - 1$ could be written as 3ℓ for some ℓ and use this to derive a contradiction.

The backwards direction states that if $x^2 - 1$ is not divisible by 3, then x is divisible by 3, but the issue is that we have no way going directly from information about $x^2 - 1$ to information about x alone, since there is no way to "solve" for x in a way which only uses integers. So, we instead look at the contrapositive, so this "if and only if" statement gives an example where we can prove on direction directly and the other by contrapositive. The contrapositive is: if x is not divisible by 3, then $x^2 - 1$ is divisible by 3. This is more in line with the types of things we've done before, where we use information about x to derivative information about x^2. The key thing we need here is to understand what it means for x to not be divisible by 3. One way to do so is to see that this means we can write x as either $3k + 1$ or $3k + 2$ for some $k \in \mathbb{Z}$, since any integer is either a multiple of 3, one more than a multiple of 3, or two more than a multiple of 3. This is good, since with a concrete expression for x (as either $3k+1$ or $3k+2$), we can work out a concrete expression for $x^2 - 1$.

Claim. Suppose $x \in \mathbb{Z}$. Then x is divisible by 3 if and only if $x^2 - 1$ is not divisible by 3.

Proof. Suppose x is divisible by 3. Then $x = 3k$ for some $k \in \mathbb{Z}$, so

$$x^2 - 1 = 9k^2 - 1 = 3(3k^2) - 1.$$

An integer which is one less than a multiple of 3 cannot be a multiple of 3 itself, so $x^2 - 1$ is not divisible by 3 as required.

For the converse direction we instead prove its contrapositive, which is: if x is not divisible by 3, then $x^2 - 1$ is divisible by 3. If x is not divisible by 3, then x is either of the form $x = 3k + 1$ or $x = 3k + 2$ for some $k \in \mathbb{Z}$. If $x = 3k + 1$ for some $k \in \mathbb{Z}$, then

$$x^2 - 1 = (3k + 1)^2 - 1 = 9k^2 + 6k = 3(3k^2 + 2k),$$

which is divisible by 3. If $x = 3k + 2$ for some $k \in \mathbb{Z}$, then

$$x^2 - 1 = (3k + 2)^2 - 1 = 9k^2 + 6k + 3 = 3(3k^2k + 1),$$

which is also divisible by 3. Thus in either case $x^2 - 1$ is divisible by 3 as claimed. □

Complements. Now that we've spoken about negations, we can consider one more basic set-theoretic construction. Suppose A and B are sets. The *complement* of B in A is the set of all things in A which are not in B. We use the notation $A - B$ for the complement of B in A, so

$$A - B = \{x \in A \mid x \notin B\}.$$

Alternatively, you might also see the notation $A \backslash B$ for this complement, although we'll stick with the notation $A - B$ the book uses. If the set A is implicitly given, we also use the notation $\overline{B}$ to denote this complement, which is the set of all things *not* in B. You can find material on complements in Section 1.6 and Chapter 8 of the book.

For instance, the complement of the set of even integers in $\mathbb{Z}$ is the set of odd integers:

$$\mathbb{Z} - \{\text{set of even integers}\} = \{\text{set of odd integers}\}.$$

Also, the notaton $\mathbb{R} - \mathbb{Q}$ denotes the set of *irrational* numbers, which are real number which are not rational. (The set $\mathbb{Q}$ of *rational* numbers consists of those real numbers which can be written as the fraction $\frac{a}{b}$ of integers with nonzero denominator.)

DeMorgan's Laws. We finish with two basic set equalities summarizing how unions and intersections behave under the operation of taking complements. Suppose A, B, C are sets. Then

$$A - (B \cup C) = (A - B) \cap (A - C) \text{ and } A - (B \cap C) = (A - B) \cup (A - C).$$

These are called *DeMorgan's Laws*, and hence state that the complement of a union is the intersection of complements, and that the complement of an intersection is the union of complements. These are quite useful equalities to know. We'll give a proof of the first one only, but you should prove the second one for practice. The proof is a basic element chase, where we show that an element of one side is in the other. The key step is in recognizing what it means to say that $x \notin B \cup C$: this is the negation of $x \in B \cup C$, so since $x \in B \cup C$ means $x \in B$ or $x \in C$, the negation thus means $x \notin B$ *and* $x \notin C$.

Proof of first equality. Let $x \in A-(B\cup C)$. Then $x \in A$ and $x \notin B\cup C$ by definition of complement. Since $x \notin B \cup C$, $x \notin B$ and $x \notin C$. Since $x \in A$ and $x \notin B$, $x \in A - B$, and since $x \in A$ and $x \notin C$, we also have $x \in A-C$. Thus we have $x \in A-B$ and $x \in A-C$, so $x \in (A-B)\cap(A-C)$. Hence $A - (B \cup C) \subseteq (A - B) \cap (A - C)$.

Conversely suppose $x \in (A-B)\cap(A-C)$. Then $x \in A-B$ and $x \in A-C$. Since $x \in A-B$, $x \in A$ and $x \notin B$, and since also $x \in A - C$, we have $x \notin C$. Since $x \notin B$ and $x \notin C$, we get $x \notin B \cup C$, so $x \in A - (B \cup C)$. Hence $(A - B) \cap (A - C) \subseteq A - (B \cup C)$, so $A - (B \cup C) = (A - B) \cap (A - C)$ as claimed. □

Lecture 9: Contradictions

Warm-Up 1. Let

$$A = \{n \in \mathbb{Z} \mid \text{there exists } k \in \mathbb{Z} \text{ such that } n = 4k + 1\}$$

and

$$B = \{n \in \mathbb{Z} \mid \text{there exists } \ell \in \mathbb{Z} \text{ such that } n = 8\ell + 1\}.$$

We show that $A - B \neq \emptyset$. This requires showing that there exists $x \in A - B$, which means there exists $n \in A$ such that $n \notin B$. The point is that all we need is the existence of one such n. We claim that 5 works. First, since $5 = 4(1) + 1$, we indeed have $5 \in A$. Now we show that $5 \notin B$. The key is to note that the only possible number ℓ satisfying

$$5 = 8\ell + 1$$

is $\ell = \frac{1}{2}$, and that $\frac{1}{2} \notin \mathbb{Z}$. Thus, it is not possible to find an *integer* ℓ satisfying $5 = 8\ell + 1$, so 5 does not satisfy the property required to belong to B, so $5 \notin B$. Thus $5 \in A - B$, so $A - B$ is nonempty.

Warm-Up 2. Suppose that we have a collection of sets S_α indexed by elements α of some other set. We show that

$$\overline{\bigcap_\alpha S_\alpha} = \bigcup_\alpha \overline{S_\alpha}.$$

To be clear of the notation, we are assuming that each S_α belongs to some unspecified larger set, and the bar notation used here indicates complements in that larger set: on the let we have the complement of the intersection of all S_α, whereas on the right we have the union of the individual complements $\overline{S_\alpha}$.

Let $x \in \overline{\bigcap_\alpha S_\alpha}$. Then $x \notin \bigcap_\alpha S_\alpha$ by the definition of complement. Now, the definition of $x \in \bigcap_\alpha S_\alpha$ is

$$\text{for all } \alpha,\ x \in S_\alpha,$$

so negating this gives what it means to say that $x \notin \bigcap_\alpha S_\alpha$:

$$\text{there exists } \alpha \text{ such that } x \notin S_\alpha.$$

Thus there exists some index β such that $x \notin S_\beta$, which means that $x \in \overline{S_\beta}$. Hence $x \in \bigcup_\alpha \overline{S_\alpha}$ since x in particular belongs to $\overline{S_\beta}$, so

$$\overline{\bigcap_\alpha S_\alpha} \subseteq \bigcup_\alpha \overline{S_\alpha}.$$

Conversely suppose $x \in \bigcup_\alpha \overline{S_\alpha}$. Then there is some index β such that $x \in \overline{S_\beta}$, which means that $x \notin S_\beta$. Hence $x \notin \bigcap_\alpha S_\alpha$ since in particular x does not belong to S_β, so $x \in \overline{\bigcap_\alpha S_\alpha}$. Thus

$$\overline{\bigcap_\alpha S_\alpha} \supseteq \bigcup_\alpha \overline{S_\alpha},$$

so equality holds as claimed.

Just as a sanity check, let us work out both sides in an explicit example to verify that we indeed get the same thing. For each $n \in \mathbb{N}$ we set $S_n = \left(\frac{1}{n}, 5 + \frac{1}{n}\right)$, which is a subset of $\mathbb{R}$. The equality proved above says that the following should be true:

$$\overline{\bigcap_{n\in\mathbb{N}} \left(\tfrac{1}{n}, 5 + \tfrac{1}{n}\right)} = \bigcup_{n\in\mathbb{N}} \overline{\left(\tfrac{1}{n}, 5 + \tfrac{1}{n}\right)},$$

where all complements are taken inside of $\mathbb{R}$. To work out the left side, we start with the intersection of which the complement is being taken. For $n = 1$ the first left endpoint of the given intervals is 1, and as n increases this endpoint moves left closer to 0. Thus only numbers larger than 1 are to the right of all such endpoints. Similarly, the right endpoints $5 + \frac{1}{n}$ start at 6 and move left closer to 5 as n increases, so only numbers smaller than or equal to 5 should be the left of all such endpoints. Thus

$$\bigcap_{n\in\mathbb{N}} \left(\tfrac{1}{n}, 5 + \tfrac{1}{n}\right) = (1, 5].$$

(Note that we can prove this precisely by showing that each side is a subset of the either side, but for this problem we'll skip these details.) Note that 5 is included in the intersection since $5 \leq 5 + \frac{1}{n}$ for all $n \in \mathbb{N}$. Thus the left side of our claimed equality is

$$\overline{\bigcap_{n\in\mathbb{N}} \left(\tfrac{1}{n}, 5 + \tfrac{1}{n}\right)} = \overline{(1, 5]} = (-\infty, 1] \cup (5, \infty).$$

Now we work out the right side. First, the complement of $\left(\frac{1}{n}, 5 + \frac{1}{n}\right)$ is:

$$\overline{\left(\tfrac{1}{n}, 5 + \tfrac{1}{n}\right)} = \left(-\infty, \tfrac{1}{n}\right] \cup \left[5 + \tfrac{1}{n}, \infty\right).$$

As we vary through all possible $n \in \mathbb{N}$ getting larger and larger, the point $\frac{1}{n}$ moves from 1 towards (but not reaching) 0, while the point $5 + \frac{1}{n}$ moves from 6 towards (but not reaching) 5. Hence, in the union, the intervals $(-\infty, \frac{1}{n}]$ gives $(-\infty, 1]$ and the intervals $\left[5 + \frac{1}{n}, \infty\right)$ give $(5, \infty)$, so

$$\bigcup_{n \in \mathbb{N}} \overline{\left(\tfrac{1}{n}, 5 + \tfrac{1}{n}\right)} = \bigcup_{n \in \mathbb{N}} \left(-\infty, \tfrac{1}{n}\right] \cup \left[5 + \tfrac{1}{n}, \infty\right) = (-\infty, 1] \cup (5, \infty),$$

which agrees with the answer we computed for the left side. Thus

$$\overline{\bigcap_{\alpha} S_\alpha} = \bigcup_{\alpha} \overline{S_\alpha}$$

is true in this example as expected.

Proof by contradiction. We're nearing the end of outlining some basic proof techniques! Next up are proofs by contradiction, which work by assuming that what you want to show is actually false and showing that this leads to an impossibility. In other words, if we can show that that *negation* of what we're trying to show cannot be true, then our original statement must in fact be true. In the book, material on proofs by contradiction can be found in Chapter 6.

The wrinkle here is that it is usually not so clear how a proof by contradiction should actually work, mainly because it is usually not so clear what the contradiction we're aiming for should actually be. As opposed to a proof by contrapositive, where we prove $P \Rightarrow Q$ by assuming Q is false and showing that P is false, in a proof by contradiction our ending point only comes about after playing around with our conditions in the hope of deriving a contradiction. Definitely, recognizing the types of contradictions one can shoot for only comes about by practice and seeing enough examples worked out.

Example 1. We justify a claim made in an example last time, that an integer of the form $3k - 1$ for $k \in \mathbb{Z}$ is not divisible by 3. Arguing by contradiction, suppose $n = 3k - 1$ for some $k \in \mathbb{Z}$ and that n *is* divisible by 3. The goal now is to show that this assumption cannot be possible by showing that it leads to a contradiction. In this case, if n is divisible by 3, we know that $n = 3\ell$ for some $\ell \in \mathbb{Z}$, so we have that

$$3k - 1 = 3\ell.$$

This, however, implies that $3(k - \ell) = 1$, which says that 1 is divisible by 3. This, then, is our sought after contradiction, so our assumption that $n = 3k - 1$ is divisible by 3 could not have been correct, so we conclude that an integer of the form $3k - 1$ for $k \in \mathbb{Z}$ is in fact not divisible by 3.

We can also phrase the contradiction above as saying that $3k - 1 = 3\ell$ implies $\frac{1}{3} = k - \ell$ is an integer, which is not true. In general, there are different possible contradictions one can obtain, but again is not always so clear what they are at the outset.

Example 2. We prove that $\sqrt{2}$ is *irrational*, which means that it is not rational. Since we are trying to prove a negative (that something is NOT rational), contradiction is a natural choice. So, for the sake of contradiction, we suppose $\sqrt{2}$ is rational, which means that it can be written as a fraction of integers:

$$\sqrt{2} = \frac{p}{q} \text{ for some } p, q \in \mathbb{Z} \text{ with } q \neq 0.$$

Again, note that at this point we don't know what the contradiction will end up being, and is something which comes about by manipulating what we have to gather new information.

From the given equality we get that $2 = \frac{p^2}{q^2}$, and hence that

$$2q^2 = p^2.$$

But this then implies that p^2 is even, and hence that p is even. Notice that this is not something we knew at the start and is something new we can work with. Since p is even we can write it as $p = 2k$ for some $k \in \mathbb{Z}$. But then $2q^2 = p^2$ becomes

$$2q^2 = 4k^2, \text{ so } q^2 = 2k^2.$$

But now this implies that q^2 is even, and hence that q is also even. Again, this is not something we knew at the start.

So, where is our contradiction? At this point we have to think about whether what we have derived so far—that p and q are both even—is possible, or whether we can rephrase what we've done in order to get a contradiction. In this case, there is no reason that p and q could not both be even, *except* if we go back and make one more crucial observation: any rational number can be written as a fraction of integers in *reduced* form, meaning integers with no common factors. If we had made this assumption about p and q at the start, we would now have a contradiction since p and q being even would imply that they were both divisible by 2. This then will be our contradiction; it was clear at the start that this is what the contradiction would be, but is something we figured out in the course of playing around with what we had. Here is a final proof:

Claim. $\sqrt{2}$ is irrational.

Proof. Aiming for a contradiction, suppose $\sqrt{2}$ is rational. Then there exist $p, q \in \mathbb{Z}$ with no common factors and $q \neq 0$ such that

$$\sqrt{2} = \frac{p}{q}.$$

This gives $2 = \frac{p^2}{q^2}$, so $2q^2 = p^2$. Hence p^2 is even, so p is even. Thus we can write $p = 2k$ for some $k \in \mathbb{Z}$, in this case $2q^2 = p^2$ becomes

$$2q^2 = 4k^2, \text{ so } q^2 = 2k^2.$$

Thus q^2 is even, so q is even, contradicting the assumption that p and q had no common factors. Thus we conclude that $\sqrt{2}$ is irrational as claimed. □

Lecture 10: More on Upper Bounds

Warm-Up 1. We show that $x^2 = 4y + 3$ has no integer solutions, or in other words that there do not exist integers x and y such that $x^2 = 4y + 3$. We do so by contradiction, so suppose there are integers x, y satisfying $x^2 = 4y + 3$. Since $4y + 3$ is odd, x^2 is also odd and hence x is odd. Thus for some $k \in \mathbb{Z}$ we have $x = 2k + 1$. Thus

$$(2k + 1)^2 = 4y + 3, \text{ so } 4k^2 + 4k + 1 = 4y + 3.$$

This gives $2 = 4(k^2 + k - y)$, which implies that 2 is divisible by 4, which is a contradiction. Hence no such x and y exist.

To emphasize, what the contradiction was going to be was not obvious at the start, and only came about by working with what we had, deriving consequences along the way until we saw

something that couldn't possibly be true. In this case, the thing to consider is what further information about x or y can we derive from knowing that $x^2 = 4y + 3$.

Warm-Up 2. We show that if a, b, c are integers satisfying $a^2 + b^2 = c^2$, then a or b is even. (Just to put this fact into the right context: a, b, c satisfying $a^2 + b^2 = c^2$ form the sides of a right triangle—at least if a, b, c are all positive—and so this fact says that in any right triangle with sides of integer length, one side must have even length. This is useful in deriving formulas for all possible integer values of a, b, c which can form the sides of a right triangle.)

Aiming for a contradiction, suppose a and b are both odd. Then $a = 2k + 1$ and $b = 2\ell + 1$ for some $k, \ell \in \mathbb{Z}$, so

$$c^2 = a^2 + b^2 = (2k + 1)^2 + (2\ell + 1)^2 = 4(k^2 + k + \ell^2 + \ell) + 2.$$

Thus c^2 is even, so c is as well. Hence there exists $p \in \mathbb{Z}$ such that $c = 2p$, so

$$4p^2 = 4(k^2 + k + \ell^2 + \ell) + 2.$$

But this implies that 2 is divisible by 4 since $2 = 4(p^2 - k^2 - k - \ell^2 - \ell)$, which is a contradiction. Thus at least one of a or b must be even as claimed.

Back to upper bounds. Now that we've spent time talking about negations and applications, we can go back to learn more about supremums. Recall that previously we showed that $\sup[0, 2] = 2$ using only the definition of supremum, but that the same technique we used there doesn't apply to show that $\sup(0, 2) = 2$: in this latter case, 2 is not in the set $(0, 2)$, so we cannot directly say that another upper bound u of $(0, 2)$ is larger than 2 only from knowing that u is larger than everything in $(0, 2)$.

So, we needed another way to argue that 2 was the supremum of $(0, 2)$. The idea we mentioned there, and which we can now make precise, is to show that nothing *smaller* than 2 can be an upper bound of $(0, 2)$. Indeed, 2 is an upper bound of $(0, 2)$ and if it is true that nothing smaller than 2 can be an upper bound of $(0, 2)$, it makes sense intuitively that 2 should be the smallest upper bound. We make the follow claim:

Alternate characterization of supremum. Suppose b is an upper bound of a subset S of $\mathbb{R}$. Then $b = \sup S$ if and only if for all $\epsilon > 0$, there exists $s \in S$ such that $b - \epsilon < s$.

Let us digest this. The statement "there exists $s \in S$ such that $b - \epsilon < s$" is precisely the negation of what it means to say that $b - \epsilon$ *is* an upper bound of S; i.e. negating "for all $s \in S$, $s \leq b - \epsilon$" gives "there exists $s \in S$ such that $s > b - \epsilon$". Thus "there exists $s \in S$ such that $b - \epsilon < s$" says that $b - \epsilon$ is not an upper bound of S. But if $\epsilon > 0$, $b - \epsilon < b$, so by considering *all* possible positive ϵ, we are also considering *all* possible number $b - \epsilon$ *smaller* than b. Thus, "for all $\epsilon > 0$, there exists $s \in S$ such that $b - \epsilon < s$" is indeed a precise way of saying that nothing smaller than b can be an upper bound of S, which should intuitively mean that b is the least upper bound. We can now prove that our intuition is correct using the techniques we've built up.

For the forward direction we suppose that $b = \sup S$ and show that for any $\epsilon > 0$, there exists $s \in S$ such that $b - \epsilon < s$. But this comes from what we said above: if $\epsilon > 0$, $b - \epsilon < b$, so $b - \epsilon$ cannot be an upper bound of S, meaning that such $s \in S$ exists. For the backwards direction we can argue by contrapositive: if $b \neq \sup S$, then there exists $\epsilon > 0$ such that for all $s \in S$, $b - \epsilon \geq s$. (Note that in class we phrased this direction as a proof by contradiction, but I think it's cleaner to phrase it as a contrapositive argument instead.)

How do we get this? Assuming $b \neq \sup S$, we have to produce some positive ϵ satisfying $s \leq b - \epsilon$ for all $s \in S$. But this latter condition would precisely say that $b - \epsilon$ is an upper bound of S, so the question is whether there is an upper bound of S which is smaller than b. This is true: if $b \neq \sup S$, b is not the smallest upper bound of S, so there must be a smaller upper bound u of S. Then we ask whether we can write u as $b - \epsilon$ for some choice of ϵ, and $\epsilon = b - u$ works. Here, then, is our proof:

Proof of alternate characterization of supremum. Suppose $b = \sup S$ and let $\epsilon > 0$. Then $b - \epsilon < b$, so since b is the least upper bound of S, $b - \epsilon$ cannot be an upper bound of S. Hence by negating the definition of upper bound we get that there exists $s \in S$ such that $b - \epsilon < s$ as claimed.

To prove the converse we argue by way of contrapositive. Thus suppose $b \neq \sup S$. Then b is not the smallest upper bound of S, so there exists an upper bound u of S which is smaller than b. Then $\epsilon = b - u$ is positive and $u = b - \epsilon$. Since u is an upper bound of S, we know that $s \leq u$ for all $s \in S$, so $s \leq b - \epsilon$ for all $s \in S$, which demonstrates the contrapositive. □

Example. We can now finally justify the fact that $\sup(0, 2) = 2$. First, any $x \in (0, 2)$ satisfies $x < 2$ by definition of the interval $(0, 2)$, so 2 is definitely an upper bound of $(0, 2)$. To show that 2 is the least upper bound we use the alternate characterization of supremums derived above. That is, for $\epsilon > 0$, we must show that there exists $s \in (0, 2)$ such that $2 - \epsilon < s$. Where does this s come from? Visually on a number line, s should be to the right of $2 - \epsilon$ and to the left of 2 (since s is still in $(0, 2)$)

$$2 - \epsilon < s < 2.$$

Thus taking s to be the midpoint between $2 - \epsilon$ and 2 works for example, and this midpoint is achieved by taking $s = 2 - \frac{\epsilon}{2}$. This is then a value of s satisfying the inequality above.

However, there is one thing to be careful about: we need s to be in the interval $(0, 2)$, and simply setting $s = 2 - \frac{\epsilon}{2}$ might not work if $\epsilon > 0$ is too large. In other words, it could be that this value of s falls to the *left* of 0, in which case it would not be in the interval $(0, 2)$. One way to get around this is by considering cases: if $0 < \epsilon < 4$, $2 - \frac{\epsilon}{2}$ is larger than 0, so the approach above works; while if $4 \leq \epsilon$, then $2 - \epsilon < 0$, in which case *anything* in $(0, 2)$ will be a valid number larger than $2 - \epsilon$. We can also avoid cases by simply saying that no matter what $\epsilon > 0$ is, we will take s to be whichever of $2 - \frac{\epsilon}{2}$ or 1 is larger; this larger value will then certainly be something in $(0, 2)$ which is larger than $2 - \epsilon$. Here is our final proof:

Proof that the supremum of $(0, 2)$ is 2. For any $x \in (0, 2)$, we have $0 < x < 2$, so 2 is an upper bound of $(0, 2)$. Now let $\epsilon > 0$ and set s to be whichever of $2 - \frac{\epsilon}{2}$ or 1 is larger. Then

$$2 - \epsilon < 2 - \frac{\epsilon}{2} \leq s$$

and

$$0 < 1 \leq s < 2,$$

so s is an element of $(0, 2)$ which is larger than $2 - \epsilon$. By the alternate characterization of supremums we conclude that $\sup(0, 2) = 2$ as claimed. □

Lecture 11: More on Real Numbers

Warm-Up 1. We show that $\sup\left\{1 - \frac{1}{n} \;\middle|\; n \in \mathbb{N}\right\} = 1$. Intuitively, the idea is that as n increases the numbers $1 - \frac{1}{n}$ approach 1 from the left, so 1 should be the smallest upper bound. We can make this precise using the alternate characterization of supremums we derived last time. It is

simple enough to argue that 1 is an upper bound of the given set simply because $1 - \frac{1}{n}$ takes 1 and subtracts something positive, so we get something smaller than 1.

To say that 1 is the supremum, let $\epsilon > 0$. We must show there exists an element s of the given set such that $1 - \epsilon < s$. Now, elements of the given set look like $s = 1 - \frac{1}{n}$, so we need to show there exists $n \in \mathbb{N}$ satisfying

$$1 - \epsilon < 1 - \frac{1}{n}.$$

How do we get such an n? The point, as is often the case, is to work backwards from what it is we want to show to see how we can arrive at that point. In our case, the inequality we want to end up with can be manipulated to

$$\frac{1}{n} < \epsilon,$$

so the question becomes one of finding $n \in \mathbb{N}$ which satisfies this instead. But this inequality in turn can be written as

$$\frac{1}{\epsilon} < n,$$

and now we're in business: certainly no matter what number $\frac{1}{\epsilon}$ is, we can find a large enough positive integer n larger than it. This is the choice of $n \in \mathbb{N}$ then that will lead us to what we want. Note that the proof below starts by specifying the $n \in \mathbb{N}$ we want (i.e. one satisfying $\frac{1}{\epsilon} < n$), but doesn't indicate the thought process which led us to recognize that that was the correct choice of n to consider; this thought process came about via some scratch work were we manipulated what we wanted to end up with.

Proof that supremum of the given set is 1. For any $n \in \mathbb{N}$, $1 - \frac{1}{n} \leq 1$, so 1 is an upper bound of the given set $\left\{1 - \frac{1}{n} \;\middle|\; n \in \mathbb{N}\right\}$. Now, let $\epsilon > 0$ and pick $n \in \mathbb{N}$ such that

$$\frac{1}{\epsilon} < n.$$

Then $\frac{1}{n} < \epsilon$, so

$$1 - \epsilon < 1 - \frac{1}{n}.$$

Hence $1 - \frac{1}{n}$ is an element of the given set which is larger than $1 - \epsilon$, showing that $1 - \epsilon$ is not an upper bound of the given set. Since $\epsilon > 0$ was arbitrary, nothing smaller than 1 is hence an upper bound of the given set, so $\sup\left\{1 - \frac{1}{n} \;\middle|\; n \in \mathbb{N}\right\} = 1$ as claimed. □

Archimedean Property of $\mathbb{R}$. The fact we used above—no matter what $\frac{1}{\epsilon}$ is we can find $n \in \mathbb{N}$ which is larger—is useful enough that it goes by its own name: the *Archimedean Property* of $\mathbb{R}$. To be precise, the Archimedean Property states that:

> For any $x \in \mathbb{R}$, there exists $n \in \mathbb{N}$ such that $x < n$.

This should seem intuitively clear, stemming from the fact that there is no restriction on how larger a positive integer can be. As described above, by rewriting this inequality in the case where x is positive, you can show that the Archimedean Property is equivalent to the following:

> For any $\epsilon > 0$, there exists $n \in \mathbb{N}$ such that $\frac{1}{n} < \epsilon$.

Again, the equivalence comes from thinking of $\frac{1}{n} < \epsilon$ as $\frac{1}{\epsilon} < n$, and it is a good exercise (which I'll leave to you) to show that these two statements are indeed equivalent: "For any $x \in \mathbb{R}$, there exists $n \in \mathbb{N}$ such that $x < n$" if and only if "For any $\epsilon > 0$, there exists $n \in \mathbb{N}$ such that $\frac{1}{n} < \epsilon$".

The Archimedean Property of $\mathbb{R}$ (either version) is very useful when wanting to make estimates involving real numbers, and is something which will show up repeatedly in analysis course. The intuitive idea is that it guarantees that reciprocals of the form $\frac{1}{n}$ with $n \in \mathbb{N}$ can be made arbitrary small: no matter how small $\epsilon > 0$ is, we can find $\frac{1}{n}$ smaller than it. In our course, this property isn't going to be so important, apart from the fact that we'll now use it to prove some things. For us, the emphasis is on understanding how to work with something like the Archimedean Property to justify whatever it is we're trying to justify.

Example. We prove that

$$\bigcup_{n \in \mathbb{N}} \left(\tfrac{1}{n}, n\right) = (0, \infty).$$

The forward containment is fairly straightforward, simply because each interval $\left(\frac{1}{n}, n\right)$ only consists of positive numbers and hence is a subset of $(0, \infty)$. The backwards containment is more interesting. If $x \in (0, \infty)$, we need to show that $x \in \bigcup_{n \in \mathbb{N}} \left(\frac{1}{n}, n\right)$. To do so, we need to show there exists $n \in \mathbb{N}$ such that $x \in \left(\frac{1}{n}, n\right)$, or in other words that there exists $n \in \mathbb{N}$ such that

$$\frac{1}{n} < x < n.$$

How do we get such an n? It makes sense intuitively that such an n should exist, since as n increases the fractions $\frac{1}{n}$ move further to the left approaching zero while the numbers n move further to the right, so that at some point $x > 0$ will be between $\frac{1}{n}$ and n. To make this precise, think of the given inequality as two inequalities in one:

$$\frac{1}{n} < x \text{ and } x < n.$$

The point is that these look like the types of things we get from the Archimedean Property. Indeed, the Archimedean Property gives us the existence of $n \in \mathbb{N}$ such $\frac{1}{n} < x$ *and* the existence of $m \in \mathbb{N}$ such that $x < m$, so that:

$$\frac{1}{n} < x < m.$$

This is *almost* the inequality we want, except that there's no reason so far why n and m have to be the same, as we would like them to be. There are various ways around this, by using cases for instance, but instead we can simply take whichever of n, m is larger; this larger value will be larger than or equal to m and will have reciprocal which is smaller than or equal to $\frac{1}{n}$, so it will be the type of positive integer we need.

Proof of claimed equality. Let $x \in \bigcup_{n \in \mathbb{N}} \left(\frac{1}{n}, n\right)$. Then there exists $n \in \mathbb{N}$ such that $x \in \left(\frac{1}{n}, n\right)$, so

$$0 < \frac{1}{n} < x.$$

Thus $x \in (0, \infty)$, so $\bigcup_{n \in \mathbb{N}} \left(\frac{1}{n}, n\right) \subseteq (0, \infty)$. Conversely, let $x \in (0, \infty)$, so that $x > 0$. By the Archimedean Property of $\mathbb{R}$ there exists $k \in \mathbb{N}$ such that $\frac{1}{k} < x$ and there exists $m \in \mathbb{N}$ such that $x < m$. Then for $n = \max\{k, m\} \in \mathbb{N}$ we have

$$\frac{1}{n} \leq \frac{1}{k} < x < m \leq n,$$

so $x \in \left(\frac{1}{n}, n\right)$. Hence $x \in \bigcup_{n \in \mathbb{N}} \left(\frac{1}{n}, n\right)$, so $(0, \infty) \subseteq \bigcup_{n \in \mathbb{N}} \left(\frac{1}{n}, n\right)$. Thus equality holds as claimed. $\square$

How $\mathbb{Q}$ behaves inside $\mathbb{R}$. Suppose we want to show that $\sup\{x \in \mathbb{Q} \mid x < 3\} = 3$, which is true. Recall that $\mathbb{Q}$ here denotes the set of rational numbers, which are real numbers which can be written as a fraction of integers. The idea here is that we can still find rational numbers which are arbitrarily close to 3 from the left, so the supremum of all rationals less than 3 should still be 3.

But how do we show this? If we follow some of the techniques we've gone through for similar questions, we might take $\epsilon > 0$ and find an element of the given set larger than $3 - \epsilon$. Previously we've taken something like $3 - \frac{\epsilon}{2}$, but now this won't work since we don't know that $3 - \frac{\epsilon}{2}$ will be rational, as we need it to be in order to belong to the given set. What we actually need to know is that there is some rational between $3 - \epsilon$ and 3:

$$\text{some } r \in \mathbb{Q} \text{ such that } 3 - \epsilon < r < 3.$$

The fact that such a rational exists is important enough that we'll call it a Theorem:

Theorem ($\mathbb{Q}$ is dense in $\mathbb{R}$). Suppose $x, y \in \mathbb{R}$ and $x < y$. Then there exists $r \in \mathbb{Q}$ such that $x < r < y$. (This statement is what means to say that $\mathbb{Q}$ is *dense* in $\mathbb{R}$.)

So, this result says that no matter what real numbers we take, even ones which are incredibly close to one another, there will always exist a rational between them. This says that the rationals can be found everywhere we look throughout all of $\mathbb{R}$: no matter what interval we take, there will be rationals inside. This fact is not so crucial for this course, but what is important for us is thinking about how one might prove this.

To prove this we need to come up with a rational between the given x and y. A rational is a fraction of integers, so what we really need is to show there exist $a, b \in \mathbb{Z}$ (with $b \neq 0$) such that

$$x < \frac{a}{b} < y.$$

We can also assume that $b > 0$, since any negative can always be absorbed into the numerator a. So, seeing that this is what we want to end up with, we think about how to get to this point. We can rewrite this inequality (using the fact that $b > 0$ as

$$bx < a < by.$$

Now we think about a way to guarantee that there will be an integer a between bx and by; in other words, what kind of property of the numbers bx and by could possibly guarantee that there is an integer between them? One way to guarantee the existence of such an integer is between is knowing that bx and by are *far enough* apart from one another, say is bx and by were further than 100 from one another. If $by - bx > 100$, for sure there will be an integer between them. (Of course, $by - bx > 1$ would also guarantee that there is an integer between bx and by, but I'm using 100 just to show something that works, not necessarily the most efficient choice.)

So, if know that we can make $by - bx$ larger than 100, we're good to go. Now the question is: do we know there is a $b \in \mathbb{N}$ for which $by - bx > 100$? After rewriting this as $b > \frac{100}{y-x}$, we now see tha the Archimedean Property comes to the rescue; this will give us the existence of the b we need, which unwinding what we did will in the end give us the rational we are looking for. This is not an easy argument to come up with on your own, and is again not so relevant for what we'll do going forward, but is a nice example of the thought process which goes into coming up with proofs.

Proof that $\mathbb{Q}$ is dense in $\mathbb{R}$. Suppose $x < y$. Then $y - x > 0$, so there exists $b \in \mathbb{N}$ such that

$$b > \frac{100}{y - x}$$

by the Archimedean Property of $\mathbb{R}$. Then $by - bx > 100$, so there must exist an integer a between bx and by:

$$bx < a < by.$$

Since $b > 0$, this gives $x < \frac{a}{b} < y$, so there exists a rational $r = \frac{a}{b}$ between x and y as claimed. $\square$

Corollary. To finish up we derive one consequence (or "corollary") of the fact that $\mathbb{Q}$ is dense in $\mathbb{R}$, namely the fact that the irrationals $\mathbb{R} - \mathbb{Q}$ are also dense in $\mathbb{R}$. The claim is that if $x < y$, there exists an irrational d such that $x < d < y$. The trick is to turn this into something to which the fact that $\mathbb{Q}$ is dense in $\mathbb{R}$ applies, and then see what we get.

Since $x < y$, $x - \sqrt{2} < y - \sqrt{2}$. Since $\mathbb{Q}$ is dense in $\mathbb{R}$, there exists $r \in \mathbb{Q}$ such that

$$x - \sqrt{2} < r < y - \sqrt{2}.$$

Then $x < r + \sqrt{2} < y$. If $r + \sqrt{2} = s$ was rational, then $\sqrt{2} = s - r$ would be rational as well, but it is not. Hence $r + \sqrt{2}$ is irrational, so we have an irrational between x and y as desired.

Lecture 12: Induction

Warm-Up. We prove that

$$\bigcap_{n \in \mathbb{N}} \left(\tfrac{1}{n}, 5 + \tfrac{1}{n}\right) = (1, 5].$$

This is an equality we used in a previous example, only then we didn't actually prove it.

Let $x \in (1, 5]$. Then for any $n \in \mathbb{N}$, since $n \geq 1$ we have:

$$\frac{1}{n} \leq 1 < x \leq 5 < 5 + \frac{1}{n},$$

so $x \in \left(\frac{1}{n}, 5 + \frac{1}{n}\right)$ for all $n \in \mathbb{N}$. Thus $x \in \bigcap_{n \in \mathbb{N}} \left(\frac{1}{n}, 5 + \frac{1}{n}\right)$, so

$$\bigcap_{n \in \mathbb{N}} \left(\tfrac{1}{n}, 5 + \tfrac{1}{n}\right) \supseteq (1, 5].$$

Now let $x \in \bigcap_{n \in \mathbb{N}} \left(\frac{1}{n}, 5 + \frac{1}{n}\right)$, so that $x \in \left(\frac{1}{n}, 5 + \frac{1}{n}\right)$ for all $n \in \mathbb{N}$. In order to show $x \in (1, 5]$, we must know that $1 < x$ and $x \leq 5$. Since $x \in \left(\frac{1}{n}, 5 + \frac{1}{n}\right)$ for all $n \in \mathbb{N}$, in particular for $n = 1$ we get that $x \in (1, 6)$, so $1 < x$ as desired.

Now, we also know that $x < 5 + \frac{1}{n}$ for all $n \in \mathbb{N}$. To prove that this implies $x \leq 5$, we instead prove the contrapositive: if $x > 5$, then there exists $n \in \mathbb{N}$ such that $x \geq 5 + \frac{1}{n}$. If $x > 5$, $x - 5 > 0$, so by the Archimedean Property of $\mathbb{R}$ there exists $n \in \mathbb{N}$ such that

$$x - 5 > \frac{1}{n}.$$

This n thus satisfies $x \geq 5 + \frac{1}{n}$, as required in the contrapositive. Thus since $x < 5 + \frac{1}{n}$ for all $n \in \mathbb{N}$, we conclude that $x \leq 5$, so $x \in (1, 5]$. Hence

$$\bigcap_{n \in \mathbb{N}} \left(\tfrac{1}{n}, 5 + \tfrac{1}{n}\right) \subseteq (1, 5],$$

so equality holds as claimed.

Induction. The final proof technique we will consider is *induction*, the point of which is to use what we know about one scenario to build up to a "larger" scenario. Induction is often sold as a way to prove statements claiming something will hold for all $n \in \mathbb{N}$ (or perhaps for all positive integers n past some starting point which might not be 1), and indeed many examples we'll come across are of this type. However, induction is more versatile than this, since not all claims which are susceptible to induction are necessarily phrased as "for all $n \in \mathbb{N}$, such and such is true". Often times, we also use induction to show that some *process* can be carried out indefinitely, and in such problems it is not always clear where the induction actually comes in. The common feature to all these problems is, again, how to use what we know about one scenario to build up to the next.

The book covers induction in Chapter 10. In fact, the book covers slightly more than what we will, in particular *strong induction* and *proof by smallest counterexample*. These are nice techniques to know, but in the end are nothing but a reworking of the ordinary induction technique we will describe, which is more important. Our goal is not to cover every possible type of induction problem you could ever come across, but rather to highlight the underlying them all such problems have in common, and for this ordinary induction is enough.

Example. We'll describe induction by working through the following example. Suppose $x \neq 1$. We claim that for any $n \in \mathbb{N}$, we have

$$1 + x + x^2 + \cdots + x^n = \frac{1 - x^{n+1}}{1 - x}.$$

This is an equality you *might* have seen in a calculus course when discussing series, since it gives the resulting of adding all powers of x up to some specified power. To get a sense for why this is true, think about writing it instead as:

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

If you multiply out the left-hand side you get something like:

$$1 + x + x^2 + \cdots + x^n - x - x^2 - x^3 - \cdots - x^{n+1},$$

where the point is that everything cancels out expect for the initial 1 and final $-x^{n+1}$ as claimed. However, this is not actually a proof due to the vague "$\cdots$" which appear, which are meant to say "keep doing the same thing until the end", but which cannot actually do in practice for arbitrary n; essentially, in this case induction is a way to make these "$\cdots$" precise.

So, we start by checking the so-called *base case*, which is the first instance in which our claim is meant to hold: in this example, when $n = 1$. The base case is thus the claim that

$$1 + x = \frac{1 - x^2}{1 - x}$$

is true, which we can see it is after factoring the numerator $1 - x^2$ as $(1 - x)(1 + x)$.

The idea is that now we could use the base case to build up to the $n = 2$ case, which we can then use to build up to the $n = 3$ case, which builds up to the $n = 4$ case, and so on. To do this in a general way, we suppose our claim is true for *some* n, and use that information to show that it will then be true for $n + 1$ as a consequence; in other words, we use what we know about one scenario (the case of some n) to build up to a larger scenario (the case of $n + 1$). If we know that claim being true for some n implies it is true for $n + 1$, *and* we know that our base case is true, we are done: applying "n case implies $n + 1$ case" to the base case $n = 1$ will give that $n = 2$ case is

true, then applying "n case implies $n+1$ case" to $n = 2$ gives that the $n = 3$ case is true, and so on. Induction is often described via an analogy with dominoes: if we know that one domino falling forces the next one to fall as well (i.e. being true for some positive integer implies true for the next positive integer), then making the first domino fall (i.e. the base case) will make all dominoes fall.

Coming back to our claim, suppose the claimed equality true holds for some $n \in \mathbb{N}$; that is, suppose

$$1 + x + x^2 + \cdots + x^n = \frac{1 - x^{n+1}}{1 - x}.$$

(This is called the *induction hypothesis*.) The goal is to use this information to show that the claimed equality holds for $n + 1$:

$$1 + x + x^2 + \cdots + x^{n+1} = \frac{1 - x^{n+2}}{1 - x}.$$

The key for all induction applications is in figuring out how to find the case to which the induction hypothesis applies hiding within the case we are trying to derive as a consequence. Here, we note that the portion of the left-hand side occurring before the x^{n+1} term is precisely the left side of our induction hypothesis, which we are assuming equals $\frac{1-x^{n+1}}{1-x}$. We have:

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

where we use the induction hypothesis in the final step. All that remains is to rewrite the final expression to show that it equals $\frac{1-x^{n+2}}{1-x}$. Thus, knowing that the claimed equality holds for *some* n implies that it holds for the next value $n+1$ as well, and this together with the base case implies that it holds for *all* $n \in \mathbb{N}$.

Here is a final proof:

Proof. Since

$$1 + x = \frac{(1 - x)(1 + x)}{1 - x} = \frac{1 - x^2}{1 - x},$$

the claimed equality holds for the base case $n = 1$. Suppose now that the claimed equality

$$1 + x + x^2 + \cdots + x^n = \frac{1 - x^{n+1}}{1 - x}$$

holds for some $n \in \mathbb{N}$. Then:

$$\begin{aligned} 1 + x + x^2 + \cdots + x^{n+1} &= (1 + x + x^2 + \cdots + x^n) + x^{n+1} \\ &= \frac{1 - x^{n+1}}{1 - x} + x^{n+1} \\ &= \frac{(1 - x^{n+1}) + x^{n+1}(1 - x)}{1 - x} \\ &= \frac{1 - x^{n+2}}{1 - x}, \end{aligned}$$

which gives the claimed equality for $n + 1$. Hence by induction we conclude that

$$1 + x + x^2 + \cdots + x^n = \frac{1 - x^{n+1}}{1 - x}$$

for all $n \in \mathbb{N}$ as claimed. □

Why does induction work? The reason why induction works was alluded to above via a "domino" effect: applying the induction step to the base case $n = 1$ gives the $n = 2$ case, then applying the induction step to $n = 2$ gives the $n = 3$ case, and so on. This should be a convincing argument that the given claim is indeed then true for all n, but we can be a little more precise. (We didn't discuss this in class.)

Say that $P(n)$ is some statement which depends on $n \in \mathbb{N}$. (In the example above, $P(n)$ would be the claim that $1 + x + x^2 + \cdots + x^n = \frac{1-x^{n+1}}{1-x}$.) The statement of induction is the following:

> Suppose $P(1)$ is true and that $P(n)$ being true implies that $P(n+1)$ is true. Then $P(n)$ is true for all $n \geq 1$.

This statement can be proved as follows. Aiming for a contradiction, suppose it is not the case that $P(n)$ is true for all $n \geq 1$. Let $m \geq 1$ denote the smallest positive integer for which $P(m)$ is false. Since $P(1)$ is true, m must be larger than 1, and hence $m - 1 \geq 1$. Since m is the smallest positive integer for which $P(m)$ is false, it must be the case that $P(m-1)$ is true since otherwise $m - 1$ would be a positive integer smaller than m which made $P(n)$ false. But by our assumption, if $P(m-1)$ is true then $P((m-1)+1) = P(m)$ is true, which contradicts the choice of m as the smallest m for which $P(m)$ is false. We thus conclude that $P(n)$ is true for all $n \geq 1$ as claimed.

Another example. Consider the numbers x_n defined by setting $x_1 = \sqrt{2}$ and then recursively declaring

$$x_{n+1} = \sqrt{2 + x_n} \text{ for } n \geq 1.$$

Thus, x_2 is defined to be $\sqrt{2 + x_1} = \sqrt{2 + \sqrt{2}}$, x_3 is defined to be

$$\sqrt{2 + x_2} = \sqrt{2 + \sqrt{2 + \sqrt{2}}},$$

and so on. In general, the expression for x_n will be something involving n nested square roots of 2.

We claim first that all of the resulting numbers are less than 2 (i.e. $x_n < 2$ for all $n \in \mathbb{N}$), which we prove using induction. For this, we will need to determine how the claimed inequality being true for some x_n implies that it will be true for the next x_{n+1} as a result. The base case is simple: $x_1 = \sqrt{2}$ is definitely less than 2.

Now, suppose $x_n < 2$ for some $n \in \mathbb{N}$. We want to show that this implies $x_{n+1} < 2$. Thus, we need a way of relating the case we want to the case we know, which here comes from the fact that we can explicitly express x_{n+1} in terms of x_n: $x_{n+1} = \sqrt{2 + x_n}$. Our assumption that $x_n < 2$ then gives

$$x_{n+1} = \sqrt{2 + x_n} < \sqrt{2 + 2} = 2,$$

which is the desired inequality. Together with the base case, induction then gives that $x_n < 2$ holds for all $n \in \mathbb{N}$.

Here is another property of the numbers x_n we can prove using induction, the fact that these numbers are *increasing*, meaning that each is larger than the one which came before: $x_n < x_{n+1}$ for all $n \in \mathbb{N}$. Again the base case is straightforward: $x_1 = \sqrt{2} < \sqrt{2 + \sqrt{2}} = x_2$ since $2 < 2 + \sqrt{2}$. For our induction step we then suppose that x_n is smaller than the term coming after (i.e. $x_n < x_{n+1}$) and use this to show that x_{n+1} is smaller than the term coming after (i.e. $x_{n+1} < x_{n+2}$). Again, to do so we need a way of relating the term x_{n+1} we want to know something about to the term x_n we are assuming something about, so we use the fact that $x_{n+1} = \sqrt{2 + x_n}$. If $x_n < x_{n+1}$, then

$$x_{n+1} = \sqrt{2 + x_n} < \sqrt{2 + x_{n+1}} = x_{n+2}$$

as desired. Thus, x_n being smaller than the next term implies that x_{n+1} is smaller than the next term, so together with the base case induction gives that $x_n < x_{n+1}$ for all $n \in \mathbb{N}$.

Final example. In class we started one final example dealing with conjugation of complex numbers, which we didn't finish until the Warm-Up of the next lecture. So, in order to keep that example all in one place, I'll put the whole thing in the subsequent lecture's write up.

Lecture 13: More on Induction

Warm-Up. First we recall (or introduce if you've never seen these before) some material regarding complex numbers. A *complex number* is an expression $a + ib$ where $a, b \in \mathbb{R}$ and i is taken to be a "number" satisfying $i^2 = -1$. We multiply complex numbers using regular distributive properties and using the fact that $i^2 = -1$ in order to simply expressions:

$$(a + ib)(c + id) = ac + iad + ibc + i^2bd = ac + iad + ibc - bd = (ac - bd) + i(ad + bc).$$

Given a complex number $z = a + ib$, its *conjugate* is the complex number $\overline{z} = a - ib$ obtained by changing the sign of the "imaginary" part ib.

We claim that for any $n \geq 2$ complex numbers $z_1, \ldots, z_n$, we have

$$\overline{z_1 z_2 \cdots z_n} = \overline{z_1}\,\overline{z_2} \cdots \overline{z_n},$$

which says that the conjugate of a product of complex numbers is always the same as the product of the individual conjugates. We prove this by induction. The base case, that

$$\overline{z_1 z_2} = \overline{z_1}\,\overline{z_2},$$

involves simply working out both sides and seeing that we get the same result. Setting $z_1 = a_1 + ib_1$ and $z_2 = a_2 + ib_2$ where $a_1, b_1, a_2, b_2 \in \mathbb{R}$, we have:

$$z_1 z_2 = (a_1 + ib_1)(a_2 + ib_2) = a_1a_2 + ia_1b_2 + ia_2b_1 + i^2b_1b_2 = (a_1a_2 - b_1b_2) + i(a_1b_2 + a_2b_1),$$

and thus

$$\overline{z_1 z_2} = (a_1a_2 - b_1b_2) - i(a_1b_2 + a_2b_1).$$

On the other hand:

$$\overline{z_1}\,\overline{z_2} = (a_1 - ib_1)(a_2 - ib_2) = a_1a_2 - ia_1b_2 - ia_2b_1 + i^2b_1b_2 = (a_1a_2 - b_1b_2) - i(a_1b_2 + a_2b_1).$$

Comparing this expression to the one we derived for $\overline{z_1 z_2}$ shows that they are the same, so

$$\overline{z_1 z_2} = \overline{z_1}\,\overline{z_2}$$

as the base case requires.

Now, how do we move on to verifying the claimed equality for more than two complex numbers? Notice that the algebra involved in verifying it for two complex numbers was already a little messy, and that doing the same for *more* complex numbers will be even messier. For instance, it is definitely possible to verify that

$$\overline{z_1 z_2 z_3} = \overline{z_1}\,\overline{z_2}\,\overline{z_3}$$

by setting $z_1 = a_1 + ib_1, z_2 = a_2 + ib_2, z_3 = a_3 + ib_3$ and working out both sides in full to see that they give the same result, but doing so will be tedious and unenlightening, and it only gets worse when using even more complex numbers.

Instead, we need a way to use what we know to build up to a larger case. For instance, how can we use the base case of two complex numbers to give the corresponding fact for three complex numbers? The key is in recognizing the a product $z_1z_2z_3$ of three complex numbers can be thought of as the product of *two* complex numbers z_1z_2 and z_3! We know that our claim holds for two complex numbers so that

$$\overline{(z_1z_2)z_3} = \overline{z_1z_2}\,\overline{z_3},$$

and then we can use the base case again to break up the first term further. Thus, no additional algebra is needed, since we can indeed use one scenario to get the next. For four complex numbers we would apply the base case to

$$\overline{(z_1z_2z_3)z_4} = \overline{z_1z_2z_3}\,\overline{z_4},$$

and then the $n = 3$ case to break up the first piece further.

The overall point here is that we use not only the induction hypothesis but the *base case* in the induction step as well. This is a technique we'll see come up again a few times, and is a standard way of justifying that something you know to be true for two things at a time will be true for any finite numbers of things as well. Here is our proof:

Proof. By the algebra worked out above, we know that $\overline{z_1z_2} = \overline{z_1}\,\overline{z_2}$ for any two complex numbers z_1 and z_2. Suppose now that the conjugate of the product of any $n \geq 2$ complex numbers is the product of their individual conjugates. (This is the induction hypothesis.) Let $z_1, \ldots, z_{n+1}$ be any $n + 1$ complex numbers. Then

$$\begin{aligned}\overline{z_1 \cdots z_{n+1}} &= \overline{(z_1 \cdots z_n)z_{n+1}} \\ &= \overline{z_1 \cdots z_n}\,\overline{z_{n+1}} \\ &= \overline{z_1} \cdots \overline{z_n}\,\overline{z_{n+1}},\end{aligned}$$

where in the second line we use the base case in the third the induction hypothesis. By induction we thus conclude that the conjugate of the product of any finite number of complex numbers is the product of their individual conjugates. □

Surprising examples of induction. We now look at some interesting and perhaps surprising examples of induction. These are surprising since induction is not used to show that some formula, equality, or inequality holds, but rather to show that some process can be carried out. This is closer to how induction is actually used in much of modern mathematics: not necessarily to prove some property involving numbers, but rather to *do* something.

Example 1. For any $n \geq 1$, take any chessboard of size $2^n \times 2^n$ and remove one square. We claim that all squares which remain can be covered using only the following types of pieces:

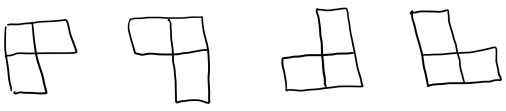

which we will refer to as the "available pieces". We will induct on the n showing up in the board size $2^n \times 2^n$. The base case is straightforward: removing one square from a 2×2 board leaves us with one of the following possibilities:

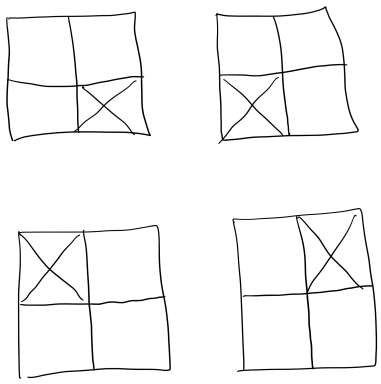

each of which can be covered with one piece of the available types.

Now, the question is: how do we use the knowledge that we can do this for a board of one size $2^n \times 2^n$ to show that we can then do it for a board of the next larger size $2^{n+1} \times 2^{n+1}$ as well? To get a feel for this we first look at the $2^2 \times 2^2$ case:

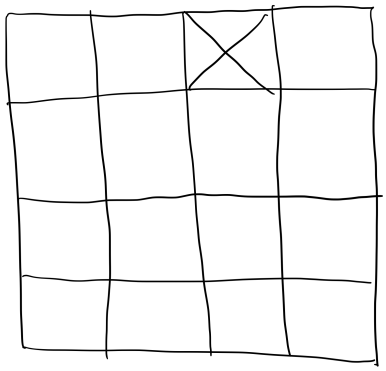

In order to build up to this, we must essentially "find" the base case somewhere in this setup. But this we can do: split up the $2^2 \times 2^2$ board into four 2×2 boards as follows:

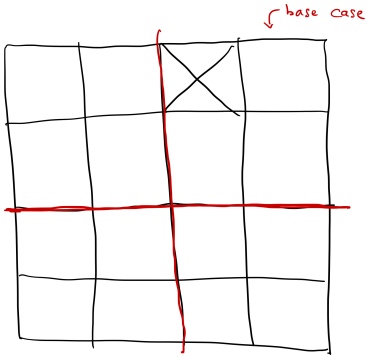

Now the base case shows up as one of the smaller boards, so we know that this smaller board can be covered using the available pieces. How do we handle the remaining three smaller boards? The base case only applies to 2×2 boards with a square removed, so to obtain such things we now remove the square from each of the remaining 2×2 boards which is closest to the center of the larger board:

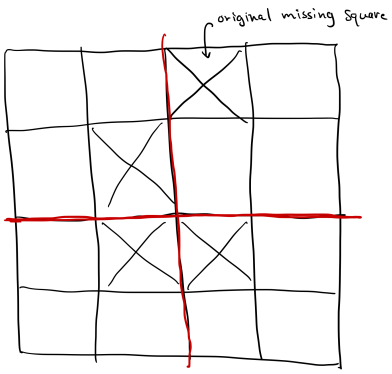

The point is that now we can apply the base case to these remaining boards, and finally use one more piece to fill in the three squares we removed near the center:

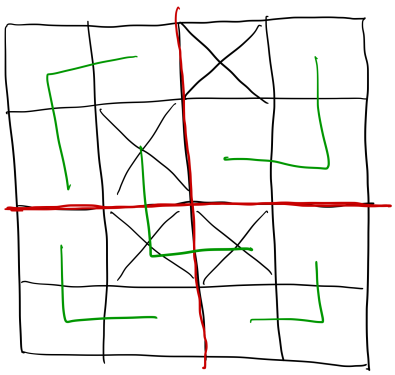

We now observe that we can use the exact same procedure to build from the $2^2 \times 2^2$ case up to the $2^3 \times 2^3$ case, then to the next case, and so on. The reason why this works is that we were able to find the "induction hypothesis" hiding within the case we were considering. Here is our proof:

Proof. The base case of a $2^1 \times 2^1$ with one square removed is worked out explicitly above. Suppose now that for some $n \geq 1$ we can cover any $2^n \times 2^n$ board with a square removed using the available shapes. Take any $2^{n+1} \times 2^{n+1}$ board with a square removed. Divide this in half both vertically and horizontally to obtain four $2^n \times 2^n$ boards. One of these smaller boards has a square removed, so we can cover this using the available pieces by the induction hypothesis.

Now remove from each of the remaining three $2^n \times 2^n$ boards the corner square which is closest to the center of the original board. This results in three $2^n \times 2^n$ boards each with a square removed, which we can cover with the available pieces using the induction hypothesis. Finally, use one single available piece to cover the squares from the smaller boards which were removed to get in the end a covering of the original $2^{n+1} \times 2^{n+1}$ board. We conclude by induction that this can thus be done for any board of size $2^n \times 2^n$ for any $n \geq 1$ with a square removed. □

Example 2. We claim that for $n \geq 3$, any convex n-sided polygon can be split up into $n - 2$ triangles. (A polygon is a shape all of whose sides are straight line segments, and saying a polygon is *convex* means that the line segment connecting two points in the polygon is itself fully contained in the polygon.) We will induct on n, the number of sides.

The base case $n = 3$ is a single triangle, so there is nothing to do since this already consists of $3 - 2 = 1$ triangle. If we consider a few examples, we see that we can clearly do what is being claimed:

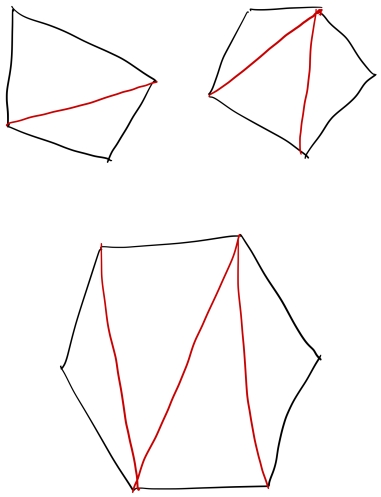

We can't very well check every possible polygon, so we now need a way to argue this can be done in general. Assuming we can do this for any n-gon (for some $n \geq 3$), we need to build up the case of an $(n+1)$-gon. Given some $(n+1)$-gon, we must thus find the "induction hypothesis" case of an n-gon hiding within our $(n+1)$-gon. But observe that if we connect two vertices like so:

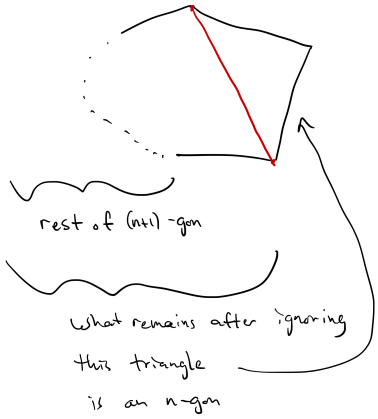

so we connect two vertices which happen to be adjacent to the *same* vertex, we end up dividing our original $(n+1)$-gon into an n-gon and a triangle; the resulting n-gon can be broken up into $n-2$ triangles by the induction hypothesis, so in the end we get that our original $(n+1)$-gon is broken up into $(n-2)+1 = (n+1)-2$ triangles as required. Note that convexity was used to guarantee that the segment we introduced to connect the two vertices above does indeed result in an n-gon and a triangle.

Example 3. Finally, we use induction to prove what's called *Euler's formula*, which is a result in the subject of *graph theory*. A *graph* is a collection of dots (i.e. vertices) and lines (i.e. edges) connecting them. A graph is *planar* if it can be drawn so that no two edges cross one another, and a graph is *connected* if it is possible to reach any vertex from any given one by moving along edges:

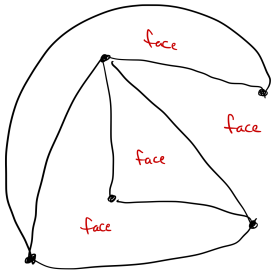

example of connected planar graph

V = 5 E = 7 F = 4

Any planar graph divides the plane up into regions we call *faces*; for instance, the planar graph drawn above has 4 faces. Euler's formula states any connected, planar graph satisfies

$$V - E + F = 2$$

where V is the number of vertices, E the number of edges, and F the number of faces. This fact has numerous applications in graph theory, but our motivation for talking about it is to give one more example of an interesting proof by induction; indeed, there is no "n" in Euler's formula at all, although there are three quantities V, E, F, so it is not so clear that this actually susceptible to induction in the first place. It is, and the key as usual is see how to build up one from scenario to something more complicated.

We will induct on the number E of edges. The simplest graph will have $E = 0$ edges, in which case we just have a single vertex:

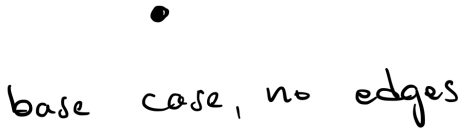

base case, no edges

Thus in this base case we have $V = 1, E = 0, F = 1$, so $V - E + F = 2$ as Euler's formula claims. Now, suppose Euler's formula holds for any connected planar graph with n edges and consider any connected planar graph with $n + 1$ edges. In order to be able to apply our induction hypothesis, we must find the n edge case hiding in our $n + 1$ edge case. But we can go from $n + 1$ edges down to n edges simply by removing an edge, which still results in a connected planar graph; Euler's formula applies to this smaller graph by the induction hypothesis, and the final thing to do is keep track of how removing an edge affects the numbers V, E, F involved. There are two things which can happen: either removing an edge leaves one vertex now isolated, or it does not

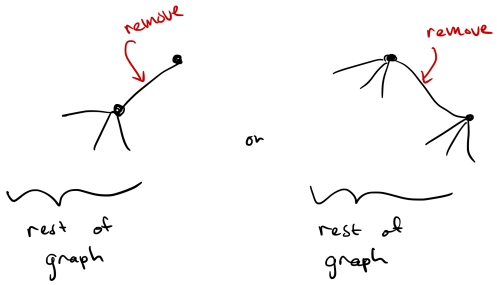

In the first case we must then remove the isolated vertex itself if our graph is to remain planar, in which case V and E both decrease by 1 while F remains unchanged, or in the second case V is left unchanged but the number of faces decreases by one, since the two faces sharing the edge we removed now combine to become a single face. Thus either way the quantity $V - E + F$ remains unchanged, and so still equals whatever it did in the induction hypothesis, which is 2.

I'll stop here and leave it to you to think about how to write this out formally, but the hard part of the work is already done above. For the purposes of learning induction, again the key point was in figuring out how to build from the induction hypothesis to the next larger case.

Lecture 14: Functions

Warm-Up. As a final induction problem, we show that any postage greater than or equal to 18 cents can be obtained using only 3 and 10 cent stamps. The base case of 18 cents can be obtained using six 3 cent stamps. We now have to think about how to build up to larger amounts.

From 18 cents we can get to 21 cents by throwing in another 3 cent stamp, or we can get to 28 cents by throwing in a 10 cent stamp, and similarly some larger quantities can be obtained starting with 18 cents. But what about 19 cents, 20 cents, 23 cents, or other values not attainable from a starting point of 18 alone? By hand we can check that 19 and 20 are attainable:

$$19 = 3(\text{3 cent stamps}) + 1(\text{10 cent stamp}), \text{ and } 20 = 2(\text{10 cent stamps}).$$

The key is now that if we also consider these to be base cases, we can then build up to all larger things, in particular since any larger value is attainable by adding enough 3 cent stamps to one of the starting values $18, 19$, or 20. Thus, this is an example of an induction with *multiple* base cases. Here is our proof:

Proof. As seen above, the base cases of $18, 19$, and 20 cents can be obtained using only 3 and 10 cent stamps. Suppose now that any postage smaller than some $n > 20$ can be obtained using only 3 and 10 cent stamps. Our goal is then to argue that n cents itself can be also obtained. Since $n - 3 < n$, $n - 3$ falls within the range of values to which the induction hypothesis (i.e. any value smaller than n) or base case ($n - 3$ could be $18, 19, 20$ for certain values of $n > 20$) applies, so this induction hypothesis gives us that $n - 3$ cents can be obtained using only 3 and 10 cent stamps. But then throwing in one more 3 cent stamps gives n cents overall as desired. We conclude by induction that any postage greater than or equal to 18 cents can be obtained. □

Functions. No doubt you've seen the notion of a function before in other courses, but here we will be a bit more formal and consider functions between arbitrary sets. A *function* $f : A \to B$ from a set A to a set B is an assignment of an element B to each element of A; in other words, a function gives to each "input" $a \in A$ a corresponding "output" $f(a) \in B$. (As in other courses, $f(a)$ is pronounced "f of a", which is the result of applying f to a.) Another common notation we'll see when defining a function is:

$$a \mapsto f(a),$$

which says that f sends a to the element $f(a)$. (The symbol $\mapsto$ means "maps to".) Given $f : A \to B$, A is called the *domain* of f, and B the *codomain* of f; so the domain is the set of possible inputs and the codomain the set where the outputs lie.

Here is a warning: the book covers functions in Chapter 12, but does so in what I think is an overly complicated manner. In particular, the book defines a function in terms of what are called *relations*, which is not something we've spoken about yet. It is true the a function is technically

a type of relation, but I think this approach obscures the basic idea behind what a function is, or rather what a function should do, which is send elements of one set to elements of another set. So, we will not use the book's approach and will instead focus on the usual (more intuitive) approach. As a result, we will do things in a different order than the book; in particular, we will first discuss the notions of *image/preimage* and then *injectivity/surjectivity*, whereas the book does this the other way around. Again, I think our approach is much simpler to follow.

Examples. We've all seen basic examples, like the function $f : \mathbb{R} \to \mathbb{R}$ defined by $f(x) = \sin x$ (i.e. the function which sends $x \in \mathbb{R}$ to $\sin x \in \mathbb{R}$), or $g : \mathbb{Z} \to \mathbb{Z}$ defined by $g(n) = 2n$ (i.e. the function which multiplies an integer input by 2).

But functions can be defined between arbitrary sets, even sets whose elements may not be numbers. To see a more interesting example we introduce the notion of a power set. Given a set A, the *power set* of A is the set $\mathcal{P}(A)$ of subsets of A:

$$\mathcal{P}(A) = \{S \mid S \subseteq A\}.$$

In other words, an *element* of $\mathcal{P}(A)$ is actually a *subset* of A, so this is our first example of a set whose elements themselves are sets! It might seem strange to think about sets whose elements are sets at first, but this is actually quite a common type of object in modern mathematics. For instance, the subsets $\{1, 2\}$ are:

$$\emptyset,\ \{1\},\ \{2\},\ \{1, 2\},$$

so the power set of $\{1, 2\}$ is:

$$\mathcal{P}(\{1, 2\}) = \{\emptyset, \{1\}, \{2\}, \{1, 2\}\}.$$

Again, note that each element of the power set, such as the empty set, is a set itself.

Consider now the function $f : \mathcal{P}(A) \to \mathcal{P}(A)$ defined by

$$S \mapsto A - S, \text{ or in other words by } f(S) = A - S.$$

This function takes as input a subset of A (i.e. an element of $\mathcal{P}(A)$) and outputs its complement in A. For instance, in the $A = \{1, 2\}$ example, f does the following:

$$f(\emptyset) = \{1, 2\},\ f(\{1\}) = \{2\},\ f(\{2\}) = \{1\},\ f(\{1, 2\}) = \emptyset,$$

in each case sending the input to its complement in $\{1, 2\}$. Properties of power sets will be more interesting when we discuss *cardinality*, but for now this is just meant to give an example of a function defined on more interesting types of sets.

Image and preimage. Suppose $f : A \to B$ is a function. Given a subset X of A, we define its *image* $f(X)$ under f to be the set of all things in B obtained by applying f to elements of X:

$$f(X) = \{f(x) \in B \mid x \in X\} = \{b \in B \mid \text{there exists } x \in X \text{ such that } f(x) = b\}.$$

To be clear, to say that $b \in f(X)$ means there is some $x \in X$ to which we can apply f to obtain b.

Given a subset Y of B, we define its *preimage* $f^{-1}(Y)$ under f to be the set of all possible inputs in A which are sent to an element of Y:

$$f^{-1}(Y) = \{a \in A \mid f(a) \in Y\}.$$

Again to be clear, to say that $a \in f^{-1}(Y)$ means by definition that $f(a) \in Y$. Visually, taking an image moves things "forward" while taking a preimage moves things "backwards":

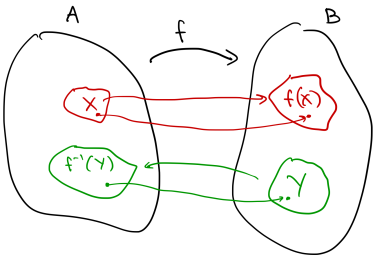

A word on notation: soon we will use f^{-1} to denote the *inverse* of the function f when it actually exists, but the f^{-1} showing up in the notation for preimage is not meant to necessarily signify an inverse function; the notion of a preimage makes sense for any function, even ones without inverses, and we still use the $f^{-1}(Y)$ notation for preimages even for non-invertible functions.

Example 1. Take the function $f : \mathbb{R} \to \mathbb{R}$ defined by $f(x) = x^2$. The image of the interval $[-1, 2]$ under f is the interval $[0, 4]$:

$$f([-1, 2]) = [0, 4].$$

This says two things: $f([-1, 2]) \subseteq [0, 4]$ means that the result of applying f to anything in $[-1, 2]$ gives something in $[0, 4]$, and $f([-1, 2]) \supseteq [0, 4]$ means that anything in $[0, 4]$ can be obtained as the result of applying f to something in $[-1, 2]$. In other words, in order to say that the image of $[-1, 2]$ is $[0, 4]$, it is not enough to know that applying f to something in $[-1, 2]$ gives something in $[0, 4]$, we also have to know that anything in $[0, 4]$ arises in this way.

To show the first containment, let $y \in f([-1, 2])$. By definition, this means that there exists $x \in [-1, 2]$ such that $f(x) = y$. But since $-1 \leq x \leq 2$, we get that $0 \leq x^2 \leq 4$, so that $y = f(x) = x^2$ is indeed in $[0, 4]$. Hence $f([-1, 2]) \subseteq [0, 4]$. Conversely, let $y \in [0, 4]$. To say that $y \in f([-1, 2])$ we need to know that there is something in $[-1, 2]$ which f sends to y. Since $0 \leq y \leq 4$, $0 \leq \sqrt{y} \leq 2$, where as usual $\sqrt{y}$ denotes the nonnegative square root of y. Since

$$f(\sqrt{y}) = \sqrt{y}^2 = y,$$

$\sqrt{y}$ is something in $[-1, 2]$ which is sent to y under f, so $y \in f([-1, 2])$. Hence $[0, 4] \subseteq f([-1, 2])$, so $f([-1, 2]) = [0, 4]$ as claimed.

As an example of a preimage, we have that the preimage of $[-1, 2]$ under f is $[-\sqrt{2}, \sqrt{2}]$:

$$f^{-1}([-1, 2]) = [-\sqrt{2}, \sqrt{2}].$$

Here we are thinking of $[-1, 2]$ as a subset of the *codomain*, which is also $\mathbb{R}$, and asking for the elements in the domain $\mathbb{R}$ which are sent into $[-1, 2]$ after applying f. Saying that $[-\sqrt{2}, \sqrt{2}] \subseteq f^{-1}([-1, 2])$ means that applying f to anything in $[-\sqrt{2}, \sqrt{2}]$ gives something in $[-1, 2]$ (which is true since if $-\sqrt{2} \leq x \leq \sqrt{2}$, then $-1 \leq x^2 \leq 2$ is true), and saying that $f^{-1}([-1, 2]) \subseteq [-\sqrt{2}, \sqrt{2}]$ means that only things in $[-\sqrt{2}, \sqrt{2}]$ can belong to $f^{-1}([-1, 2])$, or in other words that anything outside $[-\sqrt{2}, \sqrt{2}]$ is *not* sent to something in $[-1, 2]$.

Example 2. Let $f : \mathbb{Z} \to \mathbb{Z}$ be the function defined by $f(n) = 2n$. The image of the entire domain $\mathbb{Z}$ under this is the set of even integers, while the image of the set of even integers is the set of all multiples of 4. The preimage of the set of odd integers is the empty set since there is nothing in $\mathbb{Z}$ which is sent to an odd integer when applying f.

Example 3. Let $f : \mathbb{R} \to \mathbb{R}$ be the function defined by

$$f(n) = \begin{cases} x + 2 & \text{if } x < 1 \\ \frac{1}{x} & \text{if } x \geq 1, \end{cases}$$

a portion of whose graph looks like:

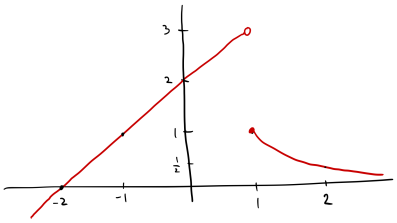

The image of the interval $[-1, 2)$ consists of everything obtained by applying f to things in $[-1, 2)$; visually in terms of the graph, this consists of the y-coordinates of points on the graph corresponding to x-coordinates in given green interval:

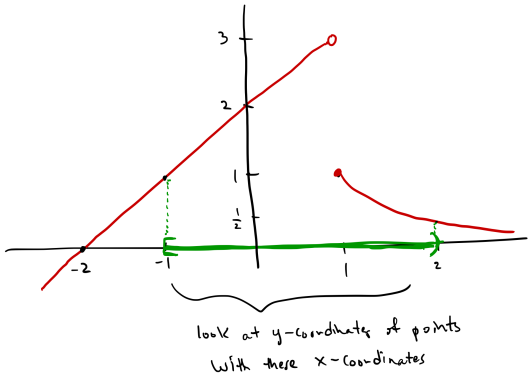

so $f([-1, 2)) = \left[\frac{1}{2}, 3\right)$. The preimage of the interval $\left[\frac{3}{4}, \frac{5}{2}\right]$, visualized as an interval on the vertical y-axis, consists of all x-coordinates corresponding to points whose y-coordinates are in $\left[\frac{3}{4}, \frac{5}{2}\right]$:

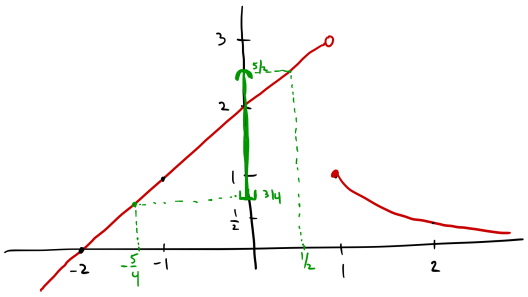

Only points between $-\frac{5}{4}$ and $\frac{1}{2}$ (not including $\frac{1}{2}$) have this property, so $f^{-1}\left[\frac{3}{4}, \frac{5}{2}\right) = \left[-\frac{5}{4}, \frac{1}{2}\right)$.

Images and unions/intersections. Suppose $f : A \to B$ is a function and that $X, Y \subseteq A$. It is a fact that

$$f(X \cup Y) = f(X) \cup f(Y),$$

so the image of a union is the union of individual images; or in other words, first taking a union and then the image gives the same result as first taking images and then a union. This will be left to you to prove on a homework assignment, but comes down to showing that each side is a subset of the other.

Similarly we can ask how intersections behave when taking images. We claim that

$$f(X \cap Y) \subseteq f(X) \cap f(Y)$$

but that the opposite containment does not hold, which reflects a property of functions we'll come back to later. For now, to prove the given containment we start with $b \in f(X \cap Y)$. To show that $b \in f(X) \cap f(Y)$ requires showing $b \in f(X)$ and $b \in f(Y)$, the first of which requires showing there is something in X which f sends to b and the second showing there is something in Y which f sends to b. But since $b \in f(X \cap Y)$, we know there exists $a \in X \cap Y$ such that $f(a) = b$, so this $a \in X \cap Y$ is already an element of X which is sent to b, so $b \in f(X)$, and an element of Y which is sent to b, so $b \in f(Y)$. Thus $b \in f(X) \cap f(Y)$, so $f(X \cap Y) \subseteq f(X) \cap f(Y)$.

Here is an example showing the opposite containment does not hold in general. Take $f : \mathbb{R} \to \mathbb{R}$ to be $f(x) = x^2$, $X = [-1, 0]$ and $Y = [0, 1]$. Then

$$f(X) = [0, 1] = f(Y) \text{ and } f(X \cap Y) = f(\{0\}) = \{0\},$$

so $f(X) \cap f(Y) = [0, 1]$ is not a subset of $f(X \cap Y)$. We'll come back to the question as to when this opposite containment *does* hold soon.

Lecture 15: Images and Preimages

Warm-Up 1. Let $f : \mathbb{Z} \to \mathbb{Z}$ be the function defined by

$$f(n) = \begin{cases} n + 2 & \text{if } n \text{ is even} \\ 2n & \text{if } n \text{ is odd.} \end{cases}$$

We show that the image of the set O of odd integers under f is the same as the image of the set M of multiples of 4:

$$f(O) = f(M).$$

First we should get a sense for what these images look like. The function f sends an even n to $n+2$ and an odd n to $2n$; note that the resulting values are then always even. Thus f only outputs even integers, so the image of either O or M should be a subset of the set of even integers. However, not all even integers will be in the image of O or M; for instance, $4 \notin f(O)$ since there does exist $n \in O$ satisfying $f(n) = 4$, because if n is odd then $f(n) = 2n$ and 4 cannot be written as $2n$ for an odd integer n. In fact, applying f to an odd $n = 2k + 1$ gives $f(n) = 2n = 4k + 2$, so only even integers of the form $4k + 2$ should be in the image of O. Similarly, if $n = 4k$ is a multiple of 4, then $f(n) = n + 2 = 4k + 2$, so only even integers of the form $4k + 2$ should in the image of M. Hence concretely, it should be the base that $f(O) = f(M)$ is the set of integers of the form $4k + 2$, a fact which will essentially follow from our proof.

Now, let $m \in f(O)$. In order to show that $m \in f(M)$ we need to show there exists $n \in M$ such that $f(n) = m$. For such $n \in M$, $f(n) = n + 2$ since n, being a multiple of 4, is even, so what we really need to show is that we can write m as $n + 2$ for some multiple n of 4. Since $m \in f(O)$, we know there exists $a \in O$ such that $f(a) = m$. This $a \in O$ can be written as $a = 2k + 1$ for some $k \in \mathbb{Z}$, and by the definition of f we then have:

$$m = f(a) = 2a = 2(2k + 1) = 4k + 2.$$

Hence $n = 4k$ is an element of M satisfying $f(n) = n + 2 = m$, so $m \in f(M)$. Thus $f(O) \subseteq f(M)$.

Conversely let $m \in f(M)$. In order to show that $m \in f(O)$ we need to show there exists $n \in O$ such that $f(n) = m$, which concretely means we need to show that $2n = m$ for some $m \in O$ since $f(n) = 2n$ for odd n. Since $m \in f(M)$, there exists $b \in M$ such that $f(b) = m$. This b is a multiple of 4, so $b = 4k$ for some $k \in \mathbb{Z}$. Then

$$m = f(b) = b + 2 = 4k + 2 = 2(2k + 1)$$

where $f(b) = b + 2$ since b is even. Thus $n = 2k + 1$ is an element of O satisfying $f(n) = 2n = m$, so $m \in f(O)$. Hence $f(M) \subseteq f(O)$, so $f(O) = f(M)$ as claimed.

Warm-Up 2. Suppose $f : A \to B$ is a function and that X, Y are subsets of B. We show that

$$f^{-1}(X - Y) = f^{-1}(X) - f^{-1}(Y),$$

which says that the operation of taking preimages "preserves" complements; i.e. taking the complement first and then the preimage gives the same result as taking preimages first and then the complement. This is a basic set equality proof where we show that each side is a subset of the other, winding the definition of preimage along the way.

Let $a \in f^{-1}(X - Y)$. By definition of preimage this means that $f(a) \in X - Y$. Hence $f(a) \in X$ and $f(a) \notin Y$. Since $f(a) \in X$, $a \in f^{-1}(X)$ by definition of preimage, and since $f(a) \notin Y$, $a \notin f^{-1}(Y)$. Thus $a \in f^{-1}(X) - f^{-1}(Y)$, so $f^{-1}(X - Y) \subseteq f^{-1}(X) - f^{-1}(Y)$.

Conversely let $a \in f^{-1}(X) - f^{-1}(Y)$. Then $a \in f^{-1}(X)$ and $a \notin f^{-1}(Y)$. Since $a \in f^{-1}(X)$, $f(a) \in X$, and since $a \notin f^{-1}(Y)$, $f(a) \notin Y$. Thus $f(a) \in X - Y$, so $a \in f^{-1}(X - Y)$. Hence $f^{-1}(X) - f^{-1}(Y) \subseteq f^{-1}(X - Y)$, so we have equality as claimed.

Back to images and intersections. Recall that last time we made the claim (and left the proof to the homework) that when $f : A \to B$ is a function and X, Y are subsets of A, the following equality holds:

$$f(X \cup Y) = f(X) \cup f(Y),$$

which says that the operation of taking an image "preserves" unions. We then proved that with respect to intersections, $f(X \cap Y) \subseteq f(X) \cap f(Y)$ was always true, but that the opposite containment is not always true by giving an explicit counterexample. (Thus, the operation of taking images does not necessarily "preserve" intersections.)

However, how could we have guessed that the opposite containment

$$f(X) \cap f(Y) \subseteq f(X \cap Y)$$

should not be true in general? *I* know this from the experience of having seen this type of thing many times before, but it is not something which one might immediately realize at the first glance. Indeed, as a student learning these things for the first time, your first instinct might be to try to prove that this containment *does* hold since, after all, why shouldn't it? Here is an attempt at a "proof" that this containment does hold, where I say "proof" because we of course know that this proof cannot be correct:

"Proof" that $f(X) \cap f(Y) \subseteq f(X \cap Y)$. Let $b \in f(X) \cap f(Y)$, so that $b \in f(X)$ and $b \in f(Y)$. Since $b \in f(X)$, there exists $a \in X$ such that $f(a) = b$, and since $b \in f(Y)$, there exists $a \in Y$ such that $f(a) = b$. Since $a \in X$ and $a \in Y$, $a \in X \cap Y$, so $b \in f(X \cap Y)$ since there is something in $X \cap Y$ mapping to b. Thus $f(X) \cap f(Y) \subseteq f(X \cap Y)$. □

So, what is wrong here? Have we not shown that anything in $f(X) \cap f(Y)$ is also in $f(X \cap Y)$? To pinpoint what goes wrong let's run through this reasoning with the counterexample we gave last time: $f : \mathbb{R} \to \mathbb{R}$ defined $f(x) = x^2$, $X = [-1, 0]$, and $Y = [0, 1]$. We have $f(X) = [0, 1] = f(Y)$, so take $b = 1 \in f(X) \cap f(Y)$. Following our attempted proof, we pick $a = -1 \in X$ such that $f(a) = b$ and $a = 1 \in Y$ such that $f(a) = b$. But now we see the problem: these are two different values of a, and so they do not give an element common to *both* X and Y which f sends to b. Indeed, since $X \cap Y = \{0\}$, there is not element in $X \cap Y$ which is sent to $b = 1$, which is the reason why this

counterexample works. The issue with our proof is that we were sloppy in our notation: once we pick $a \in X$ such that $f(a) = b$, we should not use the same a to denote the element of Y which is sent to b. Rather, we pick $a' \in Y$ such that $f(a') = b$, but now we have no way of guaranteeing that there is an element in $X \cap Y$ sent to b since a and a' could very well be different:

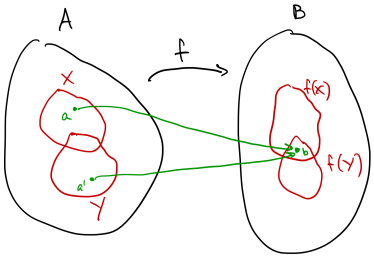

But now we ask: for which type of function f would the containment $f(X) \cap f(Y) \subseteq f(X \cap Y)$ actually be true? The issue above was that we have $a \in X$ and $a' \in Y$ sent to b, and if these are different we don't get an element of $X \cap Y$ sent to b. The underling point is that we could have *different* elements a, a' sent to the same b; whereas if we knew a and a' actually had to be the same then the proof would work and the given containment would hold. Hence, for instance, for a function with the property that you cannot have different elements sent to the same thing, the given containment would indeed be true. Such a function is said to be *injective*, which is a notion we'll come back to later.

Preimage of the image. In the same setup as before, $f : A \to B$ a function with $X \subseteq A$, consider now what happens if we take the preimage of the image of X: $f^{-1}(f(X))$. To be clear, the image $f(X)$ of X is a subset of B, whose preimage we can then take to get a subset $f^{-1}(f(X))$ of A again. We want to understand how this preimage of the image relates to the original X.

To say that $a \in f^{-1}(f(X))$ means $f(a) \in f(X)$, which in turn means there exists $x \in X$ such that $f(x) = f(a)$. Thus a is in $f^{-1}(f(X))$ when a is sent under f to the same thing as some $x \in X$ is sent to. Certainly, any $a \in X$ will have this property, since a is sent to the same thing as something in X (namely a itself) is sent to, so

$$X \subseteq f^{-1}(f(X)).$$

For a quick formal proof, take $x \in X$. Then $f(x) \in f(X)$ by definition of image, so $x \in f^{-1}(f(X))$ by definition of preimage. Hence $X \subseteq f^{-1}(f(X))$.

Now, is it true that $X = f^{-1}(f(X))$? For this it would have to be true that $f^{-1}(f(X)) \subseteq X$, but we already said previously that elements a of $f^{-1}(f(X))$ are things which are sent under f to the same thing as something in X is sent to, but this does NOT guarantee that a itself must have been in X to start with: all we know is that there exists $x \in X$ such that $f(a) = f(x)$, but not that $a \in X$.

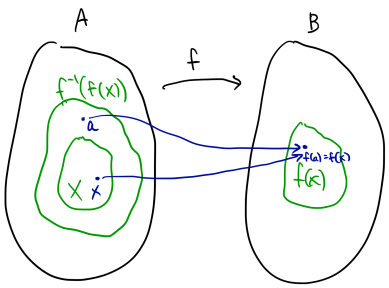

Indeed, in the $f(x) = x^2$ example, for $X = [0, 1]$ we have

$$f(X) = [0, 1], \text{ so } f^{-1}(f(X)) = f^{-1}([0, 1]) = [-1, 1]$$

since all elements in $[-1, 1]$ are sent to something in $[0, 1]$, so $f^{-1}(f(X)) \not\subseteq X$ in this case. The issue is that $-1 \in f^{-1}(f(X))$ since it is sent to the same thing as what $1 \in X$ is sent to, but $-1 \notin X$.

So, in general, we can only say that $X \subseteq f^{-1}(f(X))$ always holds, but that the opposite containment does not necessarily hold. When does the opposite containment hold? We saw before the problem arises when $f(a) = f(x)$ for some $x \in X$ and $a \notin X$, or in other words when we have two different elements being sent to the same thing. Hence, as in the case of intersections, for a function which does not send different things to the same thing (i.e. an injective function), it would in fact be true that $f^{-1}(f(X)) = X$.

Image of the preimage. As a final example, in the same setup as before, suppose now that $Y \subseteq B$ and we take the image of the preimage of Y: $f(f^{-1}(Y))$. To be clear, the preimage of Y gives a subset $f^{-1}(Y)$ of A, whose image we can then take to get a subset $f(f^{-1}(Y))$ of B. How does this compare to the original Y? We claim that

$$f(f^{-1}(Y)) \subseteq Y$$

is always true. Indeed, so say that $b \in f(f^{-1}(Y))$ means that there is some $a \in f^{-1}(Y)$ such that $f(a) = b$, but by virtue of the fact that $a \in f^{-1}(Y)$ we have that $b = f(a) \in Y$. (In other words, $f^{-1}(Y)$ is the set of all things in A which get sent to something in Y, and taking the image of this then applies f to such elements, thereby giving something in Y as a result.)

However, the opposite containment $Y \subseteq f(f^{-1}(Y))$ is not necessarily true. We'll give a counterexample in a second, but first let's try to understand what is wrong with the following "proof":

"Proof" that $Y \subseteq f(f^{-1}(Y))$. Let $y \in Y$. Pick $a \in A$ such that $f(a) = y$. Since $f(a) = y$ is in Y, $a \in f^{-1}(Y)$, and thus $y = f(a)$ is in $f(f^{-1}(Y))$. This shows that $Y \subseteq f(f^{-1}(Y))$ as claimed. □

So, what's wrong? For a counterexample, take $f(x) = x^2$ and $Y = [-1, 1]$. Then $f^{-1}(Y) = [-1, 1]$ and $f(f^{-1}(Y)) = f([-1, 1]) = [0, 1]$, which does not contain $Y = [-1, 1]$ as a subset. If we try to run through our "proof" with $y = -1$, we see the problem immediately: there is no a such that $f(a) = y$ in this case since squaring a real number cannot possibly give -1. Thus, our proof fails because there is no guarantee such $a \in A$ exists.

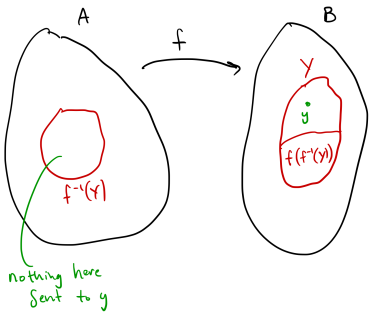

However, if we could guarantee that any $y \in Y$ was actually obtainable as the result of applying f to some element of A, our proof would work find and $f(f^{-1}(Y)) = Y$ would be true in such a case. A function with the property that everything in codomain is obtainable as an actual output is said to be *surjective*, and is a notion, along with injective, we'll come back to next time.

Lecture 16: Injectivity and Surjectivity

Warm-Up. Let $f : \mathbb{R}^2 \to \mathbb{R}^2$ be the function defined by $f(a, b) = (a + b, 2a + 2b)$. We determine the image of f (meaning the image $f(\mathbb{R}^2)$ of the entire domain) and the preimage of $\{(0, 0)\}$.

First, note that the y-coordinate $2a + 2b$ of any possible output is twice the x-coordinate $a + b$, so all possible outputs lie on the line $y = 2x$ defined by:

$$L = \big\{(x, y) \in \mathbb{R}^2 \;\big|\; y = 2x\big\}.$$

This observation so far only tells us that the image of f is contained in this line, but we claim in the fact that the image is the entire line:

$$f(\mathbb{R}^2) = L.$$

For this we need to verify that not only is anything in the image on the line (which we alluded to above) but also that anything on the line is in the image.

Let $(x, y) \in f(\mathbb{R}^2)$. Then there exists $(a, b) \in \mathbb{R}^2$ such that $f(a, b) = (x, y)$. But by the definition of f this means

$$(x, y) = f(a, b) = (a + b, 2a + 2b),$$

so $x = a + b$ and $y = 2a + 2b$. Hence $y = 2x$ is satisfied, so $(x, y) \in L$ and thus $f(\mathbb{R}^2) \subseteq L$. Conversely let $(x, y) \in L$. Then $y = 2x$ by definition of L. To show that $(x, y) \in f(\mathbb{R}^2)$ we must show that there exists a point in $\mathbb{R}^2$ to which applying f gives (x, y); that is, we need (a, b) such that

$$f(a, b) = (a + b, 2a + 2b) = (x, y).$$

But since, for instance, $f(x, 0) = (x + 0, 2x + 0) = (x, 2x) = (x, y)$, $(x, 0)$ is an element of $\mathbb{R}^2$ which f sends to (x, y), so $(x, y) \in f(\mathbb{R}^2)$. Hence $L \subseteq f(\mathbb{R}^2)$, so $f(\mathbb{R}^2) = L$ as claimed.

Now, the preimage $f^{-1}(\{(0, 0)\})$ of $\{(0, 0)\}$ consists of all points in $\mathbb{R}^2$ which are sent to something in $\{(0, 0)\}$, which in this case means all points which are sent to $(0, 0)$. In other words, the preimage consists of all (a, b) satsifying

$$f(a, b) = (a + b, 2a + 2b) = (0, 0).$$

This requires that $a + b = 0$, so that $b = -a$. Hence such a point must lie on the line $y = -x$ defined by

$$S = \big\{(x, y) \in \mathbb{R}^2 \;\big|\; y = -x\big\},$$

so we claim that $f^{-1}(\{(0, 0)\}) = S$. The proof is essentially what we did above, but let us write it out more formally.

Let $(a, b) \in f^{-1}(\{(0, 0)\})$. Then $f(a, b) \in \{(0, 0)\}$, so $f(a, b) = (0, 0)$. Hence

$$(0, 0) = f(a, b) = (a + b, 2a + 2b), \text{ so } 0 = a + b.$$

Thus $b = -a$, so $(a, b) \in S$ and hence $f^{-1}(\{(0, 0)\}) \subseteq S$. Conversely let $(x, y) \in S$, so that $y = -x$. Then $f(x, y) = (x + y, 2x + 2y) = (x - x, 2x - 2x) = (0, 0)$, so $(x, y) \in f^{-1}(\{(0, 0)\})$. Hence $S \subseteq f^{-1}(\{(0, 0)\})$, so $f^{-1}(\{(0, 0)\}) = S$ as claimed.

Injective and surjective functions. Suppose $f : A \to B$ is a function. We say that f is *injective* if whenever $x \neq y$ in A, then $f(x) \neq f(y)$; that is, different elements get sent to different things. By taking the contrapositive, we can phrase this as:

$$\text{if } f(x) = f(y), \text{ then } x = y,$$

which says that the only way two elements can be sent to the same thing is if they were actually the same element to start with. (This latter phrasing will usually be the better one to use when actually showing a given function is injective.) Another term for "injective" you'll often see is *one-to-one*, which emphasizes the idea that each output can only come from one input.

We say that f is *surjective* if for every $b \in B$ there exists $a \in A$ such that $f(a) = b$; that is, everything in the codomain can be obtained by applying f to some input. Or, said in another way using the language of images, for $f : A \to B$ to be surjective mans that the image $f(A)$ of f is all of B: $f(A) = B$. Another term for "surjective" you'll often see is *onto*.

To go back to some things we saw last time, which now we see were meant to provide possible motivations for the notions of injective and surjective, recall that the following always hold for any function:

$$f(X \cap Y) \subseteq f(X) \cap f(Y),\ X \subseteq f^{-1}(f(X)),\ f(f^{-1}(S)) \subseteq S.$$

Based on our previous discussion, we can now say that, in addiition, when f is injective it is true that $f(X \cap Y) = f(X) \cap f(Y)$ and $X = f^{-1}(f(X))$, and when f is surjective it is true that $f(f^{-1}(S)) = S$. In particular then, for injective functions, the operation of taking images *does* preserve intersections. In fact, it is also true that $f(X - Y) = f(X) - f(Y)$ for an injective function, so the operation of taking images under injective functions preserves complements, but this is not the case for non-injective functions.

Example 1. Consider the function $f : \mathbb{R}^2 \to \mathbb{R}^2, f(a, b) = (a + b, 2a + 2b)$ from the Warm-Up. We showed there that the image of f was only the line $y = 2x$, so since this is not all of $\mathbb{R}^2$ the function f is not surjective. For instance, for $(1, 1) \in \mathbb{R}^2$ there does not exist $(a, b) \in \mathbb{R}^2$ such that $f(a, b) = (1, 1)$ since there do not exist a, b satisfying $(a + b, 2a + 2b) = (1, 1)$.

The fact that the preimage of $\{(0, 0\}$ is the line $y = -x$ immediately shows that f is not injective, since there are multiple points being sent to $(0, 0)$. For instance, $f(-1, 1) = (0, 0) = f(1, -1)$ but $(-1, 1)$ and $(1, -1)$ are not the same, so f is not injective.

Example 2. Whether or not a given expression defines an injective/surjective function completely depends on the domain and codomain being considered. For instance, the function $f : \mathbb{R} \to \mathbb{R}$ defined by $f(x) = x^2$ is neither injective, since 1 and -1 both get sent to the same thing, nor surjective, since -1 is not in the image.

However, with $\mathbb{R}_{\geq 0}$ denoting the set of nonnegative real numbers, the function $g : \mathbb{R} \to \mathbb{R}_{\geq 0}$ defined by $g(x) = x^2$ *is* surjective, since by cutting the domain down to be $\mathbb{R}_{\geq 0}$ instead of all of $\mathbb{R}$ we have eliminated the elements of $\mathbb{R}$ which are not attained as outputs. Along these same lines, the function $h : \mathbb{R}_{\geq 0} \to \mathbb{R}$ defined by $h(x) = x^2$ *is* injective since for *non-negative numbers*, $x^2 = y^2$ does imply $x = y$. Finally, the function $\ell : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ defined by $\ell(x) = x^2$ is both injective and surjective—we say that such a function is *bijective*.

Example 3. Here is another example of a bijective function. For a set A, let $f : \mathcal{P}(A) \to \mathcal{P}(A)$ be the function sending a subset S of A to its complement $A - S$ in A:

$$f(S) = A - S.$$

(Recall that $\mathcal{P}(A)$ denotes the power set of A, which is the set of all subsets of A.) To see that f is injective, suppose $f(S) = f(S')$. This then means that S and S' have the same complement in A:

$$A - S = A - S'.$$

We need to show that $S = S'$. Since $A - S = A - S'$, the complements of each of these should also be the same:

$$A - (A - S) = A - (A - S').$$

But taking the complement of the complement results in the original set, so the left side above is S and the right side is S', so $S = S'$ and f is injective as claimed.

To show that f is surjective, let $S \in \mathcal{P}(A)$. We need to know there exists an element of $\mathcal{P}(A)$ which is sent to S under f, or in other words a subset of A whose complement is S. But

$$f(A - S) = A - (A - S) = S,$$

so $A - S \in \mathcal{P}(A)$ is an element sent to S, and hence f is surjective. Since f is injective and surjective, it is bijective.

Injectivity/surjective via sizes of preimages. Let us see how to rephrase the notions of injective and surjective in terms how large a preimage of the form $f^{-1}(\{b\})$ can be. Given $f : A \to B$ and $b \in B$, $f^{-1}(\{b\})$ consists of all elements of A which are sent to B. But to be injective means that, if there is such an element, there can only be one, so we have that:

$$f \text{ is injective if and only if for all } b \in B, f^{-1}(\{b\}) \text{ has at } \textit{most} \text{ one element.}$$

On the other hand, to be surjective means that given any such $b \in B$, there *is* an element which is sent to B; such an element will then be in $f^{-1}(\{b\})$, so

$$f \text{ is surjective if and only if for all } b \in B, f^{-1}(\{b\}) \text{ has at } \textit{least} \text{ one element.}$$

Thus, we can also say that f is bijective if and only if for all $b \in B$, $f^{-1}(\{b\})$ has *exactly* one element. These observations are meant to make a connection between injectivity/surjectivity and the "sizes" of certain sets.

Looking ahead to cardinality. We will later talk about the *cardinality* of a set as being the number of elements it contains, so the above facts tells us that: f is injective if the preimage of every single element has cardinality at most 1; f is surjective if the preimage of every single element has cardinality at least 1; and f is bijective if the preimage of every single element has cardinality exactly 1. The notions of injective and surjective will play a key role in understanding the cardinality of infinite sets.

To see one more glimpse of this, suppose A and B are *finite* sets and that $f : A \to B$ is a function. We are interested in understanding what f being injective or surjective tells us about the cardinality of A in relation to the cardinality of B. We first claim that if f is injective, then A must have no more elements than B does, or in other words that the cardinality of A is less than or equal to the cardinality of B:

$$|A| \leq |B|,$$

where $|S|$ is the notation for the cardinality of S, which for finite sets just means the number of elements in S. Indeed, if A had more elements than B, the necessarily we would have more than element of A being sent to the same element of B, in which case f would not be injective. If instead f is surjective, then A can have no fewer elements than does B, or in other words the cardinality of A is greater than or equal to the cardinality of B:

$$|A| \geq |B|.$$

Indeed, if A had fewer elements than B, then there would necessarily be elements of B not obtained as actual outputs since there wouldn't be enough inputs to allow that, but in this case f could not be surjective. Again, the fact that we have a relation between injectivity/surjectivity and the sizes of finite sets will pay dividends when we discuss the sizes of infinite sets.

Lecture 17: Compositions

Warm-Up 1. Let $f : \mathbb{R} - \{1\} \to \mathbb{R}$ be the function defined by $f(x) = \frac{2x+1}{x-1}$. We show that f is injective. (Note that 1 is excluded from the domain so $\frac{2x+1}{x-1}$ makes sense for all x in the domain.) Suppose $f(x) = f(y)$; we must show that $x = y$. Since $f(x) = f(y)$, we have

$$\frac{2x+1}{x-1} = \frac{2y+1}{y-1}$$

by the definition of f. Thus

$$(2x+1)(y-1) = (2y+1)(x-1), \text{ so } 2xy - 2x + y - 1 = 2yx - 2y + x - 1.$$

This gives $-3x = -3y$, so $x = y$ as desired. Hence f is injective.

Warm-Up 2. Let $f : \mathbb{R}^2 \to \mathbb{R}^2$ be the function $f(x, y) = (x + y, 2x + y)$. We show that f is surjective. Thus, for $(a, b) \in \mathbb{R}^2$, we must show there exists $(x, y) \in \mathbb{R}^2$ such that $f(x, y) = (a, b)$. Using the definition of f, this means that (x, y) must satisfy

$$(x + y, 2x + y) = (a, b), \text{ so } x + y = a \text{ and } 2x + y = b.$$

Thus the question comes down to showing that this collection of equations has a solution for any a, b. (This should remind you of things you saw in a previous linear algebra course, although here we won't require any linear algebra to proceed.) Working out some scratch work on the side, the first equation $x + y = a$ gives $x = a - y$, and then the second $2x + y = b$ gives

$$2(a - y) + y = b, \text{ so } y = 2a - b.$$

Then $x = a - y = b - a$, so this scratch work suggests that $x = b - a, y = 2a - b$ should satisfy $f(x, y) = (a, b)$. We verify this:

$$f(b - a, 2a - b) = ([b - a] + [2a - b], 2[b - a] + [2a - b]) = (a, b).$$

Thus any $(a, b) \in \mathbb{R}^2$ arises as the result of applying f to some element of the domain, so f is surjective as claimed.

Compositions. Given functions $f : X \to Y$ and $g : Y \to Z$, we can look at the function obtained by first applying f and then applying g; this is called the *composition* of f and g and is denoted by $g \circ f$. To be clear, the composition of f and g is the function $g \circ f : X \to Z$ defined by

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

The notion of a composition should be familiar from a calculus course, where for instance the chain rule tells us how to differentiate compositions, but now we're pointing out that compositions make sense in the more general setting of functions between arbitrary sets.

Note the notation: $g \circ f$ means that first we apply f and then we apply g. More generally, we can speak of the composition $f_n \circ \cdots \circ f_1$ of n functions:

$$X_1 \xrightarrow{f_1} X_2 \xrightarrow{f_2} X_3 \xrightarrow{f_3} \cdots \xrightarrow{f_{n-1}} X_n \xrightarrow{f_n} X_{n+1},$$

each one mapping into the domain of the next. ($A \xrightarrow{f} B$ is alternate notation for $f : A \to B$.)

Example. Denote by $\mathbb{R}^\infty$ the set of infinite sequences of real numbers, so that an element of $\mathbb{R}^\infty$ looks like:

$$(x_1, x_2, x_3, \ldots) \text{ where each } x_i \in \mathbb{R}.$$

Define the functions $L, R : \mathbb{R}^\infty \to \mathbb{R}^\infty$ by

$$L(x_1, x_2, x_3, \ldots) = (x_2, x_3, x_4, \ldots) \text{ and } R(x_1, x_2, x_3, \ldots) = (0, x_1, x_2, \ldots).$$

We compute the compositions $R \circ L$ and $L \circ R$, both of which give functions $\mathbb{R}^\infty \to \mathbb{R}^\infty$. First we have:

$$(R \circ L)(x_1, x_2, x_3, \ldots) = R(L(x_1, x_2, x_3, \ldots)) = R(x_2, x_3, x_4, \ldots) = (0, x_2, x_3, \ldots),$$

so $R \circ L$ is the function which just replaces the first term in $(x_1, x_2, x_3, \ldots)$ with 0. Second:

$$(L \circ R)(x_1, x_2, x_3, \ldots) = L(R(x_1, x_2, x_3, \ldots)) = L(0, x_1, x_2, \ldots) = (x_1, x_2, x_3, \ldots),$$

so $L \circ R$ is the function which sends $(x_1, x_2, x_3, \ldots)$ to itself; we say that $L \circ R$ is the *identity* function on $\mathbb{R}^\infty$. Note that $R \circ L$ is neither injective nor surjective, while $L \circ R$ is bijective.

Compositions and injectivity/surjectivity. We now give some properties of compositions in relation to injectivity and surjectivity. Suppose $f : X \to Y$ and $g : Y \to Z$ are functions. We claim that if f and g are both injective, then $g \circ f$ is injective, and that if f and g are both surjective, then $g \circ f$ is surjective.

Suppose f and g are injective. To show that $g \circ f$ is injective we must show that whenever $(g \circ f)(x) = (g \circ f)(x')$, it must be true that $x = x'$. Thus, suppose $(g \circ f)(x) = (g \circ f)(x')$, meaning

$$g(f(x)) = g(f(x')).$$

How do we use the facts that g and f are injective in order to conclude that $x = x'$? Note that g being injective only tells us something about what happens when we plug things into g, and similarly f being injective only tells us something about what happens when plugging things into f. But the equality above can be viewed as saying that plugging $f(x)$ into g gives the same result as plugging $f(x')$ into g, so injectivity of g gives

$$f(x) = f(x').$$

Then injectivity of f gives $x = x'$ as desired, so $g \circ f$ is injective.

Suppose f and g are surjective and let $z \in Z$. To show that $g \circ f$ is surjective we must show there exists $x \in X$ such that $(g \circ f)(x) = z$, or in other words $g(f(x)) = z$. But before anything else notice that this equality requires that there be something we can plug into g, which will end up being $f(x)$ for the currently unspecified x, in order to get z, and it is the surjectivity of g which says we can do this. Since g is surjective, there exists $y \in Y$ such that $g(y) = z$. Now if we know that we can write y as $y = f(x)$ for some $x \in Z$ we will be done, and we do know this because f is surjective: since f is surjective, there exists $x \in X$ such that $f(x) = y$, and this x thus satisfies

$$(g \circ f)(x) = g(f(x)) = g(y) = z$$

as required. Hence $g \circ f$ is surjective.

Extending to more functions. From this we immediately get that the composition of any number of injective functions is injective and that the composition of any number of surjective

functions is surjective. Indeed, induction gives us a way to build up from the case above of two functions to more general situations. To be clear, suppose we know that the composition of n injective functions is always injective. Then if

$$X_1 \xrightarrow{f_1} X_2 \xrightarrow{f_2} X_3 \xrightarrow{f_3} \cdots \xrightarrow{f_n} X_{n+1} \xrightarrow{f_{n+1}} X_{n+2},$$

are all injective, we get that $f_n \circ \cdots \circ f_1$ is injective by the induction hypothesis and then that

$$f_{n+1} \circ \cdots \circ f_1 = f_{n+1} \circ (f_n \circ \cdots \circ f_1)$$

is injective from the base case. A similar argument works for the composition of any number of surjective functions. (As we mentioned a while ago, using induction to extend some fact from two things at a time to any number of things at a time is quite fruitful.)

Back to compositions and injectivity/surjectivity. Here are two more final claims for functions $f : X \to Y$ and $g : Y \to Z$: if $g \circ f$ is injective, then f is injective, and if $g \circ f$ is surjective, then g is surjective. The previous facts tells us that if the individual functions have some property, the composition will as well, whereas here we are claiming that if the composition has some property, *at least one* of the individual functions will as well: in any injective composition, the first function must be injective, and in any surjective composition, the final function must be surjective.

So, suppose $g \circ f$ is injective. To show that f is injective we suppose $f(x) = f(x')$ for some $x, x' \in X$. Our goal is to show that $x = x'$. Now, are assumption that $g \circ f$ is injective doesn't tell us anything *until* we bring g into the picture somehow, but the point is that if $f(x) = f(x')$ are the same element of Y, then applying g to this common element should give the same result:

$$g(f(x)) = g(f(x')).$$

But this equality now says that $(g \circ f)(x) = (g \circ f)(x')$, so injectivity of $g \circ f$ gives $x = x'$ as desired. Note again that in order to use injectivity of $g \circ f$ we had to introduce g somehow into our given expression $f(x) = f(x')$.

Now suppose $g \circ f$ is surjective and let $z \in Z$. To show that g is surjective requires showing there exists $y \in Y$ such that $g(y) = z$—where does this y come from? The question is do we know there is something we can input into g to give z, but we at least know there is something we can input into $g \circ f$ to give z: since $g \circ f$ is surjective, there exists $x \in X$ such that

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

But this equality says precisely that $y = f(x)$ is then an element of Y which g sends to z, as needed in order to be able to conclude that g is surjective, so g is surjective as claimed.

Careful. It is NOT true in general that if $g \circ f$ is injective, g must be injective, nor is it true in general that if $g \circ f$ is surjective, f must be surjective. We'll see how to reason through the process of finding examples of these facts next time.

Lecture 18: Invertibility

Warm-Up 1. We give an example of functions $f : X \to Y$ and $g : Y \to Z$ such that $g \circ f : X \to Z$ is injective but for which g is not injective. First, let us think through the process of finding such an example. Saying that $g \circ f$ is injective means that

$$\text{if } g(f(x_1)) = g(f(x_2)), \text{ then } x_1 = x_2.$$

The key observation is that this only tells us something about what happens when applying g to *outputs* of f since the hypothesis only uses elements of Z of the form $g(f(x))$. In other words, the assumption that $g \circ f$ is injective says *nothing* about what happens when applying g to elements of Y which are *not* in the image of f, since such elements cannot be written in the form $f(x)$. Thus, the idea is that the non-injectivity of g in the sought after example should come from those elements of Y which are not in the image of f.

In whatever example we find, we know that f will have to be injective since we showed last time that injectivity of $g \circ f$ implies injectivity of f. So, to start building up an example let us start with a simple injective function, say $f : \mathbb{N} \to \mathbb{N}$ defined by $f(n) = n^2$. (Note that this is injective since $\mathbb{N}$ only consists of positive integers.) The assumption that $g \circ f$ is injective will thus only tells us something about applying g to elements of the form n^2, but for instance it will say nothing about what $g(2), g(3), g(5)$, etc should be. So, to guarantee that g is not injective we can, for instance, send all such elements to the same thing. Maintaining injectivity of g on elements of the form n^2 but forcing g to be non-injective on elements not of this form should then give a valid example.

So, take $g : \mathbb{N} \to \mathbb{N}$ defined by

$$g(m) = \begin{cases} m & \text{if } m = n^2 \text{ for some } n \in \mathbb{N} \\ 1 & \text{otherwise} \end{cases}$$

as one possible example. Then if $g(f(x_1)) = g(f(x_2))$ for some $x_1, x_2 \in \mathbb{N}$, we have:

$$g(x_1^2) = g(x_2^2), \text{ which becomes } x_1^2 = x_2^2$$

and hence gives $x_1 = x_2$ since x_1, x_2 are positive. Thus $g \circ f$ is injective. However, g is not injective since for instance $g(2) = 1 = g(3)$.

Warm-Up 2. We give an example of functions $f : X \to Y$ and $g : Y \to Z$ such that $g \circ f : X \to Z$ is surjective but for which f is not surjective. We know that in any such example, g will have to surjective since surjectivity of $g \circ f$ implies surjectivity of g. If we want f to not be surjective, let us start with some non-surjective function; say $f : \mathbb{R} \to \mathbb{R}$ defined by $f(x) = x^2$. Now, we think of an "easy" way to guarantee that $g \circ f$ will be surjective.

Surjectivity means that the image should equal the entirety of the codomain, so a simple way to make a function $X \to Z$ surjective is for Z to only consist of a single element. Then the definition surjective only asks something of this single element, so it will automatically be satisfied. Thus, take $g : \mathbb{R} \to \{0\}$ to be the function sending everything to 0. Then $g \circ f : \mathbb{R} \to \{0\}$ is also the function which sends everything to 0, so it is surjective, but f is not as required.

Invertible functions. A function $f : A \to B$ is called *invertible* if there exists a function $g : B \to A$ such that

$$g \circ f = id_A \text{ and } f \circ g = id_B,$$

where id_A and id_B denote the *identity* functions on A and B respectively, which are functions which send any element to itself. Concretely, this means that

$$g(f(a)) = a \text{ for all } a \in A \text{ and } f(g(b)) = b \text{ for all } b \in B.$$

When it exists, we call g the *inverse* of f and denote it by $g = f^{-1}$. The idea is that the inverse function is the function which does the "opposite" of what f does, meaning that whenever f sends $a \in A$ to $b \in B$, f^{-1} should send b back to a.

Example. Let $\mathbb{R}^+$ denote the set of positive real numbers. The function $f : \mathbb{R} \to \mathbb{R}^+$ given by $f(x) = e^x$ is invertible with inverse $f^{-1} : \mathbb{R}^+ \to \mathbb{R}$ defined by $f^{-1}(y) = \ln y$. Note that if we instead consider f as a function $\mathbb{R} \to \mathbb{R}$ with codomain $\mathbb{R}$, then it is not invertible since $\ln y$ is not defined by nonpositive y, meaning that the candidate for f^{-1} does not have domain $\mathbb{R}$.

Along these lines, consider the function $f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ defined by $f(x) = x^2$. This function is invertible, with inverse $f^{-1}(y) = \sqrt{y}$. However, alternating the domain or codomain of f may lead to a non-invertible function; for instance, the same $f(x) = x^2$ but considered as a function $\mathbb{R} \to \mathbb{R}$, $\mathbb{R} \to \mathbb{R}_{\geq 0}$, or $\mathbb{R}_{\geq 0} \to \mathbb{R}$ is not invertible. The function $g : \mathbb{R}_{\leq 0} \to \mathbb{R}_{\geq 0}$ defined by $g(x) = x^2$ *is* invertible with inverse $g^{-1}(y) = -\sqrt{y}$. The upshot is that domains and codomains matter.

The requirements in the definition of invertible. Here is an example showing that both requirements $g \circ f = id_A$ and $f \circ g = id_B$ are necessary in the definition of invertible, meaning that one does not automatically imply the other. Take $L, R : \mathbb{R}^\infty \to \mathbb{R}^\infty$ to be the left-shift and right-shift functions. We previously worked out that $L \circ R$ is the identity function on $\mathbb{R}^\infty$, but that $R \circ L$ is the function which replaces the first component of an element of $\mathbb{R}^\infty$ with 0, so that $R \circ L$ is not the identity function.

The point is that here neither L nor R are invertible, even though composing them in one order (but not the other) does happen to give the identity function. Technically, so far we've only shown that L and R are not inverses of one another, but we haven't ruled out the possibility that L could have some other inverse that wasn't R, or that R could have an inverse which wasn't L. That this is not possible will follow from our next fact that an invertible function must always be bijective: neither L nor R is bijective, so neither is invertible.

Invertible implies bijective. We now work towards understanding what types of functions do have inverses. First, we claim that if a function $f : A \to B$ is invertible, it must be bijective. Indeed, we get this right away based on things we've done previously: since $f \circ f^{-1} = id_B$ is surjective, f must be surjective, and since $f^{-1} \circ f = id_A$ is injective, f must be injective. Thus f is bijective. This then immediately tells us, as we pointed out before, that the functions shifting functions $L, R : \mathbb{R}^\infty \to \mathbb{R}^\infty$ are not invertible.

Bijective implies invertible. The claim is that, actually, a function $f : A \to B$ is invertible *if and only if* it is bijective. We showed the forward direction above, so we have only the backwards direction left to show.

Thus, supposing f is invertible, we must actually *construct* an inverse function $f^{-1} : B \to A$. This function should, whenever $f(a) = b$, send b back to a. But we've already seen a glimpse of this previously: if f is bijective, for any $b \in B$, there exists a unique $a \in A$ such that $f(a) = b$. Indeed, the existence comes from surjectivity of f, and the uniqueness from injectivity. Thus, we simply *define* $f^{-1} : B \to A$ to be the function sending an element of B to be this unique element of A:

$$f^{-1}(b) = \text{ the unique element } a \in A \text{ satisfying } f(a) = b.$$

With this definition, we have

$$f(f^{-1}(b)) = b$$

since, by definition, $f^{-1}(b) \in A$ has the property that applying f to it gives b, and

$$f^{-1}(f(a)) = a$$

since $f^{-1}(f(a))$ should be the unique $x \in A$ satisfying $f(x) = f(a)$, which requires that $x = a$.

So, "invertible" and "bijective" mean the same thing. A key observation, which is useful when wanting to compute an inverse explicitly, is that $f^{-1}(b) = a$ means precisely the same thing as $f(a) = b$. Thus, finding $f^{-1}(b)$ means solving $f(a) = b$ for a.

Final example. We show that the function $f : \mathbb{R}^3 \to \mathbb{R}^3$ defined by

$$f(x, y, z) = (x, x + y, x + y + z)$$

is invertible by explicitly finding its inverse. A homework problem asked to show that this function is bijective, which gives us one way of showing that it is invertible, but doesn't tell us on its own what the inverse is. The inverse function $f^{-1} : \mathbb{R}^3 \to \mathbb{R}^3$ should have the property that

$$f^{-1}(a, b, c) = (x, y, z) \text{ when } f(x, y, z) = (a, b, c).$$

Thus, to describe f^{-1}, given $(a, b, c) \in \mathbb{R}^3$ we must find $(x, y, z) \in \mathbb{R}^3$ satisfying $f(x, y, z) = (a, b, c)$, which for this particular function means:

$$(x, x + y, x + y + z) = (a, b, c).$$

Equating components gives an equation we can use to solve for x, y, z in terms of a, b, c, where we get

$$x = a, \ y = b - a, \text{ and } z = c - b.$$

Thus, the inverse of f should be the function $f^{-1} : \mathbb{R}^3 \to \mathbb{R}^3$ defined by

$$f^{-1}(a, b, c) = (a, b - a, c - b).$$

We verify that this is indeed the inverse by computing both $f \circ f^{-1}$ and $f^{-1} \circ f$:

$$\begin{aligned}(f \circ f^{-1})(a, b, c) &= f(f^{-1}(a, b, c)) \\ &= f(a, b - a, c - b) \\ &= (a, a + (b - a), a + (b - a) + (c - b)) \\ &= (a, b, c)\end{aligned}$$

and

$$\begin{aligned}(f^{-1} \circ f)(x, y, z) &= f^{-1}(f(x, y, z)) \\ &= f^{-1}(x, x + y, x + y + z) \\ &= (x, (x + y) - x, (x + y + z) - (x + y)) \\ &= (x, y, z).\end{aligned}$$

Hence $f^{-1} \circ f$ and $f \circ f^{-1}$ are both the identity functions on $\mathbb{R}^3$, so f^{-1} is the inverse of f and f is invertible as claimed.

Lecture 19: Equivalence Relations

Warm-Up 1. We find the inverse of the function $f : (\mathbb{R} - \{1\}) \times \mathbb{R} \to (\mathbb{R} - \{2\}) \times \mathbb{R}$ defined by

$$f(x, y) = \left(\frac{2x + 1}{x - 1}, x + y \right),$$

thereby verifying that f is indeed invertible. The idea, and why one should expect this function to indeed be invertible, is that we can use the first component f to determine the value of x which gives a specified output, and then use the second component to determine the value of y required. So, given any $(a, b) \in (\mathbb{R}-\{2\})\times\mathbb{R}$, it seems that it should be possible to find $(x, y) \in (\mathbb{R}-\{1\})\times\mathbb{R}$ satisfying $f(x, y) = (a, b)$. (We'll see where the restriction that the codomain be $(\mathbb{R} - \{2\}) \times \mathbb{R}$ comes from in the course of working this out.)

So, given $(a, b) \in (\mathbb{R} - \{2\} \times \mathbb{R})$ we need to find $(x, y) \in (\mathbb{R} - \{1\}) \times \mathbb{R}$ such that

$$f(x, y) = \left(\frac{2x + 1}{x - 1}, x + y\right) = (a, b),$$

which requires

$$\frac{2x + 1}{x - 1} = a \text{ and } x + y = b.$$

The equation gives $2x + 1 = ax - a$, so

$$x = \frac{-a - 1}{2 - a} = \frac{a + 1}{a - 2}.$$

(Note that here is we need $a \neq 2$ in order to ensure this fraction exists.) Then the second equation gives

$$y = b - x = b - \frac{a + 1}{a - 2}.$$

Thus, the (x, y) needed in order to satisfy $f(x, y) = (a, b)$ is

$$(x, y) = \left(\frac{a + 1}{a - 2}, b - \frac{a + 1}{a - 2}\right),$$

so the inverse of f is the function $f^{-1} : (\mathbb{R} - \{2\}) \times \mathbb{R} \to (\mathbb{R} - \{1\}) \times \mathbb{R}$ defined by

$$f^{-1}(a, b) = \left(\frac{a + 1}{a - 2}, b - \frac{a + 1}{a - 2}\right).$$

You should be able to verify that $f \circ f^{-1}$ and $f^{-1} \circ f$ are both identity functions.

Warm-Up 2. To see an instance of the types of things we can do when we do have an inverse, recall the fact that if $g \circ f$ is injective, it is not true that g must be injective. However, we now claim that if $g \circ f$ is injective *and* f is invertible, then g indeed must be injective. Just to be clear, say that f maps $A \to B$ and g maps $B \to C$.

Now, an observation we made previously is that the knowledge that $g \circ f$ is injective only tells us something about expressions of the form $g(f(a))$, but in order to show that g is injective we work with expressions of the form $g(b)$ for, at first glance, any possible $b \in B$. Indeed, supposing $g(b) = g(b')$, we must show that $b = b'$. The point is that we have to figure out a way to work f into this equality if we want to have any hope of making use of the fact that $g \circ f$ is injective.

But now the invertibility of f comes to the rescue! Since f^{-1} exists, we can write b and b' as

$$b = f(f^{-1}(b)) \text{ and } b' = f(f^{-1}(b')).$$

With this the equality $g(b) = g(b')$ becomes

$$g(f(f^{-1}(b))) = g(f(f^{-1}(b'))).$$

This is in the form needed to make use of injectivity of $g \circ f$, so we get that

$$f^{-1}(b) = f^{-1}(b'),$$

or in other words that the input $f^{-1}(b)$ into $g \circ f$ on the left is the same as the input $f^{-1}(b')$ into $g \circ f$ on the right. Finally, applying f to both sides of this new equality gives

$$f(f^{-1}(b)) = f(f^{-1}(b')),$$

which becomes $b = b'$. Thus g is injective as claimed.

Relations. Relations, and the more important notion of an *equivalence relation*, are covered in Chapter 11 in the book, before the chapter functions. I mentioned previously that the book approaches the notion of a function *using* the notion of a relation, which I find obscures the point behind functions. So, we now backtrack and say a bit about relations.

I'll preface this entire discussion by saying that it is unlikely you'll actually explicitly use the notion of a relation (or of an equivalence relation) in future courses; it is definitely true that many concepts you'll see in later courses—in particular albstract algebra, number theory, and topology—*can* be phrased in terms of equivalence relations, but the point is that this is not how these concepts are usually presented. So, you should not view our discussion of equivalence relations as essential to other courses, but rather as a way to give us more practice in working with definitions and forming mathematical arguments. Nevertheless, we will say a bit more about next time about the types of things you can do with the notion of an equivalence relation.

A *relation* on a set A is a subset of $A \times A$. This definition doesn't seem all that interesting, and indeed it's really not, except for the fact that it gives a way of formalizing various types of "relationships" one comes across in mathematics.

Examples. The best way to get a handle on what the concept of a relation is meant to do is to jump into some well-known examples. For instance, we can talk about the relation $\leq$ on $\mathbb{R}$. This particular relation captures the idea that one number can be smaller than or equal to another, or in other words that two numbers are "related" by this inequality. Formally, as a subset $\mathbb{R} \times \mathbb{R}$, the relation of which we are speaking is:

$$\{(x, y) \in \mathbb{R} \times \mathbb{R} \mid x \leq y\}.$$

To say that (x, y) is in this subset is to say that $x \leq y$. The point being, as mentioned previously, that this particular subset gives a formal way of talking about $\leq$ in set-theoretic terms. However, it is crucial to recognize that what is important here is the relationship $\leq$ itself, and not so much the subset used in the formal definition of "relation".

Here is another example, dealing with the relation of "divides" on $\mathbb{Z}$. Here, the relationship between two integers a, b we care about is that of a dividing b, so we say that "a is related to b" under this relation when a divides b. Formally, the subset in question is

$$\{(a, b) \in \mathbb{Z} \times \mathbb{Z} \mid a \text{ divides } b\}.$$

Again, it is better to place the focus on the notion of "divides" than on this particular subset itself.

Forgetting subsets. So, going forward, we will talk about relations in the sense of one element being "related" to another, where "related" depends on the specific relation being used. That is, we won't talk about a relation on A anymore as being defined by a certain subset of $A \times A$, but rather as a way to specify which elements of A should be related.

On top of this, really the only type of relations we'll actually care about are equivalence relations, which we now define. With this in mind, we will use the notation $\sim$ when denoting a relation, and write $a \sim b$ when a is related to b in whatever sense we're talking about. (The notation $\sim$ is reserved for equivalence relations, and $a \sim b$ is usually read as saying that "a is equivalent to b" under the equivalence relation being considered.)

Equivalence relations. Here are three properties a relation $\sim$ on a set A might have:

- (the reflexive property) For all $a \in A$, $a \sim a$.
- (the symmetric property) If $a \sim b$, then $b \sim a$.
- (the transitive property) If $a \sim b$ and $b \sim c$, then $a \sim c$.

The first says that anything in A is related (or equivalent) to itself; the second says that the order in which two elements are related to one another doesn't matter, so two elements are either equivalent or they're not; and the third says that two elements equivalent to a common element are actually equivalent to one another. We say that $\sim$ is an *equivalence relation* when it is reflexive, symmetric, and transitive.

Think of these three properties as capturing a type of "equality": certainly, anything should be equal to itself, equality shouldn't depend on ordering, and two things equal to a common thing are equal to each other. Of course, equivalence relations are NOT equalities, but they do *lead* to certain equalities among so-called *equivalence classes*, which we'll soon define. The overarching idea is that equivalence relations give a way to say that two non-equal things really should be thought of as being the "same", in the sense that they capture the same information in the context of whatever property in which we are interested. This is a vague description to be sure, but one which should become clearer next time.

Non-examples. The examples of $\leq$ and "divides" we mentioned previously are not equivalence relations since they are not symmetric. (They are, however, both reflexive and transitive: $x \leq x$ is true for all $x \in \mathbb{R}$, and $x \leq y$ and $y \leq z$ does imply $x \leq z$; and, "any integer divides itself" is true, and a divides b and b divides c does imply that a divides c.) The relation $\leq$ is not symmetric since $1 \leq 5$ but $5 \not\leq 1$, and divisibility is not symmetric since 2 divides 4 but 4 does not divide 2.

The relation $<$ of being strictly less than is still transitive, but is no longer reflexive. In general, none of the three properties in the definition of an equivalence relation imply the others.

Congruence mod 4. As a first example of something which *is* an equivalence relation, we define the notion of "congruence mod 4". (We chose 4 just to be specific, but more generally one can define the notion of congruence mod n for any $n \in \mathbb{N}$.)

Define $\sim$ on $\mathbb{Z}$ by saying $a \sim b$ if $a - b$ is divisible by 4. Thus, two elements are related (or "equivalent") under this relation if their difference is a multiple of 4. First let's verify that this is an equivalence relation. For any $n \in \mathbb{Z}$, $n - n = 0$ is divisible by 4, so $n \times n$ and hence $\sim$ is reflexive. (Don't get lost in all the symbols: to say that $n \times n$ means by definition of this equivalence relation that $n - n$ should be divisible by 4. The point is that whenever we see "$a \sim b$" when talking about a specific equivalence relation, we have to interpret its meaning based on that specific relation itself.)

To see that $\sim$ is symmetric, suppose $a \sim b$. This means that $a - b$ is divisible by 4, so $a - b = 4k$ for some $k \in \mathbb{Z}$. In order to be able to conclude that $b \sim a$ we would need to know that $b - a$ is also divisible by 4. But multiplying $a - b = 4k$ by -1 gives $b - a = 4(-k)$, which shows that $b - a$ is indeed divisible by 4 and hence that $b \sim a$ as required. Thus $\sim$ is symmetric. To check transitivity,

suppose $a \sim b$ and $b \sim c$ for some $a, b, c \in \mathbb{Z}$. To say that $a \sim b$ means $a - b$ is divisible by 4, and to say $b \sim c$ means $b - c$ is divisible by 4, so

$$a - b = 4k \text{ and } b - c = 4\ell$$

for some $k, \ell \in \mathbb{Z}$. The question is whether $a - c$ is divisible by 4 (which is what $a \sim c$ means in this context), but:

$$a - c = (a - b) + (b - c) = 4k + 4\ell = 4(k + \ell),$$

so $a - c$ is indeed divisible by 4. Thus $a \sim b$ and $b \sim c$ implies $a \sim c$, so $\sim$ is transitive. Hence $\sim$ is an equivalence relation as claimed.

In this specific setting, $a \sim b$ means that a and b are "congruent mod 4", and we write

$$a \equiv b \bmod 4$$

to indicate this. The resulting theory of "modular arithmetic" plays a big role in number theory and abstract algebra, and is something we'll briefly touch on next time.

Equivalence classes. Given an equivalence relation $\sim$ on A, we define the *equivalence class* of $a \in A$ to be the set $[a]$ of all things in A which are equivalent to A:

$$[a] = \{b \in A \mid b \sim a\}.$$

For instance, let us work out the equivalence classes of the "congruent mod 4" equivalence relation introduced above. To start with, the equivalence class of 0 consists of all things equivalent to 0:

$$[0] = \{n \in \mathbb{Z} \mid n \sim 0\}.$$

But according to the definition of $\sim$ in this case, $n \sim 0$ means that $n - 0$ is divisible by 4, which thus means that n should be divisible by 4. Thus n is in the equivalence class of 0 if and only if n is a multiple of 4:

$$[0] = \{4k \mid k \in \mathbb{Z}\} = \{\ldots, -12, -8, -4, 0, 4, 8, 12, \ldots\}.$$

The equivalence class of [1] consists of all n such that $n \sim 1$. But $n \sim 1$ means that $n - 1$ is divisible by 4, so $n - 1 = 4k$ for some $k \in \mathbb{Z}$. Hence $n = 4k + 1$, so things in the equivalence class of 1 are integers of the form $4k + 1$:

$$[1] = \{4k + 1 \mid k \in \mathbb{Z}\} = \{\ldots, -7, -3, 1, 5, 9, \ldots\}.$$

To say that $n \sim 2$ means $n - 2 = 4k$ for some $k \in \mathbb{Z}$, so $n = 4k + 2$ and hence

$$[2] = \{4k + 2 \mid k \in \mathbb{Z}\} = \{\ldots, -6, -2, 2, 6, 10, \ldots\}.$$

Similarly, the equivalence class of 3 consists of all things of the form $4k + 3$:

$$[3] = \{4k + 3 \mid k \in \mathbb{Z}\} = \{\ldots, -5, -1, 3, 7, 11, \ldots\}.$$

Now, what about the equivalence class of 4? To say that $n \sim 4$ means $n - 4 = 4k$ for some $k \in \mathbb{Z}$. But this is the same as saying that $n = 4(k + 1)$ is a multiple of 4, so we get that things equivalence to 4 are the same as things equivalent to 0, so $[4] = [0]$. Similarly, $n \sim 5$ means $n - 5 = 4k$ for some $k \in \mathbb{Z}$, which gives $n = 4(k + 1) + 1$, an expression indicating a number equivalent to 1. Thus $[5] = [1]$, and continuing on you'll see that

$$[6] = [2],\ [7] = [3],\ [8] = [4] = [0],\ \text{and so on.}$$

The conclusion is that this particular equivalence relation only has four distinct equivalence classes: the set $[0]$ of multiples of 4, the set $[1]$ of integers of the form $4k = 2$, those $[2]$ of the form $4k = 2$, and those $[3]$ of the form $4k + 3$. The equivalence class of any other integer will be the same as one of these four. Indeed, it is no accident that we get these redundancies: for instance, anything equivalent to 0 gives the same equivalence class as 0, so even without determining what things in $[4]$ look like separately, we could have said $[4] = [0]$ solely from the fact that $4 \in [0]$. This reflects a general property of equivalence classes we'll prove next time: $[a] = [b]$ if and only if $a \sim b$.

Equivalent things are the "same". To finish up for now, we come back to the idea that "equivalent" things should be thought of as being the "same" in a certain context. The precise version is that is, as alluded to above, equivalent things determine *equal* equivalence classes, so that things which are equivalent but non-equal to start with in a sense "become" equal after passing to equivalence classes. The point is that things in the same equivalence class are meant to share something in common, or are meant to capture the same type of information. In the congruence mod 4 case, the thing which all elements in the same equivalence class have in common is that they all give the same *remainder* when dividing by 4; if we were in a world where we only ever cared about remainders upon division by 4, then $1, 5, 9$, and anything of the form $4k + 1$ should indeed be thought of as being the "same" since they are *indistinguishable* from the point of view of taking remainders mod 4. This is the key concept behind "modular arithmetic", and we'll come back to this general point of view in various examples next time.

Lecture 20: More on Equivalences

Warm-Up. Let $\mathbb{Z}^*$ denote the set of nonzero integers. Define a relation $\sim$ on $\mathbb{Z} \times \mathbb{Z}^*$ by saying

$$(a, b) \sim (c, d) \text{ if } ad = bc.$$

We show that this is an equivalence relation. First, for any $(a, b) \in \mathbb{Z} \times \mathbb{Z}^*$ we have $ab = ba$, which is the requirement needed in order to say that $(a, b) \sim (a, b)$. Hence $\sim$ is reflexive. If $(a, b) \sim (c, d)$, then $ad = bc$, so $cb = da$ and thus $(c, d) \sim (a, b)$, so $\sim$ is symmetric. Finally, suppose $(a, b) \sim (c, d)$ and $(c, d) \sim (x, y)$. Then

$$ad = bc \text{ and } cy = dx.$$

Note that in order to conclude $(a, b) \sim (x, y)$ as transitivity requires, we need to know that $ay = bx$. Since $(a, b) \in \mathbb{Z} \times \mathbb{Z}^*$, $b \neq 0$, so $ad = bc$ gives $c = \frac{ad}{b}$. Then

$$dx = cy = \frac{ady}{b}, \text{ so } bdx = ady.$$

Since $(c, d) \in \mathbb{Z} \times \mathbb{Z}^*$, $d \neq 0$ so after dividing by d this final equality gives

$$bx = ay,$$

and hence $(a, b) \sim (x, y)$. Thus $\sim$ is transitive, so $\sim$ is an equivalence relation.

Note that the fact elements of $\mathbb{Z} \times \mathbb{Z}^*$ have nonzero second component was important here, since we needed to divide by such elements in the manipulations above. Indeed, the same relation only defined on $\mathbb{Z} \times \mathbb{Z}$ would *not* be transitive and hence would not be an equivalence relation; for instance, in the $\mathbb{Z} \times \mathbb{Z}$ case, $(1, 1) \sim (0, 0)$ and $(0, 0) \sim (3, 4)$ but $(1, 1) \not\sim (3, 4)$.

Back to equivalence classes. Let us now think about the equivalence classes of the equivalence relation in the Warm-Up. First we determine $[(1, 1)]$, which is the set of all $(a, b) \in \mathbb{Z} \times \mathbb{Z}^*$ such

that $(a, b) \sim (1, 1)$. This is true when $a \cdot 1 = 1 \cdot b$, so when $a = b$. Thus the equivalence class of $(1, 1)$ consists of all pairs where the two coordinates are the same:

$$[(1, 1)] = \{(a, a) \mid a \neq 0\}.$$

Now consider $[(1, 2)]$. A pair (a, b) is equivalent to $(1, 2)$ when $a \cdot 2 = b \cdot 1$, so when $b = 2a$. Thus the equivalence class $[(1, 2]$ consists of all pairs where the second coordinate is twice the first:

$$[(1, 2)] = \{(a, 2a) \mid a \neq 0\}.$$

As another example, you can work out that the equivalence class of $(2, 1)$ consists of all pairs where the first coordinate is twice the second:

$$[(1, 2)] = \{(2b, b) \mid b \neq 0\}.$$

What property do elements from the same equivalence class share in common? The key observation is that, if we interpret a pair (a, b) as consisting of a numerator a and a denominator b, the elements from the same equivalence class are those pairs which characterize the *same* rational number! That is, the pairs (a, a) in the equivalence class of $(1, 1)$ are those give the possible numerators and denominators of the rational number $\frac{1}{1} = 1$; the pairs $(a, 2a)$ in the equivalence class of $(1, 2)$ are those which give the rational numbers $\frac{1}{2}$; and the pairs in the equivalence class of $(2, 1)$ are those which give the rational number $\frac{2}{1} = 2$. In general, the pairs in the equivalence class of (a, b) are those which give the rational number $\frac{a}{b}$. Indeed, the condition $ad = bc$ defining this equivalence relation comes from rewriting the equality

$$\frac{a}{b} = \frac{c}{d}$$

of rational numbers in a way which does not make reference to division. The idea is that different pairs of numerators and denominators can give the same rational number, and so this equivalence relation imposes a type of "equivalence" on those which do so that they become "equal" after taking equivalence classes.

Constructing $\mathbb{Q}$. The considerations above are meant to give a way of "constructing" the set of rational numbers from the set of integers, where we essentially *define* $\mathbb{Q}$ to be the set of equivalence classes of the equivalence relation defined above on $\mathbb{Z} \times \mathbb{Z}^*$. The fact that $[(a, b)] = [(c, d)]$ precisely when $\frac{a}{b} = \frac{c}{d}$ says that equality of these equivalence classes agrees with the ordinary notion of equality of rational numbers we already have in mind.

Why would we want to do this? We won't go into this much in this course, but this discussion is all rooted in the desire to give precise meaning to all constructions one encounters in mathematics. Defining $\mathbb{Q}$ to be the set of fractions of integers:

$$\mathbb{Q} = \left\{ \frac{a}{b} \;\middle|\; a, b \in \mathbb{Z} \text{ and } b \neq 0 \right\}$$

is ambiguous unless we specify what the fraction $\frac{a}{b}$ actually means, and the problem is that to do so requires the notion of division, which implicitly assumes we already know what rational numbers are. The approach to defining $\mathbb{Q}$ via an equivalence relation, using the fact that $\frac{a}{b} = \frac{c}{d}$ can be rephrased as $ad = bc$, avoids this circularity. Of course, in practice we never actually think of rational numbers as representing equivalence classes—doing so would make working with them quite cumbersome. The point is that we only take this approach for the purpose of giving a rigorous definition to the term "rational number", but after having do so we simply work with rationals as

we're used to. This is an idea which shows up elsewhere in math, but which you probably won't see much of in your undergraduate careers.

Theorem. We now prove that, given an equivalence relation $\sim$ on a set A, two elements $a, b \in A$ determine the same equivalence class if and only if they are equivalent to one another: $[a] = [b]$ if and only if $a \sim b$. This is the basis behind the idea that equivalent things should be thought of as being the "same" in the given context since they determine the same equivalence class. To clarify the notation, an equivalence class is defined as $[a] = \{c \in A \mid c \sim a\}$.

Proof. Suppose $[a] = [b]$. Then since $a \in [a]$ (because $a \sim a$ given that $\sim$ is reflexive), we must also have $a \in [b]$. But this means that $a \sim b$, as claimed.

Conversely, suppose $a \sim b$. We must show that $[a] \subseteq [b]$ and $[b] \subseteq [a]$. To this end, let $c \in [a]$. Then $c \sim a$. Since $c \sim a$ and $a \sim b$, transitivity of $\sim$ gives $c \sim b$, so $c \in [b]$. Hence $[a] \subseteq [b]$. If instead $c \in [b]$, then $c \sim b$; since $a \sim b$, we also have $b \sim a$ by symmetry of $\sim$, and then $c \sim b$ and $b \sim a$ imply $c \sim a$ by transitivity. Hence $c \in [a]$, so $[b] \subseteq [a]$. Thus $[a] = [b]$ as claimed. $\square$

Final example. We look at one final example, where again the emphasis is on understanding the *set* of equivalence classes itself. Consider the relation $\sim$ on $\mathbb{R}$ given by

$$x \sim y \text{ if } x - y \in \mathbb{Z}.$$

A problem from Discussion 5 shows that this is an equivalence relation; we are interested in understanding the equivalence classes.

For a fixed nonnegative $y \in \mathbb{R}$, $[y]$ consists of those $x \in \mathbb{R}$ for which $x - y \in \mathbb{Z}$. Denoting this integer by k, we have

$$x = y + k \text{ for some } k \in \mathbb{Z},$$

so $[y]$ is the set of all numbers obtained by adding integers to y. The key observation is that, if we consider the decimal expansion of y:

$$y = N.y_1y_2y_3\ldots$$

where N is a nonnegative integer and each y_i a digit between 0 and 9, $y+k$ will only modify N but will maintain the same "decimal part" $0.y_1y_2y_3\ldots$. Thus $[y]$, for y nonnegative, consists of all real numbers which have the same decimal part as y. For instance, $[0.5]$ contains $2.5, 2283.5, -822.5$, and so on. This equivalence relation declares all such numbers to be the "same" in the sense of having the same decimal portion. As another example, $[0]$ is the set of all real numbers having decimal part 0.0, so in other words the set of integers.

By adding or subtracting 1 enough times, we see that any real number will be equivalent to some y between 0 and 1 inclusive, so any real number determines the same equivalence class as some $y \in [0, 1]$. Thus, we can describe all equivalence classes by focusing only on those determined by such y:

$$\text{set of equivalence classes } = \{[y] \mid y \in [0, 1]\}.$$

Moreover, numbers $0 < y < 1$ determine *different* equivalence classes since such numbers will never be equivalent to one another since their difference cannot be an integer. However, 0 and 1 determine the same equivalence class since $0 - 1$ is an integer. Thus, we can think of the set of equivalence classes as being analogous to the interval $[0, 1]$ only that think of the endpoints 0 and 1 as being the "same". Taking such an interval, visualized as a line segment, and "gluing" together the endpoints results in a circle, so the upshot is that, in some sense, we can think of the set of equivalence classes for this equivalence relation as being a circle!

This is all a lot to digest, but hints at a key way in which equivalence relations show up in other contexts, that is in using them to "construct" other sets or geometric objects. The idea here is that we take the real number line, and glue each y to all things to which it is equivalence, the resulting "space" looks like a circle. We won't be considering such things in this course any further, but such constructions show up quite in often in geometry and topology. It might seem strange to think about a circle as being constructed via this particular equivalence relation, but this point of view as to what a circle is actually does help to clarify some facts in the theory of what are called *Fourier series*, which some of you might come across later on. Again, we won't really do anything with this, apart from pointing the idea that it is often the case that you can think about equivalence classes as being "parametrized" by elements of a more well-known set. We'll see an example of this in the Warm-Up next time as well.

Lecture 21: Cardinality

Warm-Up. Consider the equivalence relation $\sim$ on $\mathbb{R}^2$ defined by

$$(x, y) \sim (a, b) \text{ if } x^2 + y^2 = a^2 + b^2.$$

(This is indeed an equivalence relation, which you should be able to verify.) We find a bijection between the set of equivalence classes and a subset of $\mathbb{R}$.

First we determine the nature of the equivalence classes themselves. The equivalence class of, say, $(1, 1)$ consists of all (a, b) such that $a^2 + b^2 = 1^2 + 1^2 = 2$, meaning that (a, b) should lie on the circle of radius $\sqrt{2}$ centered at the origin. Thus $[(1, 1)]$ is this circle of radius $\sqrt{2}$, and similarly it turns out that for $(x, y) \neq (0, 0)$, the equivalence class $[(x, y)]$ is the circle of radius $\sqrt{x^2 + y^2}$ centered at the origin, since $(a, b) \sim (x, y)$ precisely when (a, b) and (x, y) determine the same radius via $a^2 + b^2 = x^2 + y^2$. The idea is that we consider points on the same circle to be the "same" in the context of characterizing the same circle centered at the origin. The only equivalence class which does not look like a circle is that of $(0, 0)$, since $(a, b) = (0, 0)$ is the only point satisfying $a^2 + b^2 = 0^2 + 0^2 = 0$, so that $[(0, 0)]$ consists of only $(0, 0)$. Thus, the equivalence classes of this equivalence relation are circles centered at the origin together with the singleton set $\{(0, 0)\}$:

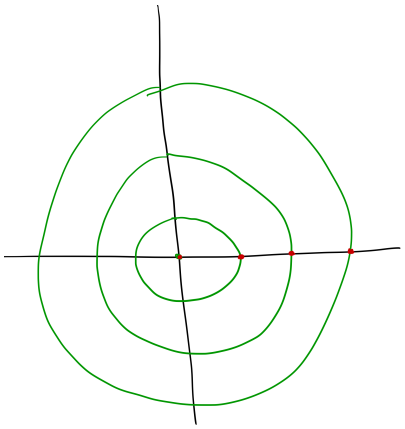

Now, we this in mind we come to our task of finding a bijection between the set of equivalence classes and a subset of $\mathbb{R}$. The question to think about is: how can be unique characterize the equivalence classes themselves using real numbers? In this case, the point is that we can associate to each equivalence class the *radius* of the circle it represents: each equivalence class determines a radius $r \geq 0$, and each such radius characterizes a single equivalence class. To think about

this another way visually, each equivalence class intersects the nonnegative x-axis in exactly one point, so the set of equivalence classes is in one-to-one correspondence with the set of points on the non-negative x-axis. (These are red dots in the picture above.) Thus, the function

$$\{\text{set of equivalence classes}\} \to [0, \infty)$$

defined by sending the equivalence class $\left\{(x, y) \;\middle|\; x^2 + y^2 = r^2\right\}$ to the number r:

$$\{(x, y) \mid x^2 + y^2 = r^2\} \mapsto r$$

gives the desired bijection. We would say that the set of equivalence classes is *parameterized* by numbers in the interval $[0, \infty)$ since to each such number corresponds a unique equivalence class.

Towards cardinality. Say, in the example above, we wanted to ask the question: how many equivalence classes are there? Of course, one answer is "infinitely many" since there are infinitely many circles centered at the origin. But we can give a more precise answer: there are as many equivalence classes as there are numbers in $[0, \infty)$. Indeed, the basic idea for everything we'll do in the remaining weeks is that the existence of a bijection between sets suggests that those sets should have the "same" "number" of elements.

To motivate this, we'll first look at the case of finite sets, where we recall some facts we previously derived. If A and B are finite sets, meaning sets with finitely many elements, we pointed out when discussing functions that having a injection $A \to B$ is equivalent to saying $|A| \leq |B|$, where $|S|$ denotes the number of elements in a set S, and that having a surjection $A \to B$ is equivalent to saying $|A| \geq |B|$. Putting these together gives that the existence of a bijection $A \to B$ is equivalent to saying $|A| = |B|$, so that, in the finite case, we can characterize whether or not two sets have the same number of elements in terms whether or not there exists a bijection between them.

Characterizing "same number of elements" in terms of the existence of a bijection might seem too abstract at first, but is actually the right way to go about it. For instance, suppose we had a set A with a million elements and another B, also with a million elements. How could we tell that these had the same number of elements? Here is the wrong way to do it: sit down and count up all the elements of A ("one, two, three, ...") and when you finish a month (!) later go count up all the elements of B. After two months you will be able to tell that, yes, A has the same number of elements as B does. This is not very efficient, and gets harder to do with larger sets, let alone infinite sets.

Instead, we can argue that A and B have the same number of elements by showing that there is some way of pairing off elements of A with elements of B in a one-to-one manner so that there is nothing left over in either set; if such a "pairing" of elements is possible, it should be the case that there were as many things in A as in B. But such a "pairing" is precisely what a bijection between A and B gives us: if $f : A \to B$ is bijective, we pair off an element $a \in A$ with its image $f(a) \in B$. The fact that f is surjective says that each element of B is paired off with something, and f injective says that each element of B is paired off with only one thing. Thus, in order for A and B to have the same number of elements, it must be possible to construct such a bijection.

The key point now is that this same intuition works even for infinite sets. We cannot literally sit down and count up all the elements in an infinite set (we would never finish!), but we *can* argue that one infinite set should have the same number of elements as another by showing that we can come up with a bijection between them.

Cardinality. We say that two sets A and B have the same *cardinality* if there exists a bijection from A to B. We use the notation $|A| = |B|$ to indicate that A and B have the same cardinality,

but we should be careful with what we actually mean by this notation: if A is finite, $|A|$ literally denotes the number of elements in A, which will be given by some nonnegative integer, but if A is infinite we are not using $|A|$ to denote the number of elements of A—at least not yet; rather, in the infinite case, we only use $|A|$ in the context of asking whether A has the same cardinality as some other set, meaning that $|A|$ in the infinite case only appears as one half of $|A| = |B|$. Showing that $|A| = |B|$ *requires* us (for now) to come up with a bijection from A to B since this is what $|A| = |B|$ means *by definition*.

We will, soon enough, attempt to assign some independent meaning to the notation $|A|$ in the infinite case, by introducing a new type of "number" which will allow us to think about the number of elements in an infinite set in a more formal way. The upshot will be that different infinite sets can have different "sizes" if we interpret "size" as "number of elements".

Cardinality of $\mathbb{Z}$. We claim that $\mathbb{Z}$ has the same cardinality as $\mathbb{N}$. This is actually something we argued back on the first day of class in the context of alluding to the types of things we would eventually be able to talk about, so here we are. To be clear, we claim that there is a bijection $\mathbb{N} \to \mathbb{Z}$, which is what it means to say that $|\mathbb{N}| = |\mathbb{Z}|$. To come up with this bijection, list the elements of $\mathbb{Z}$ in the following way:

$$0 \qquad 1 \qquad -1 \qquad 2 \qquad -2 \qquad 3 \qquad -3 \qquad \dots$$

where, after the initial zero, we list each positive integer following by its negative. Define $f : \mathbb{N} \to \mathbb{Z}$ by declaring $f(n)$ to be equal to the n-th element in this list, so $f(1) = 0, f(2) = -1$, and so on:

$$\begin{array}{cccccccc} 0 & 1 & -1 & 2 & -2 & 3 & -3 & \dots \\ f(1) & f(2) & f(3) & f(4) & f(5) & f(6) & f(7) & \dots \end{array}$$

The fact that every element of $\mathbb{Z}$ appears somewhere in the first list says that f is surjective, and that fact that each elements only appears once says that f is injective, so it is bijective. Hence $\mathbb{Z}$ has the same cardinality as $\mathbb{N}$ as claimed.

Thus, intuitively, there should be as many elements of $\mathbb{Z}$ as there are elements of $\mathbb{N}$, which can appear to be a strange thing to say! Indeed, clearly $\mathbb{N}$ is a *proper* subset of $\mathbb{Z}$, meaning a subset which is not equal to the larger set, so it would seem that there should be "more" things in $\mathbb{Z}$ than in $\mathbb{N}$; however, based on our precise definition of equal cardinality, this is not the case. In the finite case it is definitely true that a proper subset of a finite will not have the same cardinality as the larger set (it has cardinality which is "strictly less than" that of the larger set), but this is not so for infinite sets, suggesting that cardinalities of infinite sets can behave in ways we might not expect based on our intuition in dealing with finite sets. The upshot is that we have to be careful trying to applying what we know about finite sets to infinite ones, and that we must focus on the precise definition of cardinality to guide us through our endeavors.

Cardinality of intervals. Next we consider the cardinality of some intervals. For instance, we claim that $|(0, 1)| = |(0, 2)|$, so that the open intervals $(0, 1)$ and $(0, 2)$ have the same cardinality. Indeed, the function $f : (0, 1) \to (0, 2)$ defined by $f(x) = 2x$ is bijective, as you should be able to show, which proves our claim. Again, note that this can seem counterintuitive: $(0, 1)$ only makes up "half" of $(0, 2)$ in the sense of length, and yet we are saying that they have the "same" number of elements nonetheless. In effect, taking $(0, 1)$ and throwing in the infinitely many things in $[1, 2)$ doesn't actually increase the number of elements overall—amazing!

Now, $(0, 2)$ and $(3, 5)$ have the same cardinality since the function $g : (0, 2) \to (3, 5)$ defined by $g(x) = x + 3$ is bijective, so $|(0, 2)| = |(3, 5)|$. We can now ask: do $(0, 1)$ and $(3, 5)$ have the same

cardinality? Note that the function $(0,1) \to (3,5)$ defined by $x \mapsto 2x+1$ is bijective, so we indeed have $|(0,1)| = |(3,5)|$. But this function is precisely the composition $g \circ f : (0,1) \to (3,5)$ of the functions above, so the fact that is bijective follows immediately from the fact that the composition of bijective functions is always bijective. Thus we are saying that since $|(0,1)| = |(0,2)|$ and $|(0,2)| = |(3,5)|$, it follows that $|(0,1)| = |(3,5)|$.

Written in terms of this notation it seems obvious that if $|(0,1)|$ and $|(3,5)|$ are both equal to $|(0,2)|$ that they should be equal to one another, but recall that notation $|A| = |B|$ does not literally mean (in the finite case) that the size of A is equal to the size of B! The point is that since $|A| = |B|$ is defined in terms of whether or not a certain bijection exists, we have to be careful about treating it as an actual equality between the object $|A|$ and the object $|B|$. Still, the example of $(0,1)$ and $(3,5)$ of above suggests that this notation does seem to have at least one property an honest "equality" should have: if $|(0,1)|$ and $|(3,5)|$ are both "equal" to the same thing, they should be "equal" to each other.

Cardinality as an equivalence relation. In general, the reason why we use the notation $|A| = |B|$ to denote having the same cardinality is because it *does* behave as an ordinary equality should, which we now explain. First, if $|A| = |B|$ is to be interpreted as an actual equality, it should be true that $|A|$ is always equal to itself: $|A| = |A|$ for all A. But this notation says that there should exist a bijection from A to A, which there does: the identity function on A which sends anything in A to itself provides such a bijection. Thus, $|A| = |A|$ for any A, so the relation of having the same cardinality is actually *reflexive*.

Second, if $|A| = |B|$ is to be interpreted as an actual equality, it should be true that the order in which we list the two terms shouldn't matter: $|A| = |B|$ should imply $|B| = |A|$. But $|A| = |B|$ means that there exists a bijection $f : A \to B$, so in order to conclude that $|B| = |A|$ we would have to know that there exists a bijection $B \to A$... and there does! Since f is bijective, it is invertible, so the inverse function $f^{-1} : B \to A$ provides a bijection from B to A, so that $|A| = |B|$ does imply $|B| = |A|$. Hence the relation of having the same cardinality is *symmetric*.

Finally, we claim that the relation of having the same cardinality is *transitive*, as we alluded to in the interval example we looked at above. If $|A| = |B|$ and $|B| = |C|$, there exists bijections $f : A \to B$ and $g : B \to C$, and then $g \circ f : A \to C$ gives a bijection from A to C, so $|A| = |C|$. Thus, the relation is having the same cardinality is reflexive, symmetric, and transitive, so it is an equivalence relation. This is the key reason why we can, and will, interpret the notation $|A| = |B|$ as an actual equality, because it has the basic properties (reflexivity, symmetry, and transitivity) an equal should have. For us, the most beneficial property will be transitivity, since it gives us a way to show that two sets have the same cardinality without having to come up with an explicit bijection between them every single time—if we can show that they each have the same cardinality as some other common set, they themselves will then have the same cardinality as well.

Back to intervals. We showed previously that $(0,1)$, $(0,2)$, and $(3,5)$ all have the same cardinality, and in fact we now claim that all open intervals (as long as they are not empty) will have the same cardinality as one another. Indeed, take the open intervals (a,b) and (c,d) where $a < b$ and $c < d$. Draw the first on an x-axis and the second on a y-axis:

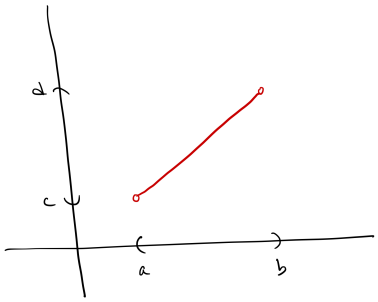

and draw the line segment connecting the point (a, c) to the point (b, d). The linear function which has this line segment as its graph will then be a bijection from (a, b) to (c, d), so $|(a, b)| = |(c, d)|$ as claimed. You can, of course, work out the explicit formula for this function by finding the equation of the line drawn above, but our picture should provide enough justification for the existence of such a function. (One thing to say about this final topic of cardinality is that we will now be a bit more lenient as to what constitutes a correct proof: as long as we can give an argument saying how we *can* do something precisely we'll consider that to often be good enough so that we don't get too bogged down in notation and overly granular details. Definitely, at this point is more important to understand the *concept* of cardinality and what it means rather than doing things at the level of formality we've done up until now. To be clear, you *should* be able to carry out this level of formality if asked to do so, but you won't always be expected to do so.)

So, all non-empty open intervals have the same cardinality as each other. Note that the function $f : (-\frac{\pi}{2}, \frac{\pi}{2}) \to \mathbb{R}$ defined by $f(x) = \tan x$ is bijective (with inverse $\arctan x$), so $\left(-\frac{\pi}{2}, \frac{\pi}{2}\right)$ and $\mathbb{R}$ have the same cardinality. Thus using transitivity we get that any non-empty open interval actually has the same cardinality as $\mathbb{R}$! Hence, for $a < b$, there are as many numbers in (a, b) as there are real numbers all together. Again, a possibly counterintuitive but true fact based on our definition of cardinality; the moral is: once we've given a certain definition, we must accept whatever consequences the mathematical gods force upon us.

Now, what about intervals which are closed, or half-open/half-closed? We claim, for instance, that $(0, 1)$ and $[0, 1)$ have the same cardinality! (So including one more number in $(0, 1)$ does not increase the number of elements you have.) But coming up with a bijection $[0, 1) \to (0, 1)$ is not as straightforward: all bijections we've described between intervals (including $\mathbb{R}$) have been given by a certain nice formula, such as $\tan x$ or some linear expression. Such "nice" functions won't work to give a bijection between $[0, 1)$ and $(0, 1)$, so we need something new. The idea is that, since 0 is the only thing in $[0, 1)$ missing from $(0, 1)$, we have to somehow "make room' available in $(0, 1)$ for this extra element. The simplest thing we could try would be to just send everything in $[0, 1)$ to itself, but of course this won't work since 0 cannot be sent to itself and still end up in $(0, 1)$.

Here is a strategy which does work: send 0 to, say, $\frac{1}{2}$. But then we cannot send $\frac{1}{2}$ to itself or else our function would not injective, so we must send $\frac{1}{2}$ somewhere else—send it to $\frac{1}{3}$. But then we need to send $\frac{1}{3}$ somewhere else, so send it to $\frac{1}{4}$, and so on. Thus, define $f : [0, 1) \to (0, 1)$ by saying

$$f(x) = \begin{cases} \frac{1}{2} & \text{if } x = 0 \\ \frac{1}{n+1} & \text{if } x = \frac{1}{n} \text{ for some } n \in \mathbb{N} \\ x & \text{otherwise.} \end{cases}$$

This function is bijective, which shows that $|[0, 1)| = |(0, 1)|$ as claimed. The idea is that we "shift" each fraction $\frac{1}{n}$ to the left a bit (which is the result of sending $\frac{1}{n}$ to $\frac{1}{n+1}$), thereby making an extra spot available where $\frac{1}{2}$ was where we can now insert the extra 0 we have available. (Note that it is not at all expected that you could have come up with this function completely on your own: this

is a key instance of where *now* that you've seen this type of argument once, you should keep it in mind as something you might be able to apply elsewhere.)

A similar strategy can be used to show that $(0, 1]$ has the same cardinality as $(0, 1)$, or that $[0, 1]$ has the same cardinality as $(0, 1)$. More generally, similar reasoning will show that (a, b), $[a, b)$, $(a, b]$, and $[a, b]$ all have the same cardinality when $a < b$. Combining this with what we know about open intervals and $\mathbb{R}$, we arrive at:

Theorem. Any interval (whether it be open, closed, half-open/half-closed) with more than one element has the same cardinality as $\mathbb{R}$. (The "with more than one element" condition is meant to exclude empty intervals like $(2, 1) = \emptyset$ or those which contain only a single point, like $[2, 2] = \{2\}$.)

The same is true of intervals which extend to infinity in one direction, like (a, ∞), $[a, \infty)$, $(-\infty, a)$, and $(-\infty, a]$, as you will show on a homework problem. Next time we'll come back to "simpler" cardinalities, like that of $\mathbb{Z}$ or $\mathbb{N}$.

Lecture 22: Countable Sets

Warm-Up. We show that the Cartesian product $[0, 1] \times [0, 1]$ has the same cardinality as $[0, 1]$. Visually, $[0, 1] \times [0, 1]$ is a square and $[0, 1]$ is a line segment, so this is another instance where it might seem counterintuitive to say that there as many points in the square as in the line segment, but there are. We must come up with a bijective function $f : [0, 1] \times [0, 1] \to [0, 1]$, which in the end will not be given by a simple "nice" formula, but requires more ingenuity.

The key fact we use is that any number between 0 and 1 has a decimal expansion of the form

$$0.x_1x_2x_3\ldots$$

where each x_i is a digit between 0 and 9 inclusive. (Note that 1 is of this form as well since $1 = 0.9999\ldots$.) So, a pair $(x, y) \in [0, 1] \times [0, 1]$ can be viewed as a pair of two such decimal expansions:

$$(0.x_1x_2x_3\ldots, 0.y_1y_2y_3\ldots).$$

To get a bijective function $f : [0, 1] \times [0, 1] \to [0, 1]$ we need a way to turn this pair of decimal expansions into a *single* decimal expansion in a way which, given the resulting output, would let us reconstruct the two original expansions from which it came. We can do this by "splicing" together the two original expansions as follows:

$$(0.x_1x_2x_3\ldots, 0.y_1y_2y_3\ldots) \mapsto 0.x_1y_1x_2y_2x_3y_3\ldots.$$

Thus, from $(0.x_1x_2x_3\ldots, 0.y_1y_2y_3\ldots)$ we create a single number by taking one digit from one or the other expansion, alternating, at a time. This gives a function $f : [0, 1] \times [0, 1] \to [0, 1]$, which is invertible with inverse $f^{-1} : [0, 1] \to [0, 1] \times [0, 1]$ given by

$$f^{-1}(0.b_1b_2b_3b_4\ldots) = (0.b_1b_3b_5\ldots, 0.b_2b_4b_6\ldots)$$

where for the first component of the output we take only the odd-indexed digits, and for the second the even-indexed digits. Both compositions $f \circ f^{-1}$ and $f^{-1} \circ f$ are indeed identities, so f is invertible and thus bijective, and hence $|[0, 1] \times [0, 1]| = |[0, 1]|$.

(To be precise, there are some subtle details in the above argument which need to be addressed, stemming from the fact that a given number can actually have to such decimal expansions; for instance, $0.0500000\ldots$ is the same as $0.049999\ldots$, which might cause trouble for injectivity of f.

There are various ways around this, say by saying that we always use decimal expansions which do not end in all 0's.)

Cardinality of $\mathbb{R}^n$. It is a fact you will prove on the homework that if A and C have the same cardinality and B and D have the same cardinality, then $A \times B$ and $C \times D$ have the same cardinality as well, which gives a nice way of determining cardinalities of various Cartesian products. (The intuition is that when considering elements (a, b) and (c, d) from $A \times B$ and $C \times D$ respectively, if there are as many possibilities for a as there are for c and as many possibilities for b as for d, then there should be as many pairs (a, b) as there are pairs (c, d).) As a consequence, since $|\mathbb{R}| = |[0, 1]|$ and $|[0, 1]^2| = |[0, 1]|$, we can immediately say that

$$|\mathbb{R}^2| = |\mathbb{R} \times \mathbb{R}| = |[0, 1] \times [0, 1]| = |[0, 1]| = |\mathbb{R}|,$$

so $\mathbb{R}^2$ and $\mathbb{R}$ have the same cardinality. But then

$$|\mathbb{R}^3| = |\mathbb{R}^2 \times \mathbb{R}| = |\mathbb{R} \times \mathbb{R}| = |\mathbb{R}|,$$

so $\mathbb{R}^3$ and $\mathbb{R}$ have the same cardinality. More generally, if $\mathbb{R}^n$ has the same cardinality as $\mathbb{R}$ for some n, then

$$|\mathbb{R}^{n+1}| = |\mathbb{R}^n \times \mathbb{R}| = |\mathbb{R} \times \mathbb{R}| = |\mathbb{R}|,$$

so induction gives that $\mathbb{R}^n$ has the same cardinality as $\mathbb{R}$ for all n. (In fact, you'll show on the homework that $\mathbb{R}^\infty$ also has the same cardinality as $\mathbb{R}$.)

This all illustrates a fact we mentioned last time: once we have enough facts about cardinality built up—such as the relation between products and cardinality, or transitivity of cardinalities—we get quick ways of showing that various sets have the same cardinality without having to explicitly construct bijections between them every single time.

Countable sets. A set S is said to be *countable* if it is finite or has the same cardinality as $\mathbb{N}$. In the case where S has the same cardinality as $\mathbb{N}$, we say that S is *countably infinite*. A set which is not countable is said to be *uncountable*.

The first basic example of a countable set, apart from $\mathbb{N}$ itself, is $\mathbb{Z}$, which we showed to be countable last time by showing that it had the same cardinality as $\mathbb{N}$. But, we recall the proof of this fact in order to see what it *really* means for a set to be countable. The idea is that we are able to come up with a list of all elements of $\mathbb{Z}$:

$$0 \quad 1 \quad -1 \quad 2 \quad -2 \quad 3 \quad -3 \quad \dots$$

and that from such a list we can come up with a bijection $\mathbb{N} \to \mathbb{Z}$ by sending n to the n-th term in this list. More generally now, for *any* infinite set S whose elements can be listed in such a way:

$$s_1 \quad s_2 \quad s_3 \quad s_4 \quad \dots$$

we can construct a bijection $\mathbb{N} \to S$ by sending n to the n-th term in this list. Thus any such set S will be countable.

Conversely, if S is countable, so that there is a bijection $f : \mathbb{N} \to S$, we can come up with a list containing all elements of S via:

$$f(1) \quad f(2) \quad f(3) \quad f(4) \quad \dots.$$

The fact that f is surjective says that each element of S will appear somewhere in this list, and the fact that f is injective says that each element of S only appears once this this list. Thus, we

conclude that S is countable if and only if it's elements can be listed in either a finite list (when S is finite) or an infinite list (when S is countable infinite). In other words, countable sets are sets whose elements we can "count", so that if we began counting its elements we would—even if it took us an infinite amount of time—be able to finish. By contrast, uncountable sets are in a sense "too large" to count in this manner; we'll talk about uncountable sets in more detail later.

So, in order to show that a given set is countable, all we need to do is show that its elements can be listed in some manner. We'll use this fact repeatedly to show that sets are countable, as opposed to having to come up with a bijection between the given set and $\mathbb{N}$ every single time.

Subsets of countable sets. For instance, we claim that any subset of a countable set is countable. (For example, the set of even integers is countable, the set of odd integers is countable, the set of prime numbers is countable, etc.) Suppose A is countable and $S \subseteq A$. Since A is countable, there is a list (either finite or infinite) containing all of its elements:

$$a_1 \quad a_2 \quad a_3 \quad a_4 \quad \dots.$$

Start with a_1 and go through this list, picking out each element which belongs to the subset S. For instance, maybe S consists of only the terms whose indices are a multiple of 3. So, a_1 would not be included, a_2 would not be, but a_3 would be so we list a_3 first in a new list. Then we skip a_4 and a_5, and list a_6 next. And so on, continuing in this manner will produce a list (which will be finite if S is finite) of all elements in S, so S is countable.

$\mathbb{Q}$ is countable. Perhaps the first surprising example of countability is $\mathbb{Q}$, the set of rational numbers. To show that $\mathbb{Q}$ is countable, we must find a way of listing all rational numbers in an infinite list. This requires some creativity, since it is not at all obvious that this should be possible. For instance, we can't just say "list 0 first and then the next rational number after 0" since there is no such thing as the next rational number after zero: the fact that $\mathbb{Q}$ is dense in $\mathbb{R}$ shows that no matter how small a positive rational we take, there will always be other rationals between it and 0, so there is no such thing as the smallest rational number.

Instead, using the fact that a rational looks like $\frac{a}{b}$ with $a \in \mathbb{Z}$ and $b \in \mathbb{N}$ (we can always take b to be positive since any negative can be absorbed into the numerator instead), we could try the following. We know we can list the elements of $\mathbb{Z}$ in an infinite list as

$$0 \quad 1 \quad -1 \quad 2 \quad -2 \quad \dots,$$

so in order to list all rationals let us first list all rationals with denominator 1 using the elements in the list of $\mathbb{Z}$ above as numerators:

$$\frac{0}{1} \quad \frac{1}{1} \quad \frac{-1}{1} \quad \frac{2}{1} \quad \frac{-2}{1} \quad \dots.$$

Next list all rationals with denominator 2 in a similar manner:

$$\frac{0}{2} \quad \frac{1}{2} \quad \frac{-1}{2} \quad \frac{2}{2} \quad \frac{-2}{2} \quad \dots,$$

then all rationals with denominator 2, and so on. If we kept going, would this not produce a huge list containing all rational numbers? The answer is no: if we started out listing all rationals with denominator 1, we would never actually reach the point at which we would start listing all rationals with denominator 1 since there are infinitely many rationals with denominator 1! In other words, what we would get is not a valid "list", since it would be something where we would have infinitely

many terms between two points, for instance infinitely many terms in the list would have to occur between $\frac{1}{1}$ and $\frac{1}{2}$. In the types of lists we are considering, we can only have finitely many terms between one spot and another, so the above procedure does not produce a valid of list of all rational numbers. This is important to have in mind when coming up with such "lists" on your own.

So, we need a different approach. The approach we take, which uses a certain "grid" to come up with the required list, is a technique which will show up elsewhere, so it is a good one to know. Create a grid, where we start by listing all integers at the top and all positive integers down the left-hand side:

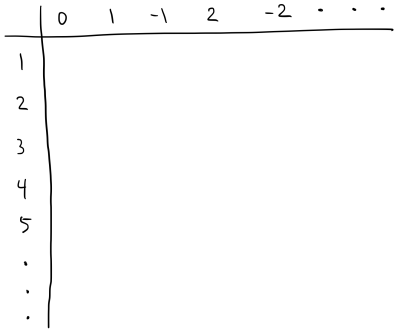

In each spot in the grid, write the rational number which has numerator the integer at the top and denominator the positive integer on the left. The resulting grid will then contain all rational numbers, more than once since each rational can be written as a fraction in more than one way. To create our required list, start by listing the upper-left entry $0/1 = 0$, then move down the next "diagonal" and list those two terms $0/2$ and $1/1$; but actually, $0/2 = 0$ has already been listed, so we skip this term and just list $1/1 = 1$ instead. Now move down to the next diagonal and list those terms, skipping over any which have been listed already, so we skip $0/3$, but list $1/2$ and $-1/1$. And so on, continue in this manner moving down one diagonal at a time and listing the rationals occuring in that diagonal, skipping over any which have been previously listed:

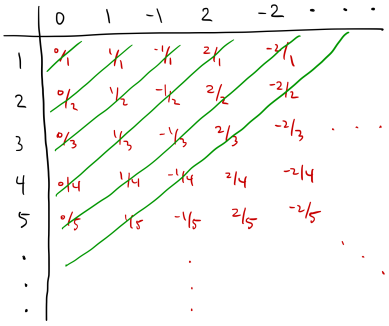

The resulting list will look like:

$$0 \quad 1 \quad \frac{1}{2} \quad -1 \quad \frac{1}{3} \quad -\frac{1}{2} \quad 2 \quad \ldots.$$

Since each rational occurs somewhere in the grid, each rational will occur somewhere in this list, so this will give an infinite list containing all rational numbers. Thus $\mathbb{Q}$ is countable as claimed.

$\mathbb{R}$ is uncountable. So, we have now many examples of sets which are countable, and we'll see more next time. But this leaves open the question as to whether every infinite set is actually countable, or whether uncountable sets actually exist. Indeed, how would even show that a set is uncountable? For now we'll just state the fact that $\mathbb{R}$ is indeed uncountable, a fact we'll prove later on. In fact, "most" infinite sets are uncountable, in a way we'll make clear soon enough.

Lecture 23: More on Countable Sets

Warm-Up. We claim that the union $\mathbb{N} \cup \{\sqrt{2}, \sqrt{3}, \sqrt{5}, \sqrt{7}, \sqrt{11}\}$ is countable. This is quick:

$$\sqrt{2}, \sqrt{3}, \sqrt{5}, \sqrt{7}, \sqrt{11}, 1, 2, 3, 4, 5, \ldots$$

is a list containing all elements of the given union, so this union is countable as claimed.

Unions of countable sets. More generally, we claim that the union of two countable sets is countable. Suppose A and B are countable. If A and B are both finite, $A \cup B$ is finite and so countable. If, say A is finite and B is countable infinite, we can as in the Warm-Up list all elements of A followed by all elements of B to get a list containing all elements of $A \cup B$, so $A \cup B$ is countable in this case as well.

But now consider the case where A and B are both countably infinite. Then there exists a listing of all elements of A:

$$a_1 \quad a_2 \quad a_3 \quad \ldots$$

and a listing of all elements of B:

$$b_1 \quad b_2 \quad b_3 \quad \ldots.$$

To get a listing of all elements in $A \cup B$ we cannot simply say "take the first list and then tack on the second list at the end" since this does not give a valid list—there would be infinitely many terms between a_1 and b_1. Instead, we can come up with a list by taking the first term from each list, then the second, the third, and so on:

$$a_1 \quad b_1 \quad a_2 \quad b_2 \quad a_3 \quad b_3 \quad \ldots.$$

If need be, cross out any any repititions to get a list containing all elements of $A_{\cup}B$. Thus $A \cup B$ is countable.

From this we can build up to unions of any finite number of countable sets. For instance, if A, B, C are all countable, $A \cup B$ is countable from what we showed above, and then

$$(A \cup B) \cup C$$

is countable since it is the union of the countable set $A \cup B$ and the countable set C. If D is countable as well, this will imply that

$$(A \cup B \cup C) \cup D$$

is also countable, and so on. To be precise, induction can be used to show (with $n = 2$ as a base case) that if $A_1, \ldots, A_n$ are all countable, then $A_1 \cup \cdots \cup A_n$ is countable.

But we can say even more: the *countable* union of countable sets is countable. That is, suppose $A_1, A_2, A_3, \ldots$ is any countable (possibly infinite) collection of sets, each of which is countable. We claim that

$$\bigcup_{n \in \mathbb{N}} A_n$$

is countable too. Indeed, since each A_n is countable, we can list the elements of each as follows:

$$\begin{array}{lllll} A_1: & a_{11} & a_{12} & a_{13} & \ldots \\ A_2: & a_{21} & a_{22} & a_{23} & \ldots \\ A_3: & a_{31} & a_{32} & a_{33} & \ldots \end{array}$$

$$\vdots \qquad \vdots$$

To be clear, a_{ij} denotes the j-th term in the listing of elements of A_i. Then every element of $\bigcup_n A_n$ will appear somewhere in this "grid", and the "move down diagonals" technique we used in showing that $\mathbb{Q}$ is countable will produce, if we skip any repeats, a list of all elements of $\bigcup_n A_n$. Thus $\bigcup_n A_n$ is countable as claimed.

This gives a nice way of showing that certain sets are countable, if we can express them as the union of countably many countable sets. We'll see examples of this shortly.

Products of countable sets. We also claim that if A and B are both countable, then $A \times B$ is countable as well. This is certainly true when both A and B are finite: if A has n elements and B has m elements, $A \times B$ has mn elements since in a pair (a, b), there are n possible choices for a and m for b. Thus the product of two finite sets is finite, so countable.

The case which really matters is the product of two countably infinite sets. If A and B are countably infinite, we create a grid:

	b_1	b_2	b_3	$\cdots$
a_1	(a_1, b_1)	(a_1, b_2)		
a_2				
a_3				
$\vdots$				

using a listing of elements of A down the left side and a listing of elements of B at the top. Fill in each spot in the grid with the pair (a_i, b_j) made from the term from A on the left and the term from B at the top, and then move down diagonals to give a listing of all elements of $A \times B$.

If A, B, C are all countable, then $A \times B$ is countable so $(A \times B) \times C$ is countable. More generally, using induction we can show that the product of any finite number of countable sets is countable: if $A_1, \ldots, A_n$ are all countable and we assume $A_1 \times \cdots \times A_n$ is countable for some n, then if A_{n+1} is also countable we get that

$$A_1 \times \cdots \times A_{n+1} = (A_1 \times \cdots \times A_n) \times A_{n+1}$$

is countable since it is the product of the countable set $A_1 \times \cdots \times A_n$ with the countable set A_{n+1}. By induction we conclude that the product of finitely many countable sets is indeed countable.

But be careful: even though the union of *countably* many countable sets is always countable, it is not true that the product of countably infinitely many countable sets is always countable. In fact, as soon as we take the product of infinitely many sets which each have at least two elements, the product becomes uncountable. We'll come back to this later.

Example 1. We finish with some examples of how to put the above facts to good use. First consider the set

$$A = \left\{ \left(p, \frac{1}{q} \right) \in \mathbb{R}^2 \;\middle|\; p, q \text{ are prime numbers} \right\}.$$

The set P of primes is countable since it is a subset of $\mathbb{N}$, and the set of reciprocals of primes is countable since the function $q \mapsto \frac{1}{q}$ is a bijection between it and P. Thus A can be expressed as the product of these two sets, A is countable as well. Alternatively, we could note that A is a subset of $\mathbb{Q} \times \mathbb{Q}$ and a subset of a countable set is always countable.

Example 2. Now consider the set S of polynomials of degree at most 2 with rational coefficients:

$$S = \left\{ a_0 + a_1 x + a_2 x^2 \;\middle|\; a_0, a_1, a_2 \in \mathbb{Q} \right\}.$$

We claim that S is countable. The idea here is that any such polynomial is determined by its coefficients a_0, a_1, a_2, and that there are countably many such coefficients since they can all together be viewed as specifying an element of $\mathbb{Q} \times \mathbb{Q} \times \mathbb{Q}$. (Intuitively, there are "$|\mathbb{Q}|$-many" choices for a_0, $|\mathbb{Q}|$-many for a_1, and $|\mathbb{Q}|$-many for a_2, so there should be $|\mathbb{Q}||\mathbb{Q}||\mathbb{Q}|$-many such polynomials. We'll talk about about what it means to multiply cardinalities, but the upshot is that this product will in fact be $|\mathbb{Q}|$ itself.) To be precise, the function $S \to \mathbb{Q}^3$ defined by

$$a_0 + a_1 x + a_2 x^2 \mapsto (a_0, a_1, a_2)$$

is bijective, so since $\mathbb{Q}^3$ (being the product of finitely many countable sets) is countable, so is S.

Example 3. Finally, we consider the set Z of all possible roots of the nonzero polynomials in the set S above:

$$Z = \{\alpha \in \mathbb{R} \mid \text{there exists nonzero } p(x) \in S \text{ such that } p(\alpha) = 0\}.$$

For instance, any rational number $\frac{a}{b}$ is in Z since such a rational is a root of $-a + bx$, and the square root $\sqrt{k}$ of any natural number k is in Z since it is a root of $x^2 - k$. We claim that Z is also countable. We already know that there are countably many possible nonzero polynomials of degree at most 2 with rational coefficients, and in addition we know that any such polynomial has at most 2 roots. Thus, given a countable listing of all such polynomials:

$$p_1(x),\ p_2(x),\ p_3(x),\ \ldots$$

we can express Z as the union of the sets of roots of these:

$$Z = \bigcup_{i \in \mathbb{N}} \{\text{roots of } p_i(x)\}.$$

This is a countable union of countable (in fact finite) sets, and so is countable itself.

Algebraic vs transcendental numbers. A real number α is said to be *algebraic* if there exists a nonzero polynomial $p(x)$ (of any degree) with rational coefficients such that $p(\alpha) = 0$. A generalization of the example above can be used to show that the set of all algebraic numbers is also countable. (This is done on a homework assignment.)

A real number which is not algebraic is said to be *transcendental*. For instance, $\pi, e, \sin 1$ are all transcendental, which is not at all straightforward to show. It will be a consequence of the fact that $\mathbb{R}$ is uncountable (which we will prove next time) that the set of transcendental numbers is also uncountable: the union of this set and the set of algebraic numbers is $\mathbb{R}$, so if this set were countable this union would be countable as well, which it is not. The point is that there are, in a sense, many more transcendental numbers than there are algebraic ones. This is somewhat surprising since it is easy to write down tons of examples of algebraic numbers: any expression made by taking sums, products, quotients, and roots of rationals will be algebraic, so for example

$$\sqrt{2} + \frac{3 - \sqrt[7]{\sqrt[3]{4} - \sqrt{2}} + \sqrt[123]{3}}{2 + \sqrt{\sqrt[3]{7} + \sqrt[6]{182}}}$$

is algebraic, but it is not so straightforward to write down examples of numbers which are transcendental, even though there are many more of these. This suggests that "most" real numbers cannot be expressed in terms of standard, well-known notation—c'est la vie.

Lecture 24: Uncountable Sets

Warm-Up. We show that the set of all subsets of $\mathbb{N}$ with at most two elements:

$$F = \{S \subseteq N \mid |S| \leq 2\}$$

is countable. Let F_0, F_1, F_2 denote the elements of F which contain $0, 1, 2$ elements respectively, so $F = F_0 \cup F_1 \cup F_2$. The only subset of $\mathbb{N}$ containing 0 elements is the empty set, so F_0 only has one element. Now, an element of F_1 looks like $\{n\}$ for some $n \in \mathbb{N}$, so there should be as many one-element subsets of $\mathbb{N}$ as there are elements of $\mathbb{N}$; more precisely, the function

$$F_1 \to \mathbb{N} \text{ defined by } \{n\} \mapsto n$$

is a bijection, so F_1 has the same cardinality as $\mathbb{N}$ and is thus countable.

A two-element subset of $\mathbb{N}$ looks like $\{n, m\}$, and so is fully determined by the two numbers n and m. Let as assume that whenever we write $\{n, m\}$ we are writing the elements in increasing order, so $n < m$. Then the function

$$F_2 \to \mathbb{N} \times \mathbb{N} \text{ defined by } \{n, m\} \mapsto (n, m)$$

is injective, so since $\mathbb{N} \times \mathbb{N}$ is countable we have that F_2 is countable as well. (Note that this function is not surjective, since by our convention that $n < m$ nothing is sent to $(2, 1)$ for instance.) Thus, $F = F_0 \cup F_1 \cup F_2$ is a countable union of countable sets, and so is countable itself. (This same strategy can be used to show that the set of all finite subsets of $\mathbb{N}$ is countable, as is done on a homework assignment. The set of *all* subsets of $\mathbb{N}$, however, is uncountable, as we'll soon see.)

$\mathbb{R}$ is uncountable. We now show that $\mathbb{R}$ is uncountable, which intuitively means that there are "too many" real numbers to list in a single infinite list. In a sense then, there are "more" real numbers than integers or rational numbers, even though all of these sets are infinite. We actually show that the interval $(0, 1)$ is uncountable; this implies that $\mathbb{R}$ is uncountable as well since if $\mathbb{R}$ were countable, any subset of it would be countable as well.

To show that $(0, 1)$ is uncountable we show that there does not exist a bijection $\mathbb{N} \to (0, 1)$. To this end, let $f : \mathbb{N} \to (0, 1)$ be *any* function. We claim that f is not surjective, which immediately implies that f is not bijective, thereby showing that no function $\mathbb{N} \to (0, 1)$ can be bijective as required. To show that f is not surjective we will come up with an explicit element of $(0, 1)$ which is not in the image of f. The following argument is known as "Cantor's diagonalization argument", and is a key tool for showing that given sets are uncountable. List the elements in the image of f in terms of their decimal expansions:

$$\begin{aligned} f(1) &= 0.x_{11}x_{12}x_{13}\ldots \\ f(2) &= 0.x_{21}x_{22}x_{23}\ldots \\ f(3) &= 0.x_{31}x_{32}x_{33}\ldots \\ \vdots &\quad \vdots \end{aligned}$$

Define the number $y = 0.y_1y_2y_3\ldots \in (0, 1)$ by taking the digit y_i to be anything different from x_{ii}; to be concrete, take

$$y_i = \begin{cases} 3 & \text{if } x_{ii} \neq 3 \\ 7 & \text{if } x_{ii} = 3. \end{cases}$$

(Note that the use of 3 and 7 here is not important—all we need to do is guarantee that y_i and x_{ii} are different. The name "diagonalization argument" comes from the use of the "diagonal" terms x_{11}, x_{22}, x_{33}, etc.) Now, this number y differs from $f(1)$ in the first decimal digit (since $y_1 \neq x_{11}$), so $y \neq f(1)$. Also, y differs from $f(2)$ in the second decimal digit (since $y_2 \neq x_{22}$), so $y \neq f(2)$. In general, $y \neq f(n)$ since y and $f(n)$ differ in the n-th decimal digit. Thus y is not in the image of f, so f is not surjective as claimed. We conclude that $(0, 1)$, and hence $\mathbb{R}$, is uncountable.

Why doesn't this show that $\mathbb{N}$ is uncountable? We should be clear about what the "diagonalization argument" does and why it works. As a test, imagine we attempted to apply it in the following scenario. Take a listing of natural numbers, for instance:

$$\begin{array}{l} 1 \\ 12 \\ 123 \\ 1234 \\ 12345 \\ \vdots \end{array}$$

Define N by choosing its i-th digit to be something different from the i-th digit of the i-th number listed above; for example

$$N = 23456\ldots$$

works since its first digit is different than the first digit of 1, its second digit is different than the second digit of 12, its third is different from the third digit of 123, and so on. This N gives something not included in the listing of natural numbers above, so if we had assumed that this listing included *all* natural numbers, would this not give a contradiction showing that $\mathbb{N}$ was in fact uncountable? The answer is no, where the point is that the resulting N does not actually represent a naturla number since it has infinitely many digits (!) whereas a natural number only has finitely many. In this case, the "diagonalization argument" *does* produce something not in our list, but not an element of $\mathbb{N}$. In the case of $(0, 1)$, the resulting number y *is* in $(0, 1)$.

Note, however, that this reasoning does show that if we took expressions consisting of infinitely many "digits" of natural numbers, that the set of such expressions is uncountable. For instance, this can be used to show that the set $\mathbb{N}^\infty$ of infinite sequences of positive integers is uncountable, as is done in a discussion section problem.

Set of binary sequences. Denote by $\{0, 1\}^\infty$ the set of *binary* sequences, which are infinite sequences

$$(x_1, x_2, x_3, \ldots)$$

where each term is either 0 or 1. Cantor's diagonalization argument shows that this is uncountable. To be concrete, suppose

$$\begin{array}{l} \mathbf{x}_1 = (x_{11}, x_{12}, x_{13}, \ldots) \\ \mathbf{x}_2 = (x_{21}, x_{22}, x_{23}, \ldots) \\ \mathbf{x}_3 = (x_{31}, x_{32}, x_{33}, \ldots) \\ \quad \vdots \quad \vdots \end{array}$$

is an infinite listing of elements of $\{0,1\}^\infty$. (So, each x_{ij} is either 0 or 1.) To be clear, the first expression $(x_{11}, x_{12}, x_{13}, \ldots)$ gives one element of $\{0,1\}^\infty$, the second $(x_{21}, x_{22}, x_{23}, \ldots)$ gives another, and so on. Define $\mathbf{y} = (y_1, y_2, y_3, \ldots)$ by setting

$$y_i = \begin{cases} 1 & \text{if } x_{ii} = 0 \\ 0 & \text{if } x_{ii} = 1. \end{cases}$$

Then $\mathbf{y} = (y_1, y_2, y_3, \ldots)$ is in $\{0,1\}^\infty$, but is not equal to any element in the listing above since it differs from $\mathbf{x}_i$ in the i-th term because $y_i \neq x_{ii}$. Thus, no infinite list of elements of $\{0,1\}^\infty$ can contain *all* elements of $\{0,1\}^\infty$ (i.e. no function $\mathbb{N} \to \{0,1\}^\infty$ can be surjective), so $\{0,1\}^\infty$ is uncountable.

The power set of $\mathbb{N}$. Finally, we use $\{0,1\}^\infty$ to show that the power set $\mathcal{P}(\mathbb{N})$ of $\mathbb{N}$ is uncountable. Recall that $\mathcal{P}(\mathbb{N})$ is the set of all subsets of $\mathbb{N}$. The Warm-Up alluded to the fact that the set of all *finite* subsets of $\mathbb{N}$ was actually countable, but now we are considering all subsets.

The idea here is that an element of $\mathcal{P}(\mathbb{N})$ can be characterized using the same type of data as an element of $\{0,1\}^\infty$, which will give us a bijection $\mathcal{P}(\mathbb{N}) \to \{0,1\}$. Indeed, suppose $S \in \mathcal{P}(\mathbb{N})$, which means that S is a subset of $\mathbb{N}$. Now, this subset either contains 1 or it doesn't, contains 2 or doesn't, and so on. Define x_i to be 0 if $i \notin S$ and 1 if $i \in S$; so if $1 \in S$ we take the first x_1 to be 1, if $2 \notin S$ we take x_2 to be 0, and so on. This results in a binary sequence

$$(x_1, x_2, x_3, \ldots)$$

characterizing those natural numbers which should belong to the subset $S \subseteq \mathbb{N}$ depending on which x_i's are 1's and which are 0's. For instance, $S = \mathbb{N}$ would produce

$$(1, 1, 1, 1, 1, 1, \ldots)$$

since every positive integers is in $\mathbb{N}$; the subset of even integers would produce

$$(0, 1, 0, 1, 0, 1, 0, 1, \ldots)$$

where the 1's show up in only the even locations; the subset of multiples of 3 would produce

$$(0, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0, 1, \ldots)$$

with 1's in the spots corresponding to a multiple of 3, and so on. This gives a function $\mathcal{P}(\mathbb{N}) \to \{0,1\}^\infty$, which is invertible since given an infinite binary sequence we can reconstruct a subset of $\mathbb{N}$ by looking at the locations where 1's occur. Again, the upshot is that an element of $\mathcal{P}(\mathbb{N})$ can be uniquely characterized using an element of $\{0,1\}^\infty$, which is what gives the desired bijection.

Since $\{0,1\}^\infty$ is uncountable, we conclude that $\mathcal{P}(\mathbb{N})$ is uncountable as well, so that there are in a sense "more" subsets of $\mathbb{N}$ than there are elements of $\mathbb{N}$. Actually, even more is true: $\mathcal{P}(\mathbb{N})$ has the same cardinality as $\mathbb{R}$. Next time we will come back to this fact as well as look at power sets in general more closely.

Lecture 25: Power Sets

Warm-Up. Set $C_0 = [0,1]$. Then take C_1 to be the set obtained by removing the "middle third" portion of C_0:

$$C_1 = \left[0, \tfrac{1}{3}\right] \cup \left[\tfrac{2}{3}, 1\right].$$

Define C_2 be the set obtained by removing the middle third portion of each interval making up C_1:

$$C_2 = \left[0, \tfrac{1}{9}\right] \cup \left[\tfrac{2}{9}, \tfrac{3}{9}\right] \cup \left[\tfrac{6}{9}, \tfrac{7}{9}\right] \cup \left[\tfrac{8}{9}, 1\right],$$

and continuing in this manner define C_n in general to be the set obtained by removing the middle portion of each interval making up C_{n-1}. The *Cantor set* is the set C consisting of what remains after we continue this process indefinitely, or equivalently the intersection of all the C_n's:

$$C = \bigcap_n C_n.$$

We show that the Cantor set is uncountable. This can seem surprising at first, since it seems to be challenging to specify precisely what real numbers belong to the Cantor set. For sure, all endpoints of all intervals in each step of the construction of C remain throughout the entire process, so all such endpoints belong to C. (So, for instance, $0, 1, \frac{1}{3}, \frac{2}{3}, \frac{1}{9}, \frac{2}{9}, \frac{7}{9}$, and $\frac{8}{9}$ are all in C.) However, all such endpoints are rational, so there are only countably many of them, and yet we are saying that C is uncountable, meaning that there are way more elements of C that aren't among these endpoints than there are endpoints. We'll take a second after the proof that C is uncountable to say more precisely what the Cantor set consists of.

Let $x \in C$. We construct an element of $\{0, 2\}^\infty$ associated to this as follows. (We'll see why I'm using $\{0, 2\}$ instead of $\{0, 1\}$ afterwards.) Since $C = \bigcap C_n$, $x \in C_n$ for all n. In particular, $x \in C_1$ so x is in one of the two intervals making up C_1; take the first element in our sequence to be 0 if x is in the "left" interval $[0, 1/3]$ and take the first element in our sequence to be 2 if x is in the "right" interval $[2/3, 1]$. Now, whichever of these intervals x is in will itself split into two smaller intervals in the construction of C_2. Since $x \in C_2$, x will be in one of these smaller intervals; take the next element in our sequence to be 0 if it is the "left" interval x is in and take it to be 2 if x is in the "right" interval. For instance, the interval $[0, 1/3]$ splits into $[0, 1/9]$ and $[2/9, 1/3]$. If $x \in [0, 1/9]$ the first two terms in the sequence we are constructing will be $0, 0$, while if $x \in [2/9, 1/3]$ we have $0, 2$ as the beginning of our sequence. Continuing in this manner, whichever interval making up C_2 that x is in will split into two smaller pieces; take 0 as the third term in our sequence if x is in the left piece and 2 if x is in the right piece, and so on. By keeping track of which interval x is in at each step in the construction of the Cantor set in this manner we get a sequence of 0's and 2's.

For instance, if we get the sequence $(0, 2, 2, 2, 0, 0, 0, \ldots)$, x is in the "left" interval of C_1, then in the "right" smaller interval which this interval splits into, then in the "right" smaller interval this splits into, then "right" again, then in the "left" smaller interval that this splits into, and so on. (This is easier to imagine if you draw a picture of this splitting into smaller and smaller intervals as we did in class. In general, a 0 means "go left" in the next step of the construction and 2 means "go right".)

This assignment of a sequence of 0's and 2's to an element $x \in C$ defines a function $C \to \{0, 2\}^\infty$. It is injective since different elements in the Cantor set produces different sequences (at some point in the construction, two different numbers in the Cantor set will belong to two different "smaller" intervals since the lengths of these smaller intervals are getting closer and closer to zero), and it is surjective since given any sequence we can use it to single out an element of the Cantor set. Thus C and $\{0, 2\}^\infty$ have the same cardinality. The "diagonalization argument" we gave last time to show that $\{0, 1\}^\infty$ is uncountable can be modified to show that $\{0, 2\}^\infty$ is uncountable by replacing 1's with 2's, so we conclude that C is uncountable as well.

What's in the Cantor set? Just for fun, let's clarify what the Cantor set actually consists of. Any real number in $[0, 1]$ has a decimal expansion, where the notation

$$0.x_1x_2x_3\ldots$$

really denotes the result of the infinite summation given by

$$\frac{x_1}{10} + \frac{x_2}{10^2} + \frac{x_3}{10^3} + \cdots .$$

By changing the "base" 10 used here, we can come up with decimal expansions with respect to other bases. In particular, any such number has a "base 3" decimal expansion

$$0.y_1y_2y_3\ldots$$

where each digit y_i is $0, 1$, or 2; this comes from expressing the given number as "base 3" infinite sum of the form:

$$\frac{y_1}{3} + \frac{y_2}{3^2} + \frac{y_3}{3^3} + \cdots .$$

If you think about how these digits relate to splitting an interval up into thirds, you can see that they precisely keep track of which third of an interval a given number belongs to when splitting it up further and further. For instance, a digit of 1 indicates that your given number should belong to the "middle third" portion of an interval. Since these middle thirds are removed in the construction of the Cantor set, we see that the Cantor set precisely consists of those numbers in $[0, 1]$ whose base 3 decimal expansions contains only 0's and 2's. For instance, the base 3 decimal expansion of $\frac{1}{4}$ looks like

$$0.0202020202020\ldots$$

with 0's and 2's alternating, so $\frac{1}{4}$ is a non-endpoint element of the Cantor set. Note, however, that $\frac{1}{4}$ is still rational, and yet it follows from what we showed before that the Cantor set contains uncountably many irrational numbers.

The power set of $\mathbb{N}$ and $\mathbb{R}$. Last time we showed that the power set of $\mathbb{N}$ was uncountable by showing it had the same cardinality as $\{0, 1\}^\infty$, but more precisely $\mathcal{P}(\mathbb{N})$ actually has the same cardinality as $\mathbb{R}$. The key idea is to still consider elements of this power set as being characterized by elements of $\{0, 1\}^\infty$, only that now we interpret a sequence of 0's and 1's as giving the digits of the *binary* (i.e. base 2) decimal expansion of a real number. The function $\{0, 1\}^\infty \to (0, 1)$ defined (via binary expansions) by

$$(x_1, x_2, x_3, \ldots) \mapsto 0.x_1x_2x_3\ldots$$

is ALMOST bijective, so $\{0, 1\}^\infty$ should have the same cardinality as $(0, 1)$, and thus $\mathcal{P}(\mathbb{N})$ should as well, which would imply that $\mathcal{P}(\mathbb{N})$ has the same cardinality as $\mathbb{R}$.

BUT, we have to be careful: the function given above is not truly injective since a given number can have more than one such binary expansion. For instance, $0.001111\ldots$ with 1's repeating is the same real number as 0.01, so two sequences of 0's and 1's can given the same real number. There are ways around this, for instance using the Cantor-Schroeder-Bernstein theorem we'll talk about next time, but we'll leave it to the book to demonstrate possible fixes. Here we're only trying to come up with a sense for why $\mathcal{P}(\mathbb{N})$ and $\mathbb{R}$ should have the same cardinality.

Power sets of finite sets. Thus, the power set $\mathcal{P}(\mathbb{N})$ of $\mathbb{N}$ has "more" elements than $\mathbb{N}$ does. This is actually true for any set in general, which is a fact we'll now build towards. As a start, we consider the cardinality of power sets of finite sets. The claim is that if A has n elements, then $\mathcal{P}(A)$ has 2^n elements. For instance, $\{1, 2, 3\}$ has $2^3 = 8$ subsets:

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

Here are three proofs of this claim.

First, we can argue by induction. The base case is the empty set with 0 elements, whose power set is $\{\emptyset\}$ and thus has $2^0 = 1$ element. Suppose now that any set with n elements has 2^n subsets, and let A be a set with $n+1$ elements: $A = \{a_1, \ldots, a_n, a_{n+1}\}$. Let P' denote the set of subsets of A not containing a_{n+1} and P'' the set of subsets of A which do contain a_{n+1}. Then P' and P'' are disjoint and $\mathcal{P}(A) = P' \cup P''$. Now, elements of P' can be viewed as subsets of $\{a_1, \ldots, a_n\}$ since these elements do not contain a_{n+1}. By the induction hypothesis, $\{a_1, \ldots, a_n\}$ has 2^n subsets, so $|P'| = 2^n$. Also, the function $P' \to P''$ defined by

$$S \mapsto S \cup \{a_{n+1}\}$$

is a bijection since any subset M of A which does contain a_{n+1} can be viewed as the complement $M - \{a_{n+1}\}$ union $\{a_{n+1}\}$. Since $|P'| = 2^n$, we get that $|P''| = 2^{n+1}$ as well. Thus

$$|\mathcal{P}(A)| = |P'| + |P''| = 2^n + 2^n = 2^{n+1}.$$

Hence by induction we conclude that if $|A| = n$, then $|\mathcal{P}(A)| = 2^n$ as claimed.

Second, we can count the number of subsets of A as follows using "n choose k" notation. The number of subsets of A with 0 elements is $\binom{n}{0}$; the number of subsets with 1 element is $\binom{n}{1}$; and so on, the number of subsets with k elements is $\binom{n}{k}$. Thus the total number of subsets of A is

$$\binom{n}{0} + \binom{n}{1} + \cdots + \binom{n}{n},$$

which equals 2^n. This can be seen, for instance, by taking the binomial formula

$$(x+y)^n = \sum_{k=0}^{n} \binom{n}{k} x^{n-k} y^k$$

and setting $x = y = 1$: the right side becomes $\binom{n}{0} + \binom{n}{1} + \cdots + \binom{n}{n}$ and the left side $(1+1)^n = 2^n$.

Finally, we give possibly the simplest proof of all. A subset of $A = \{a_1, \ldots, a_n\}$ can be characterized by specifying whether a_1 should or should not be included, whether a_2 should or should not be included, and so on. There are two choices as to what to do with a_1 (include or exclude), two choices for a_2, and so on, given $2 \cdot 2 \cdots 2 = 2^n$ ways of picking which elements of A to include in a given subset.

Comparing cardinalities. Thus, in the case of finite sets, a set always has cardinality strictly less than that of its power set. The same is true for infinite sets as well, but we now need a way of making sense of what it means for one infinite cardinality to be "smaller" than another.

We first *define* the notation $|A| \leq |B|$ to mean that there exists an injective function $A \to B$. In the case where A and B are finite, we have certainly seen that the existence of an injective function $A \to B$ means that the number of elements of A is no more than the number of elements of B, and we use this intuition as a guide in the infinite case as well. However, we should be absolutely clear: $|A| \leq |B|$ in the infinite case does NOT literally mean that the number of elements of A is less than or equal to the number of elements of B since in this case we are talking about an infinite number of elements, but rather $|A| \leq |B|$ means *by definition* that there exists an injection $A \to B$.

We then define the notation $|A| < |B|$ to mean that $|A| \leq |B|$ but $|A| \neq |B|$, which boils down to saying that there exists an injection $A \to B$ but not a bijection $A \to B$. Again, in the finite case this does capture the idea that the number of elements of A is strictly less than the number of elements of B, but in the infinite case, while can interpret $|A| < |B|$ intuitively as a statement about "number of elements", we can only interpret it rigorously in terms of injections and bijections.

For instance, we have that $|\mathbb{Q}| < |\mathbb{R}|$, since there exists an injection $\mathbb{Q} \to \mathbb{R}$ (for instance the function which sends any rational number to itself) but no bijection $\mathbb{Q} \to \mathbb{R}$ because $\mathbb{R}$ is uncountable. The claims we have derived about the power set of $\mathbb{N}$ show that $|\mathbb{N}| < |\mathcal{P}(\mathbb{N})|$, where we know that $|\mathbb{N}| \leq |\mathcal{P}(\mathbb{N})|$ since the function $\mathbb{N} \to \mathcal{P}(\mathbb{N})$ sending n to $\{n\}$ is injective.

Cardinality of power sets. We can now state our main claim about power sets, which says intuitively that the power set of a set always has "more" elements than the set of which it is the power set. The precise claim is that for any set A, $|A| < |\mathcal{P}(A)|$, which concretely means that there exists an injection $A \to \mathcal{P}(A)$ but not a bijection $A \to \mathcal{P}(A)$. The fact that $|A| \leq |\mathcal{P}(A)|$ is simple: the function $A \to \mathcal{P}(A)$ defined by $a \mapsto \{a\}$ is injective. However, showing that there is not bijection $A \to P(A)$ is much harder. In fact, the claim is that no function $A \to P(A)$ can be surjective, from which it follows that no such bijection can exist either. The proof we'll give is standard, but is definitely the type of thing you might have to go through a few times to make clear. In class we didn't give the proof until the Warm-Up the following day, but it fits better here. This is NOT a proof you have to know, but you should definitely know the fact that $|A| < |\mathcal{P}(A)|$ is always true.

Proof. Suppose $f : A \to \mathcal{P}(A)$ is any function. Define $J \subseteq A$ to be the set of all elements $a \in A$ such that the subset $f(a)$ of A does not contain a:

$$J = \{a \in A \mid a \notin f(a)\}.$$

To be clear, for each $a \in A$, $f(a) \in \mathcal{P}(A)$, meaning that $f(a)$ is a subset of A and we can then ask whether or not a is in this particular subset. We claim that the resulting subset J of A is not in the image of J, which shows that f is not surjective.

By way of contradiction, suppose there does exist $a \in A$ such that $f(a) = J$. Then there are two possibilities: either $a \in J$ or $a \notin J$. If $a \in J$, then $a \in f(a)$ since $f(a) = J$, but by the definition of J to say that $a \in J$ means $a \notin f(a)$, so we have a contradiction. If $a \notin J$, then $a \notin f(a)$ since $f(a) = J$, which by the definition of J means that $a \in J$, another contradiction. Thus we conclude that no such a exists, so J is not in the image of f and hence f is not surjective as claimed. □

No largest cardinality. We finish by pointing out that the fact we just proved now implies that there is no such thing as a "larger" cardinality. In particular, for any infinite set A, by repeatedly taking more and more power sets we have that:

$$|A| < |\mathcal{P}(A)| < |\mathcal{P}(\mathcal{P}(A))| < |\mathcal{P}(\mathcal{P}(\mathcal{P}(A)))| < \cdots$$

is a never ending sequence of strictly increasing cardinalities, so no there is no such thing as a "largest cardinality". Also note that as a consequence we now know that not all uncountable sets have the same cardinality: the power set $\mathcal{P}(\mathbb{R})$ of $\mathbb{R}$ is an uncountable set with cardinality strictly larger than that of $\mathbb{R}$.

Lecture 26: Cantor-Schroeder-Bernstein Theorem

Warm-Up. As a Warm-Up we finished the proof that $|A| < |\mathcal{P}(A)|$, which here is included at the end of the previous lecture.

Cardinal numbers and arithmetic. We now spend some time talking about "numbers" used to denote the size of infinite sets, and doing "arithmetic" with such numbers. As said in class, this

is NOT something which will be on the final, and is only included for the sake of interest and to a glimpse of the crazy awesome things we can do with the material we've built up.

The cardinality of finite sets are denoted by nonnegative integers: $0, 1, 2, 3, \ldots$. To denote the cardinalities of infinite sets we introduce a new type of "number". The smallest infinite cardinal number is denoted by $\aleph_0$ (pronounced "aleph-not") and denotes the cardinality of an infinite countable set. So, for instance,

$$|\mathbb{N}| = |\mathbb{Z}| = |\mathbb{Q}| = \aleph_0.$$

The cardinality of $\mathbb{R}$ is usually denoted by c, which stands for "continuum". The fact that $\mathbb{R}$ has larger cardinality than $\mathbb{N}$ can then be written as

$$\aleph_0 < c.$$

By analogy with the case of finite sets (where $|A| = n$ implies $|\mathcal{P}(A)| = 2^n$), the cardinality of the power set of a A is denoted by $2^{|A|}$:

$$|\mathcal{P}(A)| = 2^{|A|}.$$

Intuitively, this can be interpreted (but not in a literal sense) as saying that the number of subsets of A is given by multiplying 2 by itself "$|A|$"-many times, which comes from the idea that a subset can be characterized by specifying which elements to include and which to exclude, so that there are two possible things to do with each element. But again, this is just for the sake of intuition: $2^{|A|}$ does not literally mean that we multiply 2 by itself an infinite number of times. The fact that $\mathcal{P}(\mathbb{R})$ has cardinality larger than $\mathbb{R}$ means $c < 2^c$, and the fact that $\mathcal{P}(N)$ has the same cardinality as $\mathbb{R}$ then becomes the statement that

$$2^{\aleph_0} = c.$$

This type of equality then leads us to wonder about arithmetic operations which can be extended to the setting of *cardinal numbers*. For instance, what is $\aleph_0 + \aleph_0$, and what does this even mean? To see how to define this, we go back to the case of finite sets to get some intuition. There, if $|A| = n$, $|B| = m$, and A and B are *disjoint*, it turns out that $|A \cup B| = n + m$. Thus, we can characterize "addition" in terms of the cardinality of a union, and we thus *define*

$$\aleph_0 + \aleph_0 = |A \cup B|$$

where A and B are disjoint sets each of cardinality $\aleph_0$. Since the union of two countable sets is countable, we conclude that such a union still has cardinality $\aleph_0$, so we conclude that

$$\aleph_0 + \aleph_0 = \aleph_0,$$

demonstrating that arithmetic with cardinal numbers can be behave in ways different than what we're used to in the setting of finite numbers. We also have $\aleph_0 + c = c$ and $c + c = c$, which come from thinking about the union of $\mathbb{R}$ with a countable set or of two disjoint intervals.

Similarly, we can talk about what it means to multiply infinite cardinalities together. In the finite case, if $|A| = n$ and $|B| = m$, then $A \times B$ is a set with mn elements, so we can characterize "multiplication" in terms of Cartesian products. We thus define $|A||B|$ in general to be the cardinality of $A \times B$:

$$|A||B| = |A \times B|.$$

The fact that the product of two countably infinite sets is countably infinite gives $\aleph_0\aleph_0 = \aleph_0$, or more succinctly $\aleph_0^2 = \aleph_0$, and it can also be shown that $c\aleph_0 = c$ and $c^2 = c$. Higher-order powers can be defined similarly.

Continuum Hypothesis. Recall that $\aleph_0$ denotes the cardinality of a countably infinite set, which is the "smallest" infinite cardinality. Then, $\aleph_1$ denotes the next largest infinite cardinality, $\aleph_2$ the next largest and so on:

$$\aleph_0 < \aleph_1 < \aleph_2 < \cdots .$$

We can then ask where $|\mathbb{R}|$ shows up in this chain of cardinalities. In particular, is $|\mathbb{R}|$ in fact the next largest cardinality after $\aleph_0$, so that $\aleph_1 = c$? Or said another way, given that

$$\aleph_0 = |\mathbb{Q}| < |\mathbb{R}| = c,$$

does there exist a subset of $\mathbb{R}$ with cardinality strictly *between* that of $\mathbb{Q}$ and $\mathbb{R}$, thereby showing

$$\aleph_0 < \aleph_1 < c?$$

The *continuum hypothesis* asserts that the answer to this question is no: there does not exist a cardinality strictly between that of $\mathbb{Q}$ and $\mathbb{R}$, or equivalently $\aleph_1 = c$.

So, is the continuum hypothesis true? In the 1930's Kurt Gödel showed that the continuum hypothesis cannot be *disproven* within the currently accepted axioms of set theory. This would seem to indicate that the continuum hypothesis is true, since saying that "it cannot be proven false" *seems* to be the same as saying that "it is true", but this is not the case: in the 1960's Paul Cohen shocked the mathematical world by showing that the continuum hypothesis cannot be proven within the currently accepted axioms of set theory. Thus, the continuum hypothesis is an example of what is known as an *undecidable* statement, which is a statement which can neither be proven nor disproven. It might seem surprising (as it was when first discovered) that such statements can exist, but there you have it. This get to the heart of the limitations of mathematical reasoning and logic itself, which is the subset of current research in set theory and other foundational fields. Good stuff, no?

Cantor-Schroeder-Bernstein. After this interlude, we now come back to things which *are* relevant for our course. Our final topic, the *Cantor-Schroeder-Bernstein* theorem, gives one more way of showing that two sets have the same cardinality, without having to come up with an explicit bijection between them.

To motivate the statement, consider the notation $|A| \leq |B|$ we introduced previously. The claim we want to make is that if $|A| \leq |B|$ and $|B| \leq |A|$, then $|A| = |B|$. This seems to be obvious at first, since if two numbers are less than or equal to one another, then they must indeed be the same. HOWEVER, recall that $|A|$ and $|B|$ here are not literally numbers (except in the finite case), so that $\leq$ is not denoting an honest inequality; $|A| \leq |B|$ means *by definition* that there exists an injection $A \to B$, and $|B| \leq |A|$ means there exists an injection $B \to A$, so the question is whether the existence of these two injections implies the existence of a *bijection* $A \to B$, which is what $|A| = |B|$ means. The Cantor-Schröder- Bernstein Theorem says that this is indeed the case:

> If there exists an injective function $A \to B$ and an injective function $B \to A$, then there exists a bijective function $A \to B$.

So, symbolically, $|A| \leq |B|$ and $|B| \leq |A|$ does imply $|A| = |B|$, so it is precisely because of this theorem that we *can* treat $\leq$ as if it where an actual inequality.

We'll give the proof of this next time, but for now notice what this might actually entail. The goal is to come up with a bijective function $A \to B$ given only injective functions $A \to B$ and $B \to A$. The point is that neither of the given injective functions $A \to B$ nor $B \to A$ are assumed to themselves be bijective, and indeed don't have to be, and yet from these alone it will be possible to produce some *other* function $A \to B$ which *will* be bijective!

Example. We'll finish for now with two example applications. We previously showed that $[0, 1)$ and $(0, 1)$ have the same cardinality by constructing an explicit bijection between them. An alternate way of showing they have the same cardinality is by showing that there exists injections between them in both directions and applying Cantor-Schroeder-Bernstein. First, the function $(0, 1) \to [0, 1)$ which sends anything in $(0, 1)$ to itself is injective. Second, to get an injection $[0, 1) \to (0, 1)$ imaging "shrinking" $[0, 1)$ down to $\left[0, \frac{1}{2}\right)$ and then "shifting" this new interval over to make it fit inside $(0, 1)$. Concretely, the composition of the function $f : [0, 1) \to [0, \frac{1}{2})$ defined by $f(x) = \frac{1}{2}x$ (i.e. "shrinking") and $g : \left[0, \frac{1}{2}\right) \to \left[\frac{1}{4}, \frac{3}{4}\right)$ defined by $g(x) = x + \frac{1}{4}$ (i.e. "shifting") gives an injective function $[0, 1) \to (0, 1)$ as required.

In a similar "geometric" vein, we show that the following two subsets of $\mathbb{R}^2$ have the same cardinality:

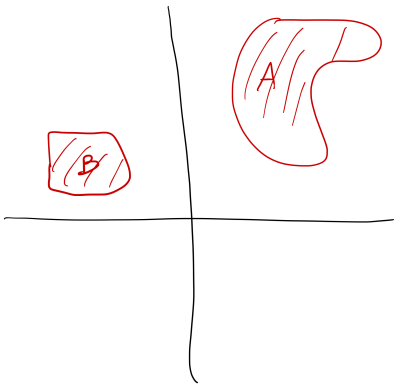

We can get an injection $A \to B$ by shrinking A (meaning scaling all points by an appropriate scalar which is less than 1) and then translating it over so that it fits inside B, and then we can get an injection $B \to A$ via a similar construction:

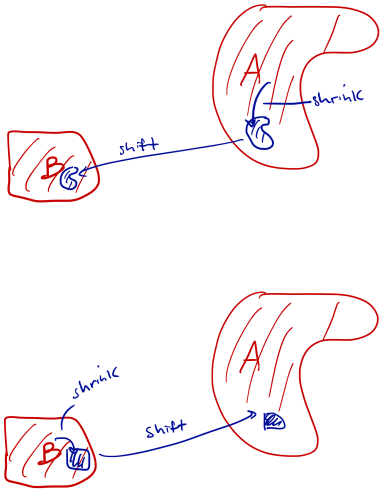

The Cantor-Schröeder-Bernstein theorem then implies that $|A| = |B|$ as claimed. Note again that we don't get an explicit (easy) bijection out of this, just that one exists.

Lecture 27: More on Cantor-Schroeder-Bernstein

Warm-Up. Suppose A, B, C are sets with $A \subseteq B \subseteq C$ and that $|A| = |C|$. We show that $|A| = |B|$ and $|B| = |C|$ as well. At first glance this might not seem to be something which requires

much in the way of proof (intuitively we are saying that if $|A| \leq |B| \leq |C|$ and $|A| = |C|$, then $|A| = |B| = |C|$), but as usual we have to be careful about how these symbols should all actually be interpreted via functions.

The function $f : A \to B$ which sends anything in A to itself is injective, as is the function $g : B \to C$ which sends anything in B to itself. This gives $|A| \leq |B|$ and $|B| \leq |C|$. Now, since $|A| = |C|$ there exists a bijective function $h : A \to C$. The composition $h^{-1} \circ g : B \to A$ is then injective, so $|B| \leq |A|$ and Cantor-Schroeder-Bernstein gives $|A| = |B|$. Also, the composition $f \circ h^{-1} : C \to B$ is injective, so $|C| \leq |B|$ and Cantor-Schroeder-Berstein again gives $|C| = |B|$ as claimed. (This all gives more evidence that $\leq$ in the context of cardinality should indeed be treated as a normal inequality.)

For one quick application, we now have a way of showing, say, that $\mathbb{R}^2$ has the same cardinality as the unit disk D centered at the origin in $\mathbb{R}^2$. We have

$$[-1, 1] \subseteq D \subseteq \mathbb{R}^2,$$

so since the interval $[-1, 1]$ and $\mathbb{R}^2$ have the same cardinality, so too does D. In general, a similar reasoning will show that any subset of $\mathbb{R}^2$ with positive area will have the same cardinality as $\mathbb{R}^2$.

Proof of Cantor-Schroeder-Bernstein. We now give a proof of the Cantor-Schroeder-Bernstein theorem, where given injective functions $f : A \to B$ and $g : B \to A$, the goal is to come up with a bijective function $A \to B$. The proof we give here is different than the proof our book gives, and is I think easier (although still tricky) to follow. The difficulty lies in having to construct a bijective function seemingly out of nowhere.

First we introduce a term which will help define the function we are after. Given $a \in A$, if there exists $b \in B$ such that $a = g(b)$, then there exists only such b by injectivity of g and we say that b is an *ancestor* of a. (The term "ancestor" is used in order to suggest that b comes "before" a since a came from b by applying g.) Similarly, given $y \in B$, if there exists $x \in A$ such that $f(x) = y$ then there is only one such x and we call it an ancestor of y. The upshot is that the "first" ancestor of $a \in A$, if there is one, is the element of B obtained by "undoing" g, and the first ancestor of $y \in B$, if there is one, is the element of A obtained by "undoing" f.

Now, suppose $a \in A$ has first ancestor $b \in B$, meaning that $a = g(b)$. We can now ask whether this $b \in B$ itself has an ancestor back in A, which is asking whether there exits $a' \in A$ such that $f(a') = b$. This a' is then also an ancestor of a (to be precise the "second" ancestor of a) since applying f and then g to it it will give a: $g(f(a')) = g(b) = a$. And so on, by backtracking as far back as possible, "undoing" f or g in an alternating manner, we can talk about other ancestors of $a \in A$ or $y \in B$. The *earliest* ancestor of $a \in A$, if it exists, is the ancestor we get by backtracking as much as possible, and similarly we can speak of the earliest ancestor of some $b \in B$.

For instance, if $a \in A$ is not in the image of g, then there is no $b \in B$ such that $a = g(b)$ and in this case a is its own earliest ancestor since we can't "backtrack" at all. If there does exist $b \in B$ such that $a = g(b)$ but where this b is itself not in the image of f, then b is the earliest ancestor of a since, although we can backtrack once to b from a, we can't backtrack any further since nothing in A will map to b under f. It can also happen that given $a \in A$, we can backtrack further and further without end: $a \in A$ comes from $b \in B$ when applying g, which in turn comes from $a' \in A$ when applying f, which in turn comes from $b' \in B$ when applying g, and so on without end. We say that such elements do not have an earliest ancestor.

With this notion of earliest ancestor we can now give our proof:

Proof. Let A_A denote the set of elements of A whose earliest ancestor is in A, A_B the set of elements of A whose earliest ancestor is in B, and A_∞ the set of elements of A with no earliest

ancestor. Similarly, let B_A, B_B, B_∞ denote the elements of B with earliest ancestor in A, in B, or nonexistence respectively. Since a given element of A either has an earliest ancestor or it doesn't, and since an earliest ancestor is unique, the sets A_A, A_B, A_∞ are disjoint and have union equal to A. Similarly, B_A, B_B, B_∞ are disjoint and have union equal to B.

Define the function $h : A \to B$ by

$$h(a) = \begin{cases} f(a) & \text{if } a \in A_A \cup A_\infty \\ \text{the first ancestor of } a & \text{if } a \in A_B. \end{cases}$$

To be clear, if $a \in A_B$, then the earliest ancestor of A is in B, which requires that it have a (unique) first ancestor in B to begin with, so that $h(a)$ for such elements is perfectly well-defined. We claim that h is bijective. First, note that if $a \in A$ has earliest ancestor in A, then so does $f(a) \in B$ since a is a first ancestor of $f(a)$ and any ancestor of a will then also be an ancestor of $f(a)$. This shows that f sends A_A into B_A. Similarly, f sends A_B into B_B and A_∞ into B_∞, again since a and $f(a)$ always have the same ancestors, and in particular if $a \in A_B$, the first ancestor of a is in B_B, so that h also sends A_A, A_B, A_∞ into B_A, B_B, B_∞ respectively. Visually, think of A and B as being made out of three pieces, with h sending each piece of A to the corresponding piece of B:

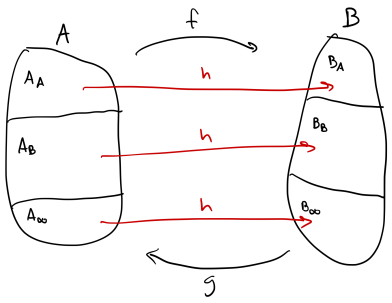

Now, to see that h is injective, suppose $h(a) = h(a')$. But what we said above, there are three possibilities: a and a' both come from A_A, or both from A_B, or both from A_∞. In other words, we can't have $a \in A_A$ and $a' \in A_B$ for instance since then $h(a) \in B_A$ and $h(a') \in B_B$ could not be equal. If $a, a' \in A_A \cup A_\infty$, then $h(a) = f(a)$ and $h(a') = f(a')$, so injectivity of f gives $a = a'$. If $a, a' \in A_B$, then $h(a) = h(a')$ means that a, a' both have the same first ancestor $b \in B$; but then $a = g(b) = a'$, so we conclude that h is injective.

To see that h is surjective, let $b \in B$. If $b \in B_A \cup B_\infty$, then there must be a first ancestor of b in A, meaning there exists $a \in A$ such that $f(a) = b$; since this a is then in $A_A \cup B_\infty$, we have $f(a) = h(b)$ so $h(a) = b$. If $b \in B_B$, then b is the first ancestor of $g(b) \in A$, so that $h(g(b)) = b$. Thus either way there exists $a \in A$ such that $h(a) = b$, so h is surjective. We conclude that h is bijective as claimed. □

Sets of functions on $\mathbb{R}$. We finished our discussion of Cantor-Schroeder-Bernstein with some final examples, both of which are included on the final homework assignment. The claims are that the set $C(\mathbb{R})$ of *continuous* functions from $\mathbb{R}$ to $\mathbb{R}$ has the same cardinality as $\mathbb{R}$, and that the set $F(\mathbb{R})$ of *all* functions from $\mathbb{R}$ to $\mathbb{R}$ has the same cardinality as the power set $\mathcal{P}(\mathbb{R})$ of $\mathbb{R}$. Since these are written up in the solutions to the final homework, we'll omit the details here. The upshot is that there are way more functions from $\mathbb{R}$ to $\mathbb{R}$ that aren't continuous than there are functions which are continuous, which might seem surprising since most all functions you've seen in previous courses are indeed continuous or at least piecewise continuous.

Practical uses of cardinality. And so our discussion of cardinality ends, but we should a bit about where all of this stuff comes up. Truth be told, it is highly unlikely you will see much in the way of cardinality in future courses or elsewhere; the *key* distinction which does show up in a few places is the distinction between what it means for a set to be countable vs uncountable. But apart from this, other aspects of cardinality don't play a big role. For instance, you will likely never again here about the cardinality of the power set of $\mathbb{R}$ or more elaborate things. We went through all this mainly as a way to introduce an interesting topic on which we could apply the mathematical reasoning and proof-writing skills we've developed. Hopefully you did find it to be interesting, even if abstract.

But, the distinction countable and uncountable sets is important, and here is one hint of this. Let A be a countable subset of, say, the unit square $[0,1] \times [0,1]$. We can imagine throwing a dart at this square and ask: what is the probability that we hit an element of A? The answer is zero! (Note that in this context, zero probability does not mean impossible, it just means highly highly highly unlikely, but you'll have to make a probability course like Math 310 to understand the distinction.) The reason comes down to the fact that any countable subset of the square, or of $\mathbb{R}^2$, has zero area. This is the type of result which plays some role in probability, analysis, and other such fields. But, we'll have to leave further discussion of this to those courses. Thanks for reading!